TECHNISCHE UNIVERSITÄT MÜNCHEN

Fakultät für Wirtschaftswissenschaften Center for Energy Markets

Essays on Stochastic Optimization with Applications to Energy Storage Valuation

Gonçalo de Almeida Terça

Vollständiger Abdruck der von der Fakultät für Wirtschaftswissenschaften der Technischen Universität München zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften (Dr. rer. pol.) genehmigten Dissertation.

Vorsitzender:	Prof. Dr. Stefan Minner
Prüfer der Dissertation:	1) Prof. Dr. David Wozabal
	2) Prof. Dr. Sebastian Schwenen

Die Dissertation wurde am 11.09.2020 bei der Technischen Universität München eingereicht und durch die Fakultät für Wirtschaftswissenschaften am 15.12.2020 angenommen.

Abstract

In light of the energy transition towards a greener generation mix, power grids are witnessing increasing shares of variable renewable sources, responsible for bringing uncertainty to the electricity supply. Stochastic optimization and energy storage represent respectively a framework and an asset to respond to this uncertainty. This dissertation focuses on issues revolving around both topics. The first essay tackles the topic of parametric sensitivity analysis in multistage stochastic linear problems. It develops the theoretical background and the methodology with which to calculate first-order derivatives of the value function with respect to parameters of the problem. It extends the classical envelope theorem to stochastic optimization, and offers a sampling method with which to calculate derivatives and perform sensitivity analysis in such problems. The second and third essays focus on the issue of energy storage. The second essay presents a framework with which to measure the added economic benefits of pairing energy storage to a variable renewable source. The essay challenges the popular belief of joint planning and proposes that in a spot market with a real-time adjustment market, there is no benefit in joint planning compared to an independent use of both assets. The third essay focuses on the operation of energy storage in the German secondary control reserve. It offers a mixed-integer bilinear description of optimal bidding in this market and analyzes the profit potential of storage as a function of its duration. It shows that this potential saturates with increasing duration, reaching a maximum at a duration value of 13 hours.

Acknowledgements

First and foremost, I would like to thank my supervisor, Prof. Dr. David Wozabal, for his mentorship and guidance throughout my years as a PhD candidate. His supervision and support cannot be overstated, having been the source of many insightful discussions over the years and given me invaluable opportunities to learn and grow as a researcher.

I would like to thank my fellow colleagues at the Center for Energy Markets, Prof. Dr. Sebastian Schwenen, Adriana Kiszka, Moritz Bohland and Vadim Gorski, for the many helpful conversations and support regarding my research.

I would like to extend my gratitude to the colleagues that assisted and helped in all logistics of everyday life as a teaching assistant at TUM, including Daniel Bias, Dr. Patrick Bielstein, Karin Papavlassopoulos and Henriett Kakonyi.

In addition to the aforementioned, I would like to acknowledge the many colleagues that made the TUM School of Management a wonderful place to work, including the colleagues over at the Department of Financial Management and Capital Markets and the Faculty administrative support staff. Special thanks to Dr. Vítor Gonçalves de Azevedo and Dr. Marcin Bartkowiak for their camaraderie and good spirits.

Finally, on a more personal note, I would like to thank my parents, Paulo and Cristina, and my wife, Irene, for their love and unwavering support.

Contents

List of Figures

List of Tables

List of Statements

Nomenclature

Chapter 3

t Decision stage, $t = 1, \ldots, T$

d Day in the planning period, $d = 1, \ldots, D$

$\alpha(d)$ First stage on day d

p_t^D Day-ahead market price for time t (in €/MWh)

p_t^R Real-time market price at time t (in €/MWh)

l_t Storage level at time t (in MWh)

$\bar{l}$ Maximum storage capacity (in MWh)

i_t Energy injection to storage at time t (in MW)

$\bar{i}$ Maximum injection (in MW)

w_t Energy withdrawal from storage at time t (in MW)

$\bar{w}$ Maximum withdrawal (in MW)

g_t Random production of the VRES at time t (in MW)

s_t^D Day-ahead trade of the storage for time t (in MW)

s_t^R Real-time trade of the storage at time t (in MW)

r_t^D Day-ahead trade of the VRES for time t (in MW)

r_t^R Real-time trade of the VRES at time t (in MW)

$V(l_T)$ Value of stored energy at the end time period T

x^J Optimal decision for problem V^J

x^S Optimal decision for problem V^S

Chapter 4

Parameters

t Time period, $t = 1, \ldots, T$

S_{Max} Maximum storage capacity

j up/down regulation index, $j \in \{\text{positive}, \text{negative}\}$

$\underline{P}_{cap}^{j}$ Capacity price soft lower bound for product j

$\bar{P}_{cap}^{j}$ Capacity price soft upper bound for product j

$\bar{V}_{cap}^{j}$ Capacity volume upper bound for product j

$\underline{P}_{t}^{j}$ Dispatch price soft lower bound for product j at time t

$\bar{P}_{t}^{+}$ Dispatch price soft upper bound for product j at time t

α Constant conversion unit factor between power and energy. Equal to $\frac{1}{3600}$ for 1 second time resolution

Data

$P_{cap}^{cut,j}$ Capacity cutoff price for product j

$P_{t}^{cut,j}$ Dispatch cutoff price for product j at time t

D_{t}^{j} Demand volume for product j at time t

Variables

P_{cap}^{+} Capacity price bid for positive balancing energy

P_{cap}^{-} Capacity price bid for negative balancing energy

V_{cap}^{+} Capacity volume bid for positive balancing energy

V_{cap}^{-} Capacity volume bid for negative balancing energy

P_{t}^{+} Dispatch price bid for positive balancing energy

P_{t}^{-} Dispatch price bid for negative balancing energy

V_{t}^{+} Dispatch volume for positive balancing energy at time t

V_{t}^{-} Dispatch volume for negative balancing energy at time t

s_{t} Storage level at time t

w_{cap}^{j} Product term $P_{cap}^{j} V_{cap}^{j}$

w_{t}^{j} Product term $P_{t}^{j} V_{t}^{j}$

Chapter 1

Introduction

1.1 Motivation

The widespread deregulation of energy markets, coupled with the increasing share of renewables in the grid has brought about considerable change in our power markets. Ever since the Energy Policy Act was passed into law in 1992 (US Congress 1992), the electricity market in the United States has been opened to competition. Until then, state-regulated utilities held the monopoly of supply and with it much control over the price of electricity. Deregulation did not come overnight, with some states taking the first steps to deregulated power markets in the late 1990s and early 2000s. As of 2018, 17 states have deregulated markets to some degree, most in the interconnected northeast, or in heavily populated regions such as California or Texas (American Public Power Association 2018). In the European Union, the liberalization of electricity markets occurred during the same period of the late 1990s and 2000s (European Parliament 2016). Concurrently, an increasing number of countries around the world are opening their markets to competition. Along with allowing third parties to access wholesale as well as retail markets, the liberalization of the electricity market ensures the unbundling of generation, transmission and distribution.

Additionally, environmental concerns and the realities of climate change have prompted the emergence and growth of renewable technologies around the world. In particular, wind and solar photovoltaic sources have seen exponential growth rates averaging respectively a 23% and 46% annual increase in installed capacity over the last 20 years (British Petroleum 2019). Nowadays, they are indisputable forces in the electricity sector, accounting for a combined 6.1% share of the world generation energy mix (International Energy Agency 2019a). In the OECD members this aggregate value is even larger, where in 2018 10% of all electricity generation came from wind and solar (International Energy Agency 2018). In Germany, wind became the main source of electricity in 2019, with a share of nearly 25% of all generated electricity (Fraunhofer ISE 2020). During the same year in Germany, solar accounted for 9% of all generated electricity.

Although the development of these renewable sources represent a great strength in the transition towards greener power markets, the increasing share of wind and solar in the electricity generation mix have introduced greater uncertainty in the power supply. Their production is inherently uncontrollable and is subject to the uncertainty of weather, namely the intensity of the wind and the amount of sunlight. The intermittency of these sources and their growing share in the generation mix have brought about considerable challenges to the markets. Owners of these sources in particular must employ measures and be equipped with appropriate modelling tools with which to manage their assets. Not only do market participants already have to deal with market uncertainties in the form of prices or demand, but they must now also contend with the uncertainty of production.

Stochastic optimization is a framework that can be adequate to tackle challenges brought about by these uncertainties. With the goal of determining the best course of action in an uncertain environment, the field of optimization under uncertainty lends itself to solve problems that arise in the market. Such problems may involve individual agents deciding on bidding strategies to maximize their performance, utilities minimizing their operating costs while serving their customers, or grid system operators attempting to minimize operating costs while guaranteeing the stability of the grid.

Another response to the intermittency of renewables is the development of energy storage. Storage technologies intended for grid operations are predicted to become an important sector and play a greater role in the energy mix of power markets. In the five years between 2013 and 2018, deployment in the sector has grown twelve-fold to an annual deployment of 1.2 GW in 2018 (International Energy Agency 2019b). Overall deployment in grid-scale but also in behind-the-meter storage is expected to increase at an exponential rate in the next few years (Bloomberg New Energy Finance 2019).

Storage technologies, as they can both deliver and consume electricity, can be used as a vehicle to regulate the amount of energy in the grid. When there is too much energy in the grid, storage can extract the excess, and when there is a shortage of power, storage can be called upon to provide it. This becomes especially important when there is a lot of uncertainty in the power supply such as when the share of renewables is high.

In fact, the flexibility brought about by storage is such an attractive notion, that this has led market participants to not only include storage to their generation portfolio, but also to pair it to their existing power plants, some of which may include intermittent renewables, and do joint planning, often by way of participating in one or several markets, with one bidding strategy. There are a few challenges to address when considering storage, such as determining in what way can storage better assist power markets. In the present document, we address two such cases.

In this thesis, the work is focused on these two aspects: first, there was a focus on optimization methods, a topic that falls in the general area of operations research; second, the thesis targets the study of storage and its role in power markets.

Regarding the former, we tackle the issue of parameter sensitivity analysis in multistage stochastic optimization problems. Given that decision variables in such problems are random, would it make sense to talk about the sensitivity of the objective function with respect to parameters in the problem? Parameters that could be deterministic or uncertain. And if it did, which is to say if they can be proven to exist, would there be a framework that would allow us to compute derivatives quantifying the sensitivity of the objective function to these parameters?

In management science, there exists the challenge to take optimal decisions. A natural problem that accompanies and extends this is to consider how sensitive are the decisions we make with respect to the parameters of the problem. What happens when the parameter is nudged just a little bit; what effect does this have on the objective function, and on the decision policy? This is the topic of the first paper in this thesis, and the object of this first part on optimization methods.

Concerning the second focus of the thesis on energy storage, we address the issue of whether the market-wide economic value that storage offers as a stabilization mechanism for market supply uncertainty also holds true for individual agents in power markets. Additionally, we address the role of storage in balancing markets and discuss what storage unit dimensions are appropriate to operate on those markets.

The first of the two essays on storage challenges the notion that the pairing of storage to variable renewable energy sources such as wind farms or photovoltaic solar plants, in what we shall refer to as *joint planning* from now on, brings added economic benefits to the holder of both assets. The essay instead makes the argument that under special conditions there are no added benefits in joint planning when compared to separate planning, which is the term we shall use to describe the independent and uncoupled use of both assets.

The second essay on storage looks into the problem of placing optimal bids on the German secondary reserve market (SRL) with a storage unit operating exclusively in that market. It models the optimization problem as a mixed-integer bilinear program, to which is then applied a relaxation method to render it an easier to solve mixed-integer linear problem (MILP). The essay focuses on the duration of a storage unit, which is the ratio between its capacity and maximum power output, and how it affects the optimal value of the bidding problem on the SRL.

1.2 Findings and Contributions

1.2.1 Optimization methods and Parametric Sensitivity Analysis

The idea for the first paper came about from trying to address a very specific issue, which was to understand whether it would be possible to calculate parameter sensitivities of the value of derivative contracts, which in the financial sector are colloquially known as Greeks, using a multistage stochastic optimization framework.

1.2.1.1 Greeks, Black-Scholes and Benchmarks

Greeks in Finance are sensitivity indicators that quantify how sensitive the price of a derivative instrument, such as an options contract, is with respect to inherent parameters characterizing the derivative instrument, such as the price of the underlying asset or the time to maturity of the derivative instrument. These indicators, besides being designated by the term "Greeks", which they became known as because of the Greek letters usually assigned to describe them, are also interchangeably called *hedge parameters* or *risk sensitivities*. They are very useful to trading professionals who use the information to make decisions on whether to buy or sell options contracts, and generally manage and hedge their portfolio of assets.

To describe an example of this, the most well known and utilized of all the Greeks in the last decades has been the Delta. In the context of a call option contract, Delta describes the sensitivity of the value of the option with respect to the price of the underlying asset. Given the asymmetric profit profile of the call option contract, it is well established that the price of the option is less sensitive to a movement in the underlying price when the price of the underlying asset is in the vicinity of the strike price of the contract rather than if it were significantly greater than the strike price. In other words, option prices are more sensitive to movements in the underlying price when they are in the money as opposed to when they are out of the money. The Delta indicator would capture this sensitivity by assigning a higher value to the latter case. The trading professional, who has access to the value of Delta on his trading screen, would interpret this in the following way: a high Delta signifies that the underlying price is currently in the money, which means that if the contract were to expire now, the option would be executed. This information would prompt the trader to make a decision regarding whether to keep or remove the option from his portfolio, depending on his portfolio objectives. If, on the one hand, the trader has no position on the underlying asset, he may have a purely speculative strategy on the option. If this is the case, and he does not believe the underlying price will change until the maturity date, he will either keep the option in his portfolio if he is long and the option is in the money, or he is short and the option is out of the money. Conversely, he will remove it from his portfolio if he is long and the option is out of the money, or if he is short and the option is in the money.

If, on the other hand, the trader has a position on the underlying asset, he may be trading options to employ a hedging strategy. This is the more usual case and the main reason Greeks indicators are used for in derivatives trading. With portfolios consisting of stocks and options, traders can reduce the risk associated to their portfolios. To do this, they rely on the information provided to them by the Greek indicators. Keeping with the example of Delta, if a trader has bought a call option (relative to a standard number of 100 shares) , and the Delta on that option stands at 0.6, he could hedge his exposure to the underlying asset by selling 60 shares of the underlying stock, thereby eliminating his exposure to small movements in the underlying price and protecting his portfolio if the underlying price drops. It should be noted that many traders implicitly associate the meaning of Delta with the probability that an option contract will expire in the money. Although this interpretation is not quite correct from a formal standpoint, it is close enough that traders can use the interpretation in practice effectively. Other Greek sensitivities such as Gamma, Vega, Theta, etc. have their own interpretation and uses to traders hedging their portfolios. For a comprehensive study of Greeks in Finance, one may consult for example Hull (2014).

The construct by which traders have so readily benefited from the information Greek indicators provided them over the years was developed in the 1970s by Fischer Black, Robert Merton, and Myron Scholes, in what became known as the Black Scholes, or Black-Scholes-Merton, model for the valuation of option contracts (Black and Scholes 1973). Before their approach, there was no mathematical framework with which to objectively determine the price of an options contract. They effectively developed a way to assign an arbitrage-free price to options contracts, and their method was universally adopted in the industry soon thereafter. Nowadays, it is considered a standard method with which to price options contracts.

One main assumption in the Black-Scholes-Merton model is to describe the price of the underlying asset as lognormally distributed. By treating the underlying asset price as a stochastic process that follows geometric Brownian motion with constant drift and volatility, the model admits a closed-form solution that possesses the important property of differentiability. The profit profile at the time of maturity of an options contract has an inherent kink at the strike price; the revenue is flat and worth zero on the unfavorable side of the strike price, and grows linearly when on the other side of the strike price where the option is executed.

This closed-form solution to the Black-Scholes-Merton model is crucial to compute the Greek indicators. As sensitivity measures, they quantify rates of change, and so by having the value of the option and the price of the underlying be characterized by differentiable functions, it is natural that these sensitivity measures be derivatives. Therefore, coming back to our example, under the Black-Scholes-Merton model, Delta is the derivative of the option value with respect to the price of the underlying asset. Other Greek quantities describe the sensitivity of the option value with respect to other parameters in the problem. One such example is for instance Vega. Although not a Greek letter *per se*, this indicator is the derivative of the option value with respect to the volatility of the underlying asset price. In the first paper, we compute one such quantity, which to be precise is the derivative of the objective function with respect to the volatility parameter of the stochastic process.

In this summarized exposition of the Black-Scholes-Merton model, we should add that the model treats time as a continuous variable.

The purpose of this detour and brief glance at Greeks in Finance was to provide a motivation for our initial interest in the topic described in the first paper. Greeks are derivatives that are very much of interest to options traders, but as they are calculated with the Black-Scholes-Merton framework, they rely on the assumptions that the model be adequate, namely that the assumption of complete markets holds. The Black-Scholes-Merton model has several limitations that are well known and exhaustively documented elsewhere, the subject of which does not fall within the scope of this thesis (see for example, Teneng 2011, for an overview).

What is relevant to these proceedings is that we had a benchmark with which to compare the results of our method, and an application with immediate real-world impact. Our initial goal and motivation was then to attempt to reproduce the calculation of Greeks by solving a multistage stochastic replication problem. Not only would a successful calculation confirm the validity of our method, but it would also mean that we could compute Greeks that would not be reliant on the assumption of a geometric Brownian motion process for the underlying asset price, but could indeed be described by a different stochastic process, presumably one deemed better suited to describe asset price movements.

This attempt would prove to be unsuccessful. We could not solve the multistage stochastic replication problem convincingly, not because of a flaw in the reasoning of our method, but rather because the algorithmic strategy we were using to solve multistage problems was ill-suited to solve the replication problem. We did nevertheless find another well-known problem with an analytical solution with which to benchmark our method, the classic newsvendor problem, which is described in detail in a numerical study in the paper.

Rather than trying to solve the very specific problem of computing Greeks in an option replication problem, our aim in more general terms with the first paper was two-fold: 1) to provide a method with which to calculate derivatives in multistage stochastic linear problems, whereby these derivatives indicate the sensitivity of the objective function to changes in model parameters, parameters that may even contribute to the characterization of the underlying uncertainty of the problem, and 2), develop a mathematical framework with which to substantiate the method and justify its validity.

1.2.1.2 Envelope Theorems

The theoretical foundation of the first paper consisted in exploring the properties of the envelope theorem and seek to extend them to multistage stochastic linear optimization.

The envelope theorem establishes a property with which to quantify the sensitivity of the objective value of an optimization problem with respect to the parameters of the problem. This topic has been extensively studied in microeconomics, being the underlying argument in central results such as Hotelling's Lemma in producer theory, Shephard's lemma and Roy's Identity in the theory of the firm and consumer choice.

In its most basic form, for unconstrained optimization, provided that the objective function is differentiable with respect to a problem parameter, the envelope theorem gives a result on the derivative of the objective value with respect to that parameter. Namely, that this derivative is equal to the partial derivative of the objective function with respect to the parameter evaluated at the optimal value. One may start proving this result by applying the chain rule to determine the total derivative of the objective value with respect to the parameter. What is remarkable about this result is that the variation of the optimal solution with respect to the parameter is irrelevant, as the first-order condition at the maximum guarantees that the derivative of the optimal value with respect to the solution vanishes. This leaves only the partial derivative of the objective value to the parameter as a non-zero term. Therefore, only a variation in the parameter is required to determine its effect on the objective value.

This idea was extended to constrained optimization by requiring further smoothness assumptions on the optimal solutions. However, these additional assumptions are too restrictive when considering that in most linear optimization problems the optimal solution is often not differentiable and sometimes not continuous or even unique with a variation in the parameter. The points in the parameter where there are changes in the optimal solution are noticeable when plotting the optimal value as a function of the parameter as they take the form of either a kink, i.e a continuous but non-differentiable point, or a jump (discontinuous point) in the objective value function. The envelope theorem would break down at these non-differentiable and even discontinuous (in the case of a jump) points of the objective value function in the parameter, thereby limiting the range of parameters points where the theorem is applicable. It seemed then necessary to identify these regions where discontinuities or kinks occurred.

These issues led us to consider the field of parametric linear programming, and specifically to evaluate the existing body of work related to sensitivity analysis in linear optimization problems. While we found significant results focused on the variation of the parameters appearing on the objective function and on the right-hand side of the set of constraints, the literature was sparser when it came to the variation of the parameters in the matrix of coefficients. In the paper, we explore and classify the set of points where the objective function is not differentiable. Indeed, one of the results in the paper is to have determined that this set is finite and has Lebesgue measure zero, which implies that the envelope theorem can be applied almost everywhere in the parameter range. This statement is of particular importance as we attempt to extend the envelope theorem to stochastic optimization.

When adapting the envelope theorem to multistage stochastic programming, we first prove a result for optimization problems where the stochastic process is discrete. We do this by showing that the traditional envelope theorem, armed with the statement regarding the finite nature of the set of non-differentiable values of the parameter, can be directly applied to the deterministic equivalent of the discrete problem. We further extend this to the case where the stochastic process is continuous.

1.2.1.3 Sampling Derivatives

The envelope theorem for multistage stochastic optimization gives us a result where the derivative of the objective function with respect to the parameter is a function of expectation operators. One of the issues in solving optimization problems where the stochastic process is continuous and we would like to evaluate an expected value, is that no matter how many discrete points are used to approximate a continuum, we can only ever claim to compute a sampled mean that approximates the expected value. This is a reasonable approximation to make as the number of discrete points increases, as we know this to be the case when we can apply the law of large numbers. In the paper, we describe a similar strategy to calculate the derivatives that we are looking for. Moreover, the algorithmic solution strategy that we use to solve multistage stochastic problems, Approximate Dual Dynamic Programming (ADDP), very naturally lends itself to such a sampling approach, as it employs a sampling strategy to solve multistage problems.

1.2.1.4 ADDP and its place in optimization theory

Linear programming in its modern formulation developed in the 1940s as a planning framework with which to solve logistics, inventory and resource allocation problems in what was initially a military setting. Kantorovich, Koopmans, von Neumann and Dantzig led the early developments and established the theoretical foundations in this field. Indeed, Dantzig is credited as having invented the simplex algorithm, which became the first numerical algorithm to efficiently solve linear optimization problems.

Soon thereafter, decomposition methods such as Benders' or Dantzig-Wolfe's were developed to tackle large-scale linear optimization problems. The requirement behind these methods is that the very large problem be structured into blocks, where decision variables can be partitioned into subsets in a way that separates a first-stage master problem from second-stage sub-problems that are solved given the solution to the first-stage problem.

This method of decomposing a large problem into smaller problems structured in several layered stages with recourse became a standard with which to solve stochastic optimization problems, optimization problems where parameters of the problem are uncertain. In these problems, the uncertainty of a particular parameter is modelled as a random variable following a specific probability distribution that is defined by the modeler. Examples of uncertain parameters in practical applications may be the price of a stock or commodity, the hourly energy production of a wind farm, the aggregate demand for electricity, etc.

1.2.1.4.1 The L-shaped Method

The L-shaped method is a standard iterative algorithm to solve two-stage, and even multistage, recourse problems, where a large linear optimization problem is decomposed in the manner alluded to above (Van Slyke and Wets 1969).

Having decomposed the original large optimization problem in a layered structure of master and subproblems, the master problem solution is passed on as input to each subproblem of the same layer. Whenever the resulting subproblem is infeasible, a linear feasibility cut is added to the master problem, so that the solution it provided to the set of subproblems is no longer allowed in the next iteration. Additionally, as feasible subproblem solutions better inform the shape of the objective of the recourse function - that we shall refer to interchangeably as the *post-decision value function* - of the master problem, optimality cuts capturing the post-decision value function behavior are introduced in the master problem. Optimality cuts, or subgradient cuts, are linear cuts where the slope coefficients and independent term are probability-weighted means across all scenarios of product terms between the vector of dual variables and the subproblem data, respectively the matrix of coefficients affecting the first-stage solution and the vector of independent terms. Together they constitute linear constraints with which to approximate the post-decision value function of the master problem. It should be noted that with each iteration of the algorithm, two estimations of the post-decision value function are evaluated. The first consists of the post-decision value function component of the optimal value of the master problem, which is computed via the set of optimality cuts collected thus far. The second consists of the evaluation of the latest optimality cut of the current iteration. The algorithm stops when these two estimations coincide, namely when the first evaluation is greater than or equal to the second in a minimization problem, and vice-versa in a maximization problem.

The solution strategy employed in the L-shaped method to solve decomposable linear problems is the basis with which we can begin to understand sampling algorithms such as stochastic dual dynamic programming (SDDP) and most notably ADDP.

1.2.1.4.2 SDDP and the curse of dimensionality

One of the issues with a decomposition approach such as the one employed in the L-shaped method is the curse of dimensionality. As the number of layered subproblems increases, or as the number of scenarios per layer increases, the number of computational operations increases, which very quickly slows down the total runtime of the program. Sampling algorithms such as SDDP (Pereira and Pinto 1991) and ADDP (Löhndorf et al. 2013) were developed to mitigate the disadvantages posed by the curse of dimensionality when solving a large multistage problem.

The L-shaped method is characterized by a two-step process where there is first a forward pass, where all subproblems across all stages are solved sequentially, each master problem providing its subproblems with its solution as input. The forward pass is then followed by a backward pass, where each subproblem is solved with the purpose of computing optimality cuts for each subproblem of the preceding stage.

The main insight of sampling algorithms such as SDDP is that the forward pass may be made leaner by evaluating only one random path generated by the stochastic process in the information structure, which is to say to solve one problem per stage, instead of solving all subproblems. In the backward pass, each subproblem receives as input the solution from the problem of the preceding stage that was chosen in the forward pass. The resulting optimality cut that is computed from considering all subproblems in a given stage is added to all subproblems of the preceding stage, instead of just the subproblem that belonged to the forward pass path.

The algorithm computes two bounds. One is deterministic and consists of the objective value calculated using the value function approximations from all nodes in the information structure. The second bound is statistical, and consists of the evaluation of the policy in forward random paths. In a minimization problem, the deterministic bound consists of a lower bound to the objective bound, and the statistical bound consists of an upper bound. The algorithm stops when the two bounds converge to sufficiently close values, where the modeler determines the criterion of "sufficiently close".

To better understand the difference between SDDP and ADDP, one must first take a look at the information structure of the two algorithms. Stochastic processes

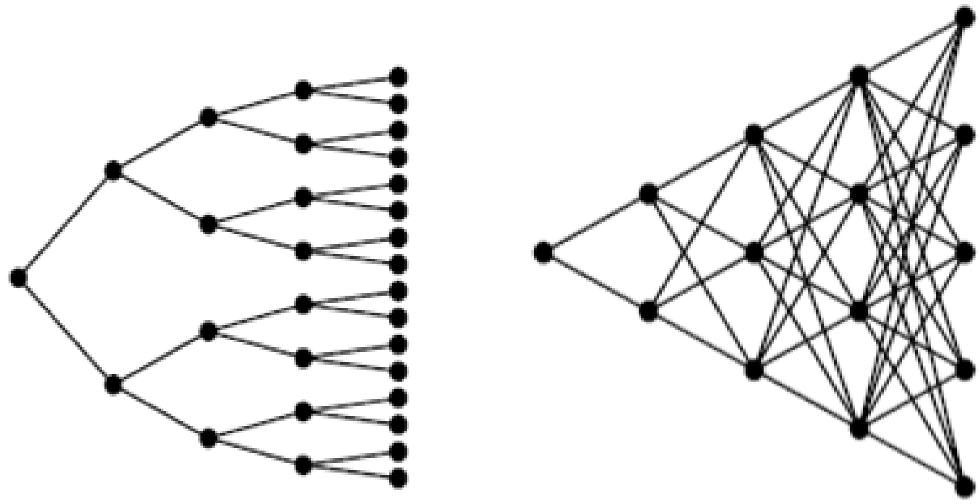

Figure 1.1: Information structure representations: scenario tree (left) and scenario lattice (right)

in multistage problems often lend themselves to being represented by tree-like structures, where each node in the tree has one unique parent node, but can itself be a parent node to a multitude of nodes. A root node represents the master problem. This root node has no uncertainty, as it represents a deterministic state, oftentimes the initial value of the stochastic process. Nodes constituting the second stage then represent multiple scenarios of the stochastic process, each joined to the root node by an arc representing the probability distribution of that scenario. Third stage scenarios representing the stochastic process movement from a given second-stage scenario then make up the subset of third stage nodes. This structure repeats itself from parent node to daughter node until the last stage of the problem.

This tree structure is the basis on which the SDDP algorithm operates. The ADDP algorithm can be considered a variant of SDDP, where the stochastic process is not represented by a scenario tree, but is instead represented by a scenario lattice, also known as a recombining tree. A scenario lattice is an information structure where each node may have one or more parent nodes. It is especially useful to represent Markov decision processes, which consist of processes where the conditional probability distribution of the stochastic process depends solely on the present state, and not on the full history of the process. Therefore, with a scenario lattice, each node represents a state of the process, but the path to get to that state is not unique. This allows for a problem in which the Markov property applies to be represented by a smaller information structure. A scenario lattice can code the same information with fewer nodes than a scenario tree can.

