Lagrangian Dual Decision Rules for Multistage Stochastic Mixed Integer Programming

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

James R. Luedtke
Department of Industrial and Systems Engineering, University of Wisconsin, Madison, Wisconsin 53706, jim.luedtke@wisc.edu

Multistage stochastic programs can be approximated by restricting policies to follow decision rules. Directly applying this idea to problems with integer decisions is difficult because of the need for decision rules that lead to integral decisions. In this work, we introduce *Lagrangian dual decision rules* (LDDRs) for multistage stochastic mixed integer programming (MSMIP) which overcome this difficulty by applying decision rules in a Lagrangian dual of the MSMIP. We propose two new bounding techniques based on stagewise (SW) and nonanticipative (NA) Lagrangian duals where the Lagrangian multiplier policies are restricted by LDDRs. We demonstrate how the solutions from these duals can be used to drive primal policies. Our proposal requires fewer assumptions than most existing MSMIP methods. We compare the theoretical strength of the restricted duals and show that the restricted NA dual can provide relaxation bounds at least as good as the ones obtained by the restricted SW dual. In our numerical study on two problem classes, one traditional and one novel, we observe that the proposed LDDR approaches yield significant optimality gap reductions compared to existing general-purpose bounding methods for MSMIP problems.

Key words: Multistage stochastic mixed integer programming, decision rules, Lagrangian dual, two-stage approximation, sampling

1. Introduction

Multistage stochastic mixed integer programming (MSMIP) is a framework to model an optimization problem involving stochastic uncertainty, where the planning horizon is divided into multiple stages, decisions are made in each stage, and some of these decisions are constrained to be integer.

The decisions in different stages cannot be made independently as they may impact subsequent stage decisions. The uncertainty is modeled as a stochastic process where the outcomes of random variables are observed over stages. In this setting, at each stage, the corresponding set of uncertain parameter values are observed, and based on this observation, the next stage decisions are made (see Figure 1). Therefore, as functions of random variables, decision variables at each stage are random variables themselves. However, they are *nonanticipative*, i.e., they only depend on the history of observations, not future realizations. Thus, in a multistage stochastic programming model, the solution is a *policy* or *decision rule* that maps all the past information to the current decisions to be made.

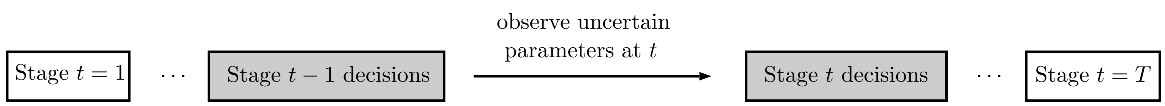

Figure 1: Dynamics of multistage stochastic programming

MSMIP naturally arises in many applications, such as unit commitment (Takriti et al. 1996), capacity expansion (Rajagopalan et al. 1998, Ahmed et al. 2003, Singh et al. 2009), generation scheduling in hydro systems (Flatabø et al. 1998, Nowak and Römisch 2000, Mo et al. 2001, Flach et al. 2010, Helseth et al. 2015), batch sizing (Lulli and Sen 2004), airline fleet composition (Listes and Dekker 2005), transmission investment planning (Newham and Wood 2007), transportation network protection (Fan and Liu 2010) and surgery planning (Gul et al. 2015).

The majority of solution methods for MSMIP require strong assumptions, including but not limited to stagewise independence, only right-hand-side uncertainty, binary state variables, and a finite (and not too large) scenario tree representation. Models lacking these conditions sometimes can be reformulated to satisfy the required assumptions, but at the expense of introducing (a potentially large number of) new variables and constraints. The most common assumption in the existing MSMIP methodologies is that the stochastic process is represented by a scenario tree. However, in general the size of a scenario tree required to obtain good quality solutions grows exponentially with the number of stages (Shapiro and Nemirovski 2005), which makes methods that rely on scenario tree models unviable when the number of stages is beyond three or four. In this paper we propose new approximation approaches for MSMIP problems that do not require these assumptions. These approaches are based on Lagrangian dual of an MSMIP model and tractability is achieved by considering restricted forms of Lagrangian multipliers, i.e., forcing them to follow some *decision rules*.

1.1. Related Literature

Even without integer decision variables, solving an MSMIP problem is theoretically and computationally challenging, due to high dimensional integration and the need to consider, while making current decisions, the (optimal) future decisions that will be made in response to the uncertain future trajectory of the stochastic process. Accordingly, all existing methods solve the problem by some form of approximation. There are three common *approximation approaches* in the literature.

First is to model the underlying stochastic process in the form of a *scenario tree* (Shapiro et al. 2009), which is the most common strategy in the existing solution methods for MSMIP. Given a finite scenario tree approximation, the MSMIP problem is converted to a (very) large-scale, structured, deterministic problem. Techniques for solving such a scenario-tree based approximation include bounding techniques (Norkin et al. 1998, Carøe and Schultz 1999, Ahmed et al. 2003, Alonso-Ayuso et al. 2003, Lulli and Sen 2004, Singh et al. 2009), cutting-plane based methods (Guan et al. 2009), aggregation approaches (Sandıkçı and Ozaltın 2014), nonanticipative Lagrangian dual approaches (Takriti et al. 1996, Chen et al. 2002), and progressive hedging algorithms (Løkketangen and Woodruff 1996, Listes and Dekker 2005, Fan and Liu 2010, Watson and Woodruff 2011, Gul et al. 2015, Gade et al. 2016). A limitation of the scenario tree approach is that, in order to obtain a good quality approximation of the stochastic process, the size of the tree in general needs to grow exponentially with the number of stages (Shapiro and Nemirovski 2005).

Stochastic dual dynamic programming (SDDP), proposed by Pereira and Pinto (1991), is a leading solution method for solving multistage stochastic linear programming (MSLP) problems as it is able to solve problems having an implicitly represented exponentially large scenario tree, under the assumption of stagewise independence. SDDP has been extended to MSMIP by considering various approximations of the non-convex cost-to-go functions (Newham and Wood 2007, Cen 2012, Cerisola et al. 2012, Löhndorf et al. 2013, Philpott et al. 2016). Proposing a new class of cutting planes, Zou (2017) introduced stochastic dual dynamic integer programming (SDDiP) algorithm, an extension of SDDP which further assumes binary state decision variables. SDDiP is able to overcome some of its restrictions with various forms of reformulations, but at the expense of introducing new decision variables, thus increasing the size of the model. For instance, bounded integer state variables can be handled via binarization schemes. Most recently, Ahmed et al. (2020) proposed stochastic Lipschitz dynamic programming for MSMIP problems with general integer variables and Lipschitz cost-to-go functions, which uses Lipschitz cuts in its backward pass. The need to explicitly convexify the value function of many mixed-integer programming problems may potentially limit the scalability of this approach.

The second approximation approach, commonly known as the *decision rule approach*, restricts the policies to follow a certain form, rather than restricting the form of the stochastic process. In the context of MSLP, Shapiro and Nemirovski (2005) presented an upper bounding technique employing primal decision rules, where the decisions at each stage are restricted to be an affine function of observed random outcomes up to that stage. This yields a static linear decision rule (LDR) policy. On the other hand, Kuhn et al. (2011) provided a lower bounding technique by applying LDRs to dual policies. Bodur and Luedtke (2018) introduced a two-stage approximation by restricting only state variables (i.e., the ones linking two consecutive stages together) to follow LDRs. By applying the so-called two-stage LDRs to primal and dual policies, the authors provided improved upper and lower bounds for MSLP problems. Decision rules in the form of polynomial (Bampou and Kuhn 2011), piecewise linear (Chen et al. 2008), bilinear and trilinear (Georghiou et al. 2015) functions have also been examined in the literature. In this approximation approach, lower bounding techniques are limited to MSLP problems. We extend this line of research to

MSMIP by proposing the use of decision rules to obtain tractable approximations of *Lagrangian* duals of MSMIP problems.

The final approximation approach includes work that does not assume a scenario tree but relies instead on exploiting *problem structure*. Brown et al. (2010) study information relaxation for stochastic dynamic programs and, similar to our approach, penalize the violation of nonanticipativity constraints (i.e., the ones that ensure consistency among the decisions when perfect information about the future is assumed). Their penalty is based on a different dual, however, and the approach requires determining problem-specific penalty functions that balance computational tractability with the strength of the obtained bound. Barty et al. (2010) propose dual approximate dynamic programming, which provides approximations for problems that can be decomposed into smaller, tractable problems when certain linking constraints are relaxed.

1.2. Our Contributions

We introduce *Lagrangian dual decision rules* (LDDRs) to obtain bounds for MSMIP problems. We design *two new lower (upper) bounding techniques* for general MSMIP problems with minimization (maximization) objective, i.e., with mixed-integer state and recourse variables, without restricting the form of the underlying stochastic process. These bounding techniques are based on two Lagrangian relaxations of the MSMIP model: stagewise (SW) and nonanticipative (NA), where the state equations linking consecutive stages and the nonanticipativity constraints are relaxed, respectively. In order to obtain tractable approximations, we restrict the associated Lagrangian multiplier policies to follow a decision rule determined by parameters which are optimized to achieve the best possible bound given the restriction. We compare the theoretical strength of the restricted duals and show that, when appropriately constructed, the restricted NA dual provides a relaxation bound at least as good as the bound obtained by the restricted SW dual.

We also develop *two LDDR driven* primal policies which can be used to obtain upper (lower) bounds for problems with minimization (maximization) objective, that incorporate the Lagrangian multipliers obtained from the lower (upper) bounding methods. We perform an extensive computational study on a stochastic multi-item lot-sizing problem with lag where demands follow an autoregressive process, and find that our SW dual based primal method returns good quality solutions, while our NA dual driven primal method provides better solutions at the expense of extra computational effort. Putting the obtained lower and upper bounds together, we observe that the LDDR restricted NA dual approach yields significant optimality gap reductions compared to standard general purpose bounding methods for MSMIP problems. Moreover, we study a novel multistage provisioning problem in optical networks, namely routing and wavelength assignment problem, and show the considerable merit of using our framework in reducing the optimality gap from classical bounding methods.

Our approximation approach is general purpose and free of strong assumptions made in the literature such as stagewise independence or existence of a tractable-sized scenario tree representation. To the best of our knowledge, this is the first approach based on decision rules that is capable of handling mixed-integer state variables. Our approach leads to subproblems in the form of deterministic mixed-integer programs, and thus can exploit the state of the art in efficiently solving them. Moreover, as the form of the restricted Lagrangian dual problems is a two-stage stochastic program, our approach enables application of theory and methods for solving two-stage stochastic integer programming models.

We summarize the main contributions of our work as follows:

- **Restricted Lagrangian dual problems**. Our proposal to apply the decision rule restriction to the two considered Lagrangian dual problems enables the use of these dual problems to obtain bounds for MSMIP problems without requiring a small-size scenario tree. In contrast, solving the full SW or NA Lagrangian dual problems is only feasible for MSMIP models in which the stochastic process is represented with a scenario tree where the total number of nodes is of manageable size, dramatically limiting both the number of stages and the number of branches per stage that can be used to approximate the underlying stochastic process. For example, in a capacity planning problem by Singh et al. (2009) it was observed that the NA dual could not be solved in a reasonable time for an instance defined by a scenario tree with four stages and eight branches per stage. We remark that while dual decision rule restrictions have previously been applied for MSLPs (Kuhn et al. 2011, Bodur and Luedtke 2018), our work is the first to use such restrictions for MSMIP.

- **Theoretical comparison of Lagrangian dual bounds**. Previous theory on the unrestricted Lagrangian SW and NA dual problems demonstrates that the NA dual is at least as strong as the SW dual. However, this theory does not apply when both dual problems are restricted as in our proposal. We extend this theory to demonstrate how the basis functions for the NA dual can be selected, relative to the ones for the SW dual, to assure that the LDDR-restricted NA dual bound is at least as good as that of the LDDR-restricted SW dual bound.

- **Dual-driven primal policies**. We develop new primal policies that use information from the solution of the restricted Lagrangian dual problems to guide decision making at each decision stage.

- **Empirical study**: We conduct computational experiments on two problem classes. We validate our theory that the bounds obtained by the restricted NA dual problem are better than those obtained by the restricted SW dual problem, and demonstrate that the former can improve significantly over the classical perfect information bound. We also find that the proposed dual-driven primal policies can outperform a policy driven based on solving deterministic approximations. Moreover, one of our applications contributes to the telecommunications literature by presenting the first MSMIP model for a widely-studied provisioning problem in optical networks.

The remainder of this paper is organized as follows. In Section 2 we formally state the MSMIP problem and two Lagrangian dual problems arising from it. In Section 3 we introduce new relaxation/bounding methods for an MSMIP model that are based on those Lagrangian relaxations. In Section 4 we provide feasible policies designed by using the information obtained from the restricted duals. In Section 5 we evaluate the proposed methodologies on two problem classes.

Notation. Random variables are represented with bold letters ($\boldsymbol{\xi}$) while their observations are regular font (ξ). 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 $(.)^\top$ for the transpose operator.

2. Problem Statement

Let T denote the number of decision stages and $\boldsymbol{\xi}_t$ be a random vector at stage $t \in [T]$ with outcomes $\xi_t \in \mathbb{R}^{\ell_t}$, and $\xi_1 = 1$ (i.e., the first stage is deterministic). The stochastic process is represented by $\{\boldsymbol{\xi}_t\}_{t=1}^T$ having probability distribution $\mathbb{P}$ and support Ξ. By $\boldsymbol{\xi}^t = (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_t)$, we denote the history of the process at stage t. By including $\xi_1 = 1$ in the history, we make sure that the stage t decisions are functions of the entire history, including the first deterministic stage.

An MSMIP problem can be modeled as follows

$$\min \quad \mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]} c_t(\boldsymbol{\xi}^t)^\top x_t(\boldsymbol{\xi}^t)\right] \tag{1a}$$

$$\text{s.t.} \quad A_t(\boldsymbol{\xi}^t)x_t(\boldsymbol{\xi}^t) + B_t(\boldsymbol{\xi}^t)x_{t-1}(\boldsymbol{\xi}^{t-1}) = b_t(\boldsymbol{\xi}^t), \quad t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \tag{1b}$$

$$x_t(\boldsymbol{\xi}^t) \in X_t(\boldsymbol{\xi}^t), \qquad t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \tag{1c}$$

where $c_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{n_t}$, $A_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\mathrm{S}} \times n_t}$, $B_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\mathrm{S}} \times n_{t-1}}$, $b_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\mathrm{S}}}$ and $\ell^t = \sum_{t'=1}^t \ell_{t'}$. $x_t(\boldsymbol{\xi}^t) \in \mathbb{Z}^{p_t} \times \mathbb{R}^{n_t - p_t}$ are the decision variables, i.e., nonanticipative policies, for $t \in [T]$. Here and throughout the paper we adopt the convention $x_0(\xi^0) \equiv 0$. $\mathbb{P}$-a.e. $\boldsymbol{\xi}^T \in \Xi$ means that the constraints are required to be almost surely satisfied with respect to $\mathbb{P}$. Constraints (1b) and (1c) are *state* and *recourse* constraints, respectively, where $X_t(\boldsymbol{\xi}^t) := \{x \in \mathbb{Z}^{p_t} \times \mathbb{R}^{n_t - p_t} : C_t(\boldsymbol{\xi}^t)x \geq d_t(\boldsymbol{\xi}^t)\}$ with $C_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\mathrm{R}} \times n_t}$ and $d_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\mathrm{R}}}$.

Throughout this work, we make the following assumptions:

Assumption 1. *Model (1) has relatively complete recourse, i.e., at every stage $t \in [T]$, given any set of feasible decisions and random variable realizations of the previous stages, there always exists a feasible set of decisions.*

Assumption 2. *$X_t(\boldsymbol{\xi}^t)$ is compact for all $t \in [T]$, $\mathbb{P}$-a.e. $\boldsymbol{\xi}^T \in \Xi$, and the expected diameter of $X_t(\boldsymbol{\xi}^t)$ is finite.*

Assumption 1 is common in the stochastic programming literature. It is naturally satisfied in many applications, even in the stronger form of complete recourse (Rockafellar and Wets 1976, Lulli and Sen 2006), e.g., Zou et al. (2017) and Ahmed et al. (2020) both assume complete continuous recourse. Otherwise, it can be achieved by modifying the model to include slack variables that measure constraint violations (Shapiro and Nemirovski 2005, Zou et al. 2017) which are then penalized in the objective. Assumption 2 refers to the boundedness of the feasibility set, which is often the case in practice. If it is not readily implied by the problem definition, it can be enforced by adding sufficiently large (small) upper (lower) bounds on decision variables.

We next describe two Lagrangian duals for MSMIP, namely SW dual (Rosa and Ruszczyński 1996) and NA dual (Rockafellar and Wets 1991).

