

# Stochastic Dual Dynamic Programming and Its Variants: A Review\*

Christian Füllner<sup>†</sup>  
Steffen Rebennack<sup>†</sup>

**Abstract.** We provide a tutorial-style review of stochastic dual dynamic programming (SDDP), one of the state-of-the-art solution methods for large-scale multistage stochastic programs. Since it was introduced about 30 years ago for solving large-scale multistage stochastic linear programming problems in energy planning, SDDP has been applied to practical problems from several fields and has been enriched by various improvements and enhancements to address broader problem classes. We begin with a detailed introduction to SDDP, with special focus on its motivation, complexity, and required assumptions. Then, we present and discuss in depth the existing enhancements as well as current research trends that allow for the alleviation of those assumptions.

**Key words.** stochastic dual dynamic programming, dynamic programming, multistage stochastic programming, sequential decision problems, large-scale optimization, linear programming, nested Benders decomposition, sampling-based optimization, global optimization

**MSC codes.** 90-02, 90C15, 90C39

**DOI.** 10.1137/23M1575093

## Contents

<b>I</b>	<b>Introduction</b>	<b>418</b>
1.1	Structure . . . . .	420
1.2	Terminology and Notation . . . . .	420
<b>2</b>	<b>Preliminaries for SDDP</b>	<b>421</b>
2.1	Modeling the Uncertainty . . . . .	422
2.2	The Decision Process . . . . .	423
2.3	Single-Problem Formulation . . . . .	426
2.4	Dynamic Programming Equations . . . . .	427
2.5	Approximations of the Value Functions . . . . .	428
<b>3</b>	<b>Standard SDDP</b>	<b>430</b>
3.1	Main Principle . . . . .	430
3.2	Forward Pass . . . . .	431
3.3	Backward Pass . . . . .	431
3.4	Bounds and Stopping . . . . .	435
3.5	Cut Properties . . . . .	435

\*Received by the editors May 25, 2023; accepted for publication (in revised form) September 27, 2024; published electronically August 7, 2025.  
<https://doi.org/10.1137/23M1575093>

<sup>†</sup>Institute for Operations Research (IOR), Stochastic Optimization (SOP), Karlsruhe Institute of Technology (KIT), Germany (christian.fuellner@kit.edu, steffen.rebennack@kit.edu).

3.6 Illustrative Example . . . . . 436

3.7 Policy Assessment . . . . . 437

**4 Convergence and Complexity 438**

4.1 Finite Convergence . . . . . 438

4.2 Complexity . . . . . 440

**5 Comparison with Related Methods 441**

5.1 Relation to SDP . . . . . 441

5.2 Relation to Nested Benders Decomposition . . . . . 442

5.3 Complexity Comparison . . . . . 443

**6 Sampling 444**

6.1 Monte Carlo Sampling . . . . . 445

6.2 Variance Reduction Techniques . . . . . 445

6.3 Importance Sampling . . . . . 446

6.4 Deterministic Sampling . . . . . 447

**7 Stopping Criteria 448**

**8 Exact Upper Bounds and Upper Approximations 450**

8.1 Primal Inner Approximation . . . . . 451

8.2 Dual SDDP . . . . . 453

**9 Applications 455**

9.1 Power System Optimization . . . . . 455

9.1.1 Long-Term Operational Planning . . . . . 457

9.1.2 Medium-Term Operational Planning . . . . . 459

9.1.3 More Energy Applications . . . . . 459

9.2 Water Resource Management . . . . . 460

9.3 Portfolio Management . . . . . 460

9.4 Further Applications . . . . . 461

**10 Software 461**

**11 SDDP for Continuous Uncertainty [Relaxing Assumption 5] 461**

11.1 Sample Average Approximation (SAA) . . . . . 462

11.2 Assessing Policy Quality . . . . . 464

11.3 Variance Reduction Techniques . . . . . 466

**12 Risk-Averse SDDP [Relaxing Assumption 8] 467**

12.1 Risk Measures . . . . . 467

12.1.1 Static Risk Measures . . . . . 468

12.1.2 Multiperiod Risk Measures . . . . . 470

12.1.3 Time Consistency . . . . . 471

12.1.4 Polyhedral Risk Measures . . . . . 472

12.2 Toward Considering Risk in SDDP . . . . . 473

12.3 SDDP with Nested Risk Measures . . . . . 473

12.3.1 Reformulating the DPE . . . . . 474

12.3.2	Forward and Backward Passes . . . . .	477
12.3.3	Upper Bound Determination and Stopping . . . . .	478
12.4	SDDP with Entropic Risk Measure . . . . .	482
12.5	SDDP with Expected Conditional AVaR . . . . .	483
12.6	Biobjective SDDP . . . . .	484
<b>13</b>	<b>SDDP with Unknown Distribution [Relaxing Assumption 3]</b>	<b>485</b>
13.1	Distributionally Robust SDDP . . . . .	485
13.1.1	Reformulation as a Risk-Averse Problem . . . . .	487
13.1.2	Solving the Inner Maximization Problem Separately . . . . .	488
13.1.3	Using a Dual Representation . . . . .	488
13.2	Partially Observable Distributions . . . . .	489
<b>14</b>	<b>Stagewise-Dependent Uncertainty [Relaxing Assumption 2]</b>	<b>490</b>
14.1	Expanding the State Space . . . . .	490
14.2	Scenario-Adaptable Cut Formulas . . . . .	492
14.3	Sensitivity of SDDP with AR Processes . . . . .	495
14.4	SDDP with Markov Chain . . . . .	495
14.4.1	Modeling . . . . .	495
14.4.2	Adaptation of SDDP . . . . .	497
14.4.3	Specific Use Cases . . . . .	498
14.5	Hybrid NBD/SDDP . . . . .	500
14.6	Saddle Cuts . . . . .	500
14.7	Applying Dual SDDP . . . . .	502
14.8	Conditional Cuts . . . . .	502
<b>15</b>	<b>Extension to Convex Programs [Relaxing Assumption 6]</b>	<b>503</b>
<b>16</b>	<b>Extensions to Mixed-Integer and Nonconvex Problems [Relaxing Assumption 6]</b>	<b>506</b>
16.1	Convexification . . . . .	506
16.2	Exact Methods . . . . .	507
<b>17</b>	<b>Infeasible Subproblems [Relaxing Assumption 9]</b>	<b>510</b>
17.1	Feasibility Cuts . . . . .	510
17.2	Penalization . . . . .	511
<b>18</b>	<b>No Block-Diagonal Structure [Relaxing Assumption 7]</b>	<b>512</b>
<b>19</b>	<b>Infinite Horizon [Relaxing Assumption 1]</b>	<b>512</b>
<b>20</b>	<b>Random Horizon [Relaxing Assumption 1]</b>	<b>513</b>
<b>21</b>	<b>Performance Improvements</b>	<b>514</b>
21.1	Speeding up SDDP Iterations . . . . .	514
21.1.1	Cut Elimination and Selection . . . . .	515
21.1.2	Sampling Schemes . . . . .	516
21.1.3	Inexact SDDP . . . . .	517
21.2	Reducing the Number of SDDP Iterations . . . . .	519
21.2.1	Multicut SDDP . . . . .	520

21.2.2 Batch Learning and Experience Replay . . . . .	521
21.2.3 Regularization . . . . .	522
21.3 Parallelization . . . . .	524
21.4 Aggregation Techniques . . . . .	526
<b>22 Outlook</b>	<b>526</b>
<b>Acknowledgments</b>	<b>527</b>
<b>References</b>	<b>528</b>

**I. Introduction.** In many decision-making situations, at least some of the data are uncertain. While this uncertainty is often disregarded, the importance of taking it into account during the decision-making process was recognized as early as 1955 by George Dantzig [48]. In stochastic programming, a common approach to taking uncertainty into account is to split the decision-making process into two different stages. In the first stage, decisions have to be taken before any uncertain data are revealed to hedge against the existing uncertainty (so-called *here-and-now* decisions). In the second stage, corrective actions, called *recourse* or *wait-and-see* decisions, can be taken, once the realization of the uncertain data is known [27]. Typically, the aim is to determine an optimal decision rule *in expectation* or with respect to some risk measure.

In many practical applications, multiple rather than two subsequent decisions have to be made [7]. If these decisions cannot be made independently, but are coupled by their effects on a system state, e.g., hydroelectric generation affecting the water level of a reservoir, or orders affecting the size of an inventory stock, this can be modeled as a multistage stochastic problem with several subsequent recourse decisions (this is also referred to as *dynamic programming* and was recently coined a *sequential decision problem* in [177]). In such a problem, trade-offs have to be made between using an existing resource immediately or saving it up for later stages, taking into account the future uncertainty.

Stochastic dual dynamic programming (SDDP) is an algorithm that tackles such multistage stochastic problems in order to compute, or at least approximate, an optimal *policy*; that is, a strategy or decision rule providing the best here-and-now decision as well as the best wait-and-see decisions for any stage and any given realization of the uncertain data. It was first proposed by Pereira and Pinto in [160] in 1991.

Historically, SDDP has its roots in two separate research streams dealing with sequential decision problems. The first stream is *stochastic dynamic programming* (SDP), which is closely related to stochastic optimal control and Markov decision processes. Here, a crucial assumption is that the uncertain data in different stages of the decision process are independent of each other (or at least Markovian). In this case, multistage stochastic problems can be expressed by dynamic programming equations (DPE), which decompose the large-scale problem by stages into several smaller subproblems. These DPE exploit the famous optimality principle of Bellman [13], which allows one to express the optimal objective value from some stage  $t$  onwards, given some state  $x_{t-1}$ , recursively by means of some stage- $t$  objective function and a so-called *expected value function*  $\mathcal{Q}_t(\cdot)$ , modeling the expected optimal objective

value from stage  $t + 1$  onwards, given the new state  $x_t$ . We introduce these concepts formally in section 2.4.

The DPE can be solved exactly by SDP solution methods, such as value iteration [13]. Basically, this method is based on traversing the stages backwards and evaluating the expected value functions  $\mathcal{Q}_t(\cdot)$  for all possible states  $x_{t-1}$  (i.e., the concept of a lookup table). Each such evaluation requires solving an optimization problem for all possible realizations of the uncertain data, which, in turn, requires finding an optimal decision over all possible actions. For this evaluation to be possible, it is assumed that the state space, the action space, and the scenario space are finite—otherwise, they would have to be discretized. However, even in the discrete case, enumerating all possible combinations is computationally intractable for all but low dimensions, as the number of evaluations suffers from combinatorial explosion. This phenomenon is known as the *curse of dimensionality* of SDP [176]. In order to circumvent this issue, *approximate dynamic programming* (ADP) methods have been developed in which expected value functions are approximated instead of being evaluated exactly (or in which optimal policies are approximated using different strategies) [176, 177]. SDDP can be regarded as one such method. Due to its close relation with SDP it also relies heavily on the assumption of stagewise independence.

A second perspective on SDDP comes from *stochastic programming*. Traditionally, in this field, multistage uncertain data are often modeled by a scenario tree which branches at each stage and consists of finitely many possible scenarios. Scenario trees do not require the stochastic data process to be stagewise-independent. Using finite scenario trees and assuming linearity, a multistage stochastic program can be reformulated as a large-scale linear programming problem [185]. However, in this extensive form such a problem is usually far too large to be solved by monolithic approaches, since the number of decision variables and constraints grows exponentially in the number of stages. To cope with this challenge, special solution techniques are required which decompose the problem. Based on the L-shaped method for solving two-stage stochastic programs [236] (a special variant of Benders decomposition [17]), one such idea is the extension of Benders-type solution methods to the multistage setting. The *nested Benders decomposition* (NBD) method by Birge [25] is such an extension and can be interpreted as a nested sequence that solves two-stage stochastic programs while traversing the scenario tree. In contrast to SDP, in NBD the functions  $\mathcal{Q}_t(\cdot)$  are not evaluated at all possible states, but are iteratively approximated by linear functions called cutting planes or *cuts*, starting from a rough initial relaxation. Such an approximation is possible, since  $\mathcal{Q}_t(\cdot)$  can be proven to be convex in  $x_{t-1}$  for linear programs. It also allows for the consideration of a continuous state space without discretization.

While NBD is a reasonable method for solving multistage stochastic linear programs of moderate time horizons (maximum four or five time steps), for larger problems, it is still computationally prohibitive as the scenario tree grows exponentially in the number of stages. As a relief, several methods have been proposed to combine the cutting-plane approximations in NBD with sampling techniques from simulation [41, 57, 109]. The most prominent among these methods is SDDP. From this perspective, SDDP can be considered to be a sampling-based variant of NBD. In order to use the sampling step in a beneficial way, in comparison to NBD, SDDP comes with the additional prerequisite that the data process is stagewise-independent.

Application-wise, the development of SDDP is closely related to hydrothermal operational planning, which attempts to determine cost-optimal generation decisions for thermal and hydroelectric power plants over several stages, while ensuring system

balance and the satisfaction of technical constraints. Since future water availability is affected by uncertain inflows into hydro reservoirs, this optimization problem can be considered to be multistage and stochastic, and thus very complex.

Prior to SDDP, various solution techniques had been proposed to tackle this type of problem. Among these were simulation models, linear programming techniques (based on either assuming inflows as deterministic or reformulating stochastic linear programs into a deterministic equivalent), special variants of dynamic programming, and SDP [238]. However, all these techniques either do not consider the uncertain nature of inflows, or suffer from the aforementioned curses of dimensionality, or do not guarantee convergence. When operating a large-scale power system dominated by hydro power, these shortcomings are severe as they prohibit cost-minimal and reliable, but at the same time computationally efficient, operational planning. The development of SDDP by Pereira and Pinto was directly driven by the endeavor to replace SDP with a more efficient optimization technique in operating the Brazilian power system. While it avoids *some* of the computational drawbacks of SDP and NBD (sometimes advertised as “breaking the curse of dimensionality”), SDDP comes with its own shortcomings, as we discuss extensively in this article.

Since its invention in 1991 SDDP has gained enormous interest, from both theoretical and application perspectives. Today, it can be considered one of the state-of-the-art solution methods for large-scale multistage stochastic problems. For this reason, it is used in various practical applications to optimize decision processes, for instance, hydrothermal operational planning, portfolio optimization, and inventory management; see section 9.

Several extensions and improvements of SDDP have now been proposed, many of them attempting to relax the original theoretical assumptions required and make SDDP applicable to broader problem classes. Others strive to improve the performance of SDDP because, despite its merits, the algorithm can take too long to converge for large problem instances.

Due to both the sheer amount and the variety of proposed enhancements, SDDP has developed into a wide-ranging research area with several subbranches that become increasingly difficult to keep track of. In this article, we give a comprehensive tutorial-style review of SDDP-related research, covering its basic principles and assumptions, its strengths and weaknesses, existing extensions, and current research trends.

**1.1. Structure.** The structure of this review is summarized in Table 1. The review can be divided into four major parts. In the first part (sections 2 to 8), we discuss the basic mechanism of SDDP. This includes formal preliminaries to formulate multistage stochastic decision problems, but also the main algorithmic steps of SDDP and a complexity analysis. In particular, we point out crucial assumptions for standard SDDP to work. In the second part (sections 9 and 10), we discuss applications which underline the practical relevance of SDDP, but also the requirement to relax some of the standard assumptions. In the third part (sections 11 to 20), we discuss various extensions of SDDP to cases where the standard assumptions are relaxed. These extensions comprise modifications of SDDP itself as well as modifications or reformulations of the considered decision problems. Finally, in the fourth part (section 21), we discuss approaches to improve the computational performance of SDDP.

**1.2. Terminology and Notation.** As mentioned above, SDDP is linked to several different research fields and communities such as stochastic programming, dynamic programming, Markov decision processes, optimal control, and reinforcement learning,

Table I Table of contents.

<b>Part I: The Mechanism of SDDP</b>		
Section 2	Preliminaries	p. 421
Section 3	Standard SDDP	p. 430
Section 4	Convergence & Complexity	p. 438
Section 5	Comparison with Related Methods	p. 441
Section 6	Sampling	p. 444
Section 7	Stopping Criteria	p. 448
Section 8	Exact Upper Bounds and Dual SDDP	p. 450
<b>Part II: Applications of SDDP</b>		
Section 9	Applications	p. 455
Section 10	Software	p. 461
<b>Part III: Extensions of SDDP</b>		
Section 11	Handling Continuous Uncertainty	p. 461
Section 12	Handling Risk-Aversion	p. 467
Section 13	Handling Distributional Uncertainty	p. 485
Section 14	Handling Stagewise-Dependent Uncertainty	p. 490
Section 15	Handling Nonlinear Convex Problems	p. 503
Section 16	Handling Mixed-Integer and Nonconvex Problems	p. 506
Section 17	Handling Infeasibility	p. 510
Section 18	Handling Non-block-diagonal Problems	p. 512
Section 19	Handling Infinite Horizons	p. 512
Section 20	Handling Random Horizons	p. 513
<b>Part IV: Accelerating SDDP</b>		
Section 21	Performance Improvements	p. 514
Section 22	Outlook	p. 526

each using different terminology and notation. This complicates the presentation of SDDP in a form that is intended to be familiar and accessible to all those interested.

To our knowledge, most active research on SDDP is conducted by researchers from the stochastic programming community. For this reason, in many sections we resort to stochastic programming language and notation. However, this review also aims to offer access to SDDP for practitioners and researchers from fields in which different perspectives and notation are standard. Therefore, we address these differences if required for understanding SDDP and attempt to avoid heavy mathematical programming notation whenever possible, especially in the early sections that introduce SDDP.

For a general non-SDDP specific attempt at unifying different disciplines related to optimization under uncertainty and sequential decision processes into a common framework, we refer to [177].

In what follows, we denote random variables by bold letters, e.g.,  $\boldsymbol{\xi}$ , and their realizations by letters in regular font, e.g.,  $\xi$ . To enhance readability, we summarize some recurring acronyms in Table 2.

**2. Preliminaries for SDDP.** In order to present SDDP in its standard form, we start by formally introducing the decision problem considered here. In particular, we point out assumptions that are crucial for the SDDP method presented to work.

We consider a multistage decision process in which decisions  $x_t$  have to be taken over some horizon  $[T] := \{1, \dots, T\}$  consisting of  $T$  stages, with the aim to minimize some objective function subject to constraints. For now, the horizon  $T$  is assumed to satisfy the following assumption.

**Table 2** Acronyms that are used throughout the text.

(P)AR	(Periodic) Autoregressive process
DPE	Dynamic programming equations
LP	Linear program
MI(N)LP	Mixed-integer (non)linear program
MSLP	Multistage stochastic linear programming problem
NBD	Nested Benders decomposition
RHS	Right-hand side
SDP	Stochastic dynamic programming
SDDP	Stochastic dual dynamic programming

*Assumption 1* (finite and deterministic horizon). The number  $T \in \mathbb{N}$  of stages is finite and deterministic.

In what follows, we discuss how SDDP may be applied to cases where this assumption is not satisfied; see sections 19 and 20.

**2.1. Modeling the Uncertainty.** The data in the decision process considered can be subject to uncertainty that is revealed over time. To this end, we consider a filtered probability space  $(\Omega, \mathcal{F}, \mathbb{P})$  with sample space  $\Omega$ ,  $\sigma$ -algebra  $\mathcal{F}$ , and probability measure  $\mathbb{P}$ , which models the uncertainty over the horizon  $[T]$ . Further, let  $\mathcal{F}_1, \dots, \mathcal{F}_T$  with  $\mathcal{F}_T := \mathcal{F}$  be a sequence of  $\sigma$ -algebras containing the events observable up to time  $t$ , thus defining a filtration with  $\mathcal{F}_1 \subseteq \mathcal{F}_2 \cdots \subseteq \mathcal{F}_T$ , and let  $\Omega_t$  be the sample space restricted to stage  $t \in [T]$ . We then define a stochastic process  $(\boldsymbol{\xi}_t)_{t \in [T]}$  with random vectors  $\boldsymbol{\xi}_t : \Omega_t \rightarrow \mathbb{R}^{\kappa_t}$ ,  $\kappa_t \in \mathbb{N}$ , over the probability space. These random vectors are assumed to be  $\mathcal{F}_t$ -measurable functions. We denote their support by  $\Xi_t \subseteq \mathbb{R}^{\kappa_t}$  for all  $t \in [T]$ . For the first stage, the data are assumed deterministic, i.e.,  $\Xi_1$  is a singleton. For each random vector  $\boldsymbol{\xi}_t$ , we denote a specific realization by  $\xi_t$ .

As a crucial ingredient for SDDP to work, we assume that the uncertainties on different stages do not depend on one another.

*Assumption 2* (stagewise independence). For all  $t \in [T]$ , the random vector  $\boldsymbol{\xi}_t$  is independent of the history  $\xi_{[t-1]} := (\xi_1, \dots, \xi_{t-1})$  of the data process.

Under Assumption 2, the random vectors  $\boldsymbol{\xi}_t$  are often referred to as *noises*. This assumption is common in dynamic programming, but is not standard in stochastic programming. In practical applications it may not be satisfied. We address how to apply SDDP to problems with stagewise-dependent uncertainty in section 14.

In addition, we make the following assumptions for the stochastic process.

*Assumption 3* (known distribution). The probability distribution  $F_{\boldsymbol{\xi}}$  of the data process  $(\boldsymbol{\xi}_t)_{t \in [T]}$  is known.

*Assumption 4* (exogeneity). The random variables  $\boldsymbol{\xi}_t$  are exogenous, i.e., the distribution  $F_{\boldsymbol{\xi}}$  of the data process  $(\boldsymbol{\xi}_t)_{t \in [T]}$  is independent of decisions  $(x_t)_{t \in [T]}$ .

*Assumption 5* (finite randomness). The support  $\Xi_t$  of  $\boldsymbol{\xi}_t$  is finite for all  $t \in [T]$ . The number of noise realizations at stage  $t \in [T]$  is given by  $q_t \in \mathbb{N}$  with  $q_1 = 1$ .

In section 13 we discuss how to apply SDDP if Assumption 3 is not satisfied. If Assumption 4 is not satisfied, the problem is said to have *decision-dependent* uncertainty [122]. As this case has not been covered in the literature on SDDP up to this point, we do not discuss the relaxation of this assumption.



Assumption 5 is a key assumption for SDDP and is standard in dynamic programming and stochastic programming in order to obtain computationally tractable problems. We discuss possible ways to treat such problems in section 11. As  $\xi_t$  is a discrete and finite random variable for all  $t \in [T]$ , its distribution  $F_\xi$  is defined by finitely many realizations  $\xi_{tj}, j = 1, \dots, q_t$ , and assigned probabilities  $p_{tj}$ .

The stagewise-independent and finite data process  $(\xi_t)_{t \in [T]}$  can be illustrated by a *recombining scenario tree* [185], also called *scenario lattice* [136]. On each stage  $t \in [T]$ , its nodes represent the possible noise realizations  $\xi_{tj}, j = 1, \dots, q_t$ . Due to stagewise independence (Assumption 2), all nodes at the same stage have an identical set of child nodes with the same noise realizations and associated probabilities. We call paths  $\xi = (\xi_t)_{t \in [T]}$  through the complete tree (stage- $T$ ) *scenarios* and index them by  $s \in \mathcal{S}$ . Note that, for each scenario  $\xi^s$ , there exists some  $j_s \in \{1, \dots, q_t\}$  such that  $\xi_t^s = \xi_{tj_s}$ . The total number of different scenarios modeled by the tree is  $|\mathcal{S}| = \prod_{t \in [T]} q_t$ . An example of a recombining scenario tree is presented in Figure 1.

**2.2. The Decision Process.** With the stochastic process in mind, we can now turn to the decision process. At stage 1, the *here-and-now* decision  $x_1$  is taken to hedge against the uncertainty in the following stages. At those stages, recourse decisions  $x_t \in \mathbb{R}^{n_t}, n_t \in \mathbb{N}$ , can be taken with knowledge of the realization of the data process at stage  $t$ . This decision process is illustrated in Figure 2.

In other words, the paradigm is that decisions can be taken *after* the uncertainty corresponding to stage  $t$  has unfolded (so-called *wait-and-see* decisions), making  $x_t(\xi_t)$  a function of  $\xi_t$ , and by that a random variable. We account for this using a bold symbol. Importantly,  $x_t(\cdot)$  only depends on realizations up to stage  $t$ , but does not anticipate future events or decisions. Future events are only considered using distributional information. Therefore,  $x_t(\cdot)$  is  $\mathcal{F}_t$ -measurable [208]. As we will see,  $x_t(\cdot)$  may also depend on the choice of  $x_{t-1}(\cdot)$  and so on, so that despite stagewise independence (Assumption 2),  $x_t(\cdot)$  is actually a function of the whole history  $\xi_{[t]}$  of the data process.

A sequence of decision functions  $(x_t(\xi_{[t]}))_{t \in [T]}$  is called a *policy* and provides a decision rule for all stages  $t \in [T]$  and any realization of the data process. By the previous arguments, such a policy is *nonanticipative*, modeling a sequence of nested conditional decisions. The aim of the decision process is to determine an *optimal* policy with respect to a given objective function and a given set of constraints.

In this context, the following assumptions are standard for SDDP.

*Assumption 6* (linearity). All functions occurring in the objective and the constraints are linear.

*Assumption 7* (consecutive coupling). Only decisions on consecutive stages can be linked by constraints.





Fig. 1 Recombining tree with three realizations per stage and highlighted scenario  $\xi^6$ .



**Fig. 2** Multistage decision process with uncertainty.

*Assumption 8* (risk-neutral policy). The aim is to determine an optimal risk-neutral policy.

As not all of these assumptions are guaranteed to be satisfied for an arbitrary problem in practice, we discuss possible ways to relax them in sections 15 and 16 (for Assumption 6), section 18 (for Assumption 7), and section 12 (for Assumption 8).

Under Assumptions 6 and 8, the optimization objective can be expressed as

$$(2.1) \quad \min_{x_1, x_2, \dots, x_T} \mathbb{E} \left[ \sum_{t \in [T]} (c_t(\xi_t))^\top x_t(\xi_t) \right],$$

with data vectors  $c_t \in \mathbb{R}^{n_t}$  for all  $t \in [T]$  and  $\mathbb{E}[\cdot]$  denoting the expected value.

Under Assumptions 6 and 7, for all  $t \in [T]$ , the constraints on the decisions can be expressed using the  $\mathcal{F}_t$ -measurable set-valued mappings  $\mathcal{X}_t(\cdot)$ , which for any  $x_{t-1}$  and any  $\xi_t \in \Xi_t$  are defined by

$$(2.2) \quad \mathcal{X}_t(x_{t-1}, \xi_t) := \left\{ x_t \in X_t \subset \mathbb{R}^{n_t} \mid T_{t-1}(\xi_t)x_{t-1} + W_t(\xi_t)x_t = h_t(\xi_t) \right\}.$$

Here,  $h_t \in \mathbb{R}^{m_t}$  are real data vectors (for  $m_t \in \mathbb{N}$ ),  $T_t$  and  $W_t$  are real-valued ( $m_{t+1} \times n_t$ ) and ( $m_t \times n_t$ ) data matrices, and  $X_t$  is a nonempty polyhedron, e.g., modeling nonnegativity constraints.

As stated above, some (or all) of the problem data can be subject to uncertainty. Hence, for all  $t \in [T]$ , we consider random variables  $c_t(\xi_t)$ ,  $T_{t-1}(\xi_t)$ ,  $W_t(\xi_t)$ , and  $h_t(\xi_t)$  depending on realizations of  $\xi_t$ .  $X_t$  is considered deterministic. Note again that the first stage is assumed to be deterministic and that  $T_0 \equiv 0$  and  $x_0 \equiv 0$ . Hence, we define  $\mathcal{X}_1 \equiv \mathcal{X}_1(x_0, \xi_1)$ .

*Remark 2.1.* For notational simplicity, when we deal with finite random variables  $\xi_t$ , we often index the vectors and matrices  $c_t$ ,  $T_{t-1}$ ,  $W_t$ , and  $h_t$  with  $j = 1, \dots, q_t$  if we address specific realizations, e.g.,  $c_{tj} := c_t(\xi_{tj})$ .

*Remark 2.2* (dynamic programming perspective). In dynamic programming, Markov decision processes, or optimal control, a slightly different perspective on sequential decision processes is usually chosen (see [177] for a comprehensive overview). The main difference is that the variables that occur are split into *state variables* and actual *decisions*. State variables  $s_t \in S_t$  model the system state at some stage  $t$ .  $S_t$  is called the *state space*. Importantly, state variables may comprise not only the resource state, but also the information or belief state of a system [177]. Decision variables model local decisions on a stage  $t$  given a state  $s_t$ . In dynamic programming they are usually discrete and called *actions*  $a_t \in A_t(s_t)$ ; in optimal control they are usually continuous and called *controls*  $u_t \in U_t(s_t)$ .  $A_t(s_t)$  and  $U_t(s_t)$  are the *action space* and

*control space*, respectively. The actions or controls are what an agent actually decides on given the current state  $s_t$ , whereas the new state  $s_{t+1}$  is uniquely determined as  $s_{t+1} = \mathcal{T}_t(s_t, u_t, \xi_{t+1})$  using a given transition function  $\mathcal{T}_t(\cdot)$  that captures the system dynamic. Therefore, from this perspective, a policy is a sequence of mappings  $\pi_t: S_t \rightarrow U_t$  from the state space to the control (or action) space. By proper modeling of the state variable, Assumption 7 is naturally satisfied.

In our above setting, states and actions are intertwined. We can set  $s_t = (x_{t-1}, \xi_t)$  and  $u_t = x_t$  to switch perspectives [6]. The state space, control space, and transition function are then implicitly given by (2.2) and the definition of  $\xi_t$ .

While our definitions above are prevalent in the literature on SDDP, sometimes an optimal control perspective is also adopted, e.g., in the French community working on SDDP (see, for example, [85]). However, in this case usually only the resource state  $r_t$  is explicitly considered as a state variable (while not including information on  $\xi_t$ ). Translating our above setting, this implies that  $r_t = x_{t-1}$  with state space  $R_t = X_t$ ,  $u_t = x_t$ , and, due to  $r_{t+1} = u_t$ , both the control space  $U_t(r_t, \xi_t)$  and the transition function  $\mathcal{T}_t(r_t, u_t, \xi_t)$  are given by the equations in (2.2).

It is worth mentioning that the distinction between state variables and controls (actions) is not only a matter of notation, but is also computationally relevant because the complexity of SDDP differs in the state and control dimensions (see also Remark 2.6 and subsection 4.2).

Given the constraint sets (2.2) for all  $t \in [T]$ , let  $\mathcal{X}_0 := \{x_0\}$  and recursively define

$$\mathcal{X}_t := \bigcup_{x_{t-1} \in \mathcal{X}_{t-1}} \bigcup_{\xi_t \in \Xi_t} \mathcal{X}_t(x_{t-1}, \xi_t)$$

for all  $t \in [T]$  [75]. Using these definitions, we are able to state assumptions required for the feasibility of our decision problem.

*Assumption 9* (feasibility and compactness).

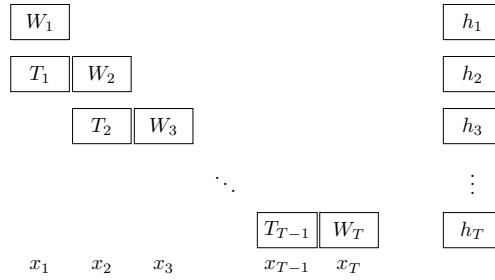
- (a) For all  $t \in [T]$ , all  $x_{t-1} \in \mathcal{X}_{t-1}$ , and almost all  $\xi_t \in \Xi_t$ , the set  $\mathcal{X}_t(x_{t-1}, \xi_t)$  is a nonempty compact subset of  $\mathbb{R}^{n_t}$  (relatively complete recourse).
- (b) The set  $\mathcal{X}_t$  is bounded for all  $t \in [T]$ .

*Remark 2.3.* Note that the linearity assumption (see Assumption 6) immediately implies that Assumption 9 (a) is satisfied not only for all  $x_{t-1} \in \mathcal{X}_{t-1}$ , but also for all  $x_{t-1} \in \text{conv}(\mathcal{X}_{t-1})$ , where  $\text{conv}(S)$  denotes the convex hull of a set  $S$ .

The set  $\mathcal{X}_t \in \mathbb{R}^{n_t}$  is called a *reachable set* in [75] and an *effective feasible region* in [124]. It may also sometimes be referred to as the state space, because in our setting  $x_t$  also takes the role of a state variable. However, in other cases the larger polyhedral set  $X_t$  may be called a state space.

The boundedness of  $\mathcal{X}_t$  in Assumption 9 (b) is required for some of the convergence results on SDDP presented in section 4. It follows naturally if  $X_t$  is bounded, since  $\mathcal{X}_t \subseteq X_t$ . Property (a) is convenient but not necessarily required and we discuss possible ways to relax it in section 17.

With all the ingredients defined, we can now model the decision problem in a form that can be tackled by SDDP. Based on its properties, in what follows we refer to this problem as a *multistage stochastic linear programming problem* “(MSLP)”. If not specified otherwise, throughout this article, we assume that (MSLP) satisfies Assumptions 1 to 9. We first discuss two different modeling approaches that are common in the literature.



**Fig. 3** Block-diagonal structure of constraints in (MSLP).

**2.3. Single-Problem Formulation.** One way to model the decision problem (MSLP) is to formulate it as a single optimization problem. This modeling approach is common in the stochastic programming community. The optimization problem can be obtained by combining (2.1) with the constraints in (2.2) for all  $t \in [T]$ .

Then, under Assumptions 1 to 9, (MSLP) can be written as

$$(2.3) \quad v^* := \begin{cases} \min_{x_1, x_2, \dots, x_T} & \mathbb{E} \left[ \sum_{t \in [T]} (c_t(\xi_t))^\top x_t(\xi_{[t]}) \right] \\ \text{s.t.} & x_1 \in \mathcal{X}_1, \\ & x_t(\xi_{[t]}) \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t) \quad \forall \xi_{[t]} \quad \forall t = 2, \dots, T, \\ & x_t(\cdot) \text{ } \mathcal{F}_t\text{-measurable} \quad \forall t = 2, \dots, T. \end{cases}$$

Importantly, the decision variables  $x_t \in \mathbb{R}^{n_t}$  depend on  $\xi_t$  (and on  $x_{t-1}$ ), so in this representation we optimize over policies. A policy  $(x_t(\xi_{[t]}))_{t \in [T]}$  is called *feasible* (or *admissible*) if it satisfies the constraints in (MSLP) for almost every realization of the random data [208].

Assumption 9 (a) implies that the feasible set of (MSLP) is compact and nonempty, and by linearity of the objective (Assumption 6) it follows that  $v^*$  is finite.

Due to its optimizing over policies, without Assumption 5, (MSLP) is an infinite-dimensional optimization problem. With Assumption 5, however, it can be reformulated in a more accessible form. More precisely, it can be reformulated as a large-scale deterministic problem, the so-called *deterministic equivalent* of (MSLP) in *extensive form* (see [208]). To this end, let  $\mathcal{S}$  denote the set of all (stage- $T$ ) scenarios. Then, for each scenario  $s \in \mathcal{S}$  a separate copy  $x_t^s$  of variables  $x_t$  can be introduced, so that the optimization over implementable policies translates to an optimization over a finite number of decision variables. However, the problem size grows exponentially in the number of stages  $T$ . Therefore, even for a finite number of scenarios, this large-scale linear program (LP) is too large to be solved by off-the-shelf solvers for all but very small instances.

A preferable solution approach is therefore to use tailored solution techniques that decompose (MSLP) into smaller subproblems. Note that from Assumption 7 and the definition of  $\mathcal{X}_t(\cdot)$  in (2.2), it is evident that the constraints of (MSLP) are block-diagonal, as only consecutive stages are coupled in the constraints. This is visualized in Figure 3.

This sequential and block-diagonal structure can be exploited to achieve the required decomposition. This is crucial for the derivation of SDDP. Interestingly, this

decomposition idea leads directly to the second common modeling approach for our decision problem.

**2.4. Dynamic Programming Equations.** An alternative, but equivalent way to model (MSLP) is to exploit the well-known optimality principle of Bellman [13] and to formulate a recursion of so-called *dynamic programming equations* (DPE), where a multistage decision process with stagewise-independent (or Markovian) uncertainty is modeled as a coupled sequence of optimization problems.

Whereas this modeling approach is often applied in stochastic programming as a way to reformulate and decompose the single problem (2.3) into a computationally tractable form, in dynamic programming it often serves as the starting point of modeling decision problems. However, in contrast to many approaches in dynamic programming, we do not discretize  $x_t$ ; see also section 5.1.

Under Assumptions 1 to 9, for  $t = T, \dots, 2$ , the DPE are given by

$$(2.4) \quad Q_t(x_{t-1}, \xi_t) := \begin{cases} \min_{x_t} & (c_t(\xi_t))^\top x_t + \mathcal{Q}_{t+1}(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t), \end{cases}$$

where

$$(2.5) \quad \mathcal{Q}_{t+1}(x_t) := \mathbb{E}_{\xi_{t+1}} [Q_{t+1}(x_t, \xi_{t+1})]$$

and  $\mathcal{Q}_{T+1}(x_T) \equiv 0$ .  $Q_t(\cdot, \cdot)$  is called the *value function* and  $\mathcal{Q}_t(\cdot)$  is called the *expected value function*, (*expected*) *cost-to-go function*, *future cost function*, or *recourse function*. For the first stage, we obtain

$$(2.6) \quad v^* = \begin{cases} \min_{x_1} & c_1^\top x_1 + \mathcal{Q}_2(x_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

For a formal proof of the equivalence of (2.3) and its DPE, we refer to [208] and section 12. Importantly, in subproblem (2.4),  $x_t$  is a deterministic variable and not a function because a fixed realization of  $\xi_t$  is considered.

We should emphasize that the equivalence of (2.3) and its DPE does not require Assumption 5. This also implies that the DPE (2.4)–(2.6) are computationally intractable in the case of general continuous random variables. While the subproblems are deterministic and finite-dimensional, there exist infinitely many value functions  $Q_t(\cdot, \cdot)$  and the evaluation of  $\mathcal{Q}_t(\cdot)$  requires the evaluation of (multidimensional) integrals. Therefore, from this perspective as well, Assumption 5 is crucial.

*Remark 2.4* (dynamic programming control perspective). Recall Remark 2.2. Using a distinction between state variables  $r_t$  and controls  $u_t$ , the DPE to (MSLP) can be formulated as

$$(2.7) \quad Q_t(r_t, \xi_t) = \min_{u_t \in U_t(r_t, \xi_t)} f_t(u_t, \xi_t) + \mathcal{Q}_{t+1}(\mathcal{T}_t(r_t, u_t, \xi_t)).$$

**Bellman Operator.** In the French literature on SDDP, in addition to taking the optimal control perspective discussed in Remarks 2.2 and 2.4, a more formal way to define the DPE is prevalent; see [75, 126], for instance. To this end, a linear Bellman operator  $\mathfrak{B}_t$  is introduced, which applied to some lower semicontinuous function  $V : \mathbb{R}^{n_t} \rightarrow \mathbb{R} \cup \{+\infty\}$  is defined as [75]

$$(2.8) \quad \widehat{\mathfrak{B}}_t(V)(x_{t-1}, \xi_t) := \min_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} (c_t(\xi_t))^\top x_t + V(x_t),$$

i.e., it maps  $(x_{t-1}, \xi_t)$  to the optimal value of an optimization problem containing function  $V(\cdot)$ . We can then further define a second operator as

$$(2.9) \quad \mathfrak{B}_t(V)(x_{t-1}) := \mathbb{E} \left[ \widehat{\mathfrak{B}}_t(V)(x_{t-1}, \xi_t) \right].$$

Setting  $V$  to  $\mathcal{Q}_t(\cdot)$  for  $t = 2, \dots, T$ , the (expected) value functions can then be recursively defined in a very compact form. We summarize the different notations for a better overview:

$$\begin{array}{l} \widehat{\mathfrak{B}}_t(\mathcal{Q}_{t+1})(x_{t-1}, \xi_t) = Q_t(x_{t-1}, \xi_t) \\ \mathfrak{B}_t(\mathcal{Q}_{t+1})(x_{t-1}) = \mathcal{Q}_t(x_{t-1}) \end{array}$$

In the remainder of this work, we stick to notation (2.4), as it is most common in the literature on SDDP referenced in this article.

We obtain the following properties for the DPE which are standard for SDDP.

**LEMMA 2.5.** *Under Assumptions 1 to 9, for the DPE defined by (2.4)–(2.6), the following properties hold:*

- (a) *We have relatively complete recourse, i.e., for any  $x_{t-1} \in \mathcal{X}_{t-1}$ , the stage- $t$  subproblem (2.4) is feasible for all  $\xi_t \in \Xi_t$ .*
- (b) *The value functions  $Q_t(\cdot, \cdot)$  and expected value functions  $\mathcal{Q}_t(\cdot)$  are finite-valued on  $\text{conv}(\mathcal{X}_{t-1})$  for all  $t = 2, \dots, T$  and all  $\xi_t \in \Xi_t$ .*
- (c) *Problem (2.6) is feasible and bounded.*

**Remark 2.6.** In addition to Remark 2.2, we should highlight that (MSLP) (in both single-problem formulation (2.3) and DPE (2.4)–(2.6)) can be straightforwardly enhanced with local decision variables  $y_t \in Y_t$  and local constraints that do not appear in different stages. In principle, they can even be incorporated without changes to our models by extending the dimension of the (state) variables  $x_t$  and adapting the matrices  $T_t$  and  $W_t$  accordingly. However, as we explain in section 4, the complexity of SDDP grows exponentially in the dimension of the state space, so this is computationally detrimental and should be avoided. Instead, purely local variables and constraints should be handled separately from those introduced above. This approach is referred to as *generalized dual dynamic programming* (GDDP) in [18].

While almost every practical application will require the introduction of these additional elements, in this work, we generally restrict their addition to coupling variables and constraints that are required to illustrate the mechanics of SDDP.

**Remark 2.7.** In general, the local objective functions may also include the states  $x_{t-1}$  instead of only depending on  $x_t$  and  $\xi_t$ . For notational simplicity, we consider a less general form of the objective function.

**2.5. Approximations of the Value Functions.** The main challenge in exploiting the DPE to solve (MSLP) is that the (expected) value functions are not known in analytical form in advance. The key idea in SDDP is to iteratively approximate them from below using linear functions, which are called *cutting planes* or short *cuts*. Together, these linear functions build polyhedral outer approximations  $\underline{\mathcal{V}}_t(\cdot)$  of  $\mathcal{Q}_t(\cdot)$

for all  $t = 2, \dots, T$ , which we refer to as *cut approximations*. In that regard, SDDP can be considered as a special variant of Kelley's cutting-plane method [118] and closely related to Benders decomposition [17]; see also section 5.2. Note that in contrast to SDP, this avoids a state discretization as  $Q_t(\cdot, \cdot)$  and  $\mathcal{Q}_t(\cdot)$  do not have to be evaluated at all possible states, but only at well-chosen trial points where new cuts are constructed; see section 5.1.

For this approximation by cuts, the following properties are crucial.

**THEOREM 2.8** ([27]). *Let  $x_{t-1} \in \text{conv}(\mathcal{X}_{t-1})$ . Then, under Assumptions 1 to 9, for all  $t = 2, \dots, T$  and a given noise realization  $\xi_t$ , the value function  $Q_t(\cdot, \xi_t)$*

- (a) *is piecewise linear and convex in  $(h_t, T_{t-1})$ ,*
- (b) *is piecewise linear and concave in  $c_t$ ,*
- (c) *is piecewise linear and convex in  $x_{t-1}$  on  $\text{conv}(\mathcal{X}_{t-1})$ .*

The main idea here is that given the definition of  $\mathcal{X}_{t-1}(\cdot)$  in (2.2),  $h_t$ ,  $T_{t-1}$ , and  $x_{t-1}$  only appear in the right-hand side (RHS) of problem (2.4). Therefore, the dual feasible set is independent of those elements and possesses finitely many extreme points. This ensures piecewise linearity of  $Q_t(\cdot, \cdot)$ , as is known from parametric optimization. The convexity follows from the linearity (Assumption 6) and the fact that all vectors and matrices are part of convex sets.

Theorem 2.8 directly implies the piecewise linearity and convexity of  $\mathcal{Q}_t(\cdot)$ .

**COROLLARY 2.9** ([27]). *Under Assumption 5 and the premises of Theorem 2.8, for all  $t = 2, \dots, T$ ,  $\mathcal{Q}_t(\cdot)$  is piecewise linear and convex in  $x_{t-1}$  on  $\text{conv}(\mathcal{X}_{t-1})$ .*

Theorem 2.8 and Corollary 2.9 also directly imply the Lipschitz continuity of the (expected) value functions.

**COROLLARY 2.10.** *Under Assumptions 1 to 9, for all  $t = 2, \dots, T$  and all  $\xi_t \in \Xi_t$ ,  $Q_t(\cdot, \xi_t)$  and  $\mathcal{Q}_t(\cdot)$  are Lipschitz continuous on  $\text{conv}(\mathcal{X}_{t-1})$ .*

Replacing the true expected value functions with cut approximations in (2.4), we can define *approximate value functions*

$$(2.10) \quad \underline{Q}_t(x_{t-1}, \xi_t) := \begin{cases} \min_{x_t} & (c_t(\xi_t))^\top x_t + \underline{V}_{t+1}(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t). \end{cases}$$

Trivially, for  $\mathcal{Q}_{T+1}(\cdot) \equiv 0$ , we have  $\underline{V}_{T+1}(\cdot) \equiv 0$ .

Note that apart from  $x_{t-1}$  and  $\xi_t$ ,  $\underline{Q}_t(\cdot, \cdot)$  is also a function of the cut approximation  $\underline{V}_{t+1}(\cdot)$ . This is especially relevant when these approximations are iteratively updated in SDDP, leading to different approximate value functions. Using the Bellman operators defined in (2.8)–(2.9) this can be expressed in a very concise way:

$$\underline{Q}_t(\cdot, \cdot) = \mathfrak{B}_t(\underline{V}_{t+1})(\cdot, \cdot).$$

Similarly, we could express this by adding an argument to  $\underline{Q}_t(\cdot, \cdot)$ , i.e., by writing  $\underline{Q}_t(x_{t-1}, \xi_t | \underline{V}_{t+1})$  or  $\underline{Q}_t(\underline{V}_{t+1})(x_{t-1}, \xi_t)$ . However, for notational simplicity we do not state this explicitly, but when dealing with SDDP we use the iteration index  $i$  for distinction. This means that  $\underline{Q}_t^i(\cdot, \cdot)$  indicates that  $\underline{Q}_t(\cdot, \cdot)$  is considered with cut approximation  $\underline{V}_{t+1}^i(\cdot)$ .

We summarize the different notations below for a better overview:

$$(2.11) \quad \boxed{\begin{aligned} \widehat{\mathfrak{B}}_t(\underline{\mathcal{V}}_{t+1})(x_{t-1}, \xi_t) &= \underline{Q}_t(x_{t-1}, \xi_t) \\ \mathfrak{B}_t(\underline{\mathcal{V}}_{t+1})(x_{t-1}) &= \underline{\mathcal{Q}}_t(x_{t-1}) := \mathbb{E}_{\xi_t} [\underline{Q}_t(x_{t-1}, \xi_t)] \end{aligned}}$$

Finally, we observe that given that the cut approximations  $\underline{\mathcal{V}}_{t+1}(\cdot)$  are polyhedral, the approximate value functions  $\underline{Q}_t(\cdot, \cdot)$  inherit the previously stated properties from  $Q_t(\cdot, \cdot)$ . In particular, the following lemma can be stated.

**LEMMA 2.11.** *Let  $\underline{\mathcal{V}}_{t+1}(\cdot)$  be a polyhedral function. Then, under Assumptions 1 to 9, for all  $t = 2, \dots, T$  and a given noise realization  $\xi_t$ ,  $\underline{Q}_t(\cdot, \xi_t)$  is piecewise linear and convex in  $x_{t-1}$  on  $\text{conv}(\mathcal{X}_{t-1})$ .*

On the other hand, as they are polyhedral, the cut approximations  $\underline{\mathcal{V}}_t(\cdot)$  for  $t = 2, \dots, T$  are nonlinear functions. Importantly for computations, subproblems (2.10) can still be formulated as LPs by using a partial epigraph reformulation and the fact that  $\underline{\mathcal{V}}_t(\cdot)$  is defined as the maximum of finitely many affine functions (modeled by some set  $\mathcal{K}$  with  $|\mathcal{K}| \in \mathbb{N}$ ):

$$(2.12) \quad \underline{Q}_t(x_{t-1}, \xi_t) = \begin{cases} \min_{x_t, \theta_{t+1}} & (c_t(\xi_t))^\top x_t + \theta_{t+1} \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t), \\ & -(\beta_{t+1,k}^i)^\top x_t + \theta_{t+1} \geq \alpha_{t+1,k}^i \quad \forall i \forall k \in \mathcal{K}. \end{cases}$$

This LP contains an additional decision variable  $\theta_{t+1}$  and finitely many additional linear constraints indexed by  $i$  and  $k$ . The structure and indexing of these constraints become clear in the next section where we present the cut generation process for SDDP.

**3. Standard SDDP.** We are now able to introduce SDDP in its standard form.

**3.1. Main Principle.** SDDP consists of two main steps in each iteration  $i$ , a *forward pass* and a *backward pass* through the stages  $t \in [T]$ .

In each forward pass, using the approximate value functions  $\underline{Q}_t^i(\cdot, \cdot)$  (recall that this implies using cut approximation  $\underline{\mathcal{V}}_{t+1}^i(\cdot)$  in (2.10)), a sequence of *trial points*  $(x_t)_{t \in [T]}$  is generated at which new cuts are then constructed in the following backward pass to improve the approximation. These trial points are also called *incumbents* or *candidate solutions*, and their sequence is called a *state trajectory* (especially in optimal control). The idea behind this approach is that the approximate value functions implicitly define a feasible (suboptimal) policy for problem (MSLP). The trial points are generated by evaluating this policy for one or several scenarios which are sampled from  $\mathcal{S}$ , i.e., by solving the respective subproblems. This has the advantage that cuts are constructed at points which (at least for some scenario) are optimal given the current cut approximation. This step can also be interpreted as a Monte Carlo *simulation* of the current policy.

In the backward pass, subgradients of  $\underline{\mathcal{Q}}_t(\cdot)$  at the trial points are used to construct cuts, passing them back to the previous stage and updating  $\underline{\mathcal{V}}_t(\cdot)$  to  $\underline{\mathcal{V}}_t^{i+1}(\cdot)$  for all  $t = 2, \dots, T$ . In this way, if not optimal, the current policy is amended (at least if the *right* scenario is sampled). In this step, a true lower bound  $\underline{v}$  for  $v^*$  is also determined.



*Remark 3.1* (statistical learning perspective). The basic principle of SDDP can also be interpreted from the perspective of supervised learning as *learning* a policy (or expected value functions  $\mathcal{Q}_t(\cdot)$  for all  $t = 2, \dots, T$ ) or *training* a model of this policy (or cut approximations  $\mathcal{V}_t(\cdot)$  for all  $t = 2, \dots, T$ ) using back-propagation. In the forward pass the inputs are propagated through the stages using the current model, and in the backward pass cuts (representing the error of the current approximation) are propagated back through the stages to update the model.

Algorithm 3.1 provides a pseudocode for SDDP. We next take a more detailed and technical look at the algorithmic steps.

**3.2. Forward Pass.** At the start of each iteration  $i$ , at first a subset  $\mathcal{K} \subseteq \mathcal{S}$  of scenarios is sampled with  $|\mathcal{K}| \ll |\mathcal{S}|$  (note that we may equivalently sample stage by stage during the forward pass). The number of samples  $|\mathcal{K}|$  may vary by iteration, but we do not state this possible dependence explicitly. Traditionally, and most commonly, in SDDP a random sampling is used, but a deterministic sampling is also possible. We discuss sampling techniques further in section 6.

At the first stage, the approximate subproblem

$$(3.1) \quad \min_{x_1 \in \mathcal{X}_1(x_0)} c_1^\top x_1 + \mathcal{V}_2^i(x_1)$$

is solved, which yields the trial point  $x_1^i = x_1^{ik}$  for all  $k \in \mathcal{K}$ . Afterwards, for each stage  $t = 2, \dots, T$  and each sample  $k \in \mathcal{K}$ , the approximate value functions  $\mathcal{Q}_t^i(x_{t-1}^{ik}, \xi_t^k)$  are evaluated recursively (this means that the subproblems (2.10) are solved for  $x_{t-1}^{ik}, \xi_t^k$ , and the current cut approximation  $\mathcal{V}_{t+1}^i(\cdot)$ ). In this way, for each sample  $k \in \mathcal{K}$ , a sequence of trial points  $(x_t^{ik})_{t \in [T]}$  is obtained.

The forward pass of SDDP is illustrated in Figure 4 for the recombining scenario tree from Figure 1 and  $\mathcal{K} = \{1, 3, 9\}$ , i.e.,  $|\mathcal{K}| = 3$ . The three sampled scenario paths are highlighted in green. The figure shows that for sample paths  $\xi^3$  and  $\xi^9$  the same node is reached at stage 3.

*Remark 3.2* (initialization). Before the first backward pass, the cut approximations  $\mathcal{V}_t(\cdot)$  are not yet initialized, so the forward pass subproblems (3.1) and (2.10) are not well-defined. This can be addressed using different strategies. First, in iteration  $i = 1$  the forward pass can be skipped and a feasible state trajectory  $(x_t^1)_{t \in [T]}$  for the backward pass can be user-defined or taken randomly instead. Second, such a state trajectory can be computed heuristically via a greedy approach where  $\mathcal{V}_t^1(\cdot) \equiv 0$  is assumed in the forward pass subproblems for all  $t \in [T]$ , so no coupling exists in the objective. Third, the cut approximations  $\mathcal{V}_t(\cdot)$  may be initialized with a valid user-defined lower bound  $\underline{\theta}_t$  for all  $t \in [T]$ . This approach is taken in the description of Algorithm 3.1.

**3.3. Backward Pass. Main Principle.** The backward pass starts at stage  $T$ . Here, for all samples  $k \in \mathcal{K}$ , we consider subproblems (2.10) for the trial point  $x_{T-1}^{ik}$  computed in the forward pass, all noise realizations  $\xi_{Tj}, j = 1, \dots, q_T$ , and  $\mathcal{V}_{T+1}^{i+1}(\cdot) \equiv 0$ . That is, we consider functions  $\mathcal{Q}_T^{i+1}(x_{T-1}^{ik}, \xi_{Tj})$  for  $j = 1, \dots, q_T$ .

As  $\mathcal{Q}_T^{i+1}(\cdot, \xi_{Tj})$  is convex in  $x_{T-1}$  by Lemma 2.11, it can be underestimated by a linear function using some subgradient  $\beta_{Tkj}^i \in \partial \mathcal{Q}_T^{i+1}(\cdot, \xi_{Tj})$  for any  $j = 1, \dots, q_T$  and any  $k \in \mathcal{K}$ :

$$\mathcal{Q}_T^{i+1}(x_{T-1}, \xi_{Tj}) \geq \mathcal{Q}_T^{i+1}(x_{T-1}^{ik}, \xi_{Tj}) + (\beta_{Tkj}^i)^\top (x_{T-1} - x_{T-1}^{ik}).$$

**Algorithm 3.1.** SDDP.

**Input:** Problem (MSLP) satisfying Assumptions 1 to 9. Bounds  $\underline{\theta}_t, t = 2, \dots, T$ .  
Stopping criterion.

**Initialization**

- 
- 1: Initialize cut approximations with  $\theta_t \geq \underline{\theta}_t$  for all  $t = 2, \dots, T$ .
  - 2: Initialize lower bound with  $\underline{v}^0 = -\infty$ .
  - 3: Set iteration counter to  $i \leftarrow 0$ .
- SDDP Loop**
- 


- 4: **while** Stopping criterion not satisfied **do**
- 5:     Set  $i \leftarrow i + 1$ .
- 6:     **Forward Pass**
- 7:     Sample a subset  $\mathcal{K} \subseteq \mathcal{S}$  of scenarios.
- 8:     Solve the approximate first-stage problem (3.1) to obtain trial point  $x_1^i = x_1^{ik}$  for all  $k \in \mathcal{K}$ .
- 9:     **for** stages  $t = 2, \dots, T$  **do**
- 10:         **for** samples  $k \in \mathcal{K}$  **do**
- 11:             Solve the approximate stage- $t$  subproblem (2.10) associated with  $\underline{Q}_t^i(x_{t-1}^{ik}, \xi_t^k)$  to obtain trial point  $x_t^{ik}$ .
- 12:         **end for**
- 13:     **Backward Pass**
- 14:     **for** stages  $t = T, \dots, 2$  **do**
- 15:         **for** samples  $k \in \mathcal{K}$  **do**
- 16:             **for** noise terms  $j = 1, \dots, q_t$  **do**
- 17:                 Solve the updated approximate stage- $t$  subproblem (2.10) associated with  $\underline{Q}_t^{i+1}(x_{t-1}^{ik}, \xi_{tj})$ . Store the optimal value and dual vector  $\pi_t^{ikj}$ .
- 18:             **end for**
- 19:             Use relations (3.4)–(3.5) and (3.7) to create an optimality cut for  $\underline{\mathcal{Q}}_t(\cdot)$ .
- 20:             Update the cut approximation  $\underline{\mathcal{V}}_t(\cdot)$  to  $\underline{\mathcal{V}}_t^{i+1}(\cdot)$  using relation (3.6).
- 21:         **end for**
- 22:     **end for**
- 23:     Solve the approximate first-stage problem (3.8) to obtain a lower bound  $\underline{v}^i$ .
- 24: **end while**

**Output:** (Approximately) optimal feasible policy for (MSLP) defined by  $x_1^i$  and cut approximations  $\underline{\mathcal{V}}_t^i(\cdot), t = 2, \dots, T$ .  $x_1^i$  defines an (approximately) optimal solution to problem (2.6) with  $\bar{v}_{\mathcal{K}}^i \approx v^*$ .

---

Since  $\underline{Q}_T^{i+1}(\cdot, \xi_{Tj})$  is a lower approximation of the true value function  $Q_T(\cdot, \xi_{Tj})$ , this directly implies

$$Q_T(x_{T-1}, \xi_{Tj}) \geq \underline{Q}_T^{i+1}(x_{T-1}^{ik}, \xi_{Tj}) + (\beta_{Tkj}^i)^\top (x_{T-1} - x_{T-1}^{ik}).$$



**Fig. 4** Illustration of SDDP forward pass for  $|\mathcal{K}|=3$ .

Taking expectations with respect to  $\xi_T$  on both sides, we obtain

$$\begin{aligned}
 \mathcal{Q}_T(x_{T-1}) &\geq \mathbb{E}_{\xi_T} [\underline{Q}_T^{i+1}(x_{T-1}^{ik}, \xi_T)] + \mathbb{E}_{\xi_T} [(\beta_{Tk}^i)^\top (x_{T-1} - x_{T-1}^{ik})] \\
 &= \mathbb{E}_{\xi_T} [\underline{Q}_T^{i+1}(x_{T-1}^{ik}, \xi_T) - (\beta_{Tk}^i)^\top x_{T-1}^{ik}] + \left( \mathbb{E}_{\xi_T} [\beta_{Tk}^i] \right)^\top x_{T-1} \\
 (3.2) \quad &= \underbrace{\sum_{j=1}^{q_T} p_{Tj} \left( \underline{Q}_T^{i+1}(x_{T-1}^{ik}, \xi_{Tj}) - (\beta_{Tkj}^i)^\top x_{T-1}^{ik} \right)}_{=:\alpha_{Tk}^i} + \underbrace{\left( \sum_{j=1}^{q_T} p_{Tj} \beta_{Tkj}^i \right)^\top}_{=:\beta_{Tk}^i} x_{T-1},
 \end{aligned}$$

where we exploit the finiteness of  $\xi_T$  (Assumption 5).  $\alpha_{Tk}^i$  is called the *cut intercept* and  $\beta_{Tk}^i$  is called the *cut gradient*. Defining

$$\phi_{Tk}^i(x_{T-1}) := \alpha_{Tk}^i + (\beta_{Tk}^i)^\top x_{T-1},$$

we can express (3.2) as

$$(3.3) \quad \mathcal{Q}_T(x_{T-1}) \geq \phi_{Tk}^i(x_{T-1}).$$

Inequality (3.3) defines a cut for  $\mathcal{Q}_T(\cdot)$ . Such a cut is constructed for each  $k \in \mathcal{K}$ . With these new cuts, the cut approximation  $\underline{\mathcal{V}}_T^i(\cdot)$  is updated to

$$\underline{\mathcal{V}}_T^{i+1}(x_{T-1}) := \max \left\{ \underline{\mathcal{V}}_T^i(x_{T-1}), \phi_{T1}^i(x_{T-1}), \dots, \phi_{T|\mathcal{K}|}^i(x_{T-1}) \right\}.$$

Thus, assuming that  $|\mathcal{K}|$  does not change over the iterations,  $\underline{\mathcal{V}}_T^{i+1}(\cdot)$  consists of  $i|\mathcal{K}|$  affine functions  $\phi_{Tk}^i(\cdot)$ ; cf. formulation (2.12).