The problems that we tackle in the numerical section of the first paper consist of a two-stage newsvendor problem, and a multistage gas valuation problem. Note that in two-stage problems, SDDP and ADDP are no different from one another, as there is no difference in the information structure between a tree and a lattice where only two stages exist. However, in multistage problems where the stochastic process is markovian, the ADDP algorithm is a suitable strategy, and it is what we use to solve the gas valuation problem.

1.2.1.5 Contributions to Parametric Sensitivity Analysis in Multistage stochastic linear problems

This condensed presentation of these sampling algorithms is meant to justify the approach we take in the paper to calculate derivatives. We compute sampled means of the derivative by taking evaluations in sampled forward passes. The same method is used in sampling algorithms to compute the statistical bound with which to make convergence checks and thus provoke the stopping point for the algorithm. For this reason, derivatives can be calculated on the fly in the normal execution of the algorithm when performing convergence checks.

In the numerical section of the paper, we test the performance of this method against two benchmarks. The first benchmark is available in the newsvendor problem, and consists of the analytical solution to the problem. The newsvendor problem is one of the few stochastic optimization problems having a closed-form solution. This problem is very convenient to use as a benchmark for this reason. The second benchmark we use is the calculation of well-known finite differences numerical derivatives. The idea consists of taking the difference between two evaluations of the objective value of the problem, solved for different values of the parameter. This requires that the optimization problem be solved twice per evaluation of one finite difference, which is impractical compared to the on-the-fly calculation of our method.

1.2.2 Storage and its applications in Power Markets

The second and third papers revolve around the use of storage in electricity markets. Consisting of a group of emerging technologies that is expected to represent a large share in the energy mix of electricity markets, storage is especially seen as a welcome stability mechanism to regulate the grid, as grid variability brought about by an increasing share of intermittent renewables is expected to increase. Furthermore, storage offers flexibility to supply and demand in a market that is traditionally rigid, but that nevertheless must always be balanced.

1.2.2.1 Coordinated or uncoordinated participation in electricity markets

Although there are unchallenged benefits for overall grid stability and flexibility, academics are taking this logic one step further and are looking to bring storage into renewable energy portfolios with the intention of coupling storage operations with their remaining portfolio in joint market bidding. This is particularly the case with wind farms and solar plants. The inherent argument is that as these renewable sources are characterized by a variable and uncontrollable production which depends on how much the wind blows or whether the skies are clear, storage would offer the possibility to provide a controllable and stable production output, stepping in to provide energy when the wind farm is short of its commitments, and storing the extra energy when the wind farm is producing too much. Wind farms would thereby avoid paying imbalance costs for deviating from their scheduled commitments to the grid.

Furthermore, storage offers the possibility to couple the portfolio's supply with market demand and the electricity price, by deferring sales to periods of higher prices, and charging in times of lower demand or when prices are low.

It is the purpose of the paper to challenge these notions and explore under which conditions it is demonstrably true that there is an advantage to joint planning, and coincidentally, under which conditions there is no advantage.

1.2.2.1.1 Market Structures

Deregulated wholesale electricity markets have different design structures across the world, but one common element is that they all have some form of day-ahead market that makes up part of the spot market. If we do not account for futures markets or bilateral contracts that tend to involve large industrial consumers, the purchase and sale of electricity is regularly done via the spot market. The day-ahead auction is the first point of contact, where producers and consumers place bids to sell and purchase power volumes to be delivered on the next day. The manner with which these bids are cleared depends on the individual market rules, but the day-ahead auction aims to maximize the amount of volume traded while defining a uniform price for electricity across a defined region. The auction defines a planning schedule for each producer, who commits to deliver constant power volumes over discretized periods of time, generally one hour intervals, over the course of the following day. With 24 auctions held, the hourly schedule for each market participant is determined on the day before delivery.

Given the introduction of variability in the power supply through the emergence of renewable sources in the last two decades, markets across the world have had to adapt their design such as to give an adequate response and accommodate for this variability. Indeed, the increase in variability of supply has engendered the need for a mechanism with which to correct and adjust deviations between the expected production volumes planned at the moment bids are placed on the day preceding the time of delivery and the actual production volumes that are realized. The response to this issue, however different the implementation in each market, invariably takes the form of some kind of intraday market, where trading can be made on the day of delivery up to a time much closer to delivery. The format of this market varies across regions. Typical implementations consist of an intraday market that is either structured as an auction or is continuously cleared, and may have shorter delivery periods, such as 15-minute increment intervals. Examples of such markets are the German intraday continuous market representative of many European countries, the real-time markets generally found in American power markets or the intraday markets in Spain or Italy, which consist of not just one but several markets, the difference between them being the closing times before delivery. The addition of an intraday market to the spot market allows the producer to update his schedule as he has more accurate information on account of the closer time to delivery. These markets will generally have much lower aggregate trading volumes than the day-ahead market. From the viewpoint of a producer facing a variable and uncontrollable supply, the day-ahead market bid is a first-order approximation, which is then adjusted in the intraday-market.

Although the variability between expected and realized production is decreased in intraday market bids, it is not however zero, as there is still uncertainty at the closing time of the market before delivery. When power plants deviate from their production schedule, either by falling short of their production commitments or by generating too much energy, modern-day grids must procure the appropriate balancing energy in reserve markets, while passing on the procurement costs to the producers that caused the deviation. These costs, which are also known as imbalance costs or imbalance penalties, take into account whether the producer is short of his commitment to the market or has produced above its planned volume, and whether the market as a whole has an aggregate shortage or has too much energy in the grid. The relative position of the trader vis-à-vis the aggregate position of the market determines whether he stands to pay or receive a payment to or from the grid. The following table captures the four possible combinations.

Table 1.1: Imbalance position

	Market Aggregate Position	
Imbalance Position	Short	Long
Short	Cost	Revenue
Long	Revenue	Cost

If the aggregate position of the market is short, a market participant with a short position pays a penalty to the grid system operator, and a participant with a long position stands to receive a payment for his opposite position. Likewise, in the event of an aggregate market oversupply, the participant with a short position receives a payment, and the participant with a long position pays a penalty. Notice that the penalty itself need not be always positive. It may in fact be negative, reversing the payment directions established in the table above.

The nature of these imbalance settlement schemes vary from market to market. However, most markets usually fall in one of two categories: either the imbalance costs are symmetric, where the penalty for being long or short is the same, or they are asymmetric, in which case they are not the same and long and short imbalance positions are resolved at different prices. Asymmetric imbalance settlements are practiced in British, French, Italian and Spanish markets, and symmetric imbalances are practiced in the German spot market.

Considering the general structure of the spot market, the goal sought out by producers when pairing a storage unit with a variable renewable source consists in reducing the imbalance penalties they face on account of the variability of production.

1.2.2.1.2 Joint Planning and the proposition of equality of profits

In the paper, we seek to define a market setting in which there is no advantage to this combination of assets, and simultaneously characterize existing market elements that do make the combination worthwhile.

The main argument made in the paper against joint planning is that any deviation incurred by the variable production source will disrupt the schedule of the storage unit. This setting assumes that the storage unit may participate freely and fully in the same market where the variable source operates. If this is the case, the storage unit will bid in the market, and will have a schedule it must honor. If there is a deviation in the variable source, the storage unit will incur a symmetric and opposite deviation of its own in order to hedge the deviation of the variable source and cancel the aggregate exposure to the imbalance settlement price. The avoided costs from hedging the deviation internally in the manner described above are the same as the loss in profit that the storage unit would otherwise have had, had it not deviated from its schedule. The overall profit and loss from operating both assets is therefore unchanged. This scenario assumes that the imbalance settlement prices are symmetric, such as in the German market.

It should be noted that in many instances found in the literature where the advantage of joint planning is quantified explicitly, the comparison is often between a variable source operating alone and a variable source with a tethered storage unit that does not have market access. In these instances, the advantage of the latter system is unchallenged, as the storage unit offers flexibility that the variable source alone would otherwise not have. However, the role of the storage unit is to merely assist the variable source in its commitments to the market, not participating in the market directly. This specific comparison does not address the issue discussed in the paper, which is to consider the case where both assets, the variable source and the storage unit, can participate in the market.

There are however many cases where the equality of profits proposed in the paper between joint and separate planning cannot be achieved.

We assumed beforehand that the imbalance settlement prices were symmetric. If the price a market participant pays for a positive deviation is different from the price for a negative deviation, the offset correction in volume employed by the storage unit to balance out the variable source deviation is not financially equivalent and therefore the equality in profits breaks down. Two cases may emerge from an asymmetry in imbalance prices. Either the hedging operation brought about by the storage unit is valued at a larger price than the original imbalance contracted by the variable source, or it is valued at a lower price. One case would represent a relative profit, while the other would mean a loss. The case is further complicated by adding to this the overall aggregate position of the market and whether the imbalance prices are positive or negative, elements that affect the magnitude of the imbalance prices and the direction of payment. Regardless of what the preferred case is in any instance, the decision maker does not know at the moment he or she must decide to hedge a deviation with the storage unit what the exact imbalance price or prices will be. Therefore, the decision maker cannot take advantage of imbalance prices.

Two-price imbalance settlement schemes are not the only possible source of price asymmetries. Intraday markets are the markets where hedging adjustments to spot market positions can be made, but they generally have lower trading volumes than the day-ahead market. This makes them prone to illiquidity, and are thus susceptible to bid-ask spreads. Bid-ask spreads define by their very nature different prices for sellers and buyers.

Furthermore, grid fees are another source of asymmetry for producers and consumers. In some markets, grid operators charge maintenance and operating costs to consumers but not to producers.

These are but a few examples of elements that introduce an asymmetry in pricing that would lead to a violation of the equality proposed in the paper. The paper further highlights the effect of market power and how gaming the market with a price maker strategy can benefit a joint and coordinated use of variable source and storage unit. Additionally, we mention that any regulatory or technical issues that restrict the use of the storage unit in market operations can disrupt the equality of profits established in the paper. Such examples include market entry barriers on the basis of the nature of the technology or minimum power bid requirements.

1.2.2.1.3 Defense for independent use storage in practice

Although we explore in the paper many instances in which the assumptions that lead to the proposition of invariant profits are violated, it was nevertheless our goal to demonstrate that tethering a storage unit to a variable source with the purpose of internal balancing may not constitute the better use of this asset. With this in mind, we carried out a numerical exercise, in which we proceed with a comparison of the profits that a storage unit joint to a wind farm would have generated operating on the German spot market with the profits that the same storage unit would have had operating on its own on the German secondary control reserve market in 2019. Specifically, it was the point of this comparison to show that if the profits generated by the storage unit operating alone with a suboptimal strategy exceeded the added value in profits generated by the storage unit when tied to a wind farm and operating optimally, then the argument for an independent use of a storage unit could be made.

To that end, we would make a comparison of the following two metrics. We would compute an upper bound of profits for joint planning, and compare it to a lower bound in separate planning. If the lower bound for separate planning exceeds the upper bound for joint planning, it would prove our point.

On the one hand, we would solve a deterministic bidding problem on the day-ahead and intraday markets, where market prices are known, and the wind farm bids the expected production forecast on the day-ahead market. We assume that the actual wind production is known in the second stage, where bids are placed in the intraday market. The storage unit would make its own bids on both markets, taking advantage of the perfect foresight of market prices. Deviations between expected wind forecast and actual production would be settled at the imbalance price, the reBAP. These would represent the costs that would be avoided with storage balancing. The upper bound would consist of the sum between the profits made by the storage and the avoided imbalance costs. This metric is an upper-bound for two reasons: First, the storage unit operates on both markets with perfect foresight of prices; Second, the imbalance costs that are avoided with the intervention of the storage unit would require the storage unit to adjust its schedule and therefore reduce its aggregate profits. In this upper bound, they added as a net profit.

On the other hand, we compute the profit that a storage unit would have operating on the German secondary control reserve by bidding a fixed percentile of the previous day's auction. To guarantee that the storage unit can fulfill its commitment to the market without ever needing to charge because it is empty or discharge because it is full, we allow storage level management operations to be resolved on the intraday market. We partitioned the storage capacity and output in a way that half of the total capacity would be dedicated to bidding on the control reserve, and the other half would be used to manage the overall storage level. The rule governing the intraday bidding is that the storage level must always be brought back to half full, regardless of what the price is on the intraday market. This strategy is both non-anticipative and suboptimal, as the bidder does not have knowledge of the auction market bids, and does not attempt to determine an optimal policy. For this reason, it is a lower bound of the profits that can be had in this market.

The experiment showed that despite the cost of maintaining feasible storage levels on the intraday market, the overall profits of bidding on the secondary control reserve were greater than the profits from tethering the storage unit to the wind farm and bidding on the spot market.

From a literature perspective, the goal of the paper is to bring awareness to the issue that it is not always a good idea to use storage as a mechanism with which to balance deviations of a variable production source. The paper also aims to provide a framework with which to analyze the merits of joint planning and to define the conditions under which it offers no advantage.

1.2.2.2 Storage Duration in Balancing Markets

The third paper, entitled "Optimal Battery Duration on the German Secondary Control Reserve", takes a look at the problem of dimensioning the size of a storage unit fit to bid exclusively on the German secondary control reserve (SCR).

In particular, it is a study of the profit potential of operating on the SCR by adopting an optimal bidding strategy, and how it evolves as a function of the duration of the battery. Duration, also known as the energy-to-power ratio, is a metric that indicates how long a battery can sustain delivering energy at its maximum power output before becoming empty. This metric is of particular importance when considering to use a storage unit on the balancing markets, where a market participant may be asked to provide a constant power volume for a period of time spanning several hours.

In the SCR, auctions are held the day before delivery, and apply to delivery periods of four hours. For a storage unit to participate exclusively on a market with a schedule that is beyond its control, it is especially important to understand what profits it can expect to generate given its duration level. This paper aims to quantify the maximum profits a storage unit can hope to achieve. These profits constitute an upper bound that no storage unit can reasonably be expected to realize, as it is the result of a clairvoyant bidding strategy where all auction data is known to the decision-maker. However, they do inform the decision-maker in the following two ways: first, its nature as an upper bound determines a ceiling to the expectations of possible investors; secondly, and perhaps more importantly, as this upper bound is computed for storage units of different durations, we can chart its evolution and observe if there is a preferred value for the duration of a storage unit operating on this market. The study of this curve, which is drawn for constant power but increasing storage size, shows that profits increase with an increase in duration, which is consistent with our expectations, as a larger storage capacity allows the access to a larger spectrum of power levels. This curve also shows that profits saturate and converge to a limit value with increasing duration, suggesting that from a certain duration value onwards, there are no more profits to be had, as the potential for profits will have maxed out. This duration level is quantified, and is found to be around 13 hours for most power levels in the range of bids commonly found in the market.

The upper bound that we calculate is the result of solving a sequence of deterministic mixed-integer bilinear optimization problems. Each problem corresponds to finding the optimal bidding strategy to maximize the profits in a four-hour auction. This involves bidding for positive and negative balancing energy, which includes a power volume and a price. In order to simplify the problem, we apply McCormick envelopes to relax and convert the bilinear terms in the objective function, which consist of product terms between price and volume, both of which are decision variables, into linear terms. We use a parametric search method to solve the mixed-integer problem.

We solve each problem sequentially for a total of 2190 auctions, which corresponds to an entire year's worth of data, encompassing the period between August of 2018 and July of 2019. This time span is judged to be a reasonable time horizon to cancel out spurious effects in the data.

1.3 Structure of the Thesis

The Thesis is organized with the following structure. Chapter 2 provides the article on "Envelope Theorems for Multi-Stage Linear Stochastic Optimization". This essay has been submitted and accepted to the academic journal *Operations Research*, having previously undergone two major revisions and one minor revision. Chapter 3 proceeds with the work on the "Economies of Scope for Electricity Storage and Variable Renewables". This paper has been submitted to the engineering journal *IEEE Transactions in Power Systems*. After a first assessment to reject and resubmit and a second assessment to revise and resubmit, the article has been revised and submitted again, where it is currently under review. Chapter 4 presents the article on "Optimal Battery Duration on the German Secondary Control Reserve". Finally, the thesis concludes with Chapter 5, with a reiteration of the main results to take away from all three studies, general conclusions and a discussion on the possible paths for further research.

Chapter 2

Envelope Theorems for Multi-Stage Linear Stochastic Optimization

written in collaboration with Prof. Dr. David Wozabal[1]

We propose a method to compute derivatives of multi-stage linear stochastic optimization problems with respect to parameters that influence the problem's data. Our results are based on classical envelope theorems and can be used in problems directly solved via their deterministic equivalents as well as in stochastic dual dynamic programming for which the derivatives of the optimal value are sampled. We derive smoothness properties for optimal values of linear optimization problems, which we use to show that the computed derivatives are valid *almost everywhere* under mild assumptions. We discuss two numerical case studies, demonstrating that our approach is superior, both in terms of accuracy as well as computationally, to naïve methods of computing derivatives that are based on difference quotients.

[1]**Publication History**: Initially submitted on 18.06.2018. Accepted for publication on *Operations Research* on 23.04.2020.

2.1 Introduction

In most planning and decision problems some of the relevant parameters are not known with certainty at the time when decisions are taken. Stochastic optimization deals with such uncertainties in problems that lend themselves to treatment in an optimization framework. More specifically, stochastic optimization is concerned with the solution of optimization problems for which some parameters in the problem formulation are uncertain (Birge and Louveaux 2011, Shapiro et al. 2014).

Most of stochastic optimization theory is concerned with obtaining optimal values and optimal decisions in the face of uncertainty. In this paper, we investigate the less researched question of how to calculate derivatives of optimal values with respect to the parameters of the problem. We devise a method to determine the sensitivity of the optimal value of a multi-stage stochastic linear optimization problem (MSLP) with respect to changes in model parameters, i.e., the data of the model. Our theory enables the decision maker to directly compute sensitivities and extract useful information about the nature of the solution of a MSLP.

Sensitivities of the optimal value with respect to parameters like costs, prices, characteristics of technical systems, parameters of the underlying stochastic process, or the amount of available resources are of interest for at least two reasons. Firstly, as in deterministic optimization, derivatives can be interpreted as *shadow prices* of the resource or parameter in question. Secondly, derivatives provide information on the inherent risks implied by changes to parameters such as prices, volatilities, and contract provisions.

There is a large number of papers that use MSLPs to either calculate optimal operational decisions or the value of an asset or a contractual right. In many of these applications, sensitivities with respect to some parameters of the underlying problem are of interest. Recent examples include studies in fields as diverse as production planning (Zhang et al. 2011, Bollapragada et al. 2015), inventory management (Mardan et al. 2015), supply chain management (Wu 2011, Fan et al. 2017), energy (Khor et al. 2008, Lai et al. 2010, Bruno et al. 2016), water resource management (Huang and Loucks 2000), project planning (Rafiee and Kianfar 2011), asset liability management (Valladão et al. 2014, Duarte et al. 2017), and mobile computing (Tham and Cao 2018).

Another prominent example from financial risk management are the so-called *Greeks*, i.e., derivatives of option prices or contract values with respect to underlying parameters of the problem (Hull 2014). For standard cases in option pricing and a handful of parameters, Greeks can be computed explicitly. However, for the large set of analytically intractable problems, for example in commodity risk management, derivatives have to be numerically approximated.

In the finance literature most authors rely on difference quotients to approximate derivatives for problems that have no analytical solutions such as American options (see Clewlow and Strickland 2000, Jäckel 2002, Glasserman 2004, Geman 2009). However, there also exist other methods to approximate derivatives. In particular, the pathwise method, discussed for example in Glasserman (2004), is interesting in this context. The ideas are somewhat similar to the ideas presented in this paper, although in a slightly different setup, which is focused on options pricing.

Taking derivatives of optimal values has a long history in the optimization literature. In particular, computing sensitivities with respect to parameters of a deterministic optimization problem is one of the main questions in the field of parametric optimization. In this paper, we build on the extensive literature on linear parametric optimization (see Gal and Greenberg 1997, for an overview), in particular, when studying the question of smoothness of the optimal value of linear optimization problems.

In stochastic optimization, determining sensitivities of optimal values is at the core of many decomposition methods like the L-shaped method (Van Slyke and Wets 1969, Birge and Louveaux 2011). These approaches determine how the recourse function of the next period varies with a change in here-and-now decisions, which usually appear only in the right-hand side of the next stage's problem. Hence, the sensitivities are evaluated with respect to decisions variables in previous stages and not the data of the problem, which is different to the focus of this paper. Furthermore, decomposition approaches are usually much more restrictive when it comes to which parameters of the problem are allowed to vary.

We also mention a stream of literature in stochastic optimization that is concerned with the sensitivity of the optimal value with respect to perturbations in the distribution of the involved random variables (Dupačová 1987, 1990, Branda and Dupačová 2012). However, the aim of these efforts is rather specific and mainly directed at robustifying solutions with respect to errors made in the estimation of the distributional model for the randomness in the problem.

Our approach differs from the existing literature in the sense that we use classical envelope theorems for linear programs to prove envelope theorems for MSLPs and calculate derivatives with respect to arbitrary parameters of the problem. The proposed approach computes derivatives based on samples from the optimal policy. We prove the validity of our approach by showing that, under weak regularity conditions, the optimal value of an MSLP is almost everywhere differentiable.

Although our results are applicable to general MSLPs, we wrote this paper with Markov decision processes (MDPs) in mind. In particular, we intend to calculate derivatives for problems solved using a Markovian variant of *stochastic dual dynamic programming* (SDDP), called *approximate dual dynamic programming* (ADDP). ADDP approximately solves stochastic optimization problems with fully continuous state and action spaces. Note that the classical SDDP algorithm is a special case of ADDP, which implies that all results in the paper also hold for SDDP. Using novel asymptotic results for ADDP, we show that, under certain regularity conditions, the sensitivities computed from discrete approximations of the problems converge to the true sensitivities of the original problem as the approximation of the continuous randomness by discrete processes gets finer.

In a numerical study, we show that our method works well for practical problems. To this end, we set up two case studies: a simple two-stage newsvendor problem which can be solved analytically, and a multi-stage gas storage optimization problem. We use the newsvendor problem to test the convergence of our sampled derivatives to the real ones and the gas storage optimization to demonstrate that the proposed method scales well for larger problem instances with many decision stages. Furthermore, the gas storage example illustrates the value of our approach in calculating Greeks in (energy) finance applications. We use a naïve approach based on difference quotients as a benchmark, and show that our method outperforms this benchmark both in accuracy as well as computationally.

The paper is structured as follows: In Section 2.2, we give a short review of classical envelope theorems and discuss the specific case of linear optimization problems. In particular, we study the question of differentiability and prove that the optimal value of a linear optimization problem is differentiable almost everywhere, jointly in all its data. In Section 2.3, we apply these results to derive an envelope theorem for MSLPs and in Section 2.4 we show how derivatives can be sampled in SDDP-type algorithms for problems with discrete randomness and demonstrate that derivatives of discrete approximations of problems with continuous randomness are asymptotically correct as the approximating discrete process gets finer. Section 2.5 is devoted to a numerical study demonstrating the practical viability and computational tractability of the proposed approach. Section 2.6 concludes the paper and discusses avenues for further research.

2.2 Envelope Theorems

In this section, we introduce envelope theorems for linear optimization problems, which will be the basis for computing sensitivities for MSLPs in the next section. In particular, we investigate when the optimal value of a linear optimization problem is differentiable and show that non-differentiability is in some sense *rare*.

Envelope theorems quantify how much the optimal value of an optimization problem changes for an incremental change in some of the parameters of the problem. Results of this type are important tools in microeconomic theory and have been extensively used to prove central results in comparative statics analysis such as Shepard's lemma, Roy's identity, or Slutsky's equation relating to questions in the theory of the firm as well as consumer choice (see Varian 1992). The first statement of the envelope theorem in its modern form arguably appears in Samuelson (1947). Dynamic envelope theorems in optimal control can, for example, be found in LaFrance and Barney (1991) and the most general results known to the authors appeared in Milgrom and Segal (2002).

The most basic form of the envelope theorem concerns maximizing a sufficiently smooth function $f(x,\theta)$ that depends on a parameter θ as well as on a decision x. Denote by $x^*(\theta)$ the optimal decision for a given parameter θ and by $V(\theta)$ the corresponding optimal value. Then clearly,

$$\begin{aligned} V'(\theta) &= \frac{d}{d\theta} \max_x f(x,\theta) = \frac{d}{d\theta} f(x^*(\theta),\theta) \\ &= \frac{\partial}{\partial x} f(x^*(\theta),\theta) \frac{d}{d\theta} x^*(\theta) + \frac{\partial}{\partial \theta} f(x^*(\theta),\theta) = \frac{\partial}{\partial \theta} f(x^*(\theta),\theta), \end{aligned}$$

where the last equality follows from the first-order condition

$$\frac{\partial}{\partial x} f(x^*(\theta),\theta) = 0$$

at the optimal point $x^*(\theta)$.

The above shows that the derivative of the optimal value can be computed knowing only $\frac{\partial}{\partial\theta} f$ and the optimal solution $x^*(\theta)$. In particular, it is not required to know $\frac{d}{d\theta} x^*(\theta)$, i.e., how the optimal solution changes in θ, which is usually much harder to calculate. This property makes the envelope theorem extremely useful in the analysis of optimization problems.

The outlined approach can be extended to constrained optimization problems but generally requires smoothness assumptions on the optimal primal and dual solutions. We base our results on the following classical envelope theorem, which extends the simple derivation above to the case of constrained optimization by considering the Lagrangian of the problem.

Theorem 2.1. (Takayama (1985), Theorem 1.F.1.). *Let* $\Theta \subseteq \mathbb{R}$ *be open and consider differentiable functions* $A : \Theta \to \mathbb{R}^{m\times n}$, $b : \Theta \to \mathbb{R}^m$, $c : \Theta \to \mathbb{R}^n$ *and a linear optimization problem of the form*

$$V(\theta) = \begin{cases} \max_x & \langle c(\theta), x\rangle \\ \text{s.t.} & A(\theta)x \le b(\theta). \end{cases} \tag{2.1}$$

If, for a given $\theta \in \Theta$, *the optimal primal and dual solutions* $x^*(\theta)$ *and* $\lambda^*(\theta)$ *are continuously differentiable in* θ, *then*

$$V'(\theta) = \langle \nabla c(\theta), x^*(\theta)\rangle + \left\langle \lambda^*(\theta), \nabla b(\theta) - \frac{d}{d\theta}A(\theta)x^*(\theta) \right\rangle. \tag{2.2}$$

Note that in the above, we use $\frac{d}{d\theta}A(\theta)$ as short-hand for the componentwise derivative with respect to θ and $\langle \cdot, \cdot \rangle$ as notation for the inner product.

Looking at the formulation of Theorem 2.1, it is obvious that, in the case of linear optimization problems, the condition on the smoothness of x^* and λ^* is restrictive and therefore the result is, in general, not applicable for all values of θ. More specifically, it is well known that the solution of linear optimization problems need not be differentiable in the parameters of the problem. The following problem is a modification of an example given in Martin (1975), which illustrates this point:

$$V(\theta) = \begin{cases} \max\limits_{x,y} & \frac{1}{2}\theta x + y \\ \text{s.t.} & x + y \le 1 \\ & x + \theta y \ge 1 \\ & x, y \ge 0. \end{cases} \tag{2.3}$$

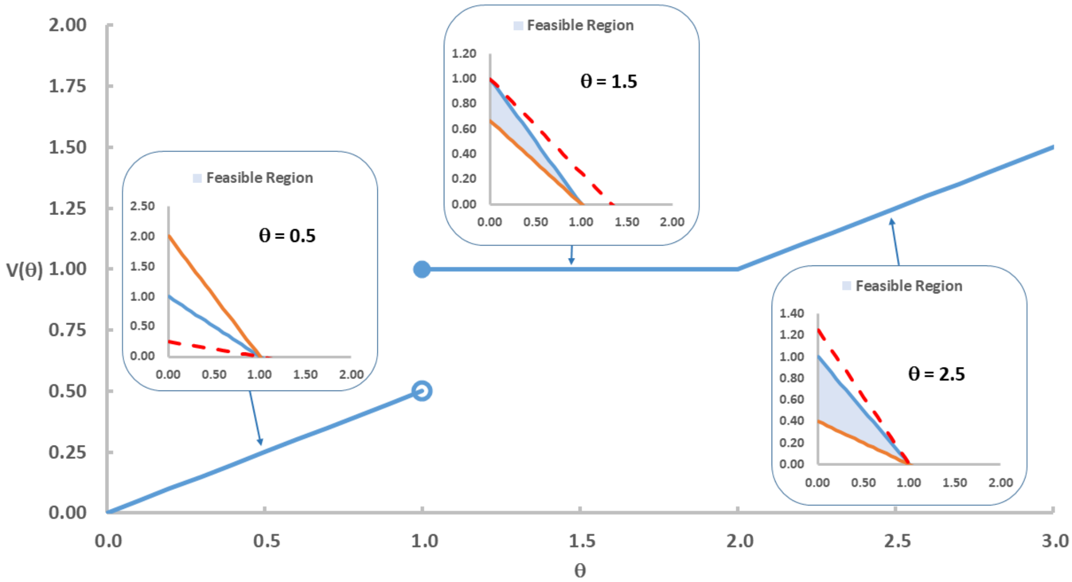

Figure 2.1: Value function and typical feasible regions for problem (2.3).

