Two-stage and Lagrangian Dual Decision Rules for Multistage Adaptive Robust Optimization

Maryam Daryalal
Department of Decision Sciences, HEC Montréal, Montréal, Québec H3T 2A7, Canada,
maryam.daryalal@hec.ca

Ayşe N. Arslan
Univ. Bordeaux, CNRS, INRIA, Bordeaux INP, IMB, UMR 5251, F-33400 Talence, France,
Centre Inria de l'Universite de Bordeaux, F-33405 Talence, France,
ayse-nur.arslan@inria.fr

Merve Bodur
Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario M5S 3G8, Canada,
bodur@mie.utoronto.ca

In this work, we design primal and dual bounding methods for multistage adjustable robust optimization (MSARO) problems by adapting two decision rules rooted in the stochastic programming literature. This approach approximates the primal and dual formulations of an MSARO problem with two-stage models. From the primal perspective, this is achieved by applying two-stage decision rules that restrict the functional forms of a certain subset of decision variables. We present sufficient conditions under which the well-known constraint-and-column generation algorithm can be used to solve the primal approximation with finite convergence guarantees. From the dual side, we introduce a distributionally robust dual problem for MSARO models using their nonanticipative Lagrangian dual and then apply linear decision rules on the Lagrangian multipliers. For this dual approximation, we present a monolithic bilinear program valid for continuous recourse problems, and a cutting-plane method for mixed-integer recourse problems. Our framework is general-purpose and does not require strong assumptions such as a stage-wise independent uncertainty set, and can consider integer recourse variables. Computational experiments on newsvendor, location-transportation, and capital budgeting problems show that our bounds yield considerably smaller optimality gaps compared to the existing methods.

Key words: Optimization under uncertainty, Robust optimization, Decision rules

---

1. Introduction

Many practical planning, design and operational problems involve making decisions under uncertainty at consecutive stages, where the decisions in one stage affect the decisions of the future stages. In such sequential decision-making problems, first-stage (*here-and-now*) decisions are the ones that are immediately implementable. Subsequent recourse (*wait-and-see*) decisions depend on the state of the system, which is a result of previous decisions and observations of the uncertain parameters. A solution is then an adaptable *policy* or *decision rule* that takes the previous decisions and history of uncertainty realizations as an input, and returns a new implementable decision. The dynamics of a sequential decision-making problem is depicted in Figure 1.

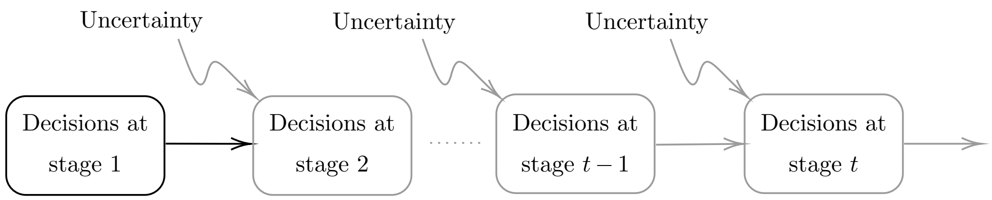

Figure 1 Sequential decision-making under uncertainty

There are several modeling frameworks for sequential decision-making problems under uncertainty. With a known probability distribution, these problems may be addressed by the *multistage stochastic programming* (MSP) paradigm, with the goal of optimizing some statistical performance measure over the planning horizon. There is an extensive body of research on MSP problems of various structures, with a rich literature on problems with continuous decision variables. Without a precise distribution, Bertsimas and Thiele (2006) point out that an *assumed* probability distribution might heavily impact the results if it is different than the *true* one, even if both distributions share the same first and second moments. To mitigate this effect, *distributionally robust optimization* (DRO) models are proposed for making decisions that are based on a family of probability distributions, often defined by using historical data (Goh and Sim 2010). These models aim to hedge against tuning decisions to a perceived distribution. However, while some recent studies have proposed tractable solution methods for linear DRO problems under certain conditions (Philpott et al. 2018, Bertsimas et al. 2019), they remain largely challenging to solve. To the best of our knowledge, the only such methodology for a multistage mixed-integer problem is developed under the assumption of endogenous uncertainty (Yu and Shen 2020), meaning that the decisions impact the underlying stochastic structure of the problem, an assumption that might not always hold.

Multistage adaptive robust optimization (MSARO) is another framework for modeling sequential decision-making problems under uncertainty, when we do not have any knowledge about the underlying probability distribution or when the underlying uncertainty is not stochastic in nature. In designing such robust solutions, we protect ourselves against worst-case outcomes. While *static* (i.e., single-stage) robust optimization is tractable for a considerable number of problem structures (e.g., a linear problem with ellipsoidal uncertainty set, yielding a conic quadratic program (Ben-Tal and Nemirovski 1999); also see Ben-Tal et al. (2009), Bertsimas et al. (2011a, 2015)), MSARO in general remains computationally difficult to solve, even for two stages (Guslitser 2002, Ben-Tal et al. 2004, Feige et al. 2007). Consequently, most solution methods rely on an approximation of MSARO. In this work we present a new solution paradigm for MSARO problems by employing primal and dual decision rules that result in two-stage approximations of MSARO and its dual problem, respectively. This leads to a solution framework that returns an adaptable policy, and an optimality gap to assess the quality of the proposed policy. In what follows, we provide an overview of the existing literature on MSARO, followed by a summary of our contributions.

1.1. Literature Review

Static robust optimization is computationally less demanding than MSARO. As such, a number of studies attempt to find problem structures where the static optimal solution is also guaranteed to be an optimal solution of the MSARO (Ben-Tal et al. 2004, Marandi and Den Hertog 2018, Bertsimas et al. 2015). For example, this is the case for linear MSARO with constraint-wise uncertainty and a compact uncertainty set, provided that it is also linear with respect to uncertain parameters. In general, however, the solutions to the static and adaptive variants differ and the MSARO solutions are preferable by virtue of being more flexible. Figure 2 presents a summary of existing solution methods for obtaining exact, approximate, and dual bounds on the MSARO, with methods developed specifically for two-stage adaptive robust optimization (2ARO) separately categorized. In the following, we briefly discuss each method and the specific problem structure it can address.

Exact solution methods are scarce in the MSARO literature and the existing studies mostly focus on two-stage adaptive robust optimization (2ARO). Bertsimas et al. (2012) designed a Benders decomposition and outer approximation approach for a two-stage unit commitment problem. Zeng and Zhao (2013) developed constraint-and-column generation for 2ARO problems with fixed recourse and finite or polyhedral uncertainty set. For small linear problems, Zhen et al. (2018) used the Fourier-Motzkin elimination iteratively to remove the second-stage decisions in a 2ARO problem with continuous fixed recourse, eventually forming a static robust optimization problem that recovers optimal first-stage solutions. They extended their method to the multistage setting by modifying the algorithm to make sure that at every stage, no decision depends on the future information (these are the so-called nonanticipativity constraints). In the case of mixed-binary recourse and only objective uncertainty, Arslan and Detienne (2022) proposed an exact method based on a Dantzig-Wolfe reformulation of the recourse problem. Similarly using Dantzig-Wolfe reformulation, Hashemi Doulabi et al. (2021) proposed a static formulation which is amenable to Benders decomposition for a subclass of 2ARO problems with mixed-integer recourse and a finite uncertainty set. In a rare instance of studies addressing the exact solution of multistage problems in

Figure 2 Solution methods for the MSARO

the literature of robust optimization, Georghiou et al. (2019) developed robust dual dynamic programming (RDDP) for continuous MSARO problems with a stage-wise rectangular uncertainty set, and proved finite/asymptomatic convergence for various problem sub-classes. RDDP is an adaptation of the stochastic dual dynamic programming algorithm from the stochastic programming literature (Pereira and Pinto 1991), to the robust setting. Georghiou et al. (2020) proposed a convergent method based on enumeration of the extreme points of the uncertainty set combined with affine decision rules to provide gradually improving primal and dual bounds for 2ARO problems with continuous recourse.

A commonly used approximation approach restricts adaptable/adjustable decisions to follow a certain functional form, known as *decision rules*. Ben-Tal et al. (2004) proposed the first decision rule for MSARO, linear decision rule (LDR), where continuous recourse decisions are expressed as affine functions of uncertain parameters where the parameters of this function are to be optimized. This approach was further generalized by introducing polynomial decision rules to construct more flexible policies. For the MSARO with continuous fixed recourse, Bertsimas et al. (2011b) showed that polynomial policies with a predetermined degree can be computed using a semidefinite program (comparable to the work of Bampou and Kuhn (2011) for MSP). Xu and Hanasusanto (2021) proposed tighter approximations than that of Ben-Tal et al. (2004) and Bampou and Kuhn (2011) (when using the same degree), by presenting exact copositive programming reformulations of the LDR-restricted MSARO with random recourse, and the quadratic decision rule-restricted problem with fixed recourse. Nonlinear decision rules were also proposed for the MSARO with continuous fixed recourse, such as deflected and segregated affine (Chen et al. 2008), extended affine (Chen and Zhang 2009), piecewise affine (Goh and Sim 2010), truncated linear (See and Sim 2010), and piecewise affine with exponentially many pieces (Ben-Tal et al. 2020). For a comprehensive list of nonlinear decision rules, interested reader may refer to a recent survey by Yanıkoğlu et al. (2019).

With mixed-integer recourse, LDRs and most of its extensions are not applicable. For the mixed-*binary* case, Bertsimas and Georghiou (2015) endogenously designed piecewise linear/constant decision rules. Bertsimas and Georghiou (2018) introduced binary decision rules for binary MSARO problems with fixed or random recourse. A popular line of research for general mixed-integer problems uses the notion of *finite adaptability*, first introduced by Bertsimas and Caramanis (2010) for linear MSARO. In finite or K-adaptability, the decision-maker a priori commits to K recourse decisions (while making the first-stage decisions), and then chooses among them after observing the uncertain parameters. Hanasusanto et al. (2015) tailored the method for problems with binary recourse, and Subramanyam et al. (2020) generalized it for mixed-integer 2ARO. A related approach is based on the concept of uncertainty set partitioning which returns a sequence of improving approximations (Bertsimas and Dunning 2016, Postek and Hertog 2016, Romeijnders and Postek 2020). Although these latter methods can be used to approximate mixed-integer MSARO problems, they do not provide a fully adaptable policy.

The computational advantage obtained from the described approximation methods often comes at the expense of optimality. In a few special cases, LDRs are shown to be optimal for the MSARO. Bertsimas et al. (2010) showed that in a one-dimensional MSARO problem with no first-stage decisions, linear second-stage costs, and a hyperrectangle uncertainty set, affine policies are optimal. In a series of studies on 2ARO problems it was shown that the result still holds if we have only random right-hand side and/or technology matrix and a simplex uncertainty set (Bertsimas and Goyal 2012, Iancu et al. 2013, Zhen et al. 2018). Nevertheless, for the majority of the MSARO problems, decision rules yield suboptimal policies.

In order to assess the quality of the solutions of an LDR-restricted problem, Kuhn et al. (2011) applied LDRs on the dual variables associated with the linear programming dual of the original problem with no integer decisions, obtaining a dual bound and thus an optimality gap. From another perspective, Hadjiyiannis et al. (2011) solved an MSARO problem for a finite set of scenarios sampled from the uncertainty set to reach a dual bound. Bertsimas and de Ruiter (2016) proposed a reformulation of MSARO models with polyhedral uncertainty, by dualizing over the recourse decisions and then uncertain parameters, leading to an equivalent problem that achieves a dual bound (potentially different than the one obtained from the original nested formulation) for a finite set of scenarios, namelythose binding in an LDR-restricted approximation of the problem.

1.2. Contributions

We present decision rule-based approaches to obtain primal and dual bounds for an MSARO problem. On the primal side, we apply *two-stage decision rules* that approximate the multistage problem with a 2ARO, and leverage existing solution methods from the 2ARO literature to solve the obtained model. This results in feasible adjustable policies. On the dual side, we use Lagrangian duality and *Lagrangian dual decision rules* for deriving a dual problem for an MSARO problem, for which we develop reformulations and solution methods. Accordingly, we obtain an optimality gap for measuring the quality of the proposed policies. Figure 3 depicts an overview of our framework. More specifically, our new policies are inspired by two numerically promising decision rules recently developed in the stochastic programming community for MSP: two-stage linear decision rules (Bodur and Luedtke 2018) and Lagrangian dual decision rules (Daryalal et al. 2021), that approximate the MSP with a two-stage problem in the primal and dual space, respectively. On the primal side we have generalized the idea of two-stage LDRs and designed a new decision rule with a broader application, including some classes of MSARO with mixed-integer recourse. Bringing this idea to the context of robust optimization makes it possible to benefit from the rich literature of 2ARO in (approximately) solving multistage problems. On the dual side, the literature of stochastic programming and distributionally robust optimization (DRO) prove to be useful in the

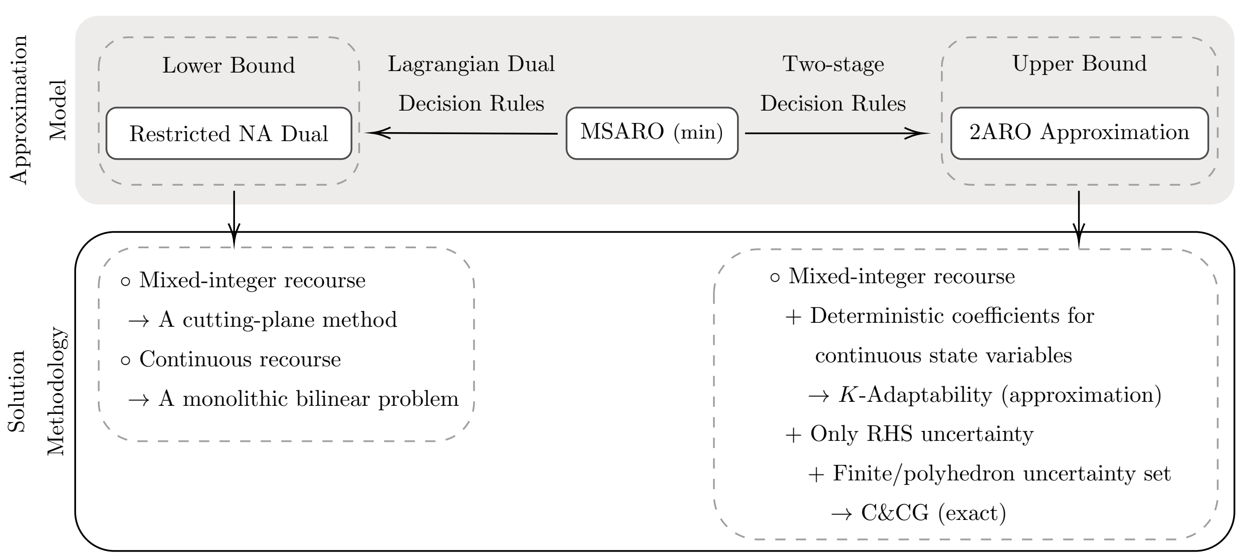

Figure 3 Proposed bounding framework for MSARO

development of a new dual problem for MSARO, building a bridge between these close communities. As in (Daryalal et al. 2021), dual solutions can also be potentially used to derive primal policies, which we leave for future research. A summary of our contributions is given below:

- For a broad class of multistage adaptive robust optimization problems, we propose a solution framework that returns adaptable policies as well as a dual bound measuring the quality of these policies. We do this by employing two-stage and Lagrangian dual decision rules, leading to novel techniques that reveal new theoretical and practical avenues.
- We present the class of two-stage decision rules that impose a predefined structure on a subset of the adaptable/adjustable decision variables, with the goal of building an approximate model that returns feasible first-stage solutions. For the resulting 2ARO approximation, we show that in some special cases, constraint-and-column generation algorithm converges to the exact solution of the 2ARO approximation, otherwise the K-adaptability approach is applicable.
- For a general MSARO, we present a dual problem and prove that, under certain conditions, strong duality holds. This is one of the few approaches for obtaining dual bounds for MSARO with integer recourse. To the best of our knowledge, the only other approach that applies to mixed-integer recourse is discretization/sampling. For this dual, we introduce Lagrangian dual decision rules, resulting in a two-stage problem. We design a cutting-plane algorithm that solves the restricted dual problem for MSARO with mixed-integer recourse. We also show that in the case of continuous recourse, the problem can be reformulated as a monolithic bilinear program. Additionally, we present an alternative dual problem that is decomposable by scenario and offers computational efficacy.
- We have evaluated the performance of our solution framework over multistage versions of three classical problems in the literature of MSARO: (i) the newsvendor problem, (ii) the location-transportation problem, and (iii) the capital budgeting problem. Each of these problem classes is suitable for a different solution method developed in this work and our analysis over various instances attests to the quality of the returned primal and dual bounds.

The remainder of the paper is organized as follows. In Section 2 we introduce the notion of a state variable and design two-stage decision rules for retrieving feasible solutions and adaptable policies. In Section 3 we reformulate the problem by employing the *nonanticipativity* constraints, and then propose a dual problem through their Lagrangian relaxation. To solve this dual problem, we apply linear decision rules on the Lagrangian multipliers of the relaxed nonanticipativity constraints. We formalize the problem of obtaining the best such dual bound, and develop solution methods depending on the nature of the recourse problem. This is followed by numerical experiments in Section 4 and concluding remarks. We note that all proofs are deferred to Appendix A.

Notation. Throughout, we use $[a] := \{1, 2, \ldots, a\}$ and $[a, b] := \{a, a+1, \ldots, b\}$ for positive integers a and b (with $a \leq b$), and $(\cdot)^\top$ for the transpose operator.

2. Primal Bounding

In this work, we study MSARO problems of the following form:

$$\nu^\star = \min \ z \tag{1a}$$

$$\text{s.t.} \ \sum_{t\in[T]} c_t(\boldsymbol{\xi}^t)^\top x_t(\boldsymbol{\xi}^t) \le z \qquad \boldsymbol{\xi}^T \in \Xi \tag{1b}$$

$$A_t(\boldsymbol{\xi}^t)x_t(\boldsymbol{\xi}^t) + B_t(\boldsymbol{\xi}^t)x_{t-1}(\boldsymbol{\xi}^{t-1}) \le b_t(\boldsymbol{\xi}^t) \qquad t\in[2,T], \ \boldsymbol{\xi}^t \in \Xi^t \tag{1c}$$

$$D_t(\boldsymbol{\xi}^t)x_t(\boldsymbol{\xi}^t) \le d_t(\boldsymbol{\xi}^t) \qquad t\in[T], \ \boldsymbol{\xi}^t \in \Xi^t \tag{1d}$$

$$x_t(\boldsymbol{\xi}^t) \in \mathbb{R}^{n_t - n_t^{\text{i}}} \times \mathbb{Z}^{n_t^{\text{i}}} \qquad t\in[T], \ \boldsymbol{\xi}^t \in \Xi^t, \tag{1e}$$

where T denotes the number of decision stages, $c_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{n_t}, b_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\text{S}}}, d_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\text{R}}}, A_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\text{S}} \times n_t}, B_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\text{S}} \times n_{t-1}}, D_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\text{R}} \times n_t}$. The uncertainty is represented by $\{\boldsymbol{\xi}_t\}_{t=1}^T$ with support Ξ where $\boldsymbol{\xi}_t$ is an uncertain vector of random variables at stage t with outcomes $\xi_t \in \Xi_t = \text{proj}_{\boldsymbol{\xi}_t}(\Xi) \subseteq \mathbb{R}^{\ell_t}$ (where by convention $\xi_1 = 1$). A realization of uncertainty is referred to as a *scenario*, $\{\xi_t\}_{t=1}^T$, which is a set of outcomes for all random variables at every stage such that $\xi_t \in \Xi_t$. History of uncertainty at stage t is characterized by $\boldsymbol{\xi}^t = (\xi_1, \dots, \boldsymbol{\xi}_t)$. Decision variables are denoted by $x_t(\boldsymbol{\xi}^t) \in X_t(\boldsymbol{\xi}^t) = \{x_t \in \mathbb{R}^{n_t - n_t^{\text{i}}} \times \mathbb{Z}^{n_t^{\text{i}}} : \ D_t(\boldsymbol{\xi}^t)x_t \le d_t(\boldsymbol{\xi}^t)\}$, i.e., nonanticipative policies that only depend on the history of uncertainty. Throughout, we make the following assumptions:

Assumption 1. *The problem has relatively complete recourse, i.e., for all $t \in [1,T]$, $\boldsymbol{\xi}^t$ and a history of feasible decisions $x_t(\boldsymbol{\xi}^t)$ made up to t, there always exists a feasible decision at stage t.*

Assumption 2. *For $t \in [T]$ and $\boldsymbol{\xi}^t \in \Xi$, the maximum diameter of $X_t(\boldsymbol{\xi}^t)$ is finite, i.e., the feasibility sets are bounded.*

When the set $X_t(\boldsymbol{\xi}^t), t \in [T]$ does not have integrality restrictions ($n_t^{\text{i}} = 0$), we have *continuous* recourse, otherwise we are dealing with an MSARO problem with *mixed-integer* recourse. Together with constraints (1b), the objective function (1a) minimizes the worst outcome. Constraints (1c) and (1d) are *state* and *recourse* constraints, respectively: while the former link different stages, the latter are local restrictions for a specific stage.

For an MSARO with continuous recourse, it is quite common to restrict the decisions $x_t(\boldsymbol{\xi}^t)$ to an affine form, which, by breaking the temporal dependencies between stages, approximates problem (1) with a static robust optimization problem. Let ν^{LDR} be the bound obtained by applying such linear decision rules (LDRs) to the set of *all* decision variables. Our goal in this section is to employ a new paradigm where a specific *subset* of the decision variables $x_t(\boldsymbol{\xi}^t)$ are enforced to follow a structured decision rule, leading to an approximation in the form of a 2ARO problem.