In the same way, for stages  $t = T-1, \dots, 2$ , cuts for  $\mathcal{Q}_t(\cdot)$  can be constructed by solving subproblems (2.10) for the trial points  $x_{t-1}^{ik}$  and all noise realizations  $\xi_{tj}, j = 1, \dots, q_t$ . Importantly, by going backwards through the stages, at stage  $t$  we can already factor in the cuts that have been constructed at the following stage  $t+1$ , thus using a better approximation as the basis to construct a new cut. This means that we consider  $\underline{\mathcal{V}}_{t+1}^{i+1}(\cdot)$  and therefore  $\underline{Q}_t^{i+1}(\cdot, \cdot)$  with index  $i+1$  in the backward pass of iteration  $i$ .

As for stage  $T$ , we obtain

$$(3.4) \quad \mathcal{Q}_t(x_{t-1}) \geq \underbrace{\sum_{j=1}^{q_t} p_{tj} \left( \underline{Q}_t^{i+1}(x_{t-1}^{ik}, \xi_t) - (\beta_{tkj}^i)^\top x_{t-1}^{ik} \right)}_{=:\alpha_{tk}^i} + \underbrace{\left( \sum_{j=1}^{q_t} p_{tj} \beta_{tkj}^i \right)^\top}_{=:\beta_{tk}^i} x_{t-1},$$

where  $\beta_{tkj}^i$  denotes a subgradient of  $\underline{Q}_t^{i+1}(\cdot, \xi_{tj})$  at  $x_{t-1}^{ik}$  for  $k \in \mathcal{K}, j = 1, \dots, q_t$ . Again, by defining

$$\phi_{tk}^i(x_{t-1}) := \alpha_{tk}^i + (\beta_{tk}^i)^\top x_{t-1},$$

we can obtain a cut

$$(3.5) \quad \mathcal{Q}_t(x_{t-1}) \geq \phi_{tk}^i(x_{t-1})$$

for each  $k \in \mathcal{K}$  and can update the cut approximation to

$$(3.6) \quad \underline{\mathcal{V}}_t^{i+1}(x_{t-1}) := \max \left\{ \underline{\mathcal{V}}_t^i(x_{t-1}), \phi_{t1}^i(x_{t-1}), \dots, \phi_{t|\mathcal{K}|}^i(x_{t-1}) \right\}.$$

**Computing Subgradients.** Up to this point, we have discussed the main idea of the cut generation process in the backward pass of SDDP, which is based on evaluating approximate value functions  $\underline{Q}_t^{i+1}(\cdot, \cdot)$  and using subgradients for them at trial points  $x_{t-1}^{ik}$ . For the interested reader, we now address how to compute those subgradients in more detail. This step uses *dual information*, i.e., it is based on the duality theory of convex programs. For simplicity, we assume  $X_t = \{x_t \in \mathbb{R}^{n_t} \mid x_t \geq 0\}$  for all  $t \in [T]$ .

Consider stage  $T$ , some  $k \in \mathcal{K}$ , and some  $j \in \{1, \dots, q_T\}$ . Then, the dual problem to the linear stage- $T$  subproblem (2.10) is

$$\begin{cases} \max_{\pi_T} & (h_{Tj} - T_{T-1,j} x_{T-1}^{ik})^\top \pi_T \\ \text{s.t.} & W_{Tj}^\top \pi_T \leq c_{Tj}. \end{cases}$$

Let  $\pi_T^{ikj}$  be an optimal dual basic solution. Such a solution always exists by relatively complete recourse and boundedness (see Assumption 9 and Lemma 2.5). By the strong duality of LPs, it follows that

$$\begin{aligned} \underline{Q}_T^{i+1}(x_{T-1}^{ik}, \xi_{Tj}) &= (h_{Tj} - T_{T-1,j} x_{T-1}^{ik})^\top \pi_T^{ikj} \\ &= -(\pi_T^{ikj})^\top T_{T-1,j} x_{T-1}^{ik} + (\pi_T^{ikj})^\top h_{Tj}. \end{aligned}$$

Importantly, the dual feasible set does not depend on  $x_{T-1}$ , but remains unchanged for all trial points. In particular,  $\pi_T^{ikj}$  is dual feasible, but not necessarily dual optimal for all  $x_{T-1}$ . Therefore, and because of minimization, it follows that

$$\begin{aligned} \underline{Q}_T^{i+1}(x_{T-1}, \xi_{Tj}) &\geq -(\pi_T^{ikj})^\top T_{T-1,j} x_{T-1} + (\pi_T^{ikj})^\top h_{Tj} \\ &= -(\pi_T^{ikj})^\top T_{T-1,j} (x_{T-1} + x_{T-1}^{ik} - x_{T-1}^{ik}) + (\pi_T^{ikj})^\top h_{Tj} \\ &= \underline{Q}_T^{i+1}(x_{T-1}^{ik}, \xi_{Tj}) - (\pi_T^{ikj})^\top T_{T-1,j} (x_{T-1} - x_{T-1}^{ik}). \end{aligned}$$

Hence,

$$\beta_{tkj}^i = -(\pi_T^{ikj})^\top T_{T-1,j}$$

is a subgradient of  $\underline{Q}_T^{i+1}(\cdot, \xi_{Tj})$  at  $x_{T-1}^{ik}$ .

The previous derivation provides some additional insight. Since the dual feasible set is polyhedral and does not depend on  $x_{T-1}$ , for each noise term  $\xi_{Tj}, j = 1, \dots, q_T$ , there exist only finitely many dual extreme points (dual basic solutions) that can be attained. Therefore, only finitely many different cut coefficients can be generated. This is crucial for some convergence proofs of SDDP, as we discuss in what follows.

For earlier stages  $t = T - 1, \dots, 2$ , the dual problem to subproblem (2.10) looks a bit more sophisticated as the cut approximations  $\underline{\mathcal{V}}_{t+1}^{i+1}(\cdot)$  have to be taken into account, which requires additional dual multipliers  $\rho_t^r$  for all cuts  $r \in \Gamma_{t+1}$ , where  $\Gamma_{t+1}$  denotes the index set of cuts generated for the following stage. However, the derivation is completely analogous and, again, we arrive at

$$\underline{Q}_t^{i+1}(x_{t-1}, \xi_{tj}) \geq \underline{Q}_t^{i+1}(x_{t-1}^{ik}, \xi_{tj}) - (\pi_t^{ikj})^\top T_{t-1,j}(x_{t-1} - x_{t-1}^{ik}),$$

so that

$$(3.7) \quad \beta_{tkj}^i = -(\pi_t^{ikj})^\top T_{t-1,j}$$

is a subgradient of  $\underline{Q}_t^{i+1}(\cdot, \xi_{tj})$  at  $x_{t-1}^{ik}$ . Interestingly, the optimal dual multipliers  $\rho_t^{rikj}$  are not explicitly required in this formula.

**3.4. Bounds and Stopping.** At the first stage, the subproblem

$$(3.8) \quad \underline{v}^i := \min_{x_1 \in \mathcal{X}_1(x_0)} c_1^\top x_1 + \underline{\mathcal{V}}_2^{i+1}(x_1)$$

is solved. As  $\underline{\mathcal{V}}_2^{i+1}(\cdot)$  is a lower approximation of  $\mathcal{Q}_2(\cdot)$ ,  $\underline{v}^i$  is a valid lower bound to the optimal value  $v^*$  of (MSLP). This bound can be initialized with  $\underline{v}^0 = -\infty$  or any a priori known lower bound for  $v^*$ .

In contrast, we are not guaranteed to obtain a valid upper bound for  $v^*$  during iterations of standard SDDP, as we only consider a small subset  $\mathcal{K} \subseteq \mathcal{S}$  of all scenarios. This means that in the forward pass, the feasible policy for (MSLP), which is implicitly defined by the current cut approximations  $\underline{\mathcal{V}}_t^i(\cdot)$ ,  $t = 2, \dots, T$ , is only evaluated for a subset of all scenarios. By evaluating these scenarios in the objective of (MSLP) and taking the sample average

$$(3.9) \quad \bar{v}_{\mathcal{K}}^i := \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \underbrace{\sum_{t=1}^T (c_t(\xi_t^k))^\top x_t^{ik}}_{=: v^i(\xi^k)},$$

we only obtain an unbiased estimator of the true upper bound  $\bar{v}^i$  (a statistical upper bound) associated with the current policy; see section 7 for more details.

After each iteration of SDDP, one or several stopping criteria are checked, which may or may not be based on  $\bar{v}_{\mathcal{K}}^i$ . We discuss different stopping criteria in detail in section 7. If SDDP does not stop, a new iteration  $i + 1$  is started with a forward pass.

**3.5. Cut Properties.** We discuss convergence of SDDP in section 4. It relies on three key properties of the derived cuts.

**LEMMA 3.3.** *For any stage  $t = 2, \dots, T$  and any  $k \in \mathcal{K}$ , the functions  $\phi_{tk}^i(\cdot)$  are*

- (a) *valid lower approximations of  $\mathcal{Q}_t(\cdot)$ ,*
- (b) *tight for  $\underline{\mathcal{Q}}_t^{i+1}(\cdot)$  (as defined in (2.11)) at  $x_{t-1}^{ik}$ , and*
- (c) *finite, i.e., only finitely many different cuts can be generated, if we restrict to dual basic solutions to generate cuts.*

*Proof.* Property (a) follows immediately from (3.3) and (3.5). (b) holds because of strong duality for LPs and by taking expected values over the obtained optimal values. Alternatively, we can rearrange the RHS of inequality (3.4) to obtain

$$(3.10) \quad \phi_{tk}^i(x_{t-1}) = \underline{\mathcal{Q}}_t^{i+1}(x_{t-1}^{ik}) + \sum_{j=1}^{q_t} p_{tj}(\beta_{tkj}^i)^\top (x_{t-1} - x_{t-1}^{ik}).$$

Inserting  $x_{t-1}^{ik}$  yields  $\phi_{tk}^i(x_{t-1}^{ik}) = \underline{\mathcal{Q}}_t^{i+1}(x_{t-1}^{ik})$ . Property (c) follows by induction using the arguments on the dual feasible region previously discussed for stage  $T$ .  $\square$

Note that  $\phi_{tk}^i(\cdot)$  is not necessarily tight for the true expected value function  $\mathcal{Q}_t(\cdot)$  in early iterations for  $t \neq T$ , and it might provide a loose cut only. However, by the finiteness and tightness properties it can be shown recursively that the derived cuts eventually become tight for  $\mathcal{Q}_t(\cdot)$  as well. In fact, after finitely many steps, the polyhedral function  $\mathcal{Q}_t(\cdot)$  is represented exactly for all  $t = 2, \dots, T$ . This is a key property for the convergence of SDDP.

**3.6. Illustrative Example.** To illustrate the key steps of SDDP, we present a simple example.

*Example 3.4.* Consider the three-stage (MSLP)

$$(3.11) \quad \begin{aligned} \min \quad & x_1 + x_2 + x_{31} + x_{32} \\ \text{s.t.} \quad & x_1 \leq 6, \\ & x_2 \geq \xi_2 - x_1, \\ & x_{31} - x_{32} = \xi_3 - x_2, \\ & x_1, x_2, x_{31}, x_{32} \geq 0, \end{aligned}$$

which is inspired by Example 2 in Chapter 5 of [27]. The uncertain data in the RHS is stagewise-independent and uniformly distributed with  $\xi_2 \in \{4, 5, 6\}$  and  $\xi_3 \in \{1, 2, 4\}$ .

Problem (3.11) does not have entirely the same structure as problem (MSLP), but can be easily converted into it by introducing slack variables. However, for illustrative purposes, we do not do this. The problem can be expressed by means of the value functions

$$(3.12) \quad Q_3(x_2, \xi_3) = \begin{cases} \min_{x_3} & x_{31} + x_{32} \\ \text{s.t.} & x_{31} - x_{32} = \xi_3 - x_2, \\ & x_{31}, x_{32} \geq 0, \end{cases}$$

and

$$Q_2(x_1, \xi_2) = \begin{cases} \min_{x_2} & x_2 + \mathcal{Q}_3(x_2) \\ \text{s.t.} & x_2 \geq \xi_2 - x_1, \\ & x_2 \geq 0. \end{cases}$$

The first-stage problem is then

$$v^* = \begin{cases} \min_{x_1} & x_1 + \mathcal{Q}_2(x_1) \\ \text{s.t.} & x_1 \in [0, 6]. \end{cases}$$

The optimal solution is given by  $x_1^* = 3$  with  $v^* = \frac{53}{9}$ .

As shown in [27], the stage-3 value functions can be written in closed form as  $Q_3(x_2, \xi_3) = |\xi_3 - x_2|$  for all scenarios. Taking expectations, a closed-form expression for  $\mathcal{Q}_3(\cdot)$  can be derived, and by recursion we obtain

$$\mathcal{Q}_2(x_1) = \begin{cases} \frac{23}{3} - \frac{16}{9}x_1, & x_1 \in [0, 1], \\ \frac{67}{9} - \frac{10}{9}x_1, & x_1 \in [1, 2], \\ \frac{59}{9} - \frac{10}{9}x_1, & x_1 \in [2, 3], \\ \frac{47}{9} - \frac{2}{3}x_1, & x_1 \in [3, 4], \\ \frac{31}{9} - \frac{2}{9}x_1, & x_1 \in [4, 5], \\ \frac{7}{3}, & x_1 \in [5, 6]. \end{cases}$$

The optimal value is  $v^* = \frac{56}{9}$ .


We apply SDDP for illustration. We assume loose initial bounds  $\theta_2, \theta_3 \geq -10$  for simplicity. In the forward pass, we sample one scenario path per iteration, i.e.,  $|\mathcal{K}| = 1$ . In iteration 1, let  $(\xi_2, \xi_3) = (5, 4)$  define this path. Solving the approximate subproblems (2.10) for all stages  $t = 1, 2, 3$  and  $(\xi_2, \xi_3) = (5, 4)$ , we obtain  $\bar{v}_{\mathcal{K}}^1 = 6$ . In fact, this is not a valid upper bound for  $v^*$ .

In the backward pass, cuts for  $\mathcal{Q}_t(\cdot), t = 2, 3$ , are derived at the trial points. For stage 3, the cut gradient is  $\beta_3(5) = 1$ . Moreover,  $\mathcal{Q}_3^2(5) = \frac{8}{3}$ . With formulas (3.5) and (3.10) this yields the cut  $\mathcal{Q}_3(x_2) \geq -\frac{7}{3} + x_2$ , which is incorporated into the stage-2 subproblems. Solving these problems yields the cut  $\mathcal{Q}_2(x_1) \geq \frac{23}{3} - 2x_1$ . At the first stage, the lower bound computes to  $\underline{v}^1 = \frac{5}{3}$ .


The expected value functions and the cuts obtained for three iterations are depicted in Figure 5. In the second and third iterations, the same scenario path  $(\xi_2, \xi_3) = (6, 1)$  is sampled in the forward pass.

Figure 6 displays the bounds  $\underline{v}^i$  and  $\bar{v}_{\mathcal{K}}^i$  for ten iterations of SDDP. It shows that the lower bounds stabilize quickly at  $v^*$ , whereas the values of  $\bar{v}_{\mathcal{K}}^i$  oscillate around  $v^*$ .

**3.7. Policy Assessment.** As mentioned above, in standard SDDP no valid upper bound  $\bar{v}$  for  $v^*$  is determined. While in each iteration a statistical upper bound (3.9) can be computed, the number of samples  $|\mathcal{K}|$  can often be too small to appropriately assess the quality of the current policy. In particular,  $|\mathcal{K}|$  is often chosen to be 1 in practice, and thus  $\bar{v}_{\mathcal{K}}^i$  is not a meaningful estimate for  $\bar{v}$ .



**Fig. 5** Expected value functions for Example 3.4 with cuts obtained in first three iterations depicted in blue, green, and red.



**Fig. 6** Bounds for 10 iterations of SDDP applied to Example 3.4.

Therefore, to assess the obtained policy, usually an additional forward simulation is conducted once SDDP has terminated. For this simulation a much higher number of sample paths through the scenario tree is used, e.g.,  $|\mathcal{K}| \in \{1000, 10000\}$ , leading to a reasonable estimator  $\bar{v}_K$ . In this step, the simulation can be performed either *in-sample* (using sample paths through the recombining scenario tree) or *out-of-sample* (using the true underlying distribution, e.g., if  $\xi_t$  is a continuous random variable that is discretized to satisfy Assumption 5; see section 11).

*Remark 3.5.* In the light of Remark 3.1, this policy assessment step can also be interpreted from a statistical learning perspective. After the model has been trained, a model validation (using in-sample data) or a model test (using out-of-sample data) is performed.

**4. Convergence and Complexity.** The convergence behavior of SDDP has been thoroughly analyzed over the years. We discuss the main convergence results in this section. We first focus on *finite* convergence of SDDP and then discuss the actual convergence *rate*, i.e., the computational complexity of SDDP. Our overview is loosely based on the literature review in [75].

**4.1. Finite Convergence.** The first convergence analyses related to SDDP were conducted in [41] and [132], which implicitly assumed independence of sampled random variables and convergent subsequences of algorithm iterates. A first complete convergence proof was given by Philpott and Guan in [171] for the case where uncertainty only enters the RHS of (MSLP) (in fact, they consider a more general algorithm than SDDP, including sampling in the backward pass). The same reasoning is used by Shapiro [205] for the case where  $W_t$ ,  $c_t$ , and  $T_{t-1}$  are also uncertain.

The convergence behavior of SDDP can be explained using two main arguments: First, as stated in Lemma 3.3, only finitely many different cuts and thus only finitely many different cut approximations  $\mathcal{V}_t(\cdot)$  can be constructed for all  $t = 2, \dots, T$ . This result requires linearity (Assumption 6) and finite random variables (Assumption 5). Moreover, those finitely many cuts also satisfy some tightness property, which implies that they are sufficient to exactly represent the polyhedral (expected) value functions (see Theorem 2.8). For a deterministic algorithm, this would result in finite convergence to the true optimal point and value (see the convergence properties of Benders decomposition [17] and Kelley’s cutting-plane method [118]).



For SDDP, it must be taken into account that scenarios are sampled in the forward pass. This means that the cut approximations might not further improve for some iterations if the *wrong* scenarios are sampled. Therefore, the second key argument for many proofs of finite convergence of SDDP is that each scenario is visited infinitely many times with probability 1 given that the algorithm does not terminate. Intuitively, this means that after finitely many iterations the *right* scenarios will be sampled with probability 1, leading to the construction of a new cut. This requirement is satisfied under *independent sampling*, that is, if the sampling in the forward pass of Algorithm 3.1 is random and independent of previous iterations. It is also satisfied for an exhaustive enumeration of all scenarios in the sampling process. We should emphasize that this argument is purely theoretical in order to establish convergence results for SDDP. When applying SDDP in practice, it is usually not possible to sample each scenario even *once* in reasonable time.

Using these two arguments, the following main convergence result can be obtained.

**THEOREM 4.1** (almost-sure finite convergence of SDDP). *Under Assumptions 1 to 9 and using an independent random sampling procedure in the forward pass, SDDP converges with probability 1 to an optimal policy of (MSLP) in a finite number of iterations.*

Importantly, almost-sure finite convergence to an optimal policy of (MSLP) does not imply that the trajectories  $(x_t^{ik})_{t \in [T]}, k \in \mathcal{K}$ , and the corresponding sample averages  $\bar{v}_{\mathcal{K}}^i$  obtained in SDDP converge, since both are random and depend on the current sample  $\mathcal{K}$ . However, the lower bounds  $\underline{v}^i$  obtained in SDDP converge to  $v^*$ .

**Deterministic Sampling.** Recently, convergence analyses of SDDP and related algorithms have often made use of deterministic sampling techniques instead of random sampling in line 6 of Algorithm 3.1 [10, 11]. Here, the idea is that the approximation error in SDDP can be controlled and guided to zero in a deterministic way if in each iteration scenarios are sampled for which the current approximation gap is maximized. This, however, requires that the approximation gap itself can be bounded rigorously. Therefore, in addition to the lower cut approximation  $\underline{V}_t(\cdot)$  an upper approximation  $\bar{V}_t(\cdot)$  is also constructed and iteratively refined [10, 240], so that deterministic lower bounds  $\underline{v}^i$  and upper bounds  $\bar{v}^i$  are computed in each iteration. For more details on deterministic sampling and deterministic upper bounds, we refer to sections 6 and 8.

**Generalizations.** It is shown in the literature that some of the basic assumptions (Assumptions 1 to 9) can be relaxed without compromising the convergence of SDDP. Girardeau, Leclère, and Philpott [85] analyze the case where SDDP is applied to multistage problems with nonlinear convex subproblems, i.e., when Assumption 6 is relaxed. In this case, the value functions  $Q_t(\cdot)$  are no longer polyhedral, but are still convex. The authors show that almost-sure convergence is still satisfied as long as some convexity and compactness assumptions and some tightened recourse assumptions are satisfied. We discuss this result in detail in section 15 when we formally introduce convex multistage stochastic nonlinear problems. The main idea is that even without polyhedrality,  $Q_t(\cdot)$  can be guaranteed to be Lipschitz continuous, so that the approximations of  $\mathcal{Q}_t(\cdot)$  become better in a whole neighborhood of the trajectories  $(x_t^{ik})_{t \in [T]}, k \in \mathcal{K}$ .

Guigues [92] generalizes this convergence result to the risk-averse case where Assumption 8 is relaxed. Forcier and Leclère [75] prove convergence for (MSLP) without finite randomness, i.e., dropping Assumption 5. Further convergence proofs are pro-

vided for multicut SDDP [8], SDDP with cut selection [8, 94], adaptive partition-based SDDP [214] (see also section 21) using SDDP with saddle cuts [59] (see also section 14), and variants of distributionally robust SDDP [69, 169] (see also section 13). Another proof of almost-sure finite convergence for extensions to nonconvex problems is provided in [240].

**4.2. Complexity.** Theorem 4.1 guarantees almost-sure finite convergence of SDDP. While this result is of theoretical interest, it may not be very relevant in practical applications as it does not provide any result on the rate of convergence. As pointed out in [75] and mentioned above, the argument based on scenarios being sampled repeatedly (infinitely many times) is, in particular, almost never applicable to SDDP in practice due to the sheer number of scenarios in  $\mathcal{S}$ . Important for the rate of convergence are the computational cost per iteration and the required number of iterations.

**Cost per Iteration.** For the computational cost per iteration, the number of LPs to be solved in the backward pass is crucial. Per sample  $k \in \mathcal{K}$  in the forward pass,  $q_t$  subproblems are solved for each stage except for  $t = 1$  in the backward pass. Therefore, the total number of LPs solved is  $1 + |\mathcal{K}| \sum_{t=2}^T q_t$ . Hence, the number of problems to be solved grows linearly in the number of stages  $T$ , the number of samples  $|\mathcal{K}|$ , and the number of noise terms  $q_t$  [185].

**Expected Number of Iterations.** The computational bottleneck for SDDP is the expected required number of iterations to achieve convergence. Recently, there has been active research on computing theoretical bounds for this number, with Lan [124] as well as Zhang and Sun [240] publishing similar results using slightly different approaches. In both cases, the authors start by considering some case of deterministic sampling before enhancing their results to the random sampling variant of SDDP (in [124] the associated algorithm is referred to as explorative dual dynamic programming (EDDP)). We discuss deterministic sampling in more detail in section 6. The main idea behind deriving iteration bounds is as follows: By exploiting Lipschitz continuity of  $\mathcal{V}_t(\cdot)$  and  $\mathcal{Q}_t(\cdot)$ , it is possible to also control the approximation error at points where no cuts are constructed, as long as they lie in a neighborhood of some trial point  $x_t^{ik}$ . As long as the state space is bounded for all  $t \in [T]$  (cf. Assumption 9), it can be completely covered by finitely many such neighborhoods [240]. Similar reasoning is applied in [75].

More formally, Lan [124] introduces the notion of saturated points  $\bar{x}_{t-1}$ , in which the approximation of  $\mathcal{Q}_t(\cdot)$  is already  $\varepsilon$ -close for some predefined tolerance  $\varepsilon > 0$ , i.e.,

$$\mathcal{Q}_t(\bar{x}_{t-1}) - \mathcal{V}_t^i(\bar{x}_{t-1}) \leq \varepsilon,$$

and *distinguishable points*  $\bar{x}_{t-1}$ , which have at least a  $\delta$ -distance to the set  $X_{t-1}^{sat}$  of already saturated points for some  $\delta > 0$ ; that is,

$$\|\bar{x}_{t-1} - x_{t-1}\| > \delta \quad \forall x_{t-1} \in X_{t-1}^{sat}.$$

If some trial point  $x_t^{ik}$  is saturated and distinguishable, the iteration  $i$  can be called *effective* [75]. Using deterministic sampling, all iterations in SDDP can be shown to be effective, and thus the number of iterations can be bounded in the aforementioned way. For random sampling, this is not true, but the probability of an effective iteration is at least  $\frac{1}{N}$  with  $N := \prod_{t=2}^{T-1} n_t$ .

In the light of Assumption 9 (b), for any  $t \in [T]$ , we call the bound  $D_t$  satisfying

$$\|x_t - x'_t\| \leq D_t \quad \forall x_t, x'_t \in \mathcal{X}_t$$

the *diameter* of the state space. Additionally, let  $L$  denote a Lipschitz constant for the objective function of (MSLP), which exists due to Corollary 2.10.

The following complexity results are then satisfied by SDDP.

**THEOREM 4.2** (complexity of SDDP [124, 240]). *Let  $D_t \leq D$  for all  $t \in [T]$ . For some arbitrary  $\varepsilon > 0$ , the (expected) number of required iterations of SDDP (Algorithm 3.1) to obtain*

- *an  $\varepsilon$ -optimal solution using deterministic sampling is*
  - *polynomial in  $T, (\frac{1}{\varepsilon}), L$  and  $D$ ,*
  - *exponential in  $n_t$ ;*
- *an  $\varepsilon$ -optimal solution using deterministic sampling, given that  $\mathcal{X}_t$  is finite with cardinality  $|\mathcal{X}_t| \leq \bar{X}$ , is*
  - *linear in  $T$  and  $\bar{X}$ ;*
- *a  $(T\varepsilon)$ -optimal solution using deterministic sampling is*
  - *linear in  $T$ ,*
  - *polynomial in  $T, (\frac{1}{\varepsilon}), L$ , and  $D$ ,*
  - *exponential in  $n_t$ ;*
- *an  $\varepsilon$ -optimal solution using random sampling is*
  - *polynomial in  $q_t, (\frac{1}{\varepsilon}), L$ , and  $D$ ,*
  - *exponential in  $T$  and  $n_t$ .*

This means that for standard SDDP (using random sampling) the expected number of iterations grows exponentially in the horizon  $T$  and the dimension  $n_t$  of the state space. This is computationally important. The exponential complexity with respect to the state dimension is not that surprising, as it is well known for cutting-plane methods [153] and inherited by SDDP. Similarly, the exponential complexity with respect to the number of stages directly follows from the exponential number of scenarios that may have to be sampled in the worst case. Interestingly, under deterministic sampling, the complexity is independent of the number  $q_t$  of noise terms per stage, as this number only affects the computational cost per iteration.

We see that using some deterministic sampling scheme a polynomial or even linear iteration complexity in  $T$  can be achieved, whereas the iteration complexity in the state space cannot be alleviated [240].

The complexity results in [124, 240] are further generalized by Forcier and Leclère [75]. They provide results for a generalized framework of SDDP-related algorithms, including SDDP with inexact cuts or regularization (see also section 21), risk-averse SDDP (see also section 12), and extensions to convex nonlinear or nonconvex mixed-integer (nonlinear) problems (see also sections 15 and 16).

**5. Comparison with Related Methods.** We briefly compare SDDP to solution methods that it is (historically) related to, as discussed in section 1.

**5.1. Relation to SDP.** SDDP is closely related to stochastic dynamic programming (SDP). SDP usually is applied in a setting in which not only state variables, but also additional local variables are considered; see Remarks 2.2 and 2.4. Therefore, the DPE and value functions are considered in the form of (2.7), which we repeat here for convenience:

$$Q_t(x_{t-1}, \xi_t) = \min_{u_t \in U_t(x_{t-1}, \xi_t)} f_t(u_t, \xi_t) + \mathcal{Q}_{t+1}(\mathcal{T}_t(x_{t-1}, u_t, \xi_t)).$$

The main idea of SDP is to explicitly evaluate the (expected) value functions for all possible cases during a forward or backward iteration through the stages  $t \in [T]$ .


This is only possible if the support  $\Xi_t$  of  $\xi_t$  and the state space  $\mathcal{X}_t \subset X_t$  are finite for all  $t \in [T]$ . Otherwise, infinitely many evaluations would be required. In addition, it is also required that the action space  $U_t(x_{t-1}, \xi_t)$  is finite for all  $x_{t-1} \in X_{t-1}, \xi_t \in \Xi_t$ , so that the minimum in (2.7) can be computed by finitely many evaluations. For this reason, all these sets may have to be discretized first [176].

The computational effort of SDP scales linearly in  $T$  and in the cardinalities  $|X_t|, |U_t(x_{t-1}, \xi_t)|$ , and  $|\Xi_t|$ . The three sets might be multidimensional and thus required to be discretized in each dimension  $n_t, \tilde{n}_t$ , and  $\kappa_t$ . Hence, their cardinality grows exponentially in these dimensions, which is computationally prohibitive for high-dimensional problems. This is known as the curse of dimensionality of SDP; see also section 1.

SDDP avoids the requirements of state space and action space discretization by not evaluating  $\mathcal{Q}_t(\cdot), t \in [T]$ , exactly for all (finitely many) possible actions and states, but instead it approximates them using an iteratively refined polyhedral outer approximation  $\underline{\mathcal{V}}_t(\cdot)$ , constructed by linear cuts. It can thus be considered an *approximate dynamic programming* (ADP) method.

**5.2. Relation to Nested Benders Decomposition.** In stochastic programming, it is common practice to consider problems (MSLP) with finite randomness (Assumption 5), but without the requirement of stagewise independence of  $\xi_t$  (Assumption 2). In that case the uncertainty can be modeled by a finite scenario tree, which compared to the recombining tree from section 2 exhibits some path dependence and satisfies the usual *tree* property that each node  $n$  has a finite set of child nodes  $\mathcal{C}(n)$ , but a unique parent node  $a(n)$ . An example of a scenario tree with  $T = 3$  and  $|\mathcal{S}| = 9$  is illustrated in Figure 7. This scenario tree represents the same number of scenarios  $|\mathcal{S}|$  as the recombining tree in Figure 1, but requires  $\sum_{t=2}^T q_t^{t-1} + 1$  instead of  $\sum_{t=2}^T q_t + 1$  nodes.

To solve the (MSLP) associated with a general scenario tree, in principle the same approximation approach as in SDDP can be used. However, due to the path dependence, the value functions  $Q_t(\cdot, \cdot)$  and expected value function  $\mathcal{Q}_t(\cdot)$  depend on the history  $\xi_{[t-1]}$  of the data process  $(\xi_t)_{t \in [T]}$ . In other words, each node  $n$  has its own value function  $Q_n(\cdot)$ , and with each node (except for leaf nodes) is associated an expected value function  $\mathcal{Q}_{\mathcal{C}(n)}(\cdot)$ . Therefore, to update the approximations  $\underline{\mathcal{V}}_{\mathcal{C}(n)}^i(\cdot)$  of all  $\mathcal{Q}_{\mathcal{C}(n)}(\cdot)$  in each iteration, all nodal subproblems have to be solved in the backward



**Fig. 7** Scenario tree with 9 scenarios and  $\xi^6$  highlighted.

pass, which in turn requires the computation of trial points  $x_{a(n)}^i$  for all nodes, i.e., solving all nodal subproblems in the forward pass as well.

Because of its close relation to the L-shaped method for solving two-stage stochastic LPs [236] and to Benders decomposition [17], this solution method is often called nested Benders decomposition (NBD). It was first proposed by Birge in 1985 [26] and can be interpreted as a decomposition method for the extensive form of the deterministic equivalent of (MSLP). Contrary to SDDP, NBD guarantees that valid lower bounds  $\underline{v}$  and upper bounds  $\bar{v}$  of  $v^*$  are determined in each iteration, which allows for a deterministic stopping criterion in a straightforward way. The upper bounds can be computed as

$$\bar{v}^i := \mathbb{E} \left[ \sum_{n \in \mathcal{T}} c_n^\top x_n^i \right],$$

where  $\mathcal{T}$  is the set of all nodes in the scenario tree.

However, due to the sheer number of subproblems to be solved in each iteration, which grows exponentially in  $T$ , NBD is only computationally tractable for problems of moderate size. By moderate we mean instances with some hundred or a few thousand scenarios, and four or five stages at maximum [234].

Furthermore, for general scenario trees, sampling scenarios from  $\mathcal{S}$  in the forward pass also does not necessarily help to reduce the computational burden and speed up the solution process, as it reduces the computational effort per iteration, but simultaneously implies that the cut approximations  $\underline{\mathcal{V}}_{\mathcal{C}(n)}^i(\cdot)$  are only improved for some  $\mathcal{Q}_{\mathcal{C}(n)}(\cdot)$  in each iteration. Under stagewise independence (Assumption 2), however, the outcome is different. The scenario tree collapses to a recombining tree. This means that for any stage  $t$ , many differing scenarios share the same nodes, and thus value functions. In particular, there exists only one expected value function  $\mathcal{Q}_t(\cdot)$  for each  $t = 2, \dots, T$ . Therefore, even if only a sample  $\mathcal{K} \subset \mathcal{S}$  of scenarios is considered in each iteration  $i$ , the cut approximations  $\underline{\mathcal{V}}_t^i(\cdot)$  for all  $\mathcal{Q}_t(\cdot)$  are still updated with new cuts. From this perspective, SDDP can be interpreted as a sampling variant of NBD which reduces the computational effort per iteration significantly [185], but which relies heavily on stagewise independence of  $(\xi_t)_{t \in [T]}$  in order to leverage the sampling with respect to value function approximations.

*Remark 5.1 (cut-sharing).* In the literature, the aforementioned property of SDDP is often referred to as *cut-sharing*. This is best understood by representing the stagewise-independent data process  $(\xi_t)_{t \in [T]}$  using a standard scenario tree. In this case, at any stage  $t$ , all nodes have the same set of successor nodes. If now only a sample  $\mathcal{K}$  of scenarios is considered in iterations of SDDP, only a subset of nodes is visited. Nonetheless, the cuts constructed for a specific node are valid for all equivalent nodes in the tree as well, so they are *shared* with other nodes/scenarios.

As mentioned above, a recombining scenario tree provides a more precise picture. For each stage, scenarios *share nodes* in the recombining tree and there exists only *one* function  $\mathcal{Q}_t(\cdot)$  that is to be approximated. Therefore, the phrase *cut-sharing* is sometimes considered to be misleading.

**5.3. Complexity Comparison.** We summarize the main complexity results for SDDP and the related methods in Table 3.

In contrast to SDP, SDDP does not require a state space and action space discretization. In particular, the latter is computationally important in practice, whereas the former may yield computational improvements, but does not translate into an improvement of the worst-case complexity class. On the other hand, SDDP does not have linear complexity in  $T$ .

**Table 3** Complexity of SDDP and related solution methods.

	Det.	Equiv.	NBD	SDDP	SDP
<b>Requirements:</b>					
stagewise independence	no		no	yes	yes (*)
state discretization	no		no	no	yes
action discretization	no		no	no	yes
noise discretization	yes		yes	yes	yes
<b>Complexity (#):</b>					
in no. of stages $T$	expon.		expon.	expon.	linear
in state dimension $n_t$	polyn.		expon.	expon.	expon.
in action dimension $\tilde{n}_t$	polyn.		polyn.	polyn.	expon.
in no. of realizations $q_t$ (§)	polyn.		polyn.	polyn.	linear
<b>Progressivity:</b>					
of bounds	yes		yes	yes	no

(\*) Markovian uncertainty is possible as well.  
(#) Comprises number of iterations, effort per iteration, subproblem sizes, subproblem solution, etc.  
(§) Note that  $q_t$  in itself is exponential in the noise dimension  $\kappa_t$  and polynomial in the discretization precision per noise component.

Another difference is that SDDP (as do NBD and most solvers for the deterministic equivalent) approximates  $v^*$  with improving lower (and upper) bounds. This means that if the computation time is increased, the quality of the approximation also improves. On the contrary, standard solution methods for SDP, such as backward induction, either manage or not to solve a problem in a given time limit, but do not use improving approximations. In particular, stopping SDP prematurely does not provide valid bounds for  $v^*$ .

Compared to NBD, SDDP mainly reduces the computational effort per iteration significantly, but it does not remove the exponential growth of the computational cost with respect to  $T$ . In return, it heavily relies on stagewise independence (Assumption 2).

We can conclude that SDDP, while mitigating some of the weaknesses of SDP and NBD (sometimes advertised as “breaking the curse of dimensionality”), does not manage to leave the respective worst-case complexity classes. On the contrary, it inherits some of the complexity drawbacks of both methods. Still, in many applications (where worst-case complexity is not an issue) it shows considerable performance improvements compared to SDP and NBD, especially for problems with continuous action space, a medium number of stages  $T$ , and a moderate state dimension  $n_t$ . While Theorem 4.2 indicates that convergence may take an extremely long time in large-scale applications, too long to be computationally tractable, SDDP has shown good performance for large-scale instances of (MSLP) in many applications, as discussed in section 9. This is due to various improvements which we address in the following sections.

**6. Sampling.** Sampling is a central element of SDDP (see section 3) and, in particular, line 6 of Algorithm 3.1. In the forward pass, a finite number  $|\mathcal{K}|$  of scenarios is sampled to simulate the current policy and compute a trajectory of trial points  $(x_t^{ik})_{t \in [T]}$  for all  $k \in \mathcal{K}$ . Often, this sampling is done from a finite set of scenarios  $\mathcal{S}$  (see Assumption 5), with  $|\mathcal{K}| \ll |\mathcal{S}|$ . Alternatively, it is possible to directly sample from a given (continuous) distribution.

In this section, we discuss different sampling techniques that can be used in SDDP. As indicated in sections 3 and 4, we can distinguish between *random sampling* and *deterministic sampling* methods. In standard SDDP, as originally proposed in [159], random sampling is used. Here, the main requirement is that the samples should be independent and identically distributed (i.i.d.). This is important for two reasons:

- (1) In this way, almost sure finite convergence of SDDP can be ensured, as any scenario is sampled infinitely many times with probability 1, assuming that the algorithm does not terminate; see section 4.
- (2) In the originally proposed stopping criterion of SDDP a confidence interval is used, which is built using the sample mean  $\bar{v}_{\mathcal{K}}^i$  (3.9); see section 7. However, by the central limit theorem, even an approximate confidence interval can only be obtained for a sequence of i.i.d. random variables.

**6.1. Monte Carlo Sampling.** The simplest sampling method satisfying the above requirement is Monte Carlo (MC) sampling. Here, samples are drawn randomly from the probability distribution of  $\xi_t$  in each iteration by first sampling from a uniform distribution and then using appropriate transforms. Under stagewise independence (Assumption 2), this is done independently for each stage  $t \in [T]$ .

As the quantities  $v^i(\xi^k)$  are i.i.d., the value  $\bar{v}_{\mathcal{K}}^i$  (3.9) that can be computed in the SDDP forward pass is an unbiased estimator of  $\bar{v}^i$  and, according to the strong law of large numbers, it converges to  $\bar{v}^i$  for  $|\mathcal{K}|$  approaching infinity. Still, the sampling error can be significant. The variance of  $\bar{v}_{\mathcal{K}}^i$  can be estimated by  $\frac{1}{|\mathcal{K}|} (\sigma_{\bar{v}, \mathcal{K}}^i)^2$ . This means that the variance can be reduced either by increasing the number of samples  $|\mathcal{K}|$  or by reducing the sample variance  $(\sigma_{\bar{v}, \mathcal{K}}^i)^2$ . Increasing the sample size might look promising at first glance, but it can become computationally intractable in practice [157]. Recall that for every sample  $k \in \mathcal{K}$ ,  $1 + \sum_{t=2}^T q_t$  subproblems have to be solved in the backward pass of each iteration. Therefore, the more promising approach is the combination of MC sampling with variance reduction techniques [157].

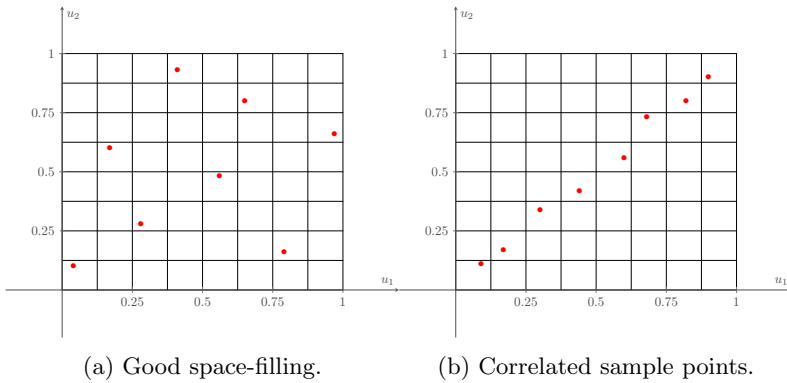
**6.2. Variance Reduction Techniques.** Incorporating variance reduction techniques into sampling in SDDP is studied extensively in [112, 157]. For a review of sampling techniques in stochastic programming in general, we refer to [111].

**Randomized QMC Sampling.** In [112] the use of quasi-Monte Carlo (QMC) sampling within SDDP is proposed. In this case, instead of randomly sampling from the uniform distribution, a deterministic sequence of points  $u^1, \dots, u^N$  from  $(0, 1)^{\kappa_t}$  is chosen. This is done in such a way that the sampled points fill  $(0, 1)^{\kappa_t}$  as homogeneously as possible (so the empirical distribution is as close to a uniform distribution as possible). After an appropriate transformation, these points provide a better representation of  $\xi_t$  than randomly sampled points.

Drawbacks of QMC methods are that the sample points are not random, the obtained estimator is biased, and no confidence interval can be established. Randomized QMC (RQMC) methods, where the choice of QMC points is combined with some kind of randomness, avoid this drawback and allow for standard error estimation [112].

Compared to MC sampling, RQMC methods achieve better convergence rates of  $\mathcal{O}(|\mathcal{K}|^{-1}(\log|\mathcal{K}|)^{\kappa_t})$ , and thus are considered more efficient. However, the convergence rate depends on the dimension  $\kappa_t$  of  $\xi_t$  [112].

**Latin Hypercube Sampling.** In Latin hypercube sampling (LHS) [148], the space  $(0, 1)^{\kappa_t}$  is divided into equidistant subintervals and then scenarios are sampled from each subinterval in such a way that in each row and column of the grid only one point is sampled. This is illustrated in Figure 8 (a).



**Fig. 8** LHS for two dimensions.

In this way, again a more homogeneous distribution of the sample points can be obtained, and compared to MC sampling the variance can be reduced. On the flipside, poor space-filling or correlation between the sample points has to be ruled out (see Figure 8 (b)), which requires significant additional effort.

**Incorporation into SDDP.** It is important to notice that while reducing the variance compared to the classical MC estimators, scenarios sampled by RQMC and LHS are no longer i.i.d. Therefore, neither sampling technique can be incorporated into SDDP without modification, if convergence properties are not to be compromised. Homem-de-Mello, De Matos, and Finardi [112] therefore suggest building sampling blocks. This means that the total number of samples  $|\mathcal{K}|$  is divided into  $M$  blocks  $\ell = 1, \dots, M$  with  $M \geq 5$  a divisor of  $|\mathcal{K}|$ . Then, for each block  $\ell$ ,  $|\mathcal{K}'| := |\mathcal{K}|/M$  scenarios are obtained using conditional sampling with RQMC or LHS that are not independent. For each  $k' \in \mathcal{K}'$ , values  $v^i(\xi^{k'})$  are determined and averaged to  $\bar{v}^{i,\ell}$ .

This is repeated for each block  $\ell$ . Then, the mean  $\bar{v}_{\mathcal{K}}^i$  of all values  $\bar{v}^{i,\ell}$ ,  $\ell = 1, \dots, M$ , and the sample variance are determined. As the scenarios of different blocks are independent, this still yields a useful confidence interval to stop the algorithm.

Another challenge reported in [112] is that it is computationally expensive to generate samples using RQMC for high dimensions. To reduce the computational effort, it might be reasonable to apply RQMC only to important components, e.g., to early stages in  $[T]$ , and standard MC or LHS to the others. This strategy is called *padding* and is applied after 6 or 12 stages for numerical tests in [112].

Experiments in [112] imply that RQMC and LHS both lead to upper bounds  $\bar{v}_{\mathcal{K}}$  oscillating around the lower bound  $\underline{v}$  more quickly than in MC sampling.

**6.3. Importance Sampling.** In [157], Parpas, Ustun, and Tran propose incorporating importance sampling into SDDP. In contrast to the previously described techniques, theirs can be used to obtain i.i.d. samples in the forward pass.

In general, the main idea of importance sampling is to attach different importances to subregions of the sample space and to sample more often from subregions of higher importance. In the context of SDDP, this means that it is sampled with priority from scenarios that contribute more to the value of the expected value functions  $\mathcal{Q}_t(\cdot)$ .

This is achieved by sampling from a distribution different from the original one, the so-called *importance sampling distribution*, but then correcting the bias introduced by this difference. Then, an importance sampling estimate of  $\bar{v}$  can be calculated as



$$\bar{v}_{\mathcal{K}}^{IS,i} := \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} v^i(\xi^k) \Lambda(\xi^k)$$

with  $\Lambda(\xi) := \frac{f(\xi)}{g(\xi)}$ , where  $f$  denotes the original distribution and  $g$  the importance sampling distribution. The likelihood function  $\Lambda(\cdot)$  is used to correct for sampling from the wrong distribution. It can be shown that importance sampling can reduce the variance of sampling estimators significantly. In the SDDP case, as shown in [157], the variance is minimized for the choice

$$g_t^*(\xi_t) := \frac{|Q_t(x_{t-1}^{ik}, \xi_t)|}{\mathbb{E}_f |Q_t(x_{t-1}^{ik}, \xi_t)|} f_t(\xi_t).$$

However, this *zero-variance distribution* is clearly a theoretical construct and is not known, which is referred to as the curse of *circularity*. Therefore, it is proposed to first approximate  $g^*$  using a framework including kernel density estimation [157].

In numerical experiments, SDDP with importance sampling is shown to outperform MC and QMC sampling-based methods, given that it is difficult to sample from the original probability distribution and that the original problem has moderate or high variance [157].

**6.4. Deterministic Sampling.** As discussed in section 4, in line 6 of SDDP (Algorithm 3.1) some deterministic sampling can also be used. In this case,  $|\mathcal{K}| = 1$ . In the literature, two different approaches are considered.

**Worst Approximation Sampling.** The first approach requires that in addition to the (lower) cut approximation  $\mathcal{V}_t(\cdot)$  of  $\mathcal{Q}_t(\cdot)$  an upper approximation  $\bar{\mathcal{V}}_t(\cdot)$  is also constructed and iteratively refined in SDDP. Assume that in the forward pass on stage  $t-1$  the trial point  $x_{t-1}^i$  has been computed. Then, for stage  $t$  the approximate subproblem (2.10) is solved for  $x_{t-1}^i$  and for all noise terms  $\xi_{tj}, j = 1, \dots, q_t$ , yielding optimal states  $x_{tj}$ . For the next stage, the trial point  $x_t^i = x_{tj'}$  is chosen such that

$$j' \in \arg \max_{j=1, \dots, q_t} \left\{ \bar{\mathcal{V}}_t^i(x_{tj}) - \mathcal{V}_t^i(x_{tj}) \right\},$$

i.e., that the gap between the current upper and lower approximations is maximized. This corresponds to sampling noise term  $\xi_{tj'}$  on stage  $t$ .

This form of deterministic sampling is used for SDDP in [240]. Its computational drawback is that at each stage  $q_t$  subproblems have to be solved instead of only  $|\mathcal{K}| \ll q_t$ . A similar approach is proposed by Baucke, Downward, and Zakeri in [10, 11] and called *problem-child node selection*. However, their setting differs somewhat from original SDDP, as each subproblem contains specific variables  $x_{tj}, j = 1, \dots, q_t$ , for all random outcomes, and therefore in their case only one subproblem has to be solved in the sampling step. Another related sampling scheme is used in robust *dual dynamic programming* (RDDP) [83]. In that case,  $\xi_{tj'}$  is determined by solving a special upper bounding problem containing  $\bar{\mathcal{V}}_t^i(\cdot)$

**Explorative Sampling.** Explorative deterministic sampling is proposed in [124] as part of EDDP. It is based on the concepts of saturated and distinguishable points introduced in section 4.2. As for the previous sampling scheme, the idea is to solve the forward pass subproblems for all  $\xi_{tj}, j = 1, \dots, q_t$ . Instead of maximizing an approximation gap, however, the trial point  $x_t^i = x_{tj'}$  is chosen such that

$$j' \in \arg \max_{j=1, \dots, q_t} \min_{x_t \in X_t^{\text{sat}}} \|x_{tj} - x_t\|,$$

i.e., the minimum distance to already saturated points is maximized. In other words, a maximum distinguishable point is chosen.

As shown in [75], worst approximation sampling and explorative sampling are equivalent in the sense that both approaches are guaranteed to lead to effective iterations; see section 4.2.

**7. Stopping Criteria.** In each iteration  $i$  of SDDP, a valid lower bound  $\underline{v}^i$  for the optimal value  $v^*$  is determined. In addition, a statistical upper bound  $\bar{v}_{\mathcal{K}}^i$  can be computed. Since the latter is not necessarily valid, an important question is when to consider an obtained policy  $(\mathbf{x}_t(\xi_{[t]}))_{t \in [T]}$  to be (approximately) *optimal* and to stop the SDDP method; see line 4 of Algorithm 3.1. If the stopping criterion is too conservative, the algorithm may iterate much longer than required, but if it is too optimistic, then SDDP may stop prematurely.

**Confidence Stopping Criteria.** In their seminal work on SDDP, Pereira and Pinto proposed using a confidence interval based stopping criterion [160]. An approximate confidence interval for a true valid upper bound  $\bar{v}^i$  is determined as follows using the estimates  $v^i(\xi^k)$  from (3.9).

Under random independent sampling, the values  $v^i(\xi^k)$  are i.i.d. random variables with expected value  $\bar{v}^i$  and variance  $(\sigma^i)^2$ . Moreover, knowing the sample mean  $\bar{v}_{\mathcal{K}}^i$  (3.9), we can define a standardized random variable

$$(7.1) \quad Z_{\mathcal{K}}^i := \frac{\bar{v}_{\mathcal{K}}^i - \bar{v}^i}{\frac{\sigma^i}{\sqrt{K}}}.$$

According to the central limit theorem, asymptotically, that is, for  $|\mathcal{K}| \rightarrow \infty$ , this random variable follows a standard normal distribution  $\mathcal{N}(0, 1)$ . This implies that for sufficiently large  $|\mathcal{K}|$ ,  $Z_{\mathcal{K}}^i$  is approximately standard normal distributed.

Due to symmetry of the standard normal distribution, it follows that

$$\mathbb{P}(-z_{1-\alpha/2} \leq Z_{\mathcal{K}}^i \leq z_{1-\alpha/2}) \approx 1 - \alpha,$$

where  $z_{1-\alpha/2}$  denotes  $(1 - \frac{\alpha}{2})$ -quantiles of  $\mathcal{N}(0, 1)$  for some level  $\alpha \in (0, 1)$ .

Inserting (7.1) and rearranging yields an approximate  $(1 - \alpha)$ -confidence interval for the true upper bound  $\bar{v}^i$ :

$$\left[ \bar{v}_{\mathcal{K}}^i - z_{1-\frac{\alpha}{2}} \frac{\sigma^i}{\sqrt{|\mathcal{K}|}}, \bar{v}_{\mathcal{K}}^i + z_{1-\frac{\alpha}{2}} \frac{\sigma^i}{\sqrt{|\mathcal{K}|}} \right].$$

As  $\sigma^i$  is unknown, it can be replaced by the sample standard distribution  $\sigma_{\bar{v}, K}^i$ , which is defined by the sample variance

$$(\sigma_{\bar{v}, K}^i)^2 := \frac{1}{|\mathcal{K}| - 1} \sum_{k \in \mathcal{K}} (v^i(\xi^k) - \bar{v}_{\mathcal{K}}^i)^2.$$

In that case, the standardized variable approximately follows a Student's  $t$ -distribution with degree of freedom  $|\mathcal{K}| - 1$ . In the literature on SDDP, even in this case, the  $(1 - \alpha)$ -confidence interval for the true upper bound  $\bar{v}^i$  is usually approximated using a standard normal distribution [208] though, which yields

$$(7.2) \quad \left[ \bar{v}_{\mathcal{K}}^i - z_{1-\frac{\alpha}{2}} \frac{\sigma_{\bar{v}, K}^i}{\sqrt{|\mathcal{K}|}}, \bar{v}_{\mathcal{K}}^i + z_{1-\frac{\alpha}{2}} \frac{\sigma_{\bar{v}, K}^i}{\sqrt{|\mathcal{K}|}} \right].$$

Pereira and Pinto propose choosing  $\alpha = 0.05$ , which implies  $z_{1-\alpha/2} = 1.96$ , and stopping SDDP if the lower bound  $\underline{v}^i$  is included in this confidence interval [160].

As pointed out by Shapiro [205], this stopping criterion has several flaws. The higher the sample variance  $(\sigma_{\underline{v}, \mathcal{K}}^i)^2$ , the earlier  $\underline{v}^i$  exceeds the lower end of the confidence interval, which provides a misguided incentive to increase  $(\sigma_{\underline{v}, \mathcal{K}}^i)^2$ . The same is true for increasing the confidence  $1 - \alpha$ , which contradicts the intuition behind  $\alpha$ . In addition, faster stopping can be achieved by reducing the sample size  $|\mathcal{K}|$ . Finally, the above stopping criterion may favor premature stopping, as it is rather unlikely that  $\bar{v}^i$  is located exactly at the lower bound of the confidence interval. For these reasons, Shapiro proposes a more conservative stopping criterion where SDDP terminates if the difference between the upper bound of the confidence interval (7.2) and  $\underline{v}^i$  is sufficiently small.

Sometimes including values  $v^j(\xi^k)$  from previous iterations  $j < i$  in (3.9) is also suggested, for instance, if  $|\mathcal{K}|$  is too small to obtain a reasonable bound. However, this destroys the independence between the different samples and thus the central limit theorem can no longer be applied and the confidence-based stopping criteria are not applicable. [53].

**A Hypothesis Test Perspective.** Considering that hypothesis tests and confidence intervals are closely related, the above stopping criterion can also be interpreted in terms of a hypothesis test with hypotheses [112]:

$$H_0 : \bar{v}^i = \underline{v}^i, \quad \text{against} \quad H_1 : \bar{v}^i \neq \underline{v}^i.$$

The null hypothesis  $H_0$  is tested using the test statistic  $\bar{v}_{\mathcal{K}}^i$ , which is assumed to be approximately normal distributed. This can again be reasoned using the central limit theorem for sufficiently large  $|\mathcal{K}|$ . Then, the region of acceptance for  $H_0$  in iteration  $i$  is given by the interval (7.2): If the lower bound  $\underline{v}^i$  does not exceed the lower bound of this region, then optimality is rejected. Otherwise, there is no compelling reason to reject it, so it is retained. By choosing  $\alpha$ , the type I error (rejecting optimality although SDDP has converged) can be controlled. However, this comes at the cost of a possibly high type II error (stopping the algorithm prematurely) [112].

**Different Hypothesis Tests.** To avoid stopping prematurely, a modified hypothesis test can be used that controls type I and type II errors simultaneously [112]. The basic principle is very similar to the above, even if it is presented for a one-sided hypothesis test with  $H_0 : \bar{v}^i \leq \underline{v}^i$ . The key difference is that if  $\underline{v}^i$  lies inside the region of acceptance, the hypothesis of optimality is not necessarily retained, but still may be rejected. In particular, stopping SDDP should be prevented if the true upper bound  $\bar{v}^i$  exceeds the lower bound  $\underline{v}^i$  considerably. As  $\bar{v}^i$  is not known, we cannot observe when this event occurs, but we can predefine a bound  $\gamma > 0$  on the probability of stopping given that it happens. For fixed  $\gamma$  and  $\alpha$  and given sample estimates, we can then compute a percentage difference  $\delta^i$  between  $\bar{v}^i$  and  $\underline{v}^i$ , for which the probability of a type II error (premature stopping) is bounded by  $\gamma$ :

$$(7.3) \quad \delta^i = (z_{1-\alpha} + z_{1-\gamma}) \frac{\sigma_{\bar{v}, \mathcal{K}}^i}{\underline{v}^i \sqrt{|\mathcal{K}|}}.$$

If  $\delta^i$  is below some predefined threshold  $\bar{\delta}$ , the sample estimates guarantee that for deviations larger than  $\bar{\delta}$ , the type II error is under control. Therefore, SDDP stops. Otherwise, the control of the type II error is not considered sufficient, and the algorithm proceeds. In other words, SDDP only terminates when  $\underline{v}^i$  lies inside the region

of acceptance and when the type II error is bounded by  $\gamma$  for a sufficiently small percentage difference  $\delta^i$ .

Summarized, the following procedure is used in each iteration  $i$  [112]:

1. Compute  $\rho^i$  as the ratio of the left interval boundary in (7.2) and  $\underline{v}^i$ .
2. If  $\rho^i \leq 1$ , then compute  $\delta^i$  according to (7.3).
  - a) If  $\delta^i < \bar{\delta}$ , then stop SDDP.
  - b) Otherwise, start a new iteration  $i + 1$  of SDDP (or adapt  $\alpha, \gamma$ , or  $\mathcal{K}$ ).

Otherwise, start a new iteration  $i + 1$  of SDDP.

Computational experiments with  $\bar{\delta} = 0.1$  and  $\gamma = 0.05$  indicate that this stopping criterion is effective in preventing SDDP from premature stopping [112]. Still, it is an heuristic and, thus far, no proposed statistical testing procedure guarantees that the probability of stopping prematurely is bounded by some  $\gamma > 0$  in general.

**Predefined Criteria.** The previous statistical stopping criteria are computationally demanding and require  $|\mathcal{K}|$  to be sufficiently large to yield reasonable approximate confidence intervals. Furthermore, in practical applications (MSLP) is often too large to achieve convergence in reasonable time. Finally, the statistical stopping criteria do not necessarily generalize to extensions of SDDP, such as risk-averse variants; see section 12. Therefore, in practice often more convenient stopping criteria are used for SDDP. For instance, it is common to stop SDDP after a fixed number of iterations  $I \in \mathbb{N}$ , after a fixed number of cuts  $|\mathcal{K}|I$ , after a predefined time, or if the lower bounds  $\underline{v}^i$  have stalled. None guarantees that an optimal policy is determined, though.

**Deterministic Stopping.** Finally, SDDP can be stopped deterministically as long as valid upper bounds  $\bar{v}^i$  for  $v^*$  are computed in addition to lower bounds  $\underline{v}^i$ . In that case, for some predefined optimality tolerance  $\varepsilon > 0$ , SDDP stops with an (approximately) optimal policy if  $\bar{v}^i - \underline{v}^i \leq \varepsilon$ . This stopping criterion requires significant additional computational effort to determine true upper bounds  $\bar{v}^i$ . Hence, there is a trade-off between achieving a more reasonable stopping criterion and spending computational resources on computations outside the core elements of SDDP. We address how such exact upper bounds can be computed in the next section.


In summary, despite various attempts at developing reasonable termination criteria for SDDP, optimally stopping SDDP remains an open challenge.