2.1. Stagewise Lagrangian Dual

In the SW dual, the constraints linking consecutive stages are relaxed, namely the *state equations*. Let $\pi_t(\boldsymbol{\xi}^t) \in \mathbb{R}^{m_t^{\mathrm{s}}}, t \in [T]$ be the dual variables associated with constraints (1b). For fixed dual functions $\pi_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\mathrm{s}}}, t \in [T]$, the SW Lagrangian relaxation problem is defined as

$$\mathfrak{L}^{\mathrm{SW}}(\pi_1,\ldots,\pi_T) = \min \; \mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]} c_t(\boldsymbol{\xi}^t)^\top x_t(\boldsymbol{\xi}^t) + \pi_t(\boldsymbol{\xi}^t)^\top \left(A_t(\boldsymbol{\xi}^t)x_t(\boldsymbol{\xi}^t) + B_t(\boldsymbol{\xi}^t)x_{t-1}(\boldsymbol{\xi}^{t-1}) - b_t(\boldsymbol{\xi}^t)\right)\right] \tag{2a}$$

$$\text{s.t. } x_t(\boldsymbol{\xi}^t) \in X_t(\boldsymbol{\xi}^t), \quad t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi, \tag{2b}$$

which decomposes by t and ξ, thus yields a deterministic mixed integer program (MIP) per stage and scenario. As the Lagrangian relaxation problem provides a valid lower bound on the optimal value of the MSMIP problem, the stagewise Lagrangian dual problem aims to find the dual functions providing the best bound

$$\nu^{\mathrm{SW}} := \max_{\{\pi_t\}_{t\in[T]}} \mathfrak{L}^{\mathrm{SW}}(\pi_1,\ldots,\pi_T). \tag{3}$$

2.2. Nonanticipative Lagrangian Dual

The NA dual is based on a reformulation of the MSMIP problem where we create a copy of every decision variable for every realization and explicitly enforce nonanticipativity. We introduce the copy variables $y(\boldsymbol{\xi}^T) = (y_1(\boldsymbol{\xi}^T),\ldots,y_T(\boldsymbol{\xi}^T))$ as perfect information variables, meaning that they depend on the entire sample path $\boldsymbol{\xi}^T = (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_T)$. For every sample path $\boldsymbol{\xi}^T \in \Xi$, we define the set

$$Y(\boldsymbol{\xi}^T) = \left\{ y \in \mathbb{Z}^{p_t} \times \mathbb{R}^{n_t - p_t} : A_t(\boldsymbol{\xi}^t) y_t + B_t(\boldsymbol{\xi}^t) y_{t-1} = b_t(\boldsymbol{\xi}^t), y_t \in X_t(\boldsymbol{\xi}^t), t \in [T] \right\}.$$

Then, the MSMIP problem (1) can be reformulated as

$$\min \ \mathbb{E}_{\boldsymbol{\xi}^T} \left[\sum_{t \in [T]} c_t(\boldsymbol{\xi}^t)^\top y_t(\boldsymbol{\xi}^T) \right] \tag{4a}$$

$$\text{s.t.} \ \ y(\boldsymbol{\xi}^T) \in Y(\boldsymbol{\xi}^T), \qquad \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \tag{4b}$$

$$y_t(\boldsymbol{\xi}^T) = \mathbb{E}_{\boldsymbol{\xi}'^T} \left[y_t(\boldsymbol{\xi}'^T) \middle| \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t \right], \quad t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi. \tag{4c}$$

Constraints (4c) are *nonanticipativity constraints* which make sure that for every partial realization of a sample path $\boldsymbol{\xi}^t$ at stage t, the decisions made at stage t are consistent (i.e., the decisions made in all sample paths $\boldsymbol{\xi}^T$ that share the history $\boldsymbol{\xi}^t$ are the same). Associating the dual functions $\gamma_t(\cdot) \in \Gamma_t, t \in [T]$, where $\Gamma_t = \{\gamma_t : \mathbb{R}^{\ell^T} \to \mathbb{R}^{n_t} \mid \mathbb{E}[\gamma_t(\boldsymbol{\xi}^T)] < \infty\}$, the nonanticipativity constraints are relaxed to obtain the NA Lagrangian dual problem

$$\mathfrak{L}^{\mathrm{NA}}(\gamma_1, \ldots, \gamma_T) = \min \ \mathbb{E}_{\boldsymbol{\xi}^T} \left[\sum_{t \in [T]} c_t(\boldsymbol{\xi}^t)^\top y_t(\boldsymbol{\xi}^T) + \gamma_t(\boldsymbol{\xi}^T)^\top \left(y_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t] \right) \right] \tag{5a}$$

$$\text{s.t.} \ \ y(\boldsymbol{\xi}^T) \in Y(\boldsymbol{\xi}^T), \quad \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi. \tag{5b}$$

The objective function in (5a) can be simplified using the following lemma whose proof can be found in the e-companion. This lemma allows us to swap the Lagrangian multipliers $\gamma_t(\boldsymbol{\xi}^T)$ and decision variables inside the expectation $y_t(\boldsymbol{\xi}'^T)$, which makes it possible to compute the coefficients ahead of the time.

Lemma 1. *Assume that the expected diameter of $X(\boldsymbol{\xi}^t)$ is finite. Then, the following equality holds*

$$\mathbb{E}_{\boldsymbol{\xi}^T} \left[\sum_{t \in [T]} c_t(\boldsymbol{\xi}^t)^\top y_t(\boldsymbol{\xi}^T) + \gamma_t(\boldsymbol{\xi}^T)^\top \left(y_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t] \right) \right] = \tag{6}$$

$$\mathbb{E}_{\boldsymbol{\xi}^T} \left[\sum_{t \in [T]} \left(c_t(\boldsymbol{\xi}^t) + \gamma_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[\gamma_t(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t] \right)^\top y_t(\boldsymbol{\xi}^T) \right].$$

Using Lemma 1, the nonanticipative Lagrangian relaxation for fixed dual functions $\gamma_t : \mathbb{R}^{\ell^T} \to \mathbb{R}^{n_t}, t \in [T]$ can be written as

$$\begin{aligned} \mathfrak{L}^{\mathrm{NA}}(\gamma_1, \dots, \gamma_T) = \min \ & \mathbb{E}_{\boldsymbol{\xi}^T} \left[\sum_{t \in [T]} \left(c_t(\boldsymbol{\xi}^t) + \gamma_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[\gamma_t(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t] \right)^\top y_t(\boldsymbol{\xi}^T) \right] \\ \text{s.t.} \ & y(\boldsymbol{\xi}^T) \in Y(\boldsymbol{\xi}^T), \quad \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi, \end{aligned}$$

which is decomposable by sample path, but not by stage. Finally the NA Lagrangian dual problem is

$$\nu^{\mathrm{NA}} := \max_{\{\gamma_t \in \Gamma_t\}_{t \in [T]}} \mathfrak{L}^{\mathrm{NA}}(\gamma_1, \dots, \gamma_T). \tag{7}$$

2.3. Primal Characterizations and Bound Comparison

Using Lagrangian duality theory, primal characterizations of the Lagrangian duals can be obtained, which in turn can be used to compare their strength (Dentcheva and Römisch 2004). For the SW Lagrangian dual problem, the primal characterization is as follows

$$\nu^{\mathrm{SW}} = \min \ \mathbb{E}_{\boldsymbol{\xi}^T} \left[\sum_{t \in [T]} c_t(\boldsymbol{\xi}^t)^\top x_t(\boldsymbol{\xi}^t) \right] \tag{8a}$$

$$\text{s.t.} \ A_t(\boldsymbol{\xi}^t) x_t(\boldsymbol{\xi}^t) + B_t(\boldsymbol{\xi}^t) x_{t-1}(\boldsymbol{\xi}^{t-1}) = b_t(\boldsymbol{\xi}^t), \quad t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \tag{8b}$$

$$x_t(\boldsymbol{\xi}^t) \in \mathrm{conv}(X_t(\boldsymbol{\xi}^t)), \quad t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi. \tag{8c}$$

That is, for each stage and sample path, the feasible set of the recourse problem on that stage is relaxed and replaced with its convex hull. The primal characterization of the NA Lagrangian dual problem is given below

$$\begin{aligned} \nu^{\mathrm{NA}} = \min \ & \mathbb{E}_{\boldsymbol{\xi}^T} \left[\sum_{t \in [T]} c_t(\boldsymbol{\xi}^t)^\top y_t(\boldsymbol{\xi}^T) \right] \\ \text{s.t.} \ & y_t(\boldsymbol{\xi}^T) \in \mathrm{conv}(Y(\boldsymbol{\xi}^T)), \quad \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \\ & y_t(\boldsymbol{\xi}^T) = \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t], \quad t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi. \end{aligned}$$

That is, for every sample path, the feasible set of the T-stage deterministic problems is relaxed and replaced with its convex hull.

Using these characterizations, Dentcheva and Römisch (2004) show, in the case when the stochastic process is represented by a finite scenario tree, that the NA dual is not worse than the SW dual, i.e.,

$$\nu^{\mathrm{SW}} \le \nu^{\mathrm{NA}}.$$

3. Lagrangian Dual Decision Rules for MSMIP

The SW dual functions π_t and the NA dual functions γ_t are policies which map every possible history of observations to a dual decision vector, making direct solution of the respective Lagrangian dual problems (3) and (7) intractable in general. We propose restricting these dual multipliers to follow decision rules, referred to as LDDRs. As such, we obtain restricted problems that have finitely many decision variables. In what follows, we first explain the LDDR approach for the SW and NA duals (Sections 3.1 and 3.2), then provide their primal characterizations for a strength comparison as in the unrestricted case (Section 3.3), and lastly provide an algorithmic framework for their solutions (Section 3.5).

3.1. Restricted Stagewise Lagrangian Dual

For $t \in [T]$, we restrict the dual variables $\pi_t(\boldsymbol{\xi}^t)$ to follow an LDDR by enforcing

$$\pi_t(\boldsymbol{\xi}^t) = \Phi_t(\boldsymbol{\xi}^t)\beta_t$$

where $\Phi_t : \mathbb{R}^{\ell^t} \to \mathbb{R}^{m_t^{\mathrm{s}} \times K_t}$ are the set of *basis functions* and $\beta_t \in \mathbb{R}^{K_t}$ is a vector of *LDDR decision variables* (i.e., the weights associated with the basis functions). For the ease of presentation, we use the matrix form of basis function outputs – i.e., one can think of Φ_t as consisting of K_t basis functions, each of which maps the history $\boldsymbol{\xi}^t$ to a vector of size m_t^{s}, the number of state equations in stage t. The set of basis functions is a model choice, and so K_t depends on this choice. We define the *LDDR-restricted SW Lagrangian dual* problem which aims to find the optimal choice of LDDR variable values to maximize the obtained lower bound

$$\nu_R^{\mathrm{SW}} := \max_{\{\beta_t\}_{t\in[T]}} \mathfrak{L}^{\mathrm{SW}}(\Phi_1\beta_1, \ldots, \Phi_T\beta_T) = \max_{\{\beta_t\}_{t\in[T]}} \mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]} \mathcal{L}_t^{\mathrm{SW}}(\beta_t, \boldsymbol{\xi}^t) - \beta_t^\top \Phi_t^\top(\boldsymbol{\xi}^t)\, b_t(\boldsymbol{\xi}^t)\right] \tag{9}$$

where $\mathcal{L}_t^{\mathrm{SW}}(\beta, \boldsymbol{\xi}^t) := \min\left\{ \left(c_t(\boldsymbol{\xi}^t) + \beta_t^\top \Phi_t^\top(\boldsymbol{\xi}^t)\ A_t(\boldsymbol{\xi}^t) + \beta_{t+1}^\top \mathbb{E}_{\boldsymbol{\xi}'^T}\left[\ \Phi_{t+1}^\top({\boldsymbol{\xi}'}^{t+1})\ B_{t+1}({\boldsymbol{\xi}'}^{t+1}) \mid \boldsymbol{\xi}^t\right]\right)^\top x_t : x_t \in X_t(\boldsymbol{\xi}^t)\right\}$. Observe that, for fixed $\boldsymbol{\xi}^t$, calculating the coefficient of β_{t+1} requires evaluating the conditional expectation $\mathbb{E}_{\boldsymbol{\xi}'^T}[\Phi_{t+1}^\top({\boldsymbol{\xi}'}^{t+1})B_{t+1}({\boldsymbol{\xi}'}^{t+1})|\boldsymbol{\xi}^t]$. The tractability of this calculation depends on the form of Φ_t and the dependence structure of the stochastic process. For instance, they can be directly calculated if Φ_t consists of affine functions and the conditional distribution of ${\boldsymbol{\xi}'}^{t+1}$ given $\boldsymbol{\xi}^t$ is known, such as when the random variables follow an autoregressive process. In general, these coefficients can be estimated via sampling.

3.2. Restricted Nonanticipative Lagrangian Dual

Letting $\Psi_t : \mathbb{R}^{\ell^T} \to \mathbb{R}^{n_t \times K'_t}$ for $t \in [T]$ to be a set of basis functions, we restrict the dual variables $\gamma_t(\boldsymbol{\xi}^T)$ to follow an LDDR

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

where $\alpha_t \in \mathbb{R}^{K'_t}, t \in [T]$, is the vector of LDDR decision variables. Then, we obtain the *LDDR-restricted NA Lagrangian dual* problem as

$$\nu_R^{\mathrm{NA}} := \max_{\{\alpha_t\}_{t\in[T]}} \mathfrak{L}^{\mathrm{NA}}(\Psi_1\alpha_1, \ldots, \Psi_T\alpha_T) = \max_{\{\alpha_t\}_{t\in[T]}} \mathbb{E}_{\boldsymbol{\xi}^T}\left[\mathcal{L}^{\mathrm{NA}}(\alpha, \boldsymbol{\xi}^T)\right], \tag{10}$$

where

$$\begin{aligned}
\mathcal{L}^{\mathrm{NA}}(\alpha, \boldsymbol{\xi}^T) &= \min\left\{ \sum_{t\in[T]} \left(c_t(\boldsymbol{\xi}^t) + \Psi_t(\boldsymbol{\xi}^T)\alpha_t - \mathbb{E}_{\boldsymbol{\xi}'^T}[\Psi_t({\boldsymbol{\xi}'}^T)\alpha_t | {\boldsymbol{\xi}'}^t = \boldsymbol{\xi}^t]\right)^\top y_t : y \in Y(\boldsymbol{\xi}^T)\right\} \\
&= \min\left\{ \sum_{t\in[T]} \left(c_t(\boldsymbol{\xi}^t) + \left(\Psi_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[\Psi_t({\boldsymbol{\xi}'}^T) | {\boldsymbol{\xi}'}^t = \boldsymbol{\xi}^t]\right)\alpha_t\right)^\top y_t : y \in Y(\boldsymbol{\xi}^T)\right\}.
\end{aligned} \tag{11}$$

Since basis functions $\Psi_t(\boldsymbol{\xi}^T)$ are functions of the entire sample path, depending on their structure we might need to evaluate the h-step ahead conditional expectations, for $1 \le h \le T-t$, while for the SW restricted dual we have $h = 1$. For example, if $\Psi_{tk'}(\boldsymbol{\xi}^T) = \boldsymbol{\xi}^{t+h}$, then $\mathbb{E}_{\boldsymbol{\xi}'^T}[\boldsymbol{\xi}^{t+h} | {\boldsymbol{\xi}'}^t = \boldsymbol{\xi}^t]$ needs to either be calculated by an analytical form (if possible), or estimated by means of sampling. See Sections 5.3.1 and 5.4.1 for some examples of computing the explicit form of the above conditional expectations.

3.3. Primal Characterizations and Bound Comparison

In this section we assume that the stochastic process is represented by a finite scenario tree, implying that for every stage $t \in [T]$, the random vector $\boldsymbol{\xi}_t$ has a discrete distribution with finite support, and also that the set of possible sample paths, Ξ, is finite. We make this assumption in this section only for the sake of simplicity in deriving primal characterizations for the proposed restricted Lagrangian duals, and note that we do not require the size of the scenario tree representation to be tractable as this assumption is used only for theoretical analysis, not for computational purposes. We emphasize that our overall methodological framework does not restrict the stochastic process to a finite scenario tree model, and rather relies on sampling to approximately solve the restricted approximations. The finite scenario tree assumption of this section is in line with the comparison of the bounds from unrestricted duals by Dentcheva and Römisch (2004).

We next derive a general primal characterization of the Lagrangian dual of a MIP when dual multipliers are restricted to a linear form. Since an MSMIP is a large-scale (structured) MIP under the finite scenario tree assumption, this analysis is sufficient to obtain the primal characterizations of our restricted duals of the MSMIP problem.

Consider a MIP $\min\{c^\top x : Dx = d,\ x \in X\}$ where $X \subseteq \mathbb{R}^n$ is a set defined by linear constraints and integer constraints on some of the decision variables, $D \in \mathbb{R}^{m\times n}$, $c \in \mathbb{R}^n$, and $d \in \mathbb{R}^m$. For $\lambda \in \mathbb{R}^m$ define

$$z(\lambda) = \min\{c^\top x + \lambda^\top (d - Dx) : x \in X\}. \tag{12}$$

The standard Lagrangian dual is the problem $z^{\mathrm{LD}} = \max_{\lambda} z(\lambda)$, and its primal characterization is

$$z^{\mathrm{LD}} = \min\{c^\top x : x \in \mathrm{conv}(X), Dx = d\}. \tag{13}$$

For a given matrix $G \in \mathbb{R}^{m\times K}$, define the the *restricted* Lagrangian dual

$$z^{\mathrm{RLD}} \;=\; \max_{\lambda,\alpha}\{z(\lambda) : \lambda = G\alpha\}. \tag{14}$$

The following Lemma provides a primal characterization of this restricted dual problem.

Lemma 2. *The restricted Lagrangian dual satisfies*

$$
\begin{aligned}
z^{RLD} \; &= \min_x \; c^\top x \\
&\quad s.t. \; x \in \operatorname{conv}(X) \\
&\qquad G^\top(Dx - d) = 0.
\end{aligned}
$$

In other words, in comparison to the unrestricted Lagrangian dual (13), the restricted Lagrangian dual still replaces the set X with its convex hull, but the constraints $Dx = d$ are relaxed to $G^\top(Dx - d) = 0$. The proof of Lemma 2 can be found in the e-companion. Using this lemma, we derive the primal characterizations of the restricted Lagrangian duals of the MSMIP.

As the stochastic process is assumed to be represented by a finite scenario tree, the probability of each sample path can be assumed to be positive. In order to apply Lemma 2, we scale the state equations in the MSMIP formulation (1) by the corresponding probabilities of these scenarios denoted by $p(\xi^T)$

$$
\begin{aligned}
\min \;\; & \mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]} c_t(\boldsymbol{\xi}^t)^\top x_t(\boldsymbol{\xi}^t)\right] \\
\text{s.t.} \;\; & p(\xi^T)(A_t(\xi^t)x_t(\xi^t) + B_t(\xi^t)x_{t-1}(\xi^{t-1})) = p(\xi^T)b_t(\xi^t), && t \in [T], \mathbb{P}\text{-a.e. } \xi^T \in \Xi \\
& x_t(\xi^t) \in X_t(\xi^t), && t \in [T], \mathbb{P}\text{-a.e. } \xi^T \in \Xi.
\end{aligned}
$$

Recall the LDDR restrictions for the duals

$$
\pi_t(\xi^t) - \Phi_t(\xi^t)\beta_t = 0, \quad t \in [T], \mathbb{P}\text{-a.e. } \xi^T \in \Xi.
$$

These are the new set of constraints added to the SW Lagrangian dual problem (3). Consequently, using Lemma 2, the primal characterization of the restricted SW Lagrangian dual problem is obtained as

$$
\nu_R^{\text{SW}} = \min \;\; \mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]} c_t(\boldsymbol{\xi}^t)^\top x_t(\boldsymbol{\xi}^t)\right] \tag{16a}
$$

$$
\text{s.t.} \;\; x_t(\xi^t) \in \operatorname{conv}(X_t(\xi^t)), \qquad t \in [T], \mathbb{P}\text{-a.e. } \xi^T \in \Xi \tag{16b}
$$

$$\mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_t(\boldsymbol{\xi}^t)^\top\left(A_t(\boldsymbol{\xi}^t)x_t(\boldsymbol{\xi}^t)+B_t(\boldsymbol{\xi}^t)x_{t-1}(\boldsymbol{\xi}^{t-1})-b_t(\boldsymbol{\xi}^t)\right)\right]=0, t\in[T]. \tag{16c}$$

Note that, for each $t\in[T]$, there are exactly K_t expected value constraints, one for each basis function vector in the matrix Φ_t. Compared to the primal characterization of the unrestricted SW dual (8), while the recourse problem feasible sets are still convexified (as in (16b)), the state equations are not almost surely satisfied, rather for every basis function, their expectation is enforced to be equal to zero. Therefore, unlike the unrestricted version (8), even if the original problem is convex (such as an MSLP), a duality gap might exist.