2.1. Two-stage Decision Rules

In a similar manner to constraints (1c) and (1d), let us partition the decision variables $x_t(\boldsymbol{\xi}^t), t \in [2,T]$ into $x^{\mathsf{s}}_t(\boldsymbol{\xi}^t) \in \mathbb{R}^{q_t}$ and $x^{\mathsf{r}}_t(\boldsymbol{\xi}^t) \in \mathbb{R}^{p_t}$, *state* and *recourse* variables, respectively, with $q_t + p_t = n_t$, with integrality restrictions, if any, embedded in the set $X_t(\boldsymbol{\xi}^t)$. Then the MSARO can be reformulated as follows:

$$
\begin{aligned}
\nu^\star = \min\ & z && && \text{(2a)}\\
\text{s.t. } & \sum_{t\in[T]} c^{\mathsf{s}}_t(\boldsymbol{\xi}^t)^\top x^{\mathsf{s}}_t(\boldsymbol{\xi}^t) + c^{\mathsf{r}}_t(\boldsymbol{\xi}^t)^\top x^{\mathsf{r}}_t(\boldsymbol{\xi}^t) \le z && \boldsymbol{\xi}^T \in \Xi && \text{(2b)}\\
& A^{\mathsf{s}}_t(\boldsymbol{\xi}^t) x^{\mathsf{s}}_t(\boldsymbol{\xi}^t) + B^{\mathsf{s}}_t(\boldsymbol{\xi}^t) x^{\mathsf{s}}_{t-1}(\boldsymbol{\xi}^{t-1}) + A^{\mathsf{r}}_t(\boldsymbol{\xi}^t) x^{\mathsf{r}}_t(\boldsymbol{\xi}^t) \le b_t(\boldsymbol{\xi}^t) && t\in[2,T],\ \boldsymbol{\xi}^t \in \Xi^t && \text{(2c)}\\
& (x^{\mathsf{s}}_t(\boldsymbol{\xi}^t), x^{\mathsf{r}}_t(\boldsymbol{\xi}^t)) \in X_t(\boldsymbol{\xi}^t) && t\in[2,T],\ \boldsymbol{\xi}^t \in \Xi^t && \text{(2d)}\\
& x_1 \in X_1, && && \text{(2e)}
\end{aligned}
$$

where $x^{\mathsf{s}}_1(\boldsymbol{\xi}^1) = x_1$, and $c^{\mathsf{s}}_t(\boldsymbol{\xi}^t), A^{\mathsf{s}}_t(\boldsymbol{\xi}^t), B^{\mathsf{s}}_t(\boldsymbol{\xi}^t)$ are sub-arrays/sub-matrices of $c_t(\boldsymbol{\xi}^t), A_t(\boldsymbol{\xi}^t), B_t(\boldsymbol{\xi}^t)$ associated with the state variables with appropriate dimensions, while $c^{\mathsf{r}}_t(\boldsymbol{\xi}^t), A^{\mathsf{r}}_t(\boldsymbol{\xi}^t)$ have the same role for the recourse variables. Let $x^{\mathsf{s}}_t(\boldsymbol{\xi}^t)$ be approximated by a decision rule, i.e., $x^{\mathsf{s}}_t(\boldsymbol{\xi}^t) = \Theta_t(\boldsymbol{\xi}^t, \beta_t)$, where $\Theta_t : \mathbb{R}^{\ell^t} \times \mathbb{R}^{K_t} \to \mathbb{R}^{q_t}$ represents the rule, and $\beta_t \in \mathbb{R}^{K_t}$ is its vector of design parameters. By substituting this rule in problem (2), we obtain an approximation that can be reformulated as:

$$
\begin{aligned}
\nu^{2\mathrm{S}} = \min\ & c_1^\top x_1 + \max_{\boldsymbol{\xi}^T\in\Xi} \min_{x^{\mathsf{r}}\in\mathcal{X}(x_1,\beta,\boldsymbol{\xi}^T)} \sum_{t\in[2,T]} c^{\mathsf{s}}_t(\boldsymbol{\xi}^t)^\top \Theta_t(\boldsymbol{\xi}^t,\beta_t) + c^{\mathsf{r}}_t(\boldsymbol{\xi}^t)^\top x^{\mathsf{r}}_t && \text{(3a)}\\
\text{s.t. } & x_1 \in X_1 && \text{(3b)}\\
& \beta_t \in \mathbb{R}^{K_t} \qquad t\in[2,T], && \text{(3c)}
\end{aligned}
$$

where:

$$
\begin{aligned}
\mathcal{X}(x_1,\beta,\boldsymbol{\xi}^T) = \Big\{ & \big(x^{\mathsf{r}}_t(\boldsymbol{\xi}^t)\big)_{t\in[2,T]} \in \mathbb{R}^{p_2}\times\mathbb{R}^{p_3}\times\cdots\times\mathbb{R}^{p_T} : \\
& A^{\mathsf{r}}_t(\boldsymbol{\xi}^t) x^{\mathsf{r}}_t \le b_t(\boldsymbol{\xi}^t) - \Big(A^{\mathsf{s}}_t(\boldsymbol{\xi}^t)\Theta_t(\boldsymbol{\xi}^t,\beta_t) + B^{\mathsf{s}}_t(\boldsymbol{\xi}^t)\Theta_{t-1}(\boldsymbol{\xi}^{t-1},\beta_{t-1})\Big) \quad t\in[2,T]\\
& (\Theta_t(\boldsymbol{\xi}^t,\beta_t), x^{\mathsf{r}}_t) \in X_t(\boldsymbol{\xi}^t) \quad t\in[2,T] \Big\}.
\end{aligned}
$$

Note that the decision rules are solely applied to the state variables, whereas the recourse variables remain fully adjustable to the uncertain parameters (see Figure 4). Problem (3) is a 2ARO and the temporal dependency between stages is removed thanks to the application of the two-stage decision rule. A practical result of such an approximation is that the resulting problem can be solved using the existing solution methods for the 2ARO. In the following sections, we study two special cases that directly benefit from this literature. The first one addresses MSARO problems with continuous state variables, and the second can be applied to those problems that have only bounded integer state variables. Together, they permit the approximation of a (bounded) mixed-integer MSARO with a 2ARO.

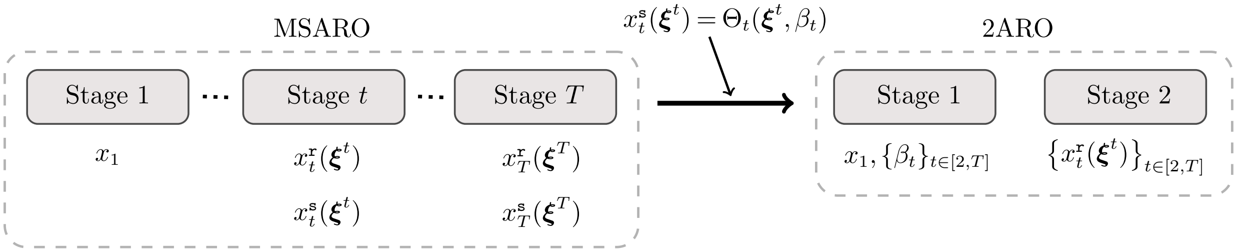

Figure 4 Two-stage decision rules

2.2. Two-stage Linear Decision Rules

If the state variables are continuous, we can approximate them with an affine function. Let $\Phi_t(\boldsymbol{\xi}^t) = \big(\Phi_{t1}(\boldsymbol{\xi}^t), \ldots, \Phi_{tK_t}(\boldsymbol{\xi}^t)\big) : \mathbb{R}^{\ell^t} \to \mathbb{R}^{K_t}$ be a vector of basis functions. The *two-stage LDR* is enforced by using $\Theta_t(\boldsymbol{\xi}^t, \beta_t) = \Phi_t(\boldsymbol{\xi}^t)^\top \beta_t$ in (3). The result is a two-stage mixed-integer *linear* program:

$$\nu^{\text{2S-LDR}} = \min \quad c_1^\top x_1 + \max_{\boldsymbol{\xi}^T \in \Xi} \; \min_{x^{\mathsf{r}} \in \mathcal{X}(x_1, \beta, \boldsymbol{\xi}^T)} \sum_{t\in[2,T]} c^{\mathsf{s}}_t(\boldsymbol{\xi}^t)^\top \Phi_t(\boldsymbol{\xi}^t)^\top \beta_t + c^{\mathsf{r}}_t(\boldsymbol{\xi}^t)^\top x^{\mathsf{r}}_t \tag{4a}$$

$$\text{s.t.} \quad x_1 \in X_1 \tag{4b}$$

$$\beta_t \in \mathbb{R}^{K_t} \qquad t \in [2, T]. \tag{4c}$$

As recourse variables $x^{\mathsf{r}}_t(\boldsymbol{\xi}^t)$ in (4) are fully adjustable, as opposed to the bound ν^{LDR} obtained by restricting them to LDRs, it immediately follows for MSARO with continuous state variables, that

$$\nu^\star \le \nu^{\text{2S-LDR}} \le \nu^{\text{LDR}}.$$

In general, we can approximate the 2ARO problem (4) with the K-adaptability approach (see Section 1.1 for an overview of the related literature), while under certain conditions it can be solved exactly. The focus of the following is on the latter case, where we assume that the original MSARO problem has fixed recourse, i.e., $c_t(\boldsymbol{\xi}^t) = c_t, A_t(\boldsymbol{\xi}^t) = A_t, B_t(\boldsymbol{\xi}^t) = B_t, D_t(\boldsymbol{\xi}^t) = D_t, t \in [2, T]$.

For solving a 2ARO problem with Ξ either a finite discrete set or a polyhedron, and only right-hand side uncertainty, Zeng and Zhao (2013) proposed the constraint-and-column generation (C&CG) method, by drawing on the fact that not all realizations in Ξ contribute to the optimal objective value. The C&CG strives to identify the most significant scenarios by starting from a smaller uncertainty set and gradually expanding it, which leads to the generation of new columns and constraints, respectively corresponding to recourse variables and second-stage constraints for the newly identified scenario. The C&CG algorithm of Zeng and Zhao (2013) assumes fixed technology matrix and first-stage costs, i.e., that the coefficients of the first-stage variables in the objective function and constraints are deterministic. However, in problem (4), basis functions $\Phi_t(\boldsymbol{\xi}^t)$ appear as coefficients of β_t, both in the objective function (4a) (in the term $c^{\mathsf{s}}_t(\boldsymbol{\xi}^t)^\top \Phi_t(\boldsymbol{\xi}^t)$) and the second-stage constraints (terms $A^{\mathsf{s}}_t \Phi_t(\boldsymbol{\xi}^t), B^{\mathsf{s}}_t \Phi_{t-1}(\boldsymbol{\xi}^{t-1})$ and $D^{\mathsf{s}}_t \Phi_t(\boldsymbol{\xi}^t)$). Next, we present the adaptation of the C&CG algorithm for solving problem (4) when the technology matrix and first-stage costs are not fixed, followed by a sufficient condition for finite convergence of the C&CG algorithm.

Constraint-and-column Generation for Two-stage LDRs Consider a relaxation of problem (4) where instead of the uncertainty set Ξ, a potentially empty subset $\hat{\Xi} \subseteq \Xi$ is used. With an attempt to build a tighter relaxation, the C&CG gradually expands $\hat{\Xi}$ with what we call *significant scenarios*. The optimal objective value of the following problem is a lower bound on $\nu^{\text{2S-LDR}}$:

$$\begin{aligned}
\mathcal{MP}(\hat{\Xi}):\ \min\ & c_1^\top x_1 + \eta && && \text{(5a)}\\
\text{s.t.}\ & \eta \geq \sum_{t\in[2,T]} \left(c_t^{\mathbf{s}\top} \Phi_t(\xi^t)^\top \beta_t + c_t^{\mathbf{r}\top} x^{\mathbf{r}}_{t,\xi^t} \right) && \xi^T \in \Xi^{\text{Feas}} && \text{(5b)}\\
& A_t^{\mathbf{r}} x^{\mathbf{r}}_{t,\xi^t} + A_t^{\mathbf{s}} \Phi_t(\xi^t)^\top \beta_t + B_t^{\mathbf{s}} \Phi_{t-1}(\xi^{t-1})^\top \beta_{t-1} \leq b_t(\xi^t) && t\in[2,T], \xi^t \in \hat{\Xi}^t && \text{(5c)}\\
& (\Phi_t(\xi^t)^\top \beta_t, x^{\mathbf{r}}_{t,\xi^t}) \in X_t(\xi^t) && t\in[2,T], \xi^t \in \hat{\Xi}^t && \text{(5d)}\\
& x_1 \in X_1, \quad \eta \in \mathbb{R} && && \text{(5e)}\\
& \beta_t \in \mathbb{R}^{K_t} && t\in[2,T], && \text{(5f)}
\end{aligned}$$

where $\hat{\Xi}^t = \text{proj}_{\boldsymbol{\xi}^t}(\hat{\Xi})$, and $\Xi^{\text{Feas}} \subseteq \hat{\Xi}$ is the set of identified significant scenarios ξ^T for which there exists a feasible first-stage solution $\hat{\beta}$ such that the feasibility space $\mathcal{X}(x_1, \hat{\beta}, \xi^T)$ is not empty. In order to have an exact solution to problem (4), we need to expand $\hat{\Xi}$ with significant scenarios (and consequently the set of recourse variable copies $x^{\mathbf{r}}_{t,\xi^t}$), until it can be proven that adding another scenario does not improve the optimal objective value of $\mathcal{MP}(\hat{\Xi})$, at which point its optimal objective value would be equal to $\nu^{\text{2S-LDR}}$. This implies that, at convergence, solving $\mathcal{MP}(\hat{\Xi})$ returns an optimal solution $(\hat{x}_1, \hat{\beta}, \hat{\eta})$ such that $\hat{\eta}$ accurately measures the worst-case second-stage cost over the complete uncertainty set Ξ. To check whether the convergence criterion is satisfied, we solve the adversarial problem for a given first-stage solution $\hat{\beta}$, which results in the following subproblem:

$$\mathcal{SP}(\hat{\beta}) = \max_{\boldsymbol{\xi}^T \in \Xi} \left\{ \sum_{t\in[2,T]} c_t^{\mathbf{s}\top} \Phi_t(\boldsymbol{\xi}^t)^\top \hat{\beta}_t + \min_{x^{\mathbf{r}} \in \mathcal{X}(x_1, \hat{\beta}, \boldsymbol{\xi}^T)} \sum_{t\in[2,T]} c_t^{\mathbf{r}\top} x_t^{\mathbf{r}} \right\}. \tag{6}$$

If $\hat{\eta} = \mathcal{SP}(\hat{\beta})$, then $\hat{\eta}$ exactly measures the worst-case cost of the second-stage problem and $(\hat{x}_1, \hat{\beta}, \hat{\eta})$ is an optimal solution to problem (4), i.e., $\nu^{\text{2S-LDR}} = c_1^\top \hat{x}_1 + \hat{\eta}$. Otherwise, let $\hat{\xi}^T$ be the optimal solution of subproblem (6) if it is feasible, or a scenario such that $\mathcal{X}(x_1, \hat{\beta}, \hat{\xi}^T)$ is an empty set. We create new variables $x^{\mathbf{r}}_{t,\hat{\xi}^t}, t \in [2,T]$ and update the set of significant scenarios with $\hat{\Xi} = \hat{\Xi} \cup \{\hat{\xi}^T\}$, and accordingly update (5c) and (5d). If subproblem (6) is feasible, then we update $\Xi^{\text{Feas}} = \Xi^{\text{Feas}} \cup \{\hat{\xi}\}$ and (5b) as well. In this case, constraints (5b)-(5d) make up the optimality cuts. If subproblem (6) is infeasible, then (5c)-(5d) act as feasibility cuts. We repeat solving the master problem and the subproblem in a cutting-plane fashion and add the appropriate cuts, until $\hat{\Xi}$ includes all the significant scenarios and $\hat{\eta} = \mathcal{SP}(\hat{\beta})$. The following proposition provides a sufficient condition for finite convergence of the C&CG algorithm (proof is given in Appendix A).

Proposition 1. *Assume that the uncertainty set Ξ is either a finite discrete set or a compact polyhedron, and the basis functions $\Phi_t(\boldsymbol{\xi}^t)$ are affine in $\boldsymbol{\xi}^t$. Then the C&CG algorithm converges to $\nu^{2S\text{-}LDR}$ of an MSARO with continuous and fixed recourse in a finite number of iterations.*

In case the conditions of Proposition 1 are not satisfied, the C&CG algorithm can still converge asymptotically to the optimal solution of problem (4). This is because the lower bound (optimal objective value of $\mathcal{MP}(\hat{\Xi})$) improves with each iteration, and under the Assumptions 1 and 2, the optimal value of problem (4) is finite.

Remark 1. With continuous recourse (i.e., $n_t^{\mathrm{i}} = 0, t \in [2,T]$), the subproblem (6) can be reformulated as a bilinear monolithic problem, using LP duality or Karush–Kuhn–Tucker conditions (Zeng and Zhao 2013). For details, refer to Appendix B.1.

2.3. Two-stage Piecewise-constant Decision Rules

In this section we study the application of two-stage decision rules in another special case of MSAROs, where each state variable $x_{ti}^{\mathrm{s}}(\boldsymbol{\xi}^t) \in \mathbb{Z}, i \in [q_t], t \in [2,T]$ is a bounded integer with a given domain $[\underline{\kappa}_{ti}, \overline{\kappa}_{ti}]$. We define the *two-stage piecewise-constant decision rule* (PCDR) as follows:

$$\Theta_{ti}(\boldsymbol{\xi}^t, \beta_t) = \begin{cases} \underline{\kappa}_{ti} & \Upsilon_{ti}(\boldsymbol{\xi}^t, \beta_t^i) \in \mathcal{K}_{t1} \\ \underline{\kappa}_{ti} + 1 & \Upsilon_{ti}(\boldsymbol{\xi}^t, \beta_t^i) \in \mathcal{K}_{t2} \\ \vdots & \vdots \\ \overline{\kappa}_{ti} & \Upsilon_{ti}(\boldsymbol{\xi}^t, \beta_t^i) \in \mathcal{K}_{tJ}, \end{cases}$$

where $\mathcal{K}_{tj} \subset \mathbb{R}, j \in [J]$ are disjoint sets with $\bigcup_{j\in[J]} \mathcal{K}_{tj} = \mathbb{R}$, and $\Upsilon_{ti}(\boldsymbol{\xi}^t, \beta_t^i) : \mathbb{R}^{\ell_t} \times \mathbb{R}^{K_t} \to \mathbb{R}$ are functions defining the policy. We remark that the PCDR is presented for each *individual* variable $x_{ti}^{\mathrm{s}}(\boldsymbol{\xi}^t)$, and has its own decision variables β_t^i. Semantically, the PCDR partitions the set $\mathbb{R}$ into J subsets, and then assigns an integer value to each partition.

A special case of PCDRs can be defined by restricting the form of all mappings $\Upsilon_{ti}(\boldsymbol{\xi}^t, \beta_t^i)$ to be a linear function of the history. In the following we show that such a decision rule results in a model that is structurally very similar to (4), thus it can be solved via the C&CG method. Let $\Upsilon_{ti}(\boldsymbol{\xi}^t, \beta_t^i) = \Upsilon_{ti}^{\mathrm{A}}(\boldsymbol{\xi}^t)^\top \beta_t^i$ be an affine function of basis functions $\Upsilon_{ti}^{\mathrm{A}}(\boldsymbol{\xi}^t)$, and $\mathcal{K}_{tj} = (a_{tj}, b_{tj}], j \in [J]$ be intervals (in practice $\mathbb{R}$ is replaced by $[-M, M]$ for a sufficiently large M). Then problem (3) becomes:

$$\nu^{\text{2S-PCDR}} = \min_{x_1, \beta} \left\{ c_1^\top x_1 + \mathcal{SP}^{\text{PCDR}}(x_1, \beta) \;\middle|\; x_1 \in X_1, \quad \beta_t \in \mathbb{R}^{q_t \times K_t}, t \in [2,T] \right\},$$

where:

$$\mathcal{SP}^{\text{PCDR}}(x_1, \beta) = \max_{\boldsymbol{\xi}^T \in \Xi} \sum_{t\in[T]} \sum_{i\in[q_t]} c_{ti}^{\mathrm{s}} x_{ti}^{\mathrm{s}} + \min_{x^{\mathrm{r}} \in \mathcal{X}(x_1, \beta, \boldsymbol{\xi}^T)} \sum_{t\in[2,T]} c_t^{\mathrm{r}\top} x_i^{\mathrm{r}} \tag{7a}$$

$$\text{s.t.} \sum_{j\in[J]} (\underline{\kappa}_{ti} + j - 1)\upsilon_{tij} = x^{\mathrm{s}}_{ti} \qquad t\in[2,T], i\in[q_t] \tag{7b}$$

$$\sum_{j\in[J]} \omega_{tij} = \Upsilon^{\mathrm{A}}_{ti}(\boldsymbol{\xi}^t)^\top \beta^i_t \qquad t\in[2,T], i\in[q_t] \tag{7c}$$

$$(a_{tj}+\epsilon)\upsilon_{tij} \le \omega_{tij} \le b_{tj}\upsilon_{tij} \qquad t\in[2,T], i\in[q_t], j\in[J] \tag{7d}$$

$$\sum_{j\in[J]} \upsilon_{tij} = 1 \qquad t\in[2,T], i\in[q_t] \tag{7e}$$

$$\upsilon_{tij}\in\{0,1\} \qquad t\in[2,T], i\in[q_t], j\in[J]. \tag{7f}$$

The PCDR is modeled using the auxiliary variables $\upsilon_{tij}, \omega_{tij}$ and constraints (7b)-(7f) (ϵ is added to lower bound on ω_{tij} to ensure that partitions $\mathcal{K}_{tj}$ are disjoint). Note that, similar to Remark 1, if recourse decision variables $x^{\mathrm{r}}_t, t\in[2,T]$ are continuous, we can reformulate problem (7) as a monolithic mixed-integer linear program. Additionally, the arguments discussed in Section 2.2 show that the C&CG algorithm converges asymptotically to the optimal solution of $\nu^{\text{2S-PCDR}}$.