**8. Exact Upper Bounds and Upper Approximations.** Recently, the idea of computing deterministic upper bounds  $\bar{v}$  for  $v^*$  and deterministic upper approximations  $\bar{\mathcal{V}}_t(\cdot)$  of  $\mathcal{Q}_t(\cdot)$  has drawn a lot of interest in the research community, in terms of both analyzing the convergence behavior of SDDP (see section 4) and developing deterministic stopping criteria (see section 7).

An intuitive way to determine upper approximations  $\bar{\mathcal{V}}_t(\cdot)$  of  $\mathcal{Q}_t(\cdot)$  is based on the observation that due to convexity of  $\mathcal{Q}_t(\cdot)$ , all its secants lie above or on its graph. Therefore, an upper approximation is possible by a convex combination of points  $(x_{t-1}, \mathcal{Q}_t(x_{t-1}))$ . To obtain an approximation on the whole state space, it can be extended using a regularization with a Lipschitz constant  $L_t$  of  $\mathcal{Q}_t(\cdot)$ . Such a constant exists according to Corollary 2.10. In this light,  $\bar{\mathcal{V}}_t(\cdot)$  can be constructed as [240]

$$(8.1) \quad \bar{\mathcal{V}}_t(x_{t-1}) = \text{co} \left( \min_{m=1, \dots, M_t} \left\{ \mathcal{Q}_t(x_{t-1}^m) + L_t \|x_{t-1} - x_{t-1}^m\| \right\} \right),$$

where  $\text{co}(f)$  denotes the convex envelope of function  $f$ . This is illustrated in Figure 9.



**Fig. 9** Inner and outer approximations of  $\mathcal{Q}_t(\cdot)$ .

Alternatively, by interpreting this idea from a set perspective, the convex epigraph  $\text{epi}(\mathcal{Q}_t)$  of  $\mathcal{Q}_t(\cdot)$  can be approximated by the convex hull  $\text{conv}(w_{t-1}^1, \dots, w_{t-1}^{M_t})$  of finitely many points  $w_{t-1} := (x_{t-1}, \mathcal{Q}_t(x_{t-1}))$  in  $\text{epi}(\mathcal{Q}_t)$ .

In principle, there are two different approaches to realizing this idea. One uses the above perspectives, which we refer to as *primal*, and one is related to a dual perspective on SDDP and its value functions [104, 126].

**8.1. Primal Inner Approximation.** Similar to subproblems (2.10), based on upper approximations  $\bar{\mathcal{V}}_t(\cdot)$  of  $\mathcal{Q}_t(\cdot)$ , approximating subproblems can be defined by replacing  $\mathcal{Q}_t(\cdot)$  with  $\bar{\mathcal{V}}_t(\cdot)$  in the DPE for all  $t \in [T]$ . This idea was first introduced by Philpott and de Matos [168]. Since they consider only the RHS of (MSLP) to be uncertain, we adopt this assumption, although it is not required.

For stages  $t = T - 1, \dots, 2$ , each element  $m$  in a given set of points  $x_t^1, \dots, x_t^{M_{t-1}}$ , and each  $\xi_{tj}, j = 1, \dots, q_t$ , the following subproblem can be solved by backward recursion:

$$(8.2) \quad \bar{\mathcal{Q}}_t(x_{t-1}^m, \xi_{tj}) := \begin{cases} \min_{x_t} & c_t^\top x_t + \bar{\mathcal{V}}_{t+1}(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}^m, \xi_{tj}). \end{cases}$$

Here, as indicated in (8.1), the upper approximation  $\bar{\mathcal{V}}_{t+1}(\cdot)$  is defined as a convex combination of points  $(x_t^m, \bar{\mathcal{Q}}_{t+1}(x_t^m)), m = 1, \dots, M_t$ . The key difference is that instead of  $\mathcal{Q}_{t+1}(x_t^m)$ , here  $\bar{\mathcal{Q}}_{t+1}(x_t^m) := \mathbb{E}[\bar{\mathcal{Q}}_{t+1}(x_t^m, \xi_{t+1})]$  is used, as  $\mathcal{Q}_{t+1}(\cdot)$  is not known:

$$(8.3) \quad \bar{\mathcal{V}}_{t+1}(x_t) := \begin{cases} \min_w & \sum_{m=1}^{M_t} w_m \bar{\mathcal{Q}}_{t+1}(x_t^m) \\ \text{s.t.} & \sum_{m=1}^{M_t} w_m x_t^m = x_t, \\ & \sum_{m=1}^{M_t} w_m = 1, \\ & w_m \geq 0, \quad m = 1, \dots, M_t. \end{cases}$$

Furthermore, in comparison to (8.1), no regularization is used.

By recursion, it can be shown that

$$\overline{Q}_t(x_{t-1}^m, \xi_{tj}) \geq Q_t(x_{t-1}^m, \xi_{tj})$$

for all  $m = 1, \dots, M_{t-1}$  and  $j = 1, \dots, q_t$ . This implies that

$$\overline{\mathcal{Q}}_t(x_{t-1}^m) \geq \mathcal{Q}_t(x_{t-1}^m).$$

The first-stage problem then yields

$$(8.4) \quad \overline{v}^{IA} := \begin{cases} \min_{x_t} & c_1^\top x_1 + \overline{\mathcal{V}}_2(x_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1, \end{cases}$$

with  $\overline{v}^{IA}$  an exact valid upper bound to  $v^*$ .

Based on these definitions, in principle, the backward pass of SDDP can be enhanced at each stage  $t = T-1, \dots, 2$  by solving subproblems (8.2) for all  $m = 1, \dots, M_t$  and each  $\xi_{tj}, j = 1, \dots, q_t$ , and updating  $\overline{\mathcal{V}}_t(\cdot)$  according to (8.3). At the first stage, by solving subproblem (8.4), an upper bound can then be computed.

A key challenge with this approach is the appropriate choice of the set of points  $x_{t-1}^m, m = 1, \dots, M_{t-1}$ , for each stage  $t$  and iteration  $i$ . On the one hand, they should be chosen such that as much of  $\mathcal{X}_{t-1}$  is spanned as possible. On the other hand, choosing (at least some of) those points as extreme points leads to  $M_t \geq 2^{n_t}$  points, i.e., the number of points grows exponentially in the dimension of the state space.

An alternative is to use the trial points from the SDDP forward pass [168]. However, even when using these points the computational effort may become excessive. Similarly to the SDDP backward pass, subproblems (8.2) have to be solved for each stage  $t \in [T]$ , each point  $x_{t-1}^m, m = 1, \dots, M_{t-1}$ , and each noise term  $\xi_{tj}, j = 1, \dots, q_t$ . However, the number  $M_{t-1}$  of points to be considered grows with each iteration, as it contains all previous trial solutions. Therefore, a suggestion is to use the upper bound computation only every few hundred iterations and to not permanently incorporate it into the backward pass [168]. In return, this hinders using the upper bounds  $\overline{v}^{IA}$  in the stopping criterion of SDDP in each iteration.

Moreover, the obtained bounds  $\overline{v}^{IA}$  may be very loose, especially in problems (MSLP) with a large number of stages. Computational tests are required to assess whether or not the information gain justifies the additional computational effort and possibly a higher number of iterations.

Baucke, Downward, and Zakeri provide a different perspective on the inner approximation idea above [10]. Instead of (8.3), they use its dual representation

$$(8.5) \quad \overline{\mathcal{V}}_{t+1}(x_t) = \begin{cases} \max_{\mu, \lambda} & x_t^\top \lambda + \mu \\ \text{s.t.} & (x_t^m)^\top \lambda + \mu \leq \overline{\mathcal{Q}}_{t+1}(x_t^m), \quad m = 1, \dots, M_t. \end{cases}$$

This shows that  $\overline{\mathcal{V}}_{t+1}(\cdot)$  can be equivalently described by maximizing over the coefficients of all supporting hyperplanes for points  $(x_t^m, \overline{\mathcal{Q}}_{t+1}(x_t^m)), m = 1, \dots, M_t$ .

In [10], the dual problem is additionally regularized, i.e., enhanced by the constraint

$$\|\lambda\| \leq L_t,$$

with  $L_t$  denoting a Lipschitz constant of  $\overline{\mathcal{Q}}_t(\cdot, \cdot)$ . This is equivalent to regularizing the primal problem (8.3) with the dual norm to  $\|\cdot\|$ ; cf. (8.1). In this way, a reasonable

approximation is also achieved for points outside the convex hull of the set defined by the points  $x_t^m, m = 1, \dots, M_t$ .

Using this expression for the inner approximation functions, Baucke, Downward, and Zakeri propose a deterministic algorithm for multistage stochastic convex programs. In their case, subproblems (8.2) are solved in each backward pass iteration and  $\bar{V}_{t+1}^i(\cdot)$  is updated by adding constraint  $(x_t^{\tilde{m}})^\top \lambda + \mu \leq \bar{\mathcal{Q}}_{t+1}^i(x_t^{\tilde{m}})$  for the current iterate  $x_t^{\tilde{m}}$  (to incorporate it into the subproblems (8.2), however, its dual (8.3) is again used). The proposed algorithm differs in further aspects from standard SDDP; for instance, it requires a multicut approach (see section 21) and uses a worst approximation sampling (see section 6.4). Moreover, choosing a reasonable and valid value for  $L_t$  can be very challenging, but is crucial for the proposed method to work as intended.

**8.2. Dual SDDP.** With respect to computing deterministic upper bounds  $\bar{v}$  for  $v^*$ , a dual perspective on SDDP and the DPE (2.4) has recently gained attention.

**Using Convex Conjugates of Value Functions.** The first proposal in this context, by Leclère et al. [126], exploits convex conjugates and the related duality concepts to derive *dual value functions* for (MSLP) in which uncertainty only appears in the RHS  $h_t(\xi_t)$ .

Let  $f: \mathbb{R}^n \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ ; then its *convex conjugate*  $f^*(\cdot)$  is defined as [192]

$$f^*(\lambda) := \sup_{x \in \mathbb{R}^n} \lambda^\top x - f(x).$$

For (MSLP), the convex conjugates  $D_t(\cdot) := Q_t^*(\cdot)$  of the value functions  $Q_t(\cdot)$  can be considered as dual value functions for  $t = 2, \dots, T$ . It can be shown that these functions also satisfy some DPE with linear subproblems on each stage. Whereas Leclère et al. consider a more general setting including control variables  $u_t$  (see Remark 2.2), for (MSLP) as defined in section 2 (and especially under Assumption 5) for  $t = 2, \dots, T$ , these subproblems can be expressed by

$$(8.6) \quad D_t(\lambda_{t-1}) := \begin{cases} \min_{\lambda_t, \mu_t, \gamma_t} & \sum_{j=1}^{q_t} p_{tj} \left( -h_{tj}^\top \mu_{tj} + D_{t+1}(\lambda_{tj}) \right) \\ \text{s.t.} & T_{t-1}^\top \left( \sum_{j=1}^{q_t} p_{tj} \mu_{tj} \right) - \sum_{j=1}^{q_t} p_{tj} \gamma_{tj} + \lambda_{t-1} = 0, \\ & W_t^\top \mu_{tj} = \lambda_{tj} + c_t, \quad j = 1, \dots, q_t, \\ & \gamma_{tj} \leq 0, \quad j = 1, \dots, q_t. \end{cases}$$

For the first stage, we obtain a deterministic problem which by  $T_0 \equiv 0$  simplifies to

$$D_1(\lambda_0) = \min_{\mu_1} h_1^\top \mu_1 + D_1(W_1^\top \mu_1 - c_1)$$

for some arbitrary initial  $\lambda_0 \leq 0$  (note that more general formulations of (MSLP) may lead to a dependence on  $\lambda_0$ ).

Using this dynamic recursion, it is possible to apply an SDDP-type algorithm, called dual SDDP, to  $D_t(\cdot)$ , using iteratively improving outer approximations  $\mathfrak{D}_t^i(\cdot)$  for  $D_t(\cdot)$ . Analogously to SDDP, this iterative method yields a converging deterministic lower bound for the first-stage optimal value, i.e.,  $\mathfrak{D}_1^i(\lambda_0) \leq D_1(\lambda_0)$ . Again applying conjugacy theory, we obtain

$$\bar{v}^i = (\mathfrak{D}_1^i)^*(x_0) \geq D_1^*(x_0) = Q_1^{**}(x_0) = Q_1(x_0) = v^*.$$

Hence, deterministic upper bounds for  $v^*$  can be obtained as conjugates of the first-stage approximations  $\mathfrak{D}_t^i(\cdot)$  evaluated at  $x_0 = 0$ , and  $(\bar{v}^i)_i$  defines a sequence converging to  $v^*$  [126].

**Using the Dual of (MSLP).** Guigues and Sagastizábal propose an alternative way to define dual value functions and DPE that can be exploited in a dual SDDP algorithm [104]. Instead of working with conjugates of the primal value functions  $Q_t(\cdot)$ , they first derive the dual to (MSLP) formulated as a single problem (2.3) and then show that this dual problem can be decomposed using DPE and dual value functions:

$$(8.7) \quad \tilde{D}_t(\pi_{t-1}) := \begin{cases} \max_{\pi_t} & \sum_{j=1}^{q_t} p_{tj} \left( -h_{tj}^\top \pi_{tj} + \tilde{D}_{t+1}(\pi_{tj}) \right) \\ \text{s.t.} & \sum_{j=1}^{q_t} p_{tj} \left( T_{t-1,j}^\top \pi_{tj} \right) + W_{t-1}^\top \pi_{t-1} \leq c_{t-1}. \end{cases}$$

It can be argued that these dual DPE are simpler and more intuitive, as they do not require conjugacy theory. Moreover, we immediately obtain that the first-stage optimal value  $\tilde{D}_1(\pi_0)$  equals  $v^*$  by strong duality for LPs. Therefore, using outer approximations  $\mathfrak{D}_t^i(\cdot)$  of these value functions in dual SDDP, again a sequence  $(\bar{v}^i)_i$  of deterministic and valid upper bounds can be computed that converges to  $v^*$  [104]. On the other hand, the dual value functions  $\tilde{D}_t(\cdot)$  cannot be directly related to the original value functions  $Q_t(\cdot)$ .

*Remark 8.1.* Even if the dual DPE (8.6) and (8.7) are derived using different tools and perspectives, they are still closely related. Note that subproblem (8.6) can be reformulated as

$$D_t(\lambda_{t-1}) = \begin{cases} \min_{\lambda_t, \mu_t} & \sum_{j=1}^{q_t} p_{tj} \left( -h_{tj}^\top \mu_{tj} + D_{t+1}(\lambda_{tj}) \right) \\ \text{s.t.} & T_{t-1}^\top \left( \sum_{j=1}^{q_t} p_{tj} \mu_{tj} \right) + \lambda_{t-1} \leq 0, \\ & W_t^\top \mu_{tj} = \lambda_{tj} + c_t, \quad j = 1, \dots, q_t. \end{cases}$$

Using the last constraint, the state  $\lambda_{t-1}$  can be expressed through the dual variables  $\mu_{t-1}$  from the previous stage:  $\lambda_{t-1} = W_{t-1}^\top \mu_{t-1} - c_{t-1}$ . Exploiting this, the subproblems only contain dual variables  $\mu_t$ , which have to be considered as state variables. By switching from max to min in the objective, we obtain the structure of (8.7) exactly.

We can make the following additional observations with respect to the dual DPE (8.6) and (8.7). First, in both cases the subproblems are not necessarily bounded and therefore artificial bounds are introduced. In [104] they are chosen as  $\pi_t \in [\underline{\pi}_t, \bar{\pi}_t]$ , whereas in [126] Lipschitz continuity of  $\mathcal{Q}_t(\cdot)$  is exploited to impose the bounds  $\|\lambda_t\|_\infty \leq L_t$  for Lipschitz constants  $L_t, t = 2, \dots, T$ . It is assumed that these bounds are chosen sufficiently large as to not affect the optimal solutions.

Second, even if the primal DPE (2.4) are assumed to have relatively complete recourse (see Assumption 9 and Lemma 2.5), this does not necessarily translate to the dual subproblems. To ensure feasibility, Guigues, Shapiro, and Cheng propose using either feasibility cuts (see section 17) or a penalization approach [104].

Third, in contrast to the primal perspective, the subproblems do not decompose by realizations of  $\xi_t$ , but instead contain separate dual variables  $\pi_{tj}$  (or  $\lambda_{tj}, \mu_{tj}, \gamma_{tj}$ ,



respectively) for all  $j = 1, \dots, q_t$ . In the forward pass of dual SDDP, the trial point  $\pi_t^i$  (or  $\lambda_t^i$ ) that is used as a parameter in the following stage is sampled from these variables. A decoupling can be achieved using a Lagrangian relaxation [44].

Finally, if  $W_t$  and  $c_t$  become uncertain as well, then the value functions and subproblems depend additionally on  $\xi_t$ . In fact, in formulation (8.7) the state space has to be extended to include the history  $\xi_{t-1}$  of the stochastic process, as the problem contains  $W_{t-1}$  and  $c_{t-1}$  [104].

Again, an SDDP-type algorithm, also referred to as *dual SDDP* in [104], can be applied to the DPE (8.7). This algorithm is presented in Algorithm 8.1. The two variants of dual SDDP have been extended to the risk-averse case [43] (see section 12) and to problems with infinite horizon (see section 19) [207].

**Dual Inner Approximation.** First and foremost, dual SDDP is an alternative to (primal) SDDP for approximating  $v^*$  by converging deterministic upper bounds  $\bar{v}^i$ . However, as shown in [126], if the dual DPE (8.6) are used, then the approximations  $\mathfrak{Q}_t^i(\cdot)$  obtained may be translated into inner approximations  $\bar{\mathcal{V}}_t^i(\cdot)$  of the primal value functions  $Q_t(\cdot)$ . In this way, policies  $(\mathbf{x}_t(\xi_{[t]}))_{t \in [T]}$  for (MSLP) can be computed. The inner approximations can be computed as Lipschitz regularizations (see section 17) of the convex conjugate of the outer approximations  $\mathfrak{Q}_t^i(\cdot)$ , which is shown to be equivalent to solving problem (8.5) with regularization  $\|\lambda\|_\infty \leq L_t$ . The key difference to the approach in [10] is the way the primal supporting points  $x_t^m$  are determined; that is, by the slopes of the dual outer approximation [126].

**Incorporation into SDDP.** While dual SDDP can be applied on its own to approximate  $v^*$ , and even compute policies  $(\mathbf{x}_t(\xi_{[t]}))_{t \in [T]}$ , it seems reasonable to incorporate it into (primal) SDDP in order to compute deterministic upper and lower bounds for  $v^*$ . Guigues, Shapiro, and Cheng suggest using both variants of SDDP in parallel [104]. In contrast, Leclère et al. propose a framework in which primal and dual SDDP are intertwined [126]:

1. Run a forward pass of (primal) SDDP, yielding trial solutions  $(x_t^i)_{t \in [T]}$  for the sampled scenario path (the authors choose  $|\mathcal{K}| = 1$ ).
2. Run a backward pass of (primal) SDDP using the trial solutions  $x_{t-1}^i$ , obtaining new slopes  $\pi_t^i$  from the cuts.
3. Run a backward pass of dual SDDP using the slopes  $\lambda_t^i = \pi_t^i$ , obtaining new cuts for the dual problem.
4. Run a forward pass of dual SDDP to obtain a new dual trajectory  $(\tilde{\lambda}_t^i)_{t \in [T]}$  and update the cuts along this trajectory.

One computational drawback of this framework, and of dual SDDP in general, is that each iteration of dual SDDP is much more computationally expensive than in standard (primal) SDDP. This hampers the application of a solely deterministic stopping criterion for very large problems [104, 126].

**9. Applications.** In this section, we present different application areas of SDDP. We also point out applications for which some of the Assumptions 1 to 9 are not satisfied, so either modifications of (MSLP) or algorithmic extensions are required in order to apply SDDP. These use cases can be regarded as a motivation for the enhancements of SDDP that we cover in the following sections.

**9.1. Power System Optimization.** By far the dominating application field of SDDP is power system optimization, in particular, the operational planning by a central planner of energy systems including hydro storage. This is due to the adequacy

---

**Algorithm 8.1.** Dual SDDP from [104].

---

**Input:** Dual to problem (MSLP) satisfying Assumptions 1 to 9. Appropriate multiplier bounds. Stopping criterion.

**Initialization**

---

- 1: Initialize cut approximations with bounded  $\tilde{\mathfrak{D}}_t^0(\cdot)$  for all  $t = 2, \dots, T$ .
- 2: Initialize upper bound with  $\bar{v}^0 = +\infty$ .
- 3: Set iteration counter to  $i \leftarrow 0$ .

**Dual SDDP Loop**

---

- 4: **while** Stopping criterion not satisfied **do**
- 5:   Set  $i \leftarrow i + 1$ .
- Forward Pass**

---

- 6:   Solve the first-stage problem (defined by replacing  $\tilde{D}_2(\cdot)$  with  $\tilde{\mathfrak{D}}_2^i(\cdot)$  and adding multiplier bounds in (8.7)). Store the trial point  $\pi_1^i$ .
- 7:   **for** stages  $t = 2, \dots, T$  **do**
- 8:     Solve the stage- $t$  subproblem (defined by replacing  $\tilde{D}_{t+1}(\cdot)$  with  $\tilde{\mathfrak{D}}_{t+1}^i(\cdot)$  and adding multiplier bounds in (8.7)) for  $\pi_{t-1}^i$  to obtain  $\pi_{tj}^i, j = 1, \dots, q_t$ .
- 9:     Sample  $\tilde{j}$  from  $j = 1, \dots, q_t$  and set  $\pi_t^i = \pi_{t\tilde{j}}^i$ .
- 10:   **end for**
- Backward Pass**

---

- 11:   **for** stages  $t = T, \dots, 2$  **do**
- 12:     Solve the updated stage- $t$  subproblem (2.10) (defined by replacing  $\tilde{D}_{t+1}(\cdot)$  with  $\tilde{\mathfrak{D}}_{t+1}^{i+1}(\cdot)$  and adding multiplier bounds in (8.7)) for  $\pi_{t-1}^i$ . Store the optimal value  $\bar{D}_t(\pi_{t-1}^i)$  and the optimal dual vector  $x_{t-1}^i$ .
- 13:     Compute

$$\alpha_t^{D,i} := \bar{D}_t(\pi_{t-1}^i) - (\beta_t^{D,i})^\top \pi_{t-1}^i,$$

and

$$\beta_t^{D,i} := -W_{t-1}x_{t-1}^i.$$

- 14:     Update the cut approximation of  $\tilde{D}_t(\cdot)$  to

$$\tilde{\mathfrak{D}}_t^{i+1}(x_{t-1}) := \min \left\{ \tilde{\mathfrak{D}}_t^i(x_{t-1}), \alpha_t^{D,i} + (\beta_t^{D,i})^\top \pi_{t-1}^i \right\}.$$

- 15:   **end for**
- 16:   Solve the first-stage problem (defined by replacing  $\tilde{D}_2(\cdot)$  with  $\tilde{\mathfrak{D}}_2^{i+1}(\cdot)$  and adding multiplier bounds in (3.8)) to obtain an upper bound  $\bar{v}^i$ .
- 17: **end while**

**Output:** Upper bound  $\bar{v}^i$  for  $v^*$ .

---

of SDDP for such problems, but also due to its origins in optimizing the operational planning of the Brazilian hydrothermal system [159, 160].

In general, solving power system optimization problems is a very complex task, as it requires the incorporation of various technical and economical details and uncertainties [117, 156, 188, 189, 190, 216, 217, 241]. Incorporating all these

aspects into one single problem is computationally intractable. Therefore, a hierarchy of problems dealing with different time scales and perspectives [51] is usually considered, such as short-term dispatch (a few days or weeks), mid-term operational planning (1–2 years), and long-term operational planning (3–5 years) [76, 88]. Results from a long-term model can then be incorporated into a model with a shorter time horizon but with more detail in the other modeling aspects.

**9.1.1. Long-Term Operational Planning.** SDDP is most prominently used for the long-term operational planning (LTOP) of hydrothermal power systems, also called long-term hydrothermal scheduling (LTHS). In the research literature, SDDP has been applied to the LTOP of various hydrothermal systems, the most prominent being the hydropower-dominated systems in Brazil [15, 32, 33, 34, 46, 51, 52, 53, 56, 91, 104, 112, 133, 134, 136, 140, 143, 168, 173, 210, 211, 212, 215, 223, 233], other Central or South American countries [6, 74, 185, 220], Norway [87, 194], and New Zealand [167, 169, 237]. In addition, to this day, SDDP is used in practice by the Brazilian system operator ONS [141, 142].

The aim of LTOP is to determine an optimal policy for the amount of power to be generated by thermal and hydroelectrical utilities over some planning horizon of several years (usually with monthly resolution) such that demand is satisfied, technical constraints are fulfilled, and the expected cost is minimized [167]. The main focus is on managing hydro reservoirs and thus the water resource efficiently. This is not trivial. While there is an incentive to use all the water in a reservoir immediately, so no fuel costs are incurred, the potential value of storing water for later stages also has to be considered, taking into account the uncertainty of future inflows. For this reason, it can be beneficial to retain water in wet periods for following dryer periods. The ability to store water in reservoirs leads to a temporal coupling of the stages. The number of inflow realizations  $q_t$  per stage is typically chosen in a range between 20 and 100. For  $T = 60$ , this yields a scenario tree with about  $1.15 \cdot 10^{78}$  or  $10^{120}$  scenarios. For each forward pass, either a single scenario [52] or 100 to 200 scenarios are sampled.

LTOP can be used to illustrate some of the challenges and limits of standard SDDP and thus motivate the necessity for extensions.

**Autoregressive Uncertainty.** In LTOP, the main source of uncertainty is future (usually monthly) inflows into the reservoirs. These inflows often show seasonality and a temporal or spatial coupling which has to be considered in the modeling. Therefore, *autoregressive* (AR) processes are usually used to model and forecast them, in particular, *periodic autoregressive* (PAR) [140, 141] and related models [146]. This means that for each reservoir and each month a different AR model is fitted or, in other words, that the parameters in the AR model are allowed to differ between months.

In addition, hydro reservoirs are often organized in cascade systems, where the generation of one turbine may affect the inflow of downstream reservoirs such that they cannot be managed separately. For this reason, inflows often not only show temporal correlation and seasonality, but also spatial correlation. To address this, instead of PAR, *spatial periodic autoregressive* (SPAR) models can be used [134]. These models are still linear, but instead of only autoregressive components, i.e., lags of  $\xi_{it}$  for some reservoir  $i$ , lags of the inflows of neighboring reservoirs  $i'$  are also used to explain  $\xi_{it}$ . Apart from inflow lags, different exogenous variables, such as climate indices, precipitation, or sea temperature, can also be used to explain inflows [131, 172].

Whenever an AR process is used for the uncertain data, the assumption of stage-wise independence (Assumption 2) is not satisfied. This motivates an extension of

SDDP that is able to handle stagewise-dependent uncertainty, which is discussed in section 14.

**Nonlinear Uncertainty.** When modeling hydro inflows, the error terms in the AR process are usually assumed to be i.i.d. with normal or log-normal distribution [51, 134]. In the latter case, the model is also referred to as a geometric PAR (GPAR) model [136]:

$$(9.1) \quad \ln(\xi_t) = \gamma_t + \Phi_t \ln(\xi_{t-1}) + \eta_t.$$

GPAR models are usually more accurate in modeling inflows, which often tend to positive skewness and are thus not normally distributed. Moreover, they have the advantage that the requirement for nonnegative inflows is naturally satisfied.

On the other hand, solving (9.1) for  $\xi_t$  yields an AR process with multiplicative instead of additive error terms [211], i.e., a nonlinear model. Incorporating this into the DPE destroys the convexity of  $\mathcal{Q}_t(\cdot)$ , making a direct application of SDDP impossible. Instead, the nonlinear model has to be approximated linearly [211]. Another idea is to normalize the inflows first using a Box–Cox transformation. However, as such a transformation is nonlinear, a linear approximation is still required afterwards [175]. Further strategies to avoid nonnegative inflows and nonlinearities are discussed in [51, 183]. [49] suggests applying bootstrapping to resample directly from the historical residuals instead of applying a nonlinear transformation.

**Continuous Uncertainty and Distributional Uncertainty.** As stated above, a normal or log-normal distribution is usually assumed for the error terms in the inflow models, both being continuous distributions (an exception is [178], where inflows are modeled as a continuous process with discrete random errors). For this reason, the assumption of finite discrete random variables (Assumption 5) is not satisfied. In addition, the chosen distribution for the model may not coincide with the *true distribution* of the uncertain data. This raises questions about how to handle continuous uncertainty and distributional uncertainty in SDDP, which we address in sections 11 and 13.

**Computational Performance.** Despite the amenities of SDDP, its performance may suffer for problems with a large number of state variables, due to its exponential complexity in the state dimension  $n_t$ ; see section 4.2. For instance, SDDP is computationally prohibitive for a complete model of the Brazilian energy system, which consists of about 150 thermal plants and more than 150 hydro storages [52]. This is aggravated if the state dimension is artificially increased, e.g., in order to deal with stagewise-dependent uncertainty; see section 14. As a remedy, it is common practice to aggregate reservoirs based on their region and hydrological properties into so-called *energy equivalent reservoirs* (EERs) [4], thus reducing the state dimension [142]. However, this comes with an increased abstraction and may lead to suboptimal policies. Moreover, as outlined in [51], the EER modeling may introduce some nonlinearities into the system that have to be mitigated by linearization.

The computational complexity with respect to the state space also makes indispensable general performance improvements for SDDP, which we discuss in section 21.

**End-of-Horizon Effect.** Another challenge when applying SDDP to LTOP in practice is the so-called *end-of-horizon effect*. It relates to the effect that the policies obtained do not guarantee a continuous and reliable energy supply *after* the planning period, because in an optimal policy, all energy remaining in the reservoirs will be used at the end of the planning period. A typical planning horizon for LTOP is 5 years with a monthly resolution, leading to 60 stages. A common practice to mitigate

the end-of-horizon effect is to add 60 more stages to the problem, i.e., to consider a problem with 120 stages [211], even if decisions for only the first half are about to be implemented. Alternatively, it seems natural to analyze how SDDP can be applied for problems with an infinite horizon or with a random horizon, where Assumption 1 is not satisfied. We address this in sections 19 and 20.

**Risk-Aversion.** Due to the high importance of system reliability and stability in preventing outages and electricity shortages, system planners may favor more risk-averse policies compared to the risk-neutral policies obtained by standard SDDP. Therefore, recently there has been increased interest in taking risk-aversion into account when applying SDDP to LTOP [115, 215]. However, as Assumption 8 is no longer satisfied, this requires the extension of standard SDDP to a risk-averse variant. We discuss different approaches to achieving this in section 12.

**9.1.2. Medium-Term Operational Planning.** Structurally, medium-term operational planning (MTOP) problems do not differ much from LTOP. The main difference is that a shorter one- or two-year time horizon is considered [51, 167, 168, 185].

**Price Uncertainty in the Objective.** Especially on a medium-term time horizon, SDDP has also been adapted from the traditional setting with a single system operator to more market-driven systems in which several electricity suppliers are active. In such systems, besides inflows, spot prices can also be considered uncertain. This imposes an additional challenge to SDDP, as it leads to stagewise-dependent uncertainty in the objective. We discuss this in detail in section 14. To deal with this challenge, for instance, for the operational planning of the Norwegian hydro-storage system, a combined SDP/SDDP approach is commonly used [86, 87, 88, 107, 108].

**Water Head Effect.** In LTOP the so-called *water head effect* of hydro storage is often disregarded, but it may become decision-relevant in MTOP. This effect concerns the fact that the production of a hydro plant increases with the net head of the reservoir. As this production function is multiplied with the water discharge, it introduces nonconvexities into the problem. Therefore, if this nonlinear effect is explicitly considered, suitable extensions of SDDP to nonconvex problem are required [38, 110, 170]. We cover such extensions in section 16.

**9.1.3. More Energy Applications.** We briefly summarize further applications of SDDP in power system optimization.

**Short-Term Dispatch.** SDDP is particularly suited for long-term planning, but it can also be applied to short-term economic dispatch problems [39, 55, 125, 155]. For shorter time horizons, it may be reasonable to include additional system aspects, for instance, power flow and security constraints, reserve energy, or different ancillary services [139, 223]. If security constraints are considered, usually linear DC power flow models are used, but recently AC power flow has also gained interest [119, 193].

Another research stream considers CO<sub>2</sub> emissions, which can be covered by imposing an emission quota system [14, 186, 184] or by introducing emission trading [187]. The first approach leads to an (MSLP) that has no block-diagonal structure (Assumption 7). We discuss how SDDP can be applied in this case in section 18.

Using a reasonable extension to mixed-integer programs (see section 16), unit commitment problems are also accessible using the SDDP formulation [243].

**Different Storage Systems.** As different types of storage systems can be modeled similarly to hydro storage, SDDP is also applicable to such systems, for instance, to optimize gas storage facilities [234] or energy storage in microgrids [22].

**Optimal Bidding.** Instead of minimizing expected system cost from the perspective of a central system operator, in strategic bidding problems power plant operators attempt to determine an optimal bidding policy in order to maximize their expected revenue, while taking into account information uncertainty, for example, with respect to inflows or the market-clearing price; see [219, 221] for an overview.

Since the future revenue functions of the price-maker are sawtooth shaped, the resulting problem is nonconvex [220]. Therefore, to apply the SDDP idea, tailor-made extensions are required, e.g., convexifications or approximations by saddle cuts [59] or step functions [170, 237]. For methodological details, we refer to section 16.

Recently, applying SDDP to optimize trading in continuous intraday markets has also gained attention [213].

**Investment Planning.** An important long-term optimization problem in power systems is to make optimal (risk-averse) investment decisions, with respect to either the expansion of renewables [35, 130, 222] or conventional projects.

For conventional power systems, common investment problems address the questions of generation expansion and transmission expansion. The main challenge with such problems is that they naturally impose the introduction of integer decision variables. Therefore, in such a case relaxations [154] or appropriate extensions of SDDP, e.g., SDDiP [243], have to be used (see section 16). Alternatively, SDDP can be incorporated into a larger Benders decomposition framework, where at the first-stage binary investment decisions are taken and, at the second stage, a multistage stochastic LP is solved by SDDP [184]. Similar applications are considered in [56] and [45] with a special focus on risk and reliability constraints.

**Coping with Renewable Uncertainty.** An increasing share of renewable energy sources introduces more variability to an energy system which has to be taken into account and balanced by appropriate mechanisms. The usage of distributed grid-level storage, such as batteries or electric vehicles, for smoothing out the variable generation of renewables is examined using SDDP in [70, 244].

**9.2. Water Resource Management.** In many energy applications of SDDP, managing water resources plays a key role as it couples subsequent stages. Apart from energy optimization, SDDP is also applied to more general water resource management problems, where not only energy production but also water usage for irrigation in agriculture [162, 228], flow requirements for navigation [228], groundwater [144], or ecological constraints [227] are taken into account in the operational planning of reservoirs. Also related is the problem of river basin management [195].

Additionally, SDDP is used for assessing various quantities in hydrological systems, e.g., the value of water [231], risk for dam projects [2, 230], resource vulnerabilities [196], or benefits and costs of cooperation or noncooperation [145, 229].

**9.3. Portfolio Management.** The optimal management of a portfolio of investments, also referred to as asset allocation, can be modeled as an (MSLP) [47]. The aim is to distribute a fixed investment sum among a finite number of assets with uncertain returns in such a way that the expected return at the end of the considered horizon is maximized. By selling or buying certain amounts of assets, the investor can restructure his portfolio in each time period. Generally, both operations are associated with transaction costs, which leads to a very complex problem [232].

In the literature on SDDP, asset allocation problems are quite popular in testing proposed improvements and enhancements of SDDP, such as regularization [97], cut-sharing [91], or inexact cuts [8]. Since most investors are risk-averse, asset allocation

problems are a popular application [64, 67, 68, 113, 120, 121], but they are also one of the main drivers for the development of risk-averse SDDP, which we introduce in section 12.

For applications of practical interest, asset allocation becomes very challenging, as is pointed out in [232]. First, risk-aversion parameters such as  $\lambda_t$  or  $\alpha_t$  (see section 12) cannot be chosen intuitively such that the true preferences of an investor are appropriately represented. For this reason, the authors propose solving a risk-constrained model with one-period conditional AVaR constraints instead of taking the usual risk-averse SDDP approach. Second, the assumption of stagewise independence of asset returns may prove unrealistic, requiring a more sophisticated approach such as incorporating a Markov chain; see section 14. Moreover, the large supply of potential assets leads to a high-dimensional state space.

**9.4. Further Applications.** Although the focus is on the previous applications, occasionally other types of applications are also investigated using SDDP. Among those applications are dairy farming [65, 90], newsvendor problems [5, 157], inventory management [8, 63, 94, 104], lot-sizing [225], and routing problems [64]. [54] and [243] explore airline revenue management, which is an established problem in dynamic programming but requires integer variables.

**10. Software.** Until recently, SDDP implementations have been solely restricted to closed research projects or commercial products. For commercial products, the most established is the SDDP implementation by PSR, a Brazilian energy consultancy [179]. For research projects, various different implementations exist that cover programming languages like AMPL, C++, Fortran, GAMS, Java, and MATLAB; see [61].

In the last few years, open-source implementations have gained more and more interest, aiming to increase research transparency, enhance research exchange and benchmarking, and facilitate access to SDDP in industry and science [61]. The most prominent programming language in this regard is Julia [21], which provides its own algebraic modeling language JUMP [66] and is increasingly used in operations research and, in particular, stochastic programming. Currently, with `StochDynamicProgram.jl` [127], `StructDualDynProg.jl` [128], and `SDDP.jl` [61], there exist three SDDP implementations in Julia. Similarly, SDDP packages are available in MATLAB (`FAST` [36]), C++ (`StOpt` [84]), and Python (`mscopy` [54]).

Currently, `SDDP.jl`, which is based on the concept of policy graphs [60], can be considered the most comprehensive package. It provides many of the features described in this article, such as cut selection, parallelism, Markov chain SDDP, objective states, belief states, and SDDiP, as well as different stopping criteria and sampling approaches. Moreover, it includes some of the approaches discussed for distributionally robust and risk-averse SDDP. However, as for most other packages, it requires the underlying stochastic process to be finite. Thus, if Assumption 5 is not satisfied, some discretization has to be applied a priori and then the results obtained by SDDP are valid for the discretized problem, but are not put into perspective with respect to the true problem. `mscopy`, on the other hand, integrates both the discretization by sample average approximation (SAA) and the solution by SDDP into one package, and thus can naturally be applied to problems with continuous uncertainty [54].

A more detailed comparison of the libraries currently available is presented in [61].

**11. SDDP for Continuous Uncertainty [Relaxing Assumption 5].** Up to this point, we have assumed the uncertainty in (MSLP) to be modeled by some discrete

and finite random process (see Assumption 5) in order for SDDP to be applicable. Until the recent work by Forcier and Leclère [75], all convergence proofs for SDDP also leveraged Assumption 5. However, in many practical applications, this assumption is not justified. For example, if the stochastic process governing the uncertain data is modeled by a time-series model, the random error terms are usually assumed to follow a continuous distribution [205]; see section 9. In the remainder of this section, we denote a problem with such a continuous data process by  $(\tilde{P})$ .

As pointed out in section 2.3, for sizes of practical interest, problem  $(\tilde{P})$  is considered computationally intractable. Therefore, if the true distribution  $F_{\xi}$  of the stochastic process  $(\xi_t)_{t \in [T]}$  is continuous, an approximation with finitely many scenarios is generally used. In the literature on multistage stochastic programming, a variety of techniques is proposed to generate (and reduce) scenario tree approximations of continuous stochastic processes. For an overview, we refer to [135].

Before focusing on the most prominent technique in the next section, we should note that recently Forcier and Leclère [75] presented a (theoretical) extension of SDDP that works directly with continuous uncertainty. In their work, it is shown that under linearity (Assumption 6), given that  $W_t(\xi_t)$  is finitely supported and given that  $c_t(\xi_t)$  and  $(T_{t-1}(\xi_t), h_t(\xi_t))$  are independent, the ideas of SDDP may be extended to the case of continuous random variables  $\xi_t$ . Instead of generating cuts for  $\mathcal{Q}_t(\cdot)$  by deriving them for each realization  $j = 1, \dots, q_t$  separately and then aggregating, which is only possible for finite random variables, the main idea is that cuts are derived for  $\mathcal{Q}_t(\cdot)$  directly. To achieve this, the support  $\Xi_t$  of  $\xi_t$  is partitioned. For the partition  $\mathcal{P}$ , special *partitioned expected value functions*  $V_{\mathcal{P},t}(\cdot)$  are then defined in which the corresponding original value functions  $Q_t(\cdot)$  and the expectation operator  $\mathbb{E}[\cdot]$  are exchanged. This allows them to be evaluated and approximated even for continuous  $\xi_t$ . If  $V_{\mathcal{P},t}(\cdot)$  is guaranteed to underestimate  $\mathcal{Q}_t(\cdot)$  for all  $x_{t-1}$  (validity) and to coincide with  $\mathcal{Q}_t(\cdot)$  at some point  $\bar{x}_{t-1}$  (tightness), it is called adapted to  $\bar{x}_{t-1}$ . Given an adapted partition, by deriving cuts for  $V_{\mathcal{P},t}(\cdot)$  as in standard SDDP, valid and tight cuts for  $\mathcal{Q}_t(\cdot)$  can be obtained even for problems  $(\tilde{P})$ . Forcier and Leclère also show how adapted partitions can be constructed for (MSLP) explicitly under certain assumptions. However, we are not aware of any applications and computational tests of the resulting extension of SDDP.

**11.1. Sample Average Approximation (SAA).** The most common approximation approach for continuous uncertainty is to use random sampling. This means that the distribution  $F_{\xi}$  is approximated using an empirical distribution  $\tilde{F}_N$  with a finite number  $N$  of scenarios, which is obtained by sampling from  $F_{\xi}$  [205]. This yields an approximating problem  $(\tilde{P}_N)$ , which then can be handled by SDDP. Often, this technique is referred to as *sample average approximation* (SAA), especially if classical MC sampling is used. We discuss SAA and the application of SDDP to solve an SAA problem in more detail in what follows. For a general analysis of SAA, we refer the interested reader to [208].

**SAA and SDDP.** Under stagewise independence of  $(\xi_t)_{t \in [T]}$  (Assumption 2), it is desirable to preserve this property in the SAA problem, especially if the latter should be solved by SDDP. To achieve this, random sampling can be applied to each stage  $t = 2, \dots, T$  independently with sample size  $\tilde{q}_t$  [205]. The obtained SAA has a total number of  $N = \prod_{t=2}^T \tilde{q}_t$  scenarios, i.e., the number of scenarios grows exponentially in the number of stages [205].

For the SAA problem  $(\tilde{P}_N)$ , for each stage  $t = 2, \dots, T$  and each sample  $j = 1, \dots, \tilde{q}_t$ , the DPE can be written as



$$(11.1) \quad \tilde{Q}_t(x_{t-1}, \tilde{\xi}_{tj}) := \begin{cases} \min_{x_t} & (c_t(\tilde{\xi}_{tj}))^\top x_t + \tilde{\mathcal{Q}}_{t+1}(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \tilde{\xi}_{tj}), \end{cases}$$

where

$$(11.2) \quad \tilde{\mathcal{Q}}_{t+1}(x_t) := \frac{1}{N_{t+1}} \sum_{j=1}^{\tilde{q}_{t+1}} \tilde{Q}_{t+1}(x_t, \tilde{\xi}_{tj})$$

and  $\tilde{\mathcal{Q}}_{T+1}(\cdot) \equiv 0$ . For the first stage, we obtain

$$(11.3) \quad \tilde{v}_N := \begin{cases} \min_{x_1} & c_1^\top x_1 + \tilde{\mathcal{Q}}_2(x_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

The DPE (11.1)–(11.3) can be approached by SDDP, as described in section 3. However, in contrast to the problems considered there, the SAA problems are random as they depend on a sample from the true data process  $(\xi_t)_{t \in [T]}$ .

**SAA Properties.** Since the aim is to solve the original problem  $(\tilde{P})$ , the central question is how the solution and the bounds obtained by applying SDDP to the SAA problem  $(\tilde{P}_N)$  relate to the solution of  $(\tilde{P})$ . We denote the optimal value of  $(\tilde{P})$  by  $\tilde{v}^*$  and the bounds obtained by SDDP in iteration  $i$  by  $\underline{\tilde{v}}^i$  and  $\tilde{v}_{\mathcal{K}}^i$ . We next summarize important properties of SAA.

- (P.11.1) *Consistency.* It can be shown that the optimal value  $\tilde{v}_N$  provides a consistent estimator of the true optimal value  $\tilde{v}^*$ , i.e.,  $\lim_{\tilde{q}_2, \dots, \tilde{q}_T \rightarrow \infty} \tilde{v}_N = \tilde{v}^*$  with probability 1 [205, 208]. The intuition behind this is that asymptotically the structure of the true process  $(\xi_t)_{t \in [T]}$  is recovered. In practical applications, though, increasing  $\tilde{q}_t$  to infinity is computationally intractable.
- (P.11.2) *Bias.*  $\tilde{v}_N$  is a biased estimator of  $\tilde{v}^*$ ; more precisely,  $\mathbb{E}[\tilde{v}_N] \leq \tilde{v}^*$  for all  $N$  [208], since only a subset of all scenarios is considered and the decisions are optimized with respect to these scenarios [52]. This means that solving the SAA problem provides a (converging) estimator of a lower bound for  $\tilde{v}^*$  [202].
- (P.11.3) *Lower Bounds.* In each iteration  $i$  of SDDP we have  $\underline{\tilde{v}}^i \leq \tilde{v}_N$ . Therefore,  $\mathbb{E}[\underline{\tilde{v}}^i] \leq \tilde{v}^*$  [205] and the SDDP lower bound is a statistical lower bound for  $\tilde{v}^*$ . Note, however, that both,  $\tilde{v}_N$  and  $\underline{\tilde{v}}^i$ , are lower bounds in expectation only, whereas this is not clear for one specific SAA problem  $(\tilde{P}_N)$ .
- (P.11.4) *Upper Bounds.* Applying SDDP to the DPE (11.1)–(11.3) yields a policy. Under relatively complete recourse (see Assumption 9) with respect to the true data process  $(\xi_t)_{t \in [T]}$ , this policy also yields feasible decisions if applied to any realization  $(\xi_t)_{t \in [T]}$  of this true process. By computing

$$(11.4) \quad \mathbb{E} \left[ \sum_{t=1}^T (c_t(\xi_t))^\top x_t^i(\xi_{[t]}) \right]$$

with the expectation taken with respect to the true process, a valid upper bound for  $\tilde{v}^*$  can be obtained [205].

- (P.11.5) The sample mean  $\tilde{v}_{\mathcal{K}}^i$  determined in iteration  $i$  in SDDP is an unbiased and consistent estimator of (11.4). Hence,  $\mathbb{E}[\tilde{v}_{\mathcal{K}}^i] \geq \tilde{v}^*$ .

Even with these theoretical properties, solving  $(\tilde{P})$  using SAA may be computationally intractable. Shapiro shows that even under relatively complete recourse (see

Assumption 9) and stagewise independence (Assumption 2) of the true data process  $(\xi_t)_{t \in [T]}$ , the total number of scenarios required in SAA problem  $(\tilde{P}_N)$  to solve  $(\tilde{P})$  with a reasonable accuracy  $\varepsilon > 0$  grows exponentially in the number of stages [203]. Therefore, he proposes the use of smaller sample sizes  $\tilde{q}_t$  for later stages, although then the accuracy of the solution can no longer be guaranteed [204].

Clearly, there exists a trade-off between the quality of the obtained bounds for  $\tilde{v}^*$  and the computational tractability of the SAA problem. Approximating  $F_\xi$  with  $F_N$  using very large sample sizes  $\tilde{q}_t$  for all  $t = 2, \dots, T$ , a much better representation of the original process  $(\xi_t)_{t \in [T]}$  is obtained, which leads to a better approximation of  $\tilde{v}^*$ . However, in this case, it may even be impossible to solve the SAA problem to optimality in reasonable time, as it may take too long until all scenarios are eventually sampled [205]. On the other hand, a very rough approximation yields a problem  $(\tilde{P}_N)$  which can be solved efficiently by SDDP, but does not provide reasonable information about the solution to the true problem  $(\tilde{P})$  [121].

**11.2. Assessing Policy Quality.** As it is computationally intractable to solve an SAA problem of  $(\tilde{P})$  with a sample size that guarantees a predetermined accuracy, in practice, usually moderate sample sizes are used. For example, in [52], sample sizes with branching numbers  $\tilde{q}_t$  between 5 and 200 are tested.

The bounds  $\tilde{v}^i$  and  $\tilde{v}_K^i$  in SDDP are determined using one specific sample of  $(\xi_t)_{t \in [T]}$ . Therefore, they only measure the in-sample performance of the determined feasible policy  $(x_t(\xi_{[t]}))_{t \in [T]}$ . To assess its quality for the original problem  $(\tilde{P})$ , i.e., its out-of-sample performance, it is necessary to evaluate it with respect to the original process  $(\xi_t)_{t \in [T]}$ . Such an evaluation also allows for the comparison of policies obtained for different SAA problems, which can be helpful in designing appropriate sampling techniques and sample sizes [52].

Various techniques have been proposed in stochastic programming to measure the performance of feasible policies, such as analyzing optimality conditions, assessing solution stability, or estimating the optimality gap [52]. Specifically for SDDP, Morton et al. have made substantial contributions [42, 52, 121] based on estimating the optimality gap ([121] analyzes a risk-averse variant of SDDP; see section 12). We thoroughly discuss their ideas for the risk-neutral case in the remainder of this subsection. In accordance with [52], we only consider uncertainty in the RHS of  $(\tilde{P})$ .

**Estimating the Optimality Gap.** For some feasible policy  $(x_t(\xi_{[t]}))_{t \in [T]}$ , let  $\tilde{v}(\xi) = \sum_{t=1}^T c_t x_t(\xi_{[t]})$  denote the random cost for some arbitrary scenario path  $\xi = (\xi_1, \dots, \xi_T)$ . From (P.11.4) we have  $\mathbb{E}[\tilde{v}(\xi)] \geq \tilde{v}^*$ . Therefore, the optimality gap induced by policy  $(x_t(\xi_{[t]}))_{t \in [T]}$  can be expressed as

$$\Delta := \mathbb{E}[\tilde{v}(\xi)] - \tilde{v}^* \geq 0.$$

This gap cannot be directly evaluated because the optimal value  $\tilde{v}^*$  is not known. Using some lower bound for  $\tilde{v}^*$ ,  $\Delta$  can be overestimated though. Such a lower bound is given by  $\mathbb{E}[\tilde{v}]$ ; see (P.11.3). This yields

$$(11.5) \quad \mathbb{E}[\tilde{v}(\xi)] - \mathbb{E}[\tilde{v}] \geq \Delta \geq 0.$$

Still, the left-hand side of (11.5) is computationally infeasible to evaluate and requires excessive computational effort to evaluate policy  $(x_t(\xi_{[t]}))_{t \in [T]}$  for all possible scenarios to obtain  $\mathbb{E}[\tilde{v}(\xi)]$ . Furthermore, from SDDP only one specific realization of  $\tilde{v}$  is known. Therefore, [52] proposes using estimators for both terms to derive an approximate one-sided confidence interval bounding  $\Delta$  from above.

**Upper Bound Estimation.** The SDDP policy  $(x_t(\xi_{[t]}))_{t \in [T]}$  is feasible for the original problem  $(\tilde{P})$ ; see (P.11.4). Hence, it can be evaluated for any realization of  $(\xi_t)_{t \in [T]}$  to assess its out-of-sample performance. Let us sample  $M_u$  i.i.d. scenario paths from  $(\xi_t)_{t \in [T]}$ . For each of those sampled scenarios  $\xi^\ell, \ell = 1, \dots, M_u$ , the SDDP subproblems (2.10) are solved in the forward direction, yielding  $x_t(\xi_{[t]}^\ell)$  and  $\tilde{v}(\xi^\ell)$  [52]. An upper bound estimator is then defined by the sample mean

$$(11.6) \quad U_{M_u} := \frac{1}{M_u} \sum_{\ell=1}^{M_u} \tilde{v}(\xi^\ell).$$

Similarly to the in-sample estimator, this estimator is an unbiased and consistent estimator of  $\mathbb{E}[\tilde{v}(\xi)]$ . Its sample variance is given by [52]

$$(11.7) \quad \sigma_U^2 := \frac{1}{M_u - 1} \sum_{\ell=1}^{M_u} (\tilde{v}(\xi^\ell) - U_{M_u})^2.$$

Alternatively, an upper bound estimator can be obtained by sampling a finite number of different SAA problems and applying to each one the SDDP policy  $(x_t(\xi_{[t]}))_{t \in [T]}$  [42]. This comes at the cost of increased computational effort.

**Lower Bound Estimation with Several SAA Problems.** From SDDP, only one single realization of  $\tilde{v}$  is known. Hence, it is not possible to directly determine a sampling error for this point estimate and to derive a confidence interval for  $\mathbb{E}[\tilde{v}]$ .

One approach to deriving a lower bound estimator is to solve a finite number of different SAA problems with SDDP and to determine the mean of the lower bounds  $\tilde{v}$ . To be precise,  $M_l$  different SAA problems are constructed, each by sampling  $\hat{q}_t$  realizations per stage from  $(\xi_t)_{t \in [T]}$ . Then SDDP is run, yielding the lower bounds  $\tilde{v}^\ell, \ell = 1, \dots, M_l$  [52]. The sample mean

$$(11.8) \quad L_{M_l} := \frac{1}{M_l} \sum_{\ell=1}^{M_l} \tilde{v}^\ell$$

then defines an estimator for  $\mathbb{E}[\tilde{v}]$  with sample variance

$$\sigma_l^2 := \frac{1}{M_l - 1} \sum_{\ell=1}^{M_l} (\tilde{v}^\ell - L_{M_l})^2.$$

Note that instead of lower bounds  $\tilde{v}^\ell$ , the optimal values  $\tilde{v}_N^\ell$  could be also used in estimator (11.8) [52]. However, as discussed in section 11.1, it may be computationally intractable to solve one single SAA problem to optimality. Thus, using  $\tilde{v}^\ell$  may be computationally preferable.

In principle, applying SDDP to not only one, but several SAA problems and building the mean of the obtained bounds seems very reasonable from a statistical perspective, as the outcome of one SAA problem is random. This also has another possible benefit: If SDDP is run for  $M_l$  different SAA problems  $(\tilde{P}_N^l)$ , each of these problems yields a different feasible policy. By calculating the upper bound estimator  $U_{M_u}$  (11.6) for each one,  $M_l$  different policies could be compared directly.

However, for problems with multiple stages and for sufficiently high  $\hat{N}_t$ , this becomes computationally intractable, even without solving  $(\tilde{P}_N^l)$  exactly. Therefore, de Matos, Morton, and Finardi, [52] follow the strategy of running SDDP once for

some SAA problem with larger branch size  $\tilde{q}_t$  to determine a high quality policy, and then running SDDP for  $M_l$  SAA problems with smaller branch size  $\hat{q}_t$  only to produce the lower bound estimate  $L_{M_l}$  and assess the quality of that policy. In their numerical tests, they choose values between 5 and 200 for  $\tilde{q}_t$  and 5 for  $\hat{q}_t$ . In general though, it is not clear how to choose  $\hat{q}_t$  to reach a reasonable trade-off between computational tractability and an appropriate quality of the lower bound estimator.

**Lower Bound Estimation with One SAA Problem.** An alternative and less costly lower bound estimator is derived by using only the existing SAA problem, which has been applied to determine the policy that is to be assessed [52].

The idea is then to use the SDDP outcome  $\tilde{v}$  as the point estimate  $L_{M_l}$  for the lower bound. To estimate the unknown sampling error of  $\tilde{v}$ , the sampling error of the in-sample upper bound estimator is used. This means that  $M_l$  scenarios are sampled from  $F_N$  (the SAA problem distribution), and formulas (11.6) and (11.7) with  $M_l$  in the role of  $M_u$  are used to compute an upper bound estimate  $\tilde{v}_{M_l}$  and sample error  $\sigma_l^2$ . The idea behind applying this sampling error is that  $\tilde{v}$  and  $\mathbb{E}[\tilde{v}_{M_l}]$  are equal if SDDP has been run to optimality. However, this also implies that if SDDP has not converged (or if  $\tilde{q}_t$  is not sufficiently large), the sampling error may be underestimated and thus the confidence intervals drawn become overly optimistic [52].

**Confidence Intervals.** Using the bound estimators and their sample variances, asymptotically valid confidence intervals can be derived [52]:

$$\left( -\infty, U_{M_u} + t_{M_u-1, \alpha} \frac{\sigma_U}{\sqrt{M_u}} \right]$$

is an asymptotically valid and, for finite  $M_u$  approximate,  $(1 - \alpha)\%$  confidence interval for  $\mathbb{E}[\tilde{v}(\xi)]$ . Here,  $t_{M_u-1, \alpha}$  denotes the  $(1 - \alpha)$ -level quantile of a Student's  $t$  distribution with  $M_u - 1$  degrees of freedom. Similarly,

$$\left[ L_{M_l} - t_{M_l-1, \alpha} \frac{\sigma_l}{\sqrt{M_l}}, \infty \right)$$

is an asymptotically valid and, for finite  $M_l$  approximate,  $(1 - \alpha)\%$  confidence interval for  $\tilde{v}^*$ . Using only one SAA problem, this confidence interval is only valid if SDDP has converged and if  $\tilde{q}_t$  is sufficiently large. Combining both intervals yields

$$\left[ 0, [U_{M_u} - L_{M_l}]_+ + t_{M_l-1, \alpha} \frac{\sigma_l}{\sqrt{M_l}} + t_{M_u-1, \alpha} \frac{\sigma_U}{\sqrt{M_u}} \right]$$

as a one-sided approximate confidence interval for the optimality gap  $\Delta$  [52]. Here,  $[x]_+ := \max\{x, 0\}$ .

**11.3. Variance Reduction Techniques.** Instead of MC sampling, importance sampling [157] and variance reduction techniques (see section 6.2) can be also applied to obtain SAA estimators with reduced bias and variance.

In [112], numerical tests comparing MC, LHS, and RQMC indicate that RQMC yields the most promising results when it comes to determining representative SAA problems. In [52], MC, LHS, and RMC are also compared for different branch sizes and policy evaluation strategies. The results indicate that with both LHS and RQMC, a reduction of bias and sampling error, higher policy quality, and tighter confidence intervals can be achieved in comparison with MC sampling, especially for smaller branch sizes  $\tilde{q}_t$ . For smaller branch sizes, LHS appears to be superior, while RQMC

yields better results for larger branch sizes. While showing higher variability for MC sampling, if combined with RQMC and LHS sampling, the computationally preferable lower bound estimator using only in-sample scenarios from the existing SAA yields comparable results to the approach solving several SAA problems [52].

**12. Risk-Averse SDDP [Relaxing Assumption 8].** In SDDP, as described in section 3, a risk-neutral optimal policy is determined for (MSLP) (see Assumption 8). More precisely, (MSLP) minimizes the expectation of the total objective value over all stages  $t \in [T]$  over feasible policies  $(\mathbf{x}_t(\xi_t))_{t \in [T]}$ , which satisfy nonanticipativity and all constraints. Hence, it can be formulated as the single problem (2.3) with objective

$$(12.1) \quad \min_{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T} \mathbb{E} \left[ \sum_{t \in [T]} (\mathbf{c}_t(\xi_t))^\top \mathbf{x}_t(\xi_t) \right].$$

As discussed in section 2.4, this problem can be expressed equivalently using the DPE (2.4)–(2.6). This equivalence is based on two important properties of expected values: first the so-called tower property

$$(12.2) \quad \mathbb{E}_{\xi_t}[\mathbf{Z}_t(\xi_t)] = \mathbb{E}_{\xi_{[t-1]}}[\mathbb{E}_{\xi_t|\xi_{[t-1]}}[\mathbf{Z}_t(\xi_t)]]$$

for some random variable  $\mathbf{Z}_t$ , and second its strict monotonicity (see property (R2') below for a formal definition) [206].

Recall that the objective value  $\sum_{t \in [T]} (\mathbf{c}_t(\xi_t))^\top \mathbf{x}_t(\xi_t)$  is random and its realizations depend on realizations of  $(\xi_t)_{t \in [T]}$ . For some specific realization, the SDDP policy may produce an objective value that deviates widely from the expectation in (12.1). In practice, decision-makers are often anxious not only to find a policy causing low costs on average, but also to avoid the risk of extremely high cost situations. This motivates the consideration of risk-averse approaches in stochastic programming.