Using the same approach, we obtain the following primal characterization for the restricted NA Lagrangian dual problem

$$\nu_R^{\mathrm{NA}}=\min\ \mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]}c_t(\boldsymbol{\xi}^t)^\top y_t(\boldsymbol{\xi}^T)\right] \tag{17a}$$

$$\text{s.t.}\ \ y_t(\xi^T)\in\operatorname{conv}(Y(\xi^T)), \qquad \mathbb{P}\text{-a.e. } \xi^T\in\Xi \tag{17b}$$

$$\mathbb{E}_{\boldsymbol{\xi}^T}\left[\Psi_t(\boldsymbol{\xi}^T)^\top\left(y_t(\boldsymbol{\xi}^T)-\mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t=\boldsymbol{\xi}^t]\right)\right]=0,\quad t\in[T]. \tag{17c}$$

Here, constraint sets $Y(\xi^T)$ are replaced with their convex hull, but the NA constraints are enforced in expectation rather than for almost every $\boldsymbol{\xi}^T$.

As mentioned in Section 2.3, for the unrestricted duals, it is proven that $\nu^{\mathrm{SW}}\le\nu^{\mathrm{NA}}$. The same result does not immediately hold for the restricted duals due to the relaxed constraints. However, we prove that when the basis functions for the restricted NA dual are properly selected, a similar inequality holds for the restricted duals. Let $\Phi_t=(\Phi_{t1},\dots,\Phi_{tK_t})$ and $\Psi_t=(\Psi_{t1},\dots,\Psi_{tK'_t})$ be the basis functions used in the restricted SW and NA duals, respectively. Note that $\Phi_t:\mathbb{R}^{\ell^t}\to\mathbb{R}^{m_t^{\mathrm{S}}\times K_t}$ and $\Psi_t:\mathbb{R}^{\ell^T}\to\mathbb{R}^{n_t\times K'_t}$ are functions mapping a set of random variables to a column vector for each basis function indexed by k and k', respectively.

THEOREM 1. *Given a set of basis functions $\Phi_t=(\Phi_{t1},\dots,\Phi_{tK_t}), t\in[T]$ used in restricted SW dual, there exists $\Psi_t=(\Psi_{t1},\dots,\Psi_{tK'_t}), t\in[T]$ for the restricted NA dual such that:*

$$\nu_R^{SW}\le\nu_R^{NA}.$$

Proof In the following constructive proof, we first build a set of basis functions Ψ_t given Φ_t. Then we show that having such Ψ_t guarantees $\nu_R^{\mathrm{SW}} \le \nu_R^{\mathrm{NA}}$.

For each $t \in [T]$, we define the set of basis functions Ψ_t, with cardinality $K'_t = K_t + K_{t+1}$ for $t < T$ and $K'_T = K_T$. For $k \in [K_t]$:

1. $\Psi_{tk'}(\boldsymbol{\xi}^T)^\top := \Phi_{tk}(\boldsymbol{\xi}^t)^\top A_t(\boldsymbol{\xi}^t)$ for every $\boldsymbol{\xi}^T \in \Xi$, where $k' = k$,
2. If $t > 1$, then $\Psi_{t-1k'}(\boldsymbol{\xi}^T)^\top := \Phi_{tk}(\boldsymbol{\xi}^t)^\top B_t(\boldsymbol{\xi}^t)$ for every $\boldsymbol{\xi}^T \in \Xi$, where $k' = K_t + k$.

Using these basis functions Ψ_t, let $y^\star(\boldsymbol{\xi}^T)$ be an optimal solution of the restricted NA dual (17) with the optimal value of ν_R^{NA}. The existence of $y^\star(\boldsymbol{\xi}^T)$ is guaranteed due to our initial assumption that the stochastic process is modeled as a finite scenario tree, as well as the boundedness of the feasible region. We construct a feasible solution to the restricted SW dual (16) with the same objective value ν_R^{NA}, which demonstrates the desired inequality. We claim that $\hat{x}_t(\boldsymbol{\xi}^t) = \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t^\star(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]$ satisfies these conditions.

We first show that at $\hat{x}_t(\boldsymbol{\xi}^t)$ the objective function value of the restricted SW dual (16) is equal to ν_R^{NA}. Indeed, the objective evaluates to

$$
\begin{aligned}
\mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]} c_t(\boldsymbol{\xi}^t)^\top \hat{x}_t(\boldsymbol{\xi}^t)\right] &= \mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]} c_t(\boldsymbol{\xi}^t)^\top \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t^\star(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right] \\
&= \sum_{t\in[T]} \mathbb{E}_{\boldsymbol{\xi}^t}\left[\mathbb{E}_{\boldsymbol{\xi}'^T}[c_t(\boldsymbol{\xi}'^t)^\top y_t^\star(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right] && \text{(18a)} \\
&= \sum_{t\in[T]} \mathbb{E}_{\boldsymbol{\xi}'^T}[c_t(\boldsymbol{\xi}'^t)^\top y_t^\star(\boldsymbol{\xi}'^T)] \quad = \quad \nu_R^{\mathrm{NA}} && \text{(18b)}
\end{aligned}
$$

where (18a) follows because for a fixed $\boldsymbol{\xi}^T$, $c_t(\boldsymbol{\xi}^t)$ is a constant that can be brought into the inside conditional expectation, and by swapping the order of expectation, the first equality in (18b) follows from the fact that, for random variables Z and Y, $\mathbb{E}_Y[\mathbb{E}_Z[Z|Y]] = \mathbb{E}[Z]$, and the second equality in (18b) follows because $\boldsymbol{\xi}^T$ and $\boldsymbol{\xi}'^T$ have the same distribution.

We next verify that $\hat{x}_t(\boldsymbol{\xi}^t)$ satisfies constraint (16b). By (17b) and the definition of $Y(\boldsymbol{\xi}'^T)$, $y_t^\star(\boldsymbol{\xi}'^T) \in X_t(\boldsymbol{\xi}^t)$ for each $t \in T$ and $\boldsymbol{\xi}'^T \in \Xi$ such that $\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t$. Since $\hat{x}_t(\boldsymbol{\xi}^t) = \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t^\star(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]$ is a convex combination of such $y_t^\star(\boldsymbol{\xi}'^T)$ variables, it follows that $\hat{x}(\boldsymbol{\xi}^t) \in \mathrm{conv}(X_t(\boldsymbol{\xi}^t))$.

Finally, we verify that $\hat{x}_t(\boldsymbol{\xi}^t)$ satisfies constraint (16c). We evaluate the left-hand side of (16c) for a fixed $t \in [T]$ at the defined solution $\hat{x}$. First break the expectation into three terms

$$\mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top\left(A_t(\boldsymbol{\xi}^t)\hat{x}_t(\boldsymbol{\xi}^t) + B_t(\boldsymbol{\xi}^t)\hat{x}_{t-1}(\boldsymbol{\xi}^{t-1}) - b_t(\boldsymbol{\xi}^t)\right)\right] = \tag{19}$$

$$\mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top A_t(\boldsymbol{\xi}^t)\hat{x}_t(\boldsymbol{\xi}^t)\right] + \mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top B_t(\boldsymbol{\xi}^t)\hat{x}_{t-1}(\boldsymbol{\xi}^{t-1})\right] - \mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}^\top(\boldsymbol{\xi}^t)b_t(\boldsymbol{\xi}^t)\right]. \tag{20}$$

Substituting $\hat{x}_t(\boldsymbol{\xi}^t) = \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t^\star(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]$ in the first term of (20), yields

$$\mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top A_t(\boldsymbol{\xi}^t)\mathbb{E}_{\boldsymbol{\xi}'^T}[y_t^\star(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right].$$

By construction, there exists some $k' \in [K_t']$ with $\Psi_{tk'}(\boldsymbol{\xi}^T)^\top = \Phi_{tk}(\boldsymbol{\xi}^t)^\top A_t(\boldsymbol{\xi}^t)$, thus constraint (17c) leads to the following equality

$$\mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top A_t(\boldsymbol{\xi}^t)\mathbb{E}_{\boldsymbol{\xi}'^T}[y_t^\star(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right] = \mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top A_t(\boldsymbol{\xi}^t)y_t^\star(\boldsymbol{\xi}^T)\right].$$

This follows by substituting the vector $\Phi_{tk}(\boldsymbol{\xi}^t)^\top A_t(\boldsymbol{\xi}^t)$ on the left-hand-side expression with its equivalent vector $\Psi_{tk'}(\boldsymbol{\xi}^T)^\top$ and using (17c) to show the equality of $\mathbb{E}_{\boldsymbol{\xi}^T}[\Psi_{tk'}(\boldsymbol{\xi}^T)^\top \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t^*(\boldsymbol{\xi}'^T)_j|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]] = \mathbb{E}_{\boldsymbol{\xi}^T}[\Psi_{tk'}(\boldsymbol{\xi}^T)^\top(y_t^*(\boldsymbol{\xi}^T))_j]$. By applying the same argument to the second term of (20) while employing constraint (17c) with $t-1$ (when $t = 1$ there is no second term), we can replace $\mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top B_t(\boldsymbol{\xi}^t)\mathbb{E}_{\boldsymbol{\xi}'^T}[y_{t-1}^\star(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^{t-1} = \boldsymbol{\xi}^{t-1}]\right]$ with $\mathbb{E}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top B_t(\boldsymbol{\xi}^t)y_{t-1}^\star(\boldsymbol{\xi}^T)\right]$ since by construction there exists $k' \in K_t'$ with $\Psi_{t-1k'}(\boldsymbol{\xi}^T)^\top = \Phi_{tk}(\boldsymbol{\xi}^t)^\top B_t(\boldsymbol{\xi}^t)$. Putting all three terms of (20) back together, we obtain that the expression (19) is equal to

$$\mathbb{E}_{\boldsymbol{\xi}^T}\left[\Phi_{tk}(\boldsymbol{\xi}^t)^\top\left(A_t(\boldsymbol{\xi}^t)y_t^\star(\boldsymbol{\xi}^T) + B_t(\boldsymbol{\xi}^t)y_{t-1}^\star(\boldsymbol{\xi}^T) - b_t(\boldsymbol{\xi}^t)\right)\right].$$

Since $y_t(\boldsymbol{\xi}^T) \in Y_t(\boldsymbol{\xi}^T)$, this expectation is equal to zero and the constraint (16c) is satisfied. $\square$

Thus, we conclude that if the basis functions are selected carefully, then the restricted NA dual is not worse than the restricted SW dual. Our numerical experiments in Section 5 illustrate that it can indeed provide strictly better bounds.

3.4. Choice of Basis Functions

Basis function selection appears as a design choice in many computational fields, and as such has a rich literature. As mentioned by Powell (2009) in the context of approximate dynamic programming, this choice is problem-specific and requires the knowledge of underlying features that affect the future for a proper design. A natural idea is to define one basis function per uncertain parameter, referred to as the *standard basis functions*. The following proposition shows that, when using the standard basis functions in the NA dual, the basis functions used in stage t can be restricted to use only the uncertain parameters for periods from stage t to the end of the horizon.

PROPOSITION 1. *Let* $\Psi_t(\boldsymbol{\xi}^T) = (1, \boldsymbol{\xi}_2, \boldsymbol{\xi}_3, \ldots, \boldsymbol{\xi}_T)$. *There exists an optimal solution* $\alpha^\star$ *to problem* (10) *where* $\alpha^\star_{tk'} = 0$ *for* $1 < k' \le t$.

Proof For $k' \le t$, $\mathbb{E}_{\boldsymbol{\xi}'^T}[\ \boldsymbol{\xi}'_{k'}\ |\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t] = \boldsymbol{\xi}_{k'}$. Hence the coefficient of $\alpha_{tk'}$ for $1 < k' \le t$ in the objective function of (11) is 0 and we can fix $\alpha^\star_{tk'} = 0$ for $1 < k' \le t$ in an optimal solution. $\square$

The following proposition confirms the intuition that including a basis function that is a linear combination of other basis functions would not change the bound obtained.

PROPOSITION 2. *For* $t \in [T]$, *let* $\Psi^1_t = (\Psi_{t1}, \ldots, \Psi_{tK'_t})$ *be a given set of basis functions and for* $w^t \in \mathbb{R}^{K'_t}$ *define* $\Psi^2_t = (\Psi_{t1}, \ldots, \Psi_{tK'_t}, (w^t)^\top \Psi^1_t)$. *Denoting by* $\nu^{NA}_R(\Psi)$ *the NA dual bound obtained by using basis functions* Ψ, *we have*

$$\nu^{NA}_R(\Psi^1) = \nu^{NA}_R(\Psi^2).$$

Proof In the NA primal characterization, the choice of basis functions affects constraints (17c). For the basis function $(w^t)^\top \Psi^1_t$ we have $\mathbb{E}_{\boldsymbol{\xi}^T}\left[\left((w^t)^\top \Psi^1_t(\boldsymbol{\xi}^T)\right)^\top \left(y_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right)\right] = \sum_{k \in K'_t} w^t_k \mathbb{E}_{\boldsymbol{\xi}^T}\left[\Psi_{tk}(\boldsymbol{\xi}^T)^\top \left(y_t(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right)\right] = 0$, which is a linear combination of the constraints associated with the basis functions included in Ψ^1_t, and thus is redundant. $\square$

If we have some knowledge that a certain function of the uncertain parameters impacts the model, it may be useful to use this function as a basis function. For instance, in the inventory example studied by Bodur and Luedtke (2018), demand at stage t is obtained from the multiplication of random variables observed in the prior stages. Therefore, a basis function at stage t is defined by multiplying those random variables (e.g., as opposed to using the individual random variables as basis functions). For a more complex stochastic process, the literature of approximate dynamic programming, approximate linear programming and machine learning offer more nonlinear choices, e.g., polynomial, piecewise constant/linear (Friedman et al. 2001) and radial basis functions (Buhmann 2000). Rather than fixing the form of the basis functions, some works have opted for design procedures that leverage learning methods such as neighbourhood component analysis (Keller et al. 2006) and neural networks (Hajizadeh and Mahootchi 2018) for constructing basis functions. In Sections 5.3.2 and 5.4.2 we study the impact of different options of basis functions for our two problem classes and show the importance of their proper design on the solution time and quality.

3.5. Solving the Restricted Dual Problems

The restricted dual problems (9) and (10) have the form of maximizing an expected value of a nonsmooth concave function of the decision variables (the coefficients of the LDDR policies). A wide variety of algorithms for (approximately) solving such problems exist, including stochastic approximation based methods (e.g., Robbins and Monro (1951), Nemirovski et al. (2009)), stochastic decomposition (Higle and Sen 1996), and sample average approximation (SAA) (e.g., Shapiro et al. (2009)). For concreteness, we describe an SAA approach. Let $\{\boldsymbol{\xi}_\omega^T\}_{\omega\in\Omega}$ be a given sample of sample-path scenarios and p_ω denote the probability of scenario $\omega \in \Omega$. For example, if the scenarios are generated via Monte Carlo sampling, then $p_\omega = 1/|\Omega|$ for all $\omega \in \Omega$. SAA replaces the expectations in (9) and (10) with sample averages, which respectively give the following SAA models for approximating the restricted SW and NA duals

$$\max_{\theta,\{\beta_t\}_{t\in[T]}} \sum_{\omega\in\Omega} p_\omega \theta_\omega \tag{21a}$$

$$\text{s.t.} \quad \theta_\omega \le \sum_{t\in[T]} \mathcal{L}_t^{\mathrm{SW}}(\beta, \boldsymbol{\xi}_\omega^t), \qquad \omega \in \Omega \tag{21b}$$

and

$$\max_{\eta,\{\alpha_t\}_{t\in[T]}} \sum_{\omega\in\Omega} p_\omega \eta_\omega \tag{22a}$$

$$\text{s.t.} \quad \eta_\omega \le \mathcal{L}^{\mathrm{NA}}(\alpha, \boldsymbol{\xi}_\omega^T), \qquad \omega \in \Omega. \tag{22b}$$

In order to solve (21) and (22), we use a regularized Benders method (Ruszczyński 1986). It is proven that under certain conditions (e.g., the feasible sets in (9) and (10) being nonempty and bounded, and the expectations being finite), as the sample size increases, the solutions to the SAA models converge to the ones of the respective original problems (Shapiro et al. 2009). To assure convergence of the SAA problems, it may be required to put bounds on the LDDR decision variables.

We refer the reader to (Kiwiel 1995, Ruszczyński 1986) for details and implementation strategies for the regularized Benders method. This algorithm uses a master problem that approximates the constraints (21b) and (22b), respectively, with a finite set of Benders cuts. The part of the algorithm that requires specialization for its application to problems (21) and (22) is the specification of the Benders cuts that are added to the master problem given a current master problem solution $\hat{\beta}$ or $\hat{\alpha}$, for problems (21) and (22), respectively.

For problem (21), subproblems evaluating $\mathcal{L}_t^{\mathrm{SW}}(\hat{\beta}, \boldsymbol{\xi}_\omega^t)$ are solved for $\omega \in \Omega$ and $t \in [T]$. Letting $\hat{x}_t^\omega$ denote the optimal solution of subproblem for t and scenario ω, the Benders cut is given below

$$\theta_\omega \le \sum_{t\in[T]} \mathcal{L}_t^{\mathrm{SW}}(\hat{\beta}, \boldsymbol{\xi}_\omega^t) + \sum_{t\in[T]} \sum_{k\in[K_t]} (\beta_{tk} - \hat{\beta}_{tk})\, g_{tk}(\hat{\beta}, \boldsymbol{\xi}_\omega^t)$$

where $g_{tk}(\hat{\beta}, \boldsymbol{\xi}_\omega^t) = -\Phi_{tk}(\boldsymbol{\xi}_\omega^t) b_t(\boldsymbol{\xi}_\omega^t) + \Big(\Phi_{tk}(\boldsymbol{\xi}_\omega^t) A_t(\boldsymbol{\xi}_\omega^t) + B_t(\boldsymbol{\xi}_\omega^t) \mathbb{E}_{\boldsymbol{\xi}'^T}\Big[\ \Phi_{tk}(\boldsymbol{\xi}'^t) \ \ |\ \boldsymbol{\xi}_\omega^{t-1}\Big]\Big)^\top \hat{x}_t^\omega$ is the k^{th} component of the subgradient of $\mathcal{L}_t^{\mathrm{SW}}(\beta, \boldsymbol{\xi}_\omega^t)$ with respect to β, at point $\hat{\beta}$ evaluated at $\boldsymbol{\xi}_\omega$.

For problem (22), subproblems evaluating $\mathcal{L}^{\mathrm{NA}}(\hat{\alpha}, \boldsymbol{\xi}_\omega^T)$ for each $\omega \in \Omega$ are solved. Denoting the subproblem optimal solutions by $\hat{y}^\omega$ for each $\omega \in \Omega$, the Benders optimality cut for the restricted NA Lagrangian dual is the inequality

$$\eta_\omega \le \mathcal{L}^{\mathrm{NA}}(\hat{\alpha}, \boldsymbol{\xi}_\omega^T) + \sum_{t\in[T]} \sum_{k'\in[K_t']} \Big(\Psi_{tk'}(\boldsymbol{\xi}_\omega^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[\Psi_{tk'}(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}_\omega^t]\Big) \Big[\, (\alpha_{tk'} - \hat{\alpha}_{tk'}) \hat{y}_t^\omega \Big].$$

Note that in the presence of stagewise independence or a recursive form such as an autoregressive process, $\mathbb{E}_{\boldsymbol{\xi}'^T}[\Psi_{tk'}(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}_\omega^t]$ can be computed directly. Otherwise conditional expectations need to be approximated, e.g., by sampling.