3. Dual Bounding

Our goal is to introduce a new dual problem that provides a lower bound for MSARO problems with mixed-integer recourse. Note that, the linear programming (LP) dual of continuous MSARO problems has been studied before (Kuhn et al. 2011), showing that even then, establishing strong duality is not straightforward. Recall that a random vector $\boldsymbol{\xi}\in\mathbb{R}^n$ is characterized by its probability space $(\mathbb{R}^n, \mathcal{F}, \mathbb{P})$, where $\mathcal{F}$ is a collection of subsets of $\mathbb{R}^n$ and $\mathbb{P}:\mathcal{F}\to[0,1]$ is a probability measure assigning a value between 0 and 1 to each element of $F\in\mathcal{F}$. We define the support of $\boldsymbol{\xi}$ as the smallest closed subset Ξ of $\mathbb{R}^n$ such that $\mathbb{P}(\Xi)=1$. Obtaining a strong LP dual is challenging due to a fundamental aspect of MSARO: the lack of knowledge about the true probability distribution, thus $\mathbb{P}$. To see the importance of $\mathbb{P}$ and its impact on the dual, we formulate the LP dual of model (1), with no integer variables (just to demonstrate the concept) for an *assigned* probability distribution. Let $\pi^{(1b)}(\boldsymbol{\xi}^T)$, $\pi^{(1c)}_t(\boldsymbol{\xi}^t)$, $\pi^{(1d)}_t(\boldsymbol{\xi}^t)$ be the LP dual variables associated with constraints (1b), (1c) and (1d), respectively, and $p:\Xi^t\to\mathbb{R}^+$ be the probability density function of $\boldsymbol{\xi}^t$ with respect to $\mathbb{P}$ (where we assume that the former exists whenever the latter does). First we multiply constraints (1b)-(1d) by their associated $p(\boldsymbol{\xi}^t)$, and then apply the standard LP duality theory. The LP dual of continuous MSARO for $\mathbb{P}$ is:

$$\begin{aligned}\mathcal{L}^{\mathrm{LP}}(\mathbb{P}) = \max\ & \mathbb{E}_{\boldsymbol{\xi}^T\sim\mathbb{P}} \sum_{t\in[T]} \left[b_t(\boldsymbol{\xi}^t)^\top \pi^{(1c)}_t(\boldsymbol{\xi}^t) + d_t(\boldsymbol{\xi}^t)^\top \pi^{(1d)}_t(\boldsymbol{\xi}^t)\right] \\ \text{s.t.}\ & \mathbb{E}_{\boldsymbol{\xi}^T\sim\mathbb{P}}\left[\pi^{(1b)}(\boldsymbol{\xi}^T)\right] = 1 \\ & \mathbb{E}_{\boldsymbol{\xi}^T\sim\mathbb{P}}\left[-c_1^\top \pi^{(1b)}(\boldsymbol{\xi}^T) + B_2(\boldsymbol{\xi}^2)^\top \pi^{(1c)}_2(\boldsymbol{\xi}^2)\right] + D_1\pi^{(1d)}_1 = \mathbf{0} \\ & \mathbb{E}_{\boldsymbol{\xi}^T\sim\mathbb{P}}\left[-c_t(\boldsymbol{\xi}^t)^\top \pi^{(1b)}(\boldsymbol{\xi}^T) \mid \boldsymbol{\xi}^t\right] + p(\boldsymbol{\xi}^t) A_t(\boldsymbol{\xi}^t)^\top \pi^{(1c)}_t(\boldsymbol{\xi}^t) + \end{aligned}$$

$$p(\boldsymbol{\xi}^{t+1})B_{t+1}(\boldsymbol{\xi}^{t+1})^\top \pi_{t+1}^{(1c)}(\boldsymbol{\xi}^{t+1}) + p(\boldsymbol{\xi}^t)D_t(\boldsymbol{\xi}^t)^\top \pi_t^{(1d)}(\boldsymbol{\xi}^t) = \mathbf{0} \quad t \in [2, T-1], \boldsymbol{\xi}^t \in \Xi^t$$

$$-p(\boldsymbol{\xi}^T)c_T(\boldsymbol{\xi}^T)^\top \pi^{(1b)}(\boldsymbol{\xi}^T) + p(\boldsymbol{\xi}^T)A_T(\boldsymbol{\xi}^T)^\top \pi_T^{(1c)}(\boldsymbol{\xi}^T) + p(\boldsymbol{\xi}^T)D_T(\boldsymbol{\xi}^T)^\top \pi_T^{(1d)}(\boldsymbol{\xi}^T) = \mathbf{0} \quad \boldsymbol{\xi}^T \in \Xi$$

$$\pi^{(1b)}(\boldsymbol{\xi}^T) \geq \mathbf{0} \quad \boldsymbol{\xi}^T \in \Xi$$

$$\pi_t^{(1c)}(\boldsymbol{\xi}^t) \leq \mathbf{0} \quad t \in [2, T], \boldsymbol{\xi}^t \in \Xi^t$$

$$\pi_t^{(1d)}(\boldsymbol{\xi}^t) \leq \mathbf{0} \quad t \in [T], \boldsymbol{\xi}^t \in \Xi^t.$$

Although we have used an assigned probability measure $\mathbb{P}$ in the derivation of the LP dual, as mentioned, the assumption in robust optimization is that we do not know the true probability distribution. As discussed by Kuhn et al. (2011), when this dual bound is used for robust optimization problems, the quality of the bound can be drastically different based on the choice of the assigned measure $\mathbb{P}$. Let $\mathcal{P} = \{\mathbb{P} \mid p^{\mathbb{P}}(\boldsymbol{\xi}^T) > 0, \boldsymbol{\xi}^T \in \Xi\}$, an ambiguity set including all probability measures with support Ξ, i.e., every $\mathbb{P} \in \mathcal{P}$ has a density function assigning a strictly positive value to all $\boldsymbol{\xi}^T \in \Xi$. The following DRO model returns the best lower bound over $\mathcal{P}$:

$$\sup_{\mathbb{P} \in \mathcal{P}} \mathcal{L}^{\text{LP}}(\mathbb{P}).$$

This shows that, even in the continuous setting, there exists a trade-off between the quality of the LP bound from an assigned probability distribution and the computational effort of searching over the entire ambiguity set. To the best of our knowledge, such a bounding problem has not been studied before. In what follows, we introduce the *nonanticipative* (NA) dual problem of an MSARO with *mixed-integer* recourse, which relies on the use of linear decision rules in a Lagrangian dual of problem (1), and propose to solve its associated DRO model to obtain strong bounds.

3.1. Nonanticipative Dual of the MSARO

The NA dual is based on a reformulation of the MSARO problem where we create a copy of every decision variable for every realization and explicitly enforce nonanticipativity. To this end, we introduce the copy variables $y(\boldsymbol{\xi}^T) = (y_1(\boldsymbol{\xi}^T), \ldots, y_T(\boldsymbol{\xi}^T))$ as perfect information variables depending on the entire realization $\boldsymbol{\xi}^T = (\xi_1, \ldots, \boldsymbol{\xi}_T)$. We can reformulate the MSARO problem (1) as:

$$\min \quad z \tag{8a}$$

$$\text{s.t.} \quad \sum_{t \in [T]} c_t(\boldsymbol{\xi}^t)^\top y_t(\boldsymbol{\xi}^T) \leq z \qquad \boldsymbol{\xi}^T \in \Xi \tag{8b}$$

$$A_t(\boldsymbol{\xi}^t)y_t(\boldsymbol{\xi}^T) + B_t(\boldsymbol{\xi}^t)y_{t-1}(\boldsymbol{\xi}^T) \leq b_t(\boldsymbol{\xi}^t) \qquad t \in [T],\ \boldsymbol{\xi}^T \in \Xi \tag{8c}$$

$$D_t(\boldsymbol{\xi}^t)y_t(\boldsymbol{\xi}^T) \leq d_t(\boldsymbol{\xi}^t) \qquad t \in [T],\ \boldsymbol{\xi}^T \in \Xi \tag{8d}$$

$$y_t(\boldsymbol{\xi}^T) = y_t(\boldsymbol{\xi}'^T) \qquad t \in [T],\ \boldsymbol{\xi}^T, \boldsymbol{\xi}'^T \in \Xi \text{ with } \boldsymbol{\xi}^t = \boldsymbol{\xi}'^t \tag{8e}$$

$$y_t(\boldsymbol{\xi}^T) \in \mathbb{R}^{n_t - n_t^{\mathrm{i}}} \times \mathbb{Z}^{n_t^{\mathrm{i}}} \qquad t \in [T],\ \boldsymbol{\xi}^T \in \Xi. \tag{8f}$$

Constraints (8e) are *nonanticipativity constraints* which ensure that at stage t for every partial realization of $\boldsymbol{\xi}^T$, the decisions made are consistent (i.e., the decisions made in all realizations sharing the history $\boldsymbol{\xi}^t$ are the same). Before introducing the NA dual derived from (8), we present the following lemma, which we use to reformulate the NA constraints (proof in Appendix A).

Lemma 1. *For given $\mathbb{P} \in \mathcal{P}$, constraints (8e) are equivalent to the following:*

$$y_t(\boldsymbol{\xi}^T) = \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}}\left[y_t(\boldsymbol{\xi}'^T) \;\middle|\; \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t \right], \quad t \in [T],\ \boldsymbol{\xi}^T \in \Xi. \tag{9}$$

Denote by $Y(\boldsymbol{\xi}^T)$ the scenario feasibility space defined by constraints (8b)-(8d) and (8f) for all $t \in [T]$ and some $\boldsymbol{\xi}^T$. For an assigned probability measure $\mathbb{P}$ and given dual functionals $\lambda_t(\cdot) : \mathbb{R}^{\ell^t} \to \mathbb{R}^{n_t}$, relaxing the NA constraints (9) results in the following NA Lagrangian relaxation problem:

$$\mathcal{L}^{\mathrm{NA}}_{\mathrm{LR}}(\mathbb{P}, \lambda_1(\cdot), \ldots, \lambda_T(\cdot)) = \min\ z + \sum_{t \in [T]} \mathbb{E}_{\boldsymbol{\xi}^T \sim \mathbb{P}}\left[\lambda_t(\boldsymbol{\xi}^T)^\top \left(y_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}}\left[y_t(\boldsymbol{\xi}'^T) \middle| \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t \right] \right) \right] \tag{10a}$$

$$\text{s.t. } \left(z, y_1(\boldsymbol{\xi}^T), \ldots, y_T(\boldsymbol{\xi}^T) \right) \in Y(\boldsymbol{\xi}^T) \qquad \boldsymbol{\xi}^T \in \Xi. \tag{10b}$$

which yields a lower bound for the original MSARO problem. The NA Lagrangian dual problem aims to find the best bound among all such Lagrangian relaxation bounds:

$$\mathcal{L}^{\mathrm{NA}}(\mathbb{P}) = \max_{\lambda_1(\cdot), \ldots, \lambda_T(\cdot)} \ \mathcal{L}^{\mathrm{NA}}_{\mathrm{LR}}(\mathbb{P}, \lambda_1(\cdot), \ldots, \lambda_T(\cdot)). \tag{11}$$

The following proposition shows that, regardless of the choice of $\mathbb{P}$, $\mathcal{L}^{\mathrm{NA}}(\mathbb{P})$ is an exact dual bound for *continuous* MSARO problems (proof in Appendix A).

Proposition 2. *Let $\mathbb{P}$ be any probability measure in $\mathcal{P}$. For continuous MSARO problems, $\mathcal{L}^{\mathrm{NA}}(\mathbb{P})$ is a strong dual.*

We remark that our construction of the dual postulates that we multiply the constraints with a density function whose support is Ξ. Therefore, the strictly positive density property of $\mathbb{P} \in \mathcal{P}$ is necessary for the exactness of our formulation. Proposition 2 suggests that the solution of problem (11) gives an exact dual bound for (8) (hence for (1)) if all decision variables are continuous. Although this bound is not necessarily exact in the case of mixed-integer recourse, the potential of leveraging the literature of multistage stochastic programming in achieving a dual bound for MSARO is quite appealing.

The objective function of the Lagrangian relaxation problem (10) contains (conditional) expectations of decision variables $y_t(\boldsymbol{\xi}^T)$, which is computationally challenging. Because of our initial assumption of relatively complete recourse (Assumption 1), and the fact that $X_t(\boldsymbol{\xi}^t)$ are bounded (Assumption 2), so is $Y(\boldsymbol{\xi}^T)$ and thanks to Lemma 1 of (Daryalal et al. 2021) we can replace the expectation term for t in the objective function (10a) with:

$$\mathbb{E}_{\boldsymbol{\xi}^T \sim \mathbb{P}}\Big[\Big(\lambda_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}}\Big[\lambda_t(\boldsymbol{\xi}'^T) \big| \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t\Big]\Big)^\top y_t(\boldsymbol{\xi}^T)\Big].$$

For given $\lambda_t(\boldsymbol{\xi}^T)$, this exchange allows us to compute the coefficients of $y_t(\boldsymbol{\xi}^T)$ in the Lagrangian relaxation problem. Still, the optimal form of the dual functionals $\lambda_t(\cdot)$ need to be determined, making the problem (11) intractable. In the next section, we restrict these Lagrangian multipliers to follow LDRs and obtain a restricted dual problem with decision variables of finite dimension. Furthermore, this new dual problem is amenable to well-known solution techniques from the literature of two-stage stochastic programming, such as sample average approximation, which are designed to approximately solve a problem with expectation in the objective function.

3.2. Lagrangian Dual Decision Rules

The Lagrangian multipliers $\lambda_t(\boldsymbol{\xi}^T)$ are infinite-dimensional. One way to obtain finite-dimensional variables is to enforce LDRs on the Lagrangian multipliers of problem (11) with a fixed $\mathbb{P}$, referred to as Lagrangian dual decision rules (LDDRs). For a set of pre-determined basis functions $\Psi_t : \Xi \to \mathbb{R}^{n_t \times K_t}$ and LDR decision variables $\alpha_t \in \mathbb{R}^{K_t}$, we restrict the form of $\lambda_t(\boldsymbol{\xi}^T)$ at stage t as follows:

$$\lambda_t(\boldsymbol{\xi}^T) = \Psi_t(\boldsymbol{\xi}^T)\alpha_t,$$

which gives us a restricted NA Lagrangian dual problem with respect to $\mathbb{P}$:

$$\mathcal{L}_R^{\mathrm{NA}}(\mathbb{P}) = \max_{\alpha_1,\dots,\alpha_T} \quad \mathcal{L}_{\mathrm{LR}}^{\mathrm{NA}}(\mathbb{P}, \Psi_1(\boldsymbol{\xi}^T)\alpha_1, \dots, \Psi_T(\boldsymbol{\xi}^T)\alpha_T). \tag{12}$$

Since problem (12) is a restriction of (11), we have $\mathcal{L}_R^{\mathrm{NA}}(\mathbb{P}) \le \mathcal{L}^{\mathrm{NA}}(\mathbb{P})$. An immediate question is whether by this restriction the bound gets weaker. More importantly, whether the choice of $\mathbb{P}$ affects the quality of the bound $\mathcal{L}_R^{\mathrm{NA}}(\mathbb{P})$. Using Lemma 2 in (Daryalal et al. 2021), the primal characterization of $\mathcal{L}_R^{\mathrm{NA}}(\mathbb{P})$ is:

$$\min \quad z \tag{13a}$$

$$\text{s.t.} \quad \big(z, y_1(\boldsymbol{\xi}^T), \dots, y_T(\boldsymbol{\xi}^T)\big) \in \operatorname{conv}(Y(\boldsymbol{\xi}^T)) \qquad \boldsymbol{\xi}^T \in \Xi \tag{13b}$$

$$\mathbb{E}_{\boldsymbol{\xi}^T \sim \mathbb{P}}\left[\Psi_t(\boldsymbol{\xi}^T)^\top \Big(y_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}}\big[y_t(\boldsymbol{\xi}'^T) \mid \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t\big]\Big)\right] = \mathbf{0} \qquad t \in [T] \tag{13c}$$

Comparing the problems (13) and (31), it is clear that for an assigned $\mathbb{P}$, (13) is a relaxation of (31), as constraints (13c) are an aggregation of (31c). Consequently, unlike $\mathcal{L}^{\mathrm{NA}}(\mathbb{P})$, even for continuous

MSARO, $\mathcal{L}_R^{\rm NA}(\mathbb{P})$ is not a strong bound. Furthermore, due to constraints (13c), the strength of the restricted NA Lagrangian dual bound depends on the choice of probability measure $\mathbb{P}$. Similar to the discussion on the LP dual at the beginning of the section, this observation motivates us to optimize over a set of probability distributions and find the best bound among all possible $\mathbb{P}$ and LDR variables α. As such, we propose to solve the following DRO problem with ambiguity set $\mathcal{P}$:

$$\nu_R^{\rm NA} = \sup_{\mathbb{P}\in\mathcal{P}} \; \mathcal{L}_R^{\rm NA}(\mathbb{P}). \tag{14}$$

We remark that (14) is a sup-max-min problem and is not in the familiar form of min-max-min that DRO problems are often represented with. However, as the notion of optimization over a family of probability distributions fits the field of DRO, we still refer to (14) as a DRO problem, with a slight abuse of terminology.

In linear and mixed-integer programming, strict inequalities such as the ones required for $\mathbb{P}\in\mathcal{P}$ (the strictly positive density property for all $\boldsymbol{\xi}^T\in\Xi$) often cause numerical and theoretical difficulties, thus are not desirable. To avoid these inequalities, in the following discussions we modify the DRO model to improve its numerical behavior. Denote by $\bar{\mathcal{P}}$ a superset of $\mathcal{P}$ that also admits distributions with support $\bar{\Xi}\subset\Xi$, i.e., allows $p(\boldsymbol{\xi}^T)=0$ for some $\boldsymbol{\xi}^T\in\Xi$. Consider the problem:

$$\bar{\nu}_R^{\rm NA} = \max_{\mathbb{P}\in\bar{\mathcal{P}}} \; \mathcal{L}_R^{\rm NA}(\mathbb{P}). \tag{15}$$

Because (15) is a relaxation of (14), it yields an upper bound for $\nu_R^{\rm NA}$. Thus, it does not immediately follow that such a bound is a valid lower bound for the MSARO problem optimal value $\nu^\star$. We show that problem (15) indeed leads to a valid dual (lower) bound for the MSARO problem (1) (see Appendix A for proof).

PROPOSITION 3. *$\bar{\nu}_R^{\rm NA}$ is a valid lower bound for $\nu^\star$.*

We refer to the DRO problem (15) as the *restricted NA dual*. It can be written explicitly as:

$$\bar{\nu}_R^{\rm NA} = \max_{\mathbb{P},\alpha} \quad \mathcal{Q}(\mathbb{P},\alpha) \tag{16a}$$
$$\text{s.t.} \quad \alpha_t\in\mathbb{R}^{K_t} \quad t\in[T] \tag{16b}$$
$$\mathbb{P}\in\bar{\mathcal{P}}, \tag{16c}$$

where

$$\mathcal{Q}(\mathbb{P},\alpha) = \min \quad z + \sum_{t\in[T]} \mathbb{E}_{\boldsymbol{\xi}^T\sim\mathbb{P}}\left[\left(\left(\Psi_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T\sim\mathbb{P}}\left[\Psi_t(\boldsymbol{\xi}'^T)\middle|\boldsymbol{\xi}'^t=\boldsymbol{\xi}^t\right]\right)\alpha_t\right)^\top y_t(\boldsymbol{\xi}^T)\right] \tag{17a}$$
$$\text{s.t.} \quad \left(z, y_1(\boldsymbol{\xi}^T),\ldots,y_T(\boldsymbol{\xi}^T)\right)\in\operatorname{conv}(Y(\boldsymbol{\xi}^T)) \qquad \boldsymbol{\xi}^T\in\Xi. \tag{17b}$$

Problem (17) is a *two-stage stochastic program* (2SP) and can benefit from its rich literature. In the next section, we build on well-known stochastic programming techniques to design a decomposition method for the restricted NA dual (16).

3.3. Solving the Restricted NA Dual Problem

For a discrete set of scenarios and given α, the expectation terms in the objective function of $\mathcal{Q}(\mathbb{P},\alpha)$ in (17), i.e., the coefficients of the y variables, can be calculated. In this section, we use this fact to solve or approximate the restricted NA dual problem. We first present a solution approach designed for finite discrete uncertainty sets. Then we show how it can be used for a general uncertainty set.