For $\theta \geq 0$, the optimal solution and optimal value are given by

$$\arg\max_{(x,y)} V(\theta) = \begin{cases} \{(0,1)\}, & 1 \leq \theta < 2 \\ \{(1,0)\}, & \theta \notin [1,2] \\ \{(1-\alpha,\alpha) : \alpha \in [0,1]\}, & \theta = 2 \end{cases}, \quad V(\theta) = \begin{cases} 1, & 1 \leq \theta < 2 \\ \frac{\theta}{2}, & \theta \notin [1,2). \end{cases}$$

Clearly, the optimal value jumps at $\theta = 1$ and kinks at $\theta = 2$. Figure 2.1 illustrates the optimal value function as well as typical examples of feasible regions in the three relevant sections of θ. An example featuring disjoint feasible regions can be found in Willner (1967).

Hence, for a general linear optimization problem, the optimal value does not even have to be continuous in the parameters of the problem and therefore the envelope theorem cannot be used to calculate $V'(\theta)$ for every θ. However, we observe that there are only finitely many points where $\theta \mapsto V(\theta)$ is not differentiable in the above example. As we will see below, this observation extends to the general case.

In a linear optimization problem, the issue of differentiability of the optimal value is closely related to the notion of uniqueness and non-degeneracy of the solution. Given an optimal simplex tableau, a (feasible) solution is *unique*, if none of the non-basic reduced cost coefficients vanish and it is *non-degenerate*, if and only if none of the basic variables vanish. If both conditions are met, differentiability of the objective follows (see Gal and Greenberg 1997). However, the two conditions are not necessary for the objective to be differentiable. In particular, in problem (2.3), we observe that for $\theta \notin [1, 2]$ the solution is degenerate, but the value function is differentiable with $V'(\theta) = \frac{1}{2}$.

Another more general sufficient condition for smoothness of V at θ is that the optimal basis remains unchanged in a neighborhood $\mathcal{U}$ of θ, i.e., that the set of binding constraints remains the same for all $\theta \in \mathcal{U}$. If this is the case, then, in $\mathcal{U}$, all entries of the optimal simplex tableau vary smoothly in θ, and, since the optimal value is a linear function of the entries, so does the optimal value.

Therefore, as a first step, we characterize the points where there are discontinuities or kinks in the value function by characterizing the sets of parameters θ where the optimal basis of a linear optimization problem of a given dimensionality does not change. Similar questions have been studied extensively in the literature on parametric linear programming (see Courtillot 1962, Willner 1967, Barnett 1968, Dent et al. 1973, Adler and Monteiro 1992, Ward and Wendell 1990, Gal and Greenberg 1997). However, results in these papers are usually restricted to either variations on the right-hand side, the objective function, or special variations in the matrix. For our purposes, a result that allows for simultaneous variation in all parameters of the problem is required. To the best of our knowledge, such a result does not yet exist.

For further analysis, we find it convenient to re-write problem (2.1) by introducing slack variables $s_1, \ldots, s_m$ for the inequality constraints, rewriting the constraints as $A(\theta)x + s = b(\theta)$. Defining $y = (x, s)$, we can thus write a new linear optimization problem equivalent to (2.1) in standard form as

$$V(\theta) = \begin{cases} \max_y & \langle d(\theta), y\rangle \\ \text{s.t.} & D(\theta)y = b(\theta),\ y \geq 0, \end{cases} \tag{2.4}$$

where d and D are correspondingly updated versions of c and A. Defining $\text{vec}(D)$ to be the row vector that consists of a concatenation of all the rows of D, $\mathcal{D} = (d, \text{vec}(D), b)^\top$ represents all the data of the problem as an element of $\mathbb{R}^N$ for some $N \in \mathbb{N}$. Furthermore, we define the short-hand $[n] = \{1, \ldots, n\}$.

As mentioned above, we are interested in the case where all the parameters of the linear optimization problem can vary simultaneously, i.e., we study how a variation in $\mathcal{D} \in \mathbb{R}^N$ affects the solution. We first prove a result characterizing the points where the value function of a linear optimization problem is smooth jointly in the parameters of the problem. The proof follows the idea outlined in Freund (1985) for the special case of changing only the matrix D along one line. We deal with the additional complication of arbitrarily changing all elements of $\mathcal{D}$ using results from real algebraic-geometry. To this end, we will require the following definitions.

Definition 2.1. (Semi-algebraic set, Bochnak et al. (2013), Definition 2.1.4). *A semi-algebraic subset of* $\mathbb{R}^n$ *can be written as*

$$\bigcup_{i=1}^{s} \bigcap_{j=1}^{r_i} A_{ij},$$

where the sets A_{ij} *are either of the form* $\{x \in \mathbb{R}^n : P_{ij}(x) < 0\}$ *or* $\{x \in \mathbb{R}^n : P_{ij}(x) = 0\}$ *for polynomials* P_{ij} *defined on* $\mathbb{R}^n$.

Definition 2.2. (Diffeomorphism, smooth manifold). *For* $A \subseteq \mathbb{R}^n$ *and* $B \subseteq \mathbb{R}^m$*, a mapping* $f : A \to B$ *is called a* smooth diffeomorphism *if it is invertible and* f *as well as* f^{-1} *are smooth. A* k*-dimensional* smooth manifold $X \subseteq \mathbb{R}^n$*, is a set such that for every point* $x \in X$ *there is an open set* $U \subseteq X$ *and a smooth diffeomorphism* $f : U \to \mathbb{R}^k$ *such that* $f(U) = B_k$*, where* $B_k \subseteq \mathbb{R}^k$ *is the open unit ball.*

Lemma 2.1. *Let $\mathcal{P}(\mathcal{D})$ be a problem of form* (2.4) *dependent on data $\mathcal{D} \in \mathbb{R}^N$ with a matrix D of dimensions $m \times n$, $m < n$. Define $\beta \subseteq [n]$ to be a basis with $|\beta| = m$ and the set*

$$R_\beta = \{\mathcal{D} : \beta \text{ is an optimal basis for } \mathcal{P}(\mathcal{D})\}.$$

Then R_β is the union of finitely many connected smooth manifolds of varying dimensions in $\mathbb{R}^N$.

Proof. Proof. For any D, define D_β as the square matrix resulting from D with the columns $[n] \setminus \beta$ removed. It holds that

$$R_\beta = \{\mathcal{D} : \det(D_\beta) \neq 0\} \cap \left\{\mathcal{D} : D_\beta^{-1} b \geq 0\right\} \cap \left\{\mathcal{D} : d_\beta D_\beta^{-1} D \geq d\right\},$$

where the first condition ensures that x with $x_\beta = D_\beta^{-1} b$ and $x_{[n]\setminus\beta} = 0$ is a primal basic solution, the second one restricts to $\mathcal{D}$ such that β is a primal feasible basis, while the third set contains the $\mathcal{D}$ such that β is dual feasible. The last two conditions in combination yield that x is primal optimal and, therefore, that β is an optimal basis for $\mathcal{P}(\mathcal{D})$.

Note that $D_\beta^{-1} = \det(D_\beta)^{-1} \operatorname{adj}(D_\beta)$ with adj(A) the adjugate of A, which is defined as the transpose of the cofactor matrix $C(A)$ of A, i.e., adj(A) $= C(A)^\top$. Note further that if $\det(D_\beta) \neq 0$, we can always re-arrange columns of D_β in such a way that $\det(D_\beta) > 0$. In this case, $D_\beta^{-1} b$ yields a basic solution with correspondingly permuted elements. We can therefore write

$$R_\beta = \{\mathcal{D} : \det(D_\beta) > 0\} \cap \{\mathcal{D} : \operatorname{adj}(D_\beta) b \geq 0\} \cap \{\mathcal{D} : d_\beta \operatorname{adj}(D_\beta) D \geq d \det(D_\beta)\}. \quad (2.5)$$

Let us first investigate the set $\{\mathcal{D} : \operatorname{adj}(D_\beta) b \geq 0\}$. The elements of the adjugate are determinants that can be written as polynomials in the variables of D. Defining $A_{(i,:)}$ as the i-th row of a matrix A, it follows that $P_i(\mathcal{D}) := \operatorname{adj}(D_\beta)_{(i,:)} b$ is a polynomial in the components of b and D, hence, in particular, a polynomial with coefficients from $\mathcal{D} \in \mathbb{R}^N$.

It follows that the set $\{\mathcal{D} : P_i(\mathcal{D}) \geq 0,\ \forall i = 1, \ldots, m\} \subseteq \mathbb{R}^N$ is a semi-algebraic set. By the same argument the other two sets in (2.5) are semi-algebraic and therefore so is R_β. By Proposition 2.9.10 in Bochnak et al. (2013)

$$R_\beta = \bigcup_{i=1}^{I} \mathcal{S}_i,$$

where $I \in \mathbb{N}$ and $\mathcal{S}_i$ are smooth manifolds in $\mathbb{R}^N$ and there exist diffeomorphisms $h_i : (0,1)^{n_i} \to \mathcal{S}_i$, where $0 \leq n_i \leq N$ for all $1 \leq i \leq I$. Since the h_i are continuous and $(0,1)^{n_i}$ are connected, $\mathcal{S}_i$ are connected sets in $\mathbb{R}^N$. □

The following corollary is an easy consequence of Lemma 2.1 and treats the case that not all of the data varies but only a subset of the entries of A, b, and c.

Corollary 2.1. *Consider a set $\{k_1, \ldots, k_L\} \subseteq [N]$ with $L \leq N$, $\mathcal{D}' = (\mathcal{D}'_1, \ldots, \mathcal{D}'_N) \in \mathbb{R}^N$, and the following affine L dimensional subspace of $\mathbb{R}^N$*

$$\mathcal{X} = \left\{ \mathcal{D} = (\mathcal{D}_1, \ldots, \mathcal{D}_N) \in \mathbb{R}^N : \mathcal{D}_k = \mathcal{D}'_k,\ \forall k \notin \{k_1, \ldots, k_L\} \right\}.$$

The set of points $\mathcal{D} \in \mathcal{X}$ where there is a basis change in the corresponding linear optimization problem is the union of finitely many connected smooth manifolds of varying dimensions in $\mathbb{R}^L$.

Finally, we prove a result confirming the intuition that the number of points where the optimal value of a linear optimization problem is not differentiable is *small* as long as the function mapping the parameter θ to the data of the problem is bicontinuous, i.e., continuous with a continuous inverse.

Theorem 2.2. *Let $\Theta \subseteq \mathbb{R}$ be an open, connected set and consider invertible bicontinuous functions $A : \Theta \to \mathbb{R}^{m \times n}$, $b : \Theta \to \mathbb{R}^m$, $c : \Theta \to \mathbb{R}^n$ defining the data of the* *linear problem*

$$V(\theta) = \begin{cases} \max_x & \langle c(\theta), x \rangle \\ \text{s.t.} & A(\theta)x \leq b(\theta). \end{cases} \tag{2.6}$$

Then the set of points $\theta \in \Theta$ where $V(\theta)$ is not differentiable is finite and therefore has Lebesgue measure 0 *in $\mathbb{R}$. In all other points $\theta \in \Theta$, $\theta \mapsto V(\theta)$ is smooth.*

Proof. Proof. Define $\mathcal{D} : \Theta \to \mathbb{R}^N$ as $\mathcal{D}(\theta) = (\text{vec}(A(\theta)), b(\theta), c(\theta))^\top$. By the first part of Lemma 2.1, the sets R_β are the union of finitely many connected sets $\mathcal{S}_i$ with $1 \leq i \leq I$. Since $\mathcal{D}(\Theta)$ is connected in $\mathbb{R}^N$, so are the sets $\mathcal{D}(\Theta) \cap \mathcal{S}_i$ and $\mathcal{T}_i = \mathcal{D}^{-1}(\mathcal{S}_i) \subseteq \Theta$. Since there are finitely many $\beta \subseteq [n]$, it follows that Θ can be written as the union of finitely many connected sets. Clearly, the optimal basis for the linear program with data $\mathcal{D}(\theta)$ stays constant for all $\theta \in \mathcal{T}_i$.

Since connected sets in $\Theta \subseteq \mathbb{R}$ are either points or intervals, the set of points Y where the basis of $\mathcal{D}(\theta) \subseteq \Theta$ changes is of finite cardinality and thus has Lebesgue measure zero. Since for every $\theta \in \Theta \setminus Y$ there is a neighborhood where the basis of the problem with data $\mathcal{D}(\theta)$ stays the same, the optimal simplex tableau is a smooth function of $\mathcal{D}(\theta)$ and therefore $\theta \mapsto V(\theta)$ is smooth in a neighborhood of θ. □

The result above is very much in line with intuition about linear optimization. It is nevertheless nontrivial, since it shows that under the given natural restrictions there is no way to find functions, however complicated, such that there is a significant number of points θ where the optimal value is non-smooth, even if all the data of the problem varies simultaneously.

In particular, Theorem 2.2 implies that if a point $\theta \in \Theta$ is chosen at random, then the probability that the function $V(\theta)$ is not differentiable at θ has probability zero. This means that the envelope theorem can be applied almost everywhere and the set Y, defined in the proof of Theorem 2.2, can be ignored in practical applications.

However, the analysis of points where the objective function is not differentiable might still be of interest, since the existence and the location of these points may reveal structural properties of the problem. In this paper, we focus on the computation of derivatives at exogenously given points and show in the next section that points of discontinuity are actually rare in stochastic optimization.

2.3 Envelope Theorems for Multi-Stage Stochastic Programming

In this section, we apply the results from Section 2.2 to prove envelope theorems for MSLPs. These results yield expressions for derivatives of optimal values which are independent of the method used to solve the problem. In particular, we argue that when solving discrete MSLPs as deterministic equivalents, our results can be directly applied.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space with a probability measure $\mathbb{P}$, a filtration $\mathcal{F} = (\mathcal{F}_1, \dots, \mathcal{F}_T)$, and $\xi = (\xi_1, \dots, \xi_T)$ a random process with $\xi_t : (\Omega, \mathcal{F}_t) \to \mathbb{R}^{M_t}$. We denote by $\xi^t = (\xi_1, \dots, \xi_t)$ the history of the random process until period t. Let further Θ be an open and connected set in $\mathbb{R}$.

We consider a general MSLP with T stages where the relevant data in stage t depends on the realization of the randomness and on a parameter $\theta \in \Theta$. The general form of the problem can be written as

$$V(\theta) = \begin{cases} \max\limits_{x_1, x_2, \dots, x_T} & \mathbb{E}\left[\sum_{t=1}^{T} \langle c_t(\theta, \xi^t), x_t(\theta, \xi^t)\rangle\right] \\ \text{s.t.} & W_t(\theta, \xi^t) x_t(\theta, \xi^t) + T_t(\theta, \xi^t) x_{t-1}(\theta, \xi^t) \le h_t(\theta, \xi^t), \quad t = 1, 2, ..., T, \end{cases} \tag{2.7}$$

where the expectation is with respect to $\mathbb{P}$, the constraints hold almost surely, and $W_t(\theta, \xi^t)$, $T_t(\theta, \xi^t)$, $c_t(\theta, \xi^t)$, and $h_t(\theta, \xi^t)$ are matrices and vectors of fitting dimension. We define the data of the problem in stage t to be

$$\mathcal{D}_t(\theta, \xi^t) = (\text{vec}(W_t(\theta, \xi^t)), \text{vec}(T_t(\theta, \xi^t)), c_t(\theta, \xi^t), h_t(\theta, \xi^t)) \in \mathbb{R}^{N_t}$$

for some $N_t \in \mathbb{N}$ and consider $\mathcal{D}_t(\theta, \xi^t)$ as a function $\mathcal{D}_t : \Theta \times \mathbb{R}^{M_t} \to \mathbb{R}^{N_t}$. Note that the decisions $x_1, \ldots, x_T$ depend on ξ and are therefore random variables. To avoid cluttered notation, in the following, we suppress the dependency of W_t, T_t, c_t, h_t, and x_t on ξ^t.

We formulate the problem in such a way that x_t depends on previous decisions $x_1, \ldots, x_{t-1}$ just by its dependency on x_{t-1}. This property can always be enforced by augmenting the state space of period t by the variables in $x_1, \ldots, x_{t-1}$. It follows that the problem can be rewritten in its dynamic programming formulation

$$V_t(x_{t-1}, \xi^t, \theta) = \begin{cases} \max\limits_{x_t} & \langle c_t(\theta), x_t \rangle + \mathbb{E}\left[V_{t+1}(x_t, \xi^{t+1}, \theta) | \xi^t\right] \\ \text{s.t.} & W_t(\theta) x_t + T_t(\theta) x_{t-1} \leq h_t(\theta), \end{cases} \tag{2.8}$$

with terminal condition $V_{T+1} \equiv 0$ and x_0 given. The assumption of $V_{T+1} \equiv 0$ is not essential, and all results still hold, if V_{T+1} is replaced by a piecewise linear concave function. For the following, it will be convenient to define so-called *post-decision* value functions (see Powell 2011) as

$$\bar{V}_{t+1}(x_t, \xi^t, \theta) = \mathbb{E}\big[V_{t+1}(x_t, \xi^{t+1}, \theta) | \xi^t\big], \quad \forall t \in [T].$$

In order to avoid pathological cases, we will make the following assumption throughout the paper.

Assumption 2.1. *The problems in* (2.8) *are feasible and bounded for all x_{t-1}, realizations ξ^t, and $\theta \in \Theta$.*

Furthermore, to be able to utilize the results of the previous section, we will make the following blanket smoothness assumption in the rest of the paper.

Assumption 2.2. *The function $\theta \mapsto \mathcal{D}_t(\theta, \alpha)$ is differentiable, invertible, and bicontinuous for every $\alpha \in \mathbb{R}^{M_t}$ and every $t \in [T]$.*

We start our analysis by applying the results from Section 2.2 to a problem where the image measure of ξ is made up of finitely many atoms, i.e., where all conditional distributions $\xi_{t+1}|\xi^t$ are finitely supported. When dealing with discrete processes ξ, the expectations of the value functions V_t, which are concave functions of x_t, can be written as a minimum of finitely many affine functions as stated in the lemma below (for a proof see for example Philpott and Guan 2008, Shapiro 2011, Löhndorf et al. 2013).

Lemma 2.2. *If ξ is finitely supported, then, for every realization of ξ^t and fixed $\theta \in \Theta$, $x_t \mapsto \bar{V}_{t+1}(x_t, \xi^t, \theta)$ is a concave, piecewise linear function.*

For what follows, we assume that, for a given θ, the optimal policy $x_t^*(\theta, \xi^t)$ is known.

Theorem 2.3. *Consider a problem of the form* (2.7) *based on a discrete process $\xi = (\xi_1, \ldots, \xi_T)$ with optimal value $V(\theta)$ and optimal policy x_t^*. Then for all but finitely many $\theta \in \Theta$ the derivative of the optimal value exists at θ and can be written as*

$$V'(\theta) = \mathbb{E}\left[\sum_{t=1}^{T} \langle \nabla c_t(\theta), x_t^* \rangle + \left\langle \lambda_t^*, \nabla h_t(\theta) - \frac{d}{d\theta} T_t(\theta) x_{t-1}^* - \frac{d}{d\theta} W_t(\theta) x_t^* \right\rangle \right], \qquad (2.9)$$

where $\lambda_t^* = \lambda_t^*(\theta, \xi^t)$ are the optimal dual solutions associated with the constraints in stage t. In particular, $V'(\theta)$ exists for (Lebesgue) almost all $\theta \in \Theta$.

Proof. Proof. We start our argument in the last stage T where, for given ξ^T, θ, and x_{T-1}, the optimization problem is deterministic and equal to

$$V_T(x_{T-1}, \xi^T, \theta) = \begin{cases} \max\limits_{x_T} & \langle c_T(\theta), x_T \rangle \\ \text{s.t.} & W_T(\theta) x_T + T_T(\theta) x_{T-1} \leq h_T(\theta). \end{cases}$$

By Assumption 2.2, Theorem 2.2 can be applied to the problem above and it follows from Theorem 2.1 that for almost all θ

$$\frac{\partial}{\partial\theta} V_T(x_{T-1}, \xi^T, \theta) = \langle \nabla c_T(\theta), x_T^* \rangle + \left\langle \lambda_T^*, \nabla h_T(\theta) - \frac{d}{d\theta} T_T(\theta) x_{T-1} - \frac{d}{d\theta} W_T(\theta) x_T^* \right\rangle \tag{2.10}$$

for the optimal primal and dual solutions x^* and λ^*.

By Pollard (2001), Chapter 2, Example 23 and the fact that V_T is smooth in a neighborhood of a point θ where it is differentiable and that the problem is bounded, we have that

$$\frac{\partial}{\partial\theta} \mathbb{E}\left[V_T(x_{T-1}, \xi^T, \theta) \middle| \xi^{T-1} \right] = \mathbb{E}\left[\frac{\partial}{\partial\theta} V_T(x_{T-1}, \xi^T, \theta) \middle| \xi^{T-1} \right]. \tag{2.11}$$

Note that the set of problematic θ, i.e., the points $\theta \in \Theta$ where V_T is not smooth in θ, changes with each realization of ξ^T. However, since the underlying process is finitely supported, by Theorem 2.2 there are only finitely many θ where any of the functions $\theta \mapsto V_T(x_{T-1}, \xi^T, \theta)$ are not differentiable.

Theorem 2.1 and Theorem 2.2 apply to the problem defining $V_{T-1}(x_{T-2}, \xi^{T-1}, \theta)$, since the post-decision value function is a piecewise linear concave function by Lemma 2.2 and therefore can be represented in a linear optimization formulation. Combining (2.10) and (2.11) yields

$$\begin{aligned} \frac{\partial}{\partial\theta} V_{T-1}(x_{T-2}, \xi^{T-1}, \theta) = {} & \langle \nabla c_{T-1}(\theta), x_{T-1}^* \rangle + \mathbb{E}[\langle \nabla c_T(\theta), x_T^* \rangle | \xi^{T-1}] \\ & + \left\langle \lambda_{T-1}^*, \nabla h_{T-1}(\theta) - \frac{d}{d\theta} T_{T-1}(\theta) x_{T-2} - \frac{d}{d\theta} W_{T-1}(\theta) x_{T-1}^* \right\rangle \\ & + \mathbb{E}\left[\langle \lambda_T^*, \nabla h_T(\theta) - \frac{d}{d\theta} T_T(\theta) x_{T-1} - \frac{d}{d\theta} W_T(\theta) x_T^* \rangle \middle| \xi^{T-1} \right]. \end{aligned}$$

Continuing in this fashion, we arrive at

$$V'(\theta) = \mathbb{E}\left[\sum_{t=1}^{T} \langle \nabla c_t(\theta), x_t^* \rangle + \left\langle \lambda_t^*, \nabla h_t(\theta) - \frac{d}{d\theta} T_t(\theta) x_{t-1}^* - \frac{d}{d\theta} W_t(\theta) \right\rangle \right].$$

Note that, since there are finitely many stages and in each stage there are finitely many points $\theta \in \Theta$ where the corresponding value functions are not differentiable, the union of these points still has finite cardinality and therefore is a Lebesgue null set. □

As mentioned in Section 2.2, the fact that the derivative exists (Lebesgue) almost surely guarantees that the probability of picking a θ at random where $V'(\theta)$ does not exist is zero. This implies that this possibility can be ignored in practice. Of course, this argument is predicated on the tacit assumption of a continuous distribution on Θ governing the random choice of θ. If the sampling distribution for some reason would have atoms at the discontinuities of the value function, the situation would obviously change. This might for example occur if θ is itself the outcome of a higher level optimization problem, i.e., if the system in question was optimally designed to solve the stochastic optimization problem.

The derivative in (2.9) can be readily computed, if the stochastic optimization problem is solved as one large monolithic linear program, via a deterministic equivalent formulation, for example using scenario trees. In this case all primal and dual solutions for all possible scenarios are known and (2.9) can be easily evaluated.

Next, we investigate the case when ξ has a continuous distribution. Since in this case the post-decision value functions are no longer piecewise linear, the above proof does not work. We circumvent this issue by approximating the continuous problem by discrete problems based on an increasing number independent identically distributed (i.i.d.) samples from the respective distributions. We then show that the derivatives of the approximating problems calculated using Theorem 2.3 converge to the derivatives of the continuous problem.

Theorem 2.4. *Let the following conditions hold for a problem of the form* (2.7) *with continuously distributed randomness* ξ*:*

1. *The feasible set of the problems* (2.8) *is bounded for all* x_{t-1}*, all realizations* ξ^t*, and all* $\theta \in \Theta$*.*

2. The set Θ is bounded.

3. $\mathcal{D}_t$ is Lipschitz for every α with $|\mathcal{D}_t(\theta,\alpha)-\mathcal{D}_t(\theta',\alpha)| \leq L_t(\alpha)|\theta-\theta'|$ and $\mathbb{E}[L_t(\alpha)] = L_t < \infty$.

4. For every t there exists an x_t^+ such that x_t^+ is in the interior of the feasible set of the problem $V_t(x_{t-1},\xi^t,\theta)$ for all x_{t-1}, ξ^t, and $\theta \in \Theta$.

Then $V'(\theta)$ can be computed by formula (2.9) *as in Theorem 2.3 for all but countably many $\theta \in \Theta$. In particular, the derivative $V'(\theta)$ exists for (Lebesgue) almost all $\theta \in \Theta$.*

Proof. Proof. Note that (2.10) still holds. For a given θ, the problem in stage $(T-1)$ can be written as

$$V_{T-1}(x_{T-2},\xi^{T-1},\theta) = \begin{cases} \max\limits_{x_{T-1}} & \langle c_{T-1}(\theta), x_{T-1}\rangle + \mathbb{E}[V_T(x_{T-1},\xi^T,\theta)|\xi^{T-1}] \\ \text{s.t.} & W_{T-1}(\theta)x_{T-1} + T_{T-1}(\theta)x_{T-2} \leq h_{T-1}(\theta). \end{cases}$$

We draw K i.i.d. samples $(\hat{\xi}_T^1,\ldots,\hat{\xi}_T^K)$ from the conditional distribution $\mathbb{P}_T|\xi^{T-1}$ of ξ_T given ξ^{T-1} and approximate $V_{T-1}(x_{T-2},\xi^{T-1},\theta)$ by

$$\hat{V}_{T-1}^K(x_{T-2},\xi^{T-1},\theta) = \begin{cases} \max\limits_{x_{T-1}} & \langle c_{T-1}(\theta), x_{T-1}\rangle + K^{-1}\sum_{k=1}^K V_T(x_{T-1},\hat{\xi}_T^k,\theta) \\ \text{s.t.} & W_{T-1}(\theta)x_{T-1} + T_{T-1}(\theta)x_{T-2} \leq h_{T-1}(\theta), \end{cases}$$

which, due to Theorem 2.3 has derivative

$$\begin{aligned}\frac{\partial}{\partial\theta}\hat{V}_{T-1}^K(x_{T-2},\xi^{T-1},\theta) = {} & \langle \nabla c_{T-1}(\theta), x_{T-1}^{K*}\rangle + K^{-1}\sum_{k=1}^K \langle \nabla c_T^k(\theta), x_T^*\rangle \\ & + \left\langle \lambda_{T-1}^{K*}, \nabla h_{T-1}(\theta) - \frac{d}{d\theta}T_{T-1}(\theta)x_{T-2} - \frac{d}{d\theta}W_{T-1}(\theta)x_{T-1}^{K*}\right\rangle \\ & + K^{-1}\sum_{k=1}^K \langle \lambda_T^*, \nabla h_T^k(\theta) - \frac{d}{d\theta}T_T^k(\theta)x_{T-1}^{K*} - \frac{d}{d\theta}W_T^k(\theta)x_T^{k*}\rangle,\end{aligned}$$

where x_{T-1}^{K*} and λ_{T-1}^{K*} are the here-and-now decisions for the sampled problems in stage $T-1$, while x_T^{k*} and λ_T^{k*} are the scenario dependent wait-and-see solutions in stage T.