For multistage stochastic programming, incorporating risk-aversion has been a popular research topic over the last decade. This includes theoretical fundamentals on dynamic risk measures [199] as well as algorithmic developments, such as rolling horizon approaches with chance constraints or AVaR constraints, which take risk-aversion into account in the constraints of (MSLP) [102, 103]. For SDDP, most focus has been on replacing expectations in the objective (12.1) with some multiperiod risk measure  $\mathcal{R}[\cdot]$  (see below for a formal definition). This yields the multistage risk-averse problem  $(P_{\mathcal{R}})$ :

$$(12.3) \quad \begin{aligned} & \min_{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T} \mathcal{R} \left[ (\mathbf{c}_1(\xi_1))^\top \mathbf{x}_1(\xi_1), \dots, (\mathbf{c}_T(\xi_T))^\top \mathbf{x}_T(\xi_T) \right] \\ & \text{s.t.} \quad \mathbf{x}_1 \in \mathcal{X}_1, \\ & \quad \mathbf{x}_t \in \mathcal{X}_t(\mathbf{x}_{t-1}(\xi_{[t-1]}), \xi_t) \quad \forall \xi_t \in \Xi_t \quad \forall t = 2, \dots, T, \\ & \quad \mathbf{x}_t(\cdot) \text{ } \mathcal{F}_t\text{-measurable} \quad \forall t = 2, \dots, T. \end{aligned}$$

We cover risk-averse SDDP in detail in the remainder of this section, but start with some theoretical concepts.

**12.1. Risk Measures.** In this section, we introduce some required foundations of risk measures, in particular for multistage problems. As our focus is on algorithmic aspects of SDDP, we refer to the comprehensive coverage of this topic in [206, 208] for technical definitions and derivations.

**12.1.1. Static Risk Measures.** A *static* (or *one-period*) risk measure is a function  $\rho: \mathcal{Z} \rightarrow \bar{\mathbb{R}}$  from the space  $\mathcal{Z}$  of random variables  $\mathbf{Z}$  to  $\bar{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$ . Often,  $\mathcal{Z}$  is assumed to be  $\mathcal{L}_1(\Omega, \mathcal{F}, \mathbb{P})$ , i.e., the space of all  $\mathcal{F}$ -measurable functions with finite first moments, as this ensures the well-definedness and finiteness of many common risk measures. Importantly, since random variables are functions themselves, risk measures are actually *functionals*. This is sometimes emphasized by calling them *risk functionals* or *risk mappings*.

We now summarize some well-known static risk measures:

- The expected value  $\mathbb{E}[\cdot]$  is the most common risk measure. It is completely risk-neutral.
- The *value-at-risk*  $\text{VaR}_\alpha[\cdot]$  to level  $\alpha \in (0, 1)$  is defined as the left-hand side  $(1 - \alpha)$ -quantile of the cumulative distribution of some random variable  $\mathbf{Z}$ :

$$(12.4) \quad \text{VaR}_\alpha[\mathbf{Z}] := \inf \{u \in \mathbb{R} \mid \mathbb{P}(\mathbf{Z} \leq u) \geq 1 - \alpha\}.$$

Note that this definition is not used consistently in the literature, and that the RHS of (12.4) may also be defined as  $\text{VaR}_{1-\alpha}[\mathbf{Z}]$ .

- The *average value-at-risk*  $\text{AVaR}_\alpha[\cdot]$  to level  $\alpha \in (0, 1)$  for some random variable  $\mathbf{Z}$  is defined by [191]

$$(12.5) \quad \text{AVaR}_\alpha[\mathbf{Z}] := \inf \left\{ u \in \mathbb{R} \mid u + \frac{1}{\alpha} \mathbb{E}[(\mathbf{Z} - u)_+] \right\},$$

where  $[x]_+$  is defined as  $\max\{x, 0\}$ . Note that the infimum is always attained in our SDDP setting of finite randomness (Assumption 5) and finite value functions  $Q_t(\cdot)$  (see Lemma 2.5).

*Remark 12.1.*  $\text{AVaR}_\alpha[\cdot]$  is also called *conditional value-at-risk*, *expected shortfall*, *expected tail loss*, or *superquantile*. In the literature on risk-averse stochastic programming, the first alternative is most frequently used with notation  $\text{CVaR}_\alpha[\cdot]$ , but to avoid confusion when we introduce conditional risk measures later, we stick to average value-at-risk.

It can be shown that an equivalent formulation of  $\text{AVaR}_\alpha[\mathbf{Z}]$  is given by [205]

$$(12.6) \quad \text{AVaR}_\alpha[\mathbf{Z}] = \text{VaR}_\alpha[\mathbf{Z}] + \frac{1}{\alpha} \mathbb{E}[(\mathbf{Z} - \text{VaR}_\alpha[\mathbf{Z}])_+],$$

i.e.,  $u^* = \text{VaR}_\alpha[\mathbf{Z}]$  minimizes the RHS in (12.5).

$\text{AVaR}_\alpha[\cdot]$  has some beneficial properties compared to  $\text{VaR}_\alpha[\cdot]$ . It considers not only the probability mass beyond  $\text{VaR}_\alpha[\cdot]$ , but also its distribution, e.g., whether it has fat or long tails. Moreover, it allows us to retain convexity of optimization problems, as is discussed in what follows.  $\text{VaR}_\alpha[\cdot]$  and  $\text{AVaR}_\alpha[\cdot]$  are illustrated in Figure 10.


*Remark 12.2.* For continuous random variables  $\mathbf{Z}$ ,  $\text{AVaR}_\alpha[\cdot]$  may as well be defined as

$$\text{AVaR}_\alpha[\mathbf{Z}] = \mathbb{E}[\mathbf{Z} | \mathbf{Z} \geq \text{VaR}_\alpha[\mathbf{Z}]].$$

- In stochastic programming, a convex combination of  $\mathbb{E}[\cdot]$  and  $\text{AVaR}[\cdot]$  is often considered; that is,

$$(12.7) \quad \hat{\rho}_{\alpha, \lambda}[\mathbf{Z}] := (1 - \lambda)\mathbb{E}[\mathbf{Z}] + \lambda \text{AVaR}_\alpha[\mathbf{Z}]$$

for some  $\lambda \in [0, 1]$ . The parameters  $\lambda$  and  $\alpha$  control the risk-aversion. Choosing  $\lambda = 0$  yields the standard risk-neutral model.



**Fig. 10**  $\text{VaR}_\alpha[\mathbf{Z}]$  and  $\text{AVaR}_\alpha[\mathbf{Z}]$  for a gamma distributed random variable  $\mathbf{Z}$ .

- For some  $\gamma > 0$ , the *entropic risk measure* is defined by

$$(12.8) \quad \text{ENT}_\gamma[\mathbf{Z}] := \frac{1}{\gamma} \log(\mathbb{E}[e^{\gamma \mathbf{Z}}]).$$

It generalizes  $\mathbb{E}[\cdot]$  (for  $\gamma \rightarrow 0$ ) and  $\text{esssup}[\cdot]$  (for  $\gamma \rightarrow \infty$ ), where  $\text{esssup}[\mathbf{Z}]$  denotes the essential supremum of a random variable  $\mathbf{Z}$ .

It is often required that risk measures satisfy some special properties, especially in an optimization context. First, we assume that all considered risk measures are proper. Another desired property is coherence, a concept introduced by Artzner et al. [3]. We employ a slightly different definition from [208] and state it for the general case of continuous random variables.

**DEFINITION 12.3.** A risk measure  $\rho: \mathcal{Z} \rightarrow \bar{\mathbb{R}}$  is called *coherent* if it satisfies the following properties:

- (R1) *Convexity:* For any  $\mathbf{Z}_1, \mathbf{Z}_2 \in \mathcal{Z}$  and all  $\lambda \in [0, 1]$  it holds that

$$\rho(\lambda \mathbf{Z}_1 + (1 - \lambda) \mathbf{Z}_2) \leq \lambda \rho(\mathbf{Z}_1) + (1 - \lambda) \rho(\mathbf{Z}_2).$$

- (R2) *Monotonicity:* If  $\mathbf{Z}_1 \leq \mathbf{Z}_2$  almost surely, then  $\rho(\mathbf{Z}_1) \leq \rho(\mathbf{Z}_2)$ .

- (R3) *Translation Equivariance:* If  $a \in \mathbb{R}$  and  $\mathbf{Z} \in \mathcal{Z}$ , then  $\rho(\mathbf{Z} + a) = \rho(\mathbf{Z}) + a$ .

- (R4) *Positive Homogeneity:* If  $\lambda > 0$  and  $\mathbf{Z} \in \mathcal{Z}$ , then  $\rho(\lambda \mathbf{Z}) = \lambda \rho(\mathbf{Z})$ .

A risk measure satisfying only properties (R1), (R2), and (R3) is called *convex*. In fact, a key feature of coherent risk measures is that they are convex, and thus convex objective functions as they appear in  $(P_{\mathcal{R}})$  and its DPE remain convex if  $\rho[\cdot]$  is applied to them.  $\text{VaR}_\alpha[\cdot]$  is not a coherent risk measure, but  $\text{AVaR}_\alpha[\cdot]$  is [164]. Therefore, in optimization  $\text{AVaR}_\alpha[\cdot]$  is usually preferred over  $\text{VaR}_\alpha[\cdot]$ .

As we later exploit, for every coherent risk measure there exists a dual representation as the *worst-case expectation* over some class of probability distributions over  $(\Omega, \mathcal{F})$  [3]. More precisely, let  $\mathcal{P}$  be a convex set of probability measures; then a coherent risk measure  $\rho[\cdot]$  can be expressed as

$$(12.9) \quad \rho[\mathbf{Z}] = \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\mathbf{Z}].$$

Table 4 Properties of common risk measures.

	(R1)	(R2)	(R3)	(R4)	(R2')	(R5)
$\mathbb{E}[\cdot]$	✓	✓	✓	✓	✓	✓
$\text{VaR}_\alpha[\cdot]$	-	✓	✓	✓	-	✓
$\text{AVaR}_\alpha[\cdot]$	✓	✓	✓	✓	-	✓
$\hat{\rho}_{\alpha,\lambda}[\cdot]$	✓	✓	✓	✓	✓*	✓
$\text{ENT}_\gamma[\cdot]$	✓	✓	✓	-	✓	✓

\* only for  $\lambda \in [0, 1)$ .

We introduce some additional relevant properties.

DEFINITION 12.4. Let  $\rho : \mathcal{Z} \rightarrow \bar{\mathbb{R}}$  be some risk measure. Then, the following properties can be defined.

- (R2') If the inequalities in (R2) in Definition 12.3 are strict, we call this property strict monotonicity.
- (R5) Law Invariance:  $\rho[\cdot]$  is called law invariant with respect to  $\mathbb{P}$  if for all  $\mathbf{Z}, \mathbf{Z}' \in \mathcal{Z}$  with the same distribution,  $\rho(\mathbf{Z}) = \rho(\mathbf{Z}')$  also holds.

Property (R5) implies that the risk measure  $\rho$  depends only on the distribution of the considered random variable  $\mathbf{Z}$ .

We summarize properties of the previously introduced risk measures in Table 4.

Remark 12.5. A classical approach in economics is to take risk-aversion into account by means of nondecreasing and convex disutility (or concave utility) functions  $g : \mathbb{R} \rightarrow \mathbb{R}$  that are applied to some random variable  $\mathbf{Z}$  before taking expectations. However, the obtained risk measure  $\rho[\mathbf{Z}] = \mathbb{E}[g(\mathbf{Z})]$  does not satisfy property (R3), which is required to equivalently express  $(P_{\mathcal{R}})$  using DPE.

**12.1.2. Multiperiod Risk Measures.** In a multistage setting, static, i.e., one-period, risk measures have to be extended to several periods, more precisely, to a sequence of random variables  $\mathbf{Z} := \mathbf{Z}_1, \dots, \mathbf{Z}_T$ , which in our case model the stagewise objectives of (MSLP). We define such multiperiod risk measures as functionals  $\mathcal{R} : \mathcal{Z} \rightarrow \bar{\mathbb{R}}$  with  $\mathcal{Z} = \mathcal{Z}_1 \times \mathcal{Z}_2 \times \dots \times \mathcal{Z}_T$ , where  $\mathcal{Z}_t$  denotes the space of random variables  $\mathbf{Z}_t$  for each stage  $t$ .

Choosing multiperiod risk measures in a reasonable way is a challenging task. First, it is not clear how risk should be measured in a multistage setting [113]. Several different options exist [64, 113, 208], such as

(12.10)  $\mathcal{R}[\mathbf{Z}] = \rho[\mathbf{Z}_1 + \dots + \mathbf{Z}_T]$  (end-of-horizon risk),

(12.11)  $\mathcal{R}[\mathbf{Z}] = \rho_1 \left[ \mathbf{Z}_1 + \rho_{2|\mathcal{F}_1} [\mathbf{Z}_2 + \dots + \rho_{T|\mathcal{F}_{T-1}} [\mathbf{Z}_T] \dots] \right]$  (nested risk),

(12.12)  $\mathcal{R}[\mathbf{Z}] = \rho[\mathbf{Z}_1] + \dots + \rho[\mathbf{Z}_T]$  (stagewise risk).

Here,  $\rho[\cdot]$  is some static risk measure and  $\rho_{t|\mathcal{F}_{t-1}}[\cdot], t = 2, \dots, T$ , (later also denoted by  $\rho_{t|\xi_{[t-1]}}[\cdot]$ ) is a family of conditional risk measures, each mapping from  $\mathcal{Z}_t$  to  $\mathcal{Z}_{t-1}$  and defined as the static risk measure  $\rho_t[\cdot]$  conditioned on  $\mathcal{F}_{t-1}$  (or  $\xi_{[t-1]}$ , respectively). If  $\rho_t[\cdot]$  is law-invariant (property (R5) in Definition 12.4), then  $\rho_{t|\mathcal{F}_{t-1}}[\cdot]$  can be obtained by replacing the given distribution with the corresponding conditional distribution [208]. Generally, the same static risk measure  $\rho[\cdot]$  is chosen for all  $\rho_t[\cdot], t = 2, \dots, T$ . Note that coherence of conditional risk measures can be defined completely analogously to unconditional risk measures. The idea of nested conditional risk measures goes back to Ruszczyński and Shapiro [200].



*Remark 12.6.* Under stagewise independence (Assumption 2), as we assume it for SDDP, the conditional risk measures  $\rho_t|_{\mathcal{F}_{t-1}}[\cdot]$  in (12.11) no longer depend on  $\mathcal{F}_{t-1}$ , and thus they coincide with  $\rho_t[\cdot]$  [208].

Second, in an optimization context, multiperiod risk measures have to be carefully chosen in such a way that the resulting problem  $(P_{\mathcal{R}})$  possesses desirable properties. In addition to convexity, time consistency is an especially crucial property.

**12.1.3. Time Consistency.** In the literature, various different definitions of time consistency exist; see, among others, [37, 113, 50, 165, 206] and references within. The term is ambiguous in the sense that it is used for risk measures, policies, and optimization problems. We state here only some of these concepts that are relevant for SDDP; for technical definitions and detailed discussions we refer to [68, 113, 206, 208].

A common definition is that an optimal policy  $(\bar{x}_t(\xi_{[t]}))_{t \in [T]}$  for  $(P_{\mathcal{R}})$  (see (12.3)) is called *time consistent* if, for any  $\tau \in [T]$ , the policy  $(\bar{x}_t(\xi_{[t]}))_{t=\tau, \dots, T}$  is optimal for  $(P_{\mathcal{R}})$  restricted to horizon  $t = \tau, \dots, T$  conditional on  $\mathcal{F}_{\tau-1}$  and  $\bar{x}_{\tau-1}$  [208]. This means that the optimal policy remains optimal after some of the uncertain data has been revealed. The problem  $(P_{\mathcal{R}})$  is then called *weakly time consistent* if at least one of its optimal policies is time consistent, or *time consistent* if every optimal policy is time consistent [208] (note that there exist deviating definitions in the literature).


Policies obtained using DPE (such as (2.4)–(2.6)) naturally satisfy time consistency. Therefore, the concept of time consistency is closely related to equivalently reformulating  $(P_{\mathcal{R}})$  (see (12.3)) into DPE [208]. For nested risk measures  $\mathcal{R}[\cdot]$  (see (12.11)), this equivalence holds under strict monotonicity (property (R2') in Definition 12.4) of  $\rho_t$  (or  $\rho_t|_{\xi_{[t-1]}}$ , respectively) for all  $t = 2, \dots, T$ . More precisely, under (R2'), by interchanging risk measures and minimization operators,  $(P_{\mathcal{R}})$  with nested risk can be expressed in the nested fashion [208]

$$(12.13) \quad \min_{x_1 \in \mathcal{X}_1} c_1^\top x_1 + \rho_2 \left[ \min_{x_2 \in \mathcal{X}_2(x_1)} (c_2(\xi_2))^\top x_2 + \rho_{3|\xi_{[2]}} \left[ \dots \right. \right. \\ \left. \left. \dots + \rho_{T|\xi_{[T-1]}} \left[ \min_{x_T \in \mathcal{X}_T(x_{T-1})} (c_T(\xi_T))^\top x_T \right] \dots \right] \right],$$

which naturally allows for a reformulation to DPE. Note that for stage 2, no conditional expectation is used as the first-stage data is deterministic. If  $\rho_t$  (or  $\rho_t|_{\xi_{[t]}}$ ) only satisfy (R2) instead of (R2'), then only weak consistency of  $(P_{\mathcal{R}})$  is guaranteed, as any optimal policy for the DPE is also optimal for problem  $(P_{\mathcal{R}})$  with nested risk, but not necessarily vice versa.

As indicated by Table 4,  $\text{AVaR}_\alpha[\cdot]$  is not strictly monotone. Therefore, even if applied in a nested conditional way, time consistency is not assured. In contrast, it can be ensured using risk measure  $\hat{\rho}_{\alpha, \lambda}[\cdot]$  defined in (12.7), given that  $\lambda \in [0, 1]$ . A drawback of nested risk is that it is less amenable to suitable interpretation, although some economic interpretations are possible [197].

For one-period risk measures  $\rho[\cdot]$  that are applied as an end-of-horizon risk measure (12.10), it is well known that time consistency is often not satisfied. For instance, some simple examples in [68, 113] show that using a one-period risk measure  $\rho[\cdot]$ , such as  $\text{VaR}_\alpha[\cdot]$  or  $\text{AVaR}_\alpha[\cdot]$ , in this setting leads to time-inconsistent decisions. Moreover, in [197], an illustrative example is presented in which even under stagewise independence (Assumption 2), the risk measure  $\hat{\rho}_{\alpha, \lambda}[\cdot]$  does not yield time-consistent policies from an end-of-horizon perspective. To achieve time consistency, it is required that



**Fig. 11** Different forms of  $(P_{\mathcal{R}})$  and conditions for their equivalence.

problem  $(P_{\mathcal{R}})$  (see (12.3)) with end-of-horizon risk measure  $\rho[\cdot]$  can be converted into an equivalent problem with nested risk using the corresponding conditional risk measures  $\rho_{|\xi[t]}$ . For this reason, Dowson, Morton, and Pagnoncelli [64] define time consistency (in their case referred to as *conditional consistency*) of a one-period risk measure  $\rho[\cdot]$  as an equivalence between the associated end-of-horizon risk and nested risk.

In fact, the only law-invariant coherent one-period risk measures  $\rho[\cdot]$  allowing for such an equivalent reformulation between an end-of-horizon risk and a nested risk perspective are  $\mathbb{E}[\cdot]$  and  $\text{esssup}[\cdot]$  [208]. Therefore, the coherent and law-invariant risk measure  $\text{AVaR}_{\alpha}[\cdot]$  does not even guarantee weak time consistency for  $(P_{\mathcal{R}})$  if it is applied as an end-of-horizon risk measure. It can be shown, though, that the noncoherent but convex risk measure  $\text{ENT}_{\gamma}[\cdot]$  from (12.8) is conditionally consistent, and thus it is sufficient to ensure time consistency of  $(P_{\mathcal{R}})$ . The equivalence of different formulations for problem  $(P_{\mathcal{R}})$  is illustrated in Figure 11.

*Remark 12.7.* In view of conditional consistency, note that nested risk measures  $\mathcal{R}[\cdot]$  from (12.11) can always be expressed equivalently using an associated end-of-horizon risk measure (12.10), called *composite risk measure*. However, as the previous discussion shows, this composite risk measure only equals  $\rho[\cdot]$  if the latter allows for a decomposition using its conditional analogues, similar to (12.2) [206, 208].

In addition, some notion of time consistency can be satisfied using *expected conditional risk measures*  $\mathcal{R}[\cdot]$ , which measure the risk stage by stage (see (12.12)), as long as the included (conditional) risk measures are coherent [113]. Applying such a risk measure in  $(P_{\mathcal{R}})$  (problem (12.3)), we obtain the problem

$$\begin{aligned}
 (12.14) \quad & \min_{x_1, x_2, \dots, x_T} \quad c_1^\top x_1 + \rho_2 \left[ (c_2(\xi_2))^\top x_2(\xi_{[2]}) \right] + \mathbb{E}_{\xi_{[2]}} \left[ \rho_3(\xi_{[2]}) \left[ (c_3(\xi_3))^\top x_3(\xi_{[3]}) \right] \right] \\
 & \quad + \dots + \mathbb{E}_{\xi_{[T-1]}} \left[ \rho_T(\xi_{[T-1]}) \left[ (c_T(\xi_T))^\top x_T(\xi_{[T]}) \right] \right] \\
 & \text{s.t.} \quad x_1 \in \mathcal{X}_1, \\
 & \quad x_t \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t) \quad \forall \xi_t \in \Xi_t \quad \forall t = 2, \dots, T, \\
 & \quad x_t(\cdot) \text{ } \mathcal{F}_t\text{-measurable} \quad \forall t = 2, \dots, T.
 \end{aligned}$$

**12.1.4. Polyhedral Risk Measures.** Multiperiod polyhedral risk measures  $\mathcal{R}[\cdot]$  are a special type of risk measure that for a time horizon of  $T \in \mathbb{N}$  can be formulated as the optimal value of certain  $T$ -stage stochastic LPs [71]. The arguments of the risk measure, e.g., in our case the objective function of (MSLP), enter these LPs on the RHS.

In [100], multiperiod extended polyhedral risk measures are introduced for which the corresponding LP has a slightly more general form. This class comprises polyhedral risk measures, spectral risk measures, and also  $\text{AVaR}_{\alpha}[\cdot]$ . These risk measures can be shown to be convex and coherent under certain assumptions [100].

The main strength of (extended) polyhedral risk measures is that they can be used naturally in a multistage stochastic programming setting. The linear programming

representation of  $\mathcal{R}[\cdot]$  and the original linear programming formulation of (MSLP) can be conflated as a single large-scale risk-neutral linear programming problem  $(P_{\mathcal{R}})$ , which allows for a reformulation by means of DPE [100].

**12.2. Toward Considering Risk in SDDP.** In the remainder of this section, we discuss the incorporation of risk-aversion into SDDP from an algorithmic perspective. As a first step, this requires the derivation of tractable DPE for  $(P_{\mathcal{R}})$ . Then, SDDP has to be adapted to those DPE, which potentially affects the cut generation, the upper bound computation, and the sampling.

The first two methodological studies of risk-averse SDDP were [100] for problems with end-of-horizon risk (12.10), in particular, using polyhedral risk measures, and [205] for problems with nested conditional risk mappings (12.11). Since then several extensions of SDDP have been proposed based on various risk measures. While some articles on this topic also cover SAA [121, 205, 211] (see section 11), we restrict our attention to finite random variables in what follows.

*Remark 12.8* (SDDP with polyhedral risk measures). As stated in section 12.1.4, polyhedral risk measures have the advantage that DPE can be derived in a straightforward way. These DPE can then be approached by standard risk-neutral SDDP. Guigues and Römisch derive the associated cut formulas and give a convergence proof for some special cases of extended polyhedral risk measures [100] and the special case of spectral risk measures [101]. This approach to SDDP is successfully applied for AVaR $_{\alpha}[\cdot]$  in [91].

Despite this straightforward approach, polyhedral risk measures also pose a significant challenge to SDDP. The stage- $t$  subproblems have to be enhanced with additional state variables  $z_{t-1}$  and  $y_1, \dots, y_{t-1}$ , which are required to store the history of previous decisions. In general, such a state space expansion is unfavorable, as it may lead to prohibitive computational cost [168]; see the complexity results in section 4.2. The specific computational cost depends on the extended polyhedral risk measure that is chosen.

**12.3. SDDP with Nested Risk Measures.** As mentioned in section 12.1.3, to obtain a risk-averse problem  $(P_{\mathcal{R}})$  with time-consistent solutions, it is often proposed to use (conditional) coherent one-period risk measures  $\rho[\cdot]$  (or  $\rho_{t|\xi_{[t]}}[\cdot]$ ) for all  $t \in [T]$  in a nested fashion. This yields the nested problem (12.13). We denote its optimal value by  $v_{\mathcal{R}}^*$ . As indicated above, we can derive an equivalent formulation using DPE [208]. Using Remark 12.6 the DPE become

$$(12.15) \quad Q_{\mathcal{R},t}(x_{t-1}, \xi_t) := \begin{cases} \min_{x_t} & (c_t(\xi_t))^{\top} x_t + \mathcal{Q}_{\mathcal{R},t+1}(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t) \end{cases}$$

with some *risk-adjusted value function*

$$(12.16) \quad \mathcal{Q}_{\mathcal{R},t+1}(x_t) := \rho_{t+1} [Q_{\mathcal{R},t+1}(x_t, \xi_{t+1})]$$

and  $\mathcal{Q}_{\mathcal{R},T+1}(\cdot) \equiv 0$ . The corresponding first-stage problem is

$$(12.17) \quad v_{\mathcal{R}}^* = \begin{cases} \min_{x_1} & c_1^{\top} x_1 + \mathcal{Q}_{\mathcal{R},2}(x_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

Fortunately, for coherent risk measures  $\rho_t[\cdot], t \in [T]$ , the nested risk measure  $\mathcal{R}[\cdot]$  also preserves convexity of  $\mathcal{Q}_{\mathcal{R},t+1}(\cdot)$ . Therefore, a cutting-plane approximation as in SDDP can be applied.

Nested conditional risk measures are by far the most frequently chosen approach for risk-averse extensions of SDDP [68, 105, 113, 121, 167, 168, 205, 211]. Most typically, the risk measure  $\hat{\rho}_{\alpha,\lambda}[\cdot]$  (see (12.7)) is used, which is coherent according to Table 4. For the remainder of section 12.3, we therefore set  $p_t[\cdot] = \hat{\rho}_{\alpha_t,\lambda_t}[\cdot]$  for all  $t \in [T]$ , if not specified otherwise.

**12.3.1. Reformulating the DPE.** The general DPE for  $(P_{\mathcal{R}})$  with nested risk measures are formulated in (12.15)–(12.17). To determine  $\mathcal{Q}_t(\cdot, \cdot), t \in [T]$ , for  $\hat{\rho}_{\alpha,\lambda}[\cdot]$  specifically, the AVaR of  $Q_t(\cdot, \cdot)$  has to be evaluated. Using its definition as the optimal value of an optimization problem with decision variable  $u \in \mathbb{R}$  [191] (see (12.5)), we are able to further reformulate the DPE such that only expectation operators occur.

**Additional State Variable Approach.** Using (12.5), the risk-adjusted value function (12.16) can be expressed as

$$(12.18) \quad \begin{aligned} \mathcal{Q}_{\mathcal{R},t+1}(x_t) = \min_{u_t \in \mathbb{R}} \mathbb{E}_{\xi_{t+1}} \Big[ & (1 - \lambda_{t+1}) Q_{\mathcal{R},t+1}(x_t, \xi_{t+1}) \\ & + \lambda_{t+1} \left( u_t + \frac{1}{\alpha_{t+1}} [Q_{\mathcal{R},t+1}(x_t, \xi_{t+1}) - u_t]_+ \right) \Big]. \end{aligned}$$

Recall that  $\lambda_t$  and  $\alpha_t, t = 2, \dots, T$ , are user-controlled parameters.

The minimization over  $u_t$  can be incorporated into the stage- $t$  subproblems [205], which yields

$$(12.19) \quad \tilde{Q}_{\mathcal{R},t}(x_{t-1}, \xi_t) = \begin{cases} \min_{x_t, u_t} & (c_t(\xi_t))^\top x_t + \lambda_{t+1} u_t + \tilde{\mathcal{Q}}_{\mathcal{R},t+1}(x_t, u_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t) \end{cases}$$

with some modified risk-adjusted value function

$$(12.20) \quad \begin{aligned} \tilde{\mathcal{Q}}_{\mathcal{R},t+1}(x_t, u_t) = & \mathbb{E}_{\xi_{t+1}} \Big[ (1 - \lambda_{t+1}) \tilde{Q}_{\mathcal{R},t+1}(\xi_{t+1}, x_t) \\ & + \frac{\lambda_{t+1}}{\alpha_{t+1}} [\tilde{Q}_{\mathcal{R},t+1}(\xi_{t+1}, x_t) - u_t]_+ \Big], \end{aligned}$$

$\tilde{\mathcal{Q}}_{\mathcal{R},T+1}(\cdot, \cdot) \equiv 0$ , and  $\lambda_{T+1} \equiv 0$  [205]. The first stage changes to

$$(12.21) \quad v_{\mathcal{R}}^* = \begin{cases} \min_{x_1, u_1} & c_1^\top x_1 + \lambda_2 u_1 + \tilde{\mathcal{Q}}_{\mathcal{R},2}(x_1, u_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

The risk-adjusted value functions  $\tilde{\mathcal{Q}}_{\mathcal{R},t+1}(\cdot, \cdot)$  differ from those defined in (12.18), but can be proven to be convex as well.

With (12.19)–(12.21), the risk measures  $\rho_{\alpha_t,\lambda_t}[\cdot]$  are incorporated into the subproblems such that only expectations have to be evaluated in the DPE. However, as pointed out in [121], in comparison with the DPE (2.4)–(2.6) of the risk-neutral case, we still observe some fundamental differences: First, an additional, albeit one-dimensional, state variable  $u_t \in \mathbb{R}$  is introduced at each stage to estimate the VaR-level, augmenting the state space by one. Second, the risk-adjusted value functions  $\mathcal{Q}_{\mathcal{R},t+1}(\cdot)$  not only depend on  $x_t$ , but also on  $u_t$  and parameters  $\lambda_t, \alpha_t$ . Third, they contain the nonlinear, i.e., piecewise linear, function  $[\cdot]_+$ .

Philpott and de Matos provide an alternative reformulation of the DPE, eliminating the nonlinear expression via an epigraph reformulation [167]. To this end, the

random term in the brackets in (12.20) is fully incorporated into the value functions. For  $t = 2, \dots, T - 1$ , this yields

$$(12.22) \quad \begin{aligned} & \widehat{Q}_{\mathcal{R},t}(x_{t-1}, u_{t-1}, \xi_t) \\ &= \begin{cases} \min_{x_t, u_t, w_t} & (1 - \lambda_t) \left( (c_t(\xi_t))^\top x_t + \lambda_{t+1} u_t + \widehat{\mathcal{Q}}_{\mathcal{R},t+1}(x_t, u_t) \right) + \frac{\lambda_t}{\alpha_t} w_t \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t), \\ & w_t - (c_t(\xi_t))^\top x_t - \lambda_{t+1} u_t - \widehat{\mathcal{Q}}_{\mathcal{R},t+1}(x_t, u_t) \geq -u_{t-1}. \end{cases} \end{aligned}$$

Using this formulation, the risk value function is defined more naturally as

$$(12.23) \quad \widehat{\mathcal{Q}}_{\mathcal{R},t+1}(x_t, u_t) = \mathbb{E}_{\xi_{t+1}} \left[ \widehat{Q}_{\mathcal{R},t+1}(x_t, u_t, \xi_{t+1}) \right].$$

Again,  $\widehat{\mathcal{Q}}_{\mathcal{R},T+1}(\cdot, \cdot) \equiv 0$  and  $\lambda_{T+1} \equiv 0$ .

The first-stage problem then reads

$$(12.24) \quad v_{\mathcal{R}}^* = \begin{cases} \min_{x_1, u_1, w_1} & c_1^\top x_1 + \lambda_2 u_1 + \widehat{\mathcal{Q}}_{\mathcal{R},2}(x_1, u_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

In comparison to the formulation (12.19)–(12.21) by Shapiro [205], additional variables and constraints have to be introduced. Both formulations allow the application of SDDP but share the drawback of augmenting the state space. Since the computational effort of SDDP grows exponentially in the state space dimension (see Theorem 4.2), such an increase should be avoided.

**Modifying the Probability Measure.** An alternative idea is to exploit the fact that  $u^* = \text{VaR}_\alpha[\mathbf{Z}]$  in the definition of  $\text{AVaR}_\alpha[\mathbf{Z}]$  (see (12.5)) and that  $\text{VaR}_\alpha[\mathbf{Z}]$  is the  $(1-\alpha)$ -quantile of a random variable  $\mathbf{Z}$ . As we assume finite randomness (Assumption 5) and solve the subproblems for all realizations  $\xi_{tj}, j = 1, \dots, q_t$ , in the backward pass of SDDP, this quantile can be manually determined for the value functions [211].

Without loss of generality, assume that for all  $t = 2, \dots, T$  and any fixed trial solution  $\bar{x}_{t-1}$  the values of  $Q_{\mathcal{R},t}(\bar{x}_{t-1}, \xi_{tj})$  are ordered for all  $j = 1, \dots, q_t$ . That means we have  $Q_{\mathcal{R},t}(\bar{x}_{t-1}, \xi_{t1}) \leq \dots \leq Q_{\mathcal{R},t}(\bar{x}_{t-1}, \xi_{t,q_t})$ . With this argument, in (12.18) the variable  $u_t$  can be replaced by the  $(1-\alpha)$ -quantile  $Q_{\mathcal{R},t+1}(\bar{x}_t, \xi_{t+1,j^*})$  with  $j^*$  chosen such that  $\sum_{j=1}^{j^*} p_{t+1,j} \geq 1 - \alpha_{t+1}$ :

$$(12.25) \quad \begin{aligned} \mathcal{Q}_{\mathcal{R},t+1}(x_t) = & \mathbb{E}_{\xi_{t+1}} \left[ (1 - \lambda_{t+1}) Q_{\mathcal{R},t+1}(x_t, \xi_{t+1}) + \lambda_{t+1} \left( Q_{\mathcal{R},t+1}(\bar{x}_t, \xi_{t+1,j^*}) \right. \right. \\ & \left. \left. + \frac{1}{\alpha_t} [Q_{\mathcal{R},t+1}(x_t, \xi_{t+1}) - Q_{\mathcal{R},t+1}(\bar{x}_t, \xi_{t+1,j^*})]_+ \right) \right]. \end{aligned}$$

In SDDP, relation (12.25) cannot be directly applied, since  $Q_{\mathcal{R},t+1}(\cdot, \xi_{t+1,j})$  is not known and is also not evaluated for all  $j = 1, \dots, q_{t+1}$ . However, the same principle can also be applied to the approximate value functions  $\underline{Q}_{\mathcal{R},t+1}(\cdot, \xi_{t+1,j})$ . Due to the monotonicity (R2) of  $\rho_{t+1}$ , we have  $\rho_{t+1} [Q_{\mathcal{R},t+1}(x_t, \xi_{t+1})] \geq \rho_{t+1} [\underline{Q}_{\mathcal{R},t+1}(x_t, \xi_{t+1})]$ , so this yields valid lower approximations.

In [168], this idea is considered from a dual perspective and used to reformulate the risk measure (12.7) even before formulating the DPE. The key idea is to use the dual representation of  $\text{AVaR}_\alpha[\cdot]$  (see (12.9)), which is given by

$$(12.26) \quad \text{AVaR}_\alpha[\mathbf{Z}] = \begin{cases} \sup_{\zeta} & \sum_{j=1}^q p_j \zeta_j Z(\xi_j) \\ \text{s.t.} & \sum_{j=1}^q p_j \zeta_j = 1, \\ & \zeta_j \geq 0, \quad j = 1, \dots, q, \\ & \zeta_j \leq \frac{1}{\alpha}, \quad j = 1, \dots, q. \end{cases}$$

It shows that  $\text{AVaR}_\alpha[\cdot]$  can be interpreted as some worst-case probability measure  $\tilde{\mathbb{P}}$  with  $\tilde{p}_j := p_j \zeta_j$  for all  $j = 1, \dots, q$ .

As shown in [168], using this definition and explicitly computing the supremum, risk measure (12.7) can be written as

$$(12.27) \quad \hat{\rho}_{\alpha_t, \lambda_t}[\mathbf{Z}_t] = \sum_{j=1}^{q_t} p_{tj} \zeta_{tj} Z_t(\xi_{tj})$$

with

$$(12.28) \quad \zeta_{tj} = \begin{cases} (1 - \lambda_t), & j < j^*, \\ (1 - \lambda_t) + \frac{1}{p_{tj^*}} \left( \lambda_t - \frac{\lambda_t}{\alpha_t} \sum_{n=j^*+1}^{q_t} p_{tn} \right), & j = j^*, \\ (1 - \lambda_t) + \frac{\lambda}{\alpha_t}, & j > j^*. \end{cases}$$

Again, note that the true value functions  $Q_t(\cdot)$  are not known explicitly in advance, and therefore the worst-case probability measure  $\tilde{\mathbb{P}}$  stemming from (12.26) is not known either. However, it can be approximated in SDDP. In particular, the DPE (12.15)–(12.17) and their approximations can be used with expectations as in standard SDDP, but with a modified probability measure that is iteratively updated. More precisely,  $\zeta_{tj}$  changes with  $\bar{x}_{t-1}$ , so the modified probabilities have to be recomputed for each stage  $t$ , iteration  $i$ , and sample  $k$  in SDDP. This principle is also extended to general coherent risk measures in [168].

This kind of change of the probability measure is also discussed in [133]. Instead of adapting the ordering and  $j^*$  based on  $Q_{t+1}^{i+1}$  and  $x_t^i$  in each iteration  $i$ , it is assumed that  $j^*$  is stable with respect to different values of  $x_{t-1}$ . Its stable value is then approximated by counting the number of iterations in which an index  $j$  exceeds  $\text{VaR}_\alpha[Q_{\mathcal{R}, t+1}(x_t, \xi_t)]$ . This is considered to be a good proxy for the ordering of the actual value functions. A different approximation strategy is proposed in [80]. As a key ingredient, multicut SDDP is used; see section 21.2.1. An ordering of the most current values  $\theta_{t+1,j}^i$  obtained for each cut approximation  $\underline{v}_{tj}^i(\cdot)$ ,  $j = 1, \dots, q_t$ , at stage  $t$  is then used as a proxy for the risk-adapted probability measure  $\tilde{\mathbb{P}}$ .

The ordering, and thus the probability measure  $\tilde{\mathbb{P}}$ , can be either updated dynamically within SDDP or determined by running risk-averse SDDP once in advance to identify the outcomes contributing to  $\text{AVaR}_\alpha[\cdot]$ . The latter approach has the advantage that the changed probability measure  $\tilde{\mathbb{P}}$  can be fixed for the following run, which yields a risk-neutral problem and allows for the application of standard SDDP.

For the third stage of Example 3.4, the expected risk value function  $\mathcal{Q}_{\mathcal{R},3}(\cdot)$  obtained by applying (12.27) and (12.28) to (12.16) is illustrated in Figure 12 for  $\alpha = 0.05$  and different values of  $\lambda$ . It can be seen that when choosing larger values




Fig. 12  $\mathcal{Q}_{\mathcal{R},3}(\cdot)$  from Example 3.4 for  $\alpha = 0.05$  and different values of  $\lambda$ .

Table 5 DPE formulations for  $(P_{\mathcal{R}})$  using a nested (conditional) risk measure based on  $\widehat{\rho}_{\alpha,\lambda}[\cdot]$ .

Description	Source	DPE
- general		(12.15)–(12.17)
- augmented state, sophisticated formula for $\mathcal{Q}_{\mathcal{R},t}(\cdot)$	[205]	(12.19)–(12.21)
- augmented state, additional constraints and variables	[167]	(12.22)–(12.24)
- $\text{VaR}_{\alpha_t}[Q_t(\cdot)]$ explicitly determined, sophisticated formula for $\mathcal{Q}_{\mathcal{R},t}(\cdot)$	[211]	(12.19), (12.21), (12.25)
- modified probability measure	[168]	(12.15)–(12.17), (12.27)–(12.28)
- modified probability measure	[133]	(12.19), (12.21), (12.25)

for  $\lambda$ , representing a higher risk-aversion, the stage-3 cost increases compared to the risk-neutral case ( $\lambda = 0$ ).

As an overview, the different forms of DPE for  $(P_{\mathcal{R}})$  using a nested (conditional) risk measure based on  $\widehat{\rho}_{\alpha,\lambda}[\cdot]$  are summarized in Table 5. All approaches in Table 5 to formulate the DPE allow for a solution of a risk-averse problem  $(P_{\mathcal{R}})$  using SDDP. Some approaches are more efficient, since the state space, the decision space, and the number of constraints are not augmented. Others are advantageous in the sense that  $\mathcal{Q}_{\mathcal{R},t}(\cdot)$  is expressed by a neat formula, and thus cut formulas can be derived more easily. With some epigraph reformulation, for all the approaches all subproblems can be formulated as LPs.

**12.3.2. Forward and Backward Passes.** The forward pass of SDDP basically remains the same as for risk-neutral SDDP from section 3; see Algorithm 3.1. That is,  $k \in \mathcal{K}$  scenarios are sampled and considered, with  $\mathcal{K} \subset \mathcal{S}$  and  $|\mathcal{K}| \ll |\mathcal{S}|$ . However, the subproblems and the associated approximate value functions  $\underline{Q}_{\mathcal{R},t}^i(\xi_t^k, \xi_{tj})$  differ from the risk-neutral case. Instead of subproblems (2.10), one of the DPE from Table 5 is chosen and the risk-adjusted value functions  $\mathcal{Q}_{\mathcal{R},t+1}(\cdot)$  that occur are replaced by cut approximations  $\underline{\mathcal{V}}_{\mathcal{R},t+1}^i(\cdot)$ . The sampling technique in line 6 of Algorithm 3.1 usually remains the same as in risk-neutral SDDP, meaning that random sampling with respect to the distribution of  $\boldsymbol{\xi}$  is used. However, it is also possible to sample scenarios with “bad” outcomes with higher probability based on proxies of the probability measure  $\mathbb{P}$  [80, 133]. This *biased sampling* can be considered to be similar to the importance sampling techniques presented in section 6.

In the backward pass, as in risk-neutral SDDP, at each stage  $t = T, \dots, 2$  the subproblems are solved for each trial solution  $x_{t-1}^{ik}, k \in \mathcal{K}$ , and possible stage- $t$  realization  $\xi_{tj}^k \equiv \xi_{tj}, j = 1, \dots, q_t$ , using an updated cut approximation  $\underline{v}_{\mathcal{R},t+1}^{i+1}(\cdot)$ . At stage  $t$ , a new cut for  $\mathcal{Q}_{\mathcal{R},t}(\cdot)$  is derived and handed back to stage  $t-1$ . The main difference to risk-neutral SDDP is again the definition of  $\mathcal{Q}_{\mathcal{R},t}(\cdot)$ . Therefore, the cut formulas used in line 18 of Algorithm 3.1 have to be adapted to the individual approach chosen. For the technical derivation of subgradients in such cases, we refer to the references in Table 5.

**12.3.3. Upper Bound Determination and Stopping.** A challenge in applying SDDP to risk-averse problems is to determine upper bounds for  $v_{\mathcal{R}}^*$  and allow for a reasonable stopping criterion, because most upper bound construction methods from the risk-neutral case (see sections 7 and 8) cannot be efficiently extended to the risk-averse case.

Recall that in the risk-neutral case, a feasible policy  $(x_t(\xi_{[t]}))_{t \in [T]}$  is determined in the backward pass and evaluated in the forward pass for different scenarios  $k \in \mathcal{K}$ , yielding a sequence of trial points  $(x_t^{ik})_{t \in [T]}$ . Then, a statistical upper bound  $\bar{v}_{\mathcal{K}}$  for  $v^*$  is determined as the sample average of the objective values of all these sample paths  $\xi^k$ ; see (3.9). Analogously, a true upper bound  $\bar{v}$  can be obtained by taking the expectation of such objective values for all scenarios  $\xi^s, s \in \mathcal{S}$ .

However, this is possible only due to the tower property (12.2) of expected values, which is required for the equivalence of the end-of-horizon formulation (12.1) and the nested formulation (12.13); see the discussion in section 12.1.3. For most coherent risk measures this property does not hold, and thus a direct analogue to the (statistical) upper bound (3.9) from risk-neutral SDDP cannot be constructed.

Since determining reasonable upper bounds is an important ingredient of SDDP, developing appropriate upper bound estimators has been an active research field in the last decade. In what follows, we discuss different approaches that have been proposed. In reviewing them, we mostly follow the presentation of Kozmík and Morton [121], who provide a comprehensive study within their own work on upper bound estimators.

**A Sample Average Estimator.** In section 12.3.1, we managed to formulate each  $\rho_t[\cdot]$  by means of only expectations in (12.20). Yet this does not assure the tower property, since the risk-adjusted value functions  $\mathcal{Q}_{\mathcal{R},t}(\cdot)$  contain a nested nonlinearity due to the  $[\cdot]_+$ -function. However, we can derive an estimator similar to (3.9) [121]. To this end, we remove the expectation in (12.20) to obtain

$$(12.29) \quad \begin{aligned} \hat{v}_t(\xi_t^k) &:= (1 - \lambda_t) \left( (c_t(\xi_t^k))^{\top} x_t^k + \hat{v}_{t+1}(\xi_t^k) \right) \\ &\quad + \lambda_t u_{t-1}^k + \frac{\lambda_t}{\alpha_t} \left[ (c_t(\xi_t^k))^{\top} x_t^k + \hat{v}_{t+1}(\xi_t^k) - u_{t-1}^k \right]_+, \end{aligned}$$

where we replace the value functions  $Q_{\mathcal{R},t+1}(\cdot)$  by the estimator of the following stage. For stage  $T$  it follows that  $\hat{v}_{T+1}(\xi_T^k) \equiv 0$  and for the first stage,

$$(12.30) \quad \hat{v}(\xi^k) := c_1^{\top} x_1 + \hat{v}_2(\xi_1^k).$$

Equation (12.30) provides a recursive estimator for the cost associated with sample path  $\xi^k$ . This estimator has to be evaluated by backward recursion starting with stage  $T$ . Importantly, formula (12.29) is only used for upper bound estimation, whereas the forward and backward problems in SDDP are still based on the original DPE (12.19)–(12.21). Determining estimator (12.30) for all scenarios  $\xi^k, k \in \mathcal{K}$ , sampled in the forward pass of SDDP, we can form an upper bound estimator



$$(12.31) \quad U^n := \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \hat{v}(\xi^k),$$

which resembles the sample average estimator (3.9) from risk-neutral SDDP.

It can be shown that  $\mathbb{E}_{\xi}[\hat{v}(\xi)] \geq v_{\mathcal{R}}^*$  and that  $U^n$  is an unbiased and consistent estimator of  $\mathbb{E}_{\xi}[\hat{v}(\xi)]$ , so it is a statistical upper bound [121]. However,  $U^n$  is also observed to have a large variance. Kozmík and Morton [121] identify as the main reason for this the fact that only a small portion of the sampled scenarios contributes to estimating  $\text{AVaR}_{\alpha}[\cdot]$ , while most scenarios contribute solely to the expectation. Therefore, a very large number of scenarios would be required for an appropriate estimate.

More crucially, because expectations are not taken conditionally on each stage as in (12.20), and due to division by  $\alpha_t \in (0, 1)$ , small or large values are very likely to propagate from late to earlier stages in the recursion to determine  $\hat{v}(\xi^k)$  [121]. Therefore, the upper bound  $\mathbb{E}_{\xi}[\hat{v}(\xi)]$  can significantly deviate from  $\bar{v}_{\mathcal{R}}$ , i.e., the upper bound induced by the current policy in SDDP. In computational experiments, an upward bias is observed that makes  $U^n$  practically useless for large  $T$  [209].

*Remark 12.9.* We should note that a very similar recursive upper bound estimator to  $\hat{v}(\xi^k)$  and  $U^n$  is proposed in [105], but for a general class of risk measures instead of only  $\hat{\rho}_{\alpha, \lambda}[\cdot]$ . The main difference in that paper is that SDDP is applied to a risk-averse stochastic optimal control model that deviates from our setting introduced in section 2. In particular, states and controls are explicitly distinguished and the decision on the controls is taken before  $\xi_t$  is realized. In this setting, the negative multiplicative effects observed in [121, 209] can be circumvented and a computationally efficient statistical upper bound is obtained.

**Conditional Sampling Estimator.** For the above reasons, estimator  $U^n$  in (12.31) is rarely considered in the literature on risk-averse SDDP. Instead, Shapiro discusses a conditional sampling estimator [205]. Here, the idea is to estimate the expectations (12.20) in the nested structure conditionally by sampling on each stage. Since, in principle, the upper bound estimator can be determined independently of the scenarios sampled in the forward pass, we denote the set of samples by  $\mathcal{M}$  instead of  $\mathcal{K}$ .  $\mathcal{M}_t$  denotes the corresponding scenario set for stage  $t$ .

For each stage,  $t = 2, \dots, T$ , this yields [121]

$$\begin{aligned} \hat{v}_t^c(\xi_t^k) := & \frac{1}{|\mathcal{M}_t|} \sum_{m_t \in \mathcal{M}_t} \left[ (1 - \lambda_t) \left( (c_t(\xi_t^{m_t}))^\top x_t^{m_t} + \hat{v}_{t+1}^c(\xi_t^{m_t}) \right) \right. \\ & \left. + \lambda_t u_{t-1}^{m_t} + \frac{\lambda_t}{\alpha_t} \left[ (c_t(\xi_t^{m_t}))^\top x_t^{m_t} + \hat{v}_{t+1}^c(\xi_t^{m_t}) - u_{t-1}^{m_t} \right]_+ \right], \end{aligned}$$

and for the first stage the estimator

$$U^c := c_1^\top x_1 + \hat{v}_2^c(\xi_1).$$

As Shapiro himself points out, this estimator has two significant drawbacks. It requires  $\prod_{t=2}^T |\mathcal{M}_t| + 1$  subproblems to be solved, and this grows exponentially in the number of stages. Moreover, the upper bounds obtained are typically not very tight. Therefore, estimator  $U^c$  is not useful for large-scale problems [121].

**Importance Sampling Estimators.** Some of the drawbacks of estimator  $U^n$  can also be addressed by importance sampling [120, 121]; see section 6 for an introduction.

By sampling scenarios associated with  $\text{AVaR}_\alpha[\cdot]$  with higher importance, it is possible to better represent it and thus reduce the variance of the estimator. Based on this idea, Kozmík and Morton put forward different importance sampling upper bound estimators [121], which are further enhanced in [120].

Using importance sampling with respect to  $\text{AVaR}_\alpha[\cdot]$  creates a considerable challenge, though. In order to determine the importance sampling distribution for some stage  $t$ , the scenarios that are associated with  $\text{AVaR}_\alpha[\cdot]$  on that stage have to be identified, i.e., which of them provide a value  $Q_{\mathcal{R},t}(x_{t-1}^k, \xi_{tj}^k)$  beyond the  $(1 - \alpha)$ -quantile. If we estimate this by solving subproblems for several  $\xi_{tj}^k$  and determining  $Q_{\mathcal{R},t}(x_{t-1}^k, \xi_{tj}^k)$ , we face a similar computational burden as that for conditional sampling.

Kozmík and Morton suggest the following approach: They use an *approximation function*  $d_t(x_{t-1}, \xi_t)$ , which estimates the recourse value of the decisions  $x_{t-1}$  after  $\xi_t$  has been observed [121]. Instead of solving the subproblems for several  $\xi_{tj}^k$ , they simply evaluate  $d_t(x_{t-1}^k, \xi_{tj}^k)$  and sort these values. Based on the order obtained, it can then be decided which scenarios are used to estimate  $\text{AVaR}_\alpha[\cdot]$ , i.e.,  $u_t := \text{VaR}_{\alpha_t}[d_t(x_{t-1}, \xi_t)]$  is determined.

This allows the definition of an importance sampling distribution depending on  $x_{t-1}$  [121]. For simplicity, we assume that all scenarios are equally likely in the original stage- $t$  distribution  $F_t$  of  $\xi_t$ ; that is, the corresponding probability density function  $f_t(\cdot)$  satisfies  $f_t(\xi_{tj}) = \frac{1}{q_t}$  for all  $j = 1, \dots, q_t$ . Then, it follows that

$$g_t(\xi_t | x_{t-1}) := \begin{cases} \frac{1}{2[\alpha_t q_t]}, & d_t(x_{t-1}, \xi_t) \geq u_t, \\ \frac{1}{2(q_t - [\alpha_t q_t])}, & d_t(x_{t-1}, \xi_t) < u_t. \end{cases}$$

This distribution ensures that it is equally likely to draw sample observations above and below  $u_t$ . Note that the formula presented in [121] looks a bit different, since it is presented in the context of SAA.

Defining weights

$$\Lambda_t(\xi_t | x_{t-1}) := \frac{f_t(\xi_t)}{g_t(\xi_t | x_{t-1})}$$

and multiplying them along the sample paths

$$\Lambda(\xi^k) := \prod_{t=2}^T \Lambda_t(\xi_t^k | x_{t-1}),$$

we can derive the estimator

$$(12.32) \quad U^i := \frac{1}{\sum_{k \in \mathcal{K}} \Lambda(\xi^k)} \sum_{k \in \mathcal{K}} \Lambda(\xi^k) \hat{v}(\xi^k).$$

This estimator is similar to (12.31), as the same recursive term  $\hat{v}(\xi^k)$  is used, but it is combined with importance instead of standard MC sampling.

With the assumptions of relatively complete recourse (based on Assumption 9) and stagewise independence (Assumption 2), estimator (12.32) is asymptotically valid, i.e., for  $|\mathcal{K}| \rightarrow \infty$ ,  $U^i$  converges to  $\mathbb{E}_f[\hat{v}(\xi)]$  with probability 1 (recall that  $\mathbb{E}_f[\hat{v}(\xi)] \geq v_{\mathcal{R}}^*$ ). Moreover, for a sufficiently good choice of  $d_t(\cdot)$ , it can be assumed that the variance is lower than for  $U^n$  [121].

Based on this idea, even better estimators are developed in [120, 121], for example, by not only sampling scenarios associated with  $\text{AVaR}_\alpha[\cdot]$  with higher priority, but also using only scenarios that contribute to the  $[\cdot]_+$ -term to estimate AVaR [121]:

$$\begin{aligned} \hat{v}_t^d(\xi_t^k) &:= (1 - \lambda_t) \left( (c_t(\xi_t^k))^\top x_t^k + \hat{v}_{t+1}^d(\xi_t^k) \right) \\ &\quad + \lambda_t u_{t-1}^k + \mathcal{I}[d_t(x_{t-1}, \xi_t) \geq u_d] \frac{\lambda_t}{\alpha_{t-1}} \left[ (c_t(\xi_t^k))^\top x_t^k + \hat{v}_{t+1}^d(\xi_t^k) - u_{t-1}^k \right]_+. \end{aligned}$$

Here,  $\mathcal{I}[\cdot]$  denotes an indicator function. For the first stage it follows that

$$\hat{v}^d(\xi^k) := c_1^\top x_1 + \hat{v}_2^d(\xi_1^k).$$

Combining this with (11.6), we obtain

$$U^d := \frac{1}{\sum_{k \in \mathcal{K}} \Lambda(\xi^k)} \sum_{k \in \mathcal{K}} \Lambda(\xi^k) \hat{v}^d(\xi^k).$$

The practical applicability of this estimator relies heavily on satisfaction of the following goodness assumption with respect to  $d_t(\cdot)$ :

$$Q_{\mathcal{R},t}(x_{t-1}, \xi_t) \geq \text{VaR}_{\alpha_t}[Q_{\mathcal{R},t}(x_{t-1}, \xi_t)] \Leftrightarrow d_t(x_{t-1}, \xi_t) \geq \text{VaR}_{\alpha_t}[d_t(x_{t-1}, \xi_t)],$$

which means that  $d_t(\cdot)$  correctly classifies whether or not a realization is in the upper  $\alpha$ -tail of the recourse value distribution.

It has been proven that this estimator is asymptotically valid as well, but that it also provides tighter upper bounds than  $U^i$  in expectation, as long as the goodness assumption above is satisfied. Moreover, a smaller variance should be expected [121]. Numerical results in [121] illustrate that even for a medium number of stages, estimator  $U^d$  provides significantly better upper bounds than  $U^n, U^c$ , and  $U^i$  and that the variance of the estimators is also reduced significantly. However, despite reducing the variance, even  $U^i$  and  $U^d$  may still show a considerable upward bias with respect to the upper bound  $\bar{v}_{\mathcal{R}}$  induced by the current policy [209].

Apart from the sampling estimators above, some completely different strategies may be used to obtain upper bounds for  $v_{\mathcal{R}}^*$  or to define some stopping criteria for SDDP in the risk-averse case.

**Using Deterministic Upper Bounds.** As discussed in section 8, we may circumvent the determination of sampling-based upper bound estimators completely if we resort to deterministic upper bounding procedures.

To this end, Philpott, de Matos, and Finardi [168] extend their inner approximation based upper bounding procedure from section 8 to the risk-averse case with nested (conditional) coherent risk measures. However, the main downside of this procedure, the required prohibitively large computational effort for a large number of state variables and an increasing number of cuts, also holds in this case.

The alternative deterministic upper bounding procedure based on dual SDDP [104, 126] has been extended to a risk-averse setting as well [43, 44].

**Using a Change of Probability Measure.** As discussed in section 12.3.1, following the approach of a change of probability measure (see (12.15)–(12.17) and (12.27)), it is also possible to run (risk-averse) SDDP once in advance to approximate the probability measure  $\bar{\mathbb{P}}$  and then a second time where the probability measure is fixed to the approximation of  $\bar{\mathbb{P}}$ . This is referred to as solving the *change-of-measure risk-neutral problem* in [133]. Whereas this approach has a lot of computational

overhead, the advantage is that a risk-neutral problem can be solved by SDDP and, therefore, the standard stopping, upper bounding, and policy assessment techniques can also be applied. Solving the change-of-measure risk-neutral problem is not guaranteed to yield optimal policies for  $(P_{\mathcal{R}})$ , but Liu and Shapiro [133] report that the quality of the policies is similar to those obtained by risk-averse SDDP.

A similar approach is used in [80], but without running SDDP twice. Instead, the biased sampling strategy from section 12.3.2 is applied based on approximating  $\tilde{\mathbb{P}}$  with information gained from using multicut SDDP. Again, the idea is that by changing the probability measure when sampling, statistical upper bounds can be computed in a standard fashion. Whereas this approach comes with almost no computational overhead (except for relying on multicut SDDP; see section 21.2.1), the theoretical properties of the estimators obtained are not explored.

**Fixing the Number of Iterations.** This approach is proposed by Philpott and de Matos [167]. They run a risk-neutral variant of SDDP first and then fix the number of iterations required until termination. The same number of iterations is then used in the risk-averse case, avoiding the challenge of upper bound evaluation.

In some practical applications in which it is computationally intractable to determine a sophisticated upper bound estimator, this approach may be useful. Promising results are reported in [167]. However, there is no theoretical guarantee of finding a sufficiently good solution for a risk-averse version of  $(P_{\mathcal{R}})$  in the same number of iterations as for a risk-neutral version. Additionally, for large problems it may already take a considerably long time to run SDDP once. Running it a second time for risk-averse problem  $(P_{\mathcal{R}})$  may partially annihilate the computational advantage of avoiding upper bound estimation.

**Lower Bound Stabilization.** As for risk-neutral SDDP, instead of using upper bounds at all, the algorithm can be terminated once the lower bounds  $v_{\mathcal{R}}^i$  stabilize. This does not provide a convergence guarantee but may be worthwhile in large-scale practical applications where other approaches become computationally prohibitive.