3.3.1. Discrete Uncertainty Set In a discrete uncertainty set with realizations $\xi^T \in \Xi$, $\bar{\mathcal{P}}$ can be modeled by a set of (in)equalities, that is $\bar{\mathcal{P}} = \left\{\rho \in \mathbb{R}_+^{|\Xi|} \mid \mathbf{1}^\top \rho = 1\right\}$, where a vector $\rho^{\mathbb{P}} \in \bar{\mathcal{P}}$ characterizes the probability measure $\mathbb{P}$ such that for all $\xi^T \in \Xi$, $\rho^{\mathbb{P}}_{\xi^T}$ is the probability of scenario ξ^T with respect to $\mathbb{P}$. So, the expectation term in the objective function (17a) can be written as:

$$\sum_{t\in[T]} \sum_{\xi^T\in\Xi} \rho^{\mathbb{P}}_{\xi^T} \bigg(\Big(\Psi_t(\xi^T) - \sum_{\substack{\xi'^T\in\Xi: \\ \xi'^t=\xi^t}} \rho^{\mathbb{P}}_{\xi'^T} \Psi_t(\xi'^T)\Big)\alpha_t \bigg)^\top y_t(\xi^T).$$

Let $\Psi_{tk}(\xi^T)$ be the k^{th} column of the matrix $\Psi_t(\xi^T)$. Then $\mathcal{Q}(\mathbb{P},\alpha)$ can be expressed as:

$$\begin{aligned} \min\ \ & z + \sum_{t\in[T]} \sum_{\xi^T\in\Xi} \rho^{\mathbb{P}}_{\xi^T} \bigg(\sum_{k\in[K_t]} \Big(\Psi_{tk}(\xi^T) - \sum_{\substack{\xi'^T\in\Xi: \\ \xi'^t=\xi^t}} \rho^{\mathbb{P}}_{\xi'^T} \Psi_{tk}(\xi'^T)\Big)\alpha_{tk} \bigg)^\top y_t(\xi^T) \\ \text{s.t.}\ \ & \big(z, y_1(\xi^T), \ldots, y_T(\xi^T)\big) \in Y(\xi^T) \quad \xi^T \in \Xi. \end{aligned}$$

Let $\gamma_{tk\xi^T} = \rho^{\mathbb{P}}_{\xi^T}\alpha_{tk}$ and $\beta_{t\xi^T} = \sum_{k\in[K_t]} \Big(\gamma_{tk\xi^T}\Psi_{tk}(\xi^T) - \sum_{\substack{\xi'^T\in\Xi: \\ \xi'^t=\xi^t}} \rho^{\mathbb{P}}_{\xi'^T}\gamma_{tk\xi^T}\Psi_{tk}(\xi'^T)\Big)$. Then with a change of variables in the restricted NA dual (16) we have:

$$\begin{aligned} \bar{\nu}_R^{\text{NA}} = \max_{\rho,\alpha,\gamma,\beta}\ & \mathcal{Q}(\beta) && && \text{(18a)} \\ \text{s.t.}\ & \sum_{\xi^T\in\Xi} \rho_{\xi^T} = 1 && && \text{(18b)} \\ & \gamma_{tk\xi^T} = \rho_{\xi^T}\alpha_{tk} && t\in[T], k\in[K_t], \xi^T\in\Xi && \text{(18c)} \\ & \beta_{t\xi^T} = \sum_{k\in[K_t]} \Big(\gamma_{tk\xi^T}\Psi_{tk}(\xi^T) - \sum_{\substack{\xi'^T\in\Xi: \\ \xi'^t=\xi^t}} \rho_{\xi^T}\gamma_{tk\xi^T}\Psi_{tk}(\xi'^T)\Big) && t\in[T], \xi^T\in\Xi && \text{(18d)} \\ & \rho \in \mathbb{R}_+^{|\Xi|} && && \text{(18e)} \\ & \alpha_t \in \mathbb{R}^{K_t},\ \gamma \in \mathbb{R}^{T\times K_t\times|\Xi|},\ \beta \in \mathbb{R}^{n_t\times T\times|\Xi|} && t\in[T] && \text{(18f)} \end{aligned}$$

where

$$\mathcal{Q}(\beta) = \min\Big\{ z + \sum_{t\in[T]} \sum_{\xi^T\in\Xi} \beta^\top_{t\xi^T} y_t(\xi^T) \ \Big|\ \big(z, y_1(\xi^T), \ldots, y_T(\xi^T)\big) \in Y(\xi^T),\ \xi^T \in \Xi \Big\}. \quad (19)$$

Model (18) can be solved via a cutting-plane method in which $\mathcal{Q}(\beta)$ is approximated by a set of linear inequalities. At each iteration, we solve the following bilinear program as the master problem:

$$\max_{p,\alpha,\gamma,\beta} \left\{ \eta \;\middle|\; (18b) - (18f),\; (\eta,\beta) \in \mathcal{H} \right\}, \tag{20}$$

where η is an auxiliary variable representing $\mathcal{Q}(\beta)$, and $\mathcal{H}$ is a set described by optimality cuts approximating $\mathcal{Q}(\beta)$. Note that, as β only parameterizes the objective function of $\mathcal{Q}(\beta)$, i.e., it does not impact the feasibility space, there is no need for feasibility cuts. With $(\hat{\eta}, \hat{\beta})$ returned from solving the master problem (20), we solve the subproblem (19) to compute $\mathcal{Q}(\hat{\beta})$, resulting in $\hat{y}_t(\xi^T)$ as the optimal solution. If $\hat{\eta} \leq \mathcal{Q}(\hat{\beta})$, we have found the optimal solution of the restricted NA dual. Otherwise we add the following optimality cut to the master problem:

$$\eta \leq \mathcal{Q}(\hat{\beta}) + \sum_{t\in[T]} \sum_{\xi^T \in \Xi} \left(\beta_{t\xi^T} - \hat{\beta}_{t\xi^T}\right)^\top \hat{y}_t(\xi^T). \tag{21}$$

This procedure continues until no more optimality cuts are found. The objective function of the subproblem, $\mathcal{Q}(\beta)$, is a concave function in β (pointwise minimum of linear functions with respect to β). Each cut (21) is a hyperplane approximating the subproblem from above. The cutting-plane method iteratively finds improving approximations of $\mathcal{Q}(\beta)$. Enforcing large lower and upper bounds on LDR variables α can ensure that the algorithm converges to ϵ-optimality.

Discrete Uncertainty in Continuous MSARO

As a special case, we study continuous MSARO with discrete uncertainty. Our goal here is to use this particular structure and derive a monolithic formulation as an alternative to the cutting-plane algorithm, to leverage off-the-shelf solvers. Denote by u_{ξ^T}, $v_{t\xi^T}$ and $w_{t\xi^T}$, the dual variables associated with the set of constraints described by $Y(\xi^T)$ for given $\xi^T \in \Xi$ (corresponding to (8b), (8c) and (8d), respectively). The LP dual of the inner minimization problem (19), i.e., the subproblem of the cutting-plane algorithm, is:

$$\mathcal{Q}^D(\beta) = \max \quad \sum_{t\in[T]} \sum_{\xi^T\in\Xi} b_t(\xi^T)^\top v_{t\xi^T} + \sum_{t\in[T]} \sum_{\xi^T\in\Xi} d_t(\xi^t)^\top w_{t\xi^T} \tag{22a}$$

$$\text{s.t.} \quad \sum_{\xi^T\in\Xi} u_{\xi^T} = 1 \tag{22b}$$

$$-c_T(\xi^T) u_{\xi^T} + A_T(\xi^T)^\top v_{T\xi^T} + D_{T\xi^T}^\top w_{T\xi^T} = \beta_{T\xi^T} \qquad \xi^T \in \Xi \tag{22c}$$

$$-c_t(\xi^t) u_{\xi^T} + A_t(\xi^t)^\top v_{t\xi^T} + D_t(\xi^t)^\top w_{t\xi^T} + B_{t+1}(\xi^{t+1})^\top v_{t+1,\xi^T} = \beta_{t\xi^T} \qquad t \in [T-1], \xi^T \in \Xi \tag{22d}$$

$$u_{\xi^T} \geq 0 \qquad \xi^T \in \Xi \tag{22e}$$

$$v_{t\xi^T}, w_{t\xi^T} \leq \mathbf{0} \qquad\qquad t \in [T], \xi^T \in \Xi. \tag{22f}$$

Merging the two maximization problems in (18), we get the monolithic *bilinear* program:

$$\bar{\nu}_R^{\mathrm{NA}} = \max \quad \sum_{t\in[T]} b_t(\xi^T)^\top v_{t\xi^T} + \sum_{t\in[T]} \sum_{\xi^T\in\Xi} d_t(\xi^t)^\top w_{t\xi^T} \tag{23a}$$

$$\text{s.t.} \quad (18b) - (18f) \tag{23b}$$

$$(22b) - (22f). \tag{23c}$$

Recall that, in adapting the C&CG algorithm for two-stage LDRs and Remark 1, we use KKT conditions to obtain a monolithic model of the subproblem and the bilinear terms in the constraints can be linearized by adding big-M constraints. The same technique would not lead to a linear reformulation of (18). Indeed, the source of nonlinearity in constraint (32e) and (32f) is the multiplication of a linear function and a dual variable. Thus, we are able to deal with them separately (and make sure that at least one of the two terms is evaluated to zero) and derive a linear model. However, in (18) the nonlinearity is rooted in the constraints of the outer maximization problem itself and enforcing the complementary slackness through separate constraints does not resolve it. Nevertheless, there is a large body of research on solution methods for bilinear problems that can be used in solving model (23). Furthermore, many nonlinear optimization solvers such as MOSEK (ApS 2022) include algorithms specifically tailored for large-scale quadratic problems and even a commercial MIP solver, Gurobi, has recently added a bilinear feature that solves non-convex quadratic optimization models to global optimality (Gurobi Optimization, LLC 2022).

3.3.2. Continuous Uncertainty Set If the uncertainty set of MSARO, Ξ, is not discrete, objective function (17a) includes the expectation of a nonsmooth concave function. The literature of two-stage stochastic programming addresses such a difficulty by means of sampling-based approaches that approximate the underlying problem. Sample average approximation (SAA) is a well-known method that replaces the expectations in the objective function of a 2SP with the average of a given sample and has favorable theoretical convergence results (see e.g., Shapiro et al. (2009)). In this section we show that (approximately) solving the 2SP (17) over a sample returns a valid dual bound for the original MSARO problem (1).

From the discussions in Section 3, for any $\mathbb{P} \in \bar{\mathcal{P}}$, we have:

$$\mathcal{L}_R^{\mathrm{NA}}(\mathbb{P}) \leq \bar{\nu}_R^{\mathrm{NA}} \leq \nu^\star. \tag{24}$$

The first inequality is valid because the bound $\bar{\nu}_R^{\mathrm{NA}}$ is the result of maximization over all $\mathbb{P} \in \bar{\mathcal{P}}$, and the second inequality is from Proposition 3. For a given finite sample $\Omega \subseteq \Xi$, we define a probability measure $\mathbb{P}_\Omega$ with $\mathbb{P}_\Omega(\Omega) = 1$ such that the density function assigns zero to any scenario not in Ω:

$$p^{\mathbb{P}_\Omega}(\boldsymbol{\xi}^T) = 0,\ \boldsymbol{\xi}^T \in \Xi \setminus \Omega,\ \text{and}\ p^{\mathbb{P}_\Omega}(\boldsymbol{\xi}^T) \geq 0,\ \boldsymbol{\xi}^T \in \Omega.$$

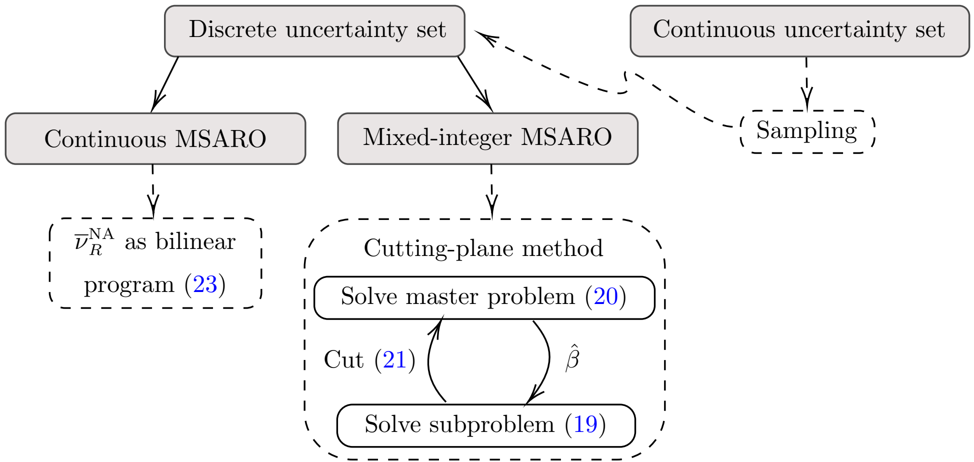

Figure 5 Solution methods for the restricted NA dual problem

Clearly $\mathbb{P}_\Omega \in \bar{\mathcal{P}}$, thus inequalities (24) hold for $\mathbb{P}_\Omega$, i.e., $\mathcal{L}_R^{\mathrm{NA}}(\mathbb{P}_\Omega) \leq \bar{\nu}_R^{\mathrm{NA}} \leq \nu^\star$, and $\mathcal{L}_R^{\mathrm{NA}}(\mathbb{P}_\Omega)$ is a valid dual bound on $\nu^\star$. As Ω is a finite set, we can represent $\mathbb{P}_\Omega$ with $\rho^{\mathbb{P}_\Omega} \in \mathbb{R}_+^{|\Omega|}$ such that $\mathbf{1}^\top \rho^{\mathbb{P}_\Omega} = 1$. Denote by $\mathcal{P}^\Omega$, the set of all such vectors ρ. Consequently, solving the DRO problem (18) with discrete uncertainty set Ω finds the best bound over all probability measures $\mathbb{P} \in \mathcal{P}^\Omega$, which is a valid dual bound on $\nu^\star$. Calling the optimal value of this new DRO problem $\bar{\nu}_R^{\mathrm{NA}}(\Omega)$, the following proposition establishes the validity of the bound $\bar{\nu}_R^{\mathrm{NA}}(\Omega)$.

Proposition 4. *For any given sample $\Omega \subseteq \Xi$, weak duality holds between $\bar{\nu}_R^{\mathrm{NA}}(\Omega)$ and $\nu^\star$, i.e.,*

$$\bar{\nu}_R^{\mathrm{NA}}(\Omega) \leq \bar{\nu}_R^{\mathrm{NA}} \leq \nu^\star.$$

So, for a continuous uncertainty set, we can get a discrete sample Ω which turns the problem into a case of discrete uncertainty. As such, we can use the arguments in Section 3.3.1 to obtain $\bar{\nu}_R^{\mathrm{NA}}(\Omega)$.

Figure 5 summarizes all the presented methods for solving the restricted NA dual model. For continuous MSARO with a discrete uncertainty set, the bilinear programming reformulation can be solved, while for mixed-integer MSARO, the devised cutting-plane method is applicable. For MSARO with a continuous uncertainty set, we use SAA, i.e., first take a sample and then approximately solve the restricted NA dual problem using the methods developed in the discrete setting.

3.4. Restricted Decomposable NA Dual Bound

In solving problem (18) using the cutting-plane method, we frequently optimize (19) to compute $\mathcal{Q}(\beta)$. This can become computationally demanding when there is a large number of scenarios (i.e., a large discrete uncertainty set or sample). Further, we cannot decompose $\mathcal{Q}(\beta)$ by the scenarios, as they are linked through the z variables. In this section, we present an alternative NA dual problem that can yield a potentially weaker bound than the one provided by (16), but is decomposable by scenarios, thus offering a computational advantage. With decision variables $y_t(\boldsymbol{\xi}^T)$ defined as before, we introduce copy variables $z(\boldsymbol{\xi}^T)$ and explicitly enforce them all to be equal. For an assigned probability measure $\mathbb{P}$ which has a strictly positive density function for every realization of the uncertainty set, the alternative NA reformulation for the MSARO problem (1) is:

$$
\begin{aligned}
\min \quad & \mathbb{E}_{\boldsymbol{\xi}^T \sim \mathbb{P}} \left[z(\boldsymbol{\xi}^T) \right] && && (25a)\\
\text{s.t.} \quad & \sum_{t \in [T]} c_t(\boldsymbol{\xi}^t)^\top y_t(\boldsymbol{\xi}^T) \le z(\boldsymbol{\xi}^T) && \boldsymbol{\xi}^T \in \Xi && (25b)\\
& z(\boldsymbol{\xi}^T) = \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}} \left[z(\boldsymbol{\xi}'^T) \right] && \boldsymbol{\xi}^T \in \Xi && (25c)\\
& (8c), (8d), (8f), (9). && && (25d)
\end{aligned}
$$

Together, (25b) and (25c) capture the semantics of the worst-case outcome, which is minimized in the objective function (25a). Relaxation of the NA constraints (9) and (25c) with the Lagrangian multipliers $\lambda_t^y(\cdot)$ and $\lambda^z(\cdot)$, leads to the decomposable NA Lagrangian dual problem $\mathcal{L}^{\text{DNA}}(\mathbb{P}) = \max_{\mathbb{P}, \lambda^y(\cdot), \lambda^z(\cdot)} \mathcal{L}_{LR}^{\text{DNA}}(\mathbb{P}, \lambda_1^y(\cdot), \dots, \lambda_T^y(\cdot), \lambda^z(\cdot))$, where

$$
\begin{aligned}
\mathcal{L}_{LR}^{\text{DNA}}(\mathbb{P}, \lambda_1^y(\cdot), \dots, \lambda_T^y(\cdot), \lambda^z(\cdot)) = \min \quad & \mathbb{E}_{\boldsymbol{\xi}^T \sim \mathbb{P}} \Big[\Big(1 + \lambda^z(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}} \left[\lambda^z(\boldsymbol{\xi}'^T) \right] \Big) z(\boldsymbol{\xi}^T) \Big] + \\
& \sum_{t \in [T]} \mathbb{E}_{\boldsymbol{\xi}^T \sim \mathbb{P}} \Big[\Big(\lambda_t^y(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}} \left[\lambda_t^y(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t \right] \Big)^\top y_t(\boldsymbol{\xi}^T) \Big] \\
\text{s.t.} \quad & \big(z(\boldsymbol{\xi}^T), y_1(\boldsymbol{\xi}^T), \dots, y_T(\boldsymbol{\xi}^T) \big) \in Y(\boldsymbol{\xi}^T) \qquad \boldsymbol{\xi}^T \in \Xi,
\end{aligned}
$$

with $Y(\boldsymbol{\xi}^T)$ the scenario feasibility space described by constraints (25b), (8c), (8d), and (8f). We remark that, in the first expectation term in the objective function the inner expectation is not conditional, since constraints (25c) (and accordingly their associated dual functions) are not defined at every stage $t \in [T]$ as decision variable $z(\boldsymbol{\xi}^T) \in \mathbb{R}$ captures the cost of an entire scenario at T. After substituting the decision rules $\lambda_t^y(\boldsymbol{\xi}^T) = \Psi_t(\boldsymbol{\xi}^T)\alpha_t^y, t \in [T]$ and $\lambda^z(\boldsymbol{\xi}^T) = \Psi_T(\boldsymbol{\xi}^T)\alpha^z$ in $\mathcal{L}^{\text{DNA}}(\mathbb{P})$, where $\alpha^z \in \mathbb{R}^{K_T}$, the restricted decomposable NA dual problem is:

$$
\nu_R^{\text{DNA}} = \max_{\mathbb{P}, \alpha^y, \alpha^z} \mathcal{L}_{LR}^{\text{DNA}}(\mathbb{P}, \Psi_1(\boldsymbol{\xi}^T)\alpha_1^y, \dots, \Psi_T(\boldsymbol{\xi}^T)\alpha_T^y, \Psi_T(\boldsymbol{\xi}^T)\alpha^z).
$$

This dual can be solved using the same methods developed in the previous sections for the non-decomposable problem. Due to the relaxation of the NA constraints on z variables, for given $\mathbb{P}$ and $\lambda^y(\cdot)$, subproblem $\mathcal{L}_{LR}^{\text{DNA}}(\cdot)$ is a relaxation of $\mathcal{L}_{LR}^{\text{NA}}(\cdot)$. Therefore $\nu_R^{\text{DNA}} \le \nu_R^{\text{NA}}$, i.e., ν_R^{NA} is a potentially stronger bound. However, the fact that $\mathcal{L}_{LR}^{\text{DNA}}(\cdot)$ is decomposable is highly desirable. Particularly for continuous uncertainty sets where we rely on sampling, and the quality of the bound varies based on the selected sample. Since the decomposable model can afford samples of larger sizes it can potentially yield better bounds compared to the bound obtained from non-decomposable model over a smaller sample. We explore this possibility in our numerical section.

4. Numerical Experiments

We evaluate the performance of the proposed bounding framework over multistage versions of three classical decision-making problems under uncertainty: (i) the newsvendor problem, (ii) the location-transportation problem, and (iii) the capital budgeting problem. Depending on their structure, each problem is solved by using the appropriate models and methods described in Sections 2 and 3, illustrating the applicability of the developed concepts to a large array of problem classes.

4.1. Benchmarks and Implementation Details

To assess the quality of the primal and dual bounds, we measure the relative distance of the bounds from the true optimal value when an exact solution of MSARO is available (in small-size instances). Otherwise, we report the optimality gap between the bounds obtained from the proposed methods, and compare it against a gap obtained via traditional bounding methods if one exists. In continuous problems, we consider LDRs (i.e., their application to all decision variables) as the benchmark for the primal decision rules. On the dual side, we use the perfect information (PI) bound (denoted by ν^{PI}) for comparison, which often can be conveniently evaluated for a general MSARO problem. The PI bound corresponds to the optimal objective value of the MSARO problem reformulated as in (8) without the nonanticipativity constraints (8e), i.e., it finds the cost of every scenario in the uncertainty set individually, and then selects the one with the worst-case cost.

The algorithms are implemented in Python and use the Gurobi Optimizer 9.5.1 (Gurobi Optimization, LLC 2022) as the MIP/bilinear program solver. The computational experiments are carried out on the Niagara supercomputer servers (Loken et al. 2010, Ponce et al. 2019). The programs stop if the running time reaches 10 hours, at which point we report the best observed lower/upper bound on the instance. As a common design choice for the basis functions of the LDRs, we use the uncertain parameters themselves, i.e., the standard basis functions. Any implementation nuances and enhancements used for improving the performance of the algorithms are discussed for each problem class in a dedicated section, along with the characteristics of the studied instances.