Our aim is to show that

$$\frac{\partial}{\partial\theta}\hat{V}^K_{T-1}(x_{T-2},\xi^{T-1},\theta) \xrightarrow{K\to\infty} \frac{\partial}{\partial\theta}V_{T-1}(x_{T-2},\xi^{T-1},\theta) \tag{2.12}$$

almost surely. According to Rudin (1964), Theorem 7.17 it is enough to show uniform convergence of $f_K(\theta) := \frac{\partial}{\partial\theta}\hat{V}^K_{T-1}(x_{T-2},\xi^{T-1},\theta)$. We start by demonstrating the uniform convergence of the term

$$g_K(\theta) = K^{-1}\sum_{k=1}^{K}\langle\lambda^{k*}_T(\theta), \frac{d}{d\theta}W^k_T(\theta)x^{k*}_T(\theta)\rangle. \tag{2.13}$$

Clearly, by the boundedness of the feasible set (assumption 1), x^{k*}_T are bounded and, by the uniform Slater condition (assumption 4), λ^{k*}_T are bounded. Consequently, there is a $B \in \mathbb{R}^+$ that bounds all the primal and dual decisions. Therefore, we have

$$|g_K(\theta) - g_K(\theta')| \le \frac{B}{K}\sum_{k=1}^{K}\sum_{ij}\left|\frac{d}{d\theta}W^k_T(\theta)_{ij} - \frac{d}{d\theta}W^k_T(\theta')_{ij}\right| \le \frac{B}{K}\sum_{k=1}^{K}\sum_{ij}|\theta-\theta'|L^k_{ij}$$
$$\xrightarrow{K\to\infty} B|\theta-\theta'|\sum_{ij}L_{ij}, \quad \mathbb{P}-a.s.,$$

where $\mathbb{P}-a.s.$ indicates almost sure convergence, L^k_{ij} is the Lipschitz constant of $\frac{d}{d\theta}W^k_T(\theta)_{ij}$, i.e., of the element in the i-th row and the j-th column of the matrix for the k-th sample and L_{ij} is its expectation (assumption 3). Note that the almost sure convergence above is due to an application of the law of large numbers. It follows that the Lipschitz constants of the term on the left are almost surely uniformly bounded, i.e., the term is uniformly Lipschitz in K. The uniform Lipschitz property of the other terms in $f_K(\theta)$ follows by analogous reasoning.

From the uniform Lipschitz property of f_K it follows that f_K are equicontinuous. Now, since Θ is bounded (assumption 2), the Arzelà-Ascoli theorem (see Rudin 1987, Theorem 11.28) establishes the uniform convergence of a subsequence f_{K_m} of f_K. Note that by Theorem 2.3 for every m, f_{K_m} is differentiable for all but finitely many $\theta \in \Theta$. It follows that all f_{K_m} are differentiable everywhere except on an unchanging countable set of points $A \subseteq \Theta$. Since the f_{K_m} also converge uniformly, f is differentiable for every $\theta \in \Theta \setminus A$ and (2.12) follows. The rest of the proof proceeds by backward induction in the same manner as the proof of Theorem 2.3. □

2.4 Sampling Derivatives

In this section, we discuss SDDP-type decomposition algorithms with an emphasis on *Approximate Dual Dynamic Programming* (ADDP), which is based on scenario lattices as discretizations for Markov processes. We show how ADDP can be used to approximate a policy for an MSLP with continuous randomness and demonstrate how the results from Section 2.3 can be used in sampling-based algorithms to sample derivatives of the objective values, which converges to the true derivatives as the approximation gets better.

2.4.1 Approximate Dual Dynamic Programming

ADDP solves linear MDPs with relatively complete recourse by iteratively approximating the value functions of the problem (Löhndorf et al. 2013, Löhndorf and Shapiro 2019) and is an extension of the popular SDDP algorithm (Pereira and Pinto 1991, Philpott and Guan 2008, Shapiro 2011). In this section, we will review the basic workings of ADDP and prove some novel results on the convergence of approximations of problems with continuous random variables by problems that are formulated using discrete representations of the underlying stochastic process and are solved using ADDP.

In the following, we will assume ξ to be a Markov process and the data of the problem c_t, W_t, T_t, and h_t to only depend on the current value ξ_t instead of the entire history ξ^t until time t. This renders problem (2.7) an MDP. We call ξ_t the environmental state of the problem, while x_{t-1} is called the resource state. Together

ξ_t and x_{t-1} constitute the state of the MDP. Note that the resource state can be influenced by the decisions in earlier stages while the environmental state evolves independent of the actions of the decision maker.

This conceptual separation enables us to approximate the value functions in two steps: First, we search for a good set of representative discrete states for the environmental state, which we organize in a *scenario lattice*. Second, we use a version of SDDP that approximates the value function at each node of the lattice by a concave, piecewise linear function of the resource state.

A scenario lattice is a graph organized in layers, each associated with a discrete point in time. A node represents a possible state of the stochastic process, and an arc represents the possibility of a state transition from one node on a given layer to a successor node on the next layer. Consequently, arcs only connect nodes in successive layers. Each arc is associated with a probability weight, and the weights of all outgoing arcs of a node model the distribution of the process conditional on that node. In contrast to scenario trees, we do not impose the restriction that every node in stage t has only one predecessor in stage $(t-1)$. For this reason, scenario lattices are sometimes called *recombining* scenario trees.

The goal of lattice construction is to build approximations of Markov processes such that the optimal policy for the lattice process yields a close to optimal policy for the true underlying process. For the purpose of this paper, we use a stochastic gradient algorithm outlined in Bally and Pagès (2003) to construct scenario lattices from simulations of the environmental state ξ of the problem. We denote the lattice process by $\tilde{\xi} = (\tilde{\xi}_1, \dots, \tilde{\xi}_T)$, where the nodes in stage t are denoted by $\tilde{\xi}_{tn}$, $n \in [C_t]$. In the following, we study the approximation properties of scenario lattices when the real underlying stochastic process is continuous and is compactly supported.

The method of Bally and Pagès (2003) chooses lattice nodes in order to minimize the Wasserstein distance between the unconditional distributions of $\xi_t : \Omega \to \mathbb{R}^{M_t}$ and $\tilde{\xi}_t : \tilde{\Omega} \to \mathbb{R}^{M_t}$. Denoting $\mathbb{P}_t$ and $\tilde{\mathbb{P}}_t$ as the image measures of ξ_t and $\tilde{\xi}_t$, respectively, the Wasserstein distance is defined as the solution of the following optimization problem (see, for example, Villani 2008)

$$W_2(\mathbb{P}_t, \tilde{\mathbb{P}}_t) = \begin{cases} \inf_{\pi} & \left(\int\limits_{\mathbb{R}^{M_t} \times \tilde{\mathbb{R}}^{M_t}} \|\xi - \tilde{\xi}\|_2^2 \, \pi\left(d\xi_t, d\tilde{\xi}_t\right) \right)^{\frac{1}{2}} \\ \text{s.t.} & \pi(A \times \mathbb{R}^{M_t}) = \mathbb{P}_t(A) \quad \forall A \in \mathcal{A}, \\ & \pi\left(\mathbb{R}^{M_t} \times B\right) = \tilde{\mathbb{P}}_t(B) \quad \forall B \in \tilde{\mathcal{A}}, \end{cases} \tag{2.14}$$

where the infimum is taken over all probability measures π on $(\mathbb{R}^{M_t} \times \mathbb{R}^{M_t}, \mathcal{A} \otimes \tilde{\mathcal{A}})$ and $\mathcal{A}$, $\tilde{\mathcal{A}}$ are the σ-algebras on $\mathbb{R}^{M_t}$ generated by ξ_t and $\tilde{\xi}_t$, respectively. The Wasserstein distance can be interpreted as the minimal effort required to move the probability mass between $\mathbb{P}_t$ and $\tilde{\mathbb{P}}_t$. In what follows, we will write $W_2(\xi_t, \tilde{\xi}_t)$ instead of $W_2(\mathbb{P}_t, \tilde{\mathbb{P}}_t)$, where no confusion can arise.

It follows from Graf and Luschgy (2000), Lemma 3.1 that the problem of finding an optimal discretization $\tilde{\xi}_t$ of ξ_t with a fixed number of nodes C_t in terms of the Wasserstein distance is equivalent to solving the C_t-center problem, i.e., finding C_t centers $\tilde{\xi}_{tn}$, $n = 1, \ldots, C_t$ that are the atoms of $\tilde{\xi}_t$ and solve

$$\inf_{\tilde{\xi}_t} W_2^2(\tilde{\xi}_t, \xi_t) = \inf_{\tilde{\xi}_t} \int_{\mathbb{R}^{M_t}} \min_{1 \le n \le C_t} ||\tilde{\xi}_{tn} - \xi_t||_2^2 \, \mathbb{P}_t(d\xi_t). \tag{2.15}$$

We define the probability of $\tilde{\xi}_{tn}$ given $\tilde{\xi}_{t-1,m}$ on the scenario lattice as

$$\tilde{\mathbb{P}}_t(\tilde{\xi}_t = \tilde{\xi}_{tn} | \tilde{\xi}_{t-1} = \tilde{\xi}_{t-1,m}) = \mathbb{P}_t \left(\operatorname*{arg\,min}_{1 \le r \le C_t} ||\tilde{\xi}_{tr} - \xi_t||_2^2 = n \,\middle|\, \operatorname*{arg\,min}_{1 \le r \le C_{t-1}} ||\tilde{\xi}_{t-1,r} - \xi_{t-1}||_2^2 = m \right).$$

For the purpose of the next results, we denote by $\tilde{\xi}_t^{C_t}$ the optimal approximation of ξ_t with C_t centers. We start by proving that with an increasing number of nodes in the lattice the optimal discretization $\tilde{\xi}_t^{C_t}$ converges almost surely to ξ_t. We show in the second part of the lemma below that this property carries over to the conditional distributions at stage t.

Lemma 2.3. *1. For every realization ξ_t*

$$\tilde{\xi}_t^{C_t}(\xi_t) := \tilde{\xi}_{tn^*}^{C_t} \xrightarrow{C_t \to \infty} \xi_t, \quad \mathbb{P} - a.s., \tag{2.16}$$

with $n^* = \arg\min \left\{ ||\tilde{\xi}_{tn}^{C_t} - \xi_t||_2 : \; n \in [C_t] \right\}$.

2. If all conditional distributions $\xi_t|\xi_{t-1}$ have finite first moments, the Wasserstein distances between the conditional distributions $\tilde{\xi}_t^{C_t}|\tilde{\xi}_{t-1}^{C_{t-1}}$ and the true conditional distributions vanishes as C_t grows, i.e.,

$$W_2\left(\tilde{\xi}_t^{C_t}|\tilde{\xi}_{t-1}^{C_{t-1}}(\xi_{t-1}), \xi_t|\xi_{t-1}\right) \xrightarrow{C_t \to \infty} 0.$$

Proof. Proof. From Graf and Luschgy (2000), Lemma 6.1 it follows that for the optimal choice of centers

$$W_2(\tilde{\xi}_t^{C_t}, \xi_t) \xrightarrow{C_t \to \infty} 0. \tag{2.17}$$

Suppose the first point would not be true, i.e., that there is a ξ_t, a subsequence $C_{tk} \xrightarrow{k \to \infty} \infty$, and an $\epsilon > 0$ for which

$$||\xi_t - \tilde{\xi}_t^{C_{tk}}(\xi_t)||_2 > \epsilon, \quad \forall C_{tk}.$$

Define $A = \left\{a \in \mathbb{R}^{M_t} : ||a - \xi_t||_2 < \frac{\epsilon}{2}\right\}$ and $\delta = \mathbb{P}(A)$. If $\delta > 0$, then

$$W_2(\tilde{\xi}_t^{C_{tk}}, \xi_t) \geq \int_A ||\xi_t - \tilde{\xi}_t^{C_{tk}}(\xi_t)||_2 \, \mathbb{P}(d\xi_t) > \delta \frac{\epsilon}{2}, \quad \forall C_{tk} \in \mathbb{N},$$

which is in contradiction to (2.17) and therefore proves 1.

To prove the second point note that

$$W_2\left(\tilde{\xi}_t^{C_t}|\tilde{\xi}_{t-1}^{C_{t-1}}(\xi_{t-1}), \xi_t|\xi_{t-1}\right) \leq \int ||\xi_t - \tilde{\xi}_t^{C_t}(\xi_t)||_2 \, \mathbb{P}_t(d\xi_t|\xi_{t-1}). \tag{2.18}$$

Additionally, we have that $||\xi_t - \tilde{\xi}_t^{C_t}(\xi_t)||_2 \leq ||\xi_t - \tilde{\xi}_{t1}^{C_t}||_2$ and

$$\int ||\xi_t - \tilde{\xi}_{t1}^{C_t}||_2 \, \mathbb{P}_t(d\xi_t|\xi_{t-1}) < \infty,$$

since the first moments of the conditional distributions are finite. Hence, using the first part, the dominated convergence theorem can be applied to show that the right side of (2.18) converges to zero implying that

$$W_2\left(\tilde{\xi}_t^{C_t}|\tilde{\xi}_{t-1}^{C_{t-1}}(\xi_{t-1}), \xi_t|\xi_{t-1}\right) \xrightarrow{C_t \to \infty} 0,$$

which establishes 2. □

Next, we show that if the involved random variables are continuous and have bounded support, the problem on the scenario lattice converges to the *real* problem on the continuous process as we increase the number of nodes in the lattice. Subsequently, we will use this result to prove that approximated derivatives of objective values of stochastic optimization problems converge to the true derivatives as the approximations get *finer*.

To this end, we consider a sequence of approximating measures $(\tilde{\mathbb{P}}_1^C, \ldots, \tilde{\mathbb{P}}_T^C)_{C \in \mathbb{N}}$ such that for each $\tilde{\mathbb{P}}_t^C$ the number of atoms $C_t(C)$ of $\tilde{\mathbb{P}}_t^C$ goes to infinity as $C \to \infty$ and the atoms of $\tilde{\mathbb{P}}_t^C$ are chosen such that they minimize the Wasserstein distance to $\mathbb{P}_t$. Further, we define $\tilde{V}_t^C$ and $\bar{\tilde{V}}_t^C$ as the value functions and post-decision value functions of the approximated problems using $(\tilde{\mathbb{P}}_1^C, \ldots, \tilde{\mathbb{P}}_T^C)_{C \in \mathbb{N}}$ for the distributions of the random parameters. Note that in the following we suppress the dependence of $\tilde{V}_t^C$ and $\bar{\tilde{V}}_t^C$ on the parameter θ to avoid cluttered notation.

Theorem 2.5. *Let the support of* $\mathbb{P}_t$ *be contained in a compact set* $X \subseteq \mathbb{R}^{M_t}$ *for all* $t \in [T]$ *and* $\mathbb{P}_t$ *be absolutely continuous with respect to the Lebesgue measure and assume that* (2.7) *has fixed recourse. Then*

$$\tilde{V}_t^C(x_{t-1}, \tilde{\xi}_t^C(\xi_t)) \xrightarrow{C \to \infty} V_t(x_{t-1}, \xi_t), \quad \mathbb{P}_t - a.s., \quad \forall x_{t-1},\ \forall t \in [T] \tag{2.19}$$

and

$$\bar{\tilde{V}}_t^C(x_t, \tilde{\xi}_t^C(\xi_t)) \xrightarrow{C\to\infty} \bar{V}_t(x_t, \xi_t), \quad \mathbb{P}_t - a.s., \ \forall x_t, \ \forall t \in [T]. \tag{2.20}$$

Proof. Proof. We start by showing (2.19) for stage T and then proceed by backward induction for $t = T-1, \ldots, 1$ assuming that the statement was already proven for $t+1$.

For the last stage T

$$\tilde{V}_T^C(x_{T-1}, \tilde{\xi}_{Tn}^C) = V_T(x_{T-1}, \tilde{\xi}_{Tn}^C), \quad \forall n \in [C_T(C)], \ \forall x_{T-1},$$

since the optimization problems that define the two functions are identical. Because of Lemma 2.3, $\tilde{\xi}_T^C(\xi_T) \to \xi_T$ and since $\tilde{\xi}_T^C \mapsto \tilde{V}_T(x_{T-1}, \tilde{\xi}_T^C)$ is continuous almost everywhere, we get

$$\tilde{V}_T^C(x_{T-1}, \tilde{\xi}_T^C(\xi_T)) \xrightarrow{C\to\infty} V_T(x_{T-1}, \xi_T), \quad \mathbb{P}_t - a.s., \quad \forall x_{T-1}.$$

The optimization problem that defines $\tilde{V}_t^C(x_{t-1}, \tilde{\xi}_t^C(\xi_t))$ is

$$\tilde{V}_t^C(x_{t-1}, \tilde{\xi}_t^C(\xi_t)) = \begin{cases} \max\limits_{x_t} & \langle \tilde{c}_t^C, x_t \rangle + \tilde{\mathbb{E}}^C \left[\tilde{V}_{t+1}^C(x_t, \tilde{\xi}_{t+1}^C) | \tilde{\xi}_t^C = \tilde{\xi}_t^C(\xi_t) \right] \\ \text{s.t.} & W_t x_t + \tilde{T}_t^C x_{t-1} \leq \tilde{h}_t^C, \end{cases}, \tag{2.21}$$

where $\tilde{\mathbb{E}}^C$ is the expectation with respect to the measure $\tilde{\mathbb{P}}_t^C$ and the data $(\tilde{c}_t^C, W_t, \tilde{T}_t^C, \tilde{h}_t^C)$ is the data stored on the lattice node closest to ξ_t.

Because of our assumption of fixed recourse, the value functions $V_{t+1}(x_t, \xi_{t+1})$ are continuous in ξ_{t+1}. The same holds for the approximated value functions $\xi_{t+1} \mapsto \tilde{V}_{t+1}^C(x_t, \tilde{\xi}_{t+1}^C(\xi_{t+1}))$, since they are almost surely constant. For a fixed ξ_t, define linear operators $L_C : \mathcal{C}(X) \to \mathbb{R}$ and $L_0 : \mathcal{C}(X) \to \mathbb{R}$ on the Banach space of almost everywhere continuous functions $(\mathcal{C}(X), ||\cdot||_\infty)$ defined on the compact set X as

$$L_C(f) = \tilde{\mathbb{E}}^C\left[f|\tilde{\xi}_t^C = \tilde{\xi}_t^C(\xi_t)\right] \text{ and } L_0(f) = \mathbb{E}\left[f|\xi_t\right],$$

it follows from Lemma 2.3 and Villani (2008), Theorem 6.9 that

$$L_C(f) \xrightarrow{C\to\infty} L_0(f), \quad \forall f \in \mathcal{C}(X), \tag{2.22}$$

implying that the L_C are pointwise bounded. The uniform boundedness principle (e.g., Bourbaki et al. 1987, Theorem III.2.1) then implies that the L_C are equi-continuous.

We now write

$$\begin{aligned}\left|L_C\left(\tilde{V}_{t+1}^C(x_t,\cdot)\right) - L_0\left(V_{t+1}(x_t,\cdot)\right)\right| \leq & \left|L_C\left(\tilde{V}_{t+1}^C(x_t,\cdot)\right) - L_C\left(V_{t+1}(x_t,\cdot)\right)\right| \\ & + \left|L_C\left(V_{t+1}(x_t,\cdot)\right) - L_0(V_{t+1}(x_t,\cdot))\right| \xrightarrow{C\to\infty} 0,\end{aligned}$$

where the first term converges to zero because of the induction hypothesis and the equi-continuity of $\{L_C\}_{C\in\mathbb{N}}$ while the second term vanishes due to (2.22). This establishes that

$$\bar{\tilde{V}}_{t+1}^C(x_t, \tilde{\xi}_t^C(\xi_t)) \xrightarrow{C\to\infty} \bar{V}_{t+1}(x_t,\xi_t), \quad \mathbb{P}_t - a.s., \ \forall x_t, \ \forall t \in [T].$$

By Lemma 2.3 we therefore have that for every realization of ξ_t the objective function of $\tilde{V}_t(x_{t-1}, \tilde{\xi}_{tn}^C(\xi_t))$ converges pointwise to the objective function of $V_t(x_{t-1},\xi_t)$ in x_{t-1}. By the concavity of the objective function in x_{t-1} this convergence is uniform on compact subsets of the feasible set (see Rockafellar 1970, Theorem 10.8). From Proposition 7.30 in Shapiro (2009), it follows that for all x_{t-1}

$$\tilde{V}_t^C(x_{t-1}, \tilde{\xi}_t^C(\xi_t)) \xrightarrow{C_t\to\infty} V_t(x_{t-1},\xi_t),$$

establishing (2.20). □

We note that given an approximating scenario lattice, the approximate problem can be solved by SDDP to obtain value function approximations for every time period t and lattice node n. Finite convergence to the true solution of the approximate problem is shown in (Löhndorf et al. 2013, Löhndorf and Shapiro 2019).

To transfer the solution from the scenario lattice back to a solution on the original process ξ and obtain a solution for an observed trajectory $\xi_1, \dots, \xi_T$ of the original process, we compute decisions as follows

$$x_t^* \in \arg\max_{x_t} \left\{ \langle c_t, x_t \rangle + \bar{\tilde{V}}_{t+1,n^*}^C(x_t, \theta) : W_t x_t + T_t \hat{x}_{t-1}^* \leq h_t \right\}, \tag{2.23}$$

whereby $n^* = \arg\min \left\{ ||\tilde{\xi}_{tn}^C - \xi_t||_2 : \; n \in [C_t(C)] \right\}$. This implies that the problem is solved with the data $\mathcal{D}_t(\theta, \xi_t) = (\mathrm{vec}(W_t(\theta, \xi_t)), \mathrm{vec}(T_t(\theta, \xi_t)), c_t(\theta, \xi_t), h_t(\theta, \xi_t))$ determined by the sample ξ_t using the post-decision value function $x_t \mapsto \bar{\tilde{V}}_{t+1,n^*}^C(x_t, \theta)$ from the lattice node $\tilde{\xi}_{tn^*}^C$ that is closest to ξ_t. We refer to this procedure as "rounding to a lattice node". For use in Algorithm 1 below, we define

$$\mathcal{S}_t(x_{t-1}, \xi_t, \theta) = \left\{ (x_t^*, \lambda_t^*) : x_t^*, \; \lambda_t^* \text{ are primal and dual optimal solutions of } (2.23) \right\}.$$

Rounding to a lattice node is made possible by the fact that each node contains a value function which can be used to make decisions for all possible resource states in t.

2.4.2 An Algorithm to Sample Derivatives

As described in the last section, ADDP solves MSLPs time step by time step for a specific realization of ξ using the dynamic programming principle (2.23). The expressions for the derivatives found in Theorem 2.3 and Theorem 2.4, which are based on primal and dual solutions for *all* scenarios, therefore cannot be evaluated directly. Also, evaluating the derivative by solving the deterministic equivalent of a discrete MSLP is in most cases computationally intractable due to the large number of scenarios typically represented in a scenario lattice. However, the approximate value functions effectively define a policy that can be used to generate samples of optimal decisions, which can be used to approximate (2.9).

In particular, if the conditions for Theorem 2.3 or Theorem 2.4 are fulfilled and the true optimal policy of the problem is known, then the derivative $V'(\theta)$ can be approximated based on an i.i.d. sample $(\hat{\xi}_t^k)_{t\in[T],k\in[K]}$ of size $K \in \mathbb{N}$ from the process $\xi = (\xi_1, \ldots, \xi_T)$. To that end, denote the resulting data as $\mathcal{D}_t^k = (\text{vec}(W_t^k(\theta, \hat{\xi}_t^k)), \text{vec}(T_t^k(\theta, \hat{\xi}_t^k)), c_t(\theta, \hat{\xi}_t^k), h_t^k(\theta, \hat{\xi}_t^k))$, let x_{tk}^* and λ_{tk}^* be the optimal primal and dual solution for the corresponding scenarios, and define the sample average estimator

$$\hat{V}_K'(\theta) = K^{-1} \sum_{k=1}^{K} \sum_{t=1}^{T} \langle \nabla c_t^k, x_{tk}^* \rangle + \lambda_{tk}^* \left(\nabla h_t^k - \frac{d}{d\theta} T_t^k x_{t-1,k}^* - \frac{d}{d\theta} W_t^k x_{tk}^* \right), \qquad (2.24)$$

where we suppress the dependency of W_t, T_t, c_t and h_t on θ to avoid cluttered notation.

We formalize the preceding discussion in Algorithm 1. Apart from $\theta \in \Theta$ and an initial resource state x_0, the algorithm takes as inputs the optimal policy found by ADDP in the form of the lattice discretization $(\tilde{\xi}_{tn})_{t\in[T],n\in[C_t]}$ of the process and a set of value functions $(\bar{\tilde{V}}_{tn})_{t\in[T],n\in[C_t]}$ – one per node of the scenario lattice. The algorithm samples trajectories $(\hat{\xi}_t^k)_{k\in[K],t\in[T]}$ of the stochastic process (line 7) starting from the deterministic root node $\xi_1 = (W_1, T_1, c_1, h_1)$, which is common to all sampled scenarios. In particular, the function $\texttt{sample}_{t+1}(\hat{\xi}_t^k)$ returns a realization $\hat{\xi}_{t+1}^k$ of ξ_{t+1} conditional on $\hat{\xi}_t^k$. Subsequently, we solve problem (2.23) by finding the lattice node closest to the sampled scenario (line 5) and then using the respective post-decision value function stored at that node to make a decision. Note that, as in (2.23), the rest of the problem data is determined by $\hat{\xi}_t^k$.

In some cases the original process is assumed to be a discrete Markov process. This is for example the case in some applications of the classical SDDP algorithm or

Data: $K \in \mathbb{N}$, x_0, θ, $(\tilde{\xi}_{tn})_{t\in[T],n\in[C_t]}$, $(\hat{\tilde{V}}_{tn})_{t\in[T],n\in[C_t]}$
Result: $\hat{V}'_K(\theta)$

$\hat{V}'_K(\theta) \leftarrow 0$
for $k \leftarrow 1$ **to** K **do**
    $\hat{\xi}_1^k = \xi_1 = (W_1, T_1, c_1, h_1)$
    for $t \leftarrow 1$ **to** T **do**
        Find $(x^*_{tk}, \lambda^*_{tk}) \in \mathcal{S}_t(x^*_{t-1,k}, \hat{\xi}_t^k, \theta)$.
        $\hat{V}'_K(\theta) \leftarrow \hat{V}'_K(\theta) + \frac{1}{K}\left[\langle \nabla \hat{c}_t^k, x^*_{tk}\rangle + \lambda^*_{tk}\left(\nabla \hat{h}_t^k - \frac{d}{d\theta}\hat{T}_t^k x^*_{t-1,k} - \frac{d}{d\theta}\hat{W}_t^k x^*_{tk}\right)\right]$
        $\hat{\xi}_{t+1}^k = (\hat{W}_{t+1}^k, \hat{T}_{t+1}^k, \hat{c}_{t+1}^k, \hat{h}_{t+1}^k) \leftarrow \texttt{sample}_{t+1}(\hat{\xi}_t^k)$
    end
end

Algorithm 1: Pseudo-code for the computation of derivatives in ADDP.

if a scenario lattice is directly estimated from data without the intermediary step of a statistical model that is sampled to generate the lattice (see for example Löhndorf et al. 2013). In this case, by the results in Theorem 2.3 and by the law of large numbers, $\hat{V}'_K(\theta)$ converges almost surely to $V'(\theta)$ for (almost) all $\theta \in \Theta$ as $K \to \infty$.

Note that in these cases the problems can be represented as deterministic equivalents as discussed in Section 2.3. However, even for medium sized lattices with more than a few stages the corresponding deterministic equivalents are too large. Since the approximate value functions found by ADDP converge to the true value functions of the problem in finitely many iterations (Löhndorf et al. 2013, Löhndorf and Shapiro 2019), (2.24) can directly be used to approximate derivatives of the optimal value.

When approximating the continuous process ξ by a scenario lattice $\tilde{\xi}$, we obtain policies that are optimal for $\tilde{\xi}$, instead of the actual process ξ. Hence, the resulting derivatives will be an approximation of the real ones, which, if the approximation of ξ by $\tilde{\xi}$ is good, can be reasonably expected to be sufficiently accurate. We formalize this intuition for the case of compactly supported randomness in the following theorem showing that the approximated derivatives of the approximated problem $\left(\hat{\tilde{V}}^{C(r)}_{K(r)}\right)'(\theta)$ converge to the real derivatives of the problem with continuous randomness as the number of lattice nodes $C_t(C)$ in every stage of the approximate problem and the number samples for the computation of (2.24) go to infinity.

Theorem 2.6. *Suppose the conditions of Theorem 2.4 and Theorem 2.5 hold and define sequences* $C(r) \xrightarrow{r\to\infty} \infty$ *and* $K(r) \xrightarrow{r\to\infty} \infty$*. Then*

$$\lim_{r\to\infty} \left(\hat{\tilde{V}}^{C(r)}_{K(r)}\right)'(\theta) = V'(\theta), \quad \mathbb{P}-a.s.$$

Proof. Proof. For a given realization $\hat{\xi}_1, \ldots, \hat{\xi}_T$ of the continuous stochastic process ξ, we define the rounded problems as in (2.23)

$$\tilde{V}_t^{C(r)}(x_{t-1}, \hat{\xi}_t) = \max_{x_t} \left\{ \langle \hat{c}_t, x_t \rangle + \bar{\tilde{V}}^{C(r)}_{t+1,n^*}(x_t) : \hat{W}_t x_t + \hat{T}_t x_{t-1} \leq \hat{h}_t \right\}, \tag{2.25}$$

with $n^* = \arg\min \left\{ ||\tilde{\xi}^{C(r)}_{tn} - \tilde{\xi}_t||_2 : \; n \in [C_t(C(r))] \right\}$ and compare their solutions $\tilde{x}_t^{C*}$ with the solutions x_t^* of the continuous problem for the sample path.