	Stagewise dual (SW)		Nonanticipative dual (NA)
Relaxation	State equations		Nonanticipativity constraints
Dual variables	π_t		γ_t
LDDR decision variables	β_t		α_t
Basis Functions	Φ_t		Ψ_t
Subproblem per	t and $\boldsymbol{\xi}^t$		$\boldsymbol{\xi}^T$
Optimal value	ν^{SW}	$\leq$	ν^{NA}

Table 1: Summary of the characteristics of the SW and NA duals

Solving (21) and (22) yields candidate LDDR solutions, say $\beta^\star_\Omega$ and $\alpha^\star_\Omega$, respectively. Let $\{\boldsymbol{\xi}^T_\omega\}_{\omega\in\Omega'}$ be an independent evaluation sample with $|\Omega'| >> |\Omega|$. In order to get statistically valid lower bounds on the optimal value of the original MSMIP problem, subproblems $\mathcal{L}^{\text{SW}}_t(\cdot)$ and $\mathcal{L}^{\text{NA}}(\cdot)$ respectively given in (21) and (22), are solved with fixed $\beta^\star_\Omega$ and $\alpha^\star_\Omega$ for every scenario in the evaluation sample, i.e., solving $\mathcal{L}^{\text{SW}}_t(\beta^\star_\Omega, \boldsymbol{\xi}^t_\omega)$ and $\mathcal{L}^{\text{NA}}(\alpha^\star_\Omega, \boldsymbol{\xi}_\omega)$ for all $\omega \in \Omega'$. The lower end of a confidence interval based on the obtained values is a statistically valid lower bound, regardless of how $\beta^\star_\Omega$ and $\alpha^\star_\Omega$ were obtained (e.g., these need not be optimal solutions).

A summary of the notation used in the two proposed restricted Lagrangian duals is given in Table 1.

4. Primal Policies

We next describe how restricted dual solutions can be used to obtain primal policies. We present a classical policy based on solving deterministic approximations, and two policies that use solutions from the restricted duals.

All the primal policies follow a rolling horizon framework: at every stage we make a decision respecting the current stage constraints, then after observing the value of the random parameters, we *roll the horizon*, meaning that the decisions of the current stage are passed as parameters to the next stage and new decisions are made. More formally, let $\{\xi^T_\omega\}_{\omega\in\Omega'}$ be a collection of sample path observations (scenarios) that we will use to evaluate a policy. For each scenario ξ^T_ω, the cost of using this policy for this scenario is computed as $U_\omega = \sum_{t\in[T]} c_t(\xi^t_\omega)^\top \hat{x}_t(\xi^t_\omega)$, where $\hat{x}_t(\xi^t_\omega), t \in [T]$ are the decisions chosen by the policy. In each of the policies we consider, the solution $\hat{x}_t(\xi^t_\omega)$ is obtained by solving an optimization model that takes as input $\hat{x}_{t-1}(\xi_\omega^{t-1})$ (when $t>1$) and the current random parameter observations ξ_ω^t. The optimization model solved at stage t includes, at a minimum, the constraints:

$$\begin{aligned} &A_t(\xi_\omega^t)x_t = b_t(\xi_\omega^t) - B_t(\xi_\omega^t)\hat{x}_{t-1}(\xi_\omega^{t-1}), \\ &x_t \in X_t(\xi_\omega^t). \end{aligned}$$

After solving the optimization problem at stage t, $\hat{x}_t(\xi_\omega^t)$ is set to the x_t component of the optimal solution. These constraints, together with the relatively complete recourse assumption (Assumption 1) ensure that a feasible solution exists at every stage t, and hence each of the presented policies is feasible. A confidence interval of the expected cost of each policy can then be computed using the values in $\{U_\omega\}_{\omega\in\Omega'}$, since each of these values represents an independent observation of the cost of the simulated policy. The upper end of this confidence interval provides a statistical upper bound on the cost of the associated feasible policy, and therefore an upper bound on the optimal value of the MSMIP model.

In the following sections we discuss three primal policies that share the same rolling horizon framework and differ from each other in how the decisions of the current stage are made, i.e., the optimization problems that need to be solved. More specifically, the key difference in these approaches is in how they approximate the impact of the current decisions on future costs.

4.1. Classical Approach

We first review a classical approach which obtains a policy by replacing the uncertain future parameter values with their conditional expected value. For a partial sample path $\boldsymbol{\xi}^t$, denote by $\bar{\xi}^s_{|\boldsymbol{\xi}^t}$ the conditional expected value of ξ^s given $\boldsymbol{\xi}^t$, for $s \geq t$. Within the rolling horizon framework, at stage t, given the solution of the previous stage $\hat{x}_{t-1}(\xi_\omega^{t-1})$, we solve the following *deterministic* problem (initialized at $t=1$ with $\hat{x}_1$ as the vector of zeros):

$$\min \;\; c_t(\xi_\omega^t)^\top x_t + \sum_{s\in[t+1,T]} c_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t})^\top x_s \tag{23a}$$

$$\text{s.t.}\quad A_t(\xi^t_\omega)x_t = b_t(\xi^t_\omega) - B_t(\xi^t_\omega)\hat{x}_{t-1}(\xi^{t-1}_\omega), \tag{23b}$$

$$A_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t})\ x_s + B_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t})\ x_{s-1} = b_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t}), \quad s \in [t+1, T] \tag{23c}$$

$$x_s \in X_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t}), \quad s \in [t, T]. \tag{23d}$$

We refer to the policy obtained from this upper bounding procedure as the *conditional expected value policy*. Note that, for a sample path ξ^T_ω, only the state equations of the current stage are enforced. For future stages, the state equations are only written for $\bar{\xi}^s_{|\boldsymbol{\xi}^t}$ and not almost every scenario. Therefore, $\hat{x}_t(\xi^t_\omega)$ might not satisfy the state equations in the next stage for ξ^{t+1}_ω. In the next section, we take advantage of our LDDR solutions with the hope to identify solutions that provide lower cost over the full horizon.

4.2. Restricted Stagewise Lagrangian Dual Driven Policy

To improve the conditional expected value policy, an idea is to add a penalty obtained from the SW Lagrangian dual for violating the state equations at the next immediate stage for ξ^{t+1}_ω. Given a restricted SW dual solution $\hat{\pi}_t(\boldsymbol{\xi}^t) = \Phi_t(\boldsymbol{\xi}^t)\ \hat{\beta}_t$, $t \in [T]$, we obtain the *SW dual driven policy* by following the same rolling horizon procedure above with the objective function as the convex combination of (23a) and the coefficient of x_t decision variables in $\mathcal{L}^{\text{SW}}_t$:

$$\begin{aligned}
\min\ & \Big(c_t(\xi^t_\omega) + \lambda\ \big(\hat{\beta}^\top_t \Phi^\top_t(\xi^t_\omega)\ A_t(\xi^t_\omega) + \hat{\beta}^\top_{t+1}\ \mathbb{E}_{\boldsymbol{\xi}'^T}\big[\Phi^\top_{t+1}(\boldsymbol{\xi}'^{t+1})\ B_{t+1}(\boldsymbol{\xi}'^{t+1}) \mid \xi^t_\omega\big]\big)\Big)^\top x_t \qquad (24a)\\
& + (1-\lambda) \sum_{s\in[t+1,T]} c_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t})^\top x_s\\
\text{s.t.}\ & A_t(\xi^t_\omega)x_t = b_t(\xi^t_\omega) - B_t(\xi^t_\omega)\hat{x}_{t-1}(\xi^{t-1}_\omega),\\
& A_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t})\ x_s + B_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t})\ x_{s-1} = b_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t}), \qquad s \in [t+1, T]\\
& x_s \in X_s(\bar{\xi}^s_{|\boldsymbol{\xi}^t}), \qquad s \in [t, T].
\end{aligned}$$

where $\lambda \in [0, 1]$ is the weight of information from the SW Lagrangian dual and a parameter of the policy. Setting $\lambda = 0$ leads to the conditional expected value policy, while $\lambda = 1$ corresponds to a modified myopic policy where the only information about the future that we pass to the primal algorithm is the penalty associated with the next immediate stage.

4.3. Restricted Nonanticipative Lagrangian Dual Driven Policy

Given a restricted NA dual policy $\hat{\gamma}_t(\boldsymbol{\xi}^T) = \Psi_t(\boldsymbol{\xi}^T)\hat{\alpha}_t$, for all $t \in [T]$, we define the *NA dual-driven policy* as follows. Using the discussed rolling horizon framework, at stage t, with observed history ξ^t_ω and previous stage decisions $\hat{x}_{t-1}(\xi^{t-1}_\omega)$, choose $x_t(\xi^t_\omega)$ as an approximate solution of the following *two-stage stochastic* program

$$\min_{x_t} \quad c_t(\xi^t_\omega)^\top x_t + \mathbb{E}_{\boldsymbol{\xi}^T}\left[g_t(x_t, \boldsymbol{\xi}^T) \mid \xi^t_\omega\right] \tag{25a}$$

$$\text{s.t.} \quad A_t(\xi^t_\omega)x_t = b_t(\xi^t_\omega) - B_t(\xi^t_\omega)\hat{x}_{t-1} \tag{25b}$$

$$x_t \in X_t(\xi^t_\omega) \tag{25c}$$

where

$$g_t(x_t, \boldsymbol{\xi}^T) = \min_{\{x_s\}_{s\in[t+1,T]}} \quad \sum_{s\in[t+1,T]} \Big(c_s(\boldsymbol{\xi}^s) + \hat{\gamma}_s(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[\hat{\gamma}_s(\boldsymbol{\xi}'^T) \mid \boldsymbol{\xi}^t]\Big)^\top x_s \tag{26a}$$

$$\text{s.t.} \quad A_s(\boldsymbol{\xi}^s)x_s + B_s(\boldsymbol{\xi}^{s-1})x_{s-1} = b_s(\boldsymbol{\xi}^s), \qquad s \in [t+1, T] \tag{26b}$$

$$x_s \in X_s(\boldsymbol{\xi}^s), \qquad s \in [t+1, T]. \tag{26c}$$

In order to solve the problem (25), for each realization ξ^t_ω, a sample $\{\xi_{|t,\omega}\}_{\omega\in\Omega_{|t}}$ is generated, where $\boldsymbol{\xi}_{|t}$ is a random variable representing the (conditional) scenarios after the history up to stage t has been observed. Then we solve it by replacing the expectation in (25a) with a sample average using this sample. Since this scheme is used in a rolling horizon fashion, in order to estimate the expected cost of this policy, multiple sample paths ξ^T are generated, and for each one a cost is recorded by the scheme outlined above. Finally, a confidence interval on upper bound value is obtained over these costs.

As the second stage problem (26) is a MIP, given a set of conditional scenarios for the future, we can solve the extensive form of this two-stage problem, by creating a copy of the second-stage variables in problem (26) for each conditional scenario and embedding them in (25) for all scenarios. Note that although two-stage stochastic MIPs are computationally difficult to solve, there has been significant research into methods for solving this problem class. This includes the integer L-shaped method (Laporte and Louveaux 1993) and a large literature on methods for generating improved cuts within such a solution approach, e.g., (Angulo et al. 2016, Bodur and Luedtke 2017, Qi and Sen 2017, Tanner and Ntaimo 2008, Gade et al. 2014, Sen and Higle 2005, Sherali and Zhu 2007, Sen and Sherali 2006, Zhang and Küçükyavuz 2014, Rahmaniani et al. 2020). Another general approach is dual decomposition (Carøe and Schultz 1999), see (Kim et al. 2019) for extensions and a recent parallel implementation. See (Küçükyavuz and Sen 2017) for a more extensive overview of exact and approximate solution methods for two-stage stochastic MIPs.

Although two-stage stochastic MIP problems are challenging, when implementing this policy it needs to only be solved once per time stage. On the other hand, estimating the expected value of this policy is computationally challenging since it is required to simulate many sample paths to construct a confidence interval. Thus, it may be necessary to heuristically limit the effort in solving the two-stage stochastic MIP, e.g., by using a small number of scenarios, or by terminating the solution process after the solver processes just a limited number of branch-and-bound nodes. Another approach is to generate a larger number of scenarios and then representing them by a smaller size sample, e.g., using a scenario reduction technique like clustering.

The overall framework for providing bounds on MSMIP is illustrated in Figure 2. Using LDDRs, one can obtain a Lagrangian dual policy. The parameters of this policy are then passed to an evaluator to provide confidence intervals (CIs) on the out of sample scenarios. When comparing upper bounds obtained from different methods, the same set of scenarios is used in their evaluation.

5. Computational Experiments

We illustrate our proposed approach on two problems: (*i*) multi-item lot-sizing problem with backlogging and production lag (MSLot), and (*ii*) routing and wavelength assignment problem (MSRWA). We present our analysis on the choice of basis functions, as well as comparison between the bounds of various proposed methods.

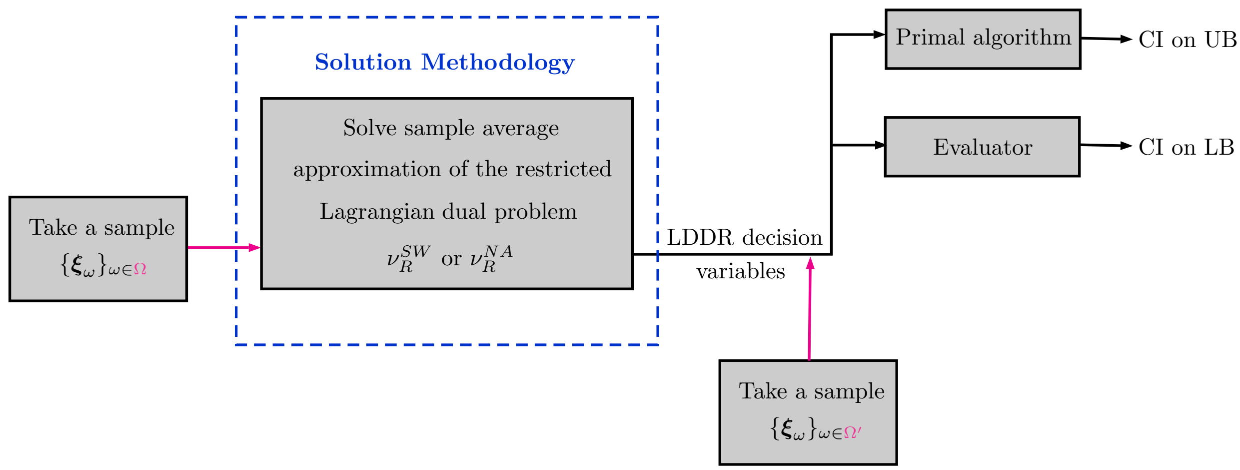

Figure 2: Solution framework

5.1. Implementation Details

To solve the SAA problems (21) and (22) for the MSLot and MSRWA, we have used the regularized Benders method, with all the parameters of the algorithm being exactly as provided by Lubin et al. (2013). The convergence tolerance of the algorithm is set to 0.001. The algorithms are implemented in C++ using IBM ILOG CPLEX 12.8 for solving MIPs, and the experiments are conducted on a MacOS X with 2.3 GHz Intel Core i5 and 16 GB RAM. For all bounding methods we have considered a maximum running time of 6 hours.

For the SW dual driven policy, in the objective function (24a), we use $\lambda = 0.25$. This value is selected based on some preliminary analysis on the lotsizing problem (whose results are provided in Table C.4 of the e-companion). For the NA dual-driven primal policy, the two-stage stochastic MIP problem (25) is approximately solved using a sample with $|\Omega_{|t}| = 25$. This set is generated by randomly sampling 100 scenarios, then clustering these into 24 groups and using the mean of each group as a scenario in the sample, and adding the conditional expectation $\bar{\xi}^s_{|\boldsymbol{\xi}^t}$ as the last scenario in the sample.

5.2. Benchmarks

The same set of scenarios is used for estimating lower bounds and evaluating primal policies. The perfect information (PI) bound is used as our benchmark for lower bounds, which is obtained similar to the restricted NA lower bound, with the LDDR variables, α, fixed to 0. This is the bound obtained by solving the version of the problem where the complete information about the future is assumed to be known. We use the conditional expected value policy as the benchmark upper bounding policy. Since this formulation is not based on a scenario-tree model (and approximating it with a scenario tree would require an intractable number of nodes), to the best of our knowledge, these are the best available comparison methods for the general MSMIP problem class.

5.3. Multi-item Stochastic Lot-sizing Problem

The MSMIP formulation for the MSLot problem is as follows

$$\min\ \mathbb{E}\left[\sum_{t\in[T]}\Big(\sum_{j\in[J]}\Big(C^{i^+}_{tj}(\boldsymbol{\xi}^t)i^+_{tj}(\boldsymbol{\xi}^t)+C^{i^-}_{tj}(\boldsymbol{\xi}^t)i^-_{tj}(\boldsymbol{\xi}^t)+C^{y}_{tj}(\boldsymbol{\xi}^t)y_{tj}(\boldsymbol{\xi}^t)\Big)+C^{o}_{t}(\boldsymbol{\xi}^t)o_t(\boldsymbol{\xi}^t)\Big)\right] \tag{27a}$$

$$\text{s.t. } i^-_{tj}(\boldsymbol{\xi}^t)-i^+_{tj}(\boldsymbol{\xi}^t)+i^+_{t-1,j}(\boldsymbol{\xi}^{t-1})-i^-_{t-1,j}(\boldsymbol{\xi}^{t-1})+x_{t-1,j}(\boldsymbol{\xi}^{t-1})=D_{tj}(\boldsymbol{\xi}^t), \quad t\in[T], j\in[J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{27b}$$

$$\sum_{j\in[J]}(TS_j y_{tj}(\boldsymbol{\xi}^t)+TB_j x_{tj}(\boldsymbol{\xi}^t))-o_t(\boldsymbol{\xi}^t)\le C_t, \quad t\in[T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{27c}$$

$$M_{tj}\, y_{tj}(\boldsymbol{\xi}^t)-x_{tj}(\boldsymbol{\xi}^t)\ge 0, \quad t\in[T], j\in[J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{27d}$$

$$i^+_{tj}(\boldsymbol{\xi}^t)\le I_{tj}, \quad t\in[T], j\in[J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{27e}$$

$$i^+_{tj}(\boldsymbol{\xi}^t)+x_{tj}(\boldsymbol{\xi}^t)\le I_{t+1,j}, \quad t\in[T], j\in[J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{27f}$$

$$0\le o_t(\boldsymbol{\xi}^t)\le O_t, \quad t\in[T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{27g}$$

$$x_{tj}(\boldsymbol{\xi}^t), i^+_{tj}(\boldsymbol{\xi}^t), i^-_{tj}(\boldsymbol{\xi}^t)\ge 0, \quad t\in[T], j\in[J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{27h}$$

$$y_{tj}(\boldsymbol{\xi}^t)\in\{0,1\}, \quad t\in[T], j\in[J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{27i}$$

where $x_{tj}(\boldsymbol{\xi}^t), i^+_{tj}(\boldsymbol{\xi}^t), i^-_{tj}(\boldsymbol{\xi}^t)$ are decision variables representing production level, inventory and backlog of product type j at stage t, respectively. Binary decision variable $y_{tj}(\boldsymbol{\xi}^t)$ is equal to 1 if production of item j is setup at stage t, 0 otherwise. Decision variable $o_t(\boldsymbol{\xi}^t)$ measures the overtime at stage t. C_t, I_{tj}, O_t are the production capacity, inventory capacity of product j, and the overtime bound at stage t, respectively. $D_{tj}(\boldsymbol{\xi}^t)$ is the demand of product j at stage t. $C^{i^+}_{tj}(\boldsymbol{\xi}^t), C^{i^-}_{tj}(\boldsymbol{\xi}^t), C^{y}_{tj}(\boldsymbol{\xi}^t)$

are respectively the costs of holding, backlog and fixed setup for product j, and $C_t^o(\boldsymbol{\xi}^t)$ is the overtime cost at stage t. For product j, TS_j is the setup time, while TB_j is the production time per unit. Constraints (27b) are the state equations, linking the inventory, backlog and production of consecutive stages. Note that there is a production lag of 1, meaning that the amount that is produced at stage t is not available until stage $t+1$. Overtime is measured by constraints (27c), while the same set of constraints ensure the production capacity is respected. Constraints (27d) link the production and setup decisions, using sufficiently large big-M values in the absence of a given limit on the production quantity. The rest of the constraints determine the bounds and integrality constraints on the decision variables. The objective function (27a) is the total expected cost, including the costs of holding, backlog, setup and overtime.