**Using Benefit Factors.** Instead of the lower bounds  $v_{\mathcal{R}}^i$ , it is also possible to condition termination of SDDP on the improvements of the cut approximations  $\mathcal{V}_{\mathcal{R},t}^i(\cdot)$ ,  $t = 2, \dots, T$ . For that purpose, Brandi et al. define a benefit factor

$$\mathcal{B}_{t,k}^i = \min \left\{ 1, \frac{\delta(x_{t-1}^{ik})}{\delta_{t,\max}^i} \right\},$$

which determines how much a new cut improves the current cut approximation  $\mathcal{V}_{\mathcal{R},t}^i(\cdot)$  at  $x_{t-1}^{ik}$  [32].  $\delta(x_{t-1}^{ik})$  is the absolute increase, while  $\delta_{t,\max}^i$  is a proxy for the maximum improvement possible. For each sample path  $k \in \mathcal{K}$ , a total benefit factor can be determined by

$$\mathcal{B}_k^i = \max \{ \mathcal{B}_{2,k}^i, \mathcal{B}_{3,k}^i, \dots, \mathcal{B}_{T,k}^i \}.$$

The risk-averse SDDP method is then stopped if the values  $\mathcal{B}_k^i$  for all  $k \in \mathcal{K}$  are below a predefined tolerance, either for one iteration or, alternatively and more robustly, for a predefined larger number of iterations.

**12.4. SDDP with Entropic Risk Measure.** As discussed before, nested risk measures come with some drawbacks. Computationwise, upper bound determination is very challenging. In addition, applying a standard one-period risk measure  $\rho[\cdot]$ , e.g.,  $\text{AVaR}_{\alpha}[\cdot]$ , as an end-of-horizon risk measure (12.10) and (possibly conditionally) in a nested risk measure (12.11) does not yield equivalent policies [64] (this is only the case

if we take the composite risk measure associated with the nested risk measure as end-of-horizon risk; however, this risk measure is usually not known explicitly; see Remark 12.7). This makes nested risk measures difficult to interpret from an end-of-horizon perspective.

For this reason, Dowson, Morton, and Pagnoncelli [64] propose applying one-period *conditionally consistent* risk measures in the context of SDDP [64]; see also [11, 165]. It can be proven that under some technical assumptions, the class of entropic risk measures  $\mathbb{ENT}_\gamma[\cdot]$  (see (12.8)) is the only class of risk measures that is conditionally consistent.

As  $\mathbb{ENT}_\gamma[\cdot]$  can be applied in a nested fashion, the DPE (12.15)–(12.17) are valid in this case. Moreover, since  $\mathbb{ENT}_\gamma[\cdot]$  is a convex risk measure, the (risk-adjusted) value functions are convex. Therefore, SDDP can be applied to derive polyhedral outer approximations.

As for standard SDDP, first, for each scenario  $k \in \mathcal{K}$  and all possible stage- $t$  realizations  $\xi_{tj}^k \equiv \xi_{tj}, j = 1, \dots, q_t$ , approximate versions of subproblems (12.15) are solved to obtain  $\underline{Q}_{\mathcal{R},t}^i(x_{t-1}^k, \xi_{tj})$ . Then, based on the dual form of  $\mathbb{ENT}_\gamma[\cdot]$ , the following auxiliary problem can be solved to evaluate the risk-adjusted value function:

$$(12.33) \quad \begin{aligned} & \mathbb{ENT}_\gamma[\underline{Q}_{\mathcal{R},t}^i(x_{t-1}^k, \xi_t)] \\ &= \begin{cases} \max_{\tilde{p}_t} & \sum_{j=1}^{q_t} \tilde{p}_{tj} \underline{Q}_{\mathcal{R},t}^i(x_{t-1}^k, \xi_{tj}) - \frac{1}{\gamma_t} \sum_{j=1}^{q_t} \tilde{p}_{tj} \cdot \log\left(\frac{\tilde{p}_{tj}}{p_{tj}}\right) \\ \text{s.t.} & \sum_{j=1}^{q_t} \tilde{p}_{tj} = 1, \\ & \tilde{p}_{tj} \geq 0, \quad j = 1, \dots, q_t. \end{cases} \end{aligned}$$

Here, parameter  $p_{tj}$  denotes the nominal probabilities of realizations  $\xi_{tj}$ , which usually equal  $\frac{1}{q_t}$ , and the decision variable  $\tilde{p}_{tj}$  denotes an alternative probability based on the entropic risk measure. In this way, problem (12.33) can be regarded as building the expectation based on some modified probability measure and with some additional penalty term. Problem (12.33) can be solved algorithmically, but as stated in [64], a closed form for  $\tilde{p}_{tj}^*$  can also be derived. Using  $\tilde{p}_{tj}^*$  and  $\mathbb{ENT}_\gamma[\underline{Q}_{\mathcal{R},t}^i(x_{t-1}^k, \xi_t)]$ , cuts can then be constructed and handed back to the previous stage.

The entropic risk measure not only ensures conditional consistency of the obtained policies, but also allows for upper bound computation as in standard SDDP, because the tower property can be employed for  $\mathbb{ENT}[\cdot]$ . However, these advantages come at the cost of an aggravated interpretation of the risk measure compared to AVaR-based risk measures. In this context, it is particularly difficult to make a reasonable choice for the parameter  $\gamma_t > 0$  [64].

**12.5. SDDP with Expected Conditional AVaR.** Another class of multiperiod risk measures that can be used as an alternative to nested risk measures is that of *expected conditional risk measures*, which we briefly introduced in section 12.1.3 [68, 113]. Here, conditional expectations are used to avoid the risk measure nesting, which proves beneficial in determining upper bounds in SDDP as it avoids the aforementioned computational difficulties while time consistency is still ensured.

Recall the risk-averse problem  $(P_{\mathcal{R}})$  using expected conditional risk measures stated in (12.14). Using  $\rho_t[\cdot] = \text{AVaR}_{\alpha_t}[\cdot]$  yields the so-called  $\mathbb{E}$ -AVaR or *multiperiod average value-at-risk* [113], which goes back to Pflug and Ruszczyński [166].

As stated in [113], by some lengthy reformulations, the objective function of problem (12.14) can be expressed in a nested manner. Therefore, equivalent DPE

can be derived and time consistency is assured. Moreover, the  $[\cdot]_+$ -function can be reformulated using an epigraph approach. Then, for  $t = 2, \dots, T$ , the DPE read

$$(12.34) \quad \check{Q}_{\mathcal{R},t}(x_{t-1}, u_t, \xi_t) = \begin{cases} \min_{x_t, u_{t+1}, w_t} & \frac{1}{\alpha_t} w_t + u_{t+1} + \check{\mathcal{Q}}_{\mathcal{R},t+1}(x_t, u_{t+1}) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t), \\ & w_t - (c_t(\xi_t))^\top x_t \geq -u_t, \\ & w_t \geq 0, \end{cases}$$

with

$$(12.35) \quad \check{\mathcal{Q}}_{\mathcal{R},t+1}(x_t, u_{t+1}) = \mathbb{E}_{\xi_{t+1}} \left[ \check{Q}_{\mathcal{R},t}(x_{t-1}, u_t, \xi_t) \right],$$

$\check{\mathcal{Q}}_{\mathcal{R},T+1}(\cdot, \cdot) \equiv 0$ , and first stage

$$(12.36) \quad v_{\mathcal{R}}^* = \begin{cases} \min_{x_1, u_2} & c_1^\top x_1 + u_2 + \check{Q}_{\mathcal{R},2}(x_1, u_2, \xi_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

In contrast to using nested conditional risk measures, the DPE here only depend on nested sums of (conditional) expectations, i.e., they have the same structure as in the risk-neutral case and hence standard SDDP can be applied. This has the advantage of allowing the use of upper bounding techniques developed for risk-neutral SDDP.

**12.6. Biobjective SDDP.** An alternative to risk-averse formulations that allows one to achieve a trade-off between obtaining the best policy in expectation (e.g., the policy with the lowest expected costs) and avoiding bad extreme outcomes (e.g., power outages or load shedding in an electricity network) is to formulate an (MSLP) with multiple competing objectives that are optimized simultaneously. Recently, a variant of SDDP for biobjective problems was put forward by Dowson, Morton, and Downward [62].

Let  $\tilde{c}_t(\xi_t)$  and  $\hat{c}_t(\xi_t)$  denote the objective coefficients for stage  $t \in [T]$  and the two competing objectives. For all but trivial cases, there exists no policy that yields the best objective value with respect to both objectives

$$\tilde{v}^* := \min_{x_1, \mathbf{x}_2, \dots, \mathbf{x}_T} \underbrace{\mathbb{E} \left[ \sum_{t \in [T]} (\tilde{c}_t(\xi_t))^\top \mathbf{x}_t(\xi_{[t]}) \right]}_{=: \tilde{v}(\mathbf{x})}$$

and

$$\hat{v}^* := \min_{x_1, \mathbf{x}_2, \dots, \mathbf{x}_T} \underbrace{\mathbb{E} \left[ \sum_{t \in [T]} (\hat{c}_t(\xi_t))^\top \mathbf{x}_t(\xi_{[t]}) \right]}_{=: \hat{v}(\mathbf{x})},$$

meaning that the two objectives are truly conflicting.

For this reason, if there is no clear preference for one of the objectives, usually the aim is to compute *Pareto-optimal* policies. A policy  $(\bar{\mathbf{x}}_t(\xi_{[t]}))_{t \in [T]}$  is Pareto-optimal if it cannot be improved in one objective without getting worse in the other one, i.e., if there exists no other policy  $(\mathbf{x}_t(\xi_{[t]}))_{t \in [T]}$  such that  $\tilde{v}(\mathbf{x}) \geq \tilde{v}(\bar{\mathbf{x}})$  and  $\hat{v}(\mathbf{x}) > \hat{v}(\bar{\mathbf{x}})$

(or the other way round). Pareto-optimal solutions are also called *nondominated*, and the set of nondominated objective vectors is called the *Pareto front* [62].

A standard approach to computing Pareto-optimal solutions in optimization is to use some *scalarization* approach in which both conflicting objectives are combined into a weighted sum that is then optimized in a deterministic single-objective problem. In our case, the DPE (2.4)–(2.6) can be adapted to

$$(12.37) \quad Q_t(x_{t-1}, \xi_t, \lambda) := \begin{cases} \min_{x_t} & (\lambda \tilde{c}_t(\xi_t) + (1 - \lambda) \hat{c}_t(\xi_t))^\top x_t + \mathcal{Q}_{t+1}(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t), \end{cases}$$

where

$$(12.38) \quad \mathcal{Q}_{t+1}(x_t, \lambda) := \mathbb{E}_{\xi_{t+1}} [Q_{t+1}(x_t, \xi_{t+1}, \lambda)]$$

and  $\mathcal{Q}_{T+1}(x_T) \equiv 0$ . For the first stage, we obtain

$$(12.39) \quad v^*(\lambda) = \begin{cases} \min_{x_1} & (\lambda \tilde{c}_1 + (1 - \lambda) \hat{c}_1)^\top x_1 + Q_2(x_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

SDDP can then be applied to these DPE. In the proposed variant,  $\lambda$  is adapted dynamically. To this end, in each iteration  $i$ , after the backward pass, one stage  $t \in [T]$  is randomly and independently sampled and the corresponding subproblem is solved again for  $x_{t-1}^k$ ,  $\xi_t^k$ , and  $\lambda^i$ . Then,  $\lambda^i$  is updated to  $\lambda^{i+1}$ , where the latter is determined as being the closest  $\lambda$  to  $\lambda^i$  such that the optimal basis of the constraint equation system changes.

It has been proven that this variant of SDDP converges almost surely to the Pareto front of biobjective (MSLP) in finitely many iterations. Note that, technically speaking, not all Pareto-optimal policies are guaranteed to be identified by SDDP because for some  $\lambda$ , multiple optimal policies may exist. However, all Pareto-optimal policies for which  $(\hat{v}(\mathbf{x}), \tilde{v}(\mathbf{x}))$  cannot be represented as a strict convex combination of other nondominated objective vectors are identified under weak assumptions [62].

**13. SDDP with Unknown Distribution [Relaxing Assumption 3].** In section 3 we introduced SDDP assuming that the probability distribution  $F_\xi$  of the data process  $(\xi_t)_{t \in [T]}$  governing the uncertainty in problem (MSLP) is known; see Assumption 3. This allowed us to sample from this specific distribution in the forward pass of SDDP or, in the case of continuous random vectors, to obtain a finite sample average approximation, as described in section 11.

However, in general the true distribution  $F_\xi$  is not usually known in practical applications. Often, only historical data is available, i.e., some realization of an unknown true distribution. This data is then used to determine a reasonable estimate for the true distribution, from which the required samples are taken. However, using such an estimation imposes the risk of *overfitting* the SDDP policies to this specific distribution and thus the available data. Philpott, de Matos, and Kapelevich [169] identify this problem as particularly noteworthy if the number of possible outcomes  $q_t$  per stage is small. For this reason, it may be reasonable to take a more robust approach and factor in the distributional uncertainty. The consideration of this type of uncertainty in SDDP is a young research area.

**13.1. Distributionally Robust SDDP.** One way to consider distributional uncertainty in SDDP is by integrating ideas from robust optimization [16, 20] into (multistage) stochastic programming. More precisely, a set of potential distributions is

considered called the *distributional uncertainty set* or the *ambiguity set* and denoted by  $\mathcal{P}$ . The expected cost is then minimized over the worst-case probability distribution from this set. This is called *distributionally robust optimization* (DRO).

Usually, the outcomes of the random variables  $\xi_t$  are fixed to a finite number of realizations observed in the historical data. The ambiguity set  $\mathcal{P}_t$  then models a variety of potential probability measures  $\mathbb{P}_t \in \mathcal{P}_t$  supported on this finite set  $\Xi_t$ .

In what follows, we restrict to DRO in the SDDP context specifically. For a general introduction to DRO, we refer to the review [182] and the tutorial [206]. We assume that all assumptions from section 3 hold, except for Assumption 3. Furthermore, we only consider uncertainty in the RHS.

Then, the distributionally robust version of (MSLP) can be written as

$$(13.1) \quad \begin{aligned} \min_{x_1, x_2, \dots, x_T} \max_{\mathbb{P} \in \mathcal{P}} \quad & \mathbb{E} \left[ \sum_{t \in [T]} (c_t(\xi_t))^\top x_t(\xi_t) \right] \\ \text{s.t.} \quad & x_1 \in \mathcal{X}_1, \\ & x_t \in \mathcal{X}_t(x_{t-1}(\xi_{[t-1]}), \xi_t) \quad \forall \xi_t \in \Xi_t \quad \forall t = 2, \dots, T. \end{aligned}$$

*Remark 13.1.* Distributionally robust stochastic programming is closely related to risk-averse stochastic programming. In particular, the operator  $\max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}[\cdot]$  can be interpreted as a multiperiod risk measure  $\mathcal{R}[\cdot]$ . This risk measure is coherent [206].

For SDDP it is necessary to reformulate problem (13.1) by means of DPE. This requires that each distribution  $\mathbb{P}$  in the ambiguity set  $\mathcal{P}$  can be expressed as the cross-product of the respective marginal distributions of random vectors  $\xi_t$  [206]. Formally,

$$\mathcal{P} := \{ \mathbb{P} = \mathbb{P}_1 \times \dots \times \mathbb{P}_T \mid \mathbb{P}_t \in \mathcal{P}_t, t \in [T] \}.$$

The ambiguity sets  $\mathbb{P}_t$  are assumed to be independent of one another. This property is called *rectangularity* of  $\mathcal{P}$  and is reminiscent of the stagewise independence assumption for vectors  $\xi_t$ . Note that  $\mathcal{P}_1$  is a singleton containing one distribution with one possible realization.

With the ambiguity sets  $\mathcal{P}_t$ , the DPE can then be written as

$$(13.2) \quad Q_{DR,t}(x_{t-1}, \xi_t) := \begin{cases} \min_{x_t} & c_t^\top x_t + \mathcal{Q}_{DR,t+1}(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t) \end{cases}$$

with

$$(13.3) \quad \mathcal{Q}_{DR,t+1}(x_t) := \max_{\mathbb{P}_{t+1} \in \mathcal{P}_{t+1}} \mathbb{E}_{\mathbb{P}_{t+1}} [Q_{DR,t+1}(x_t, \xi_{t+1})]$$

and  $\mathcal{Q}_{DR,T+1}(x_T) \equiv 0$ . Compared to section 3, here an inner maximization problem is introduced when defining  $\mathcal{Q}_{DR,t+1}(\cdot)$  to obtain the expected cost over the worst-case probability measure in  $\mathcal{P}_{t+1}$ . The first-stage problem reads

$$(13.4) \quad v_{DR}^* = \begin{cases} \min_{x_1} & c_1^\top x_1 + \mathcal{Q}_{DR,2}(x_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

How  $v_{DR}^*$  and a corresponding optimal policy can be computed algorithmically depends heavily on the specific choice of the ambiguity sets  $\mathcal{P}_t, t = 2, \dots, T$ . Various ambiguity sets are proposed in the literature. Usually, these sets are defined in such



a manner that they contain all distributions that are in *some sense* within a given range of some nominal distribution. This nominal distribution, denoted by  $\bar{\mathbb{P}}_t$ , is in turn defined by probabilities  $\bar{p}_{tj} = \frac{1}{q_t}$  for all  $j = 1, \dots, q_t$ , where  $q_t$  denotes the number of historical data samples. Based on the measure employed to evaluate the distance between two distributions or probability measures, respectively, different classes of ambiguity sets can be defined.

For SDDP, the following three distance measures have been used up to this point. In [114], the  $\ell_\infty$  metric with parameter  $r > 0$  is used to define the ambiguity set

$$(13.5) \quad \mathcal{P}_t = \left\{ \mathbb{P}_t \mid \sum_{i=1}^{q_t} p_{ti} = 1, p_{ti} \geq 0, \|p_t - \bar{p}_t\|_\infty \leq r \right\}.$$

A similar metric, but with the  $\ell_2$ -norm, is used in [169] to define the ambiguity set

$$(13.6) \quad \mathcal{P}_t = \left\{ \mathbb{P}_t \mid \sum_{i=1}^{q_t} p_{ti} = 1, p_{ti} \geq 0, \|p_t - \bar{p}_t\|_2 \leq r \right\}.$$

This is a special case of the class of  $\phi$ -divergence distances; see [12]. Both these distance measures are only applicable to discrete distributions supported on the observed data points, and contrary to the Wasserstein distance discussed in what follows, they do not go outside this initial support.

On the contrary, the Wasserstein distance allows for the comparison of general distributions (see, for instance, [224]). In our case with finite distributions  $\mathbb{P}_t$  and  $\bar{\mathbb{P}}_t$ , the Wasserstein distance can be defined by the minimization problem

$$\begin{aligned} d_W(\bar{\mathbb{P}}_t, \mathbb{P}_t) := \min_z \quad & \sum_{i=1}^{q_t} \sum_{j=1}^{q_t} \|\xi_t^i - \xi_t^j\| z_{ij} \\ \text{s.t.} \quad & \sum_{j=1}^{q_t} z_{ij} = \bar{p}_{ti} \quad \forall i = 1, \dots, q_t, \\ & \sum_{i=1}^{q_t} z_{ij} = p_{tj} \quad \forall j = 1, \dots, q_t, \\ & z_{ij} \geq 0 \quad \forall i, j = 1, \dots, q_t, \end{aligned}$$

where different choices are possible for the norm. It can be interpreted as the amount of probability mass that has to be moved between the distributions. This distance is used in [69] to define the Wasserstein ambiguity set

$$(13.7) \quad \mathcal{P}_t = \left\{ \mathbb{P}_t \mid \sum_{i=1}^q p_{ti} = 1, p_{ti} \geq 0, d_W(\bar{\mathbb{P}}_t, \mathbb{P}_t) \leq r \right\}.$$

In all three cases, very different strategies are chosen to apply SDDP to the nested min-max structure defined by the DPE (13.2)–(13.4).

**13.1.1. Reformulation as a Risk-Averse Problem.** As shown in [114], using the  $\ell^\infty$ -ambiguity set (13.5), the DPE (13.2)–(13.4) can be reformulated to those of a risk-averse multistage problem with nested conditional AVaR $_\alpha$ [·]; that is, (12.19)–(12.21) with

$$\lambda_{t+1} = 1 - \mu_{t+1}^\ell, \quad \alpha_{t+1} = \frac{\lambda_{t+1}}{\mu_{t+1}^u - \mu_{t+1}^\ell}.$$

Here,  $\mu_{t+1}^\ell := \sum_{i=1}^{q_{t+1}} p_{t+1,i}^\ell$  and  $\mu_{t+1}^u := \sum_{i=1}^{q_{t+1}} p_{t+1,i}^u$ , where  $p_{t+1}^\ell := \bar{p}_{t+1} - r$  denotes the probabilities associated with the probability measure at the lower bound of ambiguity set (13.5) (and  $p_{t+1}^u := \bar{p}_{t+1} + r$  is defined analogously for the upper bound). Therefore, SDDP can be applied as in this risk-averse setting.

**13.1.2. Solving the Inner Maximization Problem Separately.** In [169] an  $\ell^2$ -ambiguity set (13.6) is considered. The authors then exploit that for convex ambiguity sets, such as the  $\ell^2$ -ambiguity set,  $\max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[\cdot]$  can be interpreted as a coherent risk measure (see (12.9)), and thus the value functions in the DPE (13.2)–(13.4) remain convex.

To derive linear cuts to approximate these value functions, it is proposed to solve the inner maximization problem identifying the worst-case distribution separately. In the backward pass, for some stage  $t$ , first the subproblems are solved for all  $j = 1, \dots, q_t$  as usual (see line 16 of Algorithm 3.1). Then, using the obtained values of  $Q_t^i(x_{t-1}^{ik}, \xi_{tj})$ , the inner maximization problem is solved. This can be done algorithmically and in some cases even analytically, as is shown in [169]. The worst-case probability measure  $\mathbb{P}^*$  obtained can then be used to compute subgradients and cut coefficients. Even though these coefficients are determined based on cut approximation  $\underline{V}_t^{i+1}(\cdot)$  and on  $\mathbb{P}^*$ , which does not necessarily coincide with the worst-case probability measure in the true DPE, valid cuts are constructed and convergence is ensured [169].

**13.1.3. Using a Dual Representation.** If we use the Wasserstein ambiguity set (13.7) in SDDP, we obtain the inner maximization problem

$$\begin{aligned} \max_{z_t, p_{t+1}} \quad & \sum_{j=1}^{q_t} p_{t+1,j} Q_{t+1}(x_t, \xi_{t+1,j}) \\ \text{s.t.} \quad & \sum_{i=1}^{q_t} \sum_{j=1}^{q_t} d_{t+1,ij} z_{tij} \leq 1, \\ & \sum_{j=1}^{q_t} z_{tij} = \bar{p}_{ti} \quad \forall i = 1, \dots, q_t, \\ & \sum_{i=1}^{q_t} z_{tij} = p_{tj} \quad \forall j = 1, \dots, q_t, \\ & z_{tij} \geq 0 \quad \forall i, j = 1, \dots, q_t, \end{aligned}$$

with  $d_{t+1,ij} = \|\xi_{t+1}^i - \xi_{t+1}^j\|$ . Duque and Morton [69] suggest replacing this problem with its dual problem, so that the value functions can be evaluated by solving the single-level minimization problem

$$Q_{DR,t}(x_{t-1}, \xi_t) := \begin{cases} \min_{x_t, \gamma_t, \nu_t} & c_t^\top x_t + r\gamma_t + \sum_{i=1}^{q_{t+1}} q_{t+1}^i \nu_t^i \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t), \\ & d_{t+1,ij}\gamma_t + \nu_{ti} \geq Q_{DR,t+1}(x_t, \xi_{t+1,j}) \quad \forall i, j = 1, \dots, q_{t+1}, \\ & \gamma_t \geq 0, \end{cases}$$

with dual variables  $\gamma_t$  and  $\nu_t$ . As proven in [69], these value functions are piecewise linear and convex on  $\mathcal{X}_{t-1}$  and therefore can be represented by finitely many linear cuts. However, this approach requires the use of multicut SDDP (see section 21.2.1), because otherwise bilinear terms occur.

For all the strategies presented in sections 13.1.1–13.1.3, the forward pass of distributionally robust SDDP remains basically the same as in standard SDDP, with additional flexibility in the sampling step. More precisely, the sampling can be done from the nominal distribution associated with  $\bar{\mathbb{P}}_t, t = 2, \dots, T$ , from the current worst-case distribution associated with  $\mathbb{P}_t^*$ , or more generally from a convex combination of both, defined by some parameter  $\beta \in [0, 1]$ . If independent sampling is conducted, for  $\beta \in [0, 1]$ , convergence of distributionally robust SDDP can be established as for standard SDDP [69]. However, challenges in determining valid upper bounds are similarly prevalent as in the risk-averse case.

Computational results indicate that when taking the dual reformulation approach, better approximations are achieved for multicut SDDP than when solving the inner maximization in a side computation [69]. Furthermore, out-of-sample tests by Philpott, de Matos, and Kapelevich [169] imply that distributionally robust SDDP yields policies which are better suited, e.g., induce lower costs, in periods with a substantial risk of high costs.

**13.2. Partially Observable Distributions.** A different approach to dealing with distributional uncertainty is introduced by Dowson, Morton, and Pagnoncelli in [63] and is referred to as partially observable multistage stochastic programming. The idea is to consider a finite number of potential distributions by combining problem (MSLP) with a hidden Markov model. More precisely, in each stage  $t \in [T]$ , different nodes can be reached, with each node representing one Markov state. Let  $\mathcal{N}$  denote the set of all these nodes except for the root node. Each node reflects a different candidate distribution, possibly with identical realizations  $\xi_j, j = 1, \dots, q$ , but different associated probabilities.

As a key idea, consider a partition  $\mathcal{A}$  of nodes in  $\mathcal{N}$  into ambiguity sets  $A \in \mathcal{A}$ , satisfying  $\bigcup_{A \in \mathcal{A}} A = \mathcal{N}$ . For example, this partition can be chosen such that there is one ambiguity set  $A$  for each stage.

To model the distributional uncertainty, it is now assumed that at any point, only the current ambiguity set is known, while the specific node within it cannot be observed. However, for each node  $i$ , a probability  $b_i$  is available. In other words, each candidate distribution is considered to be the most accurate representation of the true underlying distribution with a certain probability. These probabilities are stored in a so-called *belief state*  $b$ . Each time an ambiguity set  $A$  is entered and a particular realization  $\xi$  of the random data is observed, the belief state is updated componentwise by applying Bayes' theorem [63].

In contrast to (MSLP) with perfect distribution information (see Assumption 3), the value functions  $Q_t(\cdot)$  have to incorporate this belief state. To this end, let  $p_{i\ell}$  be the probability of observing  $\xi_{i\ell}$  conditional on being in node  $i$  with  $\ell = 1, \dots, q^i$ . Let  $\mathcal{N}$  describe all nodes including the root node,  $\omega_{jk}$  the transition probability from node  $j$  to  $k$ , and  $B_k(b, \xi)$  the update rule for the belief state being in (unobservable) node  $k$ . Furthermore, let  $x'$  denote the current trial solution. Then, the expected value function can be written as

$$(13.8) \quad \mathcal{Q}_B(x', b) := \sum_{j \in \mathcal{N}} b_j \sum_{k \in \mathcal{N}} \omega_{jk} \sum_{\ell=1}^{q^k} p_{k\ell} Q_k(x', B_k(b, \xi_{k\ell}), \xi_{k\ell}).$$

This means that the value function  $Q_k(\cdot, \cdot)$  depends on a node and an updated belief state, and in (13.8) it is looped over all nodes, weighing the corresponding expected value with the current belief and the transition probabilities between the nodes.

In the forward pass, for each stage  $t = 2, \dots, T$ , first a new node is sampled conditionally on the (unobserved) current node. Then, a realization of  $\xi$  is sampled conditionally on the obtained node and the associated candidate distribution. Based on these samples, the ambiguity set  $A$  and the belief state  $b$  are updated. The latter can be considered as an additional state. For a more detailed description, see [63].

As proven in Theorem 1 in [63], the expected value functions  $\mathcal{Q}_B(\cdot)$  are saddle functions as they are convex in  $x$  for fixed  $b$ , but concave in  $b$  for fixed  $x$ . Therefore, to apply SDDP, the cut generation has to be adapted to this property; see also subsection 14.6. This can be achieved by using an outer approximation for  $x$  and an inner approximation for  $b$  [63]. Apart from the changed cut formula, line 15 of Algorithm 3.1 loops over all nodes in the current ambiguity set  $A$  and then conditionally on all possible realizations of  $\xi$  on the following stage. In line 18 of Algorithm 3.1, in addition to taking conditional expectations with respect to  $\xi$ , the cut components are summed over all nodes in  $A$  and weighed with the current belief state for these nodes, similar to (13.8) [63].

A different method of combining SDDP with a hidden Markov model is given in [70]. One general drawback of such hidden Markov approaches is that transition probabilities between the nodes have to be properly defined a priori.

**14. Stagewise-Dependent Uncertainty [Relaxing Assumption 2].** As explained in sections 2 and 3, stagewise independence (Assumption 2) is a standard assumption in dynamic programming, and thus also for SDDP. It is also crucial for the computational tractability of SDDP compared to NBD because it ensures that there exists only one expected value function  $\mathcal{Q}_t(\cdot)$  per stage and that cuts can be shared between scenarios; see section 5.2. However, in many applications, the uncertain data in (MSLP) (e.g., demand, fuel prices, electricity prices, inflows) shows correlations over time and so the assumption of stagewise independence is not appropriate.

If the uncertainty in problem (MSLP) is stagewise-dependent, the expected value functions  $\mathcal{Q}_t(\cdot)$  for  $t = 2, \dots, T$  not only depend on  $x_{t-1}$ , but also depend implicitly on the history  $\xi_{[t-1]}$  of the process  $(\xi_t)_{t \in [T]}$ . In order to apply SDDP, this dependence has to be taken into account, for instance, by reformulating the model or adapting the algorithmic steps in SDDP. In this section, we consider different cases of stagewise-dependent uncertainty and ways that SDDP can be applied in these cases.

**14.1. Expanding the State Space.** As a first case of stagewise-dependent uncertainty, let us assume that the data process  $(\xi_t)_{t \in [T]}$  is a simple linear AR process with lag 1 defined by appropriately chosen coefficient vectors  $\gamma_t$ , matrices  $\Phi_t$ , and stagewise-independent and i.i.d. error terms  $\eta_t$ :

$$(14.1) \quad \xi_t = \gamma_t + \Phi_t \xi_{t-1} + \eta_t.$$

*Remark 14.1.* If we still assume finite randomness (Assumption 5), now applied to  $\eta_t$ , then  $\xi_t$  can be modeled by a classical scenario tree; see section 5.2.

The most natural approach to dealing with this case is to reformulate (MSLP) in such a way that it exhibits stagewise-independent uncertainty [161]. This can be achieved by including  $\xi_{t-1}$  as an additional state variable. Then, as shown in [136],

$$\begin{aligned} \mathbb{E}_{\xi_t | \xi_{t-1}} [Q_t(x_{t-1}, \xi_t)] &= \mathbb{E}_{\eta_t | \xi_{t-1}} [Q_t(x_{t-1}, \gamma_t + \Phi_t \xi_{t-1} + \eta_t)] \\ &= \mathbb{E}_{\eta_t} [Q_t(x_{t-1}, \gamma_t + \Phi_t \xi_{t-1} + \eta_t)], \end{aligned}$$

where the second equality holds because  $\eta_t$  and  $\xi_{t-1}$  are statistically independent.

By introducing (14.1) as a constraint and defining a new value function

$$(14.2) \quad \widehat{Q}_t(x_{t-1}, \xi_{t-1}, \eta_t) := Q_t(x_{t-1}, \gamma_t + \Phi_t \xi_{t-1} + \eta_t),$$

and the corresponding expected value function

$$(14.3) \quad \widehat{\mathcal{Q}}_t(x_{t-1}, \xi_{t-1}) := \mathbb{E}_{\eta_t} [\widehat{Q}_t(x_{t-1}, \xi_{t-1}, \eta_t)]$$

for all  $t = 2, \dots, T$ , it follows that

$$\mathbb{E}_{\xi_t | \xi_{t-1}} [Q_t(x_{t-1}, \xi_t)] = \widehat{\mathcal{Q}}_t(x_{t-1}, \xi_{t-1}).$$

The state variables then consist of the resource state  $x_{t-1}$  and the information state  $\xi_{t-1}$ , while the stagewise-independent uncertainty is modeled by  $\eta_t$ . Importantly,  $\xi_t$  is regarded as a decision variable in the reformulated problem, augmenting the dimension of the decision space.

*Remark 14.2.* It is worth emphasizing that this approach is presented in different ways in the literature. In some cases, as outlined, (14.1) is explicitly incorporated into the DPE as an additional constraint [180, 211]. In some cases, each occurrence of  $\xi_t$  in the subproblems is simply replaced by the RHS of (14.1). In other cases, the dependence on  $\xi_{t-1}$  is only expressed by writing  $\widehat{Q}_t(\cdot, \cdot, \cdot)$  and  $\widehat{\mathcal{Q}}_t(\cdot, \cdot)$  as functions of  $\xi_{t-1}$ , whereas the explicit relation (14.1) is only considered in the cut generation process [91, 136, 185]. We revisit this observation in the next subsection.

Using the procedure presented above, stagewise independence (Assumption 2) is recovered for (MSLP). However, in order to apply SDDP, it also must be ensured that valid linear cuts for  $\widehat{\mathcal{Q}}_t(\cdot, \cdot)$  can be derived as functions in both types of state variables. This requires that  $\widehat{\mathcal{Q}}_t(\cdot, \cdot)$  is convex in both  $x_{t-1}$  and  $\xi_{t-1}$ . Similarly to Theorem 2.8, it can be shown that under certain assumptions this property is satisfied.

**THEOREM 14.3** ([185]). *Let  $\xi_t$  be described by (14.1) and let  $\xi_{t-1}$  be contained in some convex set. Then, under Assumptions 1 and 3–9, the expected value function  $\widehat{\mathcal{Q}}_t(\cdot, \cdot)$  is piecewise linear and*

- (a) *convex in  $x_{t-1}$  on  $\mathcal{X}_{t-1}$  for fixed  $\xi_{t-1}$ ;*
- (b) *convex in  $\xi_{t-1} = (T_{t-2}, h_{t-1})$  for fixed  $x_{t-1}, W_{t-1}, c_{t-1}$ ;*
- (c) *concave in  $\xi_{t-1} = c_{t-1}$  for fixed  $x_{t-1}, W_{t-1}, T_{t-2}, h_{t-1}$ ;*
- (d) *convex jointly in  $x_{t-1}$  and in  $\xi_{t-1} = h_{t-1}$  for fixed  $W_{t-1}, T_{t-2}, c_{t-1}$ .*

Theorem 14.3 shows that convexity in both types of state variables is only guaranteed if the stagewise-dependent part of the uncertainty only enters the RHS  $h_t(\xi_t)$  of problem (MSLP). Note that this still allows for additional stagewise-independent uncertainty in  $c_t, W_t$ , and  $T_{t-1}$ . The result also requires linearity of (MSLP) (Assumption 6) and of the AR process (14.1) defining the random variable  $\xi_t$ .

Under certain assumptions, Theorem 14.3 can be generalized to convex problems (MSLP) and stagewise dependence in the RHS defined by a convex function [91]. Moreover, the result is not limited to lag-1 processes, but can be enhanced to AR processes with higher lag order [91]. This is important for practical applications, as often several lags are required to explain a time series appropriately. In contrast, for general nonlinear stochastic processes or for uncertainty in  $W_t, c_t$ , or  $T_{t-1}$ , such a generalization seems to be impossible. In order to cover such cases, different approaches are required, which we discuss later in this section.

For simplicity, assume that  $X_t = \{x_t \in \mathbb{R}^{n_t} \mid x_t \geq 0\}$  for all  $t \in [T]$  and recall the definition of the approximate subproblem (2.10):

$$(14.4) \quad \underline{Q}_t(x_{t-1}, \xi_t) = \begin{cases} \min_{x_t, \theta_{t+1}} & (c_t(\xi_t))^\top x_t + \theta_{t+1} \\ \text{s.t.} & W_t(\xi_t)x_t = h_t(\xi_t) - T_{t-1}(\xi_t)x_{t-1}, \\ & x_t \geq 0, \\ & -(\beta_{t+1}^r)^\top x_t + \theta_{t+1} \geq \alpha_{t+1}^r \quad \forall r \in \Gamma_{t+1}, \end{cases}$$

where  $\Gamma_{t+1}$  is the index set of previously generated cuts. Then, the result in Theorem 14.3 can be illustrated by means of the feasible region of the linear programming dual to (14.4), which can be written as

$$(14.5) \quad \begin{aligned} \max_{\pi_t, \mu_t} & \quad (h_t(\xi_t) - T_{t-1}(\xi_t)x_{t-1})^\top \pi_t + a_{t+1}^\top \mu_t \\ \text{s.t.} & \quad (W_t(\xi_t))^\top \pi_t - B_{t+1}^\top \mu_t \leq c_t(\xi_t), \\ & \quad e^\top \mu_t = 1, \\ & \quad \mu_t \geq 0. \end{aligned}$$

Here, we collect all cut gradients  $\beta_{t+1}^r$  in a matrix  $B_{t+1}$  and all cut intercepts  $\alpha_{t+1}^r$  in a vector  $a_t$  for compact representation.  $\pi_t$  denotes the dual variable to the original constraints, and  $\mu_t$  denotes the dual variable to the previously generated cuts.

In the case of linear AR processes in the RHS  $h_t(\xi_t)$ , the dual feasible region is not affected by the new state variable  $\xi_{t-1}$  (and also remains polyhedral). This means that the extreme solutions obtained for one state  $\xi_{t-1}$  remain valid, although not necessarily optimal, for all other states  $\xi_{t-1}$  as well. In contrast, in other cases of stagewise dependence, the dual feasible region and its extreme solutions may change for different states, which affects the properties of  $\hat{\mathcal{Q}}_t(\cdot, \cdot)$  [185].

In summary, for affine and convex AR processes occurring in the RHS, expanding the state recovers stagewise independence (Assumption 2), but simultaneously convexity of  $\hat{\mathcal{Q}}_t(\cdot, \cdot)$  in all state variables is preserved. Therefore, SDDP can be used as it was introduced in section 3. In this case, the cuts obtained are functions of both state variables and can be formulated with a cut gradient for each of them (cf. (3.5)), i.e.,

$$\phi_t(x_{t-1}, \xi_{t-1}) = \alpha_t + (\beta_t^x)^\top x_{t-1} + (\beta_t^\xi)^\top \xi_{t-1}.$$

Unfortunately, depending on the dimension  $\kappa_{t-1}$  of  $\xi_{t-1}$ , the state space dimension can increase significantly. This effect is amplified for higher lag orders. As the computational complexity of SDDP grows exponentially in this dimension (see section 4.2), augmenting the state space is detrimental and should be avoided if possible.

**14.2. Scenario-Adaptable Cut Formulas.** The previously described adverse effect can be alleviated to some degree by a special cut generation approach that was first proposed by Infanger and Morton [116] and later enhanced by de Queiroz and Morton [180] and Guigues [91]. In all these cases, the process model, such as (14.1), is not explicitly incorporated into the subproblems; see Remark 14.2. Instead, it is merely considered within the cut generation process. The main idea is to derive scenario-adaptable closed-form cut formulas, given AR processes with a specific structure, which allow the adaptation of the cut generated for one specific history  $\bar{\xi}_{[t-1]}$  to different histories  $\xi_{[t-1]}$  of the stochastic process, and thus to different scenarios. In this way, the cuts can be shared between scenarios (see section 5.2) without the need

to incorporate (14.1) into (MSLP) as a constraint. Importantly, these cut formulas lead to exactly the same cuts as the previously described approach.

To illustrate this idea, consider a cut derived using dual problem (14.5) without paying any particular attention to the stagewise dependence. For convenience, but without loss of generality, we assume  $T_{t-1}$  to be deterministic and the RHS uncertainty to be defined by

$$(14.6) \quad \mathbf{h}_t(\boldsymbol{\xi}_t) = \Phi_t \mathbf{h}_{t-1}(\xi_{t-1}) + \boldsymbol{\eta}_t$$

with stagewise-independent error terms  $\boldsymbol{\eta}_t$ , similarly to (14.1). We obtain

$$(14.7) \quad \begin{aligned} \widehat{\mathcal{Q}}_t(x_{t-1}, \xi_{t-1}) &\geq \mathbb{E}_{\boldsymbol{\xi}_t | \xi_{t-1}} [-\boldsymbol{\pi}_t^\top T_{t-1} x_{t-1} + \boldsymbol{\pi}_t^\top \mathbf{h}_t(\boldsymbol{\xi}_t) + \boldsymbol{\mu}_t^\top \mathbf{a}_{t+1}] \\ &= \mathbb{E}_{\boldsymbol{\xi}_t | \xi_{t-1}} [-\boldsymbol{\pi}_t^\top T_{t-1}] x_{t-1} + \mathbb{E}_{\boldsymbol{\xi}_t | \xi_{t-1}} [\boldsymbol{\pi}_t^\top \mathbf{h}_t(\boldsymbol{\xi}_t) + \boldsymbol{\mu}_t^\top \mathbf{a}_{t+1}]. \end{aligned}$$

We can make the following observations:

- (i) Since the probabilities in  $\mathbb{E}_{\boldsymbol{\xi}_t | \xi_{t-1}}[\cdot]$  are assumed to not depend on  $\xi_{t-1}$  (recall that  $\boldsymbol{\eta}_t$  is stagewise-independent) and since all scenarios share the same dual feasible region, the cut gradient

$$(14.8) \quad \beta_t = \mathbb{E}_{\boldsymbol{\xi}_t | \xi_{t-1}} [-\boldsymbol{\pi}_t^\top T_{t-1}]$$

derived for one specific scenario  $\bar{\xi}_{t-1}$  is valid for all other scenarios as well.

- (ii) According to (14.6), the RHS  $\mathbf{h}_t(\boldsymbol{\xi}_t)$  depends on  $\xi_{t-1}$ . Therefore, to evaluate the cut for a specific scenario, this term has to be adapted to that scenario, otherwise the cut may become invalid. By (14.6), this term can be split into a scenario-dependent part depending on  $\xi_{t-1}$  and a scenario-independent part depending on  $\boldsymbol{\eta}_t$  only.
- (iii) The last term  $\mathbf{a}_{t+1}$  in (14.7) is the cut intercept of the following stage. As we face stagewise dependence, this intercept is no longer scenario-independent, but should denote  $\mathbf{a}_{t+1}(\xi_t)$ . Moreover, it is defined recursively: The stage- $t$  intercept includes the stage- $(t+1)$  intercept, which includes the stage- $(t+2)$  intercept, and so on. This implies that to evaluate  $\mathbf{a}_{t+1}(\xi_t)$  for a specific scenario, it is basically required to recursively traverse the whole scenario tree starting from stage  $t$ . This is computationally intractable.

To address these observations, the main idea of Infanger and Morton [116] is to express the cut intercept  $\alpha_t(\xi_{t-1})$  as the sum of a stagewise-independent term  $\alpha_t^{\text{ind}}$  and a stagewise-dependent term  $\alpha_t^{\text{dep}}(\xi_{t-1})$ :

$$(14.9) \quad \alpha_t(\xi_{t-1}) = \alpha_t^{\text{ind}} + \alpha_t^{\text{dep}}(\xi_{t-1}).$$

Let  $\bar{\pi}_t = \mathbb{E}_{\boldsymbol{\eta}_t}[\boldsymbol{\pi}_t]$  and  $\bar{\mu}_{t,t} = \mathbb{E}_{\boldsymbol{\eta}_t}[\boldsymbol{\mu}_t]$  denote the expected values of the dual variables obtained for realizations of  $\boldsymbol{\eta}_t$ . As explained above, these dual values are valid for any history of the stochastic process due to the structure of the dual feasible set. Let  $\bar{\mathcal{P}}_t$  define the  $(|\Gamma_t| \times m_t)$ -matrix containing the values of  $\bar{\pi}_t$  and  $\bar{\mathcal{R}}_t$  the  $(|\Gamma_t| \times |\Gamma_{t-1}|)$ -matrix containing the values of  $\bar{\mu}_t$  for the previously determined cuts. Furthermore, let the matrix  $D_t$  be defined recursively by

$$(14.10) \quad D_t = [\bar{\mathcal{P}}_{t+1} + \bar{\mathcal{R}}_{t+1} D_{t+1}] \Phi_t, \quad D_T = 0.$$

Then, as shown in [116], the stagewise-dependent cut intercept is given by

$$(14.11) \quad \alpha_t^{\text{dep}}(\xi_{t-1}) = [\bar{\pi}_t + \bar{\mu}_t D_t] \Phi_t \mathbf{h}_{t-1}(\xi_{t-1}).$$

**Table 6** *RHS and uncertainty models considered in the literature on SDDP with stagewise dependence to derive scenario-adaptable closed-form cut formulas.*

RHS $h_t(\xi_t)$	Autoregressive model for $\xi_t$				Source
	Model	Type	Lag	Formula	
const.	AR	L	1	$\xi_t = \Phi_t \xi_{t-1} + \eta_t$	[116]
L	AR	L	1	$\xi_t = \Phi_t \xi_{t-1} + \eta_t$	[180]
const.	PAR	L	1	$\xi_t = \varphi_t(\xi_{t-1} - \psi_{t-1}) + \mu_t + \sigma_t \eta_t$	[218]
const.	AR	L	$\geq 1$	$\xi_t = \sum_{k=1}^{t-1} (\Phi_k^t \xi_k + \Psi_k^t \eta_k) + \eta_t$	[116]
L/C*	AR	L	$\geq 1$	$\xi_t = \Phi_t \xi_{[t-1]} + \eta_t$	[91]
L/C*	AR	L	$\geq 1$	$\xi_t = \Phi_t \xi_{[t-1]} + \Psi_t \eta_t + \Theta_t$	[91]
const.	SPAR	L	$\geq 1$	$\xi_{ti} = \sum_{i'} \sum_{k=1}^{t-1} \Phi_{ii'k}^t \xi_{ti'} + \eta_{ti}$	[134]
const.	AR	NL	1	$\xi_t = \Phi_t(f_t(v_{t-1}) + \xi_{t-1}) + \eta_t$	[116]
const.	AR	NL	$\geq 1$	$\xi_t = \sum_{k=1}^{t-1} (\Phi_k^t \xi_k + f_k^t(v_k)) + \eta_t$	[116]
C	AR	C	$\geq 1$	$\xi_t = f_t(\xi_{[t-1]}, \eta_t)$	[91]

L = affine/linear function, C = convex function, NL = general nonlinear function.  
\* only in the case of inequality constraints.

This means that in line 18 of Algorithm 3.1 a cut can be constructed by using formula (14.8) for the gradient and formulas (14.9), (14.10), and (14.11) for the intercept. The stagewise-independent term can be determined either by an additional formula or by subtracting (14.11) from  $\alpha_t(\xi_{t-1})$  [116]. For a cut to be shared with a different scenario at stage  $t - 1$  only requires the adaption of the stagewise-dependent intercept (14.11) to this specific scenario. In other words, a given cut can be corrected to be valid for a different history of the stochastic process. In particular, it is not necessary to add (14.6) as a constraint to the stage- $t$  subproblem or to traverse the whole scenario tree (see Remark 14.1). Instead, only the cut gradient, the stagewise-independent part of the intercept, and the cumulative expected dual vector  $[\bar{\mathcal{P}}_{t+1} + \bar{\mathcal{H}}_{t+1} D_{t+1}] \Phi_t$  have to be stored [116].

Whereas we have so far limited our explanations to a very simple AR process, similar cut formulas can be derived for more complex processes [91, 116, 180, 185]. We give an overview of different cases covered in the literature in Table 6. Some of the process formulas in Table 6 are presented in a simplified form for reasons of clarity, e.g., by omitting standardization and the incorporation of seasonal or periodical effects. For example, this is true for the SPAR processes considered in [134] (also see section 9), where spatial dependencies between locations  $i$  and  $i'$  are taken into account.

Importantly, all processes for which scenario-adaptable closed-form cut formulas can be derived require a specific structure, such as linearity, convexity, or separability. As shown by Guigues [91], a generalization to convex AR processes and more complex structures in the RHS is possible. For instance, the RHS  $h_t$  does not have to be directly described by the stochastic process (constant  $h_t \equiv \xi_t$ ), but may also be defined as some function  $h_t(\cdot)$  of  $\xi_t$ . Moreover, for the affine case, alternative formulas to those provided by Infanger and Morton are presented by Guigues in [91]. The main difference is that only a minimal subset of coefficients is used, due to defining the process  $(\xi_t)_{t \in [T]}$  componentwise and not in vectorial form compared to (14.1) or (14.6). On the other hand, no recursive formula as in (14.10) is provided to compute the cut coefficients. Finally, Guigues shows that for feasibility cuts as well (section 17), scenario-adaptable cut formulas can be derived.

It is important to emphasize that the presented approach only partially mitigates the drawbacks of augmenting the state space. First, the history of the stochastic process has to be stored to compute  $\xi_t$ , even if such a computation is possible



outside the subproblems. Guigues provides a detailed discussion on how state vectors of minimal size can be defined in order to keep the stored information as small as possible [91]. Additionally, due to their dependence on  $\xi_{t-1}$ , or  $\xi_{[t-1]}$  in general, the expected value functions  $\hat{\mathcal{Q}}_t(\cdot, \cdot)$  live in a higher-dimensional space. Therefore, more iterations and cuts may be required to achieve convergence compared to the stagewise-independent case, as discussed in section 4.2.

**14.3. Sensitivity of SDDP with AR Processes.** Let the uncertainty in (MSLP) be modeled by an AR process. Consider the approach of expanding the state, which leads to two types of state variable:  $x_t$  and  $\xi_{[t]}$ . Both contain information on future resource availability (e.g., hydro storage volume and hydro inflow history affecting future inflows), but they differ in several aspects [215]. First, whereas the information provided by the state  $x_{t-1}$  is certain, the information provided by  $\xi_{[t-1]}$  enters an AR model predicting future realizations, which still involves uncertainty. Second, the parameters of this AR model are estimated from data and thus can be subject to estimation errors. Third, in practice it can often be observed that the values in  $(\xi_t)_{t \in [T]}$  show higher variability over short time than the values of  $(x_t)_{t \in [T]}$ . This uncertainty and variability raise the question of how much the solutions obtained in SDDP react to changes in  $\xi_{[t-1]}$ . This can be examined in a sensitivity analysis.

A general approach for sensitivity analysis in SDDP is presented in [104] and applied to an inventory problem with AR demand. The sensitivity with respect to AR model parameters  $\Phi_t$  and  $\gamma_t$  is also discussed.

For a hydrothermal problem, in [215] it is shown that the solutions obtained in SDDP are more sensitive to changes in the initial information state  $\xi_1$  than to changes in the initial resource state  $x_0$ . Based on previous observations, this leads to the unfavorable side effect of expanding the state space for which solutions of SDDP exhibit larger variability. This may have severe consequences in economic applications, such as increasing risk, unpredictability of prices, or distorted investment signals.

To address this issue, Soares, Street, and Valladão [215] present different mitigation heuristics, such as regularizing changes in  $x_t$  over time or using the accurate AR model in the forward pass of SDDP, but predefined unconditional samples in the backward pass in order to avoid the dependence of cuts on  $\xi_{[t-1]}$ . While they report positive computational results, the authors provide no theoretical results on reasonable parameter choice, cut validity, and convergence for their heuristics.

**14.4. SDDP with Markov Chain.** According to Theorem 14.3, a natural extension of SDDP to stagewise-dependent uncertainty by expanding the state space is only possible for linear (or at least convex) AR processes appearing in the RHS of problem (MSLP). In all other cases, the convexity of the expected value functions  $\hat{\mathcal{Q}}_t(\cdot, \cdot)$  is destroyed. Therefore, in such cases, different approaches are required. One such approach is to incorporate a discrete Markov chain into the uncertainty modeling. This approach is quite established in the literature on SDDP and is used in practical applications in various forms.

**14.4.1. Modeling.** Consider a Markov chain with finitely many possible states  $\zeta_{\ell}, \ell = 1, \dots, L$ , with  $L \in \mathbb{N}$ . At each stage  $t \in [T]$ , we denote the current state of the Markov chain as  $\psi_t$  (again, we assume that  $\psi_1$  is deterministic). The transition probabilities between state  $\psi_{t-1} = \zeta_{\ell}$  at stage  $t-1$  and  $\psi_t = \zeta_{\ell'}$  at stage  $t$  are then denoted by  $\omega_{\ell\ell'}$  for  $\ell, \ell' \in \{1, \dots, L\}$ . For simplicity, we assume the Markov chain to be time-homogeneous, such that  $\omega_{\ell\ell'}$  does not depend on  $t$ , even though this is not required.

We now assume that the distribution of random variable  $\xi_t$  at stage  $t \in [T]$  may depend on the state  $\psi_t$  of the Markov chain. In other words, for each possible state  $\zeta_\ell, \ell = 1, \dots, L$ , the distribution of  $\xi_t$  may differ. We emphasize this by writing  $\xi_t^\ell$ .

The value functions  $Q_t(\cdot, \cdot)$  for (MSLP) then not only depend on  $x_{t-1}$  and the realization  $\xi_t$  of  $\xi_t$ , but also on the current Markov state  $\psi_t$ . As this state can only take finitely many values, we indicate this by  $Q_{t\ell}(x_{t-1}, \xi_t)$ , where index  $\ell$  indicates conditioning on  $\psi_t = \zeta_\ell$ . Based on this definition, the expected value functions can be expressed as

(14.12) 
$$\mathcal{Q}_{t\ell}(x_{t-1}) := \sum_{\ell'=1}^L \omega_{\ell\ell'} \mathbb{E}_{\xi_t|\ell'} \left[ Q_{t\ell'}(x_{t-1}, \xi_t^{\ell'}) \right].$$


The index  $\ell$  of the expected value function refers to the previous Markov state  $\psi_{t-1} = \zeta_\ell$ . Compared to standard SDDP, the expectation is not only taken over the realizations of  $\xi_t^{\ell'}$ , but the state transitions from  $\psi_{t-1}$  to  $\psi_t$  are also taken into account. Using this definition, the DPE for stages  $t = 2, \dots, T$  can be written as

(14.13) 
$$Q_{t\ell}(x_{t-1}, \xi_t^\ell) := \begin{cases} \min_{x_t} & \zeta_\ell^\top x_t + \mathcal{Q}_{t+1,\ell}(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(\xi_t^\ell). \end{cases}$$

Note that the dependence on  $\zeta_\ell$  in (14.12) resembles the expanding-the-state approach from section 14.1. However, there are important differences.  $\psi_{t-1}$  does not necessarily enter the subproblems, but can be an underlying (hidden) Markov state. Moreover, it can only take a finite number of different values, whereas  $\xi_{[t-1]}$ , even if discrete, is treated like a continuous state variable when expanding the state space. Furthermore, as the transition probabilities  $\omega_{\ell\ell'}$  may differ for each  $\zeta_\ell$ , the cut components are weighted differently and cuts cannot be shared between different Markov states. Consequently, it is necessary to store separate expected value functions  $\mathcal{Q}_{t\ell}(\cdot)$  for each  $\ell = 1, \dots, L$ . In return, the nonconvexity of these functions is circumvented, since each  $\mathcal{Q}_{t\ell}(\cdot)$  remains convex and is approximated on its own; see also the discussion in section 14.4.

As an example, consider a problem borrowed from [167] with  $L = 2$  Markov states and  $q^\ell = 2$  realizations for  $\xi_t^\ell$  for each of them. The corresponding scenario tree with underlying Markov chain is illustrated in Figure 13. For the transition probabilities let  $\omega_{11} = q, \omega_{12} = 1 - q, \omega_{21} = 1 - p$ , and  $\omega_{22} = p$ . For all  $t$  and  $\ell \in \{1, 2\}$ , the distribution of  $\xi_t^\ell$  is given by  $p_{tj} = \frac{1}{2}$  for  $j \in \{1, 2\}$ .

As an alternative to the scenario tree in Figure 13, the stochastic process with underlying Markov chain can be represented by a Markovian policy graph with finitely



**Fig. 13** Scenario tree with underlying Markov chain (state 1 printed in black, state 2 printed in white). Replicated from [167].

many nodes per stage [60]. This approach is also included in the software package SDDP.jl; see section 10.

**14.4.2. Adaptation of SDDP.** Let us now address how SDDP works in this case. Different approaches for the forward pass are used in the literature. The most natural one is, for each stage  $t$  and each sample path  $k \in \mathcal{K}$ , to sample first from the Markov states and then conditionally from  $\xi_t^\ell$  in line 6 of Algorithm 3.1 [168]. Sometimes it is also proposed to use historical values here, e.g., true inflow spot-price combinations [86]. In such a case, it is possible that a spot price is drawn which is not a valid state of the Markov chain. Then, one strategy is to use the (in some sense) closest state from the Markov chain [86]. Another strategy is to use a linear interpolation between the hyperplanes of neighboring states [88, 237].

For stages  $t = 2, \dots, T$ , states  $\ell = 1, \dots, L$ , and samples  $k \in \mathcal{K}$ , based on (14.13), the approximate subproblems solved in the forward pass of SDDP have the form

$$(14.14) \quad \underline{Q}_{t\ell}^i(x_{t-1}^{ik}, \xi_t^{\ell k}) := \begin{cases} \min_{x_t} & (c_t(\xi_t^\ell))^\top x_t + \underline{V}_{t+1,\ell}^i(x_t) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}^{ik}, \xi_t^\ell). \end{cases}$$

Importantly, each function  $\underline{Q}_{t\ell}(\cdot)$ ,  $\ell = 1, \dots, L$ , is approximated by an individual cut approximation  $\underline{V}_{t\ell}(\cdot)$ .

In the backward pass of some iteration  $i$ , the stages are traversed in a backward direction as usual to improve the cut approximations. At each stage  $t$ , the subproblems (14.14) including  $\underline{V}_{t\ell}^{i+1}(\cdot)$  are solved for each trial state  $x_{t-1}^{ik}$ ,  $k \in \mathcal{K}$ , each stage- $t$  Markov state  $\psi_t = \zeta_\ell$ ,  $\ell = 1, \dots, L$ , and all realizations  $\xi_{tj}^\ell$ ,  $j = 1, \dots, q_t^\ell$ . This means that in lines 13–21 of Algorithm 3.1 the iterations are not over  $t$ ,  $k$ , and  $j$ , but over  $t$ ,  $k$ ,  $\ell$ , and  $j(\ell)$ .

Then, for each  $x_{t-1}^{ik}$  and  $\psi_{t-1} = \zeta_\ell$ ,  $\ell = 1, \dots, L$ , a valid cut can be derived for  $\underline{Q}_{t\ell}(\cdot)$ . Let  $\beta_{t\ell k j}^i$  denote a subgradient for  $\underline{Q}_{t\ell}^i(\cdot, \cdot)$  at  $x_{t-1}^{ik}$ . In accordance with (3.4), but also taking into account the Markov chain transition probabilities, we can then define cut coefficients

$$(14.15) \quad \begin{aligned} \beta_{t\ell k}^i &:= \sum_{\ell'=1}^L \omega_{\ell\ell'} \left( \sum_{j=1}^{q_{t\ell}} p_{t\ell j} \left( \underline{Q}_{t\ell}^{i+1}(x_{t-1}^{ik}, \xi_{tj}^{\ell'}) - (\beta_{t\ell k j}^i)^\top x_{t-1}^{ik} \right) \right), \\ \alpha_{t\ell} &:= \sum_{\ell'=1}^L \omega_{\ell\ell'} \left( \sum_{j=1}^{q_{t\ell}} p_{t\ell j} \beta_{t\ell k j}^i \right), \end{aligned}$$

where  $q_{t\ell}$  and  $p_{t\ell j}$  denote the numbers of realizations and probabilities of  $\xi_t^\ell$ .

A cut (3.5) for  $\underline{Q}_{t\ell}(\cdot)$  is then given by function

$$\phi_{t\ell k}^i(x_{t-1}) := \alpha_{t\ell}^i + (\beta_{t\ell k}^i)^\top x_{t-1}$$

and can be used to update  $\underline{V}_{t\ell}^i(\cdot)$ . Philpott, de Matos, and Finardi derive similar formulas for the multicut and risk-averse case [168].

The approach described allows for the incorporation of even nonlinear stagewise-dependent uncertainty into SDDP, but it also gives rise to some challenges. Among those challenges is the assumption of the Markov property, which may not always be appropriate. Moreover, it is necessary to define useful values  $\zeta_\ell$ ,  $\ell = 1, \dots, L$ , and transition probabilities  $\omega_{\ell\ell'}$  for the Markov states [88, 150]. Most importantly, cuts cannot be shared between Markov states, but only within them, so that separate expected value functions have to be considered for each  $\ell = 1, \dots, L$ . Therefore, the number of Markov states should be rather small to preserve computational tractability.