From Theorem 2.5 we know that

$$\bar{\tilde{V}}^{C(r)}_{t+1}(x_t, \tilde{\xi}_t^{C(r)}(\xi_t)) \xrightarrow{r\to\infty} \bar{V}_{t+1}(x_t, \xi_t), \quad \forall x_t, \; \mathbb{P}-a.s.,$$

which, together with Lemma 2.3, implies that the objective function of $\tilde{V}_t^{C(r)}(x_{t-1}, \tilde{\xi}_t^{C(r)}(\xi_t))$ converges almost surely to the objective function of $V_t(x_{t-1}, \xi_t)$ as the lattice approximation gets finer. Since the objective functions are concave they converge uniformly on compact sets (see Rockafellar 1970, Theorem 10.8).

For the first stage, the constraints of problem (2.25) and $V_1(x_0, \xi_1)$ are the same and therefore it follows from Proposition 7.30 in Shapiro (2009) that the first stage solutions $\tilde{x}_1^{C*}$ converge to a first stage solution x_1^*.

Proceeding by induction over t, we notice that for a given sample $\hat{\xi}_1, \ldots, \hat{\xi}_T$ the feasible sets of the problems $\tilde{V}_t^{C(r)}$ and V_t differ due to the difference in the respective last stage decisions x_{t-1} and $\tilde{x}_{t-1}^{C*}$ resulting in different right hand sides of the problems. Denote the feasible set of V_t by $\mathcal{C}_t$ and the feasible set of $\tilde{V}_t^{C(r)}$ by $\tilde{\mathcal{C}}_t^r$. Consider an arbitrary point x with distance $\epsilon > 0$ to the boundary of $\mathcal{C}_t$. Because of

Hoffman's Lemma (see Shapiro et al. 2014), we get that, eventually, $x \in \mathcal{C}_t \cap \tilde{\mathcal{C}}_t^r$ or $x \notin \mathcal{C}_t \cup \tilde{\mathcal{C}}_t^r$, due to the fact that by the induction hypothesis $\tilde{x}_{t-1}^{C(r)*}$ converge almost surely to x_{t-1}^*. Since the complement of the boundary of $\mathcal{C}_t$ is dense in $\mathbb{R}^{N_t}$, it follows by Proposition 7.31 in Shapiro (2009) that

$$\langle c_t, x_t \rangle + \bar{\tilde{V}}_{t+1,k}^{C(r)}(x_t) + \mathbb{1}_{\tilde{\mathcal{C}}_t^r}(x_t)$$

epi-converges to

$$\langle c_t, x_t \rangle + \bar{V}_{t+1,k}(x_t) + \mathbb{1}_{\mathcal{C}_t}(x_t)$$

as $r \to \infty$. Therefore, $\tilde{x}_t^{C*}$ converge to an optimal solution x_t^* due to Proposition 7.30 in Shapiro (2009).

The realizations of randomness where there is a basis change are the only ones where the optimal solution x_t^* is potentially not unique. This implies that the solutions converge to a unique solution of the continuous problem almost surely. An analogous argument shows that the dual solutions $\lambda_t^{C(r)*}$ converge to the dual solutions of λ_t^*.

Now we can write

$$\left| V'(\theta) - \left(\hat{\tilde{V}}_{K(r)}^{C(r)} \right)' (\theta) \right| \leq |V'(\theta) - (\tilde{V}^{C(r)})'(\theta)| + \left| (\tilde{V}^{C(r)})'(\theta) - \left(\hat{\tilde{V}}_{K(r)}^{C(r)} \right)' (\theta) \right|. \quad (2.26)$$

The first term on the right side of (2.26) converges to zero by an application of the dominated convergence theorem, since the primal and dual solutions $\tilde{x}_t^{C(r)*}$, $\tilde{\lambda}_t^{C(r)*}$ converge almost surely to x_t^* and λ_t^* respectively and the derivatives ∇c, $\frac{d}{d\theta} T_t$, $\frac{d}{d\theta} W_t$, and ∇h_t are bounded by our assumption on the compact support of all involved random variables and the smoothness of the coefficient functions.

To see that the second term converges to zero almost surely, note that by the boundedness of the random variables, Hoeffding's inequality establishes that

$$\mathbb{P}\left(\left| (\tilde{V}^{C(r)})'(\theta) - \left(\hat{\tilde{V}}_{K(r)}^{C(r)} \right)' (\theta) \right| > K(r)^{-1/4} \right) \leq 2e^{-\gamma\sqrt{K(r)}}$$

for some $\gamma > 0$. Since $\sum_r e^{-\gamma\sqrt{K(r)}} < \infty$, it follows from the Borel-Cantelli lemma (Billingsley 1986, Theorem 4.3) that

$$\hat{\hat{V}}^{C(r)}_{K(r)} \xrightarrow{r\to\infty} (\tilde{V}^{C(r)})'(\theta), \quad \mathbb{P} - a.s.,$$

which finishes the proof. □

Finally, we remark that if the second moment of $V_K'(\theta)$ exists and is positive, then, by the central limit theorem,

$$\frac{\sqrt{K}}{\sigma}(\hat{V}_K'(\theta) - V'(\theta)) \xrightarrow{d} \mathcal{N}(0,1), \text{ as } K \to \infty, \tag{2.27}$$

where $\xrightarrow{d}$ denotes convergence in distribution and $\sigma > 0$ is the standard deviation of $\hat{V}_K'$. This allows for the construction of confidence regions around the sampled derivatives $\hat{V}_K'(\theta)$. The existence of second moments, which is a condition for the application of the central limit theorem, can, for example, be established by the boundedness of optimal solutions in combination with $\nabla_\theta \mathcal{D}_t \in \mathcal{L}^1(\Omega, \mathcal{F}_t)$ for all $t \in [T]$. These conditions are usually fairly easy to check in real-world examples as will be demonstrated in Section 2.5.

2.5 Numerical Examples

In this section, we calculate derivatives for optimal values in two examples. As benchmarks, we use analytical solutions in one of the problems and naïve estimates for the derivatives, which are based on difference quotients.

More specifically, for a problem with an optimal value $V(\theta)$, we calculate the symmetric difference quotient

$$\hat{V}_{DQ}'(\theta) = \frac{V(\theta+\epsilon) - V(\theta-\epsilon)}{2\epsilon} \tag{2.28}$$

as an approximation of the derivative. We use the symmetric difference, since it yields a more stable estimate of the derivative than the one-sided difference quotient. Nevertheless, the expression (2.28) is susceptible to several distortions. Firstly, for any given $\epsilon > 0$, $\hat{V}'_{DQ}(\theta)$ is only a (biased) approximation of $V'(\theta)$, which gets better as $\epsilon \to 0$. Secondly, for small ϵ, errors in the evaluation of $\hat{V}'_{DQ}$ are amplified due to the numerical instability of the expression, which arises because of the division by the small number ϵ. Thirdly, the values of $V(\theta \pm \epsilon)$ are calculated using a sampling-based algorithm and are therefore random, which introduces another level of inaccuracy.

Dealing with these problems involves finding an ϵ which represents a good trade-off between the bias introduced by larger values of ϵ and the variance introduced by smaller levels of ϵ. Unfortunately, there is no exact method to find an optimal trade-off between these two opposing effects. We choose ϵ using the step size estimator

$$\epsilon = \Psi S^{-\frac{2}{5}}$$

discussed in Section 7.1 of Glasserman (2004) under the name of $\epsilon_{C,ii}$ by setting the unknown constant Ψ to 1. In the above formula n is the number of samples used to calculate $V(\theta + \epsilon)$ and $V(\theta - \epsilon)$ by the ADDP lower bound and is set to $S = 10,000$ in all our examples. Furthermore, we mitigate the problem of randomness in the estimate by using the same sample $\hat{\xi}$ for the evaluation of $V(x + \epsilon)$ and $V(x - \epsilon)$. As argued in Shapiro et al. (2014), this reduces the variance of the estimate $\hat{V}'_{DQ}$, provided that the covariance of $V(x + \epsilon)$ and $V(x - \epsilon)$ is positive, which is obviously true in our case.

Furthermore, we benchmark our method against the Richardson extrapolation (RE) formula

$$\hat{V}'_{RE}(\theta) = \frac{1}{3}\left(4\frac{V(\theta + \frac{\epsilon}{2}) - V(\theta - \frac{\epsilon}{2})}{\epsilon} - \frac{V(\theta + \epsilon) - V(\theta - \epsilon)}{2\epsilon}\right),$$

that is an alternative to $\hat{V}'_{DQ}$ which sometimes decreases the bias of the naïve difference quotient as argued in Glasserman (2004).

All computations presented in this section have been performed with MATLAB and the stochastic optimization problems have been solved using the MATLAB interface of QUASAR (see www.quantego.com), which provides an implementation of ADDP as a JAVA library.

2.5.1 The Newsvendor Problem

We first study a two-stage newsvendor problem, which is a classical example in stochastic optimization reminiscent of basic procurement and inventory problems. Although the problem only has two stages, it is still useful to us, since it is one of the few stochastic optimization problems with a closed-form solution, allowing us to compare approximated derivatives to the true derivatives of the optimal value.

The problem is set up as follows: The newsvendor decides in the morning how many newspapers she will purchase from the publisher, not knowing the demand during the day. She sells the papers to her clients for a known price, and excess newspapers have to be discarded at the end of the day. The problem is cast as a two-stage stochastic optimization problem: in the first stage the newsvendor decides how many papers to order, in the second stage (knowing the demand) she decides how many papers to sell and how many to discard.

Let x_b be the amount of ordered newspapers, x_s be the amount of papers sold, and D be the random demand during the day. Then the problem can be written as the following two-stage stochastic linear problem

$$V = \begin{cases} \max\limits_{x_b, x_s} & -cx_b + \mathbb{E}[px_s] \\ \text{s.t.} & x_s \leq x_b, \quad \text{a.s.} \\ & x_s \leq D, \quad \text{a.s.} \\ & x_b \geq 0, \end{cases}$$

where p is the sales price and c is the price the newsvendor pays to the publisher.

The Lagrangian of the problem is

$$\mathcal{L}(x_b, x_s, \lambda_1, \lambda_2, \lambda_3) = -cx_b + \lambda_1 x_b + \mathbb{E}\left[px_s + \lambda_2(x_b - x_s) + \lambda_3(D - x_s)\right]. \tag{2.29}$$

Assuming that $D \sim F$, it follows from elementary calculation (see Birge and Louveaux 2011) that the optimal solution and optimal value of the problem are

$$x_b^* = F^{-1}\left(\frac{p-c}{p}\right), \quad V = -cx_b^* + p\int_{-\infty}^{\infty} \min(x, x_b^*)\, dF(x), \tag{2.30}$$

respectively.

For our computations, we will use $p = 1$, $c = 0.2$ and a normal distribution with mean $\mu = 100$ and $\sigma = 20$ to model the random demand D. For later use, we write $D = \mu + \sigma X$, where $X \sim N(0,1)$. We discretize the continuous demand process into a scenario lattice with 100 nodes in the second stage using the subgradient method in Bally and Pagès (2003) for 100,000 iterations. The ADDP algorithm is terminated after 100 iterations for all the calculations, which yields an average gap of below 1% between the ADDP upper bound and sampled value of the policy (lower bound) based on 10,000 samples.

We start by calculating derivatives with respect to the sales price p, which is an example where $\theta = p$ appears in the objective function coefficients. For this purpose, we use the linear function $p \mapsto p$, which clearly fulfills the requirements of Theorem 2.3 and Theorem 2.4. Furthermore, the setup of the problem together with the choice of the normal distribution guarantees the existence of second moments required to compute confidence intervals. According to these results, we have

$$V'(p) = \frac{\partial}{\partial p}\mathcal{L}(x_b^*, x_s^*, \lambda_1, \lambda_2, \lambda_3) = \mathbb{E}[x_s^*],$$

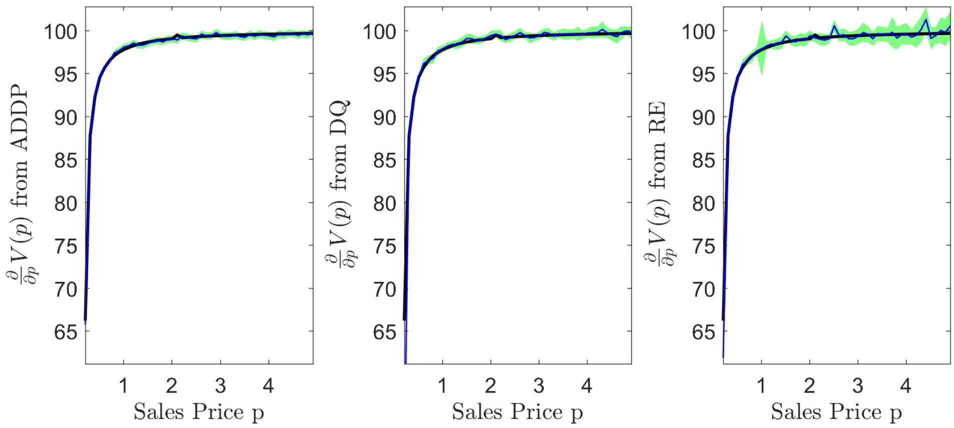

Figure 2.2: Derivative of the value function V with respect to sales price p as a function of p evaluated at 48 equally spaced points between 0.1 and 5. The bold black line is the true derivative.

where x_s^* are the optimal sales decisions. We obtain the approximation

$$V'(p) = \mathbb{E}[x_s^*] \approx K^{-1} \sum_{k=1}^{K} x_{ks}^* = \hat{V}'_K(p)$$

by drawing K =10,000 demand samples from $D \sim N(\mu, \sigma)$ and using the optimal policy found by ADDP in the scenario lattice to calculate optimal solutions x_{ks}^*. Note that the symbol $\approx$ signifies that the right hand side of the equation approximates the left hand side.

Figure 2.2 shows the derivative of the value function with respect to the sales price. The bold black curves correspond to the true derivative computed using the analytical solution. The plot on the left illustrates the results using the sampling-based method, the plot in the middle displays the numerical derivative computed using the difference quotient method, while the right plot shows the results of the Richardson extrapolation. The green regions are the confidence intervals around the respective solutions, which are calculated using (2.27) for $\hat{V}'_K(p)$ and normal confidence regions for the difference quotient with standard deviations calculated from 30 independent evaluations of $\hat{V}'_{DQ}$, each using common random numbers to compute $\hat{V}'_{DQ}$ for the range of displayed θs.

The shape of the derivative with respect to p can be interpreted as follows: As the sales price of the newspapers increases, the number of newspapers bought by the optimal policy increases. The value of this decision at first grows quite quickly, as the revenue in the second stage grows superlinearly as long as the probability of ordering too much is still relatively small. Asymptotically, the probability of having ordered too much goes to 1 while $\frac{\partial}{\partial p} x_b^* \to 0$ as $p \to \infty$ and therefore the increase in objective is dominated by the linear growth of the second stage revenue, which in turn leads to a constant $V'(p)$. More specifically, the optimal value converges to

$$-x_b^* c + p\mathbb{E}[D],$$

with its derivative with respect to p converging to $\mathbb{E}[D] = 100$, i.e., a constant.

The plot shows that all three methods perform reasonably well in capturing this pattern, but $\hat{V}'_K(p)$ is most precise in the sense that absolute deviations from the actual derivative are smaller for the ADDP-based method (0.1542 on average) than for the method based on the difference quotients (0.5394 on average) and for the Richardson extrapolation (0.3544 on average). Furthermore, we see that the confidence bounds for $\hat{V}'_K(p)$ are narrower than those of $\hat{V}'_{DQ}(p)$, indicating that the sampling-based estimate is more reliable. The Richardson extrapolation performs worst as it increases the variance of the estimate without any conceivable reduction in bias.

Turning to the issue of computational cost, we note that the computation of $\hat{V}'_{DQ}$ requires us to solve the problem twice, while the sampling based method in (2.24) requires only one solution of the stochastic optimization problem. Therefore, we expect the runtime of the computation of $\hat{V}'_{DQ}$ to be roughly twice of the runtime required to compute $\hat{V}'_K$. To obtain a measurement, we recorded the runtimes required for both estimates for the 48 measurements of $\hat{V}'_{DQ}$ and $\hat{V}'_K$ depicted in Figure 2.2. In line with the above, we find the ratio between the average computation time of $\hat{V}'_{DQ}$ and $\hat{V}'_K$ to be $r = 1.91 \pm 0.21$. This confirms that $\hat{V}'_K$ is computationally superior to $\hat{V}'_{DQ}$. Since the Richardson extrapolation formula is based in two difference

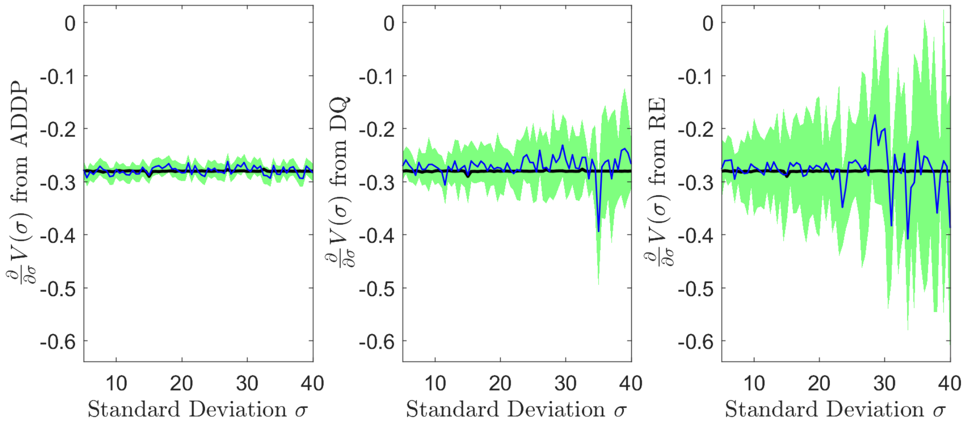

Figure 2.3: Derivative of the value function V with respect to σ as a function of σ evaluated at 71 equally spaced points between 5 and 40. The bold black line is the true derivative.

quotients, it requires the solution of four stochastic optimization problems and thus quadruples the computational effort required to compute $\hat{V}'_K$.

Next, we calculate the derivative of the optimal value function with respect to the standard deviation of the random demand, which is a parameter on the right-hand side of the constraints of the optimization problem. To that end, we use the linear function in $\sigma \mapsto \mu + \sigma X$, which fulfills the requirements of Theorem 2.3 and Theorem 2.4 as well as the conditions on the moments required for (2.27).

Taking the derivative of the Lagrangian with respect to σ yields

$$V'(\sigma) = \frac{\partial}{\partial\sigma}\mathcal{L}(x_b^*, x_s^*, \lambda_1^*, \lambda_2^*, \lambda_3^*) = \mathbb{E}[\lambda_3^* X], \tag{2.31}$$

where λ_3^* is the optimal dual solution corresponding to the constraint $x_s \leq D$. Hence, the derivative with respect to σ can be obtained by multiplying λ_3^* with the realizations $X \sim N(0,1)$, which define D. As before, to calculate derivatives, we approximate

$$V'(\sigma) = \mathbb{E}[\lambda_3^* X] \approx K^{-1}\sum_{k=1}^{K}\lambda_{3tk}^* X^k = \hat{V}'_K(\sigma).$$

The results of the comparison between $\hat{V}'_K(\sigma)$, $\hat{V}'_{DQ}(\sigma)$ and $\hat{V}'_{RE}(\sigma)$ is shown in Figure 2.3. We again use $K = 10{,}000$ samples for the estimation of $\hat{V}'_K(\sigma)$ and computed the confidence intervals for $\hat{V}'_{DQ}(\sigma)$ and $\hat{V}'_{RE}(\sigma)$ based on 30 independent repetitions of the calculation using common random numbers for sampling within each repetition.

The derivative is constant and negative with a value of about -0.28. Inspecting (2.31), we notice that the dual multiplier λ_3^* is either 0, when $x_b^* > D$ or equal to p if $x_b^* \leq D$. Hence, multiplication with λ_3^* effectively skews the symmetric distribution of X and results in an overall negative value of the expectation. The fact that the derivative is constant, follows from (2.30), which implies that the optimal order quantity is always the same quantile of the demand distribution, i.e., λ_3^* does not change with increasing volatility and consequently the expectation remains constant.

Looking at Figure 2.3, we see that our method clearly outperforms both benchmarks in terms of accuracy and variance of the estimator. In particular, the average absolute errors of $\hat{V}'_K(\sigma)$ of 0.0064 compare favorably to 0.0151 for the difference quotient and 0.0228 for the Richardson extrapolation. Lastly, we note that the results for the difference quotient are slightly upward biased, which is not the case for the Richardson extrapolation. However, since this comes at the cost of a significantly increased variance and a higher computational cost, we disregard this method in our further calculations.

As mentioned above, we solve MSLPs using scenario lattices as discretizations of the underlying stochastic process but use simulations from the actual process, in this case from the normal distribution, to calculate the derivatives by rounding to the nearest lattice node as described in Section 2.4. In the following, to evaluate the merit of this strategy, we systematically compare the errors (relative to the true derivatives), when computing $\hat{V}'_K(p)$ and $\hat{V}'_K(\sigma)$ based on simulations from the lattice and the true process for a varying number of nodes in the lattice and a varying number of samples used to calculate the derivatives. In particular, we use 50, 100, 200 lattice nodes per stage and 1,000, 10,000, and 100,000 samples to compute derivatives.

Since the calculated derivatives are subject to random variations, we repeat calculations 30 times, each time varying the parameter p from 0.3 to 5 in steps of 0.1 and the parameter σ from 5 to 40 in steps of 1. We repeat this analysis for all combinations of nodes and number of samples.

We define

$$\Delta_{lat}^{p,s} = \sum_{l=1}^{L} |\hat{V}'^{s}_{K,lat}(p_l) - V'(p_l)|, \quad \Delta_{proc}^{p,s} = \sum_{l=1}^{L} |\hat{V}'^{s}_{K,proc}(p_l) - V'(p_l)|,$$

where p_l are the prices for which the derivatives $\hat{V}'^{s}_{K,lat}$ and $\hat{V}'^{s}_{K,proc}$ are estimated based on samples from the lattices and the processes, respectively. We report the differences in average errors, i.e.,

$$\Delta^p = \frac{1}{30}\sum_{s=1}^{30} \Delta_{lat}^{p,s} - \Delta_{proc}^{p,s}$$

and test whether these quantities are significantly different from 0 using unpaired two-sample t-tests. The values for Δ^σ are computed in an analogous fashion.

Furthermore, we report the frequency $F^{lat,proc}$ with which the true derivative of the optimal value is outside the symmetric 98% confidence bounds in (2.27) around the estimates $\Delta_{lat,proc}^{p,s}$ and $\Delta_{lat,proc}^{\sigma,s}$. Deviations from the expected 2% violations occur for two reasons: Firstly, we solve an approximate (discrete) problem, which introduces a bias to the estimates. Secondly, when sampling from the lattice the samples do not come from the true distribution but from an approximation. By sampling from the *actual* normal distribution, we can avoid the second source of distortion.

The results of the analysis are reported in Table 2.1. Inspecting the outcomes for $\Delta^{p,\sigma}$, we see that the numbers are mostly positive, implying that the errors are larger when sampling from the scenario lattice. The only exception are two values for 50 and 100 nodes and 1,000 samples, which are the cases where the variance of the estimates is the highest, due to the relatively small number of samples. Correspondingly, we note that these two values are not significantly different from 0. As the number of samples increases the estimates from the process perform markedly better and differences are significant at least at the 5% level and unambiguously at the 0.1% level for estimates based on 100,000 samples. We conclude that the bias in the estimates can be consistently reduced when sampling from the true underlying stochastic process by rounding to the next lattice node as described in (2.23).

Turning our attention to the violations $F^{lat,proc}$ in the two right panels of Table 2.1, we observe that both F^{lat} and F^{proc} are decreasing in the number of nodes. Having more nodes improves the quality of the approximation by the scenario lattice and thereby reduces the bias between the true and the approximated solution, bringing F closer to its correct value of 2%. Furthermore, we observe that F increases in the number of samples used to calculate the derivatives. The reason for this is that confidence bands get narrower when the number of samples increases, and the discretization bias is therefore more often *detected*. Comparing the values for lattices and the process, we notice that, as expected, confidence intervals for the estimates based on the samples from the process are more accurate than those computed from lattice samples.

We also note that for both lattice- and process-based estimates, the absolute difference between the analytical derivative and the sampled derivative reduces with increasing sample size and with an increased number of nodes in the scenario lattice, both of which improve the approximation for obvious reasons. Results of this analysis are available upon request.

Summarizing, we remark that, if possible, it is preferable to sample from the true underlying process instead of from the discretization used to solve the stochastic optimization problem. We view the possibility to do so as a major advantage of SDDP-type solution methods, which yield policies that can be evaluated for arbitrary stochastic processes by rounding.

	Nodes	$\Delta^{p,\sigma}$			$100 \times F^{lat}$			$100 \times F^{proc}$		
Samples		1k	10k	100k	1k	10k	100k	1k	10k	100k
Price	50	-0.1238	0.2953*	0.5875***	2.8	7.3	21.7	2.6	6.1	17.6
	100	-0.2516	0.2904*	0.7348***	2.3	4.7	16.7	2.7	3.9	8.5
	200	0.1843	0.4501***	1.1861***	2.2	3.5	15.8	1.9	2.4	4.4
Std	50	0.0177	0.0725***	0.0821***	5.7	14.9	55.0	5.3	8.5	42.6
	100	0.0216	0.0362***	0.0568***	5.3	8.9	40.7	4.6	6.3	22.6
	200	0.0135	0.0099*	0.0261***	4.9	5.8	19.9	5.0	5.3	9.0

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.1: Comparison of average absolute errors for derivatives with respect to sales price and standard deviation of demand when sampling from the scenario lattice and the process in the first panel. Percentage violations of the confidence bounds in panel 2 and 3.

2.5.2 Gas Storage

In this section, we study the problem of gas storage pricing and operation on the spot market for natural gas hosted on the *National Balancing Point* (NBP) in the United Kingdom. The aim is to compute the value of owning a storage plant for one year based on optimal operational and trading strategies, which capitalize on the yearly seasonality of gas prices as well as on statistical arbitrage from short-term fluctuations in prices.

We formulate the problem as the following MSLP with random spot prices for gas P_t in stage t

$$V = \begin{cases} \max\limits_{x_t^b, x_t^s, l_t} & \sum_{t=1}^T \mathbb{E}\left[P_t(x_t^s - x_t^b)\right] \\ \text{s.t.} & l_t = l_{t-1} + \eta x_t^b - \eta^{-1} x_t^s, \quad t = 2, ..., T \\ & l_t \leq C, \quad t = 1, ..., T \\ & x_t^b,\ x_t^s,\ l_t \geq 0, \quad t = 1, ...T, \end{cases},$$

where x_t^s and x_t^b correspond to withdrawn and injected (sold and bought) quantities in stage t, respectively, l_t denotes the storage level, and all the constraints are required to hold almost surely. Furthermore, we assume that the storage has a maximal storage capacity of C and an efficiency factor $\eta \in [0, 1]$. For the purpose of our computations, we use $C = 1$ and $\eta = 1$, if not stated otherwise.

We model the stochastic gas spot prices as a geometric Ornstein-Uhlenbeck process (see Schwartz 1997), with time-dependent mean reversion levels μ_t given by the following stochastic differential equation

$$\begin{aligned} dP_s &= \kappa\big(\mu_s - \log(P_s)\big)P_s ds + \sigma P_s dW_s, \\ \mu_s &= A\sin(2\pi s + \phi), \end{aligned}$$

where time is normalized in such a way that one year corresponds to a time difference of 1. Note that the model captures yearly seasonalities in the gas price by mean reversion to the trigonometric function μ_s, where A is the amplitude and ϕ the phase of the function, while the periodicity is fixed to one year.