5.3.1. Base Data The dataset of our experiments is loosely based on the work of Helber et al. (2013). We consider the following autoregressive process model for representing the correlation between demands of different stages

$$Y_{t+1,j} = \rho\, Y_{tj} + (1-\rho)\, \epsilon_{t+1,j},$$

where ϵ_{tj} is a lognormal random variable with mean 1 and standard deviation 0.5. Demands are modeled as

$$D_{tj} = \rho^Y Y_{tj}\, \mu_{tj} + (1-\rho^Y)\, \delta_{tj},$$

where μ_{tj} is the mean of the demands of product type j at stage t, and δ_{tj} is a random variable having a mean of μ_{tj}. Thus, in this model there is an underlying autoregressive process (Y_{tj}) and the demand in each period is partially driven by this and also by an external random variable. This assures that demands are nonnegative. We consider a lognormal distribution for δ_{tj} whose standard deviation is $0.2\, t\, \mu_{tj}$, reflecting higher demand uncertainty further in the future. For μ_{tj} values, we use the means of the demands provided by Helber et al. (2013). Using consecutive substitutions, we can deduce that the following holds for the conditional expectation

$$\mathbb{E}[D_{t+h,j}|\epsilon_{tj}, \delta_{tj}] = \mu_{t+h,j} \left(\rho^Y \rho^h \left(\frac{D_{tj} - (1-\rho^Y)\, \delta_{tj}}{\rho^Y\, \mu_{tj}} - 1 \right) + 1 \right).$$

We have generated instances with $T=2,\ldots,10$ stages, and $J=3,6$ product types. There are two sets of data instances for $J=3$ with $\rho=0.2,\rho^Y=0.6$ and $\rho=0.6,\rho^Y=0.2$ leading to different levels of correlation and variation among demands. The combination $\rho=0.6,\rho^Y=0.2$ is additionally used in the generation of instances with $J=6$. For solving the SW and NA dual problems, $\lceil\frac{50}{T}\rceil|\beta|_\#$ and $\lceil\frac{100}{T}\rceil|\alpha|_\#$ scenarios are generated respectively, and the solutions are evaluated using $\lceil\frac{250}{T}\rceil|\alpha|_\#$ scenarios, where $|\cdot|_\#$ operator denotes the total number of variables of a given type (which depends on the choice of basis functions, discussed in Section 5.3.2). The rest of the configurations can be found in the e-companion.

5.3.2. Basis Function Selection As mentioned in Section 3.4, the choice of basis functions is not restricted by any predetermined form. Therefore, in choosing the basis functions for the MSlot problem we have options. Assuming a linear form, an exhaustive choice is to use all the uncertain information we have, i.e., all demand observations. At stage $\hat{t}$, for the SW dual this means using D_{tj} for all j and $t\le\hat{t}$, while in the NA dual it involves the complete set of demands for all stages in the planning horizon and all products. Other reasonable options consist of various subsets of these full sets of basis functions. Table 2 summarizes some of the alternatives considered in our experiments. Note that, Option 1 in the NA dual does not span over the entire planning horizon and it ignores the observed history. According to Proposition 1, these two are equivalent.

Option	SW dual	NA dual
1	$1\,;D_{tj},\ \forall j,\quad \forall t\le\hat{t}$	$1\,;D_{tj},\ \forall j,\quad \forall t\le T:\hat{t}<t$
2	$1\,;D_{tj},\ \forall j,\quad t=\hat{t}$	$1\,;D_{tj},\ \forall j,\quad t=\hat{t}+1$
3	$1\,;D_{tj},\ j=\hat{j},\ \forall t\le\hat{t}$	$1\,;D_{tj},\ j=\hat{j},\ \forall t\le T:\hat{t}<t$
4	$1\,;D_{tj},\ j=\hat{j},\ t=\hat{t}$	$1\,;D_{tj},\ j=\hat{j},\ t=\hat{t}+1$

Table 2: Various options to be used in basis functions, at stage $\hat{t}$ for product $\hat{j}$

In Table 2, Options 2 and 4 in the SW dual only use the stage $\hat{t}$ demands as basis functions for stage $\hat{t}$ constraints, and in the case of the NA dual, they only use the stage $\hat{t}+1$ demands for basis functions used to relax stage $\hat{t}$ constraints. Options 3 and 4 only use product $\hat{j}$ demands as basis functions for constraints associated with product $\hat{j}$. Thus, Option 1 has the most basis functions, and hence should have the best bound, whereas Option 4 uses the fewest basis functions. We compare these four options to determine which one gives the best trade-off between the quality of the bound and the computational effort. For instances with $J=3, \rho=0.6, \rho^Y=0.2$, Tables 3 and 4 present the solution time in seconds and the lower bounds returned by the SW and NA duals, respectively. For this comparison, three instances with $T=4,6,8$ are solved using the four basis function options and only penalizing the NA constraints on x variables (see Section 5.3.3). The bounds are the means of the confidence intervals over the objective values for all the scenarios in the evaluation sample. They are scaled such that 100 is the best known bound obtained for that instance. Note that the best bounds for these instances are obtained by an option of the NA dual reported in Table 5, which is why all the scaled bounds in Tables 3 and 4 are less than 100.

	# Basis functions			Time (s)			Bound (scaled)		
Option	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$
1	315	750	1365	353.8	1412.5	6273.9	60.7	57.5	56.6
2	180	300	420	91.7	243.7	632.5	57.9	55.4	53.8
3	135	300	525	109.0	288.7	795.9	58.1	55.9	54.1
4	90	150	210	157.2	795.9	323.6	54.7	54.2	46.1

Table 3: Basis function selection for the SW dual

In Table 3, Option 1 has the largest number of basis functions and LDDR variables, and the highest lower bound among all the four options. We also carried out a pairwise comparison between Option 1 and the others by performing a t-test. The test confirmed that the difference in the means is statistically significant (at 95% confidence). The higher quality of the bound using Option 1 comes at the price of a larger solution time. If the computation becomes cumbersome, the next candidates are Options 2 and 3 which do not have a significant difference in their means.

The results for the NA dual are presented in Table 4. In this case we find that Options 1 and 3, which use all future time periods rather than just the next time period, have clearly better quality of

Option	# Basis functions			Time (s)			Bound (scaled)		
	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$
1	756	1800	3276	76.9	403.2	1800.8	99.8	99.9	99.8
2	432	720	1008	44.9	232.3	770.5	93.4	94.2	93.1
3	324	720	1260	61.1	250.0	776.8	99.8	99.3	99.5
4	216	360	504	34.2	219.5	735.3	93.3	93.6	93.1

Table 4: Basis function selection for the NA dual

the bound. We conducted a pairwise t-test to determine if the differences are statistically significant. In particular, we compared all the options against Option 3, which seems to return a very high quality solution in a reasonable amount of time. The test revealed that there is no statistically distinguishable difference between the bounds obtained by Option 1 and Option 3. Option 3 is statistically significantly better than Options 2 and 4. This shows that having information about the whole planning horizon is beneficial, rather than just considering the next stage ahead, but for the NA dual there was no observed benefit to using demands for products different from the one being relaxed in the NA constraint in the basis functions. Note that these conclusions depend on the specific instances we are solving and might differ for another set of instances. In the light of the above discussions, in the rest of the experiments, Options 1 and 3 are used for the SW and NA duals, respectively.

5.3.3. NA Variable Selection In NA problem reformulation, another possibility to reduce the restricted dual problem size is to ignore altogether the NA constraints (4c) on some sets of variables (i.e., by fixing the dual variables associated with these constraints to zero). For example, in order to obtain a valid formulation it is sufficient to enforce (4c) only on *state variables*, i.e., variables that appear in more than one decision stage, and thus one may choose to only penalize the NA constraints associated with those variables. On the other hand, in the MSLot problem, once the recourse variables $x_t(\boldsymbol{\xi}^t)$ are determined, the optimal values of the state variables $i^-(\boldsymbol{\xi}^t)$ and $i^+(\boldsymbol{\xi}^t)$ are immediate from the equations (27b), suggesting it may be beneficial to penalize violation of the NA constraints on these variables. (A similar approach is used by Lulli and Sen

	Time (s)			Bound		
Option	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$
x	61.1	250.0	776.8	99.8	99.3	99.5
i^+, i^-	68.1	188.0	653.1	92.5	91.0	90.4
x, i^+, i^-	175.4	849.9	1887.3	100.0	100.0	100.0

Table 5: NA variable selection

(2004).) A natural question for our LDDR approach is, can enforcing nonanticipativity additionally on the other variables (which is redundant before the relaxation) improve the bound?

Table 5 examines this question by considering three options for instances with $J=3, \rho=0.6, \rho^Y=0.2$: having either x or state variables i^+, i^- as the NA variables, or considering a combination of them. Results of Table 5 show that the combination of x, i^+, i^- obtains the best bounds in a longer running time. The difference between the means of the options was tested using a t-test. We find that the bounds obtained by penalizing the NA constraints on only the x variables are not significantly worse than those obtained by penalizing the NA constraints on all variables. In addition, including a penalty of the NA constraints x variables leads to an improvement over just including a penalty of the NA constraints on the state variables i^+ and i^-. Since penalizing only the NA constraints on the x variables yields lower bounds that are indistinguishable from the best, in less time, we use this option in the remaining experiments.

5.3.4. Bound Comparison We now compare the quality of the bounds returned by the two Lagrangian duals, and their respective upper bounds, over a variety of test instances. Figures 3 and 4 illustrate the performance of the lower and upper bounding algorithms for instances with different stages, product types, and levels of stochasticity in demands. The reported numbers are normalized as follows. The lower bounds obtained for each instance are divided by the PI lower bound for that instance, so that the reported value is the lower bound relative to the PI lower bound and values over 1.0 indicate an improvement in the lower bound. The upper bounds are divided by the bound returned by the conditional expected value policy, so that a value below 1.0

indicates an improvement in the upper bound. For both cases the actual numbers are available in the e-companion.

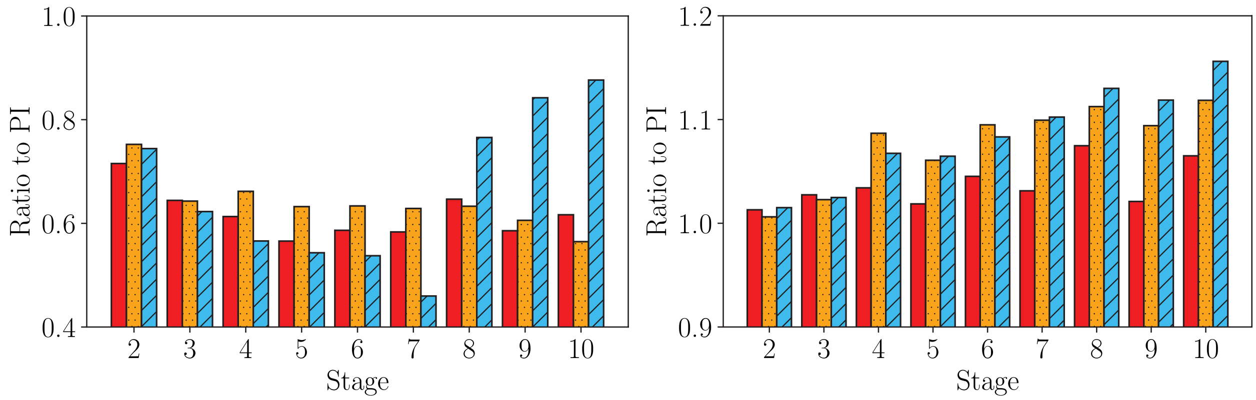

(a) SW lower bound

(b) NA lower bound

Figure 3: Performance of lower bounding algorithms with respect to each other and the PI. Solid bars represent instances with $J=3, \rho=0.2, \rho^Y=0.6$, dotted bars are instances with $J=3, \rho=0.6, \rho^Y=0.2$, and in dashed bars we have $J=6, \rho=0.6, \rho^Y=0.2$.

As shown by Figure 3b, the NA dual has been able to improve upon the PI lower bound in all instances. This improvement generally becomes more evident as the number of stages grows. From Figure 3b, it can be seen that the NA dual is performing better with a higher level of stochasticity (less correlation in the autoregressive process and a higher weight for variation). This could be due to a stronger PI bound when the system is less volatile. On the other hand, as seen in Figure 3a, the SW dual provides significantly lower bounds (as expected from Theorem 1) than those from the NA dual, and even is significantly worse than the PI bound on these test instances.

Figure 4 presents the results comparing the upper bounds obtained from the dual-driven policies to that from the conditional expected value policy. We see from Figure 4a that the SW dual-driven policy provides modest but consistent improvement over the conditional expected value policy. For the NA dual-driven policy we report results only for stages $T \leq 6$, since estimating the value of the policy was too time-consuming for larger instances and takes longer than our

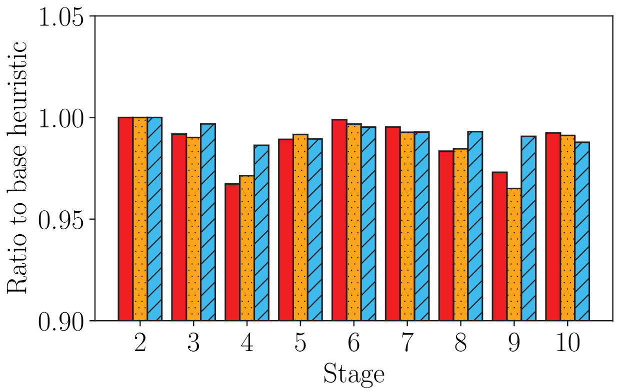

(a) SW upper bound

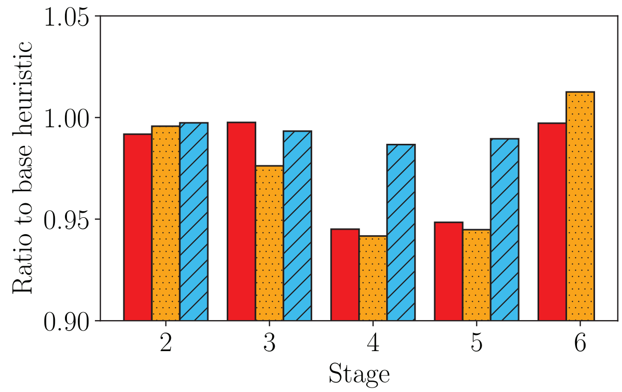

(b) NA upper bound

Figure 4: Performance of upper bounding algorithms with respect to each other and the conditional expected value policy. Solid bars represent instances with $J = 3, \rho = 0.2, \rho^Y = 0.6$, dotted bars are instances with $J = 3, \rho = 0.6, \rho^Y = 0.2$, and in dashed bars we have $J = 6, \rho = 0.6, \rho^Y = 0.2$.

maximum running time of 6 hours. It should be noted that this is caused by the very large number of scenarios in the evaluation sample, and on average obtaining solutions for one sample path does not take more than 45 seconds in the largest instances. The solution times of our lower and upper bounding techniques can be found in the e-companion. We find that for instances with $J = 3$ and $3 \leq T \leq 5$ the NA dual-driven policy was able to significantly improve upon the conditional expected value policy. We suspect that for instances with larger T or J, the lack of improvement is due to a poor approximation of the two-stage stochastic MIP (25) that is solved in this method, as we approximated this problem with just 25 scenarios considering the computational limitations. Although we do not explore this here, the NA dual-driven policy may be more practical if a decomposition algorithm is used to solve the problem (25), enabling the use of many more scenarios to approximate it.

Figure 5a presents the gap between the PI lower bound and the conditional expected value upper bound across various instances, which ranges from 3.91% to 50.67%. In Figure 5b we present the fraction of this gap that is closed using the combined improvements from the NA dual lower bound and the SW dual upper bound. We find that the gap is reduced over all instances, and is more

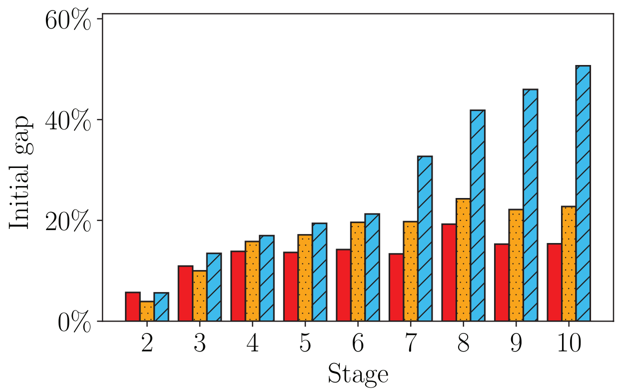

(a) Gap between PI lower bound and conditional expected value upper bound for all instances

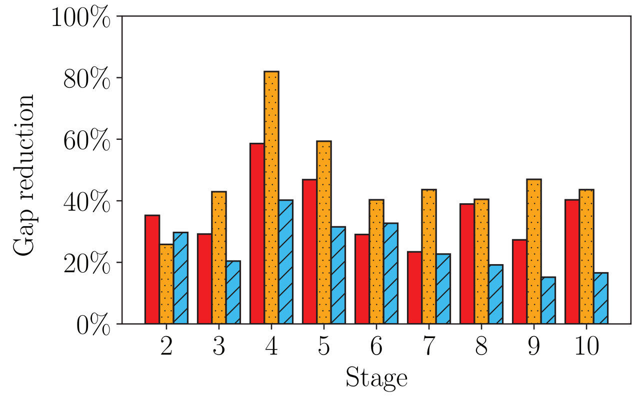

(b) Gap reduction obtained from combining the improvements of lower and upper bounding techniques

Figure 5: Gap reduction percentage with respect to the gap between the PI and the conditional expected value policy. Solid bars represent instances with $J = 3, \rho = 0.2, \rho^Y = 0.6$, dotted bars are instances with $J = 3, \rho = 0.6, \rho^Y = 0.2$, and in dashed bars we have $J = 6, \rho = 0.6, \rho^Y = 0.2$.

pronounced on instances with $J = 3$ and instances with $\rho = 0.6, \rho^Y = 0.2$ (the instances in which the demands have higher variability).