**14.4.3. Specific Use Cases.** There exist different use cases for modeling the uncertainty in (MSLP) with an integrated Markov chain.

- **Nonlinear Processes.** The data process  $(\xi_t)_{t \in [T]}$  can be modeled as a nonlinear AR process or a nonlinear transformation of a linear AR process (see section 9), which, if handled by expanding the state space, destroys the convexity of  $\hat{\mathcal{Q}}_t(\cdot, \cdot)$ . Sometimes such a nonlinear process can be approximated by assuming that the realizations  $\xi_t$  depend on an underlying system state which follows a Markov process [168], thus does not capture the nonlinearity explicitly in a formula. As the value functions are also not convex in this possibly continuous Markov state, the Markov process is approximated using a discrete Markov chain.
- **Regime-Switching Models.** Instead of a single AR process, sometimes the data process  $(\xi_t)_{t \in [T]}$  can be best modeled by a finite set of different AR processes, which are valid representations, and thus active, under different circumstances (e.g., macroeconomic, political, or ecological situations). A discrete Markov chain can then be used to model these overall system states, and AR models can be used to describe realizations of the uncertain data conditioned on these states. Such regime-switching models are very common in wind forecasting [242].
- **Hybrid SDP/SDDP.** Different parts of the data in (MSLP) exhibit stagewise-dependent uncertainty. While some of them, namely, uncertainty in the RHS  $h_t$ , can be treated by expanding the state space, for others, e.g., stagewise-dependent uncertainty in the objective coefficients  $c_t$ , that would destroy the convexity of  $\hat{\mathcal{Q}}_t(\cdot, \cdot)$ . Therefore, this part of the uncertainty may be modeled by a discrete Markov chain instead. Since one part of the uncertainty is treated as in standard SDDP (allowing for cut-sharing between scenarios), while another is treated by enumerating separate expected value functions for each  $\ell = 1, \dots, L$  (cuts cannot be shared between Markov states), this is often referred to as a *hybrid SDP/SDDP* method [86].

For instance, this setting often occurs in medium-term hydrothermal scheduling problems (see section 9) when inflow uncertainty in the RHS as well as spot-price uncertainty in the objective function are taken into account. The idea of addressing this using a Markov chain goes back to Gjelsvik et al., who modeled this kind of scheduling problem for the Norwegian power system [86, 88, 89]. Since then, this approach has been employed in several applications, for example, hydrothermal scheduling including balancing market bids [107, 108], risk management [115, 123, 149], and fuel contracts [39]. It is also applied to the modeling of fuel price uncertainty [158].

In contrast to the general approach presented above, in this case it is usually assumed that the uncertainty in the RHS and in the objective are independent of one another. Therefore, for each state  $\zeta_\ell, \ell = 1, \dots, L$ , the distribution of  $\xi_t$  is the same and marginal distributions can be used in the expectation in (14.12). Moreover, note that in this specific case the Markov chain states do not underlie the distribution of  $\xi_t$ , but instead they enter the subproblems explicitly, e.g., as objective coefficients. However, SDDP can still be applied using the same ideas as above.

- **Markov Chain SDDP.** Assume that the data process  $(\xi_t)_{t \in [T]}$  itself is Markovian, i.e., as in (14.1),  $\xi_t$  only depends on  $\xi_{t-1}$  for all  $t = 2, \dots, T$  instead of on the whole history  $\xi_{[t-1]}$ . In this case, the data process can be

represented, or at least approximated (if the random variables  $\xi_t$  are continuous), by a discrete Markov chain. This approximation can be obtained by lattice quantization techniques [31, 136]. As it contains only finitely many states per stage  $t = 2, \dots, T$ , this Markov chain can be illustrated as a recombining scenario tree or *scenario lattice* [136], just as in the case of the stagewise-independence Assumption 2; see section 2. The difference is that in the Markov chain case the probabilities of transitions to stage- $t$  nodes may differ between different stage- $(t-1)$  nodes. This also includes the possibility that some stage- $t$  nodes may not be reached from certain stage- $(t-1)$  nodes. Due to this difference, the (expected) value functions  $\mathcal{Q}_{t\ell}(\cdot)$  depend on the states  $\zeta_\ell, \ell = 1, \dots, L$ , of the Markov chain, and in SDDP cuts are derived for each of these functions separately. This idea is called Markov chain SDDP (MC-SDDP) [136] or *approximate dual dynamic programming* (ADDP) [137, 138], whereas for distinction the approach of expanding the state space is referred to as time-series SDDP (TS-SDDP).

MC-SDDP can be considered a special case of the above framework, where the Markov states explicitly enter the subproblems and completely describe the uncertainty, instead of affecting another random variable. More precisely, in (14.12) the conditional expectations and  $\xi_t^{\ell'}$  can be omitted, and only probabilities  $\omega_{\ell\ell'}$  are required. Analogously, in the cut intercept and gradient formulas (14.15), the summation over  $j$  and the probabilities  $p_{t\ell j}$  can be omitted.

For problems with moderate state space dimension to MC-SDDP, expanding the state may be computationally favorable as only one expected value function has to be approximated per stage. On the other hand, a computational advantage of MC-SDDP is that the computational effort grows linearly with the number of Markov states only [208]. In contrast, expanding the state leads to a state space dimension increase in which the complexity of SDDP grows exponentially. Moreover, MC-SDDP requires no linearity and is not limited to stagewise-dependent uncertainty, only appearing in the RHS of (MSLP). As long as the Markov property is satisfied, it allows for stagewise-dependent uncertainty in all data  $c_t, T_{t-1}, W_t$ , and  $h_t$  of (MSLP).

The main drawback of MC-SDDP lies in its relationship to the true problem  $(\tilde{P})$  in the case of a continuous data process  $(\xi_t)_{t \in [T]}$ ; see also section 11. For SDDP with AR processes, and expanding the state space, many results exist that allow for inference of the SAA solution with respect to the true problem; see section 11. One key property in this regard is that  $\xi_{t-1}$  is treated as a possibly continuous state variable in SDDP, such that the derived cuts are also valid at states which are not reached by the scenarios  $\xi^s \in \mathcal{S}$  that are considered in SDDP. Similar results are not available for MC-SDDP. In particular, the policy and lower bounds obtained are not necessarily valid for the true problem [136].

In spite of this theoretical downside, Löhdorf and Shapiro [136] report tighter lower bounds and better policies even for the true process based on computational experiments. They conjecture that this is due to a differing exploration of the state space. Expanding the state space introduces additional state variables which are not under control of the optimal policy (their trajectory is not chosen based on solving the approximate subproblems in the forward pass, but is selected randomly in the forward pass). This may lead

to the selection of states that do not provide the highest information gain. With MC-SDDP this is partially mitigated by choosing sufficiently different states in advance when constructing the Markov chain.

**14.5. Hybrid NBD/SDDP.** In the previous section, we presented a hybrid SDP/SDDP method as a tool to model different stagewise-dependent uncertain data in (MSLP) using different approaches. Instead of modeling the “complicating” part of the uncertainty by a discrete Markov chain, a scenario tree can also be used. Instead of a hybrid SDP/SDDP method, this yields a hybrid NBD/SDDP method [185]; see also section 5.2.

Assume that the random vector  $\xi_t$  modeling the uncertainty in  $c_t, W_t, T_{t-1}$ , and  $h_t$  can be separated into two separate and independent parts,  $\xi_t^S$  and  $\xi_t^T$ . The first vector  $\xi_t^S$  can either be stagewise-independent or exhibit some linear dependency if it occurs in the RHS. In the latter case, it can be handled by expanding the state space. Within SDDP, in each iteration samples of  $\xi_t^S$  are considered. The second vector  $\xi_t^T$ , on the other hand, may lead to nonconvexities in the value functions if it is approached by expanding the state space. Therefore, it is modeled by a scenario tree, which is treated exactly in SDDP. This means that for this particular part of the uncertainty, no samples are drawn, but all scenarios are considered in each iteration of SDDP, as in NBD; see section 5.2. This approach is similar to hybrid SDP/SDDP in the sense that the expected value functions  $\mathcal{Q}_t(\cdot)$  depend on the scenarios from  $\xi_t^S$  and that cuts can only be shared within, and not between, such scenarios.

By only treating the crucial part  $\xi_t^T$  of  $\xi$  as a scenario tree and the remainder  $\xi^S$  by sampling still, complex uncertainty processes can be considered, while at the same time the increase in the computational complexity is kept as small as possible [185]. To take advantage of this, the scenario tree associated with  $\xi^S$  should not be too large.

Compared to hybrid SDP/SDDP, in specific applications one or the other approach may be favorable. For instance, the Markov chain approaches allow for dependencies between different uncertainty processes. Moreover, if each realization of  $\xi_t$  is assigned to one specific Markov state  $\zeta_\ell, \ell = 1, \dots, L$ , the number of LPs to be solved per iteration can be kept the same as in standard SDDP. The scenario tree approach, by contrast, requires independence of  $\xi^S$  and  $\xi^T$ . By design, it considers all combinations of scenarios of  $\xi^T$  and  $\xi^S$ , so no assignment of realizations of  $\xi^S$  to scenarios of  $\xi^T$  is required. However, the number of LPs to be solved grows exponentially in the number of stages [185]. On the other hand, a scenario tree may be more appropriate to model very complex processes, e.g., macroeconomical, political, or structural decisions [185], for which the Markov property is not appropriate.

**14.6. Saddle Cuts.** We consider the special case of stagewise-dependent objective coefficients  $c_t(\xi_t)$  in (MSLP), as they appear for uncertain price models by AR processes. Up to this point, we have introduced SDDP with the integrated Markov chain as a suitable solution approach in this case. Next, we discuss a second approach.

As discussed in section 14.1, by expanding the state space, stagewise-independence (Assumption 2) can be recovered, but in return the expected value functions  $\mathcal{Q}_t(\cdot, \cdot)$  are no longer convex. In Theorem 14.3 it is shown that  $\hat{\mathcal{Q}}_t(\cdot, \cdot)$  is in fact convex in  $x_{t-1}$ , but concave in  $c_{t-1}$ , which yields a saddle shape. Therefore, linear cuts are not sufficient to approximate them. As a solution, which exploits the saddle shape, special saddle cuts can be used.

To derive this formally, in the vein of [59], we assume the objective coefficients are described by  $(y_t(\xi_t))^\top C_t$  instead of  $c_t(\xi_t)$ . While the matrix  $C_t$  is considered deterministic,  $y_t(\xi_t)$  is defined by the AR process

$$(14.16) \quad y_t(\xi_t) = B_t(\xi_t)y_{t-1}(\xi_{t-1}) + b_t(\xi_t)$$

for all stages  $t = 2, \dots, T$ . Here, the matrix  $B_t$  and the vector  $b_t$  are uncertain and depend on the realization of  $\xi_t$ . Thus, the sequence  $(y_t(\xi_t))_{t=1}^T$  is scenario-dependent.

Inserting relation (14.16) into the objective function and considering  $y_{t-1}$  as an additional state variable, for  $t = 2, \dots, T$ , we obtain the subproblems

$$\begin{aligned} & \hat{Q}_t(x_{t-1}, y_{t-1}, \xi_t) \\ &= \begin{cases} \min_{x_t} & (B_t(\xi_t)y_{t-1} + b_t(\xi_t))^\top C_t x_t + \hat{\mathcal{Q}}_{t+1}(x_t, B_t(\xi_t)y_{t-1} + b_t(\xi_t)) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}, \xi_t), \end{cases} \end{aligned}$$

where

$$\hat{\mathcal{Q}}_{t+1}(x_t, y_t) = \mathbb{E}_{\xi_{t+1}} [\hat{Q}_{t+1}(x_t, y_t, \xi_{t+1})]$$

and  $\hat{\mathcal{Q}}_{T+1}(x_T, Y_T) \equiv 0$ . For the first stage, we obtain

$$v^* = \begin{cases} \min_{x_1} & b_1 C_1 x_1 + \hat{\mathcal{Q}}_2(x_1, y_1) \\ \text{s.t.} & x_1 \in \mathcal{X}_1. \end{cases}$$

The additional state  $y_{t-1}$  is referred to as an *objective state*. This state is not allowed to appear in the constraints [59]. As stated above,  $\hat{\mathcal{Q}}_t(\cdot, \cdot)$  is piecewise linear and convex in  $x_{t-1}$ , but piecewise linear concave in  $y_{t-1}$  and, as such, a piecewise bilinear saddle function.

The concept of approximating saddle functions with saddle cuts goes back to Baucke, Downward, and Zakeri, who propose a deterministic algorithm to solve stochastic minimax dynamic programs [11]. A related approach is used in *robust dual dynamic programming* (RDDP), which uses an SDDP-like framework to solve multi-stage robust programs [83]. The main idea is to compute lower and upper bounding saddle functions, which combine the ideas of an outer approximation by cutting planes and an inner approximation by convex combinations of function values, the latter of which we discuss thoroughly in section 8. For stagewise-dependent objective coefficients, it is sufficient to use only the lower bounding saddle functions, so-called saddle cuts, from [11] to approximate the expected value functions in SDDP.

Let (3.4) define  $\beta_t$  and  $\alpha_t$  as in standard SDDP. Then, in line 18 of Algorithm 3.1, the  $r$ th saddle cut for  $\hat{\mathcal{Q}}_{t+1}(\cdot, \cdot)$  is defined as the solution to the optimization problem

$$(14.17) \quad \begin{aligned} & \min_{\gamma_t, \theta_{t+1}} && y_t^\top \gamma_t + \theta_{t+1} \\ & \text{s.t.} && (y_t^r)^\top \gamma_t + \theta_{t+1} \geq \alpha_{t+1}^r + (\beta_{t+1}^r)^\top x_t, \\ & && \|\gamma_t\|_\infty \leq \nu, \end{aligned}$$

where  $y_t^r = y_t^{ik}$  denotes the current objective state in iteration  $i$  and for scenario  $k \in \mathcal{K}$ . Importantly, this problem has  $x_t$  and  $y_t$  as parameters. Hence, a saddle cut gives a valid lower approximation for  $\hat{\mathcal{Q}}_{t+1}(\cdot, \cdot)$  for all  $x_t$  and  $y_t$  and can be shared between scenarios. Moreover, the saddle cuts are tight at the trial state given by  $x_t^{ik}$  and  $y_t^{ik}$ , at which they are created.

A crucial part of applying this approach is to bound the decision variable  $\gamma_t$  in (14.17) by an appropriate constant  $\nu$ . To this end, the expected value functions  $\widehat{\mathcal{Q}}_t(\cdot, \cdot)$  are required to be Lipschitz continuous with respect to  $y_{t-1}$ . As shown in [11], to ensure validity of the saddle cuts, the parameter  $\nu$  has to be chosen at least as large as the Lipschitz constant of  $\widehat{\mathcal{Q}}_t(\cdot, \cdot)$  with respect to  $y_{t-1}$  under the dual norm  $\|\cdot\|_1$  of  $\|\cdot\|_\infty$ . If it is chosen smaller, the result might be invalid cuts and suboptimal solutions. If it is chosen too large, the cuts might become very weak [59].

Incorporating the saddle cuts, for each stage  $t = 2, \dots, T$ , iteration  $i$ , and scenario  $k \in \mathcal{K}$ , the SDDP subproblems can be formulated as

$$\widehat{\mathcal{Q}}_t^i(x_{t-1}^{ik}, y_{t-1}^{ik}, \xi_{tj}) = \begin{cases} \min_{x_t, \gamma_t, \theta_{t+1}} & (y_t^{ik})^\top C_t x_t + (y_t^{ik})^\top \gamma_t + \theta_{t+1} \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}^{ik}, \xi_{tj}), \\ & (y_t^r)^\top \gamma_t + \theta_{t+1} - (\beta_{t+1}^r)^\top x_t \geq \alpha_{t+1}^r, \quad r \in \Gamma_{t+1}, \\ & \|\gamma_t\|_\infty \leq \nu, \end{cases}$$

where  $y_t^{ik} = B_t(\xi_t^k) y_{t-1}^{ik} + b_t(\xi_t^k)$ . Aside from considering the additional objective state  $y_t^{ik}$ , the forward pass remains of SDDP remains unchanged.

It can be shown that only finitely many different saddle cuts can be constructed. As a consequence, the convergence results are the same as for standard SDDP [59].

**14.7. Applying Dual SDDP.** A third approach tailored to stagewise-dependent objective coefficients  $c_t(\xi_t)$  in (MSLP) is to apply dual SDDP [104], as presented in section 8. Recall the value functions derived from the dual problem of (MSLP):

$$(14.18) \quad \widetilde{D}_t(\pi_{t-1}) := \begin{cases} \max_{\pi_t} & \sum_{j=1}^{q_t} p_{tj} \left( -h_{tj}^\top \pi_{tj} + \widetilde{D}_{t+1}(\pi_{tj}) \right) \\ \text{s.t.} & \sum_{j=1}^{q_t} p_{tj} \left( T_{t-1,j}^\top \pi_{tj} \right) + W_{t-1}^\top \pi_{t-1} \leq c_{t-1}. \end{cases}$$

These value functions are concave in  $\pi_{t-1}$ . Crucially, here the objective coefficients  $c_{t-1}$  appear in the RHS. If  $(c_t)_{t \in [T]}$  is described as a linear AR process, we can expand the state space as for the primal subproblems in section 14.1, and the new state variable  $c_{[t-2]}$  appears in the RHS. Therefore, the value functions obtained are also concave in  $c_{[t-2]}$  and can be approximated from above by linear cuts. This can be done by applying dual SDDP [104]; see section 8.

**14.8. Conditional Cuts.** The previously discussed approaches all have in common the fact that they require us to expand the state space or to set up a scenario tree or a discrete Markov chain from the true (continuous) data process (or from existing historical data). van Ackooij and Warin propose an alternative approach that works without these requirements [234]. Their approach is based on established methods in mathematical finance and optimal stopping theory. A crucial assumption is that the data process  $(\xi_t)_{t \in [T]}$  is Markovian.

Assume that a finite set  $\mathcal{S}$  of scenarios  $\xi^s, s \in \mathcal{S}$ , is given, e.g., historical observations of the data. This set is chosen in advance and is not changed within SDDP. The first key ingredient of the proposed variant of SDDP is to partition the set of possible values of  $\xi_t$  for each stage  $t \in [T]$  into a finite number  $|L_t|$  of hypercubes



$D_{t\ell}$ ,  $\ell = 1, \dots, |L_t|$ , also called meshes. This partitioning is done in such a way that approximately a uniform distribution of the samples is achieved [234].

As we explain in what follows, the main idea is now to compute cuts conditioned on specific meshes, i.e., for each mesh a different set of cuts is considered.

In the forward pass of SDDP, a subset  $L_t \subseteq \mathcal{S}_t$  of scenarios is sampled for each stage (see line 6 of Algorithm 3.1). This is done with the aim of obtaining a trial solution  $x_t^\ell$  for each mesh in expectation for all  $t = 2, \dots, T$ . Each of these trial solutions is then used in the backward pass to derive cuts.

In the backward pass, for any sequence  $(x_t^{i\ell})_{t \in [T]}$  of trial solutions, let  $(d(t)^{i\ell})_{t \in [T]}$  denote the sequence of corresponding meshes, i.e.,  $x_t^{i\ell}$  has been determined in the forward pass for  $\xi_t^\ell \in D_{t,d(t)^{i\ell}}$ . At each stage  $t = T, \dots, 2$ , the SDDP subproblems are now solved for all scenarios  $\xi_t^s$  for which  $\xi_{t-1}^s \in D_{t-1,d(t-1)^{i\ell}}$ . This means that for each trial solution, all scenarios are considered that share the same mesh with the scenario used to obtain the trial solution. Line 15 in Algorithm 3.1 is adapted accordingly.

After solving these subproblems, the obtained solutions are used to construct cuts. In contrast to standard SDDP, however, the cut coefficients are determined as estimates of the corresponding conditional expectations [234]:

$$\alpha_{t\ell}^i(\xi_{t-1}) = \hat{\mathbb{E}}_{|\xi_{t-1}}^S \left[ (\boldsymbol{\pi}_t^{i\ell s})^\top \mathbf{h}_t(\xi_t) + \sum_{r \in \Gamma_{t+1}} \mu_t^{i\ell sr} \alpha_{t+1}^r \right]$$

and

$$\beta_{t\ell}^i(\xi_{t-1}) = -\hat{\mathbb{E}}_{|\xi_{t-1}}^S \left[ (\boldsymbol{\pi}_t^{i\ell s})^\top T_{t-1} \right].$$

These estimates are computed by linearly regressing the terms for each considered scenario  $\xi_t^s$  on a finite number of local base functions, e.g., monomials in  $\mathbb{R}^{p_t}$ , with support on the considered mesh. Importantly, they are zero outside this mesh. The idea is that in this manner a cut of form

$$(14.19) \quad \mathcal{Q}_t(x_{t-1}, \xi_{t-1}) \geq \phi_{t\ell}^i(x_{t-1}, \xi_{t-1}) = (\beta_{t\ell}^i(\xi_{t-1}))^\top x_{t-1} + \alpha_{t\ell}^i(\xi_{t-1})$$

can be constructed, which provides a local update of the cut approximation in the current mesh  $D_{t-1,d(t-1)}$  and is zero, otherwise. Hence, the cut is associated with this specific mesh and stored in a corresponding index set. In following iterations of SDDP, for each subproblem only the set of cuts that is associated with the currently explored mesh [234] is then taken into account. Therefore, these cuts are called *conditional cuts*.

The main drawback of this approach is that the cuts described are not guaranteed to be valid underestimators, so the inequality in (14.19) is not guaranteed to be satisfied, because their formula relies on *estimators* that may deviate from the true conditional expectations.

However, for problems with a low-dimensional vector  $\boldsymbol{\xi}_t$  and Markovian dependency, the policies obtained using conditional cuts are reported to be competitive with those obtained by expanding the state space, but without the increase in the state dimension and without the need to set up a scenario tree [234].

**15. Extension to Convex Programs [Relaxing Assumption 6].** A natural extension of SDDP can be achieved by relaxing the assumption of linearity, i.e., Assumption 6, but assuming a multistage stochastic convex problem (MSCP). In the same vein as problem (2.3), this problem can be formulated in the general form

$$(15.1) \quad v_C^* := \begin{cases} \min_{x_1, x_2, \dots, x_T} & \mathbb{E} \left[ \sum_{t \in [T]} f_t(x_t(\xi_{[t]}), \xi_t) \right] \\ \text{s.t.} & g_1(x_1) \leq 0, \\ & g_t(x_{t-1}(\xi_{[t-1]}), x_t(\xi_{[t]}), \xi_t) \leq 0 \quad \forall \xi_{[t]} \quad \forall t = 2, \dots, T, \\ & x_t \in X_t \quad \forall t \in [T], \\ & x_t(\cdot) \text{ } \mathcal{F}_t\text{-measurable} \quad \forall t \in [T], \end{cases}$$

with  $f_t(\cdot)$  and  $g_t(\cdot, \cdot)$  some  $\mathcal{F}_t$ -measurable functions with respect to  $\xi$ .

We make the following assumptions [85, 92].

*Assumption 10.* For fixed  $\xi_t \in \Xi_t$ , let  $f_t(\cdot, \xi_t)$  and  $g_t(\cdot, \cdot, \xi_t)$  (componentwise) be proper, convex, lower semicontinuous, and differentiable functions and  $X_t$  be nonempty convex compact sets for all  $t \in [T]$ .

Under stagewise independence (Assumption 2), finite randomness (Assumption 5), and Assumption 10, the (MSCP) in (15.1) can be expressed using its DPE in the following form for  $t = 2, \dots, T$ :

$$(15.2) \quad Q_{t,C}(x_{t-1}, \xi_t) := \begin{cases} \min_{x_t} & f_t(x_t, \xi_t) + \mathcal{Q}_{t+1,C}(x_t) \\ \text{s.t.} & g_t(x_{t-1}, x_t, \xi_t) \leq 0, \\ & x_t \in X_t, \end{cases}$$

with expected value functions defined as usual by

$$(15.3) \quad \mathcal{Q}_{t+1,C}(x_t) := \mathbb{E}_{\xi_{t+1}} [Q_{t+1,C}(x_t, \xi_{t+1})]$$

and  $\mathcal{Q}_{T+1,C}(x_T) \equiv 0$ . For the first stage, this yields

$$(15.4) \quad v_C^* = \begin{cases} \min_{x_1} & f_1(x_1) + \mathcal{Q}_{2,C}(x_1) \\ \text{s.t.} & g_1(x_1) = 0, \\ & x_1 \in X_1. \end{cases}$$

Applying SDDP to (MSCP) with convergence guarantees requires a more strict recourse assumption compared to Assumption 9.

*Assumption 11* (extended relatively complete recourse [85]). Let  $\text{aff}(\mathcal{X}_t)$  be the affine hull of the reachable set  $\mathcal{X}_t$  and  $B_t(\delta_t) = \{y \in \text{aff}(\mathcal{X}_t) \mid \|y\| < \delta_t\}$  for some  $\delta_t > 0$  and some norm  $\|\cdot\|$ .

For all  $t \in t = 2, \dots, T$ , all  $x_{t-1} \in \mathcal{X}_{t-1} + B_t(\delta_t)$ , and all  $\xi_{tj}, j = 1, \dots, q_t$ , the feasible set of subproblems (15.2) is nonempty.

Intuitively, Assumption 11 demands that feasibility of the subproblems is also ensured for  $x_{t-1}$  slightly outside  $\mathcal{X}_t$ . This is required in order to guarantee Lipschitz continuity of all value functions  $Q_{t,C}(\cdot, \cdot)$  and expected value functions  $\mathcal{Q}_{t,C}(\cdot)$  [85]. In addition, all value functions are convex and thus can be approximated by linear cuts. Such cuts can be generated using Lagrangian duality. More precisely, for all  $t = 2, \dots, T$ ,  $x_{t-1} \in \mathcal{X}_{t-1}$ , and some multipliers  $\pi_t \in \mathbb{R}^{m_t}$  (with  $m_t$  the dimension of  $g_t(\cdot, \cdot)$ ), we introduce the Lagrangian function

$$(15.5) \quad L_{t,C}(\pi_t; x_{t-1}, x_t, \xi_t) = f_t(x_t, \xi_t) + \pi_t^\top g_t(x_{t-1}, x_t, \xi_t),$$

the corresponding dual function

$$(15.6) \quad \mathcal{L}_{t,C}(\pi_t; x_{t-1}, \xi_t) = \min_{x_t \in X_t} L_t(\pi_t; x_{t-1}, x_t, \xi_t),$$

and the corresponding Lagrangian dual problem

$$(15.7) \quad \max_{\pi_t \geq 0} \mathcal{L}_t(\pi_t; x_{t-1}, \xi_t).$$

Further, we make the following assumption which ensures no duality gap between the primal subproblems (15.2) and their dual problems (15.7) [92]. Here,  $\text{ri}(S)$  denotes the relative interior of some set  $S$ .

*Assumption 12* (Slater condition [92]). For all  $x_{t-1} \in \mathcal{X}_{t-1}$  and all  $\xi_{tj}, j = 1, \dots, q_t$ , there exists  $x_t \in \text{ri}(X_t)$  such that  $g_t(x_{t-1}, x_t, \xi_{tj}) < 0$ .

Then, exploiting differentiability, a subgradient of  $\mathcal{Q}_{t,C}(\cdot)$  at  $\bar{x}_{t-1}$  is given by

$$\bar{\beta}_t = \partial \mathcal{Q}_t(\bar{x}_{t-1}) = \sum_{j=1}^{q_t} p_{tj} \nabla_{x_{t-1}} L_{t,C}(\bar{\pi}_{tj}; \bar{x}_{t-1}, \bar{x}_{tj}, \xi_{tj}),$$

where  $\bar{x}_{tj}$  is an optimal solution to the primal problem (15.2) and  $\bar{\pi}_{tj}$  is an optimal solution to the dual problem (15.7) given  $\xi_{tj}$ . Moreover,  $\nabla_x h(\cdot)$  denotes the gradient of some function  $h(\cdot)$  with respect to  $x$ . Using this subgradient, a cut for  $\mathcal{Q}_t(\cdot)$  is given by [92]

$$(15.8) \quad \mathcal{Q}_t(x_{t-1}) \geq \mathcal{Q}_t(\bar{x}_{t-1}) + \bar{\beta}_t^\top (x_{t-1} - \bar{x}_{t-1}).$$

Under Assumption 11, the norm of the obtained subgradients can be shown to be bounded [92].

In summary, to adapt SDDP to the convex setting, in line 10 of Algorithm 3.1 the convex subproblems (15.2) are solved, in line 16 the dual problems (15.7) are solved, and in line 18 cuts (15.8) are constructed.

The cut derivation can be generalized to DPE including  $\mathcal{V}_t(\cdot)$  instead of  $\mathcal{Q}_t(\cdot)$ . The results can also be generalized to cost functions  $f_t(x_{t-1}, x_t, \xi_t)$  depending on the state  $x_{t-1}$ ; see [92] for details.

Contrary to the linear case, however, the expected value functions  $\mathcal{Q}_{t,C}(\cdot)$  are no longer polyhedral. As a consequence, they cannot be represented exactly by a finite number of cuts. However, it can be shown that given the above assumptions and Assumptions 1–8, almost sure asymptotic convergence of SDDP is ensured. In [85] this is proven for the case that  $x_{t-1}$  only enters the subproblems (15.2) in linear constraints, that is, when  $g_t(\cdot)$  is a linear function. In [92] the convergence proof is extended to the more general setting presented above. For both convergence proofs the differentiability requirement can also be dropped. As shown in [75], almost sure *finite* convergence can be achieved for  $\varepsilon$ -optimal policies, for some predefined  $\varepsilon > 0$ .

In [99], Guigues and Monteiro present a slightly different algorithmic approach, called StoDCuP (*stochastic dynamic cutting plane*), in which not only  $\mathcal{Q}_t(\cdot), t = 2, \dots, T$ , but also some or all nonlinear functions  $f_t(\cdot)$  and  $g_t(\cdot)$  are iteratively approximated by affine functions at the trial points visited in the forward pass.

Another variant of SDDP is DASC (*decomposition algorithm for multistage stochastic programs with strongly convex cost functions*), which is introduced in [93]. It can be applied when the (expected) value functions in (MSCP) are strongly convex. For this type of problem, their approximation is proposed using functions  $\mathcal{V}_t(\cdot)$  which are defined as the pointwise maximum of quadratic cuts instead of affine cuts. In contrast to standard SDDP, this means that the subproblems to be solved in SDDP become nonlinear, but in return good approximations of the expected value functions are obtained much quicker, and thus fewer iterations are expected [93].

While most research on SDDP deals with problems (MSLP), some of the extensions presented above and in the following sections have also been enhanced to the convex case, e.g., risk-aversion [92], inexact cuts [95], regularization [97], and exact upper bounding procedures [10, 126]. [92] contains an extension of the convergence proof from [85] to the risk-averse case. Furthermore, the idea of using inexact cuts is generalized to convex nondifferentiable problems [98]; see section 21.

**16. Extensions to Mixed-Integer and Nonconvex Problems [Relaxing Assumption 6].** In many practical applications, multistage stochastic problems involve integer decision variables or nonlinear but nonconvex terms in the objective function or constraints; see section 9. In general, such programs can be formulated in the same way as in the convex case, but it is possible that the functions  $f_t(\cdot)$  and  $g_t(\cdot)$  can be nonconvex. Moreover, in this case,  $X_t$  is the intersection of a convex compact set, e.g., representing box constraints, with possible integer constraints, i.e.,  $X_t \subset \mathbb{R}^{n_{t1}} \times \mathbb{Z}_+^{n_{t2}}$  with  $n_t = n_{t1} + n_{t2}$ . We denote the optimal value by  $v_{NC}^*$ .

Under stagewise independence (Assumption 2), the DPE can be written as (15.2)–(15.4), but for distinction we denote the value functions by  $Q_{t,NC}(x_{t-1}, \xi_t)$  and the expected value functions by  $\mathcal{Q}_{t,NC}(x_{t-1})$  for all  $t = 2, \dots, T$ . Both integer variables and nonconvex functions make this a nonconvex multistage stochastic programming problem (MSNCP). Importantly,  $Q_{t,NC}(\cdot, \cdot)$  and  $\mathcal{Q}_{t,NC}(\cdot)$  are no longer ensured to be convex, but become nonconvex functions in  $x_{t-1}$ . They are also not guaranteed to be (Lipschitz) continuous. This poses significant challenges for approximation algorithms such as SDDP, as linear cuts are not sufficient to approximate  $\mathcal{Q}_{t,NC}(\cdot)$ .

To approach (MSNCP) by SDDP, different strategies can be used. As nonlinear or mixed-integer stochastic programming are large research areas in their own right, we give a brief overview here and refer to the cited literature for methodological details.

**16.1. Convexification.** A standard approach in practice is to solve a static convex relaxation ( $\hat{P}_{NC}$ ) of (MSNCP) which is associated with convex expected value functions  $\hat{\mathcal{Q}}_t(\cdot)$  for all  $t \in [T]$ . Such relaxation can be achieved by relaxing the integrality constraints and replacing nonconvex functions with convex relaxations, such as McCormick envelopes [147]. In this case, though, the Benders cuts determined by SDDP can be very loose. Therefore, only some rough underapproximation  $\hat{v}_{NC}^*$  of the optimal value  $v_{NC}^*$  may be determined. However, sometimes this is considered sufficient to obtain reasonable policies for practical implementation. Also note that even if convex relaxations are considered when running SDDP to compute a policy, the simulation of this policy afterwards can be executed including integrality constraints and nonconvex functions; see, for instance, [193].

A second strategy is to keep the subproblems in SDDP nonconvex, but to convexify the expected value functions  $\mathcal{Q}_{t,NC}(\cdot)$  in some sense. Often, in this case, the nonlinearities in (MSNCP) are first relaxed by piecewise linear approximations such that all subproblems are mixed-integer linear programs (MILPs) [38, 226]. In the backward pass, more specifically, in line 16 of Algorithm 3.1 for all  $t = T, \dots, 2$  and all  $\xi_{tj}, j = 1, \dots, q_t$ , instead of solving a linear programming relaxation of the subproblems (2.10) (or its linear programming dual), a Lagrangian relaxation is solved in which the coupling constraints  $g_t(x_{t-1}, x_t, \xi_{tj}) \leq 0$  are relaxed. For any trial point  $x_{t-1}^{ik}$  and any multiplier  $\pi_t \in \mathbb{R}^{m_t}$ , this relaxation can be written as

$$\begin{aligned} \mathcal{L}_t^{i+1}(\pi_t; x_{t-1}^{ik}, \xi_{tj}) &:= \min_{x_t} f_t(x_t, \xi_{tj}) + \mathcal{Q}_{t+1}(x_t) + \pi_t^\top g_t(x_{t-1}^{ik}, x_t, \xi_{tj}) \\ \text{s.t. } x_t &\in \mathcal{X}_t. \end{aligned}$$

In the Lagrangian dual, this dual function is maximized over all multipliers  $\pi_t$ :

$$(16.1) \quad v_{t,LD}^{i+1}(x_{t-1}^{ik}, \xi_{tj}) := \max_{\pi_t \geq 0} \mathfrak{L}_t^{i+1}(\pi_t; x_{t-1}^{ik}, \xi_{tj}).$$

It is known from the theory on Lagrangian relaxation that the optimal value  $v_{t,LD}^{i+1}(x_{t-1}^{ik}, \xi_{tj})$  coincides with the lower convex envelope of  $Q_{t,NC}^{i+1}(\cdot, \xi_{tj})$  at  $x_{t-1}^{ik}$  [82]. Therefore, cuts obtained based on (16.1) are associated with a convexification of the value function. In order to derive utilizable cut formulas from (16.1), some specific conditions have to be satisfied by the constraints. Suppose the constraints  $g_t(x_{t-1}, x_t, \xi_t) \leq 0$  can be rewritten as

$$\widehat{g}_t(x_{t-1}) - \bar{g}_t(x_t, \xi_t) \leq 0, \quad \tilde{g}_t(x_t, \xi_t) \leq 0,$$

i.e., with the nonlinear function being separable with respect to  $x_{t-1}$ , and let  $\pi_t^{ikj}$  denote optimal multipliers in (16.1). Then, in line with section 3.3, *Lagrangian cuts* can be derived as [220]

$$\mathcal{Q}_{t,NC}(x_{t-1}) \geq \alpha_{tk}^i + (\beta_{tk}^i)^\top \widehat{g}_t(x_{t-1}),$$

with

$$\alpha_{tk}^i = \sum_{j=1}^{q_t} p_{tj} \left( \mathfrak{L}_t(\pi_t^{ikj}; x_{t-1}^{ik}, \xi_{tj}) - (\pi_t^{ikj})^\top \widehat{g}_t(x_{t-1}^{ik}) \right),$$

$$\beta_{tk}^i = \sum_{j=1}^{q_t} p_{tj} \pi_t^{ikj}.$$

For linear functions  $\widehat{g}_t(\cdot)$  and  $\bar{g}_t(\cdot, \cdot)$ , a similar result is derived in [38].

The obtained Lagrangian cuts provably dominate standard Benders cuts, which can be obtained by solving linear programming relaxations [220]. However, convergence of SDDP is not guaranteed, since there may still exist a duality gap between  $v_{t,LD}^{i+1}(x_{t-1}^{ik}, \xi_{tj})$  and  $Q_{t,NC}^{i+1}(x_{t-1}^{ik}, \xi_{tj})$ . Moreover, the generation of Lagrangian cuts can be computationally costly. Various methods have been proposed to solve the Lagrangian dual (16.1), such as cutting-plane methods [118], subgradient methods [73, 174], and bundle methods [129], but all of them can take considerable time compared to solving a linear programming relaxation. Advantageously, even suboptimal Lagrangian multipliers  $\pi_t$  yield valid cuts for  $\mathcal{Q}_{t,NC}(\cdot)$ .

Instead of a static convexification approach [38], Steeger and Rebennack [218, 220] also apply the above principle in a dynamic fashion by considering the DPE for the Lagrangian relaxations in the backward pass.

**16.2. Exact Methods.** Recently, there has been more research on directly applying the SDDP idea to (MSNCP) problems to avoid the requirement of convexification and to close the optimality gap.

**Step Functions.** Given that the value functions  $Q_{t,NC}(\cdot)$  are monotonically increasing or decreasing, they can be approximated by special step functions instead of affine functions. This idea is incorporated into the SDDP framework in the so-called *mixed-integer dynamic approximation scheme* (MIDAS) [170]. To determine the step functions, mixed-integer linear subproblems have to be solved exactly at each stage and in each iteration. In contrast to the previous approaches, convergence of MIDAS to an approximately optimal policy for (MSNCP) is guaranteed.

**SDDiP.** For the mixed-integer linear case, the *stochastic dual dynamic integer programming* (SDDiP) approach by Zou, Ahmed, and Sun [243] allows for the computation of optimal policies for (MSNCP) as long as all state variables  $x_t$  are binary (or bounded integer).

Consider the subproblems (2.10), but with binary state variables  $x_t \in \{0, 1\}^{n_t}$ . Similarly to the approaches in [38, 220, 226], Lagrangian dual problems are solved in the backward pass (line 16 of Algorithm 3.1) to derive valid cuts. However, in SDDiP a new class of Lagrangian cuts is proposed. The crucial idea is to introduce local copies  $z_t$  of the state variables  $x_{t-1}$  and to relax the corresponding copy constraints in the Lagrangian relaxation:

$$\begin{aligned} \mathcal{L}_t^{i+1}(\pi_t; x_{t-1}^{ik}, \xi_{tj}) &:= \min_{x_t, z_t} (c_t(\xi_{tj}))^\top x_t + \Omega_{t+1}(x_t) + \pi_t^\top (x_{t-1}^{ik} - z_t) \\ \text{s.t.} \quad x_t &\in \mathcal{X}_t(z_t, \xi_t), \\ z_t &\in [0, 1]^{d_a(n)}. \end{aligned}$$

In the Lagrangian dual, this dual function is maximized over all multipliers  $\pi_t$ :

$$\tilde{v}_{t,LD}^{i+1}(x_{t-1}^{ik}, \xi_{tj}) := \max_{\pi_t} \mathcal{L}_t^{i+1}(\pi_t; x_{t-1}^{ik}, \xi_{tj}).$$

Then, in line 18 of Algorithm 3.1, Lagrangian cuts can be determined as

$$(16.2) \quad \mathcal{Q}_{t,NC}(x_{t-1}) \geq \alpha_{tk}^i + (\beta_{tk}^i)^\top x_{t-1},$$

with

$$\begin{aligned} \alpha_{tk}^i &= \sum_{j=1}^{q_t} p_{tj} \left( \mathcal{L}_t(\pi_t^{ikj}; x_{t-1}^{ik}, \xi_{tj}) - (\pi_t^{ikj})^\top x_{t-1}^{ik} \right), \\ \beta_{tk}^i &= \sum_{j=1}^{q_t} p_{tj} \pi_t^{ikj}. \end{aligned}$$

These cuts can be proven to be valid and, in particular, tight, as defined in Lemma 3.3. The key aspect behind this tightness property is that for  $x_{t-1} \in \{0, 1\}^{n_t}$  the value functions  $Q_{t,NC}(\cdot)$  coincide with their lower convex envelopes at all  $x_{t-1}$ . Therefore, Lagrangian cuts recovering the latter are also tight for the former.

Moreover, if only dual basic solutions are considered, the cuts (16.2) are also finite in the sense of Lemma 3.3. Therefore, almost sure finite convergence of SDDiP to an optimal policy of (MSNCP) is guaranteed [243].

If the state variables  $x_t$  are bounded and general integer or even continuous, they can be componentwise approximated by a (weighted) sum of binary variables in order to apply SDDiP [243]:

$$x_{tj} \approx \sum_{k=1}^{K_{tj}} 2^{k-1} \beta_{tj} \lambda_{tkj},$$

with discretization precision  $\beta_{ti}$  (for integer  $x_t$ ,  $\beta_t = 1$ ), binary variables  $\lambda_{tkj}$ ,  $k = 1, \dots, K_{tj}$ , and  $K_{tj} \in \mathbb{N}$  for all  $j = 1, \dots, n_t$ . Under some recourse assumptions, it can be proven that for a sufficiently fine binary expansion, an approximately optimal policy for (MSNCP) is computed. However, in practice it may be challenging to choose an appropriate precision in advance.

A more generalized framework to generate Lagrangian cuts in SDDiP is presented in [78] based on ideas from [40]. Additionally, the theory behind Lagrangian cuts and their properties is explored in more detail in [79].


SDDiP is applied in the case studies [110], [181], and [243]. In the latter, additional nonconvex functions occur in (MSNCP), which are linearized.

**Nonconvex Lipschitz Cuts.** As long as the value functions are assured to be Lipschitz continuous (e.g., because the complete continuous recourse [243] property is satisfied), the requirement of binary state variables can be dropped. This is exploited in the *stochastic Lipschitz dynamic programming* (SLDP) method proposed by Ahmed, Cabral, and da Costa in [1], which enhances SDDiP to general MILPs. In contrast to the Lagrangian cuts (16.2), here two types of nonconvex, but Lipschitz continuous cuts are derived to approximate  $\mathcal{Q}_{t,NC}(\cdot)$ : reverse-norm cuts, which are constructed using Lipschitz constants, and augmented Lagrangian cuts, which are based on (16.2) but contain an additional penalization term  $-\mu\|x_{t-1} - x_{t-1}^i\|$ , where  $\mu$  denotes some user-controlled parameter and  $\|\cdot\|$  some arbitrary norm.

This idea is further refined by Zhang and Sun in [240] who propose a new framework to solve multistage nonconvex stochastic mixed-integer nonlinear programs (MINLPs) as part of their complexity analysis of SDDP-like algorithms; see section 4. The first key ingredient of their framework is to consider Lipschitz regularizations of the value functions; see section 17.2. This ensures that the value functions considered are Lipschitz continuous without the requirement of restricting recourse assumptions for (MSNCP). The second idea is to construct nonlinear *generalized conjugacy cuts* by solving conjugate dual problems, similar to the approach in SLDP. Despite being of theoretical interest, this method has not yet been applied in computational experiments. In particular, it is not clear how to solve the conjugate dual problems efficiently in general. Moreover, the framework requires the costly solution of MINLP subproblems in each iteration.

Based on concepts from SDDiP and [240], Füllner and Rebennack present a new framework to solve multistage (stochastic) nonconvex MINLPs [77]. Here, the original MINLP is outer approximated by MILPs using piecewise linear relaxations, which are iteratively improved in an outer loop. In an inner loop, those MILPs are solved by an SDDP- and NBD-like decomposition scheme, which combines the Lipschitz regularization approach from [240] with binary approximation to generate nonconvex cuts. In contrast to SDDiP, the binary approximation is applied only temporarily to derive linear cuts in the lifted binary space, which are then projected back to the original state space. The pointwise maximum of this projection yields a Lipschitz continuous nonconvex cut for the value functions. The projection is computationally important, as it allows the construction of cuts which are guaranteed to also be valid for the outer loop MINLPs. The binary approximation is dynamically refined within the algorithm, instead of a static choice in advance. Another key difference compared to the approach from [240] is that it is not necessary to solve MINLPs in each iteration to derive cuts. The cut projection closure for a nonconvex and discontinuous value function is illustrated in Figure 14.

Similar to SLDP [1], however, it is necessary to introduce a potentially large number of auxiliary variables and constraints to express the nonconvex approximations by mixed-integer linear constraints. While the framework in [77] is presented for deterministic problems, the inner loop decomposition method can be enhanced to the stochastic case [79]. Therefore, by appropriate modifications of the refinement and stopping criteria, the larger framework may also be enhanced to stochastic problems.



**Fig. 14** Nonconvex and discontinuous value function with tight nonconvex cut.

Finally, the use of more general, possibly nonconvex dual functions to approximate  $\mathcal{Q}_t(\cdot)$  is proposed in [30].

**Scaled Cuts.** A recently proposed hybrid between using classical linear cuts and using nonconvex cuts is the use of so-called *scaled cuts* [235]. These cuts are linear, thus computationally preferable, but are derived by scaling nonconvex approximations in such a way that they are guaranteed to asymptotically recover  $\overline{\text{co}}((\cdot)\mathcal{Q}_t(\cdot)(\cdot))$ . This is sufficient to guarantee convergence for MILPs. However, even if linear, so far scaled cuts have only been applied for two-stage stochastic programs, as computing them comes with a lot of additional effort.

**17. Infeasible Subproblems [Relaxing Assumption 9].** Under relatively complete recourse (see Assumption 9), it is guaranteed that any subproblem occurring in the DPE (2.4)–(2.6) and their approximations (2.10) has a feasible solution. As we also assume boundedness, for each of these subproblems there exists some optimal point with finite optimal value. Moreover, all value functions are finite-valued.

In some practical applications, Assumption 9 may not be satisfied. For instance, variable bounds may prevent equality constraints from being satisfied for all  $x_{t-1}$  and all realizations of  $\xi_t$ , as is illustrated by a toy example in [91]. In such a case, the primal subproblems become infeasible and the corresponding dual problems become unbounded. Different measures can be taken to cope with infeasibilities.

**17.1. Feasibility Cuts.** One approach is to approximate the effective domains  $\text{dom}(\mathcal{Q}_t)$  of  $\mathcal{Q}_t(\cdot)$  by cutting away states  $x_{t-1}^{ik} \in \mathcal{X}_t$  leading to infeasible subproblems on stage  $t$ . This can be achieved by generating so-called *feasibility cuts* in addition to the *optimality* cuts derived in section 3. These cuts have the form  $(\beta_t^f)^\top x_{t-1} \leq \alpha_t^f$ , with cut gradient  $\beta_t^f$ , cut intercept  $\alpha_t^f$ , and the superscript  $f$  signifying the cut as a feasibility cut. They can be derived as follows [91].

Consider some stage- $t$  subproblem

$$(17.1) \quad \underline{Q}_t^i(x_{t-1}^{ik}, \xi_t^k) = \begin{cases} \min_{x_t} & (c_t(\xi_t^k))^\top x_t + \underline{\mathcal{V}}_{t+1}^i(x+1) \\ \text{s.t.} & x_t \in \mathcal{X}_t(x_{t-1}^{ik}, \xi_t^k), \\ & (\beta_{t+1}^{fr})^\top x_t \leq \alpha_{t+1}^f, \quad r \in \Gamma_{t+1}^f, \end{cases}$$



in the forward pass of SDDP. This problem may already contain some feasibility cuts, which are indexed by  $r \in \Gamma_{t+1}^f$ . To assess the feasibility of problem (17.1) and construct a feasibility cut if required, we consider the auxiliary feasibility problem

$$v_t^f(x_{t-1}^{ik}, \xi_t^k) := \begin{cases} \min_{x_t, y_t^+, y_t^-, z_t} & e^\top y_t^+ + e^\top y_t^- + e^\top z_t \\ \text{s.t.} & W_t(\xi_t^k)x_t + Iy_t^+ - Iy_t^- = h_t(\xi_t^k) - T_{t-1}(\xi_t^k)x_{t-1}^{ik}, \quad (\sigma_t) \\ & (\beta_{t+1}^{fr})^\top x_t + Iz_t \leq \alpha_{t+1}^{fr}, \quad r \in \Gamma_{t+1}^f, \quad (\omega_t^r) \\ & x_t \geq 0, \\ & y_t^+, y_t^-, z_t \geq 0. \end{cases}$$

Here, slack variables  $y_t^+$ ,  $y_t^-$ , and  $z_t$  are introduced to (17.1) to ensure feasibility. The symbol  $I$  denotes the identity matrix and  $e$  denotes a vector of ones. If we have  $v_t^f(x_{t-1}^{ik}, \xi_t^k) = 0$ , the subproblem (17.1) is feasible; otherwise, it is infeasible.

By strong duality of linear programs,  $v_t^f(x_{t-1}^{ik}, \xi_t^k)$  can be expressed as

$$(17.2) \quad v_t^f(x_{t-1}^{ik}, \xi_t^k) = (h_t(\xi_t^k) - T_{t-1}(\xi_t^k)x_{t-1}^{ik})^\top \sigma_t + \sum_{r \in R_{t+1}^f} (\alpha_{t+1}^{fr})^\top \omega_t^r$$

with optimal dual vectors  $\sigma_t^{ik}$  and  $\omega_t^{ikr}$ ,  $r \in R_{t+1}^f$ . Then, in the case of infeasibility it follows that the term in (17.2) is larger than 0.

To avoid the observed infeasibility at stage  $t$  in future iterations, the stage- $(t-1)$  trial point  $x_{t-1}^{ik}$  should be removed from the feasible set at stage  $t-1$ . This can be achieved by adding the feasibility cut

$$(17.3) \quad -(\sigma_t^{ik})^\top T_{t-1}(\xi_t^k)x_{t-1} + (\sigma_t^{ik})^\top h_t(\xi_t^k) + \sum_{r \in R_{t+1}^f} (\omega_t^{ikr})^\top \alpha_{t+1}^{fr} \leq 0$$

to stage  $t-1$ . By defining

$$\alpha_{t-1}^f := -(\sigma_t^{ik})^\top h_t(\xi_t^k) - \sum_{r \in R_{t+1}^f} (\omega_t^{ikr})^\top \alpha_{t+1}^{fr}$$

and

$$\beta_{t-1}^f := -(\sigma_t^{ik})^\top T_{t-1}(\xi_t^k),$$

the cut (17.3) can be expressed in the previously stated form.

An important question when using feasibility cuts in SDDP is to ask how to proceed once an infeasible subproblem has been detected and a new feasibility cut (17.3) has been generated. For example, it is possible to stop the forward pass and traverse the stages in the backward direction until the root node of the scenario tree is reached. Alternatively, the current subproblem can be resolved to obtain a new trial point  $x_{t-1}^{ik}$  and the forward pass can be continued. For SDDP, no assessment and comparison of these strategies has yet been conducted.

Another drawback is that feasibility cuts do not necessarily prevent infeasibilities when the policy obtained is simulated outside of SDDP [91]. For this reason, the construction of feasibility cuts is usually circumvented in SDDP.

**17.2. Penalization.** Another common approach is to artificially enforce relatively complete recourse for a problem at hand, even if it is not satisfied initially. This can be achieved by using *soft constraints*, that is, introducing slack variables to relax certain

constraints and then penalizing their violation in the objective function. In some applications, this may even be practically justifiable; for example, in load balance equations in power optimization, slack variables can be used to model load shedding or curtailment. However, a reasonable choice of the penalty parameters is not trivial and may distort the expected value functions [91].

**Lipschitz Regularization.** A specific penalization approach is to consider Lipschitz regularizations, also called *Pasch–Hausdorff envelopes* of the value functions. More precisely, let  $\|\cdot\|$  denote some norm,  $\sigma_t > 0$  some constant, and  $z_t$  a local stage- $t$  copy of  $x_{t-1}$ . Then, by allowing  $z_t$  to deviate from the incumbent  $x_{t-1}^{ik}$  and penalizing such deviations in the objective, for all  $t = 2, \dots, T$  and the approximate value functions (2.10) we obtain the approximate Lipschitz-regularized value functions

$$(17.4) \quad \underline{Q}_t^{R;i+1}(x_{t-1}^{ik}, \xi_t; \|\cdot\|) := \min_{z_t \geq 0} \left\{ \underline{Q}_t^{i+1}(z_t, \xi_t) + \sigma_t \|z_t - x_{t-1}^{ik}\| \right\}.$$

These functions have been proven to be Lipschitz continuous on  $\mathbb{R}^{d_a(n)}$  with Lipschitz constant  $\sigma_t$ . Moreover, for sufficiently large  $\sigma_t$  for all  $t \in [T]$ , it can be shown that by considering the regularized problems the original (MSLP) is still solved to optimality [72, 240]. However, how to choose  $\sigma_t$  in a sufficient manner is an open challenge in practice. If  $\sigma_t$  is chosen too small, it may even happen that the Lipschitz-regularized value function is constant  $-\infty$ , given that the subproblem associated with  $\underline{Q}_t^{i+1}(\cdot, \xi_t)$  for some fixed  $\xi_t$  is unbounded for any  $z_t$ .

When Lipschitz-regularized subproblems (17.4) are solved in the forward pass of SDDP (line 10 of Algorithm 3.1), then subgradients of  $\underline{Q}_t^{R;i+1}(\cdot; \|\cdot\|)$  can be derived in the backward pass (line 18) by solving a dual problem (line 16), where the dual variables are bounded by  $\sigma_t$  in the dual norm  $\|\cdot\|^*$  to  $\|\cdot\|$  [240].

**18. No Block-Diagonal Structure [Relaxing Assumption 7].** A key element of dynamic programming methods is that in the multistage decision process only subsequent stages are linked in the constraints, which allows one to express (MSLP) using the DPE (2.4)–(2.6). In the single-problem formulation (2.3) of (MSLP), this coincides with a block-diagonal structure; see Assumption 7.

In some cases, it may be relevant to include constraints spanning multiple stages instead. One example is the incorporation of emission quotas that are not allowed to be exceeded for a given time horizon in energy optimization problems [14, 184, 186].

In order to apply SDDP, the considered (MSLP) has to be reformulated as a problem satisfying Assumption 7. This can be achieved by aggregating stages [58], even though this changes the structure, solution, and interpretability of (MSLP). An alternative approach is to augment the state space. For emission quotas, for instance, instead of summing emissions over several stages and comparing them with the upper bound, at a given stage the remaining emission allowances can be considered as an additional state variable [184]; see section 9.

**19. Infinite Horizon [Relaxing Assumption 1].** Up to this point, we have considered problems (MSLP) with a finite time horizon  $T < \infty$  (Assumption 1). In some practical applications, however, repeated decisions have to be modeled without a clear bound on the horizon. Considering such infinite-horizon problems is, for instance, common for Markov decision processes [19]. In such a case, to ensure that  $v^*$  is finite, a geometric discount factor  $\delta < 1$  is introduced for the cost at each stage.

Since SDDP performs a forward pass and a backward pass through all stages in each iteration, it is not directly applicable to such problems as no iteration would ever be completed. Therefore, different solution methods are often utilized in such

a setting; see, for example, [9]. However, recently there has been some focus on enhancing the SDDP idea to problems with infinite time horizon.

One approach, called Benders squared or  $B^2$ , is based on limiting each iteration of SDDP to a finite horizon of  $\tau$  stages and then dynamically increasing  $\tau$  per iteration, e.g., by 1, until convergence is reached [152]. By presuming that the uncertainty occurs in the RHS and is not only stagewise-independent, but is also i.i.d. for all stages  $t \in [T]$ , almost sure convergence to an approximately optimal policy is assured. The reason for this is that under this special assumption,  $\mathcal{Q}_t(\cdot)$  are the same for all stages, so cuts computed at stage  $t$  can be incorporated not only at stage  $t - 1$ , but at *all stages* [152].

A different option for adapting SDDP to infinite-horizon problems exists if such problems possess some kind of periodical behavior [210]. Assume that for some period  $m \in \mathbb{N}$ , the distributions of  $\xi_t$  as well as the data  $c_t, W_t, h_t$ , and  $T_{t-1}$  are the same for  $t = \tau$  and  $t = \tau + m$  for all  $\tau = 2, \dots$ . Then, under Assumption 9, the functions  $\mathcal{Q}_t(\cdot)$  and  $\mathcal{Q}_{t+m}(\cdot)$  are equivalent as well. This means that it is sufficient to derive cuts for  $\mathcal{Q}_{t+m}(\cdot)$  at stages  $t = 2, \dots, m + 1$  in order to obtain valid cuts for all stages.

In the forward pass of SDDP, it is proposed to only consider a finite number of  $T$  stages starting from stage 1 (see line 8 of Algorithm 3.1), with  $T \geq m + 1$  in order to determine at least one trial point for each of the differing expected value functions. In the case of  $T > m + 1$ , multiple candidate trial points exist at which cuts can be constructed in the backward pass. Before starting the backward pass, the trial points used can be chosen from such a candidate set randomly or by some heuristic.

For both approaches,  $B^2$  and periodic SDDP, for discount factors  $\delta$  close to 1, the influence of late stages on  $v^*$  may be substantial, and thus policy evaluation and upper bound determination can become very challenging and computationally costly. Despite that, Shapiro and Ding [210] propose some proxies based on some finite but sufficiently large  $T$ . However, they do not provide a convergence proof.

A big advantage of SDDP for periodical problems is that it can also be applied to increase the performance for problems with a finite but very large number of stages, given that they satisfy some notion of periodicity. The authors present an example in which for a 60-month horizon, exploiting the periodical structure of the problem, instead of a 60-stage problem only a 13-stage problem has to be solved [210]. This can make even large problems amenable to SDDP and computationally tractable. It is also considered to mitigate the so-called *end-of-horizon effect* discussed in section 9.

On a different note, the policy graph approach introduced by Dowson [60] to model (MSLP) provides a natural extension to infinite-horizon problems, as it allows for cyclic graphs. Solving such problems, similarly to [152], relies on a discount factor and a truncation after a finite number of nodes in the graph. Then, approximate convergence can be proven.

**20. Random Horizon [Relaxing Assumption 1].** Another way to relax Assumption 1 is to assume that the horizon  $T$  is random. For simplicity, we discuss this aspect for the linear case only, even though the convex case is presented in [96].

Consider (MSLP) from section 2.3, satisfying Assumptions 2–8 but with  $T$  not fixed. Instead, we make the following assumption:

*Assumption 13.* The time horizon  $T$  is a discrete random variable taking values in  $\{2, \dots, \bar{T}\}$  with known  $\bar{T} \in \mathbb{N}$ .

Then, the horizon  $T$  induces the Bernoulli process  $(D_t)_{t \in [\overline{T}]}$  with realizations

$$D_t = \begin{cases} 0 & \text{if the optimization period ended at } t \text{ or before,} \\ 1, & \text{otherwise,} \end{cases}$$

and therefore  $T$  can be written as

$$T = \min \{t \in [1, \overline{T}] \mid D_t = 0\}.$$

Under stagewise-independence (Assumption 2), the decisions  $\mathbf{x}_t(\cdot)$  are functions of  $\xi_t$ ,  $\mathbf{D}_t$ ,  $\mathbf{D}_{t-1}$ , and  $x_{t-1}$  (see  $x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)$ ). In other words,  $\mathbf{x}_t$  is  $\mathcal{F}_t$ -measurable with  $\mathcal{F}_t$  the sigma-algebra  $\sigma(\xi_t, \mathbf{D}_j, j \leq t)$  [96].