4.2. Robust Budgeted Newsvendor Problem

In this section, we extend the two-stage newsvendor problem studied in Xu and Hanasusanto (2021) to the multistage setting. In this problem, a decision-maker (the newsvendor) needs to order from a set of items to be sold (only) at the next decision stage, with the objective of maximizing the worst-case profit over the planning horizon. Let $d_{it}(\boldsymbol{\xi}^t)$ be the uncertain demand of item $i \in [I]$ at stage $t \in [2, T]$, c_i and s_i the purchase and shortage costs of item i, respectively, and r_i its sale price. To meet the customers' demands of stage t, at stage $t-1$ the decision-maker decides on the amounts to be ordered from each item, such that the total spending over the T stages does not exceed a predetermined budget of B. Denote by $x_{it}(\boldsymbol{\xi}^t)$ the decision variable for the amount of item i ordered at stage $t \in [T-1]$. The multistage multi-item budgeted newsvendor problem is:

$$
\begin{aligned}
\max\ & z && && \text{(27a)}\\
\text{s.t. } & z \le \sum_{i\in[I]}\sum_{t\in[2,T]} y_{it}(\boldsymbol{\xi}^t) && \boldsymbol{\xi}^T \in \Xi && \text{(27b)}\\
& y_{it}(\boldsymbol{\xi}^t) \le \big(r_i - c_i\big)x_{i,t-1}(\boldsymbol{\xi}^{t-1}) - r_i\big(x_{i,t-1}(\boldsymbol{\xi}^{t-1}) - d_i(\boldsymbol{\xi}^t)\big) && i\in[I], t\in[2,T], \boldsymbol{\xi}^t \in \Xi^t && \text{(27c)}\\
& y_{it}(\boldsymbol{\xi}^t) \le \big(r_i - c_i\big)x_{i,t-1}(\boldsymbol{\xi}^{t-1}) - s_i\big(d_i(\boldsymbol{\xi}^t) - x_{i,t-1}(\boldsymbol{\xi}^{t-1})\big) && i\in[I], t\in[2,T], \boldsymbol{\xi}^t \in \Xi^t && \text{(27d)}\\
& \sum_{i\in[I]}\sum_{t\in[T-1]} x_{it}(\boldsymbol{\xi}^t) \le B && \boldsymbol{\xi}^{T-1} \in \Xi^{T-1} && \text{(27e)}\\
& x_t(\boldsymbol{\xi}^t) \in \mathbb{R}^I_+ && t\in[T-1], \boldsymbol{\xi}^t \in \Xi^t. && \text{(27f)}
\end{aligned}
$$

where auxiliary variable $y_{it}(\boldsymbol{\xi}^t)$ captures the profit from item i at stage t, by means of (27c)- (27d). Constraints (27e) impose a budget of B over the order amounts throughout the planning horizon. The objective function is the worst-case profit of the newsvendor, modeled via (27a) and (27b).

4.2.1. Problem Instances Our instance generation loosely follows the procedure described in Ardestani-Jaafari and Delage (2021) for the two-stage robust newsvendor problem. r_i, s_i and c_i are drawn uniformly from the intervals $[140,160]$, $[80,90]$ and $[50,70]$, respectively. We consider a discrete uncertainty set modeled as a stagewise-dependent scenario tree with branching factor BR (i.e., every node of the tree prior to the leaves has BR many child nodes). The demand realizations $d_{it}(\boldsymbol{\xi}^t), i\in[I], t\in[2,T]$ at a child node are drawn uniformly from $[\mu_{it}-\sigma_{it}, \mu_{it}+\sigma_{it}]$, where μ_{it} and σ_{it} are uniformly drawn from the intervals $[20,40]$ and $[10,20]$. We have generated 26 small-size instances with $T\in[3,5]$, $I\in[2,5]$, $B\in\{100,150,200,250,300\}$, and $\text{BR}\in\{2,3,4,5,10\}$, such that the number of scenarios $|\Xi| = \text{BR}^{T-1}$ is less than 150. Additionally, we have generated 18 large-size instances with $T\in[4,8]$. For $T=4$, the number of items I lies in the set $[3,5]$, with a budget $B\in\{200,300\}$ and $\text{BR}\in\{10,15,20\}$. For $T\in[5,8]$, our instances have $I\in\{3,4\}$ items, budget of $B\in\{300,400,500,600\}$, and $\text{BR}\in[3,6]$, restricted to the cases with $|\Xi|\le 3000$.

4.2.2. Quality of the Bounds The small-size instances can be conveniently optimized by solving model (8) over all the scenarios in the uncertainty set. Using this, we can examine the quality of a primal/dual bound by measuring its relative distance to the optimal objective value $\nu^\star$. For this problem, all primal and dual problems are solved by the extensive form (i.e., monolithic) of their respective models. Figures 6 and 7 present the gap between the bound and the optimal value of the exact solution, defined as $100\big(\frac{\nu^\star-\nu^{(\cdot)}}{\nu^\star}\big)$ and $100\big(\frac{\nu^{(\cdot)}-\nu^\star}{\nu^\star}\big)$ for primal and dual bounds, respectively, and presented as a percentage (detailed results are given in the e-companion). In each figure, the solid bars depict the performance of the newly proposed bounds, while the hatched

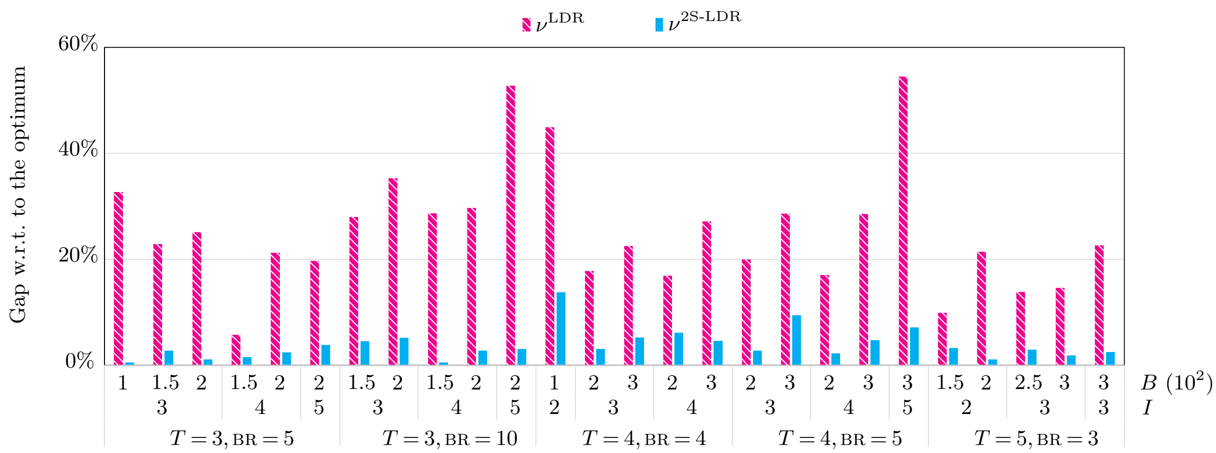

Figure 6 Quality of the bounds from LDRs and two-stage LDRs for small-size newsvendor instances

bars represent the benchmarks. The results show that $\nu^{\text{2S-LDR}}$ and $\bar{\nu}_R^{\text{NA}}$ outperform the benchmark bounds by orders of magnitude. More precisely, $\nu^{\text{2S-LDR}}$ on average achieves 84% improvement over ν^{LDR}, with reductions in relative distance ranging from 64% to 98%. Interestingly, the quality of the bound $\nu^{\text{2S-LDR}}$ remains rather stable with changes in the number of items I and budget B, compared to the drastic changes of ν^{LDR} with variations in the inputs. The notable performance of the two-stage LDRs for the newsvendor problem contrasted with the LDRs can be explained by the nature of the recourse variables. In model (27), $y_{it}(\boldsymbol{\xi}^t)$ determines the profit of a given realization at stage t for item i, which for the newsvendor problem is by definition nonlinear. In fact, $y_{it}(\boldsymbol{\xi}^t)$ is an auxiliary variable, defined to linearize the following net profit at stage t from item i:

$$r_i \min\left\{x_{i,t-1}(\boldsymbol{\xi}^{t-1}), d_{it}(\boldsymbol{\xi}^t)\right\} - c_i x_{i,t-1}(\boldsymbol{\xi}^{t-1}) - s_i \max\left\{d_{it}(\boldsymbol{\xi}^t) - x_{i,t-1}(\boldsymbol{\xi}^{t-1}), 0\right\}.$$

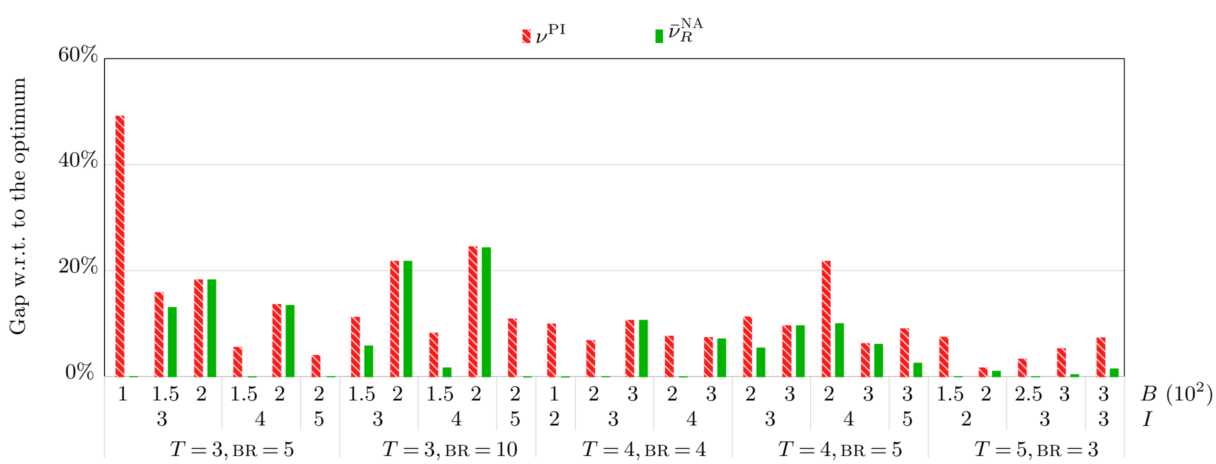

Figure 7 Quality of the bounds from the PI and LDDRs for small-size newsvendor instances

Therefore, for the newsvendor problem LDRs *always* return suboptimal decisions as they restrict the form of the nonlinear profit function to be affine, while two-stage LDRs allow them to take any form, giving them an immediate advantage over LDRs.

From the dual perspective, $\bar{\nu}_R^{\text{NA}}$ achieves an average improvement of 55% compared to ν^{PI}, and in 9 instances fully closes the gap. From the results of Figure 7, a common observation is that LDDRs return a bound of higher quality for smaller values of the ratio $\frac{B}{(T-1)\times I}$, which is an estimate of the average budget available per product at every stage. This trend suggests that the restricted NA dual bound performs better with tighter budget, when there is a higher dependency between stages, in which case the importance of intermediate decisions becomes more pronounced.

4.2.3. Optimality Gap For instances with larger number of scenarios, we compare the optimality gap from the benchmark methods ($\text{OPT}^{\mathtt{T}}$) with the gap obtained by applying the newly proposed methods, namely two-stage LDRs and LDDRs ($\text{OPT}^{\mathtt{N}}$):

$$\text{OPT}^{\mathtt{T}} = 100\Big(\frac{\nu^{\text{PI}} - \nu^{\text{LDR}}}{\nu^{\text{LDR}}}\Big), \quad \text{OPT}^{\mathtt{N}} = 100\Big(\frac{\bar{\nu}_R^{\text{NA}} - \nu^{\text{2S-LDR}}}{\nu^{\text{2S-LDR}}}\Big).$$

Table 1 presents the optimality gaps from the benchmark and proposed models (whose running times are provided in Appendix C). Achieving an average gap reduction of 82%, there is a considerable value in using the two-stage LDRs and LDDRs in devising policies for the multistage newsvendor problem. The optimality gaps $\text{OPT}^{\mathtt{N}}$ range from 4% to 46%, with some interesting cases such as instances 3, 5 and 14 where LDR policies can even lead to a profit loss in the worst-case. In

Table 1 Optimality gaps for larger instances of the newsvendor problem

Instance	T	BR	$\lvert\Xi\rvert$	I	B	Primal Bounds		Dual Bounds		Optimality Gap		Gap reduction
						ν^{LDR}	ν^{2S}	$\bar{\nu}_R^{\text{NA}}$	ν^{PI}	$\text{OPT}^{\mathtt{N}}$	$\text{OPT}^{\mathtt{T}}$	
1				3	200	5648.8	8142.9	8876.4	9353.0	9.0%	65.6%	86.3%
2				3	300	9143.5	13687.6	17853.0	17853.0	30.4%	95.3%	68.1%
3	4	10	1000	4	200	-104.4	440.0	642.4	919.0	46.0%	980.4%	95.3%
4				4	300	11029.5	15854.7	18400.7	18432.0	16.1%	67.1%	76.1%
5				5	300	724.9	6125.5	6748.6	7368.0	10.2%	916.5%	98.9%
6	4	15	3375	3	200	5111.3	7072.0	7646.6	8222.0	8.1%	60.9%	86.7%
7	4	20	8000	3	300	9030.9	13040.7	17615.0	17615.0	35.1%	95.1%	63.1%
8	5	5	625	3	300	11134.5	15053.4	16191.6	16680.0	7.6%	49.8%	84.8%
9				4	300	3651.5	9494.8	10553.8	10628.0	11.2%	191.1%	94.2%
10	5	6	1296	3	400	15013.4	20271.7	25157.0	25157.0	24.1%	67.6%	64.3%
11				3	400	15492.3	21457.9	28163.0	28163.0	31.2%	81.8%	61.8%
12	6	4	1024	4	400	7124.3	14805.2	15426.2	15473.0	4.2%	117.2%	96.4%
13				4	500	14445.4	24887.7	32946.9	32973.0	32.4%	128.3%	74.8%
14				3	300	24.3	1495.9	2137.3	2383.0	42.9%	9692.2%	99.6%
15	7	3	729	3	400	12994.7	17774.1	19554.5	19983.0	10.0%	53.8%	81.4%
16				4	400	3267.5	4303.1	5354.3	5427.0	24.4%	66.1%	63.0%
17	8	3	2187	3	500	16149.1	25321.4	30259.0	30259.0	19.5%	87.4%	77.7%
18				4	600	16892.3	27608.1	28717.7	28747.0	4.0%	70.2%	94.3%

the majority of the instances, both two-stage LDRs and LDDRs contribute to the improvement of the gap, although the primal side clearly has the larger impact, to the point that in five instances, the dual bound $\bar{\nu}_R^{\mathrm{NA}}$ and the benchmark ν^{PI} are the same. In some instances, the PI bound might already be strong enough so that using LDDRs does not make a tangible difference in strengthening it. This is, for instance, the case for instances 8, 12 and 18 that attain a small optimality gap without significantly improving ν^{PI}. As discussed for the small-size instances, depending on the number of stages, items, and available budget, in some instances the PI bound can indeed be weak. In such cases $\bar{\nu}_R^{\mathrm{NA}}$ can be potentially improved by a better design of the basis functions for LDDRs that better captures the correlation among all items imposed by the budget restriction.

4.3. Robust Location-Transportation Problem

The two-stage robust location-transportation problem studied by Zeng and Zhao (2013) is as follows. Given a set of I potential facilities with building cost f_i and unit capacity cost $a_i, i \in [I]$, we have to meet the uncertain demand of a set of customers J with unit transportation cost $c_{ij}, i \in [I], j \in [J]$. The goal is to decide which facilities to open and their initial capacities, such that the worst-case total cost of facility deployment and future transportation is minimized. Letting $d_{jt}(\boldsymbol{\xi}^t)$ be the demand of customer j at stage t, we define the MSARO extension of the problem:

$$
\begin{aligned}
\min \quad & z && && \text{(28a)}\\
\text{s.t.} \quad & z \geq \sum_{i\in[I]} f_i y_i + \sum_{t\in[T]}\sum_{i\in[I]} a_i s_{it}(\boldsymbol{\xi}^t) + \sum_{t\in[2,T]}\sum_{i\in[I]}\sum_{j\in[J]} c_{ij} x_{ijt}(\boldsymbol{\xi}^t) && \boldsymbol{\xi}^T \in \Xi && \text{(28b)}\\
& s_{i1} \leq K_i y_i && i \in [I] && \text{(28c)}\\
& s_{it}(\boldsymbol{\xi}^t) = s_{i,t-1}(\boldsymbol{\xi}^{t-1}) - \sum_{j\in[J]} x_{ijt}(\boldsymbol{\xi}^t) && i \in [I], t \in [2,T], \boldsymbol{\xi}^t \in \Xi^t && \text{(28d)}\\
& \sum_{i\in[I]} x_{ijt}(\boldsymbol{\xi}^t) \geq d_{jt}(\boldsymbol{\xi}^t) && j \in [J], t \in [2,T], \boldsymbol{\xi}^t \in \Xi^t && \text{(28e)}\\
& y \in \{0,1\}^I, \; s_1 \in \mathbb{R}_+^I && && \text{(28f)}\\
& s_t(\boldsymbol{\xi}^t) \in \mathbb{R}_+^I, \; x_t(\boldsymbol{\xi}^t) \in \mathbb{R}_+^{I\times J} && t \in [2,T], \boldsymbol{\xi}^t \in \Xi^t, && \text{(28g)}
\end{aligned}
$$

where y_i is a binary variable equal to 1 if facility i is built, s_{i1} determines the initial capacity of facility i, $s_{it}(\boldsymbol{\xi}^t)$ is the state variable calculating the remaining capacity of facility i at stage t, and $x_{ijt}(\boldsymbol{\xi}^t)$ is the amount of goods transported from facility i to customer j at stage t. Objective function (28a) together with constraint (28b) measures the worst-case cost. Constraints (28c) bound the initial capacity of the facilities, whereas constraints (28d) are the state equations calculating the remaining capacities. Constraints (28e) ensure that the customer demands are met.

4.3.1. Problem Instances Our instances are generated using the instance parameters described in Zeng and Zhao (2013) for the two-stage problem. For number of stages $T \in [3,5]$, we have five combinations for the number of facilities and customers, $(I,J) \in \{(5,5),(5,7),(5,10),(10,10),(20,20)\}$. Fixed installation, unit capacity, and unit transportation costs are drawn from $f_i \in [100,1000], a_i \in [10,100], c_{ij} \in [1,1000]$, respectively. Maximal capacity is set to $K_i = 2 \times 10^4$ based on preliminary experiments to make sure the instances are feasible. Customer demands are random parameters with support $\left[\mu_{jt},(1+\alpha^d)\mu_{jt}\right]$, where $\mu_{jt} \in [10,500]$ and $\alpha^d \in \{0.1,0.3,0.5\}$ are given, with α^d a parameter controlling the variation among demand realizations of customer j. For this problem, we have generated 32 instances with $\alpha^d = 0.5$ over small scenario trees to compare the bounds with the optimal objective value. To demonstrate the versatility of our solution methods, in addition to the described scenario-tree instances, we have generated 85 instances where the demands belong to the following budgeted uncertainty set:

$$\Xi = \left\{ d \in \mathbb{R}_+^{J \times T-1} \;\middle|\; d_{jt} = \mu_{jt} + \delta_{jt}\sigma_{jt},\ \delta_{jt} \in [0,1],\ j \in [J],\ t \in [2,T],\ \sum_{t \in [2,T]} \sum_{j \in [J]} \delta_{jt} \le \Gamma \right\}.$$

Note that Γ correlates the demands of all customers and stages together, which results in stagewise (temporal) dependence between the decision stages. We use $\Gamma = \alpha^u I$, with $\alpha^u \in \{0.1,0.4,0.7,1\}$.

4.3.2. Scenario-Tree Instances In solving the primal (2S-LDR) and exact models of the scenario-tree instances, we have used their respective extensive forms, while for computing the restricted NA dual bound we implemented the cutting-plane method described in Section 3.3.

Table 2 presents the gap between the bound and the optimal value of the exact solution, $100\left(\frac{\nu^{(\cdot)}-\nu^\star}{\nu^\star}\right)$ and $100\left(\frac{\nu^\star-\nu^{(\cdot)}}{\nu^\star}\right)$ for the primal and dual problems, respectively. Results show that, among the group of instances with similar characteristics, as the branching factor or the number of stages increases, the ν^{LDR} bound gets noticeably worse. In contrast, $\nu^{\text{2S-LDR}}$ stays very close to the optimal value, with an average relative distance of 0.2% among all the instances (compared to 8.7% for ν^{LDR}). We do not observe the same trend for the dual bounds, and their relative distance to the optimal value fluctuates even between two instances that only differ in the number of scenarios. Nevertheless, $\bar{\nu}_R^{\text{NA}}$ considerably outperforms PI, with an average improvement of 39.4%.