5.4. Stochastic Routing and Wavelength Assignment Problem

In a common optical telecommunications network, links are optical fibers that can simultaneously transmit signals on different wavelengths. The routing and wavelength assignment (RWA) problem is the classical resource allocation problem for such networks. Given a network, a set of wavelengths, and a set of connection requests (origin and destination node pairs), the RWA problem assigns a path and a wavelength to each request, with the objective of maximizing the number of granted requests. This pair of path and wavelink is called a *lightpath*. The building block of lightpaths is a *wavelink*, the pair of a link and a wavelength. Wavelinks on a valid lightpath all have the same wavelength (*wavelength continuity*). To avoid *wavelength conflict*, no two lightpaths can share a wavelink. Figure 6 presents a bidirectional optical network with three wavelengths (represented by different colors). While the dotted red and wavy blue lightpaths are valid, the green dashed lightpaths are invalid as they share the wavelink from node 1 to 3.

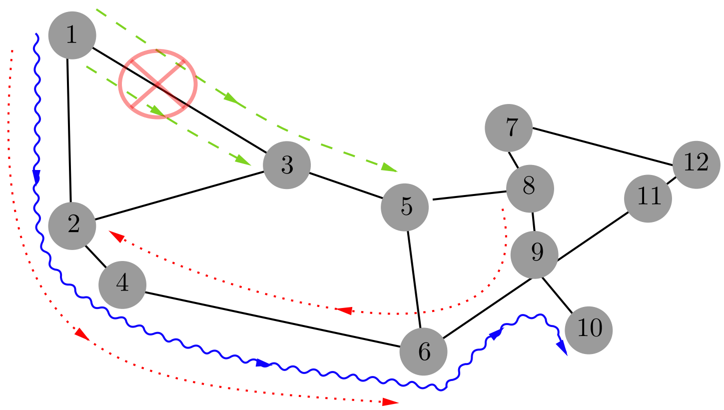

Figure 6: An optical network with valid and invalid lightpaths

Due to the scale of the problems in telecommunications, most solution methods are heuristics, an overview of which is compiled by Simmons and Rouskas (2020), with scalable exact algorithms only recently proposed by Jaumard and Daryalal (2017) for deterministic traffic. For stochastic incoming traffic, Daryalal and Bodur (2020) proposed a two-stage stochastic MIP model and solved it via a decomposition framework. In reality, the set of connection requests arrives randomly over time, suggesting that a multistage model might better capture the dynamic process.

Consider an optical network with a given physical topology represented by a multigraph $G = (\mathcal{V}, \mathcal{L})$ and a set of available wavelengths Ω indexed by ω. We denote by $\mathcal{L}(\omega) \subseteq \mathcal{L}$ the set of free fiber links on wavelength ω. We assume requests remain on the network throughout the planning horizon, and let $r^{t,(s,d)}(\boldsymbol{\xi}^t)$ denote a random parameter representing the number of requested lightpaths between node pair (s, d) at stage t. At stage t, we let $x^{t,(s,d)}_{\ell,\omega}(\boldsymbol{\xi}^t)$ be a binary (recourse) decision variable equal to 1 if wavelink (ℓ, ω) is used for a new connection request between $(s, d) \in \mathcal{SD}^t(\boldsymbol{\xi}^t) = \{(s, d) \in \mathcal{V} \times \mathcal{V} : r^{t,(s,d)}(\boldsymbol{\xi}^t) > 0\}$ and equal to 0 otherwise, and let $\mu^t_{\ell,\omega}(\boldsymbol{\xi}^t)$ be a binary (state) variable equal to 1 if wavelink (ℓ, ω) is serving any node pair at stage t (based on a current or previous accepted connection). Finally, $z^{t,(s,d)}(\boldsymbol{\xi}^t)$ is a (recourse) decision variable measuring the number of granted requests between node pair (s,d) during stage t. With these decision variables, the MSMIP model for multistage RWA (MSRWA) is:

$$\max \ \mathbb{E}_{\boldsymbol{\xi}^T}\left[\sum_{t\in[T]}\sum_{(s,d)\in\mathcal{SD}^t(\boldsymbol{\xi}^t)} z^{t,(s,d)}(\boldsymbol{\xi}^t)\right] \tag{28a}$$

$$\text{s.t.} \quad \sum_{(s,d)\in\mathcal{SD}^1} x_{\ell,\omega}^{1,(s,d)}(\boldsymbol{\xi}^1) = \mu_{\ell,\omega}^1(\boldsymbol{\xi}^1) \qquad \omega\in\Omega, \ell\in\mathcal{L}(\omega) \tag{28b}$$

$$\mu_{\ell,\omega}^t(\boldsymbol{\xi}^t) = \mu_{\ell,\omega}^{t-1}(\boldsymbol{\xi}^{t-1}) + \sum_{(s,d)\in\mathcal{SD}^t(\boldsymbol{\xi}^t)} x_{\ell,\omega}^{t,(s,d)}(\boldsymbol{\xi}^t) \qquad t\in[2,T], \omega\in\Omega, \ell\in\mathcal{L}(\omega), \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{28c}$$

$$\sum_{\substack{\ell\in\\ \delta^+(v)\cap\mathcal{L}(\omega)}} x_{\ell,\omega}^{t,(s,d)}(\boldsymbol{\xi}^t) = \sum_{\substack{\ell\in\\ \delta^-(v)\cap\mathcal{L}(\omega)}} x_{\ell,\omega}^{t,(s,d)}(\boldsymbol{\xi}^t) \qquad t\in[T], \omega\in\Omega, (s,d)\in\mathcal{SD}^t(\boldsymbol{\xi}^t), \; v\in\mathcal{V}\setminus\{s,d\}, \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{28d}$$

$$\sum_{\omega\in\Omega}\sum_{\substack{\ell\in\\ \delta^-(s)\cap\mathcal{L}(\omega)}} x_{\ell,\omega}^{t,(s,d)}(\boldsymbol{\xi}^t) = 0 \qquad t\in[T], (s,d)\in\mathcal{SD}^t(\boldsymbol{\xi}^t), \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{28e}$$

$$\sum_{\omega\in\Omega}\sum_{\substack{\ell\in\\ \delta^+(d)\cap\mathcal{L}(\omega)}} x_{\ell,\omega}^{t,(s,d)}(\boldsymbol{\xi}^t) = 0 \qquad t\in[T], (s,d)\in\mathcal{SD}^t(\boldsymbol{\xi}^t), \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{28f}$$

$$z^{t,(s,d)}(\boldsymbol{\xi}^t) = \sum_{\omega\in\Omega}\sum_{\substack{\ell\in\\ \delta^+(s)\cap\mathcal{L}(\omega)}} x_{\ell,\omega}^{t,(s,d)}(\boldsymbol{\xi}^t) \qquad t\in[T], (s,d)\in\mathcal{SD}^t(\boldsymbol{\xi}^t), \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{28g}$$

$$z^{t,(s,d)}(\boldsymbol{\xi}^t) \le r^{t,(s,d)}(\boldsymbol{\xi}^t) \qquad t\in[T], (s,d)\in\mathcal{SD}^t(\boldsymbol{\xi}^t), \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{28h}$$

$$x_{\ell,\omega}^{t,(s,d)}(\boldsymbol{\xi}^t) \in \{0,1\} \qquad t\in[T], (s,d)\in\mathcal{SD}^t(\boldsymbol{\xi}^t), \omega\in\Omega, \ell\in\mathcal{L}(\omega), \; \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{28i}$$

$$\mu_{\ell,\omega}^t(\boldsymbol{\xi}^t) \in \{0,1\} \qquad t\in[T], \omega\in\Omega, \ell\in\mathcal{L}(\omega), \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T\in\Xi \tag{28j}$$

where $\delta^{+(-)} \subseteq \mathcal{L}$ is the set of outgoing (incoming) links of node $v\in\mathcal{V}$. The objective in (28a) maximizes the expected number of granted requests. Constraints (28b) and (28c) are state equations that keep track of whether or not each wavelink has been used by stage t, which together with the binary restrictions on the state variables prevent a wavelink from being used more than once in the planning horizon. Constraints (28d) -(28g) ensure that a valid lightpath is created for each accepted connection. Constraints (28h) bound the number of lightpath assignments to a node pair by its demand. Constraints (28i)-(28j) are the integrality and bound restrictions.

5.4.1. Data and Basis Functions For the underlying topology of our instances, we have considered ABILENE, a standard long-haul network in the literature of the RWA problem (Orlowski et al. 2010), with $|\Omega| \in \{1,3\}$. The total number of request arrivals in the incoming traffic R_t follows an integer-valued GARCH process (INGARCH, Ferland et al. (2006)), where the deviates have a Poisson distribution with a parameter λ_t (Chen et al. 2015, Xiong et al. 2018) that depends on R_{t-1} and λ_{t-1}:

$$\begin{cases} R_t|\mathcal{F}_{t-1} : \text{Poisson}(\lambda_t) \\ \lambda_t = \gamma_0 + \gamma_1 R_{t-1} + \delta_1 \lambda_{t-1}, \qquad \gamma_0 > 0, \gamma_1 \geq 0, \delta_1 \geq 0, \gamma_1 + \delta_1 < 1 \end{cases}$$

and the traffic is distributed uniformly among the node pairs (Giorgetti et al. 2015) to yield the random variables $r^{t,(s,d)}(\boldsymbol{\xi}^t)$. We can show that the expectation of the process is $\bar{R} = \dfrac{\gamma_0}{1-(\gamma_1+\delta_1)}$ and the k-step ahead conditional expectation is

$$\mathbb{E}[R_{t+1}|\mathcal{F}_t\ ] = \gamma_0 \sum_{i=0}^{k-2} (\gamma_1+\delta_1)^i + (\gamma_1+\delta_1)^{k-1}(\gamma_0 + \gamma_1 R_t + \delta_1 \lambda_t).$$

For $T \in \{3,4,5\}$ we have generated instances with $\bar{R} \in \{10,30\}$, $\gamma_1 = 0.5$ and two levels of correlation between λ_t parameters of consecutive stages $\delta_1 \in \{0.2, 0.05\}$.

As in the NA variable selection discussion in Section 5.3.3 for the MSLot problem, here for the NA variables we have two options: state variables $\mu^t_{\ell,\omega}(\boldsymbol{\xi}^t)$ and recourse variables $x^{t,(s,d)}_{\ell,\omega}(\boldsymbol{\xi}^t)$. There is a clear mapping between the random parameters and the recourse decisions. However, the number of LDDR decision variables would grow with the number of connection requests. Therefore, we opt to choose $\mu^t_{\ell,\omega}(\boldsymbol{\xi}^t)$ as our NA variables as this requires fewer LDDR decision variables. We use the random request arrivals as the basis functions for the SW and NA duals. Note that there is a mismatch between the indices of $r^{t',(s,d)}(\boldsymbol{\xi}'^t)$ and the dimensions of the state equations. For the NA dual this is also the case since we choose $\mu^t_{\ell,\omega}(\boldsymbol{\xi}^t)$ as the NA variables. Therefore, deciding on the form of the basis functions is not straightforward. Table 6 presents the two options we have considered as the set of basis functions. In Option 2, we try to reduce the number of basis functions by selecting the node pairs that have a higher chance of being served by a wavelink. This is based on the observation made by Jaumard and Daryalal (2017) that RWA models rarely grant the requests of the node pairs on longer paths. For a fixed wavelink $(\hat{\ell}, \hat{\omega})$, we are considering the demands of the node pairs that belong to the following set

$$\mathcal{SD}^t_{\hat{\ell}} = \Big\{(s,d) \in \mathcal{V} \times \mathcal{V} : r^{t,(s,d)} > 0,\ \hat{\ell} \in \bigcup_{i=1}^{2} \mathrm{SP}^i(s,d)\Big\},$$

where $\mathrm{SP}^i(s,d)$ is the set of i^{th} shortest paths between s and d.

Option	SW Dual	NA dual
1	$1\ ;\ \sum_{(s,d)\in\mathcal{SD}^t} r^{t,(s,d)}, \quad \forall t \le \hat{t}$	$1\ ; r^{t,(s,d)},\ \ (s,d) \in \mathcal{SD}^t,\ \ \forall t \le T : \hat{t} < t$
2	$1\ ; r^{t,(s,d)}, \quad (s,d) \in \mathcal{SD}^t_{\hat{\ell}},\ \ \forall t \le \hat{t}$	$1\ ; r^{t,(s,d)},\ \ (s,d) \in \mathcal{SD}^t_{\hat{\ell}},\ \ \forall t \le T : \hat{t} < t$

Table 6: Various options to be used in basis functions, at stage $\hat{t}$ for wavelink $(\hat{\ell}, \hat{\omega})$

5.4.2. Bound Comparison In this section, we present the results of our lower and upper bounding methods on the MSRWA instances using basis function Options 1 and 2 (Table 6). Since our best results in the previous section were obtained by using the SW dual to drive a primal policy, and the NA dual to obtain lower bounds, we only consider these options for the RWA problem.

The results are presented in Table 7. For each option, in the first column we report the ratio of the NA upper bound to the PI upper bound, in the second column we report the ratio of the SW lower bound to the conditional expected value lower bounding policy, and in the third column we report the overall gap reduction (relative to the initial gap) obtained by considering both bounds. Using Option 1, the initial gap is considerably improved, with gap reductions varying from 7% (10^{th} instance) to 63% (1^{st} instance). Our upper and lower bounding methods both contribute to this gap reduction. Although the upper bounds from Option 1 should theoretically always be better than those from Option 2 (because the set of basis functions in Option 2 is a subset of those in Option 1), the higher difficulty in solving the Lagrangian dual for Option 1 meant that we needed to use a smaller sample size, often resulting in worse bounds than were obtained with Option 2 when evaluated out of sample. As evidenced by the results in Table 7, using Option 2, even with a subset of the standard basis functions we are able to reduce the gaps up to 15% more in comparison to the Option 1. This demonstrates the importance of the design of basis functions.

			Basis Function Option 1			Basis Function Option 2		
$(\lvert\Omega\rvert, \hat{R}, \gamma_0, \delta_1)$	T	Initial Gap	$\frac{\text{NA UB}}{\text{PI}}$	$\frac{\text{SW LB}}{\text{Exp}}$	Gap reduction	$\frac{\text{NA UB}}{\text{PI}}$	$\frac{\text{SW LB}}{\text{Exp}}$	Gap reduction
$(1, 10, 3, 0.2)$	3	27.0%	0.94	1.08	63.0%	0.94	1.05	50.4%
	4	43.1%	0.95	1.09	41.8%	0.93	1.12	56.3%
	5	45.1%	0.97	1.15	50.1%	0.96	1.16	54.5%
$(1, 10, 4.5, 0.05)$	3	22.0%	0.97	1.02	27.4%	0.97	1.05	42.8%
	4	34.8%	0.97	1.06	32.0%	0.95	1.07	41.8%
	5	51.4%	0.97	1.18	51.4%	0.94	1.18	59.8%
$(3, 30, 9, 0.2)$	3	27.7%	0.99	1.06	30.5%	0.99	1.08	39.9%
	4	24.2%	0.98	1.03	22.5%	0.98	1.04	29.1%
	5	34.0%	0.98	1.09	42.2%	0.98	1.10	44.3%
$(3, 30, 13.5, 0.05)$	3	20.3%	1.00	1.01	7.0%	1.00	1.03	18.0%
	4	18.3%	0.99	1.03	27.4%	0.99	1.03	29.0%
	5	32.2%	0.99	1.03	17.7%	0.98	1.02	15.6%

Table 7: Gap reductions for MSRWA instances, obtained from combining the improvements of NA upper bound and SW lower bound.

6. Conclusion

In this work, we introduced the idea of Lagrangian dual decision rules where decision rules are used in the Lagrangian dual of an MSMIP. The result is an approximation problem that can be solved by stochastic approximation or sample average approximation. The approximate problem does not have a multistage structure, and hence does not require a scenario tree for its approximate solution. Two lower bounding policies based on two Lagrangian duals are proposed: stagewise and nonanticipative, where the former is an easier problem consisting of single period subproblems, while the latter can potentially lead to better bounds if the basis functions are selected properly. The solutions to both of these duals can be incorporated in constructing primal policies. The lower and upper bounding methods were evaluated by solving instances of two problem classes: a multi-item stochastic lot-sizing problem and routing and wavelength assignment problem in optical networks. The results show that our methods can substantially reduce the optimality gap relative to the use of two general-purpose bounding policies. Future work includes design of more scalable dual-driven policies and investigation of other Lagrangian dual decision rule structures, such as piecewise linear form.

Acknowledgments

This work was supported by Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2018-04984], NSF award CMMI-1634597 and by the Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Contract Number DE-AC02-06CH113.

References

Ahmed S, Cabral FG, da Costa BFP (2020) Stochastic Lipschitz dynamic programming. *Mathematical Programming* 1–39.

Ahmed S, King AJ, Parija G (2003) A multi-stage stochastic integer programming approach for capacity expansion under uncertainty. *Journal of Global Optimization* 26(1):3–24.

Alonso-Ayuso A, Escudero LF, Ortuno MT (2003) BFC, a branch-and-fix coordination algorithmic framework for solving some types of stochastic pure and mixed 0–1 programs. *European Journal of Operational Research* 151(3):503–519.

Angulo G, Ahmed S, Dey SS (2016) Improving the integer L-shaped method. *INFORMS Journal on Computing* 28(3):483–499.

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.

Barty K, Carpentier P, Girardeau P (2010) Decomposition of large-scale stochastic optimal control problems. *RAIRO-Operations Research* 44(3):167–183.

Bodur M, Luedtke J (2017) Mixed-integer rounding enhanced Benders decomposition for multiclass service system staffing and scheduling with arrival rate uncertainty. *Management Science* 63:2073–2091.

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

Brown DB, Smith JE, Sun P (2010) Information relaxations and duality in stochastic dynamic programs. *Operations Research* 58(4-part-1):785–801.

Buhmann MD (2000) Radial basis functions. *Acta Numerica* 9:1–38.

Carøe CC, Schultz R (1999) Dual decomposition in stochastic integer programming. *Operations Research Letters* 24(1-2):37–45.

Cen Z (2012) Solving multi-stage stochastic mixed integer linear programs by the dual dynamic programming approach. Research Report RR-7868, INRIA.

Cerisola S, Latorre JM, Ramos A (2012) Stochastic dual dynamic programming applied to nonconvex hydrothermal models. *European Journal of Operational Research* 218(3):687–697.