As shown in [96], for (MSLP) with this type of random horizon, the following DPE can be derived. Importantly, the state space is augmented by  $D_{t-1}$ :

$$Q_t(x_{t-1}, D_t, D_{t-1}, \xi_t) = \min_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} D_{t-1}(c_t(\xi_t))^\top x_t + \mathcal{Q}_{t+1}(x_t, D_t),$$

where

$$\mathcal{Q}_{t+1}(x_t, D_t) = \mathbb{E}_{\xi_{t+1}, \mathbf{D}_t | D_{t-1}} [Q_{t+1}(x_t, \xi_{t+1})]$$

and  $\mathcal{Q}_{\overline{T}+1}(x_{\overline{T}}, D_{\overline{T}}) \equiv 0$ . For the first stage, we obtain

$$v^* = \min_{x_1 \in \mathcal{X}_1(x_0, \xi_1)} c_1^\top x_1 + \mathcal{Q}_2(x_1, D_1).$$

These DPE are the same as those that would be obtained for a problem with a fixed number of stages  $\overline{T} \in \mathbb{N}$ , but with an objective function including the stagewise-dependent stochastic process  $(D_t)_{t \in [\overline{T}]}$ . As  $(D_t)_{t \in [\overline{T}]}$  can be modeled by an inhomogeneous Markov chain with two states, SDDP for processes with Markov chains can be applied [96]; see section 14.4.

**21. Performance Improvements.** Apart from extensions to different problem classes, a lot of research on SDDP has focused on improving its computational performance, because standard SDDP can suffer from various performance issues.

As shown in section 4.2, its worst-case iteration complexity is exponential in the number of stages  $T$  and the dimension  $n_t$  of the state space, the latter being a well-known drawback of cutting-plane methods in general. Whereas SDDP is successfully applied to various large-scale problems in practice (see section 9) with the optimality gap closed in reasonable time, for problems with a large state space in particular it may fail to converge empirically. For instance, Ávila, Papavasiliou, and Löndorf [6] report instances for which the lower bounds  $\underline{v}^i$  already start to stall at a gap of about 22%.

In addition to the high number of iterations required, the computational effort in each iteration can also become substantial, even if the number of subproblems solved per iteration has linear complexity; see section 4.2. The reason for this is that with each iteration of SDDP, the subproblems (2.10) become larger as additional cuts are included. This can increase the computational effort per iteration significantly, especially if many iterations are required until convergence is achieved.

In this section, we give an overview of modifications of SDDP that address these issues and improve its performance.

**21.1. Speeding up SDDP Iterations.** We start with techniques attempting to speed up the SDDP iterations by reducing the computational effort.

**21.1.1. Cut Elimination and Selection.** As mentioned above, with each added cut, the subproblems (2.10) become larger and thus potentially harder to solve. However, computational results indicate that SDDP tends to generate a large number of similar or redundant cutting planes, which do not contribute much to the approximation quality in later iterations [6, 212]. Therefore, the computational burden of SDDP may be reduced if only a subset of all cuts is taken into account. However, this requires careful elimination of cuts that are dominated and do not contribute to the solution process, or careful selection of decisive cuts, otherwise the performance of SDDP may even become worse.

**Cut Elimination.** One way to reduce the number of cuts is to eliminate some cuts permanently. This can be done by repeatedly solving an auxiliary problem after a specified number of iterations, checking the feasibility of the system

$$\begin{cases} \theta_{t+1} \leq \alpha_{t+1}^{\tilde{r}} + (\beta_{t+1}^{\tilde{r}})^{\top} x_t, \\ \theta_{t+1} \geq \alpha_{t+1}^r + (\beta_{t+1}^r)^{\top} x_t, & r \in \Gamma_{t+1} \setminus \{\tilde{r}\}, \\ x_t \in X_t, \end{cases}$$

for each  $\tilde{r} \in \Gamma_{t+1}$ , where  $X_t$  is assumed to be a compact set [212]. If this system is infeasible, then the cut  $\theta_{t+1} \geq \alpha_{t+1}^{\tilde{r}} + (\beta_{t+1}^{\tilde{r}})^{\top} x_t$  is redundant and can be eliminated. The drawback of this method is that the auxiliary problem has to be solved for all cuts in the system.

A different approach is to permanently store all cuts for each stage  $t$ , but only select a subset of those cuts to be considered in each iteration  $i$ . Selection techniques based on this approach are introduced in [8, 53].

**Selecting Last Cuts.** In this naive strategy, only the  $\Gamma \in \mathbb{N}$  most recently added cuts are selected. Although, on average, late cuts may provide a better approximation of  $\mathcal{Q}_t(\cdot)$  than do early cuts, this strategy does not guarantee that all important cuts are considered.

**Level of Dominance.** This strategy is a heuristic to consider only nondominated cuts and avoid the computational effort of the cut elimination approach above. Using the most basic approach, only cuts are selected that yield the highest function value at one of the trial solutions considered so far within the algorithm. This is called *level 1 dominance* [53]. A similar approach is proposed in [163], but there cuts are permanently removed if they are dominated.

Let  $x_t^\ell$  be the trial solution corresponding to the  $\ell$ th cut,  $\ell \in \Gamma_{t+1}$ , and  $\phi^r(x_t^\ell)$  the corresponding function value of cut  $r$ . Then, the values  $v(\ell) := \max_{r \in \Gamma_{t+1}} \{\phi^r(x_t^\ell)\}$  and  $r(\ell) := \arg \max_{r \in \Gamma_{t+1}} \{\phi^r(x_t^\ell)\}$  can be saved in a list and updated every time a new cut is constructed. Similarly, a level  $H$  dominance strategy can be used, selecting the  $H \in \mathbb{N}$  highest cuts for all trial solutions. Using this strategy, only previous trial points are taken into consideration, though. Therefore, cuts may be excluded that provide a significant benefit at not yet visited feasible states.

Another challenge is that this strategy requires a lot of resources to store all the required cut information—especially since the number of trial points visited increases significantly in the course of SDDP. Memory requirements can even be relevant for level 1, especially if the maximum function value at the trial solutions is attained by several cuts. As a remedy, in [94], the *limited memory level 1* strategy is introduced, selecting only the oldest of such cuts. In [8] a more general cut selection strategy is applied to SDDP and almost sure convergence is proven.

**Dynamic Cut Selection.** A dynamic but also computationally more expensive strategy is to select cuts dynamically within the SDDP framework. [53] proposes

removing all cuts at the beginning of each iteration. Then, for each stage  $t$ , each scenario  $k$ , and each function  $\phi^r(\cdot), r \in \Gamma_{t+1}$ , the forward pass subproblem (2.10) is solved. If the current cut yields the highest value at the obtained trial solution, it is added to the subproblem and the next cut is considered. In this way, only those cuts are selected that contribute to the optimal solution in the current iteration. On the other hand, the additional loop in the forward pass may slow down the convergence speed. The computational effort can be reduced by some additional heuristics [53].

A similar approach is considered in [33]. Here, cuts are iteratively added as long as they induce a substantial change in the current optimal value and up to a predefined maximum number of cuts. Instead of iterating over all cuts, in each step, the cut with the highest value at the current trial point is chosen as a candidate for selection.

Numerical results for sampling about 5,000 scenarios and computing 10,000 cuts in SDDP indicate that all cut selection techniques can significantly speed up the classical SDDP method [53]. For example, the level 1 strategy is reported to be ten times faster than SDDP without cut selection. For dynamic cut selection, the reported speed-up is much smaller. It is also shown that the cut selection strategies do not have a significant impact on the quality of the determined policies and bounds. In [8], limited memory level 1 is identified as more efficient than pure level 1.

**21.1.2. Sampling Schemes.** SDDP allows the use of a variety of different sampling schemes that affect its computational performance.

**Number of Forward Samples per Iteration.** In standard SDDP (see section 3),  $|\mathcal{K}|$  scenarios are sampled in each iteration, with  $|\mathcal{K}| \ll |\mathcal{S}|$  and  $\mathcal{K} \subset \mathcal{S}$  (line 6 of Algorithm 3.1). Philpott and Guan even propose a method with only  $|\mathcal{K}| = 1$  for all iterations [171]. This strategy may be particularly efficient in earlier iterations in order to obtain a rough approximation of  $\mathcal{Q}_t(\cdot)$  quickly without wasting too much effort in regions that are likely to be far from optimal. On the other hand, if the current policy is already reasonably good, it should be beneficial to generate more than one new cut per stage and iteration [53]. Moreover, if  $|\mathcal{K}| = 1$ , then it is not possible to apply a statistical stopping criterion; see section 7.

Therefore, instead of fixing  $|\mathcal{K}|$ , a *scenario incrementation* strategy in which  $|\mathcal{K}|$  is gradually increased is a promising approach [211] that is tested in [53].

**Subsampling Trial Points.** In the *reduced sampling method* (ReSa) [109], the forward pass follows the same principle as in SDDP by sampling scenarios  $\xi_t^k, k \in \mathcal{K}$ , for  $\mathcal{K} \subset \mathcal{S}$ . In the backward pass, however, to reduce the number of subproblems to be solved, only a subsample  $\tilde{\mathcal{K}} \subset \mathcal{K}$  is considered and only  $|\tilde{\mathcal{K}}|$  cuts are generated (see line 14 of Algorithm 3.1). Considering more samples in the forward pass without additional effort in the backward pass is helpful in computing accurate statistical upper bounds in an efficient manner.

A similar approach is applied in the *abridged nested decomposition* (AND) method [57]. It is claimed that SDDP is not well designed for bushier scenario trees with large values  $q_t$  because solving  $|\mathcal{K}|q_t$  subproblems per stage may quickly become computationally costly, especially for large  $|\mathcal{K}|$ . On the other hand, a large  $|\mathcal{K}|$  may be required to obtain reliable statistical upper bounds and to incorporate information on sufficiently many scenarios in the trajectories  $(x_t^{ik})_{k \in \mathcal{K}}$ . As a remedy, an alternative sampling scheme is proposed. In the forward pass, on each stage  $t \in [T]$  a set  $\mathcal{K}_t$  of realizations is sampled and trial points  $x_t^{ik}$  are computed. On stage  $t+1$ , however, only a few branching values are used as parameters (in the forward and backward passes), which can either be sampled from or be a convex combination of all  $x_t^{ik}, k \in \mathcal{K}_t$ . The latter idea allows the computation of trial points which contain information on a

large set of scenarios, while keeping the computational effort in the backward pass at a minimum. The main drawback of AND is that the special structure of the forward pass does not allow any direct estimate of an upper bound [109].

**Sampling in the Cut Generation Process.** The computational effort in the backward pass can be reduced if the subproblems (2.10) are solved not for all noise terms  $\xi_{tj}$ ,  $j = 1, \dots, q_t$ , in each iteration, but only for a subset (see line 15 of Algorithm 3.1). The remaining elements that are required to compute a valid cut can then be used from previous iterations where the corresponding noise  $\xi_{tj}$  was sampled.

Further, if the uncertainty is restricted to the RHS  $h_t$  of (MSLP), then the dual feasible set does not depend on it. Therefore, optimal dual multipliers which correspond to dual extreme points can be reused between different realizations  $j = 1, \dots, q_t$ . This allows for the following procedure: Assume that in each iteration  $i$ , for each stage  $t \in [T]$  only one noise term  $\hat{\xi}_t^i$  is sampled and used to compute optimal dual multipliers  $\hat{\pi}_t^i$  and (scenario-specific) cut intercepts  $\hat{\alpha}_t^i$  as in (21.2.1). For each stage  $t = 2, \dots, T$ , all dual multipliers and intercepts obtained up to iteration  $i$  are then stored in a set  $\mathcal{D}_t^i$  defined by  $\mathcal{D}_t^i = \mathcal{D}_t^{i-1} \cup \{(\hat{\pi}_t^i, \hat{\alpha}_t^i, \hat{\xi}_t^i)\}$ .

For any  $\xi_{tj}$ ,  $j = 1, \dots, q_t$ , and a given incumbent  $x_{t-1}^i$ , the dual multipliers used to compute a new cut in line 19 of Algorithm 3.1 can then be determined as

$$(21.1) \quad \left( \hat{\pi}_t^{ij}, \hat{\alpha}_t^{ij}, \hat{\xi}_t^{ij} \right) = \arg \max_{(\hat{\pi}_t, \hat{\alpha}_t, \hat{\xi}_t) \in \mathcal{D}_t^i} \left\{ \hat{\alpha}_t - \hat{\pi}_t^\top T_{t-1} x_{t-1}^i + \hat{\pi}_t^\top (h_t(\xi_{tj}) - h_t(\hat{\xi}_t^i)) \right\}.$$

Hence, optimal dual multipliers for  $\xi_{tj}$  are not necessarily used, but rather those in  $\mathcal{D}_t^i$  providing the best approximation for realization  $\xi_{tj}$  at  $x_{t-1}^i$ .

Let  $\pi_{tj}^i = \hat{\pi}_t^{ij}$  and  $\alpha_{tj}^i = \hat{\alpha}_t^{ij} + (\hat{\pi}_t^{ij})^\top (h_t(\xi_{tj}) - h_t(\hat{\xi}_t^i))$  for all  $j = 1, \dots, q_t$ . Then, a cut can be defined by using subgradient formula (3.7) and taking expectations as in formula (3.4). Note that our description slightly differs from the presentation in the literature, as we have adapted it to our cut formulas in section 3.3.

This idea for the cut generation process is used in two related algorithms, which differ mainly in when cuts are constructed. The CUPPS (*convergent cutting-plane and partial-sampling*) method [41] only contains a forward pass, in which both trial points are computed and cuts are generated using some sample  $\xi_t^{k'}$ . It has the drawback that the cuts obtained are not necessarily tight. First, the dual multipliers obtained from formula (21.1) are not necessarily optimal for all  $j = 1, \dots, q_t$ . Second, no backward pass is used, and thus new information in form of cuts for stage  $t+1$  is not taken into account when deriving a new cut for stage  $t$ .

The *dynamic outer approximation sampling algorithm* (DOASA) [171] contains a forward pass and a backward pass. In the former, a trajectory of trial points  $(x_t^{ik})_{k \in \mathcal{K}}$  is computed as in SDDP (note that in [171]  $|\mathcal{K}| = 1$  is chosen, but this is not mandatory). In the backward pass, cuts are constructed using a backward sample  $\xi_t^{k'}$  and formula (21.1). It is proven that this generalization of SDDP also exhibits almost sure finite convergence [171].

**21.1.3. Inexact SDDP.** Recall Lemma 3.3 (b), which states that the cuts generated in the backward pass of SDDP are *tight* for  $\mathcal{Q}_t^{i+1}(\cdot)$  at the incumbent  $x_{t-1}^{ik}$ . This result is premised on using optimal dual multipliers in the cut formula, i.e., solving the linear programming subproblem or its dual to global optimality in line 16 of Algorithm 3.1 (exact solution). Whereas such an exact solution is the standard assumption in the literature on SDDP, computationally it may be more efficient to solve subproblems only approximately, especially early in the solution process when the cut approximations are suboptimal anyway [95].

We first introduce the notion of inexact cuts.

**DEFINITION 21.1** ( $\varepsilon$ -inexact cut). *For any  $t = 2, \dots, T$ ,  $\varepsilon > 0$ , and a trial point  $x_{t-1}^{ik}$ , let  $\phi_t : \mathbb{R}^{d_{a(n)}} \rightarrow \mathbb{R}$  be an affine function satisfying*

$$\mathcal{Q}_t(x_{t-1}) \geq \underline{\mathcal{Q}}_t^{i+1}(x_{t-1}) \geq \phi_t(x_{t-1}) \quad (\text{validity})$$

for all  $x_{t-1} \in \mathcal{X}_{t-1}$  and

$$\underline{\mathcal{Q}}_t^{i+1}(x_{t-1}^{ik}) - \phi_t(x_{t-1}^{ik}) \leq \varepsilon \quad (\varepsilon\text{-tightness}).$$

Then,  $\phi_t(\cdot)$  defines an  $\varepsilon$ -inexact cut at  $x_{t-1}^{ik}$  [95].

Importantly, inexact cuts still yield valid lower approximations of  $\mathcal{Q}_t(\cdot)$  for all  $t = 2, \dots, T$ . We now address how inexact cuts can be determined.

**Linear Problems.** For any iteration  $i$  in SDDP, any  $t = 2, \dots, T$ , and any trial point  $x_{t-1}^{ik}$ , consider the linear subproblem (2.10). For simplicity, we assume that  $X_t = \{x_t \in \mathbb{R}^{n_t} \mid x_t \geq 0\}$ .

For some  $\varepsilon > 0$ , let  $\pi_{tjk}^i$  be an  $\varepsilon$ -optimal feasible solution for the dual problem of (2.10) given  $\xi_{tj}$  and let  $\theta_{tjk}^i$  be the corresponding dual objective value for  $j = 1, \dots, q_t$ . Then, analogously to section 3.3, an  $\varepsilon$ -inexact cut can be defined by [95]

$$\mathcal{Q}_t(x_{t-1}) \geq \phi_{tk}^i(x_{t-1}) := \alpha_{tk}^i + (\beta_{tk}^i)^\top x_{t-1},$$

with intercept and subgradient defined by

$$\begin{aligned} \alpha_{tk}^i &= \sum_{j=1}^{q_t} p_{tj} (\theta_{tjk}^i - (\beta_{tkj}^i)^\top x_{t-1}^{ik}), \\ \beta_{tk}^i &= - \sum_{j=1}^{q_t} p_{tj} (\pi_{tjk}^i)^\top T_{t-1,j}. \end{aligned}$$

**Nonlinear Differentiable Problems.** Consider (MSCP) as introduced in section 15, that is, satisfying Assumptions 10–12. Moreover, recall the definitions of the Lagrangian function (15.5), the dual function (15.6), and the Lagrangian dual problem (15.7).

Then, an  $\varepsilon$ -inexact cut can be derived using a pair of approximate primal-dual solutions as follows [95]. Let  $\bar{x}_{tj}$  be an  $\varepsilon$ -optimal feasible primal solution for problem (15.2) given some noise realization  $\xi_{tj}$ ,  $j = 1, \dots, q_t$ , and some trial point  $\bar{x}_{t-1}$ , and let  $\bar{\pi}_{tj}$  be an  $\varepsilon$ -optimal feasible solution for the corresponding Lagrangian dual (15.7).

We define

$$(21.2) \quad \eta(\varepsilon) := \ell(\bar{\pi}_{tj}; \bar{x}_{t-1}, \bar{x}_{tj}, \xi_{tj}) := \max_{x_t \in X_t} \nabla_{x_t} L_{t,C}(\bar{\pi}_{tj}; \bar{x}_{t-1}, \bar{x}_{tj}, \xi_{tj})^\top (\bar{x}_{tj} - x_t).$$

Assume that  $f_t(x_t, \xi_{tj})$  takes finite values for all  $x_t \in X_t$  and that the term in (21.2) is finite. Then, an  $\varepsilon$ -inexact cut can be defined by

$$\mathcal{Q}_t(x_{t-1}) \geq \phi_{tk}^i(x_{t-1}) := \alpha_{tk}^i + (\beta_{tk}^i)^\top x_{t-1},$$

with intercept and subgradient defined by



$$\alpha_{tk}^i = \sum_{j=1}^{q_t} p_{tj} (L_{t,C}(\bar{\pi}_{tj}; \bar{x}_{t-1}, \bar{x}_{tj}, \xi_{tj}) - \eta(\varepsilon) - (\beta_{tkj}^i)^\top x_{t-1}^{ik}),$$

$$\beta_{tk}^i = \sum_{j=1}^{q_t} p_{tj} \nabla_{x_{t-1}} L_{t,C}(\bar{\pi}_{tj}; \bar{x}_{t-1}, \bar{x}_{tj}, \xi_{tj}).$$

We refer to [95] for a convergence analysis of SDDP using inexact cuts, for both the linear case and the nonlinear convex case. In particular, it is shown that the dual solutions obtained are almost surely bounded and that the error terms  $\eta(\varepsilon_t^i)$  vanish as  $i$  approaches  $+\infty$ .

**Nondifferentiable Problems.** Using SDDP with inexact cuts is generalized to nondifferentiable problems in [98]. In this paper, inexact cuts are derived using two different approaches. In the first approach, it is assumed that the objective and constraint functions have saddle-point representations. The second approach is more general but requires the introduction of additional variables and constraints.

More precisely, consider (MSCP) as introduced in section 15 and assume that it satisfies Assumptions 10 and 11 except for the differentiability properties. Using a local copy  $z_t$  of the state variable  $x_{t-1}$ , the approximate value functions can be reformulated as

$$(21.3) \quad Q_{t,C}(x_{t-1}, \xi_t) := \begin{cases} \min_{x_t, z_t} & f_t(x_t, \xi_t) + \mathcal{Q}_{t+1,C}(x_t) \\ \text{s.t.} & g_t(z_t, x_t, \xi_t) \leq 0, \\ & x_t \in X_t, \\ & x_{t-1} = z_t. \end{cases}$$

Assume that this modified subproblem satisfies a Slater condition analogous to Assumption 12. Additionally, consider the Lagrangian dual problem

$$(21.4) \quad \max_{\pi_t} \mathcal{L}_t(\pi_t; x_{t-1}, \xi_t)$$

with dual function

$$\mathcal{L}_{t,C}(\pi_t; x_{t-1}, \xi_t) = \begin{cases} \min_{x_t \in X_t} & f_t(x_t, \xi_t) + \mathcal{Q}_{t+1,C}(x_t) + \pi_t^\top (x_{t-1} - z_t) \\ \text{s.t.} & g_t(z_t, x_t, \xi_t) \leq 0, \\ & x_t \in X_t, \end{cases}$$

which is obtained by relaxing the copy constraint.

Given a trial point  $\bar{x}_{t-1}$  and a noise realization  $\xi_{tj}, j = 1, \dots, q_t$ , let  $\bar{x}_{tj}$  denote an  $\varepsilon_P$ -optimal feasible solution of problem (21.3) and let  $\bar{\pi}_{tj}$  be an  $\varepsilon_D$ -optimal feasible solution of problem (21.4). Then, an  $(\varepsilon_P + \varepsilon_D)$ -inexact cut is defined by function

$$\phi_{tk}^i(x_{t-1}) := \sum_{j=1}^{q_t} p_{tj} \left( f_t(\bar{x}_{tj}, \xi_{tj}) - (\varepsilon_P + \varepsilon_D) + \bar{\pi}_{tj}^\top (x_{t-1} - \bar{x}_{t-1}) \right).$$

For more details and a convergence analysis, we refer to [98].

**21.2. Reducing the Number of SDDP Iterations.** We now consider techniques that attempt to reduce the required number of iterations of SDDP until convergence is reached.

**21.2.1. Multicut SDDP.** In the backward pass of SDDP, for any  $t \in [T]$  and any  $x_{t-1}^{ik}, k \in \mathcal{K}$ , subproblems (2.10) are solved for all noise realizations  $\xi_{tj}$ ,  $j = 1, \dots, q_t$  (line 16 of Algorithm 3.1). By taking expected values, a cut (3.5) is derived (line 18). Such cuts are then incorporated into the stage- $(t-1)$  subproblem using a single variable  $\theta_t \in \mathbb{R}$  by

$$\phi_{tk}^i(x_{t-1}) = (\beta_{tk}^i)^\top x_{t-1} + \alpha_{tk}^i \leq \theta_t;$$

see section 3.3. This is referred to as a single-cut approach.

A different approach, called multicut, which is well studied for (nested) Benders decomposition [28, 81, 151], is to not aggregate the dual information, but to generate a separate cut for each noise realization  $\xi_{tj}$ ,  $j = 1, \dots, q_t$ . This means that a similar step to line 18 is included in the loop over  $j$  in Algorithm 3.1. In the stage- $t$  subproblem, this requires the introduction of variables  $\theta_{t,\ell}$  and cut approximations  $\underline{Q}_{t+1,\ell}^{i+1}(\cdot)$  for all  $\ell = 1, \dots, q_t$ . In this case, we obtain cuts

$$\phi_{tkj}^i(x_{t-1}) := (\beta_{tkj}^i)^\top x_{t-1} + \alpha_{tkj}^i \leq Q_t(x_{t-1}, \xi_{tj}), \quad j = 1, \dots, q_t,$$

where, analogously to the derivation in section 3.3,  $\beta_{tkj}^i$  denotes a subgradient of  $\underline{Q}_t^{i+1}(\cdot, \xi_{tj})$  at  $x_{t-1}^{ik}$  for  $k \in \mathcal{K}, j = 1, \dots, q_t$ , and  $\alpha_{tkj}^i$  is defined by

$$\alpha_{tkj}^i := \underline{Q}_t^{i+1}(x_{t-1}^{ik}, \xi_t) - (\beta_{tkj}^i)^\top x_{t-1}^{ik}.$$

The expectation is then taken in the objective function instead of the cut formula:

$$\underline{Q}_t^{i+1}(x_{t-1}^{ik}, \xi_{tj}) = \min_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} (c_t(\xi_{tj}))^\top x_t + \sum_{\ell=1}^{q_{t+1}} p_{t+1,\ell} \underline{Q}_{t+1,\ell}^{i+1}(x_t).$$

In this way, more specific information about the value functions is incorporated into the subproblems, hopefully leading to fewer iterations. Moreover, it can be shown that multicut SDDP has the same convergence properties as SDDP [8]. On the downside, the number of decision variables and cuts grows significantly compared to the single-cut approach, especially if  $q_t$  is large, which increases the computational effort for each iteration. Therefore, up to this point multicut SDDP has rarely been considered in the literature. It should be most promising when  $q_t$  is only of moderate size. For the two-stage case, a rule of thumb is that a single-cut approach should be preferred if the number of realizations is considerably larger than the number of first-stage constraints [27]. Note that, in principle, a trade-off between single-cut and multicut is also possible by partially aggregating cuts [24, 29]. Another approach to reducing the computational burden of multicut SDDP is to combine it with cut selection strategies (see section 21.1.1), as proposed in [8].

We return to Example 3.4 to illustrate the multicut approach.

*Example 21.2* (continuation of Example 3.4). Using multicut SDDP, at stage 3, instead of  $\mathcal{Q}_3(\cdot)$ , the functions  $Q_3(\cdot, \xi_3)$  are separately approximated by cuts for  $\xi_3 \in \{1, 2, 4\}$ . These value functions are displayed in Figure 15. Each one consists of only two linear pieces, so two cuts are required to represent them exactly. In contrast,  $\mathcal{Q}_3(\cdot)$  consists of four linear segments. Therefore, multicut SDDP should need fewer iterations than single-cut SDDP to achieve convergence.




Fig. 15 Stage-3 value functions for Example 3.4.

**21.2.2. Batch Learning and Experience Replay.** While SDDP is used in stochastic programming, dynamic programming, and optimal control, its methodology also shares some characteristics with  $Q$ -learning algorithms, which are studied in reinforcement learning; see Remark 3.1. This can be exploited by translating established performance enhancing techniques from reinforcement learning to SDDP.

As one such technique, Ávila, Papavasiliou, and Löhdorf [6] propose the use of a *batch learning* technique called *experience replay* in SDDP. The motivation for this is that in SDDP, the cut approximations  $\underline{V}_t(\cdot)$  of the expected value functions  $\mathcal{Q}_t(\cdot)$  are generated recursively in a backward pass through the stages  $t = T, \dots, 2$ . This means that approximation errors at later stages are propagated to earlier stages by means of the cut approximations  $\underline{V}_t(\cdot)$ , which then leads to loose cuts at these earlier stages, and so on. However, this implies that errors are accumulated at early stages. The authors identify this as a driver of the slow convergence of SDDP, as it favors the overexploring of suboptimal regions and the generation of redundant cuts throughout the iterations.

Experience replay addresses this issue by revisiting previous trial points  $x_t^i$  and updating the cut approximations  $\underline{V}_t(\cdot)$  at these points. This seems counterintuitive at first glance because cuts are generated at already-visited points instead of improving the approximation of  $\mathcal{Q}_t(\cdot)$  at regions of  $\mathcal{X}_t$  that have yet not been visited. However, by taking into account all the information currently available to update  $\underline{V}_t(\cdot)$  at  $x_t^i$ , it avoids the issue that at earlier stages  $\tau < t$  unnecessarily poor approximations of  $\mathcal{Q}_t(\cdot)$  at  $x_t^i$  are used for several more iterations.

More precisely, the proposed SDDP method works as follows. A predefined number of iterations of standard SDDP are executed and the corresponding trial points  $x_t^i$  are stored in a replay memory  $M_t$  for all  $t \in [T - 1]$  in line 10 of Algorithm 3.1. When the number of elements of the replay memory reaches a predefined cardinality  $Z$ , then the experience replay step is initiated. This step performs a backward pass through the stages  $t = T - 1, \dots, 2$ . For each stage  $t$ , first a batch  $B_t \subseteq M_t$  of trial points is selected from the replay memory (a full batch  $B_t = M_t$  is also possible). For each trial point  $\tilde{x}_t^\ell$  from this batch, with  $\ell = 1, \dots, |B_t|$ , the previously generated cut is removed from  $\underline{V}_{t+1}^{i+1}(\cdot)$  and a new cut is constructed by solving the associated subproblems (2.10) (including the experience replay updates from following stages) for  $\tilde{x}_t^\ell$ . With these cuts,  $\underline{V}_{t+1}^i(\cdot)$  is updated and then the previous stage is explored. After the replay step, a standard iteration of SDDP is started again.

It is shown that experience replay manages to improve the convergence behavior of SDDP and also the out-of-sample performance of the obtained policies [6]. However, experience replay comes at an increased computational effort, as every  $Z$  iterations an additional backward pass solving  $q_t|B_t|$  subproblems for each stage  $t = T, \dots, 2$  has to be performed. For full batches, this adds up to  $q_t|K|Z$  LPs per stage. For this reason, the authors suggest parallelizing both standard SDDP iterations as well as the experience replay. They report computational results which indicate that batch learning is better at exploiting parallelism than standard SDDP.

**21.2.3. Regularization.** As Kelley's cutting-plane method [118, 153], SDDP exhibits an iteration complexity which is exponential in the dimension  $n_t$  of the state variables; see section 4.2. An unfavorable characteristic of cutting-plane methods, and also of SDDP, in this regard is *zig-zagging* behavior. This means that trial points  $x_t^i$  and  $x_t^{i+1}$  computed in subsequent iterations can be located far from one another in different regions of the state space, and that with each new cut the minimum of the subproblems (2.10) is again attained in the respective other region. In particular, this implies that these regions of  $\mathcal{X}_t$  experience very tight, but almost redundant approximations  $\mathcal{Y}_t(\cdot)$  of  $\mathcal{Q}_t(\cdot)$ , while other regions are not properly explored and thus the approximation quality at the true optimum improves very slowly.

In convex and nonsmooth optimization, regularization techniques called *bundle methods* are shown to have faster convergence than classical cutting-plane methods [129], as they mitigate zig-zagging by stabilizing subsequent trial points around a *stability center* (also called incumbent). Hence, translating these regularization techniques to SDDP looks promising.

A common regularization approach, which is predominantly used in two-stage stochastic programming [198, 201], is convex quadratic regularization. Here, some quadratic deviation of  $x_t$  from a stability center  $\hat{x}_t$  is penalized in the objective function for stabilization. An application of quadratic regularization to SDDP is not straightforward because using a separate stability center for each scenario  $s \in \mathcal{S}$  is computationally infeasible due to the exponential growth of  $|\mathcal{S}|$  in  $T$  [5].

Therefore, Asamov and Powell [5] propose a regularization technique for linear problems in which stability centers are considered part of the state variable and thus are the same for all realizations of  $\xi_{tj}, j = 1, \dots, q_t$ . Then, in the forward pass subproblems (2.10) (see line 10 of Algorithm 3.1), the objective function is modified to

$$(21.5) \quad c_t^\top x_t + \mathcal{Y}_{t+1}^i(x+1) + \frac{\gamma^i}{2}(x_t - \hat{x}_t^{i-1})^\top H_t(x_t - \hat{x}_t^{i-1}),$$

with a positive semidefinite matrix  $H_t$  and some sequence  $(\gamma^i)_{i \in \mathbb{N}}$  satisfying  $\gamma^i \geq 0$  for all  $i$  and  $\lim_{i \rightarrow \infty} \gamma^i = 0$ . The stability centers  $\hat{x}_t^{i-1}$  are chosen as the previous forward pass solution, i.e., the solution is stabilized around a “known” region of the domain of  $\mathcal{Q}_t(\cdot)$ . This idea is generalized to nonlinear problems and improved in [97] by considering weighted averages of several previous forward pass solutions.

Using objective (21.5), a convex, continuous, and linearly constrained quadratic programming problem has to be solved in each forward pass step of SDDP, hopefully reducing the required number of iterations. Importantly, only the forward pass of SDDP is changed, while the backward pass remains the same. In particular, only linear programs have to be solved in the backward pass. As the cuts are still finite (see Lemma 3.3), almost sure finite convergence is assured. In computational tests, it is shown that this method exhibits faster convergence than SDDP, in particular for a high state dimension  $n_t$  [5]. This speed-up is especially important for regularized

DDP (dual dynamic programming); see the numerical experiments in [97]. DDP is the corresponding deterministic counterpart of SDDP (when  $\xi_t$  is deterministic for all  $t \in [T]$ ).

Whereas the above approach stabilizes the solution around a “known” region of the domain of  $\mathcal{Q}_t(\cdot)$ , in a sampling setting, it is not clear whether this is always beneficial. For the current sample  $\xi_t^k$  a region may be identified and used for stabilization that is not an appropriate indicator for all  $\xi_{tj}, j = 1, \dots, q_t$ . Additionally, as pointed out in [233], the condition  $\lim_{i \rightarrow \infty} \gamma^i = 0$  might mean that the regularization is diminished and the proposed method in [5] reduces to standard SDDP *before* convergence is obtained, although regularization may be particularly important close to the optimal solution. Therefore, this is claimed to be detrimental to convergence speed [233].

Van Ackooij and Warin [233] also address the issue that convergence of proximal bundle methods usually requires the stability centers to be feasible, which is not guaranteed for SDDP subproblems where the feasible set changes with  $x_{t-1}^i$ . Therefore, they propose combining SDDP with a level bundle method which does not face this requirement.

For stage  $t$  and scenario  $\xi_t^k$ , trial solutions  $x_t^{ik}$  are obtained by solving

$$(21.6) \quad \begin{cases} \min_{x_t} & \psi_t(x_t) \\ \text{s.t.} & x_t \in \mathbb{X}_t(x_{t-1}^{ik}; \ell_t), \end{cases}$$

with  $\psi_t(x_t) : \mathbb{R}^{n_t} \rightarrow \mathbb{R}$  a given convex function, e.g.,  $\psi_t(x_t) := x_t^\top x_t$ , and

$$(21.7) \quad \mathbb{X}_t(x_{t-1}^{ik}; \ell_t) := \begin{cases} \arg \min_{x_t \geq 0} & \max \{c_t^\top x_t + \mathcal{V}_{t+1}^i(x+1), \ell_t\} \\ \text{s.t.} & W_t x_t = h_t - T_{t-1} x_{t-1}^{ik}. \end{cases}$$

If the maximum in (21.7) is attained by the first term, then  $x_t^{ik}$  obtained by solving (21.6) is an ordinary SDDP trial point, referred to as a *normal iterate*. Otherwise, problem (21.6) reduces to a typical level bundle method subproblem, yielding a regularized *level iterate*  $x_t^{ik}$ .

The determination of a good level  $\ell_t$  and an efficient regularization for SDDP are still open questions, and heuristics are proposed in [233] to choose  $\ell_t$ .

In [23] a nonconvex regularization scheme is suggested instead of the previous convex schemes. The main motivation behind this is that the existing schemes tend to heavily penalize the exploration of the state space right from the beginning, even if it is well known that early trajectories in SDDP are far from optimal. By using specific concave penalty functions it is possible to penalize large deviations from the stability centers much less than for the quadratic regularization from [5], but at the same time to heavily penalize small deviations. One example of such a penalty function is the the minimax concave penalty

$$(21.8) \quad \sum_{j=1}^{n_t} \lambda \int_0^{|x_{tj} - \hat{x}_{tj}^i|} \max \left\{ 1 - \frac{v}{\alpha \lambda}, 0 \right\} dv$$

with parameters  $\lambda, \alpha$ , which is known from statistical learning [239].

By using this penalty function for regularization in the forward pass of SDDP, in early iterations exploration of the state space is incentivized. In addition, letting  $\lambda \rightarrow 0$  in SDDP, for later iterations, more and more focus is then put on exploitation close to the current stability center. It is shown that even for the nonconvex regularization

using (21.8) the SDDP subproblems remain tractable, as they can be reformulated as equivalent MILPs.

Computational results in [23] indicate that with this regularization, the iteration times of SDDP increase, but on the other hand, many fewer iterations are required until SDDP terminates due to bound stalling, and that the lower bounds at termination exceed those of standard SDDP and the quadratic regularization (21.5).

An alternative stabilization approach is proposed in [15] based on the concept of Chebyshev centers of polyhedrons. Here, in the forward pass of SDDP, the subproblems (2.10) are modified such that the computed trial states are defined as Chebyshev centers of the polyhedrons given by previously constructed cuts and an appropriate upper bound. It can be shown that this approach is equivalent to modifying the cut formula to

$$(21.9) \quad -(\beta_{t+1}^r)^\top x_t + \theta_{t+1} \geq \alpha_{t+1}^r + \bar{\sigma}_t \|(1, c_t + \beta_{t+1}^r)\|, \quad r \in \Gamma_{t+1}.$$

The authors use the Euclidean norm  $\|\cdot\|_2$  in (21.9); however, different choices are possible as well.

Geometrically, the additional term in (21.9) changes the cut intercept, thus lifting the cut. For  $\bar{\sigma}_t = 0$ , the usual SDDP trial point  $x_t^i$  is determined, whereas for  $\bar{\sigma}_t > 0$  an offset in the objective compared to the standard SDDP subproblem is considered, yielding a different iterate. To actually improve the performance of SDDP, choosing  $\bar{\sigma}_t$  appropriately is crucial, but not trivial. By contrast, if  $\bar{\sigma}_t$  is chosen too large, basically any feasible point can become the new trial solution. Moreover, to ensure convergence, it has to be ensured that  $\bar{\sigma}_t$  converges to zero in the course of the algorithm. In [15] heuristics are used to determine  $\bar{\sigma}_t$ , but it is not clear whether they guarantee performance gains for SDDP.

**21.3. Parallelization.** The performance of SDDP can be improved not only by modifications of the algorithm itself, but also by its implementation and computational execution. Since several computational steps in SDDP are independent of one another, performance improvement can be achieved by parallelization.

Different parallelization strategies have been proposed for SDDP. They can be classified with respect to how the workload is distributed among different processors and how the processors are synchronized. Based on this observation, Ávila, Papavasiliou, and Löhdorf [6] present a taxonomy of parallelization strategies, which we follow in this section.

**Parallelization by Scenario.** This is the predominant parallelization strategy for SDDP in the literature. Generally, a *synchronized* version is proposed. In the forward pass, for all  $t \in [T]$ , the subproblems (2.10) are solved for  $|\mathcal{K}|$  different scenarios, which are sampled independently. The uncertain data  $\xi_t^k$  and the trial solutions  $x_{t-1}^k$  in each of those problems only depend on scenario  $k$ . Therefore, the different scenarios  $\xi^k, k \in \mathcal{K}$ , can be assigned to different processors. Assuming  $P$  different processors, each processor is assigned  $\frac{P}{|\mathcal{K}|}$  scenarios and solves all corresponding subproblems. A master process is then used to aggregate the objective values and compute the upper bound estimate (3.9). This means that there is a synchronization point for all processors at the end of the forward pass.

In the backward pass, a similar approach is followed. The subproblems are again distributed among the processors by scenarios, in such way that for a specific stage  $t = T, \dots, 2$  and a scenario-based trial point  $x_{t-1}^k$ , the subproblems for all noise realizations  $\xi_{tj}, j = 1, \dots, q_t$ , are solved by the same processor. Evenly distributing the problems

between processors, in this way each processor solves  $\frac{P}{|\mathcal{K}|}q_t$  subproblems. However, it is also possible to let the master process assign new scenarios to processes once they become idle instead of using a fixed assignment scheme [173].

After solving all associated subproblems, each processor then generates a cut for  $\mathcal{Q}_t(\cdot)$  and sends it to the master process. When cut generation is finished for all  $k \in \mathcal{K}$ , the processors are synchronized so that they can all proceed with the same set of cuts on stage  $t - 1$ . As stated in [106], this synchronization can be partially relaxed to avoid waiting for single slow processors. Instead, the master process can assign stage- $(t - 1)$  subproblems to available processors even if not all cuts have yet been generated for stage  $t$ . Numerical results show that such a partial relaxation can improve the computational performance of SDDP. However, the number of cuts to wait for to achieve an optimal trade-off between faster iterations and better approximation of  $\mathcal{Q}_t(\cdot)$  is problem-dependent.

Further, an *asynchronous* approach can be used in which processors immediately return to stage  $t - 1$  after generating their cuts at stage  $t$ , using all cuts currently available without waiting for other processes to finish [61].

A major shortcoming of parallelization by scenario is that using more processors becomes more beneficial the more scenarios  $|\mathcal{K}|$  are sampled in the forward pass. However, as discussed in section 21.1.2, it is often favorable to only consider one or a few scenarios per iteration, especially in earlier iterations. Choosing large  $|\mathcal{K}|$  may lead to the accumulation of similar trial points and the generation of redundant cuts [6]. Therefore, exploiting the potential performance gains of additional processors may wrongly incentivize the sampling of more scenarios than is reasonable, thus not accelerating but in fact slowing down the solution process. In addition, Ávila, Papavasiliou, and Löhdorf [6] report computational results indicating that (synchronized) parallelization by scenario scales poorly when increasing the number of samples  $|\mathcal{K}|$  due to the combination of long waiting times between processors and low quality cuts.

**Parallelization by Node.** Using parallelization by node, the strategy is to draw only one or a few samples in the forward pass, as this is often computationally preferable. Then, the forward pass is not necessarily parallelized. In the backward pass, the work is distributed among the processors by nodes of the recombining tree (cf. section 2.1). That means that even for the same  $k \in \mathcal{K}$  and the associated trial point  $x_{t-1}^k$ , the subproblems (2.10) for different realizations  $\xi_{tj}, j = 1, \dots, q_t$ , may be solved by different processors. The processors are synchronized at each stage to generate aggregated cuts (given that a single-cut approach is used).

In [6], the authors report clear computational benefits when using parallelization by node compared to parallelization by scenario, and also better scaling properties. However, these results require that the processors can access a shared memory, otherwise the computational overhead is too large. Another drawback is that distributing subproblems for different  $\xi_{tj}$  but the same  $x_{t-1}^{ik}$  among different scenarios prevents the exploitation of warm starting techniques.

Parallelization by node can also be used in an asynchronous way, as proposed by Machado et al. [143] in their *asynchronous SDDP* method. In this method, the subproblems of all stages  $t = 1, \dots, T$  are solved simultaneously. More precisely, in each *step*, for all stages  $t = 1, \dots, T$  and scenarios  $k \in \mathcal{K}$ , the subproblems for all realizations  $\xi_{tj}, j = 1, \dots, q_t$ , are solved. Once a processor is finished, it constructs a new cut for  $\mathcal{Q}_t(\cdot)$  using all available information. If a required processor has yet not finished, multipliers  $\pi_{tkj}$  from previous steps are reused. The generated cut can then be incorporated into stage  $t - 1$  in the next step. Additionally, each processor generates

a new trial point which can be used at stage  $t$  in the next step. In contrast to SDDP iterations, this approach requires several steps to propagate information through all stages. Therefore, an ordinary forward pass can only be observed implicitly over several stages. This has to be considered in the computation of upper bounds.

Independent of the applied strategy, parallelizing SDDP in practice comes with considerable challenges, such as communication overhead, problem-dependent performance, and lack of reproducibility of results. Therefore, its potential to speed up SDDP in general is naturally limited [6].

**21.4. Aggregation Techniques.** Aggregating information in (MSLP) is another tool with the potential to speed up the SDDP solution process.

One approach is to aggregate the variables and constraints of several time periods into a single stage, thus solving a problem with a smaller horizon  $T$ . This is straightforward for NBD [58], where each node of the aggregated problem is a subtree of the original scenario tree, even though only few time periods can be aggregated to keep the subproblems tractable. However, it cannot be directly generalized to the sampling and stagewise-independent setting in SDDP. The main issue is that it is difficult to model the uncertainty appropriately, without violating nonanticipativity [58].

An alternative approach is to aggregate realizations of  $\xi_t$  on each stage (or a subset of stages) [214]. To this end, for some stage  $t$ , the support  $\Xi_t$  is partitioned into clusters  $C_t^\ell, \ell = 1, \dots, L_t$ , with  $L_t \in \mathbb{N}$ . Instead of solving subproblems associated with  $Q_t^{i+1}(x_{t-1}^{ik}, \xi_{tj})$  for all  $j = 1, \dots, q_t$  in the backward pass of SDDP (line 16 of Algorithm 3.1), subproblems associated with  $Q_t^{i+1}(x_{t-1}^{ik}, \bar{\xi}_t^\ell)$  are solved for clusters  $\ell = 1, \dots, L_t$ , with  $\bar{\xi}_t^\ell := \sum_{j \in C_t^\ell} \frac{p_{tj}}{\bar{p}_t^\ell} \xi_{tj}$ , and  $\bar{p}_t^\ell$  the probability of cluster  $C_t^\ell$ . This should be beneficial in early iterations where policies are still far from optimal and a fine information structure unnecessarily slows down the solution process.

By using subgradients and intercepts associated with clusters  $C_t^\ell, \ell = 1, \dots, L_t$ , in line 18 of Algorithm 3.1 *coarse cuts* can be generated for  $\mathcal{Q}_t(\cdot)$ . Given that  $W_t$  and  $c_t$  are deterministic, these cuts are valid underestimators for  $\mathcal{Q}_t(\cdot)$  by Jensen's inequality [214]. They are not guaranteed to be tight, though.

The authors in [214] discuss several different refinement strategies, such as refinements within the SDDP backward pass (the partition at stage  $t$  is refined as soon as a coarse cut does not improve the approximation of  $\mathcal{Q}_t(\cdot)$  at the trial point  $x_{t-1}^{ik}$ ) or refinements outside of SDDP. In the latter case, SDDP is performed on a coarse recombining tree which is iteratively refined once the algorithm has stopped. Computational results show that this latter approach performs significantly better than the first approach due to lower computational overhead. However, identifying when SDDP should best be stopped to perform a refinement remains a challenging task.

**22. Outlook.** In this tutorial-style review, we have given an overview of the motivation, theory, strengths and weaknesses, extensions, and applications of SDDP.

While many proposals have been made in the last 30 years on how to extend SDDP and improve its performance, there remain open research questions that leave room for future improvement. Among the most crucial topics are those listed in what follows.

1. *Stopping.* To this date, in many applications SDDP is stopped heuristically, e.g., based on a fixed number of iterations or stabilization of lower bounds, which leaves the task of defining a reasonable stopping criterion to the user. Recently, there has been some pioneering work on developing deterministic



upper bounding techniques and stopping criteria, but these are still limited as they require significant computational effort.

2. *Upper bounds in risk-averse SDDP.* Developing efficient upper bounding techniques is especially relevant to risk-averse variants of SDDP, where the commonly used nested risk measures do not allow for employment of their pendants from risk-neutral SDDP. Lately, different risk measures have been proposed that avoid this issue. However, such risk measures usually hamper interpretability. Therefore, the question of how risk should be optimally measured in SDDP in order to obtain a computationally tractable problem, while at the same time properly reflecting the true risk preferences of a decision-maker, can still be regarded as open.
3. *Distributionally robust SDDP.* Recently, the consideration of distributional uncertainty in SDDP has gained more interest. However, while DRO is a flourishing research area, incorporating it into SDDP is still in its early stages, with potential for further improvements.
4. *Nonconvex extensions.* In many applications, nonlinear functions or integer variables are required to appropriately model the problem at hand. As the (expected) value functions become nonconvex in this case, traditional cutting-plane techniques fail to approximate them correctly. Starting with SDDiP, recently there has been a trend to extend the NBD and SDDP frameworks to nonconvex problems. Lagrangian-type cuts, which are possibly nonconvex, show theoretical potential in approximating nonconvex functions. However, their construction is computationally costly and subject to rather strong technical assumptions, such that large-scale nonconvex problems in particular remain computationally intractable. Consequently, in the future, the trade-off between computationally efficient cut generation techniques and best possible approximations of the value functions needs to be further explored.
5. *Regularization.* As a descendant of Kelley's cutting-plane method, SDDP has a computational complexity that grows exponentially in the dimension of the state variables. Therefore, it can become computationally intractable for problems with high-dimensional state space. This is aggravated by common reformulations, e.g., in the case of stagewise-dependent uncertainty, that artificially augment the state space. For Kelley's method, regularization methods have proven helpful in accelerating the solution process. Whereas some first attempts have been made to regularize SDDP, an efficient regularization remains an open challenge.
6. *Reinforcement learning techniques.* As the case of batch learning shows, SDDP can benefit from acceleration techniques that are well known and established in reinforcement learning, but have not yet been translated to the SDDP setting. By exploiting its affinity to  $Q$ -learning, there should be a lot of potential to improve the computational performance of SDDP in practice.
7. *Decision-dependent uncertainty.* The only standard assumption for SDDP that has not yet been relaxed in the literature is to allow for stagewise-dependent stochastic processes modeling the uncertainty in (MSLP). This topic still has to be studied.

**Acknowledgments.** The authors thank the three anonymous referees and the associate editor for their detailed and insightful feedback which helped to substantially improve earlier versions of this article.

## REFERENCES

- [1] S. AHMED, F. G. CABRAL, AND B. FREITAS PAULO DA COSTA, *Stochastic Lipschitz dynamic programming*, Math. Program., 191 (2022), pp. 755–793, <https://doi.org/10.1007/s10107-020-01569-z>. (Cited on p. 509)
- [2] D. ARJOON, Y. MOHAMED, Q. GOOR, AND A. TILMANT, *Hydro-economic risk assessment in the eastern Nile River basin*, Water Resources Econ., 8 (2014), pp. 16–31, <https://doi.org/10.1016/j.wre.2014.10.004>. (Cited on p. 460)
- [3] P. ARTZNER, F. DELBAEN, J.-M. EBER, AND D. HEATH, *Coherent measures of risk*, Math. Finance, 9 (1999), pp. 203–228, <https://doi.org/10.1111/1467-9965.00068>. (Cited on p. 469)
- [4] N. V. ARVANITIDITS AND J. ROSING, *Composite representation of a multireservoir hydro-electric power system*, IEEE Trans. Power Appar. Syst., 89 (1970), pp. 319–326, <https://doi.org/10.1109/TPAS.1970.292595>. (Cited on p. 458)
- [5] T. ASAMOV AND W. B. POWELL, *Regularized decomposition of high-dimensional multistage stochastic programs with Markov uncertainty*, SIAM J. Optim., 28 (2018), pp. 575–595, <https://doi.org/10.1137/16M1072231>. (Cited on pp. 461, 522, 523)
- [6] D. AVILA, A. PAPAVALIOU, AND N. LÖHNDORF, *Batch learning SDDP for long-term hydro-thermal planning*, IEEE Trans. Power Syst., 39 (2024), pp. 614–627, <https://doi.org/10.1109/TPWRS.2023.3246724>. (Cited on pp. 425, 457, 514, 515, 521, 522, 524, 525, 526)
- [7] H. BAKKER, F. DUNKE, AND S. NICKEL, *A structuring review on multi-stage optimization under uncertainty: Aligning concepts from theory and practice*, Omega, 96 (2020), 102080, <https://doi.org/10.1016/j.omega.2019.06.006>. (Cited on p. 418)
- [8] M. BANDARRA AND V. GUIGUES, *Single cut and multicut stochastic dual dynamic programming with cut selection for multistage stochastic linear programs: Convergence proof and numerical experiments*, Comput. Manage. Sci., 18 (2021), pp. 125–141, <https://doi.org/10.1007/s10287-021-00387-8>. (Cited on pp. 440, 460, 461, 515, 516, 520)
- [9] R. BAUCKE, *An Algorithm for Solving Infinite Horizon Markov Dynamic Programmes*, preprint, [http://www.optimization-online.org/DB\\_FILE/2018/04/6565.pdf](http://www.optimization-online.org/DB_FILE/2018/04/6565.pdf), 2018. (Cited on p. 513)
- [10] R. BAUCKE, A. DOWNWARD, AND G. ZAKERI, *A Deterministic Algorithm for Solving Multistage Stochastic Programming Problems*, preprint, [http://www.optimization-online.org/DB\\_FILE/2017/07/6138.pdf](http://www.optimization-online.org/DB_FILE/2017/07/6138.pdf), 2017. (Cited on pp. 439, 447, 452, 455, 506)
- [11] R. BAUCKE, A. DOWNWARD, AND G. ZAKERI, *A Deterministic Algorithm for Stochastic Minimax Dynamic Programmes*, preprint, [http://www.optimization-online.org/DB\\_FILE/2018/02/6449.pdf](http://www.optimization-online.org/DB_FILE/2018/02/6449.pdf), 2018. (Cited on pp. 439, 447, 483, 501, 502)
- [12] G. BAYRAKSAN AND D. K. LOVE, *Data-driven stochastic programming using phi-divergences*, in *Tutorials in Operations Research: The Operations Research Revolution*, INFORMS, 2015, pp. 1–19. (Cited on p. 487)
- [13] R. E. BELLMAN, *Dynamic Programming*, Princeton University Press, Princeton, NJ, 1957. (Cited on pp. 418, 419, 427)
- [14] M. M. BELSNES, A. HAUGSTAD, B. MO, AND P. MARKUSSEN, *Quota modeling in hydro-thermal systems*, in *IEEE Bologna Power Tech Conference Proceedings*, 2003, pp. 1–6, <https://doi.org/10.1109/PTC.2003.1304382>. (Cited on pp. 459, 512)
- [15] F. BELTRÁN, E. C. FINARDI, G. M. FREDO, AND W. DE OLIVEIRA, *Improving the performance of the stochastic dual dynamic programming algorithm using Chebyshev centers*, Optim. Eng., 23 (2022), pp. 147–168, <https://doi.org/10.1007/s11081-020-09558-z>. (Cited on p. 457, 524)
- [16] A. BEN-TAL, L. EL GHAOU, AND A. NEMIROVSKI, *Robust Optimization*, Princeton University Press, 2009. (Cited on p. 485)
- [17] J. F. BENDERS, *Partitioning procedures for solving mixed-variables programming problems*, Numer. Math., 4 (1962), pp. 238–252, <https://doi.org/10.1007/BF01386316>. (Cited on pp. 419, 429, 438, 443)
- [18] J. M. VELÁSQUEZ BERMÚDEZ, *GDDP: Generalized dual dynamic programming theory*, Ann. Oper. Res., 117 (2002), pp. 21–31, <https://doi.org/10.1023/A:1021557003554>. (Cited on p. 428)
- [19] D. P. BERTSEKAS, *Dynamic programming and optimal control*, Vol. II, 4th ed., Athena Scientific, 2017. (Cited on p. 512)
- [20] D. BERTSIMAS, D. B. BROWN, AND C. CARAMANIS, *Theory and applications of robust optimization*, SIAM Rev., 53 (2011), pp. 464–501, <https://doi.org/10.1137/080734510>. (Cited on p. 485)
- [21] J. BEZANSON, A. EDELMAN, S. KARPINSKI, AND V. B. SHAH, *Julia: A fresh approach to numerical computing*, SIAM Rev., 59 (2017), pp. 65–98, <https://doi.org/10.1137/141000671>. (Cited on p. 461)

- [22] A. BHATTACHARYA, J. P. KHAROUFEH, AND B. ZENG, *Managing energy storage in micro-grids: A multistage stochastic programming approach*, IEEE Trans. Smart Grid, 9 (2018), pp. 483–496, <https://doi.org/10.1109/TSG.2016.2618621>. (Cited on p. 459)
- [23] A. BHATTACHARYA, J. P. KHAROUFEH, AND B. ZENG, *A nonconvex regularization scheme for the stochastic dual dynamic programming algorithm*, INFORMS J. Comput., 35 (2023), pp. 1161–1178, <https://doi.org/10.1287/ijoc.2021.0255>. (Cited on pp. 523, 524)
- [24] M. BIEL AND M. JOHANSSON, *Dynamic Cut Aggregation in L-Shaped Algorithms*, preprint, <https://arxiv.org/abs/1910.13752>, 2019. (Cited on p. 520)
- [25] J. R. BIRGE, *Solution Methods for Stochastic Dynamic Linear Programs*, Ph.D. dissertation, Systems Optimization Laboratory, Stanford University, 1980. (Cited on p. 419)
- [26] J. R. BIRGE, *Decomposition and partitioning methods for multistage stochastic linear programs*, Oper. Res., 33 (1985), pp. 989–1007, <https://doi.org/10.1287/opre.33.5.989>. (Cited on p. 443)
- [27] J. R. BIRGE AND F. LOUVEAUX, *Introduction to Stochastic Programming*, 2nd ed., Springer Series in Operations Research and Financial Engineering, Springer Science & Business Media, 2011. (Cited on pp. 418, 429, 436, 520)
- [28] J. R. BIRGE AND F. V. LOUVEAUX, *A multicut algorithm for two-stage stochastic linear programs*, European J. Oper. Res., 34 (1988), pp. 384–392, [https://doi.org/10.1016/0377-2217\(88\)90159-2](https://doi.org/10.1016/0377-2217(88)90159-2). (Cited on p. 520)
- [29] X. BLANCHOT, F. CLAUTIAUX, B. DETIENNE, A. FROGER, AND M. RUIZ, *The Benders by batch algorithm: Design and stabilization of an enhanced algorithm to solve multicut Benders reformulation of two-stage stochastic programs*, European J. Oper. Res., 309 (2023), pp. 202–216, <https://doi.org/10.1016/j.ejor.2023.01.004>. (Cited on p. 520)
- [30] S. BOLUSANI AND T. K. RALPHS, *A framework for generalized Benders' decomposition and its application to multilevel optimization*, Math. Program., 196 (2022), pp. 1–38, <https://doi.org/10.1007/s10107-021-01763-7>. (Cited on p. 510)
- [31] J. F. BONNANS, Z. CEN, AND T. CHRISTEL, *Energy contracts management by stochastic programming techniques*, Ann. Oper. Res., 200 (2012), pp. 199–222, <https://doi.org/10.1007/s10479-011-0973-5>. (Cited on p. 499)
- [32] R. B. S. BRANDI, A. L. M. MARCATO, B. H. DIAS, T. P. RAMOS, AND I. C. DA SILVA JUNIOR, *A convergence criterion for stochastic dual dynamic programming: Application to the long-term operation planning problem*, IEEE Trans. Power Syst., 33 (2018), pp. 3678–3690, <https://doi.org/10.1109/TPWRS.2017.2787462>. (Cited on pp. 457, 482)
- [33] R. B. S. BRANDI, T. P. RAMOS, B. H. DIAS, A. L. M. MARCATO, AND I. CHAVES DA SILVA JUNIOR, *Improving stochastic dynamic programming on hydrothermal systems through an iterative process*, Electr. Pow. Syst. Res., 123 (2015), pp. 147–153, <https://doi.org/10.1016/j.epsr.2015.02.011>. (Cited on pp. 457, 516)
- [34] A. BRIGATTO, A. STREET, AND D. M. VALLADÃO, *Assessing the cost of time-inconsistent operation policies in hydrothermal power systems*, IEEE Trans. Power Syst., 32 (2017), pp. 4541–4550, <https://doi.org/10.1109/TPWRS.2017.2672204>. (Cited on p. 457)
- [35] S. BRUNO, S. AHMED, A. SHAPIRO, AND A. STREET, *Risk neutral and risk averse approaches to multistage renewable investment planning under uncertainty*, European J. Oper. Res., 250 (2016), pp. 979–989, <https://doi.org/10.1016/j.ejor.2015.10.013>. (Cited on p. 460)
- [36] L. CAMBIER AND D. SCIEUR, *FAST (Finally An SDDP Toolbox)*, <https://github.com/leopoldcambier/FAST> (accessed April 27, 2021). (Cited on p. 461)
- [37] P. CARPENTIER, J.-P. CHANCELIER, G. COHEN, M. DE LARA, AND P. GIRARDEAU, *Dynamic consistency for stochastic optimal control problems*, Ann. Oper. Res., 200 (2012), pp. 247–263, <https://doi.org/10.1007/s10479-011-1027-8>. (Cited on p. 471)
- [38] S. CERISOLA, J. M. LATORRE, AND A. RAMOS, *Stochastic dual dynamic programming applied to nonconvex hydrothermal models*, European J. Oper. Res., 218 (2012), pp. 687–697, <https://doi.org/10.1016/j.ejor.2011.11.040>. (Cited on pp. 459, 506, 507, 508)
- [39] R. M. CHABAR, M. V. F. PEREIRA, S. GRANVILLE, L. A. BARROSO, AND N. A. ILIADIS, *Optimization of fuel contracts management and maintenance scheduling for thermal plants under price uncertainty*, in IEEE Power Systems Conference and Exposition, IEEE, 2006, pp. 923–930. (Cited on pp. 459, 498)
- [40] R. CHEN AND J. LUEDTKE, *On generating Lagrangian cuts for two-stage stochastic integer programs*, INFORMS J. Comput., 34 (2022), pp. 2332–2349, <https://doi.org/10.1287/ijoc.2022.1185>. (Cited on p. 509)
- [41] Z.-L. CHEN AND W. B. POWELL, *A convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse*, J. Optimiz. Theory App., 102 (1999), pp. 497–524, <https://doi.org/10.1023/A:1022641805263>. (Cited on pp. 419, 438, 517)