As is common in pricing, we work with the risk-neutral measure by adding a constant market price of risk λ to the deterministic trend (see Schwartz 1997), i.e.,

$$dP_s = \kappa\big(\mu_s - \log(P_s) - \lambda\big)P_s ds + \sigma P_s dW_s^*, \tag{2.32}$$

where dW_t^* is the increment of the Brownian motion under the equivalent martingale measure. Solving (2.32) using Itō's Lemma, we find that P_s is log-normally distributed with

$$P_s = P_0 \exp\left[e^{-\kappa s} + \kappa \int_0^s e^{\kappa(t-s)}\mu_t^* dt + \sigma \int_0^s e^{\kappa(t-s)} dW_t^*\right],$$

where $\mu_t^* = \mu_t - \frac{\sigma^2}{2\kappa} - \lambda$. The arbitrage-free futures prices implied by the model are given by

$$\begin{aligned} \log(F_s) &= \mathbb{E}_0[P_s] = \mathbb{E}_0[\log(P_s)] + 2^{-1}\operatorname{Var}_0(\log(P_s)) \\ &= \log(P_0)e^{-\kappa s} + \kappa \int_0^s e^{\kappa(x-s)}\mu_s^* dx + \frac{\sigma^2}{4\kappa}\left(1 - e^{-2s\kappa}\right) \\ &= \log(P_0)e^{-\kappa s} + \frac{\kappa A}{4\pi^2 + \kappa^2}\left(-2\pi\left(\cos(2\pi s + \phi) - e^{-\kappa s}\cos(\phi)\right)\right. \end{aligned}$$

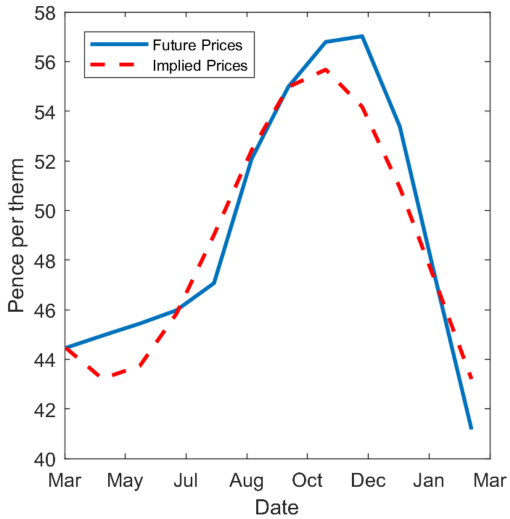

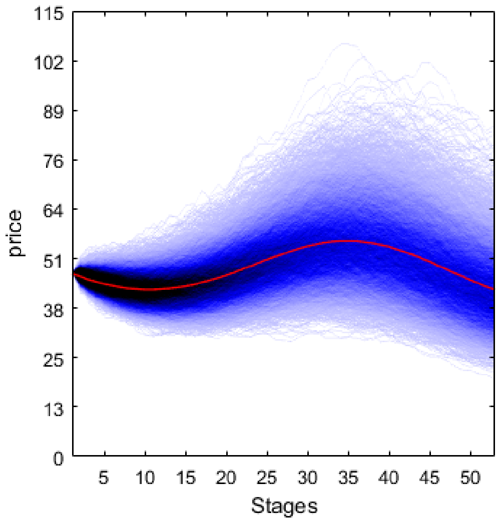

Figure 2.4: Fit of the price model (2.32) to the actual futures curve (left) and a plot of 10,000 paths from the calibrated process (right) in weekly resolution for one year with the red line representing the average price.

$$+\,\kappa\left(\sin(2\pi s+\phi)-e^{-\kappa s}\sin(\phi)\right)\right)-\left(\frac{\sigma^2}{\kappa}+\lambda\right)\left(1-e^{-\kappa s}\right)+\frac{\sigma^2}{4\kappa}\left(1-e^{-2\kappa s}\right), \tag{2.33}$$

where $\mathbb{E}_0$ and Var_0 are the expectation and the variance under the risk-neutral measure.

We calibrate the parameters of the model by fitting the implied futures prices (2.33) to an observed monthly future price curve for delivery at NBP in April 2018 to March 2019 from the 29$^{\text{th}}$ of March of 2018.[2] More specifically, we minimize the absolute difference between observed and implied prices with the function *multistart* from the MATLAB global optimization toolbox using the spot price on the 29$^{\text{th}}$ of March as P_0.

The calibration results in parameter values of

$$(\kappa, A, \sigma, \lambda, \phi) = (0.189, 4.993, 0.199, -3.407, 5.203).$$

The fit of the model to the observed futures curve seems satisfactory for our purposes

[2]Data source:

and is depicted in Figure 2.4 on the left and a plot of 10,000 simulations from the fitted process can be found in the right panel of the same figure.

For our computations, we discretize (2.32) to weekly time steps and assume a planning horizon from April 2018 to March 2019 (52 weeks) with the deterministic state being the 29.03.2018. The resulting lattice is built using $100,000$ simulations from the process and 100 nodes per non-terminal stage. This results in a rather large stochastic optimization problem with 53 stages. The ADDP algorithm was terminated after 100 iterations for all the calculations, which yields an average gap of below 1% between the ADDP upper bound and sampled value of the policy (lower bound) based on $10,000$ samples.

As a first exercise, we consider the derivative of the optimal value with respect to the amplitude A of the seasonal variation, i.e., we calculate the sensitivity of the value of the gas storage with respect to the magnitude of the summer/winter spread in gas prices. To this end, we calculate the derivative of the Lagrangian of the problem with respect to A arriving at

$$V'(A) = \mathbb{E}\left[\sum_{t=1}^{T}(x_t^{s*} - x_t^{b*})\frac{\partial}{\partial A}P_t\right] \approx K^{-1}\sum_{k=1}^{K}\sum_{t=1}^{T}(x_{kt}^{s*} - x_{kt}^{b*})\frac{\partial}{\partial A}P_t = \hat{V}'_K(A),$$

with x_{kt}^{b*} and x_{kt}^{s*} the optimal buy and sell decisions in scenario k and

$$\begin{aligned}\frac{\partial}{\partial A}P_t &= \kappa P_t \int_0^t e^{\kappa(s-t)}\sin(2\pi s + \phi)ds \\ &= \frac{\kappa P_t}{4\pi^2 + \kappa^2}\left(-2\pi\left(\cos(2\pi t + \phi) - e^{-\kappa t}\cos(\phi)\right) + \kappa\left(\sin(2\pi t + \phi) - e^{-\kappa t}\sin(\phi)\right)\right).\end{aligned}$$

Clearly, all the conditions for Theorem 2.4 are fulfilled in this example due to the smoothness of the involved functions and the distributional assumptions in (2.32).

In Figure 2.5, we compare $\hat{V}'_K(A)$ using K =10,000 samples with the estimates $\hat{V}'_{DQ}(A)$. Since A is the amplitude of the time-dependent mean reversion level μ_t, it influences the value of the storage in two ways: Firstly, an increase in A makes statistical arbitrage based on price spreads between summer and winter more prof-

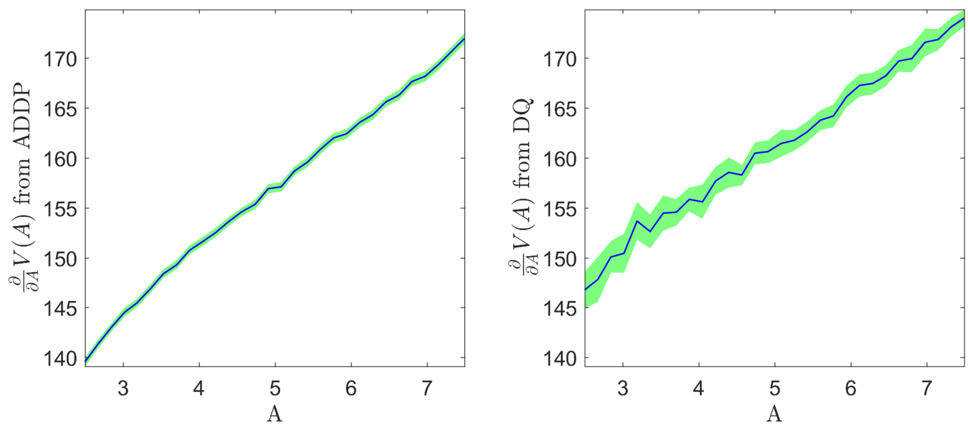

Figure 2.5: Comparing $\frac{\partial V}{\partial A}$ calculated via sampling (left) and via the difference quotient method (right).

itable. Secondly, an increase in A increases the potential for arbitrage trades: for a low A, only a few weeks can be traded profitably, while the efficiency losses cannot be recovered for the rest of the weeks. As A increases, this effect diminishes and the optimal policy trades more. Both effects in combination result in the positive and increasing derivative that can be observed in Figure 2.5.

Although, unlike in the newsvendor example, we do not have an analytic solution as reference, it seems obvious that $\hat{V}'_K(A)$ captures the true shape of $V'(A)$ better than $\hat{V}'_{DQ}(A)$ as the latter violates monotonicity and shows significant local variability, which is unlikely to be a feature of the actual derivative. This is also reflected in the wider confidence bounds for the estimate based on the difference quotient, which are again based on 30 independent evaluations using common random numbers as in the newsvendor example.

Furthermore, we note that $\hat{V}'_{DQ}(A)$ is consistently upwards biased. While the bias could be corrected with a smaller ϵ (see discussion above), this would lead to an even higher variance of the estimate. This shows that there is no choice of ϵ that leads to a result which is comparable in quality to $\hat{V}'_K(A)$ both in terms of bias and variance.

Finally, we calculate the derivative with respect to the flow efficiency factor η.

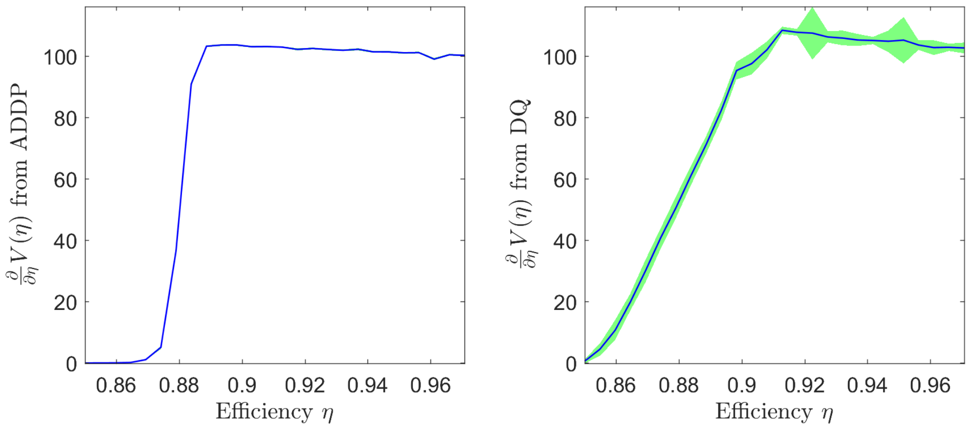

Figure 2.6: Comparing $\frac{\partial V}{\partial \eta}$ calculated via sampling (left) and via the difference quotient method (right).

Taking the derivative of the Lagrangian of the problem with respect to η yields

$$V'(\eta) = \mathbb{E}\left[\sum_{t=1}^{T} \lambda_t^*(x_t^{b*} + \frac{1}{\eta^2}x_t^{s*})\right] \approx K^{-1}\sum_{k=1}^{K}\sum_{t=1}^{T}\lambda_{kt}^*\left(x_{kt}^{b*} + \frac{1}{\eta^2}x_{kt}^{s*}\right) = \hat{V}_K'(\eta),$$

where λ_t^* are the optimal dual solutions assigned to the storage balance equations. Note that the above derivative is the derivative with respect to a parameter in the matrix, i.e., the left-hand side of the constraints. As above, we use K =10,000 to estimate $\hat{V}_K'(\eta)$ and 30 independent evaluations using common random numbers to compute the confidence bands for $\hat{V}_{DQ}'(\eta)$.

The results of the analysis are depicted in Figure 2.6. Up to an efficiency of around 87%, the losses from using the storage outweigh the potential gains that can be made by arbitraging on time spreads in the prices. Therefore, the optimal value as well as the derivative of the optimal value are 0. As η increases beyond this threshold, the storage starts expanding its operation to capitalize on more and more time spreads which get increasingly profitable. This explains the steep rise in $V'(\eta)$ in the region from 0.87 to 0.9. Once all possible time spreads are *used* by the optimal policy, further increases in η yield a more or less constant increase in $V(\eta)$. The comparison of the two methods to compute derivatives points to the same conclusion as the previous result. $\hat{V}_{DQ}'(\eta)$ is biased and at the same time has a higher variance than $\hat{V}'_K(\eta)$, which leads to the conclusion that $\hat{V}'_K(\eta)$ is clearly superior.

2.6 Conclusion

In this paper, we propose a method to calculate derivatives of value functions of MSLPs at points where these functions are differentiable. We base our results on classical envelope theorems and use some novel findings on the smoothness of the optimal values of linear optimization problems to establish that the value functions in question are differentiable almost everywhere. We use the latter property of differentiability almost everywhere to establish that the derivatives computed with our method are valid with probability one. To apply these results to SDDP-type decomposition methods for Markovian problems, we propose a discretization method based on scenario lattices. We show that, under certain regularity conditions, the sensitivities computed from discrete approximations of the problems converge to the true sensitivities of the original problem as the approximation of the continuous randomness gets finer.

The simple two-stage newsvendor example in Section 2.5 demonstrates that our method clearly outperforms a naïve computation of the derivative via the difference quotient in terms of average errors, the variance of the estimates, and the required computation time. These findings extend to the more elaborate gas storage optimization example in Section 2.5.2, which shows that the method works well on large problems with many stages.

Another interesting take-away of our numerical analysis is that the estimates improve when we use the real stochastic process instead of the scenario lattice to generate samples in Algorithm 1. In particular, we observe that the former estimates lead to significantly smaller errors on average and to more reliable confidence intervals.

This paper opens some avenues for further research. In particular, it would be interesting to use the results from Section 2.2 to study second derivatives as well as mixed derivatives of MSLPs based on existing second-order envelope theorems. This could prove useful in the context of financial risk management, which is often based on *Greeks* that represent second derivatives of option values. A prominent example is *gamma*, which is the second derivative of the value of a contingent claim with respect to the value of the underlying.

Furthermore, it would be interesting to apply the methods developed in the paper to practical problems. In particular, pricing and hedging problems in finance, where there is no analytical solution for the value of certain contracts, might be a rewarding topic for further study.

Chapter 3

Economies of Scope for Electricity Storage and Variable Renewables

written in collaboration with Prof. Dr. David Wozabal[1]

In this paper, we investigate whether and under which conditions jointly owning a variable renewable source of electricity (VRES) and an electricity storage generates economies of scope in competitive electricity markets. Using a simple stochastic optimization model that assumes frictionless markets, we analytically show that no economic benefit arises from combining the two assets. This finding is in contradiction to large parts of the literature, which claim that it is in the economic interest of owners of VRES to additionally own electricity storage. We also identify circumstances where our argument does not hold and the combination of storage and VRES could theoretically make economic sense on the level of individual agents. In the last part of the paper, we demonstrate in a numerical case study of the German market that even in cases where our theoretical results do not hold, tying together storage and VRES may produce suboptimal results.

[1]**Publication History**: Initially submitted to *IEEE Transactions on Power Systems* on 26.11.2019. Under review for a second round of revisions as of 12.05.2020.

3.1 Introduction

The increasing share of variable renewable sources of electricity (VRES) such as wind and solar power has brought about a considerable change to electricity systems in many countries. One challenge accompanying this change are unforeseen and costly imbalances created by VRES, which, in the worst case, could even threaten grid stability. Electricity storage offers the flexibility to mitigate these problems and stabilize electricity systems, potentially generating welfare gains for the respective societies (Kondoh et al. 2000).

Many authors take this logic one step further and postulate that owners of VRES can generate economies of scope from owning electricity storage, i.e., that the above argument extends from the system level to the portfolio level of individual agents (Zhao et al. 2015). One prevailing justification of this idea is that a storage would serve as a buffer, assisting the VRES to adjust its outputs in order to honor its commitments to the market thereby avoiding balancing costs (Garcia-Gonzalez et al. 2008). Another common argument for joint ownership is based on the fact that owners of variable production cannot adapt their production to take advantage of fluctuating market prices (Bathurst and Strbac 2003, Castronuovo and Peças Lopes 2004). A tethered storage could thus help to sell electricity when prices are high instead of at random times when the VRES produces.

Although the above arguments are widely accepted, little structured thought has been given to whether and under which circumstances joint planning of storage and VRES constitutes an optimal use of the storage asset. The aim of this paper is to fill this gap and provide the community with tools for a principled analysis as to when there are economics of scope for joint ownership and operation of storage and VRES.

We focus on competitive markets and price-taking agents that have no market power. Apart from this being by far the most common setting in the literature on economics of scope of energy storage and VRES, this choice is motivated by the fact that while all of renewable production combined frequently influences power prices, owners of single VRES plants are typically small players that are unlikely to be able to exert market power. Furthermore, we restrict our analysis to the profits of single players, i.e., we do not contribute to the literature on welfare implications of storage as for example in Schill and Kemfert (2011), Sioshansi (2014). However, we think that our results inform such a discussion.

In Section 3.2, we review the literature and categorize papers that claim that there are economic advantages of joint ownership and operation for single agents. In Section 3.3, we argue that in a broad range of circumstances there is no advantage in joint ownership. In particular, using a simple, yet quite general model, we show that if the storage unit has full access to all markets, markets are liquid, and the balancing market is accessible until imbalances of the VRES are known, there is provably no benefit from combining VRES with storage. In Section 3.4, we highlight that in certain circumstances such as market illiquidity and certain regulatory regimes restricting the market access of storage, the theoretical arguments in Section 3.3 cannot be applied and economies of scope are theoretically possible. Finally, in Section 3.5, we use a numerical example comparing the joint operation of a battery and a wind farm with separate operation of the two assets on the German spot and reserve markets to demonstrate that, even in settings where the absence of economies of scope cannot be formally proven, it is often suboptimal to tie the operation of a storage unit to a VRES. Section 3.6 concludes the paper and argues that the results of the paper provide a blueprint for a more careful consideration of potential economies of scope from combining energy storage with VRES.

3.2 Literature Survey

There is a large body of recent literature investigating the merits of combining VRES with electricity storage. The settings in these papers vary from optimally designing island systems without any grid connection (Ntomaris and Bakirtzis 2016, Papaefthymiou et al. 2010, Alimisis and Hatziargyriou 2013, Mason 2015) to determining the socially optimal amount of storage for a given economy (Dong et al. 2017, Pflaum et al. 2017, Ghofrani et al. 2013, Kocaman and Modi 2017), and also includes a large stream of literature that focuses on the gains that individual agents can make by the combination of these assets. In this section, we compile a reasonably large collection of papers that fall in the last category and demonstrate that the premise that jointly owning and operating VRES and storage can yield significant economies of scope for the owners of the assets is common and often goes unchallenged.

Since there are literally hundreds of papers on this topic, we do not claim our selection to be even close to complete. Instead, in order to demonstrate the pervasiveness of the idea of economies of scope, we select papers that give a representative sample of the different modeling choices, are published in leading journals, and have a good impact in terms of citations.

The selected papers can be found in Table 3.1. All listed papers claim economies of scope for the joint operation of storage and renewables which are either based on ad-hoc and often unrealistic assumptions or operate in settings where, according to the results in this paper, there should be no advantage of joint planning.

We have categorized these papers according to the type of optimization problem they address, what markets they consider, the risk preferences of the decision maker, and whether the authors explicitly compare the advantage of joint planning with separately planning storage and VRES.

More specifically, all systems consist of a VRES, in most cases a wind farm, and a storage capable of providing assistance to or store curtailed production from the VRES. Most models feature a day-ahead auction and a settlement phase where all imbalances resulting from erroneous forecasts of the production from VRES are cleared. The structure of this settlement phase varies from market to market. Additionally, some papers consider the bidding problem on multiple markets, notably on intraday or real-time markets, where each market participant can make adjustments

Table 3.1: Survey of papers.

ID	System	Method	Market	Risk	Comparison
Banshwar et al. (2019)	VPP (RES,PHS)	Det. NLP	DA,SR	neutral	
Crespo-Vazquez et al. (2018)	W,S	2s Stoch. convex	DA,ID,BM (RT)	neutral	WO / WA
Garcia-Gonzalez et al. (2008)	W,PHS	2s Stoch. MILP	DA	neutral	US / CS
Jiang et al. (2012)	W,PHS	2s Rob. MILP	DA	neutral	
Wang et al. (2013)	W,PHS,Th.	2s Stoch. MILP	DA,S	neutral	
Ding et al. (2015)	W,S	3s rol. Stoch. MIP	DA,ID	neutral	US / CS
Ding et al. (2016b)	W,S	2s Stoch. LP	DA,S	averse	US / CS
Ding et al. (2017)	W,S	2s Stoch. MILP	DA,S	neutral	
Sioshansi and Denholm (2010)	CSP,TES	Det. MILP	DA	neutral	WO / WA*
Khodayar and Shahidehpour (2013)	W,PHS	2s Stoch. MILP	DA,S	neutral	US / CS
Shu and Jirutitijaroen (2014)	W,CAES	MS,LP	DA (ERCOT)	neutral	WO / WA
Liu et al. (2015)	W,Hydro	Interval MILP	DA,ID	averse	US / CS
Khodayar et al. (2013)	W,PHS	Det. MIP	DA,S	neutral	WO / WA
Attarha et al. (2018)	W,CAES	Det. MILP, 2s Rob. MILP	DA (ERCOT)	neutral	US / CS
Varkani et al. (2011)	W,PHS	2s Stoch. MINLP	DA,S,BM (SR,RR)	neutral	US / CS
Malakar et al. (2014)	W,PHS	Stoch. MILP	DA (IND)	neutral	
Oskouei and Yazdankhah (2015)	W,PV,PHS	2s Stoch. LP	DA	neutral	
Bayón et al. (2016)	W,PHS	Other	DA	neutral	WO / WA
Al-Swaiti et al. (2017)	W,PHS,Th.	2s Stoch. MILP	DA,S,BM	averse	US / CS
Moradi et al. (2017)	W,UWCAES	Dynamic MINLP	DA,S	neutral	WO / US / CS
Mauch et al. (2012)	W,CAES	Stoch. NLP	DA	neutral	WO / WA
Jaramillo Duque et al. (2011)	W,PHS	Det. LP	DA,S	neutral	US / CS
Tan et al. (2014)	W,S,DR	2s Stoch. MINLP	Regulated DA,S	neutral	WO / WA
Parastegari et al. (2015)	W,PV,PHS,S	2s Stoch. MIP	DA,BM (SR,NSR)	neutral	US / CS
Ding et al. (2012)	W,PHS	Det. MIQCP, Stoch. Ch. constr.	Regulated DA	neutral	
Moghaddam et al. (2013)	W,CHS	2s Stoch. MILP	DA,S	averse	CS w/ IPW
Murage and Anderson (2014)	W,PHS	2s Stoch. LP	PPA (KEN)	neutral	WO / WA
Liu et al. (2015)	W,Hydro	2s Stoch. MINLP	DA,RT	averse	US / CS
Ghasemi et al. (2016)	W,EVs	2s Stoch. MINLP	DA	neutral	US / CS
Yıldıran and Kayahan (2018)	W,PHS	2s Stoch. MILP	DA,S	averse	
Su et al. (2019)	W,PHS	2s Stoch. MILP	Regulated DA	neutral	CS w/ IPW
Sun et al. (2019)	W,PV,PHS	Other	DA,ID,S	neutral	US / CS
Bathurst and Strbac (2003)	W,S	Det. MILP	DA,S	neutral	CS w/ IPW
Ghofrani et al. (2014)	W,EVs	2s Stoch. NLP	RT	neutral	
Thatte et al. (2013)	W,S	Rob. LP	DA,S	averse	Rob. / Stoch.
Márquez Angarita and Garcia Usaola (2007)	W,Hydro	2s Stoch. MIP	DA,mult. Auctions	neutral	US / CS
Castronuovo and Peças Lopes (2004)	W,PHS	Det. LP	Regulated Market	neutral	WO / WA

System W = wind power, S = storage, PHS = pumped-hydro storage, (UW)CAES = (underwater) compressed-air energy storage, Th. = thermal/conventional, VPP = virtual power plant, RES = renewable energy source, CSP = concentrated solar power, TES = thermal energy storage, PV = photovoltaic, DR = demand response, EV = electric vehicle, **Method** 2s = 2 stage, MS = multi-stage, LP = linear program, NLP = non-linear program, MIP = mixed-integer program, MILP = mixed-integer linear program, MINLP = mixed-integer non-linear program, MIQCP = mixed-integer quadratically-contrained program, Stoch. = stochastic, Rob. = robust, Det. = deterministic, Ch. constr. = chance constrained, **Market** DA = day-ahead market, S = Imbalance Settlement, ID = intraday market, (N)SR = (non-)spinning reserve market, RR = regulation reserve, BM = balancing market, RT = real-time market, BC = bilateral contract, IND = Indian market, KEN = Kenyan market, ERCOT = electric reliability council of Texas, **Risk Preference** neutral = risk-neutral (expected profits), averse = risk-averse (CVaR), **Comparison** WO = wind-only, WA = wind assisted, US = uncoordinated system, CS = coordinated system, w/ IPW = with increasing penalty weights, * = intermittent source alone v.s. joint system

to their day-ahead market bids close to delivery. Other papers include operation on the balancing markets.

As can be seen from column 4 of the table, most authors use two-stage stochastic (mixed-integer) linear programming approaches, which are very similar to the model we propose in Section 3.3. A minority of papers either use multistage stochastic optimization or deterministic models. Finally, some papers address uncertainty using robust and interval optimization.

Most papers employ risk-neutral strategies where the objective function is either the expected profit or the expected cost as indicated in column 5 of the table. However, a stream of literature explores risk-averse strategies using the Value-at-risk and the Conditional Value-at-risk (CVaR).

The last column of Table 3.1 highlights whether a paper performs an explicit comparison between jointly planning VRES and storage and planning separately or whether only the performance of the VRES operating alone is compared with the joint setup.

Finally, the issue of how imbalances are resolved plays a significant role. In regulated markets as well as in some deregulated markets any deviation in production is penalized with a constant fine (Garcia-Gonzalez et al. 2008, Ding et al. 2012, Moghaddam et al. 2013, Liu et al. 2015, Ghasemi et al. 2016). In these cases, it is common for the imbalance price to be set as a fixed surcharge of the electricity price.

In contrast, deregulated markets often set mechanisms where wind farm bidders are remunerated for their production surplus and charged for any deficit. In Mauch et al. (2012), Parastegari et al. (2015), Su et al. (2019), Ghofrani et al. (2014), Márquez Angarita and Garcia Usaola (2007), imbalance prices are set as functions of the electricity market price with overproduction remunerated at a lower price than what is charged for balancing shortages. Papers that deal with two-price imbalance settlement schemes such as the Spanish market consider these prices as another uncertain input (Sánchez de la Nieta et al. 2013, Al-Awami and El-Sharkawi 2011, Sánchez de la Nieta et al. 2015, 2016, Gomes et al. 2017). Another popular rule for imbalance prices is that both surplus and deficit imbalances are resolved at the same price. This is for example the case of the German spot market (Ding et al. 2016b, Khodayar and Shahidehpour 2013).

Generally speaking, papers that claim economies of scope for the combination of storage and VRES fall into at least one of the four following categories:

1. Papers that either only study the joint operation of storage and VRES or compare joint operation with owning only VRES, e.g., Crespo-Vazquez et al. (2018), Sioshansi and Denholm (2010), Khodayar et al. (2013) (see last column of Table 3.1). In these papers, no comparison of joint and separate planning is conducted and therefore the questions of whether the proposed use of the storage is optimal is not investigated at all.

2. Another very common feature is that storage is not allowed to participate on the market where imbalance costs arise, i.e., the storage is allowed to avoid imbalances for the VRES but not to participate in the balancing market where the flexibility of the storage could be sold to other market participants instead. This also includes the cases where authors introduce non-market based penalties for the deviation from the schedule of the VRES, e.g., Bayón et al. (2016), Tan et al. (2014), Castronuovo and Peças Lopes (2004).

3. In many papers where there is a explicit comparison between joint and separate planning the economic benefit of joint planning and therefore the economies of scope are found to be rather small and, in many cases, is most probably statistically insignificant, e.g., Khodayar and Shahidehpour (2013), Liu et al. (2015), Sánchez de la Nieta et al. (2015).

4. Finally, a rather small fraction of papers investigate situations where the arguments presented in the next section are not applicable and there might be economies of scope from combining electricity storage with VRES. Examples are papers that deal with a two-price imbalancing scheme as discussed above, e.g., Mauch et al. (2012), Parastegari et al. (2015), Su et al. (2019).

Note that all the papers listed in Table 3.1 fall in at least one of the categories 1 – 3.

3.3 A simple multi-stage model

Electricity storage generates economies of scope for the owners of VRES, if the additional profits they can generate from owning a storage exceed the profits that can be earned by an agent that only owns the storage. If storage would complement VRES in this way, then owners of VRES would be in an especially good position to build up storage capacity.

In this section, we investigate this claim and outline a simple stochastic optimization model that captures the essence of our argument and mirrors the most important aspects of a majority of the papers mentioned in the last section. Subsequently, we use the model to discuss the perceived advantages of joint operation of a storage unit and a VRES and show that, for many cases, these benefits do not exist.

3.3.1 Setting and Notation

Without loss of generality, we restrict planning to one day of operation and trading on electricity markets using a VRES and a storage. Trading takes place on a day-ahead market one day ahead of delivery and a real-time balancing market which trades up to (or close to) delivery. Correspondingly, we set up the planning problem in several stages. In the first stage, decisions on the bids on the day-ahead market are taken, while in the subsequent stages, real-time trading takes place.