4.3.3. Budgeted-Uncertainty Instances For the location-transportation instances of this section, the $\nu^{\text{2S-LDR}}$ bounds are obtained using the C&CG method of Section 2.2 and detailed models are given in Appendix B.2.1. Similar to the scenario-tree instances, we use the cutting-plane method for obtaining the dual bound $\bar{\nu}_R^{\text{NA}}$, using a sample of size $50(T-2)$. For both algorithms, we stop when the optimality gap of the method falls below 5% or we reach the 10-hour time limit. Figure 8 shows the value of $\text{OPT}^{\mathtt{N}}$, the optimality gap between $\nu^{\text{2S-LDR}}$ and $\bar{\nu}_R^{\text{NA}}$, for our 85 instances (detailed results are provided in Appendix C.2). In this figure, every bar represents an

Table 2 Quality of the bounds for the scenario-tree instances of the location-transportation problem

T	I	J	BR	$\lvert\Xi\rvert$	ν^{LDR}	ν^{2S}	$\bar{\nu}_R^{\text{NA}}$	PI
3	5	10	3	9	2.9%	0.0%	1.4%	4.7%
			4	16	0.8%	0.0%	3.4%	3.4%
			5	25	4.6%	0.0%	1.5%	13.2%
			6	36	5.6%	0.0%	17.4%	27.5%
			7	49	8.2%	0.0%	9.4%	9.4%
			8	64	8.8%	0.0%	5.2%	14.6%
			9	81	6.5%	0.0%	6.1%	6.1%
			10	100	6.6%	0.0%	8.0%	10.2%
3	10	10	3	9	2.3%	0.0%	2.8%	13.2%
			4	16	2.7%	0.1%	1.2%	8.4%
			5	25	6.1%	0.0%	12.0%	19.3%
			6	36	9.8%	0.0%	10.2%	22.6%
			7	49	13.7%	0.1%	6.7%	14.1%
			8	64	13.9%	0.3%	5.5%	5.5%
			9	81	12.2%	0.2%	8.1%	15.5%
			10	100	12.3%	0.2%	10.1%	19.0%
3	20	20	3	9	2.5%	0.7%	9.5%	11.3%
			4	16	3.9%	0.5%	7.0%	7.0%
			5	25	5.8%	0.5%	28.4%	28.4%
			6	36	6.6%	0.5%	13.0%	19.2%
			7	49	9.1%	0.5%	11.7%	31.3%

T	I	J	BR	$\lvert\Xi\rvert$	ν^{LDR}	ν^{2S}	$\bar{\nu}_R^{\text{NA}}$	PI
4	5	10	3	27	8.0%	0.0%	2.3%	8.0%
			4	64	6.8%	0.0%	2.0%	2.0%
			5	125	13.5%	0.1%	5.0%	5.0%
4	10	10	3	27	11.9%	0.1%	17.3%	17.3%
			4	64	13.3%	0.0%	10.4%	21.0%
			5	125	13.1%	0.4%	9.4%	10.0%
4	20	20	3	27	9.0%	0.9%	8.5%	18.1%
			4	64	11.9%	0.5%	9.7%	21.7%
5	5	10	3	81	14.7%	1.0%	4.4%	6.3%
			4	256	13.9%	0.5%	11.2%	19.1%
5	10	10	3	81	16.8%	0.7%	8.2%	8.2%

instance. All instances of a block (characterized by T, (I,J) and α^d) share the same parameters except α^u, and instances within a block appear in increasing order of their α^u. Our methods return an average optimality gap of 17.4% among all the instances, with a minimum gap of 5.2% and a maximum gap of 29.2%. Given the strength of the primal bound in the scenario-tree instances, and the fact that the bound remains strong despite increasing the size of the tree, we believe that the larger gaps in most instances are caused by weak $\bar{\nu}_R^{\text{NA}}$ bounds. Since the value of $\bar{\nu}_R^{\text{NA}}$ depends on the sample on which it is solved, not having a large enough sample can lead to a poor bound. However, because of the bilinear form of the cutting-plane master problem, we were not able to

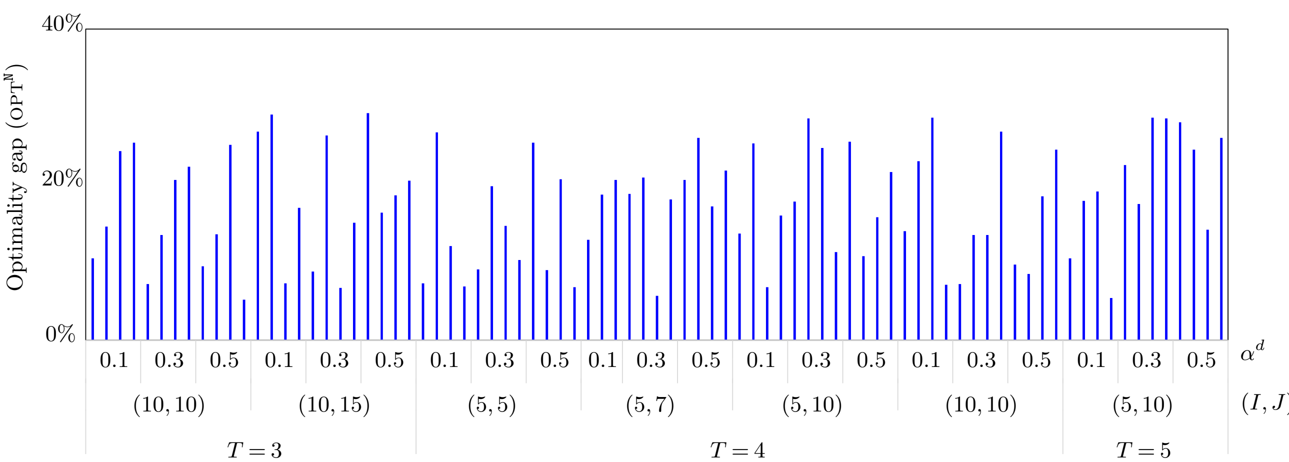

Figure 8 Optimality gaps between $\nu^{\text{2S-LDR}}$ and $\bar{\nu}_R^{\text{NA}}$ for the budgeted-uncertainty instances.

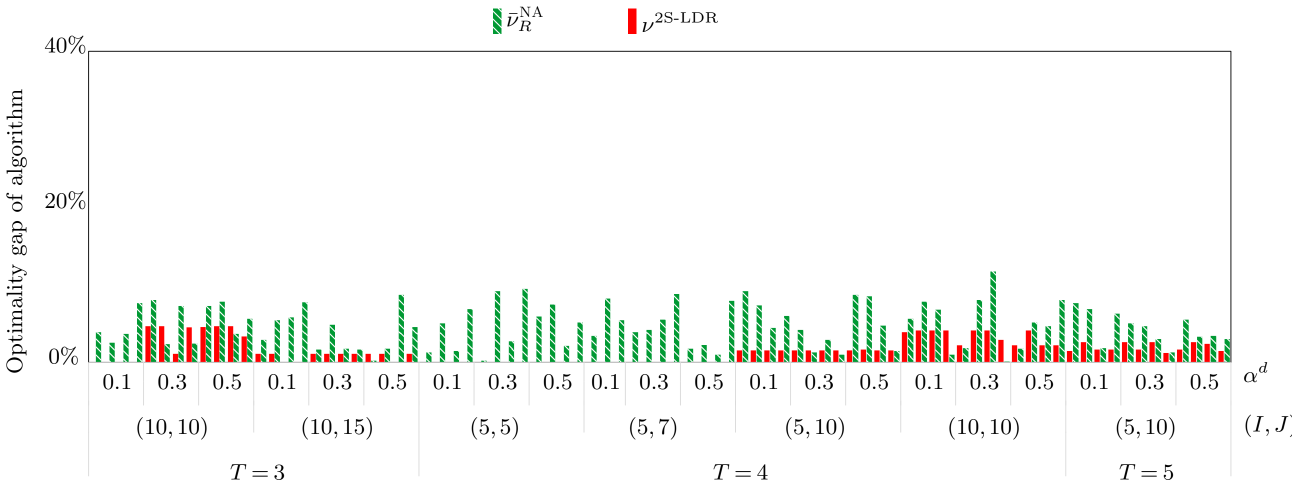

Figure 9 Optimality gaps of the solution algorithms for $\nu^{\text{2S-LDR}}$ and $\bar{\nu}_R^{\text{NA}}$ at convergence/interruption.

solve the model with large samples using off-the-shelf commercial solvers. By employing specialized algorithms developed for bilinear problems, it could be possible to increase the size of the sample and improve the $\bar{\nu}_R^{\text{NA}}$ bound, which we leave for future research.

Next we analyze the performance of the algorithms (C&CG and the cutting-plane) in solving our instances. Solution times and number of iterations are given in Appendix C.2. Here, we study their final optimality gap, that could have a negative impact on the quality of the bounds. Figure 9 shows that in all instances, the C&CG stops at a solution of the 2ARO model with less than 5% gap. In fact, the method proves to be quite powerful in detecting the significant scenarios for the 2ARO approximation, such that in 69 out of 85 instances it achieves this gap after only two iterations. On the other hand, in 39 instances the cutting-plane method is not able to reach the optimality gap of 5% within the time limit of 10 hours, meaning that we use the best lower bound on $\bar{\nu}_R^{\text{NA}}$ obtained at the end of 10 hours in calculating $\text{OPT}^{\mathbb{N}}$. This can further contribute to a weak dual bound. This suggests that, in addition to having a difficult nonlinear master problem, the cutting-plane method itself requires algorithmic enhancements, such as the design of stronger cuts.

4.4. Robust Capital Budgeting Problem

In the capital budgeting problem, a company wants to invest in a subset of I projects with uncertain cost and profit, subject to an initial budget of B that can be increased by getting a loan. A variant of the two-stage problem is studied by Subramanyam et al. (2020). In the following, we formulate the multistage capital budgeting problem as an MSARO. Over a planning horizon of T stages, let $x_{it}(\boldsymbol{\xi}^t)$ be a binary decision variable taking the value of 1 if the company decides to invest in the project $i \in [I]$ at stage $t \in [T]$, with a cost of $c_{it}(\boldsymbol{\xi}^t)$ and profit of $r_{it}(\boldsymbol{\xi}^t)$. Further, let $L_t(\boldsymbol{\xi}^t)$ be a continuous decision variable determining the amount of loan the company decides to get at stage $t \in [T]$ with a unit cost of $c_{\text{L}}\mu^{t-1}, \mu > 1$. The MSARO model is as follows:

$$
\begin{aligned}
&\max \quad z && && \text{(29a)}\\
&\text{s.t.} \quad z \le \sum_{t\in[T]}\sum_{i\in[I]} r_{it}(\boldsymbol{\xi}^t)x_{it}(\boldsymbol{\xi}^t) - \sum_{t\in[T]} c_{\text{L}}\mu^{t-1}L_t(\boldsymbol{\xi}^t) && \boldsymbol{\xi}^T \in \Xi && \text{(29b)}\\
&\quad B_t(\boldsymbol{\xi}^t) - B_{t-1}(\boldsymbol{\xi}^{t-1}) + C_{t-1}(\boldsymbol{\xi}^{t-1}) - L_t(\boldsymbol{\xi}^t) = 0 && t\in[T], \boldsymbol{\xi}^t \in \Xi^t && \text{(29c)}\\
&\quad \sum_{i\in[I]} c_{it}(\boldsymbol{\xi}^t)x_{it}(\boldsymbol{\xi}^t) = C_t(\boldsymbol{\xi}^t) && t\in[T], \boldsymbol{\xi}^t \in \Xi^t && \text{(29d)}\\
&\quad B_t(\boldsymbol{\xi}^t) - C_t(\boldsymbol{\xi}^t) \ge 0 && t\in[T], \boldsymbol{\xi}^t \in \Xi^t && \text{(29e)}\\
&\quad x_t(\boldsymbol{\xi}^t) \in \{0,1\}^I, L_t(\boldsymbol{\xi}^t) \in \mathbb{R}_+ && t\in[T], \boldsymbol{\xi}^t \in \Xi^t, && \text{(29f)}
\end{aligned}
$$

where $B_t(\boldsymbol{\xi}^t)$ is the amount of available funds at stage $t \in [T]$, determined by constraints (29c), while $C_t(\boldsymbol{\xi}^t)$ is the expenditure calculated through constraints (29d), with initial values of $B_0(\boldsymbol{\xi}^0) = B$ and $C_0(\boldsymbol{\xi}^0) = 0$. Constraints (29e) bound the expenditure amount by available funds. The objective is to maximize the worst-case profit over the planning horizon, measured by constraints (29b).

4.4.1. Problem Instances Our instance generation loosely follows the procedure of Subramanyam et al. (2020) for the two-stage problem. For $T \in [3,5]$ and $I \in \{5,8,10\}$, uncertain costs and profits $\left(c_{it}(\boldsymbol{\xi}^t), r_{it}(\boldsymbol{\xi}^t)\right)$ belong to the following uncertainty set

$$
\Xi = \left\{ (c,r) \in \mathbb{R}_+^{I\times T} \times \mathbb{R}_+^{I\times T} \;\middle|\; c_{it} = c_{it}^0 + \delta_{it}^c \sigma_{it},\ r_{it} = \frac{c_{it}^0}{5} + \delta_{it}^r \sigma_{it},\ ,\ \delta_{it}^c, \delta_{it}^r \in [0,1],\ i \in [I],\ t \in [T] \right\},
$$

where the nominal costs c^0 and perturbations σ are drawn uniformly from $[0,10]^5$ and $[0,5]^5$, respectively. Loan purchase cost is $c_L\mu^{t-1} = 0.8(1.2)^{t-1}$ per unit of loan. For each combination of T and I we consider two levels of initial budget which impacts the dependence between stages.

4.4.2. Optimality Gap Due to to presence of binary variables $x_t(\boldsymbol{\xi}^t)$, with the existing methods in the literature of 2ARO we cannot calculate $\nu^{\text{2S-LDR}}$ exactly. Therefore, we solve an approximation of it using the K-adaptability method of Subramanyam et al. (2020) with $K = 2$ and use the obtained bound ν^{K} in measuring the optimality gap. On the dual side, we study three options: (i) ν^{Ω}, the upper bound from solving the model (29) with a sample $|\Omega| = 100$, (ii) the NA bound $\bar{\nu}_R^{\text{NA}}$ with the same sample Ω, (iii) the decomposable NA bound $\bar{\nu}_R^{\text{DNA}}$ with a sample of size 300. For each option, when the model is not solved to optimality within the time limit, the best valid bound is used in the calculations. Note that, if we can solve both options (i) and (ii) to optimality, we should expect a better bound from option (i). Table 3 presents the optimality gaps of the capital budgeting instances between the bound ν^{K} and the three choices of upper bound. Results show that, even though the bound ν^{Ω} should theoretically be at least as good as the bound $\bar{\nu}_R^{\text{NA}}$, on

Table 3 Optimality gaps for the capital budgeting instances

T	I	B	$\left(\frac{\nu^{\Omega}-\nu^{\mathrm{K}}}{\nu^{\mathrm{K}}}\right)\%$	$\left(\frac{\bar{\nu}_R^{\mathrm{NA}}-\nu^{\mathrm{K}}}{\nu^{\mathrm{K}}}\right)\%$	$\left(\frac{\bar{\nu}_R^{\mathrm{DNA}}-\nu^{\mathrm{K}}}{\nu^{\mathrm{K}}}\right)\%$
3	5	50	17%	20%	16%
		100	27%	25%	18%
	8	100	24%	25%	24%
		150	34%	21%	21%
	10	100	33%	28%	16%
		150	38%	29%	19%
4	5	100	35%	26%	20%
		150	46%	33%	15%
	8	150	43%	30%	34%
		200	57%	37%	30%
	10	150	48%	36%	31%
		200	44%	33%	20%
5	5	200	65%	42%	20%
		250	42%	39%	37%
	8	250	40%	39%	31%
		300	43%	40%	21%
	10	250	55%	31%	31%
		300	65%	31%	25%

average the NA bound returns a better upper bound within the same time limit. This is testament to the difficulty of the multistage problem even when it is solved for a discrete set of scenarios. The decomposable NA bound $\bar{\nu}_R^{\mathrm{DNA}}$ further improves the results by using a larger sample which is viable because of its superior computational performance. Note that, the best gaps in Table 3, ranging between 15% to 34%, are obtained from approximations over approximations on both primal and dual side. Accordingly, these rather large gaps can be attributed to both bounding methods. ν^{K} can be further improved by using larger values of K, whereas the NA bounds might benefit from implicit enumeration techniques in solving the relaxation models without resorting to sampling.

5. Conclusion

Robust optimization models are built on a different premise than stochastic programming in the sense that they do not assume any knowledge about the probability distribution, focusing instead on optimizing the worst-case outcomes. We adapt two decision rules from the stochastic programming literature (namely two-stage LDRs and LDDRs) to MSARO, which result in two-stage approximation models for which we design several appropriate solution methods. More specifically, for two-stage linear decision rules we present sufficient conditions that make the approximation models amenable to the constraint-and-column generation method. In order to solve LDDR-induced models, we build a cutting-plane method for mixed-integer MSAROs, as well as a monolithic bilinear program for continuous problems. Our numerical experiments on the newsvendor, location-transportation and capital budgeting problems reveal that our methods considerably reduce the optimality gap from commonly used approaches in the literature.

Acknowledgments

This work was supported by Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2018-04984] and Agence Nationale de la Recherche of France [Grant ANR-22-CE048-0018]. Computations were performed on the Niagara supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto.

References

ApS M (2022) *MOSEK modeling cookbook*.

Ardestani-Jaafari A, Delage E (2021) Linearized robust counterparts of two-stage robust optimization problems with applications in operations management. *INFORMS Journal on Computing* 33(3):1138–1161.

Arslan AN, Detienne B (2022) Decomposition-based approaches for a class of two-stage robust binary optimization problems. *INFORMS Journal on Computing* 34(2):857–871.

Bampou D, Kuhn D (2011) Scenario-free stochastic programming with polynomial decision rules. *50th IEEE Conference on Decision and Control and European Control Conference*, 7806–7812.

Ben-Tal A, El Ghaoui L, Nemirovski A (2009) Robust optimization. *Robust optimization* (Princeton university press).

Ben-Tal A, El Housni O, Goyal V (2020) A tractable approach for designing piecewise affine policies in two-stage adjustable robust optimization. *Mathematical Programming* 182(1):57–102.

Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Mathematical Programming* 99(2):351–376.

Ben-Tal A, Nemirovski A (1999) Robust solutions of uncertain linear programs. *Operations Research Letters* 25(1):1–13.

Bertsimas D, Brown DB, Caramanis C (2011a) Theory and applications of robust optimization. *SIAM review* 53(3):464–501.

Bertsimas D, Caramanis C (2010) Finite adaptability in multistage linear optimization. *IEEE Transactions on Automatic Control* 55(12):2751–2766.

Bertsimas D, de Ruiter FJ (2016) Duality in two-stage adaptive linear optimization: Faster computation and stronger bounds. *INFORMS Journal on Computing* 28(3):500–511.

Bertsimas D, Dunning I (2016) Multistage robust mixed-integer optimization with adaptive partitions. *Operations Research* 64(4):980–998.

Bertsimas D, Georghiou A (2015) Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Operations Research* 63(3):610–627.

Bertsimas D, Georghiou A (2018) Binary decision rules for multistage adaptive mixed-integer optimization. *Mathematical Programming* 167(2):395–433.

Bertsimas D, Goyal V (2012) On the power and limitations of affine policies in two-stage adaptive optimization. *Mathematical Programming* 134(2):491–531.

Bertsimas D, Goyal V, Lu BY (2015) A tight characterization of the performance of static solutions in two-stage adjustable robust linear optimization. *Mathematical Programming* 150(2):281–319.

Bertsimas D, Iancu DA, Parrilo PA (2010) Optimality of affine policies in multistage robust optimization. *Mathematics of Operations Research* 35(2):363–394.

Bertsimas D, Iancu DA, Parrilo PA (2011b) A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Transactions on Automatic Control* 56(12):2809–2824.

Bertsimas D, Litvinov E, Sun XA, Zhao J, Zheng T (2012) Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Transactions on Power Systems* 28(1):52–63.

Bertsimas D, Sim M, Zhang M (2019) Adaptive distributionally robust optimization. *Management Science* 65(2):604–618.

Bertsimas D, Thiele A (2006) A robust optimization approach to inventory theory. *Operations Research* 54(1):150–168.

Bodur M, Luedtke JR (2018) Two-stage linear decision rules for multi-stage stochastic programming. *Mathematical Programming* 1–34.

Chen X, Sim M, Sun P, Zhang J (2008) A linear decision-based approximation approach to stochastic programming. *Operations Research* 56(2):344–357.

Chen X, Zhang Y (2009) Uncertain linear programs: Extended affinely adjustable robust counterparts. *Operations Research* 57(6):1469–1482.

Daryalal M, Bodur M, Luedtke JR (2021) Lagrangian dual decision rules for multistage stochastic mixed integer programming. *Operations Research* To appear.

Feige U, Jain K, Mahdian M, Mirrokni V (2007) Robust combinatorial optimization with exponential scenarios. *International Conference on Integer Programming and Combinatorial Optimization*, 439–453 (Springer).

Geoffrion AM (1974) *Lagrangean relaxation for integer programming*, volume 2, 82–114 (Berlin, Heidelberg: Springer Berlin Heidelberg).

Georghiou A, Tsoukalas A, Wiesemann W (2019) Robust dual dynamic programming. *Operations Research* 67(3):813–830.

Georghiou A, Tsoukalas A, Wiesemann W (2020) A primal–dual lifting scheme for two-stage robust optimization. *Operations Research* 68(2):572–590.

Goh J, Sim M (2010) Distributionally robust optimization and its tractable approximations. *Operations Research* 58(4-part-1):902–917.

Gurobi Optimization, LLC (2022) Gurobi Optimizer Reference Manual. URL https://www.gurobi.com.

Guslitser E (2002) *Uncertainty-immunized solutions in linear programming*. Ph.D. thesis, Citeseer.

Hadjiyiannis MJ, Goulart PJ, Kuhn D (2011) A scenario approach for estimating the suboptimality of linear decision rules in two-stage robust optimization. *50th IEEE Conference on Decision and Control and European Control Conference*, 7386–7391 (IEEE).

Hanasusanto GA, Kuhn D, Wiesemann W (2015) K-adaptability in two-stage robust binary programming. *Operations Research* 63(4):877–891.

Hashemi Doulabi H, Jaillet P, Pesant G, Rousseau LM (2021) Exploiting the structure of two-stage robust optimization models with exponential scenarios. *INFORMS Journal on Computing* 33(1):143–162.

Hendrix EM, Boglárka G, et al. (2010) *Introduction to nonlinear and global optimization*, volume 37 (Springer).