- [42] A. CHIRALAKSANAKUL AND D. P. MORTON, *Assessing policy quality in multi-stage stochastic programming*, Stochastic Programming E-Print Series, 12 (2004), <https://doi.org/10.18452/8319>. (Cited on pp. 464, 465)
- [43] B. F. P. DA COSTA AND V. LECLÈRE, *Dual SDDP for risk-averse multistage stochastic programs*, Oper. Res. Lett., 51 (2023), pp. 332–337, <https://doi.org/10.1016/j.orl.2023.04.001>. (Cited on pp. 455, 481)
- [44] B. F. P. DA COSTA AND V. LECLÈRE, *Duality of Upper Bounds in Stochastic Dynamic Programming*, preprint, [https://optimization-online.org/wp-content/uploads/2023/07/dual\\_rn.pdf](https://optimization-online.org/wp-content/uploads/2023/07/dual_rn.pdf), 2023. (Cited on pp. 455, 481)
- [45] L. C. DA COSTA, JR., F. S. THOMÉ, J. DIAS GARCIA, AND M. V. F. PEREIRA, *Reliability-constrained power system expansion planning: A stochastic risk-averse optimization approach*, IEEE Trans. Power Syst., 36 (2021), pp. 97–106, <https://doi.org/10.1109/TPWRS.2020.3007974>. (Cited on p. 460)
- [46] E. L. DA SILVA AND E. C. FINARDI, *Parallel processing applied to the planning of hydrothermal systems*, IEEE Trans. Parall. Distrib. Syst., 14 (2003), pp. 721–729, <https://doi.org/10.1109/TPDS.2003.1225052>. (Cited on p. 457)
- [47] G. B. DANTZIG AND G. INFANGER, *Multi-stage stochastic linear programs for portfolio optimization*, Ann. Oper. Res., 45 (1993), pp. 59–76, <https://doi.org/10.1007/BF02282041>. (Cited on p. 460)
- [48] G. B. DANTZIG, *Linear programming under uncertainty*, Manage. Sci., 1 (1955), pp. 197–206, <https://doi.org/10.1287/mnsc.1.3-4.197>. (Cited on p. 418)
- [49] C. M. B. DE CASTRO, A. L. M. MARCATO, R. CASTRO SOUZA, I. CHAVES SILVA JUNIOR, F. L. CYRINO OLIVEIRA, AND T. PULINHO, *The generation of synthetic inflows via bootstrap to increase the energy efficiency of long-term hydrothermal dispatches*, Electr. Pow. Syst. Res., 124 (2015), pp. 33–46, <https://doi.org/10.1016/j.epsr.2015.02.014>. (Cited on p. 458)
- [50] M. DE LARA AND V. LECLÈRE, *Building up time-consistency for risk-measures and dynamic optimization*, European J. Oper. Res., 249 (2016), pp. 177–187, <https://doi.org/10.1016/j.ejor.2015.03.046>. (Cited on p. 471)
- [51] V. L. DE MATOS AND E. C. FINARDI, *A computational study of a stochastic optimization model for long term hydrothermal scheduling*, Int. J. Electric. Power, 43 (2012), pp. 1443–1452, <https://doi.org/10.1016/j.ijepes.2012.06.021>. (Cited on pp. 457, 458, 459)
- [52] V. L. DE MATOS, D. P. MORTON, AND E. C. FINARDI, *Assessing policy quality in a multistage stochastic program for long-term hydrothermal scheduling*, Ann. Oper. Res., 253 (2017), pp. 713–731, <https://doi.org/10.1007/s10479-016-2107-6>. (Cited on pp. 457, 458, 463, 464, 465, 466, 467)
- [53] V. L. DE MATOS, A. B. PHILPOTT, AND E. C. FINARDI, *Improving the performance of stochastic dual dynamic programming*, J. Comput. Appl. Math., 290 (2015), pp. 196–208, <https://doi.org/10.1016/j.cam.2015.04.048>. (Cited on pp. 449, 457, 515, 516)
- [54] L. DING, S. AHMED, AND A. SHAPIRO, *A Python Package for Multi-Stage Stochastic Programming*, preprint, [http://www.optimization-online.org/DB\\_FILE/2019/05/7199.pdf](http://www.optimization-online.org/DB_FILE/2019/05/7199.pdf), 2019. (Cited on p. 461)
- [55] T. DING, X. ZHANG, R. LU, M. QU, M. SHAHIDEHPOUR, Y. HE, AND T. CHEN, *Multi-stage distributionally robust stochastic dual dynamic programming to multi-period economic dispatch with virtual energy storage*, IEEE Trans. Sustain. Energy, 13 (2022), pp. 1–13, <https://doi.org/10.1109/TSTE.2021.3105525>. (Cited on p. 459)
- [56] A. L. DINIZ, M. E. P. MACEIRA, C. L. V. VASCONCELLOS, AND D. D. J. PENNA, *A combined SDDP/Benders decomposition approach with a risk-averse surface concept for reservoir operation in long term power generation planning*, Ann. Oper. Res., 292 (2020), pp. 649–681, <https://doi.org/10.1007/s10479-019-03419-4>. (Cited on pp. 457, 460)
- [57] C. J. DONOHUE AND J. R. BIRGE, *The abridged nested decomposition method for multistage stochastic linear programs with relatively complete recourse*, Algorithmic Oper. Res., 1 (2006), pp. 20–30. (Cited on pp. 419, 516)
- [58] T. N. DOS SANTOS AND A. L. DINIZ, *A new multiperiod stage definition for the multistage Benders decomposition approach applied to hydrothermal scheduling*, IEEE Trans. Power Syst., 24 (2009), pp. 1383–1392, <https://doi.org/10.1109/TPWRS.2009.2023265>. (Cited on p. 512, 526)
- [59] A. DOWNWARD, O. DOWSON, AND R. BAUCKE, *Stochastic dual dynamic programming with stagewise-dependent objective uncertainty*, Oper. Res. Lett., 48 (2020), pp. 33–39, <https://doi.org/10.1016/j.orl.2019.11.002>. (Cited on pp. 440, 460, 501, 502)
- [60] O. DOWSON, *The policy graph decomposition of multistage stochastic programming problems*, Networks, 76 (2020), pp. 3–23, <https://doi.org/10.1002/net.21932>. (Cited on pp. 461, 497, 513)

- [61] O. DOWSON AND L. KAPELEVICH, *SDDP.jl: A Julia package for stochastic dual dynamic programming*, INFORMS J. Comput., 33 (2020), pp. 27–33, <https://doi.org/10.1287/ijoc.2020.0987>. (Cited on p. 461, 525)
- [62] O. DOWSON, D. P. MORTON, AND A. DOWNWARD, *Bi-objective multistage stochastic linear programming*, Math. Program., 196 (2022), pp. 907–933, <https://doi.org/10.1007/s10107-022-01872-x>. (Cited on pp. 484, 485)
- [63] O. DOWSON, D. P. MORTON, AND B. K. PAGNONCELLI, *Partially observable multistage stochastic programming*, Oper. Res. Lett., 48 (2020), pp. 505–512, <https://doi.org/10.1016/j.orl.2020.06.005>. (Cited on pp. 461, 489, 490)
- [64] O. DOWSON, D. P. MORTON, AND B. K. PAGNONCELLI, *Incorporating convex risk measures into multistage stochastic programming algorithms*, Ann. Oper. Res., (2022), pp. 1–25, <https://doi.org/10.1007/s10479-022-04977-w>. (Cited on pp. 461, 470, 472, 482, 483)
- [65] O. DOWSON, A. PHILPOTT, A. MASON, AND A. DOWNWARD, *A multi-stage stochastic optimization model of a pastoral dairy farm*, European J. Oper. Res., 274 (2019), pp. 1077–1089, <https://doi.org/10.1016/j.ejor.2018.10.033>. (Cited on p. 461)
- [66] I. DUNNING, J. HUCHETTE, AND M. LUBIN, *JuMP: A modeling language for mathematical optimization*, SIAM Rev., 59 (2017), pp. 295–320, <https://doi.org/10.1137/15M1020575>. (Cited on p. 461)
- [67] J. DUPAČOVÁ AND V. KOZMÍK, *SDDP for multistage stochastic programs: Preprocessing via scenario reduction*, Comput. Manage. Sci., 14 (2017), pp. 67–80, <https://doi.org/10.1007/s10287-016-0261-6>. (Cited on p. 461)
- [68] J. DUPAČOVÁ AND V. KOZMÍK, *Structure of risk-averse multistage stochastic programs*, OR Spectrum, 37 (2015), pp. 559–582, <https://doi.org/10.1007/s00291-014-0379-2>. (Cited on pp. 461, 471, 474, 483)
- [69] D. DUQUE AND D. P. MORTON, *Distributionally robust stochastic dual dynamic programming*, SIAM J. Optim., 30 (2020), pp. 2841–2865, <https://doi.org/10.1137/19M1309602>. (Cited on pp. 440, 487, 488, 489)
- [70] J. L. DURANTE, J. NASCIMENTO, AND W. B. POWELL, *Risk Directed Importance Sampling in Stochastic Dual Dynamic Programming with Hidden Markov Models for Grid Level Energy Storage*, preprint, <https://arxiv.org/pdf/2001.06026.pdf>, 2020. (Cited on pp. 460, 490)
- [71] A. EICHHORN AND W. RÖMISCH, *Polyhedral risk measures in stochastic programming*, SIAM J. Optim., 16 (2005), pp. 69–95, <https://doi.org/10.1137/040605217>. (Cited on p. 472)
- [72] M. J. FEIZOLLAHI, S. AHMED, AND X. A. SUN, *Exact augmented Lagrangian duality for mixed integer linear programming*, Math. Program., 161 (2017), pp. 365–387, <https://doi.org/10.1007/s10107-016-1012-8>. (Cited on p. 512)
- [73] M. L. FISHER, *An applications oriented guide to Lagrangian relaxation*, Interfaces, 15 (1985), pp. 10–21, <https://doi.org/10.1287/inte.15.2.10>. (Cited on p. 507)
- [74] B. C. FLACH, L. A. BARROSO, AND M. V. F. PEREIRA, *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 (2010), pp. 299–314, <https://doi.org/10.1049/iet-gtd.2009.0107>. (Cited on p. 457)
- [75] M. FORCIER AND V. LECLÈRE, *Trajectory following dynamic programming algorithms without finite support assumptions*, J. Convex Anal., 30 (2023), pp. 951–999. (Cited on pp. 425, 427, 438, 439, 440, 441, 448, 462, 505)
- [76] S. M. FRANK AND S. REBENNACK, *An introduction to optimal power flow: Theory, formulation, and examples*, IIE Trans., 48 (2016), pp. 1172–1197, <https://doi.org/10.1080/0740817X.2016.1189626>. (Cited on p. 457)
- [77] C. FÜLLNER AND S. REBENNACK, *Non-convex nested Benders decomposition*, Math. Program., 196 (2022), pp. 987–1024, <https://doi.org/10.1007/s10107-021-01740-0>. (Cited on p. 509)
- [78] C. FÜLLNER, X. A. SUN, AND S. REBENNACK, *A New Framework to Generate Lagrangian Cuts in Multistage Stochastic Mixed-Integer Programming*, preprint, <https://optimization-online.org/?p=27312.pdf>, 2024. (Cited on p. 509)
- [79] C. FÜLLNER, X. A. SUN, AND S. REBENNACK, *On Lipschitz Regularization and Lagrangian Cuts in Multistage Stochastic Mixed-Integer Linear Programming*, preprint, <https://optimization-online.org/?p=27295.pdf>, 2024. (Cited on p. 509)
- [80] J. DIAS GARCIA, I. LEAL DE FREITAS, R. CHABAR, AND M. V. F. PEREIRA, *A Multicut Approach to Compute Upper Bounds for Risk-Averse SDDP*, preprint, <https://arxiv.org/pdf/2307.13190.pdf>, 2023. (Cited on pp. 476, 477, 482)
- [81] H. I. GASSMANN, *MSLiP: A computer code for the multistage stochastic linear programming problem*, Math. Program., 47 (1990), pp. 407–423, <https://doi.org/10.1007/BF01580872>. (Cited on p. 520)

- [82] A. M. GEOFFRION, *Lagrangian relaxation for integer programming*, in Approaches to Integer Programming, M. L. Balinski, ed., Springer, Berlin, Heidelberg, 1974, pp. 82–114. (Cited on p. 507)
- [83] A. GEORGHIOU, A. TSOUKALAS, AND W. WIESEMANN, *Robust dual dynamic programming*, Oper. Res., 67 (2019), pp. 813–830, <https://doi.org/10.1287/opre.2018.1835>. (Cited on pp. 447, 501)
- [84] H. GEVRET, N. LANGRENÉ, J. LELONG, R. LOBATO, T. OUILLO, X. WARIN, AND A. MAHESHWARI, *STochastic OPTimization library in C++*, Technical report, EDF Innovation Lab, Palo Alto, 2018. (Cited on p. 461)
- [85] P. GIRARDEAU, V. LECLÈRE, AND A. B. PHILPOTT, *On the convergence of decomposition methods for multistage stochastic convex programs*, Math. Oper. Res., 40 (2015), pp. 130–145, <https://doi.org/10.1287/moor.2014.0664>. (Cited on pp. 425, 439, 501, 505, 506)
- [86] A. GJELSVIK, M. M. BELSNES, AND A. HAUGSTAD, *An algorithm for stochastic medium-term hydrothermal scheduling under spot price uncertainty*, in Proceedings of 13th Power Systems Computation Conference, 1999, pp. 1079–1085. (Cited on p. 459, 497, 498)
- [87] A. GJELSVIK, M. M. BELSNES, AND M. HÅLAND, *A case of hydro scheduling with a stochastic price model*, in Proceedings of the 3rd International Conference on Hydropower, E. Broch, D. K. Lysne, N. Flatabø, and E. Helland-Hansen, eds., A.A. Balkema, 1997, pp. 211–218. (Cited on pp. 457, 459)
- [88] A. GJELSVIK, B. MO, AND A. HAUGSTAD, *Long- and medium-term operations planning and stochastic modelling in hydro-dominated power systems based on stochastic dual dynamic programming*, in Handbook of Power Systems I, Energy Systems, P. M. Pardalos, S. Rebennack, M. V. F. Pereira, and N. A. Iliadis, eds., Springer-Verlag, Berlin Heidelberg, 2010, pp. 33–55. (Cited on pp. 457, 459, 497, 498)
- [89] A. GJELSVIK AND S. W. WALLACE, *Methods for Stochastic Medium-term Scheduling in Hydro-dominated Power Systems*, Technical report, Norwegian Electric Power Research Institute, Trondheim, 1996. (Cited on p. 498)
- [90] Z. GUAN AND A. B. PHILPOTT, *A multistage stochastic programming model for the New Zealand dairy industry*, Int. J. Prod. Econ., 134 (2011), pp. 289–299, <https://doi.org/10.1016/j.ijpe.2009.11.003>. (Cited on p. 461)
- [91] V. GUIGUES, *SDDP for some interstage dependent risk-averse problems and application to hydro-thermal planning*, Comput. Optim. Appl., 57 (2014), pp. 167–203, <https://doi.org/10.1007/s10589-013-9584-1>. (Cited on pp. 457, 460, 473, 491, 492, 494, 495, 510, 511, 512)
- [92] V. GUIGUES, *Convergence analysis of sampling-based decomposition methods for risk-averse multistage stochastic convex programs*, SIAM J. Optim., 26 (2016), pp. 2468–2494, <https://doi.org/10.1137/140983136>. (Cited on pp. 439, 504, 505, 506)
- [93] V. GUIGUES, *DASC: A Decomposition Algorithm for Multistage Stochastic Programs with Strongly Convex Cost Functions*, preprint, <https://arxiv.org/pdf/1711.04650.pdf>, 2017. (Cited on p. 505)
- [94] V. GUIGUES, *Dual dynamic programming with cut selection: Convergence proof and numerical results*, European J. Oper. Res., 258 (2017), pp. 47–57, <https://doi.org/10.1016/j.ejor.2016.10.047>. (Cited on pp. 440, 461, 515)
- [95] V. GUIGUES, *Inexact cuts in stochastic dual dynamic programming*, SIAM J. Optim., 30 (2020), pp. 407–438, <https://doi.org/10.1137/18M1211799>. (Cited on pp. 506, 517, 518, 519)
- [96] V. GUIGUES, *Multistage stochastic programs with a random number of stages: Dynamic programming equations, solution methods, and application to portfolio selection*, Optim. Methods Softw., 36 (2021), pp. 211–236, <https://doi.org/10.1080/10556788.2020.1800007>. (Cited on pp. 513, 514)
- [97] V. GUIGUES, M. A. LEJEUNE, AND W. TEKAYA, *Regularized stochastic dual dynamic programming for convex nonlinear optimization problems*, Optim. Eng., 21 (2020), pp. 1133–1165, <https://doi.org/10.1007/s11081-020-09511-0>. (Cited on pp. 460, 506, 522, 523)
- [98] V. GUIGUES, R. MONTEIRO, AND B. SVAITER, *Inexact cuts in SDDP applied to multistage stochastic nondifferentiable problems*, SIAM J. Optim., 31 (2021), pp. 2084–2110, <https://doi.org/10.1137/20M1330075>. (Cited on pp. 506, 519)
- [99] V. GUIGUES AND R. D. C. MONTEIRO, *Stochastic dynamic cutting plane for multistage stochastic convex programs*, J. Optim. Theory Appl., 189 (2021), pp. 513–559, <https://doi.org/10.1007/s10957-021-01842-x>. (Cited on p. 505)
- [100] V. GUIGUES AND W. RÖMISCH, *Sampling-based decomposition methods for multistage stochastic programs based on extended polyhedral risk measures*, SIAM J. Optim., 22 (2012), pp. 286–312, <https://doi.org/10.1137/100811696>. (Cited on pp. 472, 473)

- [101] V. GUIGUES AND W. RÖMISCH, *SDDP for multistage stochastic linear programs based on spectral risk measures*, Oper. Res. Lett., 40 (2012), pp. 313–318, <https://doi.org/10.1016/j.orl.2012.04.006>. (Cited on p. 473)
- [102] V. GUIGUES AND C. SAGASTIZÁBAL, *The value of rolling-horizon policies for risk-averse hydro-thermal planning*, European J. Oper. Res., 217 (2012), pp. 129–140, <https://doi.org/10.1016/j.ejor.2011.08.017>. (Cited on p. 467)
- [103] V. GUIGUES AND C. SAGASTIZÁBAL, *Risk-averse feasible policies for large-scale multistage stochastic linear programs*, Math. Program., 138 (2013), pp. 167–198, <https://doi.org/10.1007/s10107-012-0592-1>. (Cited on p. 467)
- [104] V. GUIGUES, A. SHAPIRO, AND Y. CHENG, *Duality and sensitivity analysis of multistage linear stochastic programs*, European J. Oper. Res., 308 (2023), pp. 752–767, <https://doi.org/10.1016/j.ejor.2022.11.051>. (Cited on pp. 451, 454, 455, 456, 457, 461, 481, 495, 502)
- [105] V. GUIGUES, A. SHAPIRO, AND Y. CHENG, *Risk-averse stochastic optimal control: An efficiently computable statistical upper bound*, Oper. Res. Lett., 51 (2023), pp. 393–400, <https://doi.org/10.1016/j.orl.2023.05.002>. (Cited on pp. 474, 479)
- [106] A. HELSETH AND H. BRAATEN, *Efficient parallelization of the stochastic dual dynamic programming algorithm applied to hydropower scheduling*, Energies, 8 (2015), pp. 14287–14297, <https://doi.org/10.3390/en8121431>. (Cited on p. 525)
- [107] A. HELSETH, M. FODSTAD, AND B. MO, *Optimal medium-term hydropower scheduling considering energy and reserve capacity markets*, IEEE Trans. Sustain. Energy, 7 (2016), pp. 934–942, <https://doi.org/10.1109/TSTE.2015.2509447>. (Cited on pp. 459, 498)
- [108] A. HELSETH, B. MO, M. FODSTAD, AND M. N. HJELMELAND, *Co-optimizing sales of energy and capacity in a hydropower scheduling model*, in IEEE Power Tech, Eindhoven, The Netherlands, 2015, pp. 1–6. (Cited on pp. 459, 498)
- [109] M. HINDSBERGER, *ReSa: A method for solving multistage stochastic linear programs*, J. Appl. Oper. Res., 6 (2014), pp. 2–15. (Cited on pp. 419, 516, 517)
- [110] M. N. HJELMELAND, J. ZOU, A. HELSETH, AND S. AHMED, *Nonconvex medium-term hydropower scheduling by stochastic dual dynamic integer programming*, IEEE Trans. Sustain. Energy, 10 (2019), pp. 481–490, <https://doi.org/10.1109/TSTE.2018.2805164>. (Cited on pp. 459, 509)
- [111] T. HOMEM-DE-MELLO AND G. BAYRAKSAN, *Monte Carlo sampling-based methods for stochastic optimization*, Surveys Oper. Res. Manage. Sci., 19 (2014), pp. 56–85, <https://doi.org/10.1016/j.sorms.2014.05.001>. (Cited on p. 445)
- [112] T. HOMEM-DE-MELLO, V. L. DE MATOS, AND E. C. FINARDI, *Sampling strategies and stopping criteria for stochastic dual dynamic programming: A case study in long-term hydrothermal scheduling*, Energy Syst., 2 (2011), pp. 1–31, <https://doi.org/10.1007/s12667-011-0024-y>. (Cited on pp. 445, 446, 449, 450, 457, 466)
- [113] T. HOMEM-DE-MELLO AND B. K. PAGNONCELLI, *Risk aversion in multistage stochastic programming: A modeling and algorithmic perspective*, European J. Oper. Res., 249 (2016), pp. 188–199, <https://doi.org/10.1016/j.ejor.2015.05.048>. (Cited on pp. 461, 470, 471, 472, 474, 483)
- [114] J. HUANG, K. ZHOU, AND Y. GUAN, *A Study of Distributionally Robust Multistage Stochastic Optimization*, preprint, <https://arxiv.org/abs/1708.07930>, 2017. (Cited on p. 487)
- [115] N. A. ILIADIS, V. F. PERIRA, S. GRANVILLE, M. FINGER, P. A. HALDI, AND L. A. BARROSO, *Benchmarking of hydroelectric stochastic risk management models using financial indicators*, in 2006 IEEE Power Engineering Society General Meeting, 2006, pp. 1–8. (Cited on pp. 459, 498)
- [116] G. INFANGER AND D. P. MORTON, *Cut sharing for multistage stochastic linear programs with interstage dependency*, Math. Program., 75 (1996), pp. 241–256, <https://doi.org/10.1007/BF02592154>. (Cited on pp. 492, 493, 494)
- [117] J. KALLRATH, P. M. PARDALOS, S. REBENNACK, AND M. SCHEIDT, *Optimization in the Energy Industry*, Springer, 2009. (Cited on p. 456)
- [118] J. E. KELLEY, *The cutting-plane method for solving convex programs*, J. Soc. Ind. Appl. Math., 8 (1960), pp. 703–712, <https://doi.org/10.1137/0108053>. (Cited on pp. 429, 438, 507, 522)
- [119] A. KISZKA AND D. WOZABAL, *Stochastic dual dynamic programming for optimal power flow problems under uncertainty*, European J. Oper. Res., 321 (2025), pp. 814–836. (Cited on p. 459)
- [120] V. KOZMÍK, *On variance reduction of mean-CVaR Monte Carlo estimators*, Comput. Manag. Sci., 12 (2015), pp. 221–242, <https://doi.org/10.1007/s10287-014-0225-7>. (Cited on pp. 461, 479, 480, 481)

- [121] V. KOZMÍK AND D. P. MORTON, *Evaluating policies in risk-averse multi-stage stochastic programming*, Math. Program., 152 (2015), pp. 275–300, <https://doi.org/10.1007/s10107-014-0787-8>. (Cited on pp. 461, 464, 473, 474, 478, 479, 480, 481)
- [122] V. KRASKO AND S. REBENNACK, *Two-stage stochastic mixed-integer nonlinear programming model for post-wildfire debris flow hazard management: Mitigation and emergency evacuation*, European J. Oper. Res., 263 (2017), pp. 265–282, <https://doi.org/10.1016/j.ejor.2017.05.004>. (Cited on p. 422)
- [123] T. KRISTIANSEN, *Hydropower scheduling and financial risk management*, Optim. Control Appl. Methods, 27 (2006), pp. 1–18, <https://doi.org/10.1002/oca.768>. (Cited on p. 498)
- [124] G. LAN, *Complexity of stochastic dual dynamic programming*, Math. Program., 191 (2022), pp. 717–754, <https://doi.org/10.1007/s10107-020-01567-1>. (Cited on pp. 425, 440, 441, 447)
- [125] Y. LAN, Q. ZHAI, X. LIU, AND X. GUAN, *Fast stochastic dual dynamic programming for economic dispatch in distribution systems*, IEEE Trans. Power Syst., 38 (2022), pp. 3828–3840, <https://doi.org/10.1109/TPWRS.2022.3204065>. (Cited on p. 459)
- [126] V. LECLÈRE, P. CARPENTIER, J.-P. CHANCELIER, A. LENOIR, AND F. PACAUD, *Exact converging bounds for stochastic dual dynamic programming via Fenchel duality*, SIAM J. Optim., 30 (2020), pp. 1223–1250, <https://doi.org/10.1137/19M1258876>. (Cited on pp. 427, 451, 453, 454, 455, 481, 506)
- [127] V. LECLÈRE, F. PACAUD, T. RIGAUT, AND H. GERARD, *StochDynamicProgramming.jl*, <https://github.com/JuliaStochOpt/StochDynamicProgramming.jl>. (Cited on p. 461)
- [128] B. LEGAT, *StructDualDynProg.jl*, <https://github.com/JuliaStochOpt/StructDualDynProg.jl>. (Cited on p. 461)
- [129] C. LEMARÉCHAL, A. NEMIROVSKII, AND Y. NESTEROV, *New variants of bundle methods*, Math. Program., 69 (1995), pp. 111–147, <https://doi.org/10.1007/BF01585555>. (Cited on pp. 507, 522)
- [130] C. M. LEOCADIO, C. S. G. RICHÁ, M. Z. FORTES, V. H. FERREIRA, AND B. H. DIAS, *Economic evaluation of a CHP biomass plant using stochastic dual dynamic programming*, Electr. Eng., 102 (2020), pp. 2605–2615, <https://doi.org/10.1007/s00202-020-01056-w>. (Cited on p. 460)
- [131] L. M. M. LIMA, E. POPOVA, AND P. DAMIEN, *Modeling and forecasting of Brazilian reservoir inflows via dynamic linear models*, Int. J. Forecasting, 30 (2014), pp. 464–476, <https://doi.org/10.1016/j.ijforecast.2013.12.009>. (Cited on p. 457)
- [132] K. LINOWSKY AND A. B. PHILPOTT, *On the convergence of sampling-based decomposition algorithms for multistage stochastic programs*, J. Optim. Theory Appl., 125 (2005), pp. 349–366, <https://doi.org/10.1007/s10957-004-1842-z>. (Cited on p. 438)
- [133] R. P. LIU AND A. SHAPIRO, *Risk neutral reformulation approach to risk averse stochastic programming*, European J. Oper. Res., 286 (2020), pp. 21–31, <https://doi.org/10.1016/j.ejor.2020.01.060>. (Cited on pp. 457, 476, 477, 481, 482)
- [134] T. LOHMANN, A. S. HERING, AND S. REBENNACK, *Spatio-temporal hydro forecasting of multireservoir inflows for hydro-thermal scheduling*, European Journal of Operational Research, 255 (2016), pp. 243–258, <https://doi.org/10.1016/j.ejor.2016.05.011>. (Cited on pp. 457, 458, 494)
- [135] N. LÖHNDORF, *An empirical analysis of scenario generation methods for stochastic optimization*, European J. Oper. Res., 255 (2016), pp. 121–132, <https://doi.org/10.1016/j.ejor.2016.05.021>. (Cited on p. 462)
- [136] N. LÖHNDORF AND A. SHAPIRO, *Modeling time-dependent randomness in stochastic dual dynamic programming*, European J. Oper. Res., 273 (2019), pp. 650–661, <https://doi.org/10.1016/j.ejor.2018.08.001>. (Cited on pp. 423, 457, 458, 490, 491, 499)
- [137] N. LÖHNDORF AND D. WOZABAL, *Gas storage valuation in incomplete markets*, Eur. J. Oper. Res., 288 (2021), pp. 318–330, <https://doi.org/10.1016/j.ejor.2020.05.044>. (Cited on p. 499)
- [138] N. LÖHNDORF, D. WOZABAL, AND S. MINNER, *Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming*, Oper. Res., 61 (2013), pp. 810–823, <https://doi.org/10.1287/opre.2013.1182>. (Cited on p. 499)
- [139] R. LU, T. DING, B. QIN, J. MA, X. FANG, AND Z. DONG, *Multi-stage stochastic programming to joint economic dispatch for energy and reserve with uncertain renewable energy*, IEEE Trans. Sustain. Energy, 11 (2020), pp. 1140–1151, <https://doi.org/10.1109/TSTE.2019.2918269>. (Cited on p. 459)
- [140] M. E. P. MACEIRA AND J. M. DAMÁZIO, *Use of the PAR (p) model in the stochastic dual dynamic programming optimization scheme used in the operation planning of the Brazilian*



- hydropower system*, Probab. Eng. Inform. Sci., 20 (2006), pp. 143–156, <https://doi.org/10.1017/S0269964806060098>. (Cited on p. 457)
- [141] M. E. P. MACEIRA, V. S. DUARTE, D. D. J. PENNA, L. A. M. MORAES, AND A. C. G. MELO, *Ten years of application of stochastic dual dynamic programming in official and agent studies in Brazil—Description of the NEWAVE program*, in 16th Power Systems Computation Conference (PSCC), Glasgow, Scotland, Curran Associates, 2008, pp. 1416–1422. (Cited on p. 457)
- [142] M. E. P. MACEIRA, D. D. J. PENNA, A. L. DINIZ, R. J. PINTO, A. C. G. MELO, C. V. VASCONCELLOS, AND C. B. CRUZ, *Twenty years of application of stochastic dual dynamic programming in official and agent studies in Brazil: Main features and improvements on the NEWAVE model*, in 2018 Power Systems Computation Conference (PSCC), Dublin, Ireland, IEEE, 2018, pp. 374–380. (Cited on pp. 457, 458)
- [143] F. D. R. MACHADO, A. L. DINIZ, C. L. T. BORGES, AND L. C. BRANDÃO, *Asynchronous parallel stochastic dual dynamic programming applied to hydrothermal generation planning*, Electr. Pow. Syst. Res., 191 (2021), pp. 1–14, <https://doi.org/10.1016/j.epr.2020.106907>. (Cited on pp. 457, 525)
- [144] H. MACIAN-SORRIBES, A. TILMANT, AND M. PULIBO-VELAZQUEZ, *Improving operating policies of large-scale surface-groundwater systems through stochastic programming*, Water Resour. Res., 53 (2017), pp. 1407–1423, <https://doi.org/10.1002/2016WR019573>. (Cited on p. 460)
- [145] G. F. MARQUES AND A. TILMANT, *The economic value of coordination in large-scale multi-reservoir systems: The Parana River case*, Water Resour. Res., 49 (2013), pp. 7546–7557, <https://doi.org/10.1002/2013WR013679>. (Cited on p. 460)
- [146] Y. MBEUTCHA, M. GENDREAU, AND G. EMIEL, *Benefit of PARMA modeling for long-term hydroelectric scheduling using stochastic dual dynamic programming*, J. Water Res. Plan. Manag., 147 (2021), pp. 1–12, [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0001333](https://doi.org/10.1061/(ASCE)WR.1943-5452.0001333). (Cited on p. 457)
- [147] G. P. MCCORMICK, *Computability of global solutions to factorable nonconvex programs: Part I – Convex underestimating problems*, Math. Program., 10 (1976), pp. 147–175, <https://doi.org/10.1007/BF01580665>. (Cited on p. 506)
- [148] M. D. MCKAY, R. J. BECKMAN, AND W. J. CONOVER, *A comparison of three methods for selecting values of input variables in the analysis of output from a computer code*, Technometrics, 21 (1979), pp. 239–245, <https://doi.org/10.2307/1268522>. (Cited on p. 445)
- [149] B. MO, A. GJELSVIK, AND A. GRUNDT, *Integrated risk management of hydro power scheduling and contract management*, IEEE Trans. Power Syst., 16 (2001), pp. 216–221, <https://doi.org/10.1109/59.918289>. (Cited on p. 498)
- [150] B. MO, A. GJELSVIK, A. GRUNDT, AND K. KARESEN, *Optimisation of hydropower operation in a liberalised market with focus on price modelling*, in PowerTech Proceedings, IEEE Porto, 2001, pp. 1–6. (Cited on p. 497)
- [151] D. P. MORTON, *An enhanced decomposition algorithm for multistage stochastic hydroelectric scheduling*, Ann. Oper. Res., 64 (1996), pp. 211–235, <https://doi.org/10.1007/BF02187647>. (Cited on p. 520)
- [152] G. NANNICINI, E. TRAVERSI, AND R. WOLFLER CALVO, *A Benders squared ( $B^2$ ) framework for infinite-horizon stochastic linear programs*, Math. Program. Comput., 13 (2021), pp. 645–681, <https://doi.org/10.1007/s12532-020-00195-2>. (Cited on p. 513)
- [153] Y. NESTEROV, *Introductory Lectures on Convex Optimization: A Basic Course*, Springer, Boston, 2004. (Cited on pp. 441, 522)
- [154] N. NEWHAM, *Power System Investment Planning Using Stochastic Dual Dynamic Programming*, Ph.D. thesis, University of Canterbury, 2008. (Cited on p. 460)
- [155] A. PAPAVALIOU, Y. MOU, L. CAMBIER, AND D. SCIEUR, *Application of stochastic dual dynamic programming to the real-time dispatch of storage under renewable supply uncertainty*, IEEE Trans. Sustain. Energy, 9 (2018), pp. 547–558, <https://doi.org/10.1109/TSTE.2017.2748463>. (Cited on p. 459)
- [156] P. M. PARDALOS, S. REBENNACK, M. V. F. PEREIRA, N. A. ILIADIS, AND V. PAPPU, eds., *Handbook of Wind Power Systems*, Springer, 2013. (Cited on p. 456)
- [157] P. PARPAS, B. USTUN, AND Q. K. TRAN, *Importance sampling in stochastic programming: A Markov Chain Monte Carlo approach*, INFORMS J. Comput., 27 (2015), pp. 358–377, <https://doi.org/10.1287/ijoc.2014.0630>. (Cited on pp. 445, 446, 447, 461, 466)
- [158] M. V. F. PEREIRA, N. M. CAMPODÓNICO, AND R. KELMAN, *Application of Stochastic Dual DP and Extensions to Hydrothermal Scheduling*, Technical report, PSRI Technologies, Chesterfield, MO, 1999. (Cited on p. 498)

- [159] M. V. F. PEREIRA AND L. M. V. G. PINTO, *Stochastic optimization of a multireservoir hydroelectric system: A decomposition approach*, Water Resour. Res., 21 (1985), pp. 779–792, <https://doi.org/10.1029/WR021i006p00779>. (Cited on pp. 445, 456)
- [160] M. V. F. PEREIRA AND L. M. V. G. PINTO, *Multi-stage stochastic optimization applied to energy planning*, Math. Program., 52 (1991), pp. 359–375, <https://doi.org/10.1007/BF01582895>. (Cited on pp. 418, 448, 449, 456)
- [161] M. V. F. PEREIRA, *Optimal scheduling of hydrothermal systems - An overview*, in IFAC Symposium on Planning and Operation of Electric Energy Systems., Rio de Janeiro, Brazil, 1985, pp. 1–6. (Cited on p. 490)
- [162] S. J. PEREIRA-CARDENAL, B. MO, A. GJELSVIK, N. D. RIEGELS, K. ARNBJERG-NIELSEN, AND P. BAUER-GOTTWEIN, *Joint optimization of regional water-power systems*, Adv. Water Resour., 92 (2016), pp. 200–207, <https://doi.org/10.1016/j.advwatres.2016.04.004>. (Cited on p. 460)
- [163] L. PFEIFFER, R. APPARIGLIATO, AND S. AUCHAPT, *Two Methods of Pruning Benders' Cuts and their Application to the Management of a Gas Portfolio*, preprint, [http://www.optimization-online.org/DB\\_FILE/2012/11/3683.pdf](http://www.optimization-online.org/DB_FILE/2012/11/3683.pdf), 2012. (Cited on p. 515)
- [164] G. C. PFLUG, *Some Remarks on the Value-at-Risk and the Conditional Value-at-Risk*, in Probabilistic Constrained Optimization: Methodology and Applications, S. P. Uryasev, ed., Kluwer Academic Publishers, 2000, pp. 272–281. (Cited on p. 469)
- [165] G. C. PFLUG AND A. PICHLER, *Time-consistent decisions and temporal decomposition of coherent risk functionals*, Math. Oper. Res., 41 (2016), pp. 682–699, <https://doi.org/10.1287/moor.2015.0747>. (Cited on pp. 471, 483)
- [166] G. C. PFLUG AND A. RUSZCZYŃSKI, *Measuring risk for income streams*, Comput. Optim. Appl., 32 (2005), pp. 161–178, <https://doi.org/10.1007/s10589-005-2058-3>. (Cited on p. 483)
- [167] A. B. PHILPOTT AND V. L. DE MATOS, *Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion*, European J. Oper. Res., 218 (2012), pp. 470–483, <https://doi.org/10.1016/j.ejor.2011.10.056>. (Cited on pp. 457, 459, 474, 477, 482, 496)
- [168] A. B. PHILPOTT, V. L. DE MATOS, AND E. FINARDI, *On solving multistage stochastic programs with coherent risk measures*, Oper. Res., 61 (2013), pp. 957–970, <https://doi.org/10.1287/opre.2013.1175>. (Cited on pp. 451, 452, 457, 459, 473, 474, 475, 476, 477, 481, 498)
- [169] A. B. PHILPOTT, V. L. DE MATOS, AND L. KAPELEVICH, *Distributionally robust SDDP*, Comput. Manage. Sci., 15 (2018), pp. 431–454, <https://doi.org/10.1007/s10287-018-0314-0>. (Cited on pp. 440, 457, 485, 487, 488, 489)
- [170] A. B. PHILPOTT, F. WAHID, AND F. BONNANS, *MIDAS: A mixed integer dynamic approximation scheme*, Math. Program., 181 (2020), pp. 19–50, <https://doi.org/10.1007/s10107-019-01368-1>. (Cited on pp. 459, 460, 507)
- [171] A. B. PHILPOTT AND Z. GUAN, *On the convergence of stochastic dual dynamic programming and related methods*, Oper. Res. Lett., 36 (2008), pp. 450–455, <https://doi.org/10.1016/j.orl.2008.01.013>. (Cited on pp. 438, 516, 517)
- [172] J. PINA, A. TILMANT, AND P. CÔTÉ, *Optimizing multireservoir system operating policies using exogenous hydrologic variables*, Water Resour. Res., 53 (2017), pp. 9845–9859, <https://doi.org/10.1002/2017WR021701>. (Cited on p. 457)
- [173] R. J. PINTO, C. L. T. BORGES, AND M. E. P. MACEIRA, *An efficient parallel algorithm for large scale hydrothermal system operation planning*, IEEE Trans. Power Syst., 28 (2013), pp. 4888–4896, <https://doi.org/10.1109/TPWRS.2012.2236654>. (Cited on pp. 457, 525)
- [174] B. T. POLYAK, *Minimization of unsmooth functionals*, USSR Comput. Math. Math. Phys., 9 (1969), pp. 14–29, [https://doi.org/10.1016/0041-5553\(69\)90061-5](https://doi.org/10.1016/0041-5553(69)90061-5). (Cited on p. 507)
- [175] H. POORSEPAHY-SAMIAN, V. ESPANMANESH, AND B. ZAHRAIE, *Improved inflow modeling in stochastic dual dynamic programming*, J. Water Res. Plan. Manag., 142 (2016), pp. 1–10, [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000713](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000713). (Cited on p. 458)
- [176] W. B. POWELL, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Wiley-Interscience, 2007. (Cited on pp. 419, 442)
- [177] W. B. POWELL, *Reinforcement Learning and Stochastic Optimization – A Unified Framework for Sequential Decisions*, John Wiley & Sons, 2022. (Cited on pp. 418, 419, 421, 424)
- [178] G. PRITCHARD, *Stochastic inflow modeling for hydropower scheduling problems*, European J. Oper. Res., 246 (2015), pp. 496–504, <https://doi.org/10.1016/j.ejor.2015.05.022>. (Cited on p. 458)
- [179] PSR SDDP, *Stochastic hydrothermal dispatch with network restrictions*, <https://www.psr-inc.com/software/sddp.html>? (accessed April 27, 2021). (Cited on p. 461)
- [180] A. R. QUEIROZ AND D. P. MORTON, *Sharing cuts under aggregated forecast when decomposing multi-stage stochastic programs*, Oper. Res. Lett., 41 (2013), pp. 311–316, <https://doi.org/10.1016/j.orl.2013.03.003>. (Cited on pp. 491, 492, 494)

- [181] F. QUEZADA, C. GICQUEL, AND S. KEDAD-SIDHOUM, *Combining polyhedral approaches and stochastic dual dynamic integer programming for solving the uncapacitated lot-sizing problem under uncertainty*, INFORMS J. Comput., 34 (2022), pp. 1024–1041, <https://doi.org/10.1287/ijoc.2021.1118>. (Cited on p. 509)
- [182] H. RAHIMIAN AND S. MEHROTRA, *Distributionally Robust Optimization: A Review*, preprint, <https://arxiv.org/pdf/1908.05659.pdf>, 2019. (Cited on p. 486)
- [183] L. RASO, P.-O. MALATERRE, AND J.-C. BADER, *Effective streamflow process modeling for optimal reservoir operation using stochastic dual dynamic programming*, J. Water Res. Plan. Manag., 143 (2017), pp. 1–11, [https://doi.org/10.1061/\(ASCE\)WR.1943-5452.0000746](https://doi.org/10.1061/(ASCE)WR.1943-5452.0000746). (Cited on p. 458)
- [184] S. REBENNACK, *Generation expansion planning under uncertainty with emissions quotas*, Electric Power Systems Res., 114 (2014), pp. 78–85. (Cited on pp. 459, 460, 512)
- [185] S. REBENNACK, *Combining sampling-based and scenario-based nested Benders decomposition methods: Application to stochastic dual dynamic programming*, Math. Program., 156 (2016), pp. 343–389. (Cited on pp. 419, 423, 440, 443, 457, 459, 491, 492, 494, 500)
- [186] S. REBENNACK, B. FLACH, M. V. F. PEREIRA, AND P. M. PARDALOS, *Stochastic hydro-thermal scheduling under CO<sub>2</sub> emissions constraints*, IEEE Trans. Power Syst., 27 (2012), pp. 58–68, <https://doi.org/10.1109/TPWRS.2011.2140342>. (Cited on pp. 459, 512)
- [187] S. REBENNACK, N. A. ILIADIS, M. V. F. PEREIRA, AND P. M. PARDALOS, *Electricity and CO<sub>2</sub> emissions system prices modeling and optimization*, in PowerTech, 2009 IEEE Bucharest, IEEE, 2009, pp. 1–6. (Cited on p. 459)
- [188] S. REBENNACK, P. M. PARDALOS, M. V. F. PEREIRA, AND N. A. ILIADIS, EDs., *Handbook of Power Systems I*, Springer, 2010. (Cited on p. 456)
- [189] S. REBENNACK, P. M. PARDALOS, M. V. F. PEREIRA, AND N. A. ILIADIS, EDs., *Handbook of Power Systems II*, Springer, 2010. (Cited on p. 456)
- [190] M. RESENER, S. REBENNACK, P. M. PARDALOS, AND S. HAFNER, eds., *Handbook of Optimization in Electric Power Distribution Systems*, Springer, 2020. (Cited on p. 456)
- [191] R. T. ROCKAFELLAR AND S. URYASEV, *Optimization of conditional value-at-risk*, J. Risk, 2 (2000), pp. 21–42, <https://doi.org/10.21314/JOR.2000.038>. (Cited on pp. 468, 474)
- [192] R. T. ROCKAFELLAR AND R. J. B. WETS, *Variational Analysis*, Springer, Berlin Heidelberg, 2009. (Cited on p. 453)
- [193] A. W. ROSEMBERG, A. STREET, J. DIAS GARCIA, D. M. VALLADÃO, T. SILVA, AND O. DOWSON, *Assessing the cost of network simplifications in long-term hydrothermal dispatch planning models*, IEEE Trans. Sustain. Energy, 13 (2022), pp. 196–206, <https://doi.org/10.1109/TSTE.2021.3106810>. (Cited on pp. 459, 506)
- [194] T. A. ROTTING AND A. GJELSVIK, *Stochastic dual dynamic programming for seasonal scheduling in the Norwegian power system*, IEEE Trans. Power Syst., 7 (1992), pp. 273–279, <https://doi.org/10.1109/59.141714>. (Cited on p. 457)
- [195] C. ROUGÉ AND A. TILMANT, *Using stochastic dual dynamic programming in problems with multiple near-optimal solutions*, Water Resour. Res., 52 (2016), pp. 4151–4163, <https://doi.org/10.1002/2016WR018608>. (Cited on p. 460)
- [196] C. ROUGÉ, A. TILMANT, B. ZAITCHIK, A. DEZFULI, AND M. SALMAN, *Identifying key water resource vulnerabilities in data-scarce transboundary river basins*, Water Resour. Res., 54 (2018), pp. 5264–5281, <https://doi.org/10.1029/2017WR021489>. (Cited on p. 460)
- [197] B. RUDLOFF, A. STREET, AND D. M. VALLADÃO, *Time consistency and risk averse dynamic decision models: Definition, interpretation and practical consequences*, European J. Oper. Res., 234 (2014), pp. 743–750, <https://doi.org/10.1016/j.ejor.2013.11.037>. (Cited on p. 471)
- [198] A. RUSZCZYŃSKI, *Decomposition methods in stochastic programming*, Math. Program., 79 (1997), pp. 333–353, <https://doi.org/10.1007/BF02614323>. (Cited on p. 522)
- [199] A. RUSZCZYŃSKI, *Risk-averse dynamic programming for Markov decision processes*, Math. Program., 125 (2010), pp. 235–261, <https://doi.org/10.1007/s10107-010-0393-3>. (Cited on p. 467)
- [200] A. RUSZCZYŃSKI AND A. SHAPIRO, *Conditional risk mappings*, Math. Oper. Res., 31 (2006), pp. 544–561, <https://doi.org/10.1287/moor.1060.0204>. (Cited on p. 470)
- [201] A. RUSZCZYŃSKI AND A. SWIETANOWSKI, *Accelerating the regularized decomposition method for two stage stochastic linear problems*, European J. Oper. Res., 101 (1997), pp. 328–342, [https://doi.org/10.1016/S0377-2217\(96\)00401-8](https://doi.org/10.1016/S0377-2217(96)00401-8). (Cited on p. 522)
- [202] A. SHAPIRO, *Inference of statistical bounds for multistage stochastic programming problems*, Math. Methods Oper. Res., 58 (2003), pp. 57–68, <https://doi.org/10.1007/s001860300280>. (Cited on p. 463)

- [203] A. SHAPIRO, *On complexity of multistage stochastic programs*, Oper. Res. Lett., 34 (2005), pp. 1–8, <https://doi.org/10.1016/j.orl.2005.02.003>. (Cited on p. 464)
- [204] A. SHAPIRO, *Stochastic programming approach to optimization under uncertainty*, Math. Program., 112 (2008), pp. 183–220, <https://doi.org/10.1007/s10107-006-0090-4>. (Cited on p. 464)
- [205] A. SHAPIRO, *Analysis of stochastic dual dynamic programming method*, European J. Oper. Res., 209 (2011), pp. 63–72, <https://doi.org/10.1016/j.ejor.2010.08.007>. (Cited on pp. 438, 449, 462, 463, 464, 468, 473, 474, 475, 477, 479)
- [206] A. SHAPIRO, *Tutorial on risk neutral, distributionally robust and risk averse multistage stochastic programming*, European J. Oper. Res., 288 (2021), pp. 1–13, <https://doi.org/10.1016/j.ejor.2020.03.065>. (Cited on pp. 467, 471, 472, 486)
- [207] A. SHAPIRO AND Y. CHENG, *Dual bounds for periodical stochastic programs*, Oper. Res., 71 (2022), pp. 120–128, <https://doi.org/10.1287/opre.2021.2245>. (Cited on p. 455)
- [208] A. SHAPIRO, D. DENTCHEVA, AND A. RUSZCZYŃSKI, *Lectures on Stochastic Programming: Modeling and Theory*, 2nd ed., SIAM, 2014, <https://doi.org/10.1137/1.9781611973433>. (Cited on pp. 423, 426, 427, 448, 462, 463, 467, 469, 470, 471, 472, 473, 499)
- [209] A. SHAPIRO AND L. DING, *Upper Bound for Optimal Value of Risk Averse Multistage Problems*, Technical report, Georgia Tech, 2016. (Cited on pp. 479, 481)
- [210] A. SHAPIRO AND L. DING, *Periodical multistage stochastic programs*, SIAM J. Optim., 30 (2020), pp. 2083–2102, <https://doi.org/10.1137/19M129406X>. (Cited on pp. 457, 513)
- [211] A. SHAPIRO, W. TEKAYA, J. P. DA COSTA, AND M. P. SOARES, *Risk neutral and risk averse stochastic dual dynamic programming method*, European J. Oper. Res., 224 (2013), pp. 375–391, <https://doi.org/10.1016/j.ejor.2012.08.022>. (Cited on pp. 457, 458, 459, 473, 474, 475, 477, 491, 516)
- [212] A. SHAPIRO, W. TEKAYA, M. P. SOARES, AND J. P. DA COSTA, *Worst-case-expectation approach to optimization under uncertainty*, Oper. Res., 61 (2013), pp. 1435–1449, <https://doi.org/10.1287/opre.2013.1229>. (Cited on pp. 457, 515)
- [213] P. SHINDE, I. KOUVELIOTIS-LYSIKATOS, AND M. AMELIN, *Multistage stochastic programming for VPP trading in continuous intraday electricity markets*, IEEE Trans. Sustain. Energy, 13 (2022), pp. 1–12, <https://doi.org/10.1109/TSTE.2022.3144022>. (Cited on p. 460)
- [214] M. SIDDIG AND Y. SONG, *Adaptive partition-based SDDP algorithms for multistage stochastic linear programming with fixed recourse*, Comput. Optim. Appl., 81 (2022), pp. 201–250, <https://doi.org/10.1007/s10589-021-00323-1>. (Cited on pp. 440, 526)
- [215] M. P. SOARES, A. STREET, AND D. M. VALLADÃO, *On the solution variability reduction of stochastic dual dynamic programming applied to energy planning*, European J. Oper. Res., 258 (2017), pp. 743–760, <https://doi.org/10.1016/j.ejor.2016.08.068>. (Cited on pp. 457, 459, 495)
- [216] A. SOROKIN, S. REBENNACK, P. M. PARDALOS, N.A. ILIADIS, AND M. V. F. PEREIRA, EDS., *Handbook of Networks in Power Systems I*, Springer Science & Business Media, 2012. (Cited on p. 456)
- [217] A. SOROKIN, S. REBENNACK, P. M. PARDALOS, N.A. ILIADIS, AND M. V. F. PEREIRA, EDS., *Handbook of Networks in Power Systems II*, Springer Science & Business Media, 2012. (Cited on p. 456)
- [218] G. STEEGER, T. LOHMANN, AND S. REBENNACK, *Strategic bidding for a price-maker hydroelectric producer: Stochastic dual dynamic programming and Lagrangian relaxation*, IIEE Transactions, 50 (2018), pp. 929–942, <https://doi.org/10.1080/24725854.2018.1461963>. (Cited on pp. 494, 507)
- [219] G. STEEGER AND S. REBENNACK, *Strategic bidding for multiple price-maker hydroelectric producers*, IIE Transactions, 47 (2015), pp. 1013–1031, <https://doi.org/10.1080/0740817X.2014.1001928>. (Cited on p. 460)
- [220] G. STEEGER AND S. REBENNACK, *Dynamic convexification within nested Benders decomposition using Lagrangian relaxation: An application to the strategic bidding problem*, European J. Oper. Res., 257 (2017), pp. 669–686, <https://doi.org/10.1016/j.ejor.2016.08.006>. (Cited on pp. 457, 460, 507, 508)
- [221] G. M. STEEGER, *Strategic Bidding for Price-maker Hydroelectric Producers*, Ph.D. thesis, Colorado School of Mines, 2014. (Cited on p. 460)
- [222] A. STREET, L. A. BARROSO, B. FLACH, M. V. PEREIRA, AND S. GRANVILLE, *Risk constrained portfolio selection of renewable sources in hydrothermal electricity markets*, IEEE Trans. Power Syst., 24 (2009), pp. 1136–1144, <https://doi.org/10.1109/TPWRS.2009.2022981>. (Cited on p. 460)
- [223] A. STREET, A. BRIGATTO, AND D. M. VALLADÃO, *Co-optimization of energy and ancillary services for hydrothermal operation planning under a general security criterion*, IEEE Trans.

- Power Syst., 32 (2017), pp. 4914–4923, <https://doi.org/10.1109/TPWRS.2017.2672555>. (Cited on pp. 457, 459)
- [224] X. A. SUN AND A. J. CONEJO, *Robust Optimization in Electric Energy Systems*, Internat. Ser. Oper. Res. Manag. Sci., Springer, 2022. (Cited on p. 487)
- [225] S. THEVENIN, Y. ADULYASAK, AND J.-F. CORDEAU, *Stochastic dual dynamic programming for multiechelon lot sizing with component substitution*, INFORMS J. Comput., 34 (2022), pp. 3151–3169, <https://doi.org/10.1287/ijoc.2022.1215>. (Cited on p. 461)
- [226] F. S. THOMÉ, *Representação de não-convergência no planejamento da operação hidrotérmica utilizando PDDE*, Ph.D. thesis, COPPE-UFRJ, 2013. (Cited on pp. 506, 508)
- [227] A. TILMANT, L. BEEVERS, AND B. MUYUNDA, *Restoring a flow regime through the coordinated operation of a multireservoir system: The case of the Zambezi River basin*, Water Resour. Res., 46 (2010), pp. 1–11, <https://doi.org/10.1029/2009WR008897>. (Cited on p. 460)
- [228] A. TILMANT AND R. KELMAN, *A stochastic approach to analyze trade-offs and risks associated with large-scale water resources systems*, Water Resour. Res., 43 (2007), pp. 1–11, <https://doi.org/10.1029/2006WR005094>. (Cited on p. 460)
- [229] A. TILMANT AND W. KINZELBACH, *The cost of noncooperation in international river basins*, Water Resour. Res., 48 (2012), pp. 1–12, <https://doi.org/10.1029/2011WR011034>. (Cited on p. 460)
- [230] A. TILMANT, J. LETTANY, AND R. KELMAN, *Hydrological risk assessment in the Euphrates-Tigris river basin: A stochastic dual dynamic programming approach*, Water Int., 32 (2009), pp. 294–309, <https://doi.org/10.1080/02508060708692208>. (Cited on p. 460)
- [231] A. TILMANT, D. PINTÉ, AND Q. GOOR, *Assessing marginal water values in multipurpose multireservoir systems via stochastic programming*, Water Resour. Res., 44 (2008), pp. 1–17, <https://doi.org/10.1029/2008WR007024>. (Cited on p. 460)
- [232] D. M. VALLADÃO, T. SILVA, AND M. POGGI, *Time-consistent risk-constrained dynamic portfolio optimization with transactional costs and time-dependent returns*, Ann. Oper. Res., 282 (2019), pp. 379–405, <https://doi.org/10.1007/s10479-018-2991-z>. (Cited on pp. 460, 461)
- [233] W. VAN ACKOOIJ, W. DE OLIVEIRA, AND Y. SONG, *On level regularization with normal solutions in decomposition methods for multistage stochastic programming problems*, Comput. Optim. Appl., 74 (2019), pp. 1–42, <https://doi.org/10.1007/s10589-019-00104-x>. (Cited on pp. 457, 523)
- [234] W. VAN ACKOOIJ AND X. WARIN, *On conditional cuts for stochastic dual dynamic programming*, EURO J. Comput. Optim., 8 (2020), pp. 173–199, <https://doi.org/10.1007/s13675-020-00123-y>. (Cited on pp. 443, 459, 502, 503)
- [235] N. VAN DER LAAN AND W. ROMELJNDERS, *A converging Benders' decomposition algorithm for two-stage mixed-integer recourse models*, Operations Research, 72 (2023), pp. 2190–2214, <https://doi.org/10.1287/opre.2021.2223>. (Cited on p. 510)
- [236] R. M. VAN SLYKE AND R. WETS, *L-shaped linear programs with applications to optimal control and stochastic programming*, SIAM J. Appl. Math., 17 (1969), pp. 638–663, <https://doi.org/10.1137/0117061>. (Cited on pp. 419, 443)
- [237] F. WAHID, *River Optimisation: Short-Term Hydro-Bidding under Uncertainty*, Optimization and Control, Ph.D. thesis, University of Auckland, New Zealand, 2017. (Cited on pp. 457, 460, 497)
- [238] W. W.-G. YEH, *Reservoir management and operations models: A state-of-the-art review*, Water Resour. Res., 21 (1985), pp. 1797–1818, <https://doi.org/10.1029/WR021i012.p01797>. (Cited on p. 420)
- [239] C. ZHANG, *Nearly unbiased variable selection under minimax concave penalty*, Ann. Statist., 38 (2010), pp. 894–942, <https://doi.org/10.1214/09-AOS729>. (Cited on p. 523)
- [240] S. ZHANG AND X. A. SUN, *Stochastic dual dynamic programming for multistage stochastic mixed-integer nonlinear optimization*, Math. Program., 196 (2022), pp. 935–985, <https://doi.org/10.1007/s10107-022-01875-8>. (Cited on pp. 439, 440, 441, 447, 450, 509, 512)
- [241] Q. P. ZHENG, S. REBENNACK, P. M. PARDALOS, M. V. F. PEREIRA, AND N. A. ILIADIS, EDS., *Handbook of CO<sub>2</sub> in Power Systems*, Springer Science & Business Media, 2012. (Cited on p. 456)
- [242] X. ZHU, M. G. GENTON, Y. GU, AND L. XIE, *Space-time wind speed forecasting for improved power system dispatch*, TEST, 23 (2014), pp. 1–25, <https://doi.org/10.1007/s11749-014-0351-0>. (Cited on p. 498)
- [243] J. ZOU, S. AHMED, AND X. A. SUN, *Stochastic dual dynamic integer programming*, Math. Program., 175 (2019), pp. 461–502, <https://doi.org/10.1007/s10107-018-1249-5>. (Cited on pp. 459, 460, 461, 508, 509)
- [244] L. ZÉPHYR AND C. L. ANDERSON, *Stochastic dynamic programming approach to managing power system uncertainty with distributed storage*, Comput. Manag. Sci., 15 (2018), pp. 87–110, <https://doi.org/10.1007/s10287-017-0297-2>. (Cited on p. 460)