Due to varying market designs all over the world, we do not specify the exact nature of the real-time balancing market. In our model, the crucial property of these markets is that they provide a price for electricity at a time when there is no (or very little) uncertainty about the production of the renewable asset. Candidates for such markets are real-time markets as implemented in many US systems (Ela et al. 2014) or intraday markets which trade close to delivery. Alternatively, one can think of the price on the real-time balancing market as the price of deviations from scheduled production such as the reBAP in Germany.

In the following, we denote by $t = 0, \ldots, T$ the number of decision stages in our model, where $t = 0$ represents the time of spot market bidding and the later time periods correspond to delivery periods for electricity on the considered day. Depending on the resolution of the markets, typically $T \in \{24, 48, 96\}$. To avoid conversion factors between stock and flow variables in our models, we assume an hourly resolution in this section, i.e., $T = 24$. Furthermore, we denote the day-ahead and real-time prices for period t by p_t^D and p_t^R, respectively. We assume that all markets are sufficiently liquid such that the owner of the plants is a price taker, there is no bid-ask spread, and the real-time price is symmetric, i.e., independent of whether electricity is bought or sold. In our models, we consider an idealized storage with perfect efficiency, no degradation, and no self-discharge. Relaxations of this assumption are discussed at the end of this section and in Section 3.4.

In order to model the information structure of the problem, we use a general probability space $(\Omega, \mathcal{F})$ and a filtration $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots \subseteq \mathcal{F}_T \subseteq \mathcal{F}$ with $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and assume the day-ahead prices p_t^D to be known to the decision maker at the start of planning, i.e., $p_t^D \triangleleft \mathcal{F}_0$, where $\triangleleft$ denotes measurability of a variable with respect to a sigma algebra. We furthermore assume that $p_t^R \triangleleft \mathcal{F}_t$.

The storage level in period t is modeled by l_t and is bounded above by the storage capacity $\bar{l} > 0$. Changes in l_t at time t are due to injections $i_t \geq 0$ or withdrawals $w_t \geq 0$. The random production g_t of the intermittent asset in period t does not incur any cost and cannot be controlled by the owner. We assume that g_t becomes known at the same time the real-time market for delivery in t trades, i.e., $g_t \triangleleft \mathcal{F}_t$. We denote by s_t^D and s_t^R the day-ahead and real-time trades of the storage unit and by r_t^D and r_t^R the corresponding trades of the VRES. Of course, if both assets are owned and operated by the same party, the bids can be added up for all practical purposes. However, to conceptualize the notion of separate and joint planning, we treat the trading decisions separately.

In the following, we describe the nature of the benefits that are often ascribed to the joint planning of storage and VRES in our setting. In the next section, we introduce a model that combines both of the perceived advantages of joint planning in one model.

Many authors stress that storage can be used to store production of VRES in order to sell it when prices are highest (Bathurst and Strbac 2003, Castronuovo and Peças Lopes 2004). In order to conceptualize this notion, it is easiest to ignore uncertainty and only consider the day-ahead market. The general thrust of the argument is not impacted by this assumption.

If the intermittent source of electricity trades alone, then its owner earns

$$V_1 = \sum_{t=1}^{T} g_t p_t^D.$$

Note that in the above it is implicitly assumed that $r_t^D = g_t$, which is possible, since g_t is deterministic.

If there is a storage unit, then the owner of the assets can effectively sell the production of the intermittent asset in the most profitable hours. This can be accomplished by solving the following deterministic optimization problem

$$V_2 = \begin{cases} \max\limits_{s^D, r^D, l, w, i} & \sum_{t=1}^{T} (s_t^D + r_t^D) p_t^D \\ \text{s.t} & l_t = l_{t-1} + i_t - w_t, \ \forall t \in [T] \\ & w_t - i_t = s_t^D + r_t^D - g_t, \ \forall t \in [T] \\ & 0 \leq l_t \leq \bar{l}, \ r_t^D, s_t^D, i_t, w_t \geq 0, \ \forall t \in [T] \end{cases}$$

for a fixed initial storage level $l_0 \geq 0$ and $[T] = \{1, \ldots, T\}$. Clearly, $V_1 \leq V_2$ and the difference between the two values is due to the shift in the sales of electricity to times when prices are high.

The second perceived advantage of joint planning concerns balancing of unforeseen surpluses or shortages generated by intermittent production. More specifically, it is usually assumed that $p_t^R \gg p_t^D$ in case of shortage and $p_t^R \ll p_t^D$ in case of overproduction. This is the case if there are separate real-time prices for buying and selling and can be argued even for a symmetric price if the surplus/shortage of the producer is highly correlated with the surplus/shortage of the overall system. This is a realistic assumption in systems with large capacities of VRES of the same kind.

To demonstrate the advantage of storage in this setting, we consider random real-time prices and a fixed day-ahead schedule r_t^D, which was already decided on. The expected revenue of a producer that has to balance her day-ahead market bids on the real-time balancing market can be written as

$$V_3 = \mathbb{E}\left[\sum_{t=1}^{T} r_t^D p_t^D + (g_t - r_t^D) p_t^R\right].$$

Note that the VRES has to clear the imbalance $(g_t - r_t^D)$ on the real-time market for the real time price p_t^R. For example, if $g_t < r_t^D$, then the VRES has to procure electricity to cover the shortage.

If the producer owns a storage, she can avoid potentially costly balancing on the real-time market by using the storage and solve the following problem

$$V_4 = \begin{cases} \max\limits_{l,w,i} & \sum_{t=1}^{T} \mathbb{E}\left[r_t^D p_t^D + (g_t - r_t^D - i_t + w_t) p_t^R\right] \\ \text{s.t} & l_t = l_{t-1} + i_t - w_t, \quad \forall t \in [T] \\ & 0 \leq l_t \leq \bar{l},\, r_t^D, s_t^D, i_t, w_t \geq 0, \quad \forall t \in [T] \\ & i_t,\, w_t,\, l_t \triangleleft \mathcal{F}_t, \quad \forall t \in [T] \\ & w_t \leq g_t - r_t,\; i_t \leq r_t^D - g_T, \quad \forall t \in [T], \end{cases}$$

where the last two sets of constraints ensure that the storage operation is tied to the surplus/deficit production of the renewable asset. Clearly, $V_4 \geq V_3$, i.e., the storage generates additional revenue for the owner of the VRES.

Note that in both examples above the storage plant does not have full market access and only trades with the energy that it gets via internal transfer from the VRES. In the next section we will relax this assumption.

3.3.2 Joint versus Separate Planning

In this section, we introduce two notions of planning: joint planning of the two assets, which allows for the internal transfer of energy between the VRES and the storage, and separate planning, which determines how much the two assets can earn when planned separately.

The joint planning model captures both the advantages of combining a storage with VRES that were discussed in the last section and additionally gives the storage full market access, i.e., the ability to not only store energy produced from VRES but also directly trade energy on the market. Consequently, the storage generates economies of scope for owners of VRES if and only if the profit from joint planning exceeds that from planning the two assets separately.

Defining the value of stored energy at the end of the day as $V(l_T)$, we can write the joint bidding problem as

$$
\begin{aligned}
\max_{s^D, s^R, r^D, r^R, l, w, i} \quad & \textstyle\sum_{t=1}^{T} (s_t^D + r_t^D) p_t^D \\
& + \mathbb{E}\left[(s_t^R + r_t^R) p_t^R\right] + \mathbb{E}[V(l_T)] \\
\text{s.t} \quad & l_t = l_{t-1} + i_t - w_t \\
& w_t - i_t = s_t^D + s_t^R + (r_t^D + r_t^R - g_t) \\
& 0 \le l_t \le \bar{l},\ i_t, w_t \ge 0 \\
& i_t,\ w_t,\ r_t^R,\ s_t^R,\ l_t \lhd \mathcal{F}_t,
\end{aligned}
\tag{V^J}
$$

where all the constraints hold for all $t \in [T]$ and with probability 1. Note that in the above problem trading decisions are not sign restricted, i.e., the decision maker is allowed to engage in speculative trading. Furthermore, the setup makes it possible to transfer energy directly from intermittent production to the storage unit. In particular, if $(r_t^D + r_t^R - g_t) \neq 0$, i.e., the intermittent plant generates an imbalance, then the storage unit has to deal with it, either by adapting its trades or by injections/withdrawals.

Next we consider the problem of separately planning the two assets. We can enforce separate planning, by disallowing internal transfers between intermittent production and the storage unit, i.e., impose the constraint $(r_t^D + r_t^R - g_t) = 0$. Clearly, the constraint ensures that the two problems decouple, since the objective function is separable in the decision variables for two assets and there are no constraints that involve variables for both assets. More specifically, we have

$$\left\{\begin{array}{ll} \max\limits_{s^D, s^R, l, w, i} & \sum_{t=1}^T s_t^D p_t^D + \mathbb{E}\left[s_t^R p_t^R\right] + \mathbb{E}[V(l_T)] \\ \text{s.t.} & l_t = l_{t-1} + i_t - w_t \\ & w_t - i_t = s_t^D + s_t^R \\ & 0 \leq l_t \leq \bar{l},\ i_t, w_t \geq 0 \\ & i_t,\ w_t,\ s_t^R,\ l_t \lhd \mathcal{F}_t \end{array}\right\}$$
$$+\left\{\begin{array}{ll} \max\limits_{r^D, r^R} & \sum_{t=1}^T r_t^D p_t^D + \mathbb{E}\left[r_t^R p_t^R\right] \\ \text{s.t} & r_t^D + r_t^R - g_t = 0 \\ & r_t^R \lhd \mathcal{F}_t \end{array}\right\}, \tag{V^S}$$

where the first problem is the problem of optimizing the storage alone while the second problem is the problem for the VRES. Again, all constraints in both problems hold for all $t \in [T]$ and with probability 1.

We are now in a position to prove the following proposition.

Proposition 3.1. *There are optimal solutions for V^J and V^S that yield identical profits with probability* 1.

Proof. Let $x^J = (s^D, r^D, s^R, r^R, l, i, w)$ be an optimal decision for (V^J). Clearly, the decision

$$x^S = (s^D, r^D, s^R + (r_t^D + r_t^R - g_t), r^R - (r_t^D + r_t^R - g_t), l, i, w)$$

fulfills the constraints of separate planning and is therefore feasible for (V^S). Since the term $(r_t^D + r_t^R - g_t)$ is added and subtracted in the objective function, the objective value does not change. This establishes that (V^S) yields profits at least as high as (V^J). The other direction is trivial, since (V^S) has one more constraint and therefore a smaller feasible set than (V^J) and the same objective function. □

The following corollary is an immediate consequence of Proposition 3.1.

Corollary 3.1. $V^S = V^J$, *i.e., the joint and separate planning yield the same profits for the planner.*

Note that the above result holds even if some of the simplifying assumptions in the models above are dropped. In particular, the following modifications do not change the results:

1. By Proposition 3.1, the profits of (V^J) and (V^S) the same with probability 1. Therefore, replacing the risk-neutral optimization by a risk-averse decision maker will not change the result in Corollary 3.1.

2. In most day-ahead markets, bidders can submit complex price-dependent bidding functions and bidders do not know the clearing price at the time of bidding. However, a generalization of this kind would not change the results.

3. Note that similar to the last point, it is not required to know the prices on the real-time market at the time a decision on the quantities to be traded is taken.

4. Clearly, since storage operation is the same for x^S and x^J in the proof of Proposition 3.1, a more detailed modeling of the technical characteristics of the storage such as storage losses, content-dependent efficiency, or random inflows, will not make a difference.

5. The nature or degree of uncertainty about g_t and p_t^R has no effect on the validity of Corollary 3.1.

3.4 Asymmetric Prices, Market Power, Information, and Regulatory Barriers

In this section, we present cases where the equality between separate and joint planning, established in Proposition 3.1, does not hold and joint planning may be more profitable. In particular, we discuss how asymmetries in prices in the real-time balancing market, caused by market design, market power, illiquidity, or grid fees may lead to a violation of Proposition 3.1. Furthermore, we will describe cases in which a storage unit cannot participate in markets where the VRES trades its energy or settles its imbalances.

3.4.1 Asymmetric Prices

Proposition 3.1 assumes that the prices p_t^D and p_t^R at which energy is traded are the same regardless of whether energy is sold or bought. In the case where prices for buying and selling are different, the proof for revenue equivalence between separate and joint planning fails. In this case, the imbalance generated by the intermittent source is settled at a different price than the imbalance that the storage creates. Hence, the internal transfer of energy is no longer revenue-equivalent to the transfer via the market.

This is, in particular, the case if the settlement of imbalances follows a two-price system where positive and negative balancing energy are differently priced (Ding et al. 2016a, Díaz et al. 2019, Akbari et al. 2019, Heredia et al. 2018, Sánchez de la Nieta et al. 2013, Al-Awami and El-Sharkawi 2011, Sánchez de la Nieta et al. 2015, 2016, Gomes et al. 2017).

Furthermore, asymmetric prices might also arise from bid-ask spreads in continuous trading on illiquid intraday markets. Volumes traded in European intraday markets are a fraction of the volumes traded on the day-ahead markets, which may lead to significant bid-ask spreads (Balardy 2018, European Energy Exchange 2019).

Lastly, grid fees can be a source of asymmetric prices as well. In some markets, operating and maintenance costs of the electricity grid are charged to electricity consumers but not to producers. If storage is treated as an electricity consumer while charging, it follows that it has to pay grid fees in addition to the market price when buying electricity. The ability to transact internally, instead of via the market, would thus benefit joint systems, as no grid fees would have to be paid.

In summary, conditions that would lead to an asymmetry between buy and sell prices may make joint planning more profitable than separate planning.

3.4.2 Market Power and Strategic Bidding

Until now we have assumed that firms are price takers, i.e., that p_t^R and p_t^D are independent of bidding decisions. If this assumption is violated the firms have market power and can strategically influence prices in their favor. The ability to profit from gaming the market in this way increases with the capacities of the respective firms. Hence, if players have market power, controlling more assets and therefore more capacities can be an advantage that potentially makes joint planning more profitable than separate planning (Barbry et al. 2019).

3.4.3 Balancing Markets

There are market designs for electricity markets which do not contain a real-time balancing market that clears continuously, making it impossible for the storage to trade at the same time at which the imbalances of the VRES realize. A prominent example for such settings are balancing markets where the right to participate in the market is auctioned significantly before imbalances realize. These market designs are common in Europe (ENTSO-E 2018).

In such settings, it is not possible to construct an equivalent separate planning solution as is done in the proof of Proposition 3.1, since the storage unit cannot trade on the real-time balancing market at the time the imbalances become known, and therefore cannot exactly match the imbalance profile of the VRES. Hence, a strict proof of the equivalence between the two modes of planning is not possible.

3.4.4 Regulatory & Technical Issues

In Proposition 3.1, it is assumed that storage units can participate directly in all markets. If, however, the storage is prevented from participating in the market, either on the basis of legal or technical reasons, then the role of the storage in joint systems is limited to merely assisting the intermittent source in its operations. The concept of separate planning would lose its meaning, as no independent bidding of the storage would be allowed. Examples of entry barriers are restrictions based on minimum technological requirements such as capacity or sustained power input/output over a specified period of time as is the case in some balancing markets (Netzregelverbund 2019b).

3.5 A numerical example

In the last section, we argued that there are realistic circumstances where the reasoning presented in Section 3.3 does not apply and joint planning might potentially generate economies of scope. In this section, we present a numerical example to substantiate that, even if the assumptions of our model are violated, it is plausibly better to sell the flexibility of a storage device on the market, instead of tying it to the balancing of one particular asset.

To make this point, we compare the revenues generated by the joint planning of a wind farm and a storage unit when trading on the day-ahead (DA) and intraday (ID) market in Germany with the revenues generated by separately planning the assets, whereby the storage participates in the market for secondary control reserve (SRL).

Note that the two models in this section are *not* related in the way joint and separate planning are related in Section 3.3. In particular, the separate planning model is not a restriction of the joint planning model. Instead, the joint planning model can be seen as the optimization *after* the owner of the assets decided on balancing the VRES with her storage. Bidding on the market for secondary control reserve implies that capacities can no longer be used for balancing the production of the wind park as the capacities always have to be available for the transmission system operator (TSO). This effectively prohibits the joint planner from participating in the market for secondary control reserve and forces her to the spot market.

We measure the economic performance of both systems on the 365 days from 01.01.2019 to 31.12.2019 and assume that the storage has 1 MWh of storage capacity and a maximum power input/output of 2 MW, while the wind farm has a capacity of 1 MW.

Since optimal bidding for both cases involves the solution of complicated stochastic optimization problems, we compute a lower bound on the revenues from separate planning and an upper bound for the revenues from joint bidding. Clearly, if the lower bound for separate planning exceeds the upper bounds for joint planning we can conclude that the former is more profitable. The code used to produce the numerical results in this section is available in a public *GitHub* repository.[2]

3.5.1 Strategies

We define the upper bound on the profit of joint planning by assuming that the wind park bids its generation forecasts on the day-ahead market. The storage unit trades on the day-ahead and intraday markets using perfect foresight of prices on

[2]See https://github.com/davidwozabal/TPWRS-00218-2020

both markets generating profits by optimizing over l, w, i, s^D, s^R

$$
\begin{aligned}
\max \quad & \tfrac{1}{4}\textstyle\sum_{t=1}^{T} s_t^D p_t^D + \tfrac{1}{4}\sum_{t=1}^{T} s_t^R p_t^R \\
\text{s.t.} \quad & l_t = l_{t-1} + \tfrac{1}{4}(i_t - w_t), \quad \forall t \in [T] \\
& w_t - i_t = s_t^D + s_t^R, \quad \forall t \in [T] \\
& s_t^D = s^D_{4\lfloor \frac{t-1}{4} \rfloor + 1}, \quad \forall t \in [T] \\
& \bar{i} \leq s_R,\ s_t^D \leq \bar{w}, \quad \forall t \in [T] \\
& 0 \leq l_{\alpha(d)} + \tfrac{1}{4}\textstyle\sum_{t=\alpha(d)}^{\alpha(d)+v} s_t^D \leq \bar{l},\ \forall d \in [D],\ \forall v \in [95] \\
& 0 \leq l_t \leq \bar{l}, \quad \forall t \in [T] \\
& 0 \leq i_t \leq \bar{i},\ 0 \leq w_t \leq \bar{w}, \quad \forall t \in [T],
\end{aligned}
$$

where $T = 4 \times 8760 = 35040$ is the number of 15-minute intervals in the planning period and all prices are known to the planner at the time of planning. Note that the third constraint enforces that day-ahead bids are the same for all quarter hours in one hour, since $4\lfloor \frac{t-1}{4} \rfloor + 1$ is the index of the first quarter hour of which t is a part of. The fourth and the fifth constraint enforce that the trading strategy is *asset backed*, i.e., that trades are not (entirely) motivated by arbitrage gains between the intraday and the day-ahead markets. Furthermore, we assume that the storage balances all imbalances of the wind farm thereby avoiding balancing payments without limiting the storage's power or energy capacity, i.e., restricting the optimization above. For this reason and since we use perfect foresight in the planning, the computed profits are upper bounds for the profits that can be obtained by the two assets on the spot market. Note that especially the assumption of perfect foresight generates revenues that far exceed what a storage owner can hope for under realistic circumstances when trading on the spot market.

Next, we discuss the lower bound on profits from bidding on the market for secondary control reserve. Secondary control reserve in Germany is procured in daily auctions. A bid for a specific tender consists of a capacity, a direction (up or down regulation), and two prices: a capacity price and an energy price. The TSO auctions a fixed amount of capacity and accepts bids in the order of their capacity price until the desired capacity is covered. If there is a requirement for up or down regulation, the TSO asks those accepted bidders who bid the best energy prices to produce/consume electricity. The remuneration of capacity payments as well as energy payments is conducted in a pay-as-bid fashion (Consentec 2014).

To compute the lower bound on separate operation, we assume that the wind farm bids its forecasts, and that deviations from the bids are accounted for using reBAP prices. reBAP prices are the settlement prices that have to be paid in Germany for deviations from a pre-registered schedule of production or consumption of electricity.

The storage unit uses a suboptimal non-anticipative policy to bid on all available secondary markets for control reserve for up and down regulation. It offers 1 MW capacity on the market for positive as well as negative reserve for all auctions in the planning period. The capacity and energy bids are chosen as quantiles of the distribution of bids from the power and energy prices of the respective auction on the previous day. We use the median prices for the capacity bids and the 1% and 99% quantiles for the positive and negative power prices, respectively.

We reserve the remaining 1 MW power capacity to balance the storage on the intraday market. In particular, if required, 5 minutes before delivery of every 15-minute product on the intraday market, we place a bid that guarantees that the storage neither runs empty nor overflows in the following 20 minutes, even if the TSO calls off energy at the maximum rate. In this way, we guarantee that the storage can always fulfill the commitments to the TSO resulting from its bids on the balancing market.

3.5.2 Setup & Results

For the calculations, we use day-ahead prices and quarter-hourly intraday index prices published by the European Energy Exchange (EEX) (European Energy Exchange 2019). We use wind forecasts W_t^{exp} and actual production W_t^{obs} from TenneT, consisting of quarter-hourly aggregate forecast and production from all wind power plants in the TSO's control area. Quarter-hourly imbalance prices, *reBAP*, were obtained from the German market regulator website, regelleistung.net (Netzregelverbund 2019a).

Clearly, the forecast error for aggregate production is lower than for a single wind turbine. The mean average error (MAE) of the wind forecast as a percentage of the average capacity factor for the TenneT zone is 11.96%. In Holttinen et al. (2013), the wind forecast errors for individual turbines range from 52% to 56% of MAE as a percentage of produced power. Therefore, to simulate an upper bound for the imbalances of a single wind turbine, we scale up the deviations between forecasts and production in the TenneT area by a factor of 5 in our calculations. The imbalance costs generated by a wind farm with capacity of 1MW are thus estimated by

$$V_{imb} = 5 \times \left(\frac{1}{4} \sum_{t=1}^{T} \frac{W_t^{obs} - W_t^{exp}}{P_{cap}} p_t^R \right),$$

where $P_{cap} = 27422$ MW is the total installed wind power capacity in the TenneT control area during the period of study (WindGuard 2019) and p_t^R is the *reBAP*. Note that the denominator 4 converts between power measured in MW and energy measured in MWh for 15-minute periods and the pre-multiplication with 5 is due to the above mentioned scaling of the deviations on the level of the grid zone to the level of the individual plant. The resulting imbalance costs sum up to $V_{imb} =$ €9,685.

Table 3.2 displays the results of the experiment. We observe that the upper bound for joint planning is lower than the lower bound for separate planning, demonstrating that it is preferable to use the storage to bid on the secondary control reserve market rather than pair it with a VRES with the goal of reducing imbalance costs.

Turning our attention to the results from separate planning, we notice that the largest revenue comes from positive balancing energy deployment. This is in

Table 3.2: Profits for both strategies rounded to unit Euros

Spot Upper Bound	**DA + ID**	**Reserve Lower Bound**	**SRL + ID**
Day-ahead Trading	1,533	Positive Capacity	22,447
Intraday Trading	86,392	Negative Capacity	24,910
Avoided Balancing	9,685	Positive Power	145,668
		Negative Power	-47,591
		Intraday Trading	-42,226
Total Profit	**97,608**	**Total Profit**	**103,208**

large part due to the strategy of bidding the first percentile and ensuring that the storage is frequently called. In contrast, negative energy creates a significant cost, which is expected since we are bidding the 99% quantile of the previous day's energy bids. The cost of the intraday bidding strategy of maintaining the storage state of charge, €42,226, is rather high. However, in aggregate, the strategy still outperforms the upper bound based on joint planning by more than €5,000. Note that since especially the upper bound on the profits from spot trading is quite loose, the actual difference under realistic circumstances likely substantially exceeds this figure.

This example illustrates that, although there are settings where Proposition 3.1 may not hold, in many of these situations it is plausible that the storage owner is still better off not committing itself to balancing a specific VRES, but rather operating independently in the market, which is most profitable for the storage.

3.6 Conclusion

In a time where storage technologies are looked at as a means to regulate our energy needs and to meet our climate goals, it is our intention to provide a framework for a more careful and principled study of the economics of scope arising from the combination of an energy storage with VRES.

More specifically, we pose the question whether joint planning of a VRES and an electricity storage is always the best use of these assets.

We show how in a simple setting, where bidding is made first on a day-ahead market and then corrected on a real-time balancing market, there is no advantage in joint planning. The situation becomes less clear when there are price asymmetries, market power, entry barriers restricting access to the markets, or when the storage unit participates in an auction-type balancing market.

We make the case, by way of a numerical example, that there are good economic arguments against operationally tying a storage to VRES. Instead, it seems more profitable to let the storage participate independently in the most lucrative markets.

Our results thus show that policy makers cannot count on owners of VRES to build up significant storage capacities and thereby themselves provide the flexibility required to tackle imbalances created by their intermittent production.

Chapter 4

Optimal Battery Duration on the German Secondary Control Reserve

In this paper, we investigate the effect of the energy-to-power ratio on the profits of a storage unit operating optimally and exclusively on the German secondary control reserve market. With data ranging from August 2018 to July 2019, we solve a deterministic rolling-window mixed integer bilinear optimization problem, where optimal bidding decisions are taken for every 4-hour period auction in the year. We obtain a curve detailing the aggregate profits in the year as a function of battery duration (energy-to-power ratio). Profits converge to a saturation level with increasing duration. We calculate this saturation level, which we define as the theoretical maximum profit potential a storage unit operating at a given power level output can hope to achieve in this market in the time range of analysis. At a duration level of 4 hours, which is a customary energy-to-power ratio for battery storage power stations, an optimal rolling-window policy will secure 97.5% of the profits potential. 99% is achieved at a duration of 8 hours, and 99.8% at 12 hours. Duration levels above 13 hours yield no further increase in profits, independently from the power volume bid in the market.

4.1 Introduction

The share of energy storage in our power markets is expected to increase significantly in the coming years (Mackenzie 2019). In particular, energy storage systems with adequate time response seem well-suited to provide balancing power in control reserve markets, as they can both supply and consume power when needed (Kocer et al. 2019, Medina et al. 2014). The intended use of storage in wholesale power markets brings about a few challenges for the private storage unit owner. Questions such as whether a storage unit can be profitable and honor its commitments while trading exclusively on the control reserve markets, or what the right size, power output and time responsiveness of the storage unit should be in order to operate effectively in these markets are relevant to agents looking to invest in energy storage (Nasrolahpour et al. 2016).

In this paper, we are interested in the latter question, more specifically we shall focus on determining the optimal relationship between the maximum energy capacity of a storage unit and its maximum power output when operating on the German secondary control reserve. This metric is referred to in the literature by one of several terms, such as energy-to-power ratio (E2P) or duration (Denholm et al. 2019, Fuchs et al. 2012). We shall use these terms interchangeably throughout this paper.

In contrast to the spot market, where successful bidders are scheduled to provide or consume a predetermined volume of power during a particular time slot of the following day, operation on balancing markets requires the market participant to reserve capacity and provide power when needed, without knowing whether and how often they will be called (Consentec 2014). This poses a challenge to storage operations, because the storage unit must be large enough to accommodate the amount of power bid in the market. This limitation in feasible power bids can be mitigated to a certain extent, by allowing the storage unit to procure or dispatch energy externally, such as via the intraday spot market, in such a way so as to ensure the storage unit will always be able to honor its commitment to the balancing market, either by providing energy without fully depleting, or by storing energy without ever becoming full.

However, in this setting, storage unit capacity and power must be partitioned from the start (Netzregelverbund 2019b), meaning that the storage unit will have less capacity volume to bid on the control reserve. The storage unit cannot dynamically share its capacity and volume among the control reserve and other markets.

In this paper, we will be focusing on the saturation of profits that comes with an increasing energy-to-power ratio (E2P). We will not allow a partition of the storage capacities or power outputs across markets, but will consider the exclusive operation of a storage unit on the German secondary control reserve (SCR). The insights we gain from this analysis are unaffected by this restriction. Furthermore, we shall characterize this saturation curve of profits as a function of E2P and notice that it is independent of the storage power level.

In Section 4.2, we present a brief literature survey and expose the main contribution of this study. In Section 4.3, we describe a rolling-window deterministic model to bid optimally on the German secondary control reserve. In Section 4.4, we present a parametric sweep approach to solve the optimization problem. We show our results in Section 4.5 regarding duration, and finally summarize our conclusions.

4.2 Contribution

There have been many studies involving the use of storage intended for grid-based systems. These have ranged from literature surveys reviewing the state of the art of a collection of storage technologies and their best-use applications (see for example Aneke and Wang 2016, for an overview), to profitability studies considering the installation cost of storage of different power and capacity levels (US Department of Energy 2019). In Hesse et al. (2017), the authors conduct a review of lithium-ion stationary battery systems and their applications to the grid. Staffell and Rustomji (2016) investigate the sources of revenue available to energy storage in the energy markets and ancillary services, and outline the obstacles and issues faced in each.