Iancu DA, Sharma M, Sviridenko M (2013) Supermodularity and affine policies in dynamic robust optimization. *Operations Research* 61(4):941–956.

Kuhn D, Wiesemann W, Georghiou A (2011) Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming* 130(1):177–209.

Loken C, Gruner D, Groer L, Peltier R, Bunn N, Craig M, Henriques T, Dempsey J, Yu CH, Chen J, et al. (2010) SciNet: lessons learned from building a power-efficient top-20 system and data centre. *Journal of Physics-Conference Series*, volume 256, 012026.

Marandi A, Den Hertog D (2018) When are static and adjustable robust optimization problems with constraint-wise uncertainty equivalent? *Mathematical Programming* 170(2):555–568.

Pereira MV, Pinto LM (1991) Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* 52(1-3):359–375.

Philpott AB, de Matos VL, Kapelevich L (2018) Distributionally robust SDDP. *Computational Management Science* 15(3):431–454.

Ponce M, van Zon R, Northrup S, Gruner D, Chen J, Ertinaz F, Fedoseev A, Groer L, Mao F, Mundim BC, et al. (2019) Deploying a top-100 supercomputer for large parallel workloads: The Niagara supercomputer. *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)*, 1–8 (Association for Computing Machinery, New York, NY, United States).

Postek K, Hertog Dd (2016) Multistage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. *INFORMS Journal on Computing* 28(3):553–574.

Romeijnders W, Postek K (2020) Piecewise constant decision rules via branch-and-bound based scenario detection for integer adjustable robust optimization. *INFORMS Journal on Computing* .

See CT, Sim M (2010) Robust approximation to multiperiod inventory management. *Operations Research* 58(3):583–594.

Shapiro A, Dentcheva D, Ruszczyński A (2009) *Lectures on Stochastic Programming: Modeling and Theory* (SIAM).

Subramanyam A, Gounaris CE, Wiesemann W (2020) K-adaptability in two-stage mixed-integer robust optimization. *Mathematical Programming Computation* 12(2):193–224.

Wolsey LA, Nemhauser GL (1999) *Integer and combinatorial optimization*, volume 55 (John Wiley & Sons).

Xu G, Hanasusanto GA (2021) Improved decision rule approximations for multi-stage robust optimization via copositive programming. *arXiv preprint arXiv:1808.06231* .

Yanıkoğlu İ, Gorissen BL, den Hertog D (2019) A survey of adjustable robust optimization. *European Journal of Operational Research* 277(3):799–813.

Yu X, Shen S (2020) Multistage distributionally robust mixed-integer programming with decision-dependent moment-based ambiguity sets. *Mathematical Programming* 1–40.

Zeng B, Zhao L (2013) Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters* 41(5):457–461.

Zhen J, Den Hertog D, Sim M (2018) Adjustable robust optimization via fourier–motzkin elimination. *Operations Research* 66(4):1086–1100.

Appendices

A. Proofs

In this section, we present the proofs of the three propositions and one lemma mentioned in the body of the paper, for which we also restate the claims for convenience.

PROPOSITION 1. *Assume that the uncertainty set Ξ is either a finite discrete set or a compact polyhedron, and the basis functions $\Phi_t(\boldsymbol{\xi}^t)$ are affine in $\boldsymbol{\xi}^t$. Then the C&CG algorithm converges to $\nu^{2S\text{-}LDR}$ of an MSARO with continuous and fixed recourse in a finite number of iterations.*

Proof Denote by $f(x_1, \boldsymbol{\xi}^T, \beta)$, the objective function of the inner maximization in the two-stage problem (4). Then, we can rewrite problem (4) as:

$$\nu^{\text{2S-LDR}} = \min\left\{c_1^\top x_1 + \max_{\boldsymbol{\xi}^T\in\Xi} f(x_1, \boldsymbol{\xi}^T, \beta) \mid x_1 \in X_1,\ \beta_t \in \mathbb{R}^{K_t},\ t\in[2,T]\right\}.$$

First, we show that $f(x_1, \boldsymbol{\xi}^T, \beta)$ is convex in $\boldsymbol{\xi}^T$ for given $x_1 \in X_1$. For $\boldsymbol{\xi}_1^T, \boldsymbol{\xi}_2^T \in \Xi$ and $\lambda \in [0,1]$, the following (in)equalities hold:

$$\begin{aligned}
&\lambda f(x_1, \boldsymbol{\xi}_1^T, \beta) + (1-\lambda) f(x_1, \boldsymbol{\xi}_2^T, \beta) && = \\
&\sum_{t\in[2,T]} c_t^{\mathsf{s}\top}\big(\lambda\Phi(\boldsymbol{\xi}_1^t)^\top + (1-\lambda)\Phi(\boldsymbol{\xi}_2^t)^\top\big)\beta_t + \\
&\qquad \lambda\Big(\min_{x^{\mathsf{r}}\in\mathcal{X}(x_1,\beta,\boldsymbol{\xi}_1^T)} \sum_{t\in[2,T]} c_t^{\mathsf{r}\top} x_t^{\mathsf{r}}\Big) + (1-\lambda)\Big(\min_{x^{\mathsf{r}}\in\mathcal{X}(x_1,\beta,\boldsymbol{\xi}_2^T)} \sum_{t\in[2,T]} c_t^{\mathsf{r}\top} x_t^{\mathsf{r}}\Big) && = && (30a)\\
&\sum_{t\in[2,T]} c_t^{\mathsf{s}\top}\Phi(\lambda\boldsymbol{\xi}_1^t + (1-\lambda)\boldsymbol{\xi}_2^t)^\top\beta_t + \\
&\qquad \lambda\Big(\min_{x^{\mathsf{r}}\in\mathcal{X}(x_1,\beta,\boldsymbol{\xi}_1^T)} \sum_{t\in[2,T]} c_t^{\mathsf{r}\top} x_t^{\mathsf{r}}\Big) + (1-\lambda)\Big(\min_{x^{\mathsf{r}}\in\mathcal{X}(x_1,\beta,\boldsymbol{\xi}_2^T)} \sum_{t\in[2,T]} c_t^{\mathsf{r}\top} x_t^{\mathsf{r}}\Big) && \geq && (30b)\\
&\sum_{t\in[2,T]} c_t^{\mathsf{s}\top}\Phi(\lambda\boldsymbol{\xi}_1^t + (1-\lambda)\boldsymbol{\xi}_2^t)^\top\beta_t + \min_{x^{\mathsf{r}}\in\mathcal{X}(x_1,\beta,\lambda\boldsymbol{\xi}_1^T(1-\lambda)\boldsymbol{\xi}_2^T)} \sum_{t\in[2,T]} c_t^{\mathsf{r}\top} x_t^{\mathsf{r}} && = f(x_1, \lambda\boldsymbol{\xi}_1^T + (1-\lambda)\boldsymbol{\xi}_2^T, \beta).
\end{aligned}$$

Equality (30a) holds because $\Phi(\boldsymbol{\xi}^t)$ is affine. Inequality (30b) follows from the convexity of the linear minimization problem (by the fixed recourse assumption), since $\boldsymbol{\xi}$ appears only on the right-hand-sides of constraints and x_t^r are continuous (by the continuous recourse assumption).

Now, let Ξ^{EXT} be the set of extreme points of Ξ if it is a polyhedron, or that of its convex-hull conv(Ξ) if it is a finite discrete set. In maximization of a convex function over a compact polyhedral set, there is an optimal solution that is an extreme point (Hendrix et al. 2010). Then, the two-stage problem becomes:

$$\nu^{\text{2S-LDR}} = \min\left\{c_1^\top x_1 + \eta \mid \eta \geq f(x_1, \boldsymbol{\xi}^T, \beta),\ \boldsymbol{\xi}^T \in \Xi^{\text{EXT}},\ x_1 \in X_1,\ \beta_t \in \mathbb{R}^{K_t},\ t\in[2,T]\right\}.$$

C&CG is then the process of gradually adding constraints for each extreme point. Because $|\Xi^{\text{EXT}}| < +\infty$, iterating over all extreme points takes finitely many steps, which concludes the proof. □

Lemma 1. *For given $\mathbb{P} \in \mathcal{P}$, constraints (8e) are equivalent to the following:*

$$y_t(\boldsymbol{\xi}^T) = \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}}\left[y_t(\boldsymbol{\xi}'^T) \;\middle|\; \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t \right], \quad t \in [T],\ \boldsymbol{\xi}^T \in \Xi. \tag{9}$$

Proof For a fixed t and $\boldsymbol{\xi}^t$, and every $\boldsymbol{\xi}'^T$ that shares the same history with $\boldsymbol{\xi}^T$ up to t (including $\boldsymbol{\xi}^T$), multiply both sides of (8e) by $p(\boldsymbol{\xi}'^T)$. The corresponding constraint (9) is derived by integrating over all such $\boldsymbol{\xi}'^T$ scenarios, thus it is a valid equality for problem (8). Moreover, this constraint makes sure that the decision variables under all the scenarios sharing the same history take the same value, i.e., $\mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}}\left[y_t(\boldsymbol{\xi}'^T) \;\middle|\; \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t \right]$, which enforces the semantics of the nonanticipativity constraints (8e). □

Proposition 2. *Let $\mathbb{P}$ be any probability distribution in $\mathcal{P}$. For continuous MSARO problems, $\mathcal{L}^{\mathrm{NA}}(\mathbb{P})$ is a strong dual.*

Proof The strength of a Lagrangian dual problem can be studied by its primal characterization, derived by Geoffrion (1974) for a mixed-integer linear optimization problem. Then, using standard Lagrangian duality theory (see, for instance, Wolsey and Nemhauser (1999)), the primal characterization of (11) is:

$$\min \quad z \tag{31a}$$

$$\text{s.t.} \quad \left(z, y_1(\boldsymbol{\xi}^T), \ldots, y_T(\boldsymbol{\xi}^T)\right) \in \operatorname{conv}(Y(\boldsymbol{\xi}^T)) \qquad \boldsymbol{\xi}^T \in \Xi \tag{31b}$$

$$y_t(\boldsymbol{\xi}^T) = \mathbb{E}_{\boldsymbol{\xi}'^T \sim \mathbb{P}}\left[y_t(\boldsymbol{\xi}'^T) \;\middle|\; \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t \right] \qquad t \in [T], \boldsymbol{\xi}^T \in \Xi. \tag{31c}$$

The result follows from the equivalence of (9) and (8e), as well as the fact that for continuous MSARO problems, we have that $\operatorname{conv}(Y(\boldsymbol{\xi}^T)) = Y(\boldsymbol{\xi}^T)$ for all $\boldsymbol{\xi}^T \in \Xi$. □

Proposition 3. *$\bar{\nu}_R^{\mathrm{NA}}$ is a valid lower bound for $\nu^\star$.*

Proof We have shown that $\mathcal{L}_R^{\mathrm{NA}}(\mathbb{P}) \le \nu^\star$ for any $\mathbb{P} \in \mathcal{P}$. Now we claim that, for any $\hat{\mathbb{P}} \in \bar{\mathcal{P}} \setminus \mathcal{P} = \left\{ \mathsf{P} \in \bar{\mathcal{P}} \mid \exists \boldsymbol{\xi}^T \in \Xi : \ p^{\mathsf{P}}(\boldsymbol{\xi}^T) = 0 \right\}$ where $p^{\hat{\mathbb{P}}} : \Xi \to \mathbb{R}_+$ is the density function of $\hat{\mathbb{P}}$, the inequality $\mathcal{L}_R^{\mathrm{NA}}(\hat{\mathbb{P}}) \le \nu^\star$ also holds. Consider the nonanticipativity constraints (8e) in the NA reformulation of the MSARO problem (1). For a fixed $\hat{\mathbb{P}}$ and $\boldsymbol{\xi}^T$ in the uncertainty set, and all $\boldsymbol{\xi}'^T \in \Xi$ sharing the same history with $\boldsymbol{\xi}^T$ up to t ($\boldsymbol{\xi}^t = \boldsymbol{\xi}'^t$), multiply both sides of their associated constraint (8e) by $p^{\hat{\mathbb{P}}}(\boldsymbol{\xi}'^T)$ and integrate over all such $\boldsymbol{\xi}'^T$. Note that, any $p^{\hat{\mathbb{P}}}(\boldsymbol{\xi}'^T) = 0$ removes a nonanticipativity constraint from the problem. Thus, constraints (9) obtained with respect to $\hat{\mathbb{P}}$ are not equivalent to constraints (8e) and are simply an aggregation hence a valid equality. Therefore, replacing the nonanticipativity constraints with (9) obtained with respect to $\hat{\mathbb{P}}$, results in a relaxation of (1) and not an exact reformulation. Subsequently, relaxation of (9) and construction of the Lagrangian dual problem (11) with respect to $\hat{\mathbb{P}}$ yields a relaxation of a relaxation of a minimization problem, where for each fixed $\lambda_t(\cdot)$ we have $\mathcal{L}^{\text{NA}}_{\text{LR}}(\hat{\mathbb{P}}, \lambda_1(\cdot) \ldots, \lambda_T(\cdot)) \leq \nu^\star$, and consequently $\mathcal{L}^{\text{NA}}(\hat{\mathbb{P}}) \leq \nu^\star$. From the primal characterizations (31) and (13), the inequality $\mathcal{L}^{\text{NA}}_{R}(\hat{\mathbb{P}}) \leq \mathcal{L}^{\text{NA}}(\hat{\mathbb{P}})$ follows, which completes the proof. □

B. Detailed Models

B.1. Monolithic Form of Model (6)

Let π^{S} and π^{R} be the LP dual variables associated with the state and recourse constraints in $\mathcal{X}(x_1, \hat{\beta}_t, \boldsymbol{\xi}^T)$, in the inner minimization problem of (6), respectively. Then, subproblem (6) with continuous recourse can be modeled as follows:

$$
\begin{aligned}
\max \quad & \sum_{t\in[2,T]} c_t^{\mathtt{s}\top} \Phi_t(\boldsymbol{\xi}^t)^\top \hat{\beta}_t + c_t^{\mathtt{r}\top} x_t^{\mathtt{r}} && && \text{(32a)}\\
\text{s.t.} \quad & A_t^{\mathtt{r}} x_t^{\mathtt{r}} + A_t^{\mathtt{s}} \Phi_t(\boldsymbol{\xi}^t)^\top \hat{\beta}_t + B_t^{\mathtt{s}} \Phi_{t-1}(\boldsymbol{\xi}^{t-1})^\top \hat{\beta}_{t-1} - b_t(\boldsymbol{\xi}^t) \leq 0 && t \in [2,T] && \text{(32b)}\\
& D_t^{\mathtt{s}} \Phi_t(\boldsymbol{\xi}^t)^\top \hat{\beta}_t + D_t^{\mathtt{r}} x_t^{\mathtt{r}} - d_t(\boldsymbol{\xi}^t) \leq 0 && t \in [2,T] && \text{(32c)}\\
& A_t^{\mathtt{r}\top} \pi_t^{\mathtt{S}} + D_t^{\mathtt{r}\top} \pi_t^{\mathtt{R}} = c_t^{\mathtt{r}} && t \in [2,T] && \text{(32d)}\\
& \Big(A_t^{\mathtt{r}} x_t^{\mathtt{r}} + A_t^{\mathtt{s}} \Phi_t(\boldsymbol{\xi}^t)^\top \hat{\beta}_t + B_t^{\mathtt{s}} \Phi_{t-1}(\boldsymbol{\xi}^{t-1})^\top \hat{\beta}_{t-1} - b_t(\boldsymbol{\xi}^t) \Big)^\top \pi_t^{\mathtt{S}} = 0 && t \in [2,T] && \text{(32e)}\\
& \Big(D_t^{\mathtt{s}} \Phi_t(\boldsymbol{\xi}^t)^\top \hat{\beta}_t + D_t^{\mathtt{r}} x_t^{\mathtt{r}} - d_t(\boldsymbol{\xi}^t) \Big)^\top \pi_t^{\mathtt{R}} = 0 && t \in [2,T] && \text{(32f)}\\
& \boldsymbol{\xi}^T \in \Xi, \quad x_t^{\mathtt{r}} \in \mathbb{R}^{p_t}, \quad \pi_t^{\mathtt{S}} \in \mathbb{R}_-^{m_t^{\mathtt{S}}}, \quad \pi_t^{\mathtt{R}} \in \mathbb{R}_-^{m_t^{\mathtt{R}}} && t \in [2,T]. && \text{(32g)}
\end{aligned}
$$

Inequalities (32b)-(32c) are the primal feasibility constraints at $\hat{\beta}_t$, (32d) the dual feasibility constraints, and (32e)-(32f) are the complementary slackness constraints. The latter include bilinear terms, as basis functions $\Phi_t(\boldsymbol{\xi}^t)$ are functions of $\boldsymbol{\xi}^T$ which are decision variables in (32). They can be linearized with the addition of binary decision variables via the so-called big-M constraints. A detailed example of building a mixed-integer linear model for a multistage location-allocation problem is provided in Appendix B.2.

B.2. Models for the Location-Transportation

B.2.1. Column-and-constraint Generation with Two-stage Linear Decision Rules Applying LDRs on the state variables of model (28), $s_{it}(\boldsymbol{\xi}^t) = \beta_{it}^0 + \sum_{t'\in[2,t]} \sum_{j\in[J]} d_{jt'}(\boldsymbol{\xi}^t)\beta_{it}^{jt'}, i \in [I], t \in [2,T]$, results in the following 2ARO problem in monolithic form:

$$
\begin{aligned}
\min \quad & z \\
\text{s.t.} \quad & z \geq \sum_{i\in[I]} f_i y_i + \sum_{i\in[I]} a_i s_{i1} + \sum_{t\in[2,T]} \sum_{i\in[I]} a_i \big(\beta_{it}^0 + \sum_{t'\in[2,t]} \sum_{j\in[J]} d_{jt'}(\boldsymbol{\xi}'^t)\beta_{it}^{jt'}\big) + \\
& \sum_{t\in[2,T]} \sum_{i\in[I]} \sum_{j\in[J]} c_{ij} x_{ijt}(\boldsymbol{\xi}^t) \qquad\qquad \boldsymbol{\xi}^T \in \Xi
\end{aligned}
$$

$$
\begin{aligned}
& s_{i1} \le K_i y_i && i \in [I] \\
& \beta^0_{i2} + \sum_{j\in[J]} d_{j2}(\boldsymbol{\xi}^2)\beta^{j2}_{i2} = s_{i1} - \sum_{j\in[J]} x_{ij2}(\boldsymbol{\xi}^2) && i \in [I], \boldsymbol{\xi}^2 \in \Xi^2 \\
& \beta^0_{it} + \sum_{t'\in[2,t]}\sum_{j\in[J]} d_{jt'}(\boldsymbol{\xi}^t)\beta^{jt'}_{it} = \beta^0_{i,t-1} + \sum_{t'\in[2,t-1]}\sum_{j\in[J]} d_{jt'}(\boldsymbol{\xi}^{t-1})\beta^{jt'}_{i,t-1} - \sum_{j\in[J]} x_{ijt}(\boldsymbol{\xi}^t) && \\
& && i \in [I], t \in [3,T], \boldsymbol{\xi}^t \in \Xi^t \\
& \sum_{i\in[I]} x_{ijt}(\boldsymbol{\xi}^t) \ge d_{jt}(\boldsymbol{\xi}^t) && j \in [J], t \in [2,T], \boldsymbol{\xi}^t \in \Xi^t \\
& \beta^0_{it} + \sum_{t'\in[2,t]}\sum_{j\in[J]} d_{jt'}(\boldsymbol{\xi}^t)\beta^{jt'}_{it} \ge 0 && i \in [I], t \in [2,T], \boldsymbol{\xi}^t \in \Xi^t \\
& y \in \{0,1\}^I,\ z, x \ge 0.
\end{aligned}
$$

For better clarity, let us re-write the above formulation in the more common min-max-min form of 2ARO:

$$
\begin{aligned}
\min\ & \sum_{i\in[I]} f_i y_i + \sum_{i\in[I]} a_i\big(s_{i1} + \sum_{t\in[2,T]} \beta^0_{it}\big) + \max_{\boldsymbol{\xi}^T\in\Xi} && \min \sum_{t\in[2,T]}\sum_{i\in[I]}\sum_{j\in[J]} \Big(c_{ij}x_{ijt} + \sum_{t'\in[2,t]} a_i d_{jt'}(\boldsymbol{\xi}^t)\beta^{jt'}_{it}\Big) \\
\text{s.t. } & s_{i1} \le K_i y_i \quad i \in [I] && \text{s.t. } \sum_{j\in[J]} x_{ij2} = s_{i1} - \beta^0_{i2} - \sum_{j\in[J]} d_{j2}(\boldsymbol{\xi}^2)\beta^{j2}_{i2} \qquad i \in [I] \\
& y \in \{0,1\}^I,\ z \ge 0. && \sum_{j\in[J]} x_{ijt} = \beta^0_{i,t-1} - \beta^0_{it} - \sum_{j\in[J]}\sum_{t'\in[2,t]} d_{jt'}(\boldsymbol{\xi}^t)\beta^{jt'}_{it} \\
& && \quad + \sum_{j\in[J]}\sum_{t'\in[2,t-1]} d_{jt'}(\boldsymbol{\xi}^{t-1})\beta^{jt'}_{i,t-1} \\
& && \qquad i \in [I], t \in [3,T] \\
& && \sum_{i\in[I]} x_{ijt} \ge d_{jt}(\boldsymbol{\xi}^t) \qquad j \in [J], t \in [2,T] \\
& && \beta^0_{it} + \sum_{t'\in[2,t]}\sum_{j\in[J]} d_{jt'}(\boldsymbol{\xi}^t)\beta^{jt'}_{it} \ge 0 \qquad j \in [J], t \in [2,T] \\
& && x \ge 0
\end{aligned}
$$