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, Zhu S, Jiang L, Zhu Z (2015) On spectrum efficient failure-independent path protection p-cycle design in elastic optical networks. *Journal of Lightwave Technology* 33(17):3719–3729.

Chen ZL, Li S, Tirupati D (2002) A scenario-based stochastic programming approach for technology and capacity planning. *Computers & Operations Research* 29(7):781–806.

Daryalal M, Bodur M (2020) Stochastic RWA and lightpath rerouting in WDM networks. Preprint at https://arxiv.org/pdf/2012.10084.pdf.

Dentcheva D, Römisch W (2004) Duality gaps in nonconvex stochastic optimization. *Mathematical Programming* 101(3):515–535.

Fan Y, Liu C (2010) Solving stochastic transportation network protection problems using the progressive hedging-based method. *Networks and Spatial Economics* 10(2):193–208.

Ferland R, Latour A, Oraichi D (2006) Integer-valued GARCH process. *Journal of Time Series Analysis* 27(6):923–942.

Flach B, Barroso L, Pereira M (2010) Long-term optimal allocation of hydro generation for a price-maker company in a competitive market: latest developments and a stochastic dual dynamic programming approach. *IET Generation, Transmission & Distribution* 4(2):299–314.

Flatabø N, Haugstad A, Mo B, Fosso OB (1998) Short-term and medium-term generation scheduling in the Norwegian hydro system under a competitive power market structure. *International Conference on Electrical Power System Operation and Management*.

Friedman J, Hastie T, Tibshirani R, et al. (2001) *The elements of statistical learning*, volume 1 (Springer Series in Statistics New York).

Gade D, Hackebeil G, Ryan SM, Watson JP, Wets RJB, Woodruff DL (2016) Obtaining lower bounds from the progressive hedging algorithm for stochastic mixed-integer programs. *Mathematical Programming* 157(1):47–67.

Gade D, Küçükyavuz S, Sen S (2014) Decomposition algorithms with parametric Gomory cuts for two-stage stochastic integer programs. *Mathematical Programming* 144(1-2):39–64.

Georghiou A, Wiesemann W, Kuhn D (2015) Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming* 152(1-2):301–338.

Giorgetti A, Paolucci F, Cugini F, Castoldi P (2015) Dynamic restoration with GMPLS and SDN control plane in elastic optical networks. *Journal of Optical Communications and Networking* 7(2):A174–A182.

Guan Y, Ahmed S, Nemhauser GL (2009) Cutting planes for multistage stochastic integer programs. *Operations Research* 57(2):287–298.

Gul S, Denton BT, Fowler JW (2015) A progressive hedging approach for surgery planning under uncertainty. *INFORMS Journal on Computing* 27(4):755–772.

Hajizadeh E, Mahootchi M (2018) Optimized radial basis function neural network for improving approximate dynamic programming in pricing high dimensional options. *Neural Computing and Applications* 30(6):1783–1794.

Helber S, Sahling F, Schimmelpfeng K (2013) Dynamic capacitated lot sizing with random demand and dynamic safety stocks. *OR Spectrum* 35(1):75–105.

Helseth A, Mo B, Fodstad M, Hjelmeland MN (2015) Co-optimizing sales of energy and capacity in a hydropower scheduling model. *2015 IEEE Eindhoven PowerTech*, 1–6.

Higle J, Sen S (1996) *Stochastic Decomposition: A Statistical Method for Large Scale Stochastic Linear Programming* (Boston: Springer-Verlag).

Jaumard B, Daryalal M (2017) Efficient spectrum utilization in large scale RWA problems. *IEEE/ACM Transactions on Networking* 25(2):1263–1278.

Keller PW, Mannor S, Precup D (2006) Automatic basis function construction for approximate dynamic programming and reinforcement learning. *Proceedings of the 23rd International Conference on Machine learning*, 449–456.

Kim K, Petra CG, Zavala VM (2019) An asynchronous bundle-trust-region method for dual decomposition of stochastic mixed-integer programming. *SIAM Journal on Optimization* 29(1):318–342.

Kiwiel KC (1995) Approximations in proximal bundle methods and decomposition of convex programs. *Journal of Optimization Theory and Applications* 84(3):529–548.

Knapp AW (2005) *Basic Real Analysis* (Springer Science & Business Media).

Küçükyavuz S, Sen S (2017) An introduction to two-stage stochastic mixed-integer programming. *Leading Developments from INFORMS Communities*, 1–27 (INFORMS).

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

Laporte G, Louveaux FV (1993) The integer L-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters* 13(3):133–142.

Listes O, Dekker R (2005) A scenario aggregation–based approach for determining a robust airline fleet composition for dynamic capacity allocation. *Transportation Science* 39(3):367–382.

Löhndorf N, Wozabal D, Minner S (2013) Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research* 61(4):810–823.

Løkketangen A, Woodruff DL (1996) Progressive hedging and tabu search applied to mixed integer (0, 1) multistage stochastic programming. *Journal of Heuristics* 2(2):111–128.

Lubin M, Martin K, Petra CG, Sandıkçı B (2013) On parallelizing dual decomposition in stochastic integer programming. *Operations Research Letters* 41(3):252–258.

Lulli G, Sen S (2004) A branch-and-price algorithm for multistage stochastic integer programming with application to stochastic batch-sizing problems. *Management Science* 50(6):786–796.

Lulli G, Sen S (2006) A heuristic procedure for stochastic integer programs with complete recourse. *European Journal of Operational Research* 171(3):879–890.

Mo B, Gjelsvik A, Grundt A, Karesen K (2001) Optimisation of hydropower operation in a liberalised market with focus on price modelling. *IEEE Porto Power Tech Proceedings*, volume 1, 6–pp.

Nemirovski A, Juditsky A, Lan G, Shapiro A (2009) Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization* 19(4):1574–1609.

Newham N, Wood A (2007) Transmission investment planning using SDDP. *Australasian Universities Power Engineering Conference*, 1–5.

Norkin VI, Pflug GC, Ruszczyński A (1998) A branch and bound method for stochastic global optimization. *Mathematical Programming* 83(1-3):425–450.

Nowak MP, Römisch W (2000) Stochastic Lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty. *Annals of Operations Research* 100(1-4):251–272.

Orlowski S, Wessäly R, Pióro M, Tomaszewski A (2010) SNDlib 1.0—Survivable network design library. *Networks: An International Journal* 55(3):276–286.

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

Philpott A, Wahid F, Bonnans J (2016) MIDAS: A mixed integer dynamic approximation scheme. *Mathematical Programming* 1–32.

Powell WB (2009) What you should know about approximate dynamic programming. *Naval Research Logistics* 56(3):239–249.

Qi Y, Sen S (2017) The ancestral Benders' cutting plane algorithm with multi-term disjunctions for mixed-integer recourse decisions in stochastic programming. *Mathematical Programming* 161(1-2):193–235.

Rahmaniani R, Ahmed S, Crainic TG, Gendreau M, Rei W (2020) The benders dual decomposition method. *Operations Research* 68(3):878–895.

Rajagopalan S, Singh MR, Morton TE (1998) Capacity expansion and replacement in growing markets with uncertain technological breakthroughs. *Management Science* 44(1):12–30.

Robbins H, Monro S (1951) A stochastic approximation method. *The Annals of Mathematical Statistics* 22(3):400–407.

Rockafellar RT, Wets RJ (1976) Stochastic convex programming: relatively complete recourse and induced feasibility. *SIAM Journal on Control and Optimization* 14(3):574–589.

Rockafellar RT, Wets RJB (1991) Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research* 16(1):119–147.

Rosa CH, Ruszczyński A (1996) On augmented Lagrangian decomposition methods for multistage stochastic programs. *Annals of Operations Research* 64(1):289–309.

Ruszczyński A (1986) A regularized decomposition method for minimizing a sum of polyhedral functions. *Mathematical Programming* 35(3):309–333.

Sandıkçı B, Ozaltın OY (2014) A scalable bounding method for multi-stage stochastic integer programs. *Chicago Booth Research Paper* .

Sen S, Higle JL (2005) The C^3 theorem and a D^2 algorithm for large scale stochastic mixed-integer programming: set convexification. *Mathematical Programming* 104:1–20.

Sen S, Sherali H (2006) Decomposition with branch-and-cut approaches for two-stage stochastic mixed-integer programming. *Mathematical Programming* 106:203–223.

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

Shapiro A, Nemirovski A (2005) On complexity of stochastic programming problems. *Continuous Optimization*, 111–146 (Springer).

Sherali H, Zhu X (2007) On solving discrete two-stage stochastic programs having mixed-integer first- and second-stage variables. *Mathematical Programming* 108:597–616.

Simmons JM, Rouskas GN (2020) Routing and wavelength (spectrum) assignment. *Handbook of Optical Networks*, 447–484 (Springer).

Singh KJ, Philpott AB, Wood RK (2009) Dantzig-Wolfe decomposition for solving multistage stochastic capacity-planning problems. *Operations Research* 57(5):1271–1286.

Takriti S, Birge JR, Long E (1996) A stochastic model for the unit commitment problem. *IEEE Transactions on Power Systems* 11(3):1497–1508.

Tanner M, Ntaimo L (2008) Computations with disjunctive cuts for two-stage stochastic mixed 0-1 integer programs. *Journal of Global Optimization* 58:365–384.

Watson JP, Woodruff DL (2011) Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems. *Computational Management Science* 8(4):355–370.

Xiong Y, Li Y, Zhou B, Wang R, Rouskas GN (2018) SDN enabled restoration with triggered precomputation in elastic optical inter-datacenter networks. *Journal of Optical Communications and Networking* 10(1):24–34.

Zhang M, Küçükyavuz S (2014) Finitely convergent decomposition algorithms for two-stage stochastic pure integer programs. *SIAM Journal on Optimization* 24:1933–1951.

Zou J (2017) *Large scale multistage stochastic integer programming with applications in electric power systems*. Ph.D. thesis, Georgia Institute of Technology.

Zou J, Ahmed S, Sun XA (2017) Stochastic dual dynamic integer programming. *Mathematical Programming* 1–42.

Appendix

In the following, we have provided further details on the proofs and results given in the paper.

A. Proof of Lemma 1

Expanding the objective function in (5a) and rearranging the terms yields

$$\sum_{t\in[T]} \left(\mathbb{E}_{\boldsymbol{\xi}^T}\left[c_t(\boldsymbol{\xi}^t)^\top y_t(\boldsymbol{\xi}^T)\right] + \mathbb{E}_{\boldsymbol{\xi}^T}\left[\gamma_t(\boldsymbol{\xi}^T)^\top y_t(\boldsymbol{\xi}^T)\right] - \mathbb{E}_{\boldsymbol{\xi}^T}\left[\gamma_t(\boldsymbol{\xi}^T)^\top \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right]\right). \tag{A.1}$$

For the last term in (A.1), the following equalities hold

$$\mathbb{E}_{\boldsymbol{\xi}^T}\left[\gamma_t(\boldsymbol{\xi}^T)^\top \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right] = \mathbb{E}_{\boldsymbol{\xi}^T}\left[\mathbb{E}_{\boldsymbol{\xi}'^T}[\gamma_t(\boldsymbol{\xi}^T)^\top y_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right] \tag{A.2a}$$

$$= \mathbb{E}_{\boldsymbol{\xi}'^T}\left[\mathbb{E}_{\boldsymbol{\xi}^T}[\gamma_t(\boldsymbol{\xi}^T)^\top y_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right] \tag{A.2b}$$

$$= \mathbb{E}_{\boldsymbol{\xi}'^T}\left[\mathbb{E}_{\boldsymbol{\xi}^T}[\gamma_t(\boldsymbol{\xi}^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\, y_t(\boldsymbol{\xi}'^T)\right] \tag{A.2c}$$

$$= \mathbb{E}_{\boldsymbol{\xi}^T}\left[\mathbb{E}_{\boldsymbol{\xi}'^T}[\gamma_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t] y_t(\boldsymbol{\xi}^T)\right]. \tag{A.2d}$$

First note that, in $\mathbb{E}_{\boldsymbol{\xi}^T}\left[\gamma_t(\boldsymbol{\xi}^T)^\top \mathbb{E}_{\boldsymbol{\xi}'^T}[y_t(\boldsymbol{\xi}'^T)|\boldsymbol{\xi}'^t = \boldsymbol{\xi}^t]\right]$, since the dual function $\gamma_t(\boldsymbol{\xi}^T)$ is fixed, $\gamma_t(\boldsymbol{\xi}^T)$ is a vector of numbers inside the first expectation, hence it can be pushed inside the second expectation (Equation (A.2a)). As we have assumed that $\gamma_t(\boldsymbol{\xi}^T)$ is a member of the set Γ_t with its expectation being bounded, and the expected diameter of the set $Y_t(\boldsymbol{\xi}'^t)$ to which $y_t(\boldsymbol{\xi}'^t)$ belongs is finite, using Fubini-Toneli Theorem (Knapp 2005) the order of the two expectations can be exchanged (Equation (A.2b)). Then we can take $y_t(\boldsymbol{\xi}'^T)$ out from the inside expectation, as it is just a vector of numbers inside (Equation (A.2c)). Since $\boldsymbol{\xi}^T$ and $\boldsymbol{\xi}'^T$ have the same support and distribution, the last equality (Equation (A.2d)) is satisfied. Equality (6) can then be proven by substitution.

B. Proof of Lemma 2

First observe that since the objective in (12) is linear, the restricted Lagrangian dual problem (14) can be written as

$$z^{\text{RLD}} = \max_{\lambda,\alpha} \; \min_{x} \; \{c^\top x + \lambda^\top (d - Dx) : x \in \text{conv}(X)\}$$

$$\text{s.t.} \quad \lambda - G\alpha = 0$$

Let $\{x^i\}_{i\in[M]}$ and $\{r^k\}_{k\in[K]}$ be the complete set of extreme points and extreme rays of conv(X), respectively. Then, for any fixed λ, we have

$$z(\lambda) = \begin{cases} -\infty, & \exists r^k : (c^\top - \lambda^\top D) r^k < 0 \\ \min_{i\in[M]} \{c^\top x^i + \lambda^\top (d - Dx^i)\}, \text{ otherwise.} & \end{cases}$$

Therefore we can reformulate the Lagrangian dual problem as

$$\begin{aligned} z^{\mathrm{LD}} = \max_{\lambda,\alpha} \; \min_{i\in[M]} \quad & \{c^\top x^i + \lambda^\top (d - Dx^i)\} \\ \text{s.t.} \quad & \lambda - G\alpha = 0 \\ & (c^\top - \lambda^\top D) r^k \geq 0, \quad k \in [K] \end{aligned}$$

which is equivalent to

$$z^{\mathrm{LD}} = \max_{\lambda,\alpha,\eta} \quad \eta \tag{B.3a}$$

$$\text{s.t.} \quad \lambda - G\alpha = 0 \qquad (\theta) \tag{B.3b}$$

$$\lambda^\top D r^k \leq c^\top r^k \qquad k \in [K] \qquad (\beta_k) \tag{B.3c}$$

$$\eta + \lambda^\top (Dx^i - d) \leq c^\top x^i \qquad i \in [M] \qquad (\gamma_i) \tag{B.3d}$$

where θ, β_k and γ_i are the dual variables associated with constraints (B.3b), (B.3c) and (B.3d) respectively. Now take the dual of the above problem

$$\begin{aligned} \min_{\theta,\beta,\gamma} \quad & c^\top \left(\sum_{i\in[M]} \gamma_i x^i + \sum_{k\in[K]} \beta_k r^k \right) \\ \text{s.t.} \quad & \sum_{i\in[M]} \gamma_i = 1 \\ & D \left(\sum_{i\in[M]} \gamma_i x^I + \sum_{k\in[K]} \beta_k r^k \right) + \theta = d \\ & -G^\top \theta = 0 \end{aligned}$$

$$\gamma, \beta \geq 0$$

As we know that $\mathrm{conv}(X) = \left\{ \sum_{i\in[M]} \gamma_i x^i + \sum_{k\in[K]} \beta_k r^k : \sum_{i\in[M]} \gamma_i = 1, \gamma_i, \beta_k \geq 0,\ i \in [M], k \in [K] \right\}$, we have

$$\begin{aligned} \min\ & c^\top x \\ \text{s.t.}\ & x \in \mathrm{conv}(X) \\ & Dx + \theta = d \\ & -G^\top \theta = 0. \end{aligned} \tag{B.4}$$

Eliminating the θ variables using (B.4) yields the result.

Note that if $\theta = \mathbf{0}$, then the original MIP solution, say x^*, is feasible. So, this gives a relaxation of the original MIP.