Denote by π^1, π^2, π^3 the dual variables associated with three constraint sets of the inner minimization problem. After taking its LP dual and the adding the KKT conditions, we get the following as the subproblem of the column-and-constraint generation:

$$
\begin{aligned}
\max\quad & \sum_{t\in[2,T]}\sum_{i\in[I]}\sum_{j\in[J]}\sum_{t'\in[2,t]} a_i d_{jt'}\beta^{jt'}_{it} + \sum_{t\in[2,T]}\sum_{i\in[I]}\sum_{j\in[J]} c_{ij}x_{ijt} \\
\text{s.t.}\quad & \sum_{j\in[J]} x_{ij2} = s_{i1} - \beta^0_{i2} - \sum_{j\in[J]} d_{j2}\beta^{j2}_{i2} && i \in [I] \\
& \sum_{j\in[J]} x_{ijt} = \beta^0_{i,t-1} - \beta^0_{it} - \sum_{j\in[J]}\sum_{t'\in[2,t]} d_{jt'}\beta^{jt'}_{it} + \sum_{j\in[J]}\sum_{t'\in[2,t-1]} d_{jt'}\beta^{jt'}_{i,t-1} && i \in [I], t \in [3,T] \\
& \sum_{i\in[I]} x_{ijt} \ge d_{jt} && j \in [J], t \in [2,T] \\
& \pi^1_i + \pi^3_{j2} \le c_{ij} && i \in [I], j \in [J]
\end{aligned}
$$

$$\pi^2_{it} + \pi^3_{jt} \le c_{ij} \qquad i \in [I], j \in [J], t \in [3,T]$$

$$\Big(d_{jt} - \sum_{i\in[I]} x_{ijt}\Big)\pi^3_{jt} = 0 \xrightarrow{\text{linearization}}$$

$$\pi^3_{jt} \le M(1-\ell^R_{jt}), \quad \sum_{i\in[I]} x_{ijt} - d_{jt} \le M\ell^R_{jt} \qquad j \in [J], t \in [2,T]$$

$$\big(\pi^1_i + \pi^3_{j2} - c_{ij}\big)x_{ij2} = 0 \xrightarrow{\text{linearization}}$$

$$x_{ij2} \le M(1-\ell^D_{ij2}), \quad c_{ij} - \pi^1_i - \pi^3_{j2} \le M\ell^D_{ij2} \qquad i \in [I], j \in [J]$$

$$\big(\pi^2_{it} + \pi^3_{jt} - c_{ij}\big)x_{ijt} = 0 \xrightarrow{\text{linearization}}$$

$$x_{ijt} \le M(1-\ell^D_{ijt}), \quad c_{ij} - \pi^2_{it} - \pi^3_{jt} \le M\ell^D_{ijt} \qquad i \in [I], j \in [J], t \in [3,T]$$

$$\beta^0_{it} + \sum_{t'\in[2,t]}\sum_{j\in[J]} d_{jt'}\beta^{jt'}_{it} \ge 0 \qquad j \in [J], t \in [2,T]$$

$$d_{jt} = \mu_{jt} + \delta_{jt}\sigma_{jt} \qquad j \in [J], t \in [2,T]$$

$$\sum_{t\in[2,T]}\sum_{j\in[J]} \delta_{jt} \le \Gamma$$

$$x_{ijt} \ge 0, \ \ell^D_{ijt} \in \{0,1\} \qquad i \in [I], j \in [J], t \in [2,T]$$

$$0 \le \delta_{jt} \le 1, \ 0 \le \pi^3_{jt}, \ \ell^R_{jt} \in \{0,1\} \qquad j \in [J], t \in [2,T].$$

Then, the column-and-constraint generation master problem at κ^{th} iteration is:

$$\min \sum_{i\in[I]} f_i y_i + \sum_{i\in[I]} a_i\big(s_{i1} + \sum_{t\in[2,T]} \beta^0_{it}\big) + \eta$$

$$\text{s.t. } s_{i1} \le K_i y_i \qquad i \in [I]$$

$$\eta \ge \sum_{t\in[2,T]}\sum_{i\in[I]}\sum_{j\in[J]} \Big(c_{ij}x^k_{ijt} + \sum_{t'\in[2,t]} a_i d^k_{jt'}\beta^{jt'}_{it}\Big) \qquad k \in [\kappa]$$

$$\sum_{j\in[J]} x^k_{ij2} = s_{i1} - \beta^0_{i2} - \sum_{j\in[J]} d^k_{j2}\beta^{j2}_{i2} \qquad i \in [I], k \in [\kappa]$$

$$\sum_{j\in[J]} x^k_{ijt} = \beta^0_{i,t-1} - \beta^0_{it} - \sum_{j\in[J]}\sum_{t'\in[2,t]} d^k_{jt'}\beta^{jt'}_{it} + \sum_{j\in[J]}\sum_{t'\in[2,t-1]} d^k_{jt'}\beta^{jt'}_{i,t-1}$$
$$i \in [I], t \in [3,T], k \in [\kappa]$$

$$\sum_{i\in[I]} x^k_{ijt} \ge d^k_{jt} \qquad j \in [J], t \in [2,T], k \in [\kappa]$$

$$\beta^0_{it} + \sum_{j\in[J]}\sum_{t'\in[2,t]} d^k_{jt'}\beta^{jt'}_{it} \ge 0 \qquad i \in [I], t \in [2,T], k \in [\kappa]$$

$$x^k \ge 0 \qquad k \in [\kappa]$$

$$y \in \{0,1\}^I, \ z \ge 0,$$

where x^k_{ijt} and d^k_{jt} are the column and significant scenario generated at iteration k.

B.2.2. Restricted NA Dual The $\mathcal{Q}(\beta^s, \beta^x)$ function for the location-transportation problem:

$$
\begin{aligned}
&\mathcal{Q}(\beta^s, \beta^x) = \\
&\qquad \min \quad z + \sum_{t\in[2,T]} \sum_{i\in[I]} \beta^s_{it}(\boldsymbol{\xi}^T) s_{it}(\boldsymbol{\xi}^T) + \sum_{t\in[2,T]} \sum_{i\in[I]} \sum_{j\in[J]} \beta^x_{ijt}(\boldsymbol{\xi}^T) x_{ijt}(\boldsymbol{\xi}^T) \\
&\qquad \text{s.t.} \quad z \geq \sum_{i\in[I]} y_i + \sum_{t\in[T]} \sum_{i\in[I]} a_i s_{it}(\boldsymbol{\xi}^T) + \sum_{t\in[2,T]} \sum_{i\in[I]} \sum_{j\in[J]} c_{ij} x_{ijt}(\boldsymbol{\xi}^T) && \boldsymbol{\xi}^T \in \Xi \\
&\qquad\qquad s_{i1} \leq K_i y_i, && i \in [I] \\
&\qquad\qquad s_{it}(\boldsymbol{\xi}^T) = s_{i,t-1}(\boldsymbol{\xi}^{T-1}) - \sum_{j\in[J]} x_{ijt}(\boldsymbol{\xi}^T) && i \in [I], t \in [2,T], \boldsymbol{\xi}^T \in \Xi \\
&\qquad\qquad \sum_{i\in[I]} x_{ijt}(\boldsymbol{\xi}^T) \geq d_{jt}(\boldsymbol{\xi}^t) && i \in [I], t \in [2,T], \boldsymbol{\xi}^T \in \Xi \\
&\qquad\qquad y \in \{0,1\}^I, z \in \mathbb{R}_+^{I\times T}, x \in \mathbb{R}_+^{I\times J\times(T-1)}.
\end{aligned}
$$

For this mixed-integer subproblem, the cutting-plane method presented in Section 3.3 is applicable.

C. Detailed Results

C.1. Newsvendor Problem

Table 4: Detailed results for the newsvendor problem

T	BR	$\|\Xi\|$	I	B	ν^{LDR}	$\nu^{2\text{S}}$	$\nu^{\star}$	$\bar{\nu}_R^{\text{NA}}$	ν^{PI}
3	5	25	3	100	656.6	970.2	975.3	975.3	1457.0
			3	150	6668.5	8410.4	8647.2	9794.5	10036.0
			3	200	8662.0	11446.2	11569.7	13709.0	13709.0
			4	150	5154.8	5384.4	5466.8	5466.8	5781.0
			4	200	10049.4	12452.0	12759.9	14501.9	14525.0
			5	200	5944.7	7120.6	7402.4	7402.4	7713.0
3	10	100	3	150	5752.1	7620.9	7984.6	8464.9	8895.0
			3	200	6279.6	9199.0	9700.9	11833.0	11833.0
			4	150	3502.7	4886.5	4910.2	5002.5	5323.0
			4	200	7922.9	10961.1	11273.4	14036.5	14060.0
			5	200	2865.5	5877.7	6063.1	6063.1	6737.0
4	4	64	2	100	669	1047	1214	1214	1337
			3	200	7807	9198	9491	9492	10157
			3	300	13421	16414	17317	19188	19188
			4	200	2034	2297	2447	2447	2640
			4	300	13645	17860	18721	20095	20139
4	5	125	3	200	6643.8	8070.1	8295.6	8760.0	9249.0
			3	300	11854.3	15035.4	16606.8	18240.0	18240.0
			4	200	1443.7	1700.8	1739.8	1917.7	2122.0
			4	300	13186.4	17579.4	18450.8	19618.0	19635.0
			5	300	5142.0	10494.4	11299.5	11607.6	12343.0
4	10	1000	3	200	5648.8	8142.9		8876.4	9353.0
			3	300	9143.5	13687.6		17853.0	17853.0
			4	200	-104.4	440.0		642.4	919.0
			4	300	11029.5	15854.7		18400.7	18432.0
			5	300	724.9	6125.5		6748.6	7368.0
4	15	3375	3	200	5111.3	7072.0		7646.6	8222.0
			3	300	9030.9	13040.7		17615.0	17615.0
5	3	81	2	150	2490.9	2673.5	2763.4	2763.3	2975.4
			2	200	8888.4	11192.5	11313.7	11448.6	11525.6
			3	250	7809.0	8800.3	9065.4	9065.5	9382.3
			3	300	14702.3	16898.8	17214.3	17308.7	18156.1
5	4	256	3	300	12367.7	15591.0	15989.0	16252.5	17192.6
5	5	625	3	300	11134.5	15053.4		16191.6	16680.0
			4	300	3651.5	9494.8		10553.8	10628.0
5	6	1296	3	400	15013.4	20271.7		25157.0	25157.0
6	4	1024	3	400	15492.3	21457.9		28163.0	28163.0
			4	400	7124.3	14805.2		15426.2	15473.0
			4	500	14445.4	24887.7		32946.9	32973.0
7	3	729	3	300	24.3	1495.9		2137.3	2383.0
			3	400	12994.7	17774.1		19554.5	19983.0
			4	400	3267.5	4303.1		5354.3	5427.0

Table 4: Detailed results for the newsvendor problem (continued)

T	BR	$\|\Xi\|$	I	B	ν^{LDR}	ν^{2S}	$\nu^{\star}$	$\bar{\nu}_R^{\text{NA}}$	PI
8	3	2187	3	500	16149.1	25321.4		30259.0	30259.0
			4	600	16892.3	27608.1		28717.7	28747.0

Table 5 Running time for the larger size newsvendor problem instances

T	BR	$\|\Xi\|$	I	B	Time (s)			
					ν^{LDR}	ν^{2S}	$\bar{\nu}_R^{\text{NA}}$	PI
4	10	1000	3	200	38	31	7561	1
			3	300	50	31	24	1
			4	200	47	28	625	1
			4	300	58	30	2385	1
			5	300	59	33	3311	2
4	15	3375	3	200	158	140	642	3
4	20	8000	3	300	726	675	1077	7
5	5	625	3	300	16	15	1426	1
			4	300	22	19	5053	1
5	6	1296	3	400	54	54	53	1
6	4	1024	3	400	57	53	48	1
			4	400	70	67	5911	2
			4	500	74	69	3846	2
7	3	729	3	300	37	38	525	1
			3	400	37	36	13010	1
			4	400	50	50	12874	1
8	3	2187	3	500	364	335	358	4
			4	600	462	464	37409	5

C.2. The Location-Transportation Problem

Table 6: Detailed results for the location-allocation problem

(T, I, J, α^d)	α^u	ν^{2S}	$\bar{\nu}_R^{NA}$
$(3, 10, 10, 0.1)$	0.1	763055.2	682934.4
	0.4	763053.2	651647.4
	0.7	763052.3	577630.6
	1	763054.7	569238.8
$(3, 10, 10, 0.3)$	0.1	1577694.3	1464100.3
	0.4	1577694.3	1364705.6
	0.7	1519050.9	1206126.4
	1	1575438.6	1224115.8
$(3, 10, 10, 0.5)$	0.1	1575813.5	1426111.2
	0.4	1577694.3	1363127.9
	0.7	1577694.3	1181693.0
	1	1555386.7	1474506.6
$(3, 10, 15, 0.1)$	0.1	1171035.9	857198.3
	0.4	1171035.9	831435.5
	0.7	1158076.3	1073536.8
	1	1158076.1	961203.2
$(3, 10, 10, 0.3)$	0.1	1171035.9	1067984.7
	0.4	1171035.9	863053.5
	0.7	1171035.9	1092576.5
	1	1171035.9	994209.5
$(3, 10, 10, 0.5)$	0.1	1171035.9	829093.4
	0.4	1171035.9	978986.0
	0.7	1158076.4	942674.2
	1	1171035.9	930973.5
$(3, 10, 20, 0.5)$	0.4	1672246.0	1282612.7
	1	1675102.0	1457338.7
$(4, 5, 5, 0.1)$	0.1	781934.0	724852.8
	0.4	781934.0	573157.6
	0.7	781934.0	687320.0
	1	781934.0	727980.6
$(4, 5, 5, 0.3)$	0.1	781934.0	710778.0
	0.4	781934.0	627111.1
	0.7	781934.0	666989.7
	1	781934.0	701394.8
$(4, 5, 5, 0.5)$	0.1	781934.0	583322.8
	0.4	781934.0	711559.9
	0.7	781934.0	620073.7
	1	781934.0	728762.5
$(4, 5, 7, 0.1)$	0.1	1044144.0	909449.4
	0.4	1044144.0	848889.1
	1	1044144.0	829050.3

Table 6: Detailed results for the location-allocation problem (continued)

(T, I, J, α^d)	α^u	ν^{2S}	$\bar{\nu}_R^{NA}$
(4, 5, 7, 0.3)	0.1	1044144.0	847844.9
	0.4	1044144.0	825917.9
	0.7	1044144.0	984627.8
	1	1044144.6	855154.4
(4, 5, 7, 0.5)	0.1	1044144.0	829050.3
	0.4	1044144.8	772667.2
	0.7	1044144.3	864551.5
	1	1044144.7	816521.2
(4, 5, 10, 0.1)	0.1	2327093.7	2008281.8
	0.4	2327093.7	1738339.0
	0.7	2327093.7	2168851.3
	1	2327093.7	1954758.7
(4, 5, 10, 0.3)	0.1	2327093.7	1912871.0
	0.4	2327093.7	1663872.0
	0.7	2327093.7	1752301.5
	1	2327094.0	2064132.4
(4, 5, 10, 0.5)	0.1	2327093.7	1733684.8
	0.4	2330349.6	2078671.9
	0.7	2327093.7	1959412.9
	1	2327093.7	1824441.4
(4, 10, 10, 0.1)	0.1	2408515.07	2071323.0
	0.4	2408515.14	1854556.7
	0.7	2408515.1	1719679.8
	1	2408515.4	2237510.8
(4, 10, 10, 0.3)	0.1	2359826.5	2189919.0
	0.4	2408515.1	2083365.6
	0.7	2408515.5	2083365.9
	1	2377021.4	1739979.6
(4, 10, 10, 0.5)	0.1	2359826.5	2130923.4
	0.4	2408515.1	2203791.3
	0.7	2359826.5	1923258.6
	1	2359826.5	1781669.0
(5, 5, 10, 0.1)	0.1	3436537.5	3075701.1
	0.4	3478066.7	2855492.7
	0.7	3442732.0	2785170.2
	1	3442731.5	3256824.0
(5, 5, 10, 0.3)	0.1	3478066.1	2695501.3
	0.4	3442731.5	2840253.5
	0.7	3478066.1	2483339.2
	1	3427190.6	2450441.3
(5, 5, 10, 0.5)	0.1	3442732.0	2478767.0
	0.4	3478066.2	2625940.0
	0.7	3469376.2	2976724.8
	1	3436537.5	2543037.8

Table 7: Algorithmic details for the location-allocation problem

(T, I, J, α^d)	α^u	$\nu^{\text{2S-LDR}}$, C&CG		$\bar{\nu}_R^{\text{NA}}$, Cutting-plane	
		#iterations	Time (s)	#iterations	Time (s)
$(3, 10, 10, 0.1)$	0.1	14	147.3	612	1077.9
	0.4	2	55.6	795	1109.5
	0.7	2	84.0	718	955.4
	1	2	104.8	1313	$>$ 10h
$(3, 10, 10, 0.3)$	0.1	12	834.2	1000	$>$ 10h
	0.4	2	47.9	222	937.4
	0.7	2	64.1	726	$>$ 10h
	1	2	66.6	976	541.0
$(3, 10, 10, 0.5)$	0.1	2	104.6	267	$>$ 10h
	0.4	2	102.8	655	$>$ 10h
	0.7	2	81.2	1045	679.4
	1	2	94.5	1387	$>$ 10h
$(3, 10, 15, 0.1)$	0.1	2	146.1	240	2584.8
	0.4	2	107.6	333	$>$ 10h
	0.7	8	91.6	414	$>$ 10h
	1	5	28.3	808	$>$ 10h
$(3, 10, 10, 0.3)$	0.1	2	125.7	766	4155.8
	0.4	2	93.6	959	4206.4
	0.7	2	90.6	455	4640.6
	1	2	113.9	1186	5003.7
$(3, 10, 10, 0.5)$	0.1	2	93.1	1240	3458.0
	0.4	4	151.3	802	2402.3
	0.7	2	108.0	870	$>$ 10h
	1	2	85.9	1085	3419.8
$(3, 10, 20, 0.5)$	0.4	2	3558.8	1137	8746.5
	1	2	1915.8	492	7273.0
$(4, 5, 5, 0.1)$	0.1	2	0.4	950	887.6
	0.4	2	0.5	1477	$>$ 10h
	0.7	2	0.5	626	1182.8
	1	2	0.5	1241	$>$ 10h
$(4, 5, 5, 0.3)$	0.1	2	0.4	1360	1186.9
	0.4	2	0.4	552	$>$ 10h
	0.7	2	0.4	1010	739.6
	1	2	0.4	543	$>$ 10h
$(4, 5, 5, 0.5)$	0.1	2	0.3	546	$>$ 10h
	0.4	2	0.5	282	$>$ 10h
	0.7	2	0.8	873	791.2
	1	2	0.4	900	$>$ 10h
$(4, 5, 7, 0.1)$	0.1	11	76.2	315	1184.0
	0.4	2	131.8	419	$>$ 10h
	1	7	153.8	1049	$>$ 10h
$(4, 5, 7, 0.3)$	0.1	2	113.2	851	838.7
	0.4	2	189.4	1341	916.0
	0.7	2	174.7	783	$>$ 10h
	1	2	150.8	965	$>$ 10h

Table 7: Algorithmic details for the location-allocation problem (continued)

(T, I, J, α^d)	α^u	$\nu^{\text{2S-LDR}}$, C&CG		$\bar{\nu}_R^{\text{NA}}$, Cutting-plane	
		#iterations	Time (s)	#iterations	Time (s)
$(4, 5, 7, 0.5)$	0.1	2	111.9	1226	742.0
	0.4	2	97.3	381	831.4
	0.7	17	131.8	559	704.9
	1	2	196.7	588	$>$ 10h
$(4, 5, 10, 0.1)$	0.1	5	78.7	1499	$>$ 10h
	0.4	11	87.8	721	$>$ 10h
	0.7	11	100.2	478	4062.0
	1	2	115.4	547	$>$ 10h
$(4, 5, 10, 0.3)$	0.1	2	196.3	1021	7742.2
	0.4	2	99.5	652	11814.1
	0.7	2	48.2	572	11408.6
	1	2	106.1	403	8880.7
$(4, 5, 10, 0.5)$	0.1	2	179.7	358	$>$ 10h
	0.4	2	34.2	1427	$>$ 10h
	0.7	2	71.6	608	11996.6
	1	2	39.1	737	10706.6
$(4, 10, 10, 0.1)$	0.1	2	13054.8	1822	$>$ 10h
	0.4	2	12710.6	2147	$>$ 10h
	0.7	2	4352.5	758	$>$ 10h
	1	2	7139.0	1238	13663.5
$(4, 10, 10, 0.3)$	0.1	6	10709.5	868	9490.9
	0.4	10	7592.3	746	$>$ 10h
	0.7	2	12070.4	1330	$>$ 10h
	1	2	7577.7	842	12937.5
$(4, 10, 10, 0.5)$	0.1	2	8261.0	758	10538.8
	0.4	2	8988.9	791	$>$ 10h
	0.7	2	5248.9	1461	14107.7
	1	2	3049.6	684	$>$ 10h
$(5, 5, 10, 0.1)$	0.1	2	349.6	1401	$>$ 10h
	0.4	2	978.5	248	$>$ 10h
	0.7	2	647.6	586	14988.1
	1	2	757.6	884	$>$ 10h
$(5, 5, 10, 0.3)$	0.1	2	863.5	779	$>$ 10h
	0.4	2	589.0	1053	16150.2
	0.7	2	402.9	757	12702.8
	1	2	941.7	997	14375.6
$(5, 5, 10, 0.5)$	0.1	2	1008.4	409	$>$ 10h
	0.4	4	721.7	1310	12800.1
	0.7	12	379.6	911	15108.2
	1	4	720.3	518	13741.1