C. The Lot-sizing Problem

C.1. Stagewise Lagrangian Dual

Relax state equations (27b), except for $t = 1$. Then, we have

$$\min\ \mathbb{E}\left[L(\lambda)\right] \tag{C.5a}$$

$$\text{s.t.}\ i^-_{1j}(\boldsymbol{\xi}^1) - i^+_{1j}(\boldsymbol{\xi}^1) = d_{1j}(\boldsymbol{\xi}^1), \quad j \in [J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \tag{C.5b}$$

$$(27\text{c}) - (27\text{i}) \tag{C.5c}$$

where

$$\begin{aligned} L(\lambda) := & \sum_{t\in[T]} \left(\sum_{j\in[J]} \left(C^{i+}_{tj}(\boldsymbol{\xi}^t) i^+_{tj}(\boldsymbol{\xi}^t) + C^{i-}_{tj}(\boldsymbol{\xi}^t) i^-_{tj}(\boldsymbol{\xi}^t) + C^y_{tj}(\boldsymbol{\xi}^t) y_{tj}(\boldsymbol{\xi}^t) \right) + C^o_t(\boldsymbol{\xi}^t) o_t(\boldsymbol{\xi}^t) \right) \\ & + \sum_{t\in[T]} \sum_{j\in[J]} \lambda_{tj}(\boldsymbol{\xi}^t) \left(i^-_{tj}(\boldsymbol{\xi}^t) - i^+_{tj}(\boldsymbol{\xi}^t) + i^+_{t-1,j}(\boldsymbol{\xi}^{t-1}) - i^-_{t-1,j}(\boldsymbol{\xi}^{t-1}) + x_{t-1,j}(\boldsymbol{\xi}^{t-1}) - D_{tj}(\boldsymbol{\xi}^t) \right). \end{aligned}$$

Restrict $\lambda_t(\boldsymbol{\xi}^t)$ to follow LDDR

$$\lambda_{tj}(\boldsymbol{\xi}^t) = \sum_{k\in[K_t]} \Phi_{tjk}(\boldsymbol{\xi}^t) \alpha_{tjk}.$$

Then, for fixed $\hat{\alpha}$, the objective function (C.5a) is equivalent to

$$\sum_{t\in[T]} \mathbb{E}\left[L_t(\hat{\alpha})\right] \tag{C.6}$$

where

$$\begin{aligned}
&L_t(\hat{\alpha},\boldsymbol{\xi}^t) := \\
&\sum_{j\in[J]}\Bigg[\Bigg(\sum_{k\in[K_t]}\Phi_{tjk}(\boldsymbol{\xi}^t)\alpha_{tjk}\Bigg)\big(-D_{tj}(\boldsymbol{\xi}^t)\big) \\
&+\Bigg(\mathbb{E}\Bigg[\sum_{k\in[K_{t+1}]}\Phi_{t+1,j,k}(\boldsymbol{\xi}^{t+1})\alpha_{t+1,jk}\,\Bigg|\,\boldsymbol{\xi}^t\Bigg]\Bigg)x_{tj} \\
&+\Bigg(C^{i^+}_{tj}(\boldsymbol{\xi}^t)-\sum_{k\in[K_t]}\Phi_{tjk}(\boldsymbol{\xi}^t)\alpha_{tjk}+\mathbb{E}\Bigg[\sum_{k\in[K_{t+1}]}\Phi_{t+1,j,k}(\boldsymbol{\xi}^{t+1})\alpha_{t+1,jk}\,\Bigg|\,\boldsymbol{\xi}^t\Bigg]\Bigg)i^+_{tj} \\
&+\Bigg(C^{i^-}_{tj}(\boldsymbol{\xi}^t)+\sum_{k\in[K_t]}\Phi_{tjk}(\boldsymbol{\xi}^t)\alpha_{tjk}-\mathbb{E}\Bigg[\sum_{k\in[K_{t+1}]}\Phi_{t+1,j,k}(\boldsymbol{\xi}^{t+1})\alpha_{t+1,jk}\,\Bigg|\,\boldsymbol{\xi}^t\Bigg]\Bigg)i^-_{tj} \\
&+C^y_{tj}(\boldsymbol{\xi}^t)y_{tj}\Bigg]+C^o_t(\boldsymbol{\xi}^t)o_t.
\end{aligned}$$

For fixed $\hat{\alpha}$, $\boldsymbol{\xi}^t$ and $t>1$, define $\mathcal{L}_t(\hat{\alpha},\boldsymbol{\xi}^t)$ as follows

$$\begin{aligned}
\mathcal{L}_t(\hat{\alpha},\boldsymbol{\xi}^t) := \min\quad & L_t(\hat{\alpha},\boldsymbol{\xi}^t) \\
\text{s.t.}\quad & \sum_{j\in[J]}\big(TS_j y_{tj}+TB_j x_{tj}\big)-o_t\le C_t \\
& M_{tj}\; y_{tj}-x_{tj}\ge 0, && j\in[J] \\
& i^+_{tj}\le I_{tj}, && j\in[J] \\
& i^+_{tj}+x_{tj}\le I_{t+1,j}, && j\in[J] \\
& 0\le o_t\le O_t \\
& x_{tj}, i^+_{tj}, i^-_{tj}\ge 0, && j\in[J] \\
& y_{tj}\in\{0,1\}, && j\in[J].
\end{aligned}$$

For $t=1$ we have

$$\mathcal{L}_1(\hat{\alpha},\boldsymbol{\xi}^1) := \min\quad L_1(\hat{\alpha},\boldsymbol{\xi}^1)$$

$$
\begin{aligned}
\text{s.t.}\quad & i^-_{1j} - i^+_{1j} = d_{1j}, && j \in [J] \\
& \sum_{j\in[J]} \big(TS_j y_{1j} + TB_j x_{1j}\big) - o_1 \le C_1 \\
& M_{1j}\ y_{1j} - x_{1j} \ge 0, && j \in [J] \\
& i^+_{1j} \le I_{1j}, && j \in [J] \\
& i^+_{1j} + x_{1j} \le I_{2,j}, && j \in [J] \\
& 0 \le o_1 \le O_1 \\
& x_{1j}, i^+_{1j}, i^-_{1j} \ge 0, && j \in [J] \\
& y_{1j} \in \{0,1\}, && j \in [J].
\end{aligned}
$$

Finally, the LDDR-restricted stagewise Lagrangian dual problem is defined as

$$
\begin{aligned}
\max\ & \sum_{t\in[T]} \mathbb{E}\left[\mathcal{L}_t(\alpha, \boldsymbol{\xi}^t)\right] \\
\text{s.t.}\ & \alpha_{tk} \in \mathbb{R}^J, \qquad t \in [T], k \in [K_t].
\end{aligned}
$$

C.2. Nonanticipative Lagrangian Dual

Reformulate the MSLot problem as follows

$$
\begin{aligned}
\min\ \ & \mathbb{E}\left[\sum_{t\in[T]} \left[\sum_{j\in[J]} \left[C^{i^+}_{tj}(\boldsymbol{\xi}^t) i^{+,na}_{tj}(\boldsymbol{\xi}^T) + C^{i^-}_{tj}(\boldsymbol{\xi}^t) i^{-,na}_{tj}(\boldsymbol{\xi}^T) + C^{y}_{tj}(\boldsymbol{\xi}^t) y^{na}_{tj}(\boldsymbol{\xi}^T)\right] + C^o_t(\boldsymbol{\xi}^t) o^{na}_t(\boldsymbol{\xi}^T)\right]\right] \\
\text{s.t.}\ \ & i^{-,na}_{tj}(\boldsymbol{\xi}^T) - i^{+,na}_{tj}(\boldsymbol{\xi}^T) + i^{+,na}_{t-1,j}(\boldsymbol{\xi}^T) - i^{-,na}_{t-1,j}(\boldsymbol{\xi}^T) + x^{na}_{t-1,j}(\boldsymbol{\xi}^T) = D_{tj}(\boldsymbol{\xi}^t), \\
& && t \in [T], j \in [J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \\
& \sum_{j\in[J]} (TS_j y^{na}_{tj}(\boldsymbol{\xi}^T) + TB_j x^{na}_{tj}(\boldsymbol{\xi}^T)) - o^{na}_t(\boldsymbol{\xi}^T) \le C_t, && t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \\
& M_{tj}\ y^{na}_{tj}(\boldsymbol{\xi}^T) - x^{na}_{tj}(\boldsymbol{\xi}^T) \ge 0, && t \in [T], j \in [J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \\
& i^{+,na}_{tj}(\boldsymbol{\xi}^T) \le I_{tj}, && t \in [T], j \in [J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \\
& i^{+,na}_{tj}(\boldsymbol{\xi}^T) + x^{+,na}_{tj}(\boldsymbol{\xi}^T) \le I_{t+1,j}, && t \in [T], j \in [J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi \\
& 0 \le o^{na}_t(\boldsymbol{\xi}^T) \le O_t, && t \in [T], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi
\end{aligned}
$$

$$x_{tj}^{na}(\boldsymbol{\xi}^T) = \mathbb{E}_{\boldsymbol{\xi}'^T}[x_{tj}^{na}(\boldsymbol{\xi}'^T) \,|\, \boldsymbol{\xi}^t], \qquad t \in [T], j \in [J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi$$

$$x_{tj}^{na}(\boldsymbol{\xi}^T), i_{tj}^{+,na}(\boldsymbol{\xi}^T), i_{tj}^{-,na}(\boldsymbol{\xi}^T) \geq 0, \qquad t \in [T], j \in [J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi$$

$$y_{tj}^{na}(\boldsymbol{\xi}^T) \in \{0,1\}, \qquad t \in [T], j \in [J], \mathbb{P}\text{-a.e. } \boldsymbol{\xi}^T \in \Xi$$

where for any variable $a^{na}(\boldsymbol{\xi}^T)$ the superscript na indicates the anticipative copy variable corresponding to original variable $a(\boldsymbol{\xi}^t)$. Relaxing the nonanticipativity constraints using dual variables $\gamma_t(\boldsymbol{\xi}^T)$ and enforcing LDDR on these duals as

$$\gamma_{tj}(\boldsymbol{\xi}^T) = \sum_{k' \in [K'_t]} \Psi_{tjk'}(\boldsymbol{\xi}^T)\alpha_{tjk'},$$

the LDDR-restricted nonanticipative Lagrangian dual problem is obtained as follows

$$\begin{aligned} &\max \; \mathbb{E}\left[\mathcal{L}(\alpha, \boldsymbol{\xi}^T)\right] \\ &\text{s.t. } \alpha_{tk'} \in \mathbb{R}^J, \qquad t \in [T], k' \in [K'_t] \end{aligned}$$

where for fixed $\hat{\alpha}$

$$\begin{aligned}
\mathcal{L}(\hat{\alpha}, \boldsymbol{\xi}^T) = \min \quad & \sum_{t \in [T]} \left[\sum_{j \in [J]} \left[C_{tj}^{i^+}(\boldsymbol{\xi}^t) i_{tj}^{+,na} + C_{tj}^{i^-}(\boldsymbol{\xi}^t) i_{tj}^{-,na} + C_{tj}^{y}(\boldsymbol{\xi}^t) y_{tj}^{na} \right] + C_t^{o}(\boldsymbol{\xi}^t) o_t^{na} \right] \\
& + \sum_{t \in [T]} \sum_{j \in [J]} \sum_{k' \in [K'_t]} \Big(\Psi_{tjk'}(\boldsymbol{\xi}^T) - \mathbb{E}_{\boldsymbol{\xi}'^T}[\Psi_{tjk'}(\boldsymbol{\xi}'^T) | \boldsymbol{\xi}^t] \Big) \hat{\alpha}_{tjk'} x_{tj}^{na} \\
\text{s.t.} \quad & i_{tj}^{-,na} - i_{tj}^{+,na} + i_{t-1,j}^{+,na} - i_{t-1,j}^{-,na} + x_{t-1,j}^{na} = D_{tj}(\boldsymbol{\xi}^t), && t \in [T], j \in [J] \\
& \sum_{j \in [J]} \left(TS_j y_{tj}^{na} + TB_j x_{tj}^{na} \right) - o_t^{na} \leq C_t, && t \in [T] \\
& M_{tj} \; y_{tj}^{na} - x_{tj}^{na} \geq 0, && t \in [T], j \in [J] \\
& i_{tj}^{+,na} \leq I_{tj}, && t \in [T], j \in [J] \\
& i_{tj}^{+,na} + x_{tj}^{+,na} \leq I_{t+1,j}, && t \in [T], j \in [J] \\
& 0 \leq o_t^{na} \leq O_t, && t \in [T] \\
& x_{tj}^{na}, i_{tj}^{+,na}, i_{tj}^{-,na} \geq 0, && t \in [T], j \in [J] \\
& y_{tj}^{na} \in \{0,1\}, && t \in [T], j \in [J].
\end{aligned}$$

C.3. Parameters of MSLot Instances

Our instances are generated using parameters that are loosely based on the work of Helber et al. (2013). Section 5.3.1 explains the demand generation procedure. In this section, we specify the rest of the parameters. We do not consider a production cost. The overtime cost is 100 per unit of overtime. Holding cost is 15 per unit, while the backlog cost is $(\delta_{i^-})(c^{i^+}_{tj})$, where $\delta_{i^-} = 2$. For the last stage though, we have an end of horizon effect and $c^{i^-}_{Tj} = 150$. TB_j is set to 1, and $TS_j = ts^{rel}\ \overline{\mathbb{E}[D_j]}\ TB_j$, where $\overline{\mathbb{E}[D_j]} = \frac{\sum_{t\in[T]} \mathbb{E}[D_{tj}(\boldsymbol{\xi}^t)]}{T}$ is the average expected demand of product j, and ts^{rel} is 0.25. The setup cost is $c^y_{tj} = \delta_y\ \overline{\mathbb{E}[D_j]}\ TBO^2\ c^{i^+}_{ij}$, where TBO is the processing time between orders (set to 2), and δ_y is 1.2. Production capacity is $C_t = 0.9\frac{\sum_{j\in[J]} \mathbb{E}[D_{tj}(\boldsymbol{\xi}^t)]}{Util}$, where $Util$ is 0.6. Inventory capacity is $I_{tj} = \delta_I\ \overline{\mathbb{E}[D_j]}$, with $\delta_I = 10$. The bound on overtime is $O_t = \delta_O C_t$, where $\delta_O = 0.25$. For big-M values in the MSLot formulation, we have $M_j = 6\ \overline{\mathbb{E}[D_j]}$.

C.4. Numerical Section - Detailed Results

In the next pages, the actual numbers used in Section 5 are provided, without any scaling and normalization. Also, the confidence intervals are given for the bounds, in the form of (mean ± width).

	# Basis functions			Time (s)			Bound		
Option	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$
1	315	750	1365	353.8	1412.5	6273.9	32474.6 ± 448.6	41291.3 ± 1194.2	69849.2 ± 5151.9
2	180	300	420	91.7	243.7	632.5	30947.9 ± 363.9	39835.2 ± 496.4	66388.5 ± 5159.7
3	135	300	525	109	288.7	795.9	31053.7± 371.9	40195.4 ± 671.9	66727.2 ± 2203.9
4	90	150	210	157.2	795.9	323.6	29245.5 ± 396.9	38958.3 ± 1086.0	56882.1 ± 4201.2

Table C.1: Basis function selection analysis for the SW dual

Option	# Basis functions			Time (s)			Bound		
	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$
1	756	1800	3276	76.9	403.2	1800.8	53352.5 ± 1261.5	71767.2 ± 2515.2	123173.0 ± 5101.1
2	432	720	1008	44.9	232.3	770.5	49940.7 ± 403.5	67696.9 ± 501.2	114913.0 ± 1834.5
3	324	720	1260	61.1	250	776.8	53328.5 ± 1186.6	71352.1 ± 1539.8	122755.0 ± 3085.8
4	216	360	504	34.2	219.5	735.3	49902.7 ± 407.0	67247.6 ± 442.2	114829.0 ± 1715.5

Table C.2: Basis function selection analysis for the NA dual

Option	Time (s)			Bound		
	$T=4$	$T=6$	$T=8$	$T=4$	$T=6$	$T=8$
x	61.1	68.1	175.4	53328.5 ± 1186.6	71352.1 ± 1539.8	122755.0 ± 3085.8
i^+, i^-	250.0	188.0	849.9	49456.0 ± 403.9	65348.7 ± 392.2	111605.0 ± 1778.4
x, i^+, i^-	776.8	653.1	1887.3	53459.1 ± 1012.2	71844.8 ± 939.4	123403.0 ± 2400.2

Table C.3: NA variable selection analysis

λ	Bound		
	$T=4$	$T=6$	$T=8$
0.00	58292.3 ± 481.1	81073.9 ± 672.5	143503.0 ± 1912.6
0.25	56624.5 ± 503.2	80815.3 ± 689.7	143503.2 ± 1928.4
0.50	56925.1 ± 499.4	81138.1 ± 680.5	144290.8 ± 1891.7
1.00	118093.6 ± 1357.2	175553.3 ± 1921.6	282653.2 ± 2651.1

Table C.4: Analysis of different values for λ in SW UB for the lotsizing problem

Instance	T	LB PI	LB SW	LB NA	UB Cond Exp	UB SW	UB NA
$\rho = 0.2, \rho^Y = 0.6, J = 3$	2	46584.6	33327.5	47189.4	49406.4	49406.4	49001.1
	3	45487.7	29308.6	46735.3	51076.7	50660.4	50953.9
	4	48868.7	29962.1	50537.6	56731.2	54878.6	53614.3
	5	55920.5	31638.4	56965.3	64759.2	64062.4	61419.5
	6	64844.7	38036.7	67777.5	75596.2	75515.2	75384.4
	7	72980.9	42580.4	75263.3	84230.4	83838.4	>6h
	8	104296.0	67443.0	112093.0	129166.0	127026.0	>6h
	9	102604.0	60101.1	104760.0	121131.0	117865.0	>6h
	10	114903.0	70844.6	122374.0	135756.0	134728.0	>6h
$\rho = 0.6, \rho^Y = 0.2, J = 3$	2	46587.0	35050.1	46876.1	48484.6	48484.6	48277.4
	3	45547.9	29282.9	46585.9	50607.9	50109.8	49404.2
	4	49070.8	32474.3	53328.5	58292.3	56624.5	54892.8
	5	55927.6	35362.3	59325.2	67492.5	66928.6	63767.0
	6	65167.3	41291.3	71352.1	81073.9	80815.3	82093.8
	7	73532.0	46230.9	80838.4	91635.6	90969.8	>6h
	8	110344.0	69849.2	122755.0	145747.0	143503.0	>6h
	9	105611.0	63981.7	115548.0	135670.0	130932.0	>6h
	10	120331.0	67950.0	134594.0	155783.0	154407.0	>6h
$\rho = 0.6, \rho^Y = 0.2, J = 6$	2	97422.8	72505.3	98890.1	103233.0	103233.0	102963.5
	3	90470.1	56340.1	92717.4	104545.0	104218.0	103844.1
	4	94000.1	53199.8	100338	113226.0	111675.8	111718.4
	5	107992.0	62849.1	114973.0	134017.0	132604.4	132614.5
	6	121141.0	67740.8	131226.0	153875.0	153151.2	>6h
	7	118421.0	65444.7	130542.0	175972.0	174715.2	>6h
	8	139644.0	106919.0	157804.0	240137.0	238465.8	>6h
	9	127301.0	107217.7	142417.0	235657.0	233470.2	>6h
	10	135019.0	118333.6	156096.0	273716.0	270377.2	>6h

Table C.5: Bound comparison for various instances

Instance	T	SW Time (s) LB	SW Time (s) UB (avg./sample path)	NA Time (s) LB	NA Time (s) UB (avg./sample path)
$\rho = 0.2, \rho^Y = 0.6, J = 3$	2	3.4	<0.1s	3.9	<0.1s
	3	18.3	<0.1s	24.4	0.2
	4	161.4	<0.1s	59.7	0.8
	5	338.3	<0.1s	115.9	1.7
	6	750.7	0.1	339.0	3.7
	7	960.0	0.2	437.4	7.7
	8	2184.1	0.3	605.3	9.5
	9	2460.0	0.4	714.0	12.4
	10	2711.4	0.5	1975.4	17.6
$\rho = 0.6, \rho^Y = 0.2, J = 3$	2	3.5	<0.1s	3.9	<0.1s
	3	17.6	<0.1s	25.2	0.2
	4	67.6	<0.1s	58.2	0.7
	5	115.1	<0.1s	72.5	1.8
	6	675.2	0.1	374.6	4.0
	7	883.8	0.1	413.5	7.4
	8	1499.0	0.3	645.0	9.6
	9	1725.6	0.5	696.5	12.2
	10	1919.7	0.4	1741.7	16.5
$\rho = 0.6, \rho^Y = 0.2, J = 6$	2	20.2	<0.1s	19.0	0.1
	3	288.6	<0.1s	119.8	0.6
	4	512.7	<0.1s	679.9	2.0
	5	4747.9	<0.1s	1334.4	5.0
	6	6154.1	0.1	2290.8	8.2
	7	10963.4	0.2	3749.9	14.2
	8	12859.1	0.3	7805.3	25.6
	9	13946.1	0.5	8912.0	29.9
	10	15524.3	0.5	11228.1	42.0

Table C.6: Solution time of bounding techniques in seconds for the MSLot problem. For UB techniques, the average time of one sample path is given. For NA upper bounds with $T > 6$, the reported value is the average time over an evaluation sample of size 10.