Many studies considering the operation of storage on balancing markets or ancillary services have focused on the primary control reserve (PCR). These studies target the optimal dimensioning and sizing of a storage unit intended for the PCR, by way of finding the minimum capacity required to honor the technical requirements of the grid (Oudalov et al. 2007), or by factoring in installation costs and minimum degradation of the battery over its lifetime (Engels et al. 2019). In Oudalov et al. (2007), the authors cite for their setting on the PCR an optimal storage capacity to be 0.62 hours multiplied by the nominal power rating. In Engels et al. (2019), while looking at the German PCR, a 1.6 MW power capacity with a a 1.6 MWh storage capacity is found to offer the best net present value.

Fewer studies have considered the operation of storage on the secondary control reserve. Zeh et al. (2015) describe a storage operating on the German SCR as it was before the change in the delivery periods to 4-hour blocks, and allows the storage to assist a solar photovoltaic generation source and participate in other markets.

With the current study, we offer a description of the optimal bidding problem on the SCR given the market structure as of the 2018 change (Consentec 2020). We present a deterministic bilinear formulation of the storage bidding problem on the SCR. We also perform a sensitivity analysis of the optimal profit performance as a function of the energy-to-power ratio of the battery. The results on this study inform the literature on the adequate size and power of storage for the SCR.

4.3 A simple rolling-window model

In this section, we outline the underlying optimization problem of bidding energy prices with a storage unit on the German secondary control reserve (SCR).

4.3.1 The German Secondary Control Reserve

The SCR is structured as a pay-as-bid auction, where market participants bid to provide positive and negative balancing energy for a period of four hours. All auctions take place on the day before their respective period of delivery (Consentec 2014), and positive and negative balancing power are treated as distinct products, each having their own auction. In the SCR, market participants place bids for a volume of offered power (in MW), a price to provide power capacity (in currency per MW), and a price to actually dispatch balancing power (in currency per MWh) when called upon to do so by the grid operator. Throughout the paper, we shall refer to these as capacity volume, capacity price and dispatched energy price respectively. Additionally, we shall call the volume that is actually traded as dispatched volume.

4.3.2 Mixed-Integer Bilinear Problem Formulation

The inherent optimization problem a market participant entering the SCR faces can be modelled as a mixed-integer bilinear program. The market participant wishes to maximize his profits, which consist of revenue streams for provided capacity and dispatched power. Each of these revenue streams are modelled as bilinear product terms, where price and volume are both decision variables to optimize for. Furthermore, binary variables are required to model the effect of whether the bidder will be called upon to provide dispatched energy over the course of the four-hour period, as well as whether the capacity bid is accepted by the grid operator.

We will assume that the storage unit owner places bids to provide both positive and negative balancing energy, and does not participate in any other market nor has any other means of charging or discharging his storage unit other than via the SCR. For a single auction period, the deterministic optimization problem where time is discretized in finite time units of one second is described by:

$$
\begin{aligned}
\max \quad & P^+_{cap}V^+_{cap} + P^-_{cap}V^-_{cap} + P^+\sum_{t=1}^{T} V_t^+ + P^-\sum_{t=1}^{T} V_t^- \\
\text{s.t.} \quad & s_t = s_{t-1} + \alpha(V_t^- - V_t^+), \quad \forall t \in [T] \\
& 0 \leq s_t \leq S_{Max}, \quad \forall t \in [T] \\
& \begin{bmatrix} \underline{P}^j_{cap} \leq P^j_{cap} \leq P^{cut,j}_{cap} \\ 0 \leq V^j_{cap} \leq \bar{V}^j_{cap} \\ \begin{bmatrix} \underline{P}^j \leq P^j \leq P^{cut,j}_t \\ V^j_t = \min(V^j_{cap}, D^j_t) \end{bmatrix} \vee \begin{bmatrix} P^{cut,j}_t < P^j \leq \bar{P}^j \\ V^j_t = 0 \end{bmatrix} \end{bmatrix} \vee \begin{bmatrix} P^{cut,j}_{cap} < P^j_{cap} \leq \bar{P}^j_{cap} \\ V^j_{cap} = 0 \\ P^j = 0 \\ V^j_t = 0 \end{bmatrix}. \\
& \forall t \in [T], \quad j = \{+,-\}
\end{aligned}
\tag{4.1}
$$

The objective function consists of the four terms already mentioned, which encompass revenues from capacity bids and from dispatched energy for both positive and negative balancing energy, $j = \{+,-\}$ respectively.

The set of constraints include balance continuity constraints for the storage level s_t, which dictate that the storage level at any given time t must be equal to the storage level in the previous time period $(t-1)$ plus any amounts injected V_t^- to or withdrawn V_t^+ from the storage unit during the time period t. α is a constant conversion unit factor to equate the storage level (in units of MWh) and the transacted power (in units of MW). For a finite time resolution of one second, α is equal to 1 over 3600, as there are 3600 seconds per hour. Furthermore, the storage level is bounded at all times by zero when empty, and by its maximum capacity S_{Max}. Likewise, the capacity volume bid V^j_{cap} is bounded by an upper-bound, $\bar{V}^j_{cap}$, which can be thought of as the maximum input/output power of the storage unit. Without loss of generality, we assume the storage unit to have a 100% efficiency.

There are then two levels of integer decision-making to be taken for each product. Firstly, depending on the capacity price bid chosen, the storage unit may or may not be selected by the grid operator to provide capacity in this four-hour period.

If the capacity price bid is greater than the last price to be accepted, the storage unit is not accepted, and all volumes - capacity and dispatched volume alike - are set to zero, the storage unit does not participate and the profits are zero. If the capacity price bid is within the selected price range, then there comes a second layer of integer decision-making, which occurs for every time period t. Again, whether the storage unit is called upon to dispatch balancing energy depends on the dispatch energy price bid P^j and on the volume demand D_t^j at each second. For each second, we can calculate a dispatch cutoff price $P_t^{cut,j}$, which is the last price to be selected to dispatch balancing energy as a function of the auction supply curve and its intersection with the demand volume. If the energy bid is below the cutoff price for time t, the storage unit is called to dispatch energy. If not, the volume traded at time t is zero. When the storage is called, the volume traded is the minimum between the demand volume at time t , D_t^j, and the bidden capacity volume V_{cap}^j. The former is only ever selected if the aggregate volume is low, the price bid is the lowest in the auction order book and therefore the storage unit is the first or among the first units to be called upon to dispatch. Otherwise, the unit must always provide its bidden capacity volume. Lastly, $\underline{P}_{cap}^j, \bar{P}_{cap}^j, \underline{P}^j, \bar{P}^j$ are soft lower and upper bounds for capacity and energy price bids.

4.3.3 Problem Relaxation

As we are dealing with a deterministic problem, all the data concerning capacity bids is known, which in particular means that it is known what the highest price to be accepted is for each auction, so unless there is a reason to not take part in the auction, the optimal decision for capacity price will always be the highest price to be accepted. Unless there is a scenario where no price level that would lead to being accepted in the capacity auction guarantees a positive profit, it is always preferable to have a capacity bid accepted. Even if the energy price bid is too high for the storage unit to ever be called, the storage unit will still generate a revenue for providing capacity.

Therefore, we shall focus exclusively on the decision-making part of the problem having to do with the selection of the dispatched energy price and volumes. We shall assume that the storage unit owner's capacity bid is always accepted, and that the capacity volume bid is equal to the storage maximum power throughput. This helps us further simplify the problem and remove one layer of integer decision-making as now only two possible cases remain per time period t, which are the cases where the storage unit is called or is not called to dispatch balancing energy. The simplified optimization problem then becomes:

$$\begin{aligned}
\max \quad & \sum_{t=1}^{T} w_t^{+} + \sum_{t=1}^{T} w_t^{-} \\
\text{s.t.} \quad & s_t = s_{t-1} + \alpha(V_t^{-} - V_t^{+}), \quad \forall t \in [T] \\
& 0 \le s_t \le S_{Max}, \quad \forall t \in [T] \\
& \begin{bmatrix} \underline{P}^j \le P^j \le P_t^{cut,j} \\ V_t^j = \min(V_{cap}^j, D_t^j) \\ w_t^j = P^j \min(V_{cap}^j, D_t^j) \end{bmatrix} \vee \begin{bmatrix} P_t^{cut,j} < P^j \le \bar{P}^j \\ V_t^j = 0 \\ w_t^j = 0 \end{bmatrix} \\
& \forall t \in [T], \quad j = \{+,-\}
\end{aligned} \tag{4.2}$$

This will be the problem we shall focus on in the next sections.

4.4 Parametric Sweep

In the following, we shall describe the method to solve problem (4.2).

4.4.1 Traditional solvers approach

We can solve problem (4.2) through the use of a solver capable of solving mixed-integer linear programs. Commercially-available solvers such as Gurobi, CPLEX or FICO Xpress are perfectly suited for this. However, as we are solving a problem that for every auction considers 14400 time periods, which is the number of seconds in four hours, this considerably slows down the performance of the algorithm and increases the runtime of the program. As we would like to solve the optimization problem not just for one auction but for all four-hour periods in a year, which essentially involves solving the problem 2190 times in sequential order, another method must be found for this to be a time-feasible approach.

4.4.2 Parametric Sweep Method

Instead of using a solver, we solve the problem by performing a parametric sweep over the price range for positive and negative balancing energy. Having access to the auction supply curve and demand volume, we evaluate the combinations of feasible price levels for positive and negative balancing energy, constrained by storage level bounds, and select the combination leading to the highest profits. This method not only solves the problem much faster than solving with a solver, but also provides more information on the problem, most notably on the feasible set.

Exemplifying our approach, we select the range of bidden prices in the auctions for positive and negative balancing energy corresponding to the same four-hour period, and evaluate how the storage level of the battery would change over the course of the period as it is called to dispatch in result of the price level bids. After tracking the storage level for all price combination pairs of positive and negative balancing energy over time, we keep the combinations for which the storage boundary constraints are not violated. This constitutes the feasible set of prices. From among this set, we select the combination with the highest profits, which is our optimal solution.

We then solve the problem sequentially for each four-hour period in the year, in which the initial storage level of problem k is equal to the last storage level of problem $(k-1)$ to preserve continuity. This is a myopic strategy, as it computes the optimal policy for each four-hour period, with no knowledge of what the future auctions will be like.

4.4.3 Runtime comparison between methods

To illustrate the differences in runtime between solving the optimization problem with a solver and solving it with the parametric sweep method, consider the following. Using a Dell Latitude E5450 laptop running Windows 10 with a fifth generation Core i5 CPU and 8 GB of RAM, the optimization using the Gurobi solver took around 90 minutes to solve problem (4.2), which corresponds to a single auction period of four hours. By comparison, the same machine using the parametric sweep method took 50 minutes to compute the optimal policy for an entire year's worth of data, i.e. 2190 auction periods.

4.5 Optimal battery duration for the SCR

The E2P ratio represents the amount of time a battery may last while providing its maximum power at a constant rate from full to empty. As an example, an 8 MWh battery with peak input/output power of 2 MW has a duration of 4 hours.

The purpose of this study consists in determining the E2P value beyond which a battery operating exclusively on the SCR, optimizing its price bids on a 4-hour rolling window for one year in the terms defined in Section 4.3, would no longer generate more profits. This is relevant to determine what capacity and input/output power a storage station project destined to operate in the SCR should have. More precisely, knowing this limit value for E2P sets an upper bound on the maximum size a battery intended for trading on the SCR should have. Any battery capacity larger than this limit value, with respect to its power output, will not generate any more profits under this strategy. Furthermore, we will show that for usual volume bids in the SCR, ranging from 1 MW to 100 MW, the profits as a function of duration show the same saturation point when normalized for power.

4.5.1 Case Study

We shall consider auction and volume demand data from the German SCR market ranging from the 1st of August 2018 to the 31st of July 2019. The public auction and volume demand data were acquired from the German regulator website, regelleistung.net.

Using the parametric sweep method to solve the deterministic rolling-window optimization model, we compute the annual profits a storage unit would have had during this period, and repeat this exercise for increasing values of battery capacity for a given power bid. In particular, we shall be interested in the E2P ratio between capacity and power, and how annual profits vary with varying duration.

Figure 4.1 describes the annual profits as a function of duration, for a battery with maximum input/output power of 1 MW. This maximum is the same for both positive and negative balancing energy.

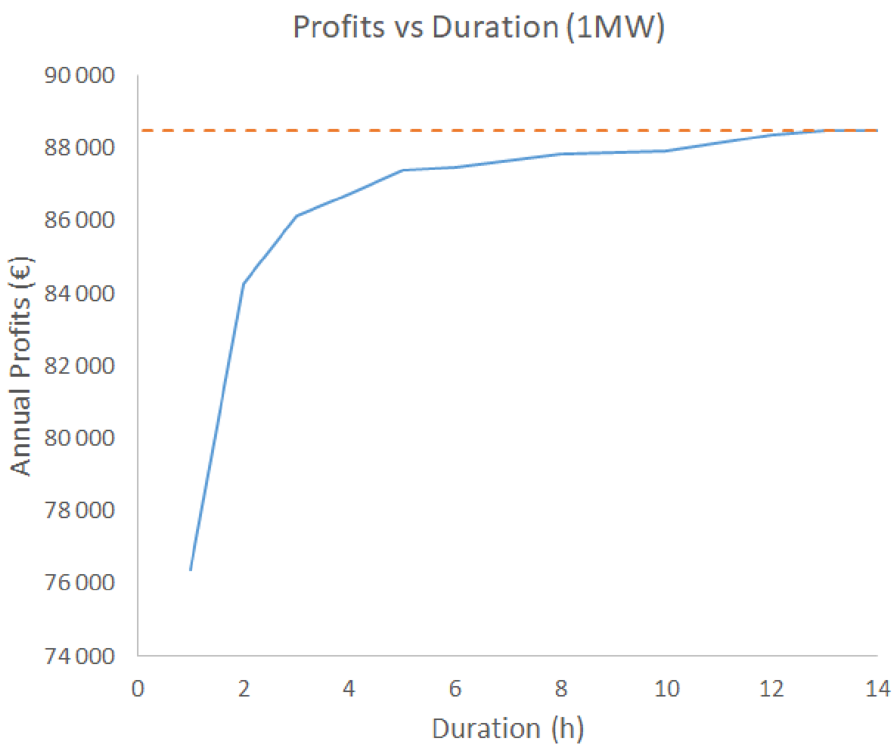

Figure 4.1: Annual Profits saturate with increasing Duration. The dotted orange line corresponds to the maximum profits potential, reached with a duration of 13 hours and beyond.

We see that for increasing battery size, the annual profits increase. This is intuitive as the feasible set of price levels to bid on increase, which means better solutions can be found. Additionally, we see that annual profits saturate as the size increases, reaching a limit profit value of 88467 EUR. No matter how much larger a battery is with respect to its power output, the profits to be gained are capped by this amount. This can be interpreted by the fact that from a certain point onwards, the maximum capacity constraint no longer is binding, and the issue that restricts the storage unit from making more profits is the maximum power output.

This result leads us to make the following proposition concerning battery operations on the SCR.

The profits a battery operating exclusively and optimally on the SCR at a given power volume converge to a limit cap with increasing duration.

Storage operation constraints limit the phase space of price bids available to the storage owner. In particular, a higher capacity translates into a relaxation of the boundary constraints restraining the battery storage level at any given time. This allows the storage owner to access lower price regions, thereby guaranteeing that the storage unit is called more often.

At some point, one may imagine an extreme case where the capacity of the storage unit is so large that its capacity size is no longer a factor in determining the price bids. To illustrate this, one could take the case where the storage owner bids the lowest price in the auction for positive balancing energy, and the market is continuously asking for positive energy during that 4-hour period. This would mean that the battery would be required to discharge continuously at bidden volume for four hours. The storage capacity boundary constraint would no longer be binding. In contrast, if the battery size is not large enough or its initial state is not high, then the lowest price level would not be available to the storage owner as it would violate storage boundary constraints. Instead, the power volume bid, which is kept constant by assumption, would become the limiting factor to attaining higher profits. Therefore, the optimal performance with perfect foresight of storage units with increasing duration sharing the same power bid increase monotonically and eventually converge.

4.5.2 Independence from power level bids

We shall see in the following that this saturation of profits is independent of the maximum power output selected in the customary range of power levels transacted on the SCR, i.e. which falls between power bids from as low as 1 MW to bids as high as 100 MW.

We repeated this exercise for increasing power levels. A summary can be found in Figure 4.2. Figure 4.2 describes annual profits as a function of duration, for different power levels ranging from 1 MW to 100 MW. The profits are normalized

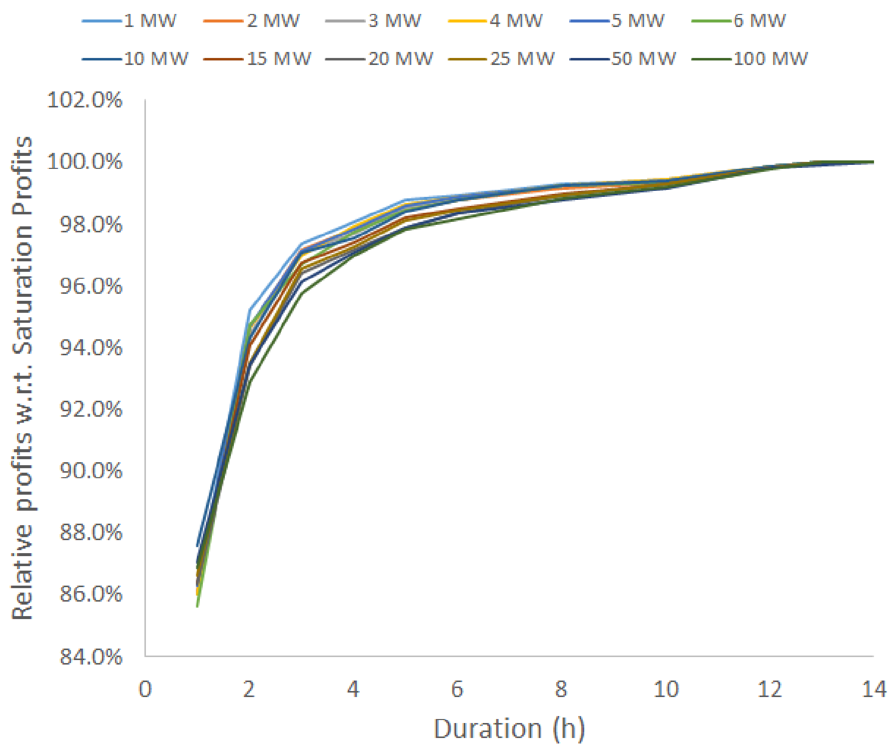

Figure 4.2: Increasing Power Levels in customary range do not alter the outcome.

with respect to their limit level. We see that for increasing power levels, the profits curve behaves the same way with respect to duration. This leads us to conclude that no matter the maximum power output of the battery, with an optimal trading policy, not only do profits saturate to a given level, but the ratio of profits to this limit cap is well defined. In this light, we observe that around 97.5% of the profits to be had on the SCR, under a deterministic rolling-window optimal policy over the course of one year, with a battery of a given maximum power output, can be achieved with a duration of 4 hours. Furthermore, the duration level beyond which profits saturate is reached at 13 hours for nearly all power volumes. A detailed breakdown of these results can be found in Table 4.1.

Table 4.1: Profits as a percentage of saturation profits per power volume

	1 MW	2 MW	3 MW	4 MW	5 MW	6 MW	10 MW	15 MW	20 MW	25 MW	50 MW	100 MW
1 h	86.3%	86.1%	86.0%	86.0%	86.3%	85.6%	87.6%	86.4%	86.3%	86.6%	87.0%	86.8%
2 h	95.2%	94.6%	94.4%	94.6%	94.7%	94.7%	94.3%	94.1%	93.4%	93.5%	93.5%	92.9%
3 h	97.3%	97.1%	97.1%	97.0%	97.1%	96.7%	97.0%	96.7%	96.4%	96.5%	96.1%	95.7%
4 h	98.0%	97.8%	97.7%	97.9%	97.8%	97.7%	97.5%	97.4%	97.1%	97.2%	97.0%	96.9%
5 h	98.7%	98.4%	98.5%	98.6%	98.6%	98.4%	98.4%	98.2%	97.9%	98.1%	97.9%	97.8%
6 h	98.9%	98.8%	98.8%	98.8%	98.8%	98.7%	98.8%	98.5%	98.4%	98.4%	98.3%	98.1%
8 h	99.3%	99.2%	99.2%	99.2%	99.2%	99.2%	99.2%	98.9%	98.9%	98.9%	98.8%	98.8%
10 h	99.4%	99.3%	99.4%	99.4%	99.4%	99.3%	99.4%	99.3%	99.3%	99.3%	99.1%	99.2%
12 h	99.9%	99.8%	99.8%	99.8%	99.8%	99.8%	99.8%	99.8%	99.8%	99.8%	99.8%	99.8%
13 h	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	99.9%	100.0%
14 h	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
15 h	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
20 h	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
30 h	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
50 h	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

4.6 Conclusion

Battery storage stations are dimensioned according to the purpose that they are intended for. E2P ratios in actual energy storage projects range from the fraction of the hour, e.g. 6 minutes, to several days, e.g. 120 hours (DOE 2020). Energy storage projects that have both the power output and size and satisfy the fast time-response required to participate in the SCR typically consist of lithium-ion batteries, and these projects have E2P ratios that do not usually go beyond four hours (International Electrotechnical Commission 2011).

In this paper, we set an upper-bound on the energy capacity a storage unit should have with respect to its maximum power throughput when operating on the German secondary control reserve. From this value onward, an increase in duration will not lead to higher profits. We have shown that this upper bound does not depend on the power volume bid, as doesn't the curve of relative profits with respect to this upper bound. These insights inform the market agent looking to invest in storage technologies with the purpose to trade on the SCR what dimensions a storage unit should have when it comes to energy capacity and to power, as well as what profit potential such a storage unit offers.

These results are obtained for a bidding optimization problem that is both deterministic and myopic. Nevertheless, the addition of uncertainty, or a more complete description of storage unit operations, such as the inclusion of storage efficiency, does not alter the insights described in this paper regarding the saturation of profits with increasing duration.

Chapter 5

Conclusion

Electricity markets across the world have changed considerably in the past two decades. The deregulation of markets in Europe was a first step towards a change in market design, which led to the opening of the markets to competition. This was further accentuated by the transition towards a greener generation mix. The development of renewable energy sources in the electricity grid, particularly the ever increasing shares of wind and solar energy, introduced variability in the electricity supply. The variability of supply brought challenges to the sector, that producers, consumers and grid system operators must contend with.

This doctoral thesis is a collection of works revolving around two answers to the challenges brought about by the variability of renewables sources in the grid.

The first is methodological in nature, and consists of the use of stochastic optimization applied to electricity markets. This framework, which deals with the task of taking best decisions under uncertainty, is adept at quantifying the variability of wind and solar production, and factoring them in to allow market participants and system operators to adopt policies aiming to maximize profits from bidding in the market, or minimize operating costs while meeting demand.

The second answer to address the variability of renewables is the development of storage technologies intended for the grid. Storage has a role to play in bringing stability to the grid. In being able to balance supply and demand, grid-based storage is expected to grow considerably in the next few years and become a larger presence in electricity markets around the world.

The research presented in this thesis concerns optimization methods on the one hand, and the valuation of storage on the other. The work was split among them in the following way: the first paper included in this thesis concerns itself with the former, interesting itself in the topic of parametric sensitivity analysis in multistage stochastic linear optimization problems. The second and third papers are in the camp of the latter, offering insights into the use of storage in electricity markets, from a market participant's perspective.

The first essay developed a theoretical foundation and a methodology with which to calculate derivatives of the value function of a multistage stochastic linear program with respect to inherent parameters of the problem.

One main contribution of this essay was to extend the literature on envelope theorems to a setting of multistage stochastic optimization. The essay lays out the conditions under which the value function is differentiable in the parameter of study, and in turn characterizes the set of parameter points where it is not. This set is shown to be finite and have Lebesgue measure 0, thereby allowing the statement that the value function of a multistage linear problem is differentiable almost everywhere.

The second main contribution of this essay was to develop a sampling method with which to calculate these derivatives. The argument uses the property of differentiability almost everywhere to justify that the derivatives are valid with probability one. The method can be employed to perform sensitivity analysis in a practical setting, where the computed derivatives are sensitivity indicators that can be used to gauge the status of a particular problem. We mentioned that the so-called financial Greeks were the initial motivation for this study, and referenced their importance as indicators that traders rely on to make decisions. In this spirit, this method stands as a useful general tool with many potential practical applications, in Finance and otherwise.

There a few paths forward to further build on what was studied here. From a methodological perspective, this study would be nicely complemented by additional analyses on the viability to compute higher-order derivatives in multistage stochastic linear problems. From a practical perspective, it would be interesting to see the development of an application where a sensitivity indicator could be as relevant in its respective field to inform decisions as the Financial Greeks are in derivatives trading.

The second essay established a framework with which to study the economic benefits of the combined operations of energy storage with an intermittent renewable source such as wind or solar.

It defines a strict contrast between joint and separate planning, and underlines the difference between a storage that can independently and directly participate in the markets, and one that can merely assist the variable renewable source in balancing its deviations in production. In particular, the essay describes that in a setting in which bidding is made on a spot market consisting of a day-ahead market and an intraday market, where an initial bid is made on the former and an adjustment is made on the latter to compensate for deviations observed between actual and planned production, there is no economic advantage in pairing a storage unit with a variable renewable source in joint planning.

This is admittedly a statement that will be seen as controversial in the community, as it goes against the widespread and often unquestioned belief that there is inherent value to be found in joint planning on an individual scale, either to reduce internal imbalances, to stabilize production output or to couple production with market price development.

The essay lays out exceptions to this statement of equality of profits in joint and separate planning. Price asymmetries to buyer and seller disrupt the statement of strict equality, either in the form of two-price imbalance settlement schemes, in illiquid markets with a bid-ask spread, or in systems charging grid fees to consumers but not to producers. The equality of profits cannot be guaranteed in real-time balancing markets where there is no continuous adjustment, particularly in auctions where scheduling rules are determined before the period of delivery. The comparison between joint and separate planning loses its validity in markets where there are entry barriers restricting the direct participation of storage, as the storage would not generate any profits on its own.

It is nevertheless a point worth making, and stands as the main contribution of the essay, which is to offer a methodology with which to analyze the actual benefits of joint planning, and determine the settings in which there is no advantage.

The third essay in this thesis analyzed the earnings potential of a storage unit taking part in the German secondary control reserve. It focused on the concept of duration in a storage unit, and established an upper limit on the profit potential of a storage unit as a function of its duration.

Solving a deterministic mixed-integer bilinear optimization problem by way of a parametric sweep method, the essay determined an upper bound on the profits a storage unit would have had with an optimal policy on the German secondary control reserve during the one-year period between August 2018 and July 2019. It establishes that profits under these conditions saturate and converge to a maximum value as the storage has higher duration values.

The main contributions of this study were two-fold. First, it offers a description of the problem of optimal bidding with a storage unit on the German secondary control reserve, which is inherently a bilinear optimization problem, and proposes a practical method to solve this problem. Second, it offers insights on the dimensioning of a storage unit, namely that the duration of a storage unit determines the earnings potential on the secondary control reserve and that a storage unit with a duration of 13 hours, which stands at just slightly above three times the 4-hour auction period of delivery in this market, will already reach the maximum potential profits to be had. A larger storage unit will not have higher potential profits.

These insights tell the story of what happens when a storage unit operates optimally with perfect information. There are ways to build upon what is discussed in this essay. It would be interesting to see how this upper bound changes when the problem is no longer deterministic. Having a stochastic formulation of this problem would enhance this study by giving a more realistic upper bound to inform investors and decision makers. Furthermore, this study focused on an exclusive operation on the secondary control reserve. A study allowing the storage unit to additionally bid on a real-time market to manage its storage levels, such as in the numerical example of the second essay, would contribute to expanding the combination of feasible price bids accessible to the storage unit, which is likely to improve the potential profits curve.

Bibliography

I. Adler and R. Monteiro. A geometric view of parametric linear programming. *Algorithmica*, 8:161–176, 12 1992.

E. Akbari, R.-A. Hooshmand, M. Gholipour, and M. Parastegari. Stochastic programming-based optimal bidding of compressed air energy storage with wind and thermal generation units in energy and reserve markets. *Energy*, 171:535 – 546, 2019.

A.T. Al-Awami and M.A. El-Sharkawi. Coordinated trading of wind and thermal energy. *IEEE Transactions on Sustainable Energy*, 2(3):277–287, July 2011. .

M.S. Al-Swaiti, A.T. Al-Awami, and M.W. Khalid. Co-optimized trading of wind-thermal-pumped storage system in energy and regulation markets. *Energy*, 138:991 – 1005, 2017.

V. Alimisis and N. D. Hatziargyriou. Evaluation of a hybrid power plant comprising used ev-batteries to complement wind power. *IEEE Transactions on Sustainable Energy*, 4(2):286–293, April 2013.

American Public Power Association. Retail Electric Rates in Deregulated and Regulated States. Technical report, APPA, 2018.

M. Aneke and M. Wang. Energy storage technologies and real life applications – a state of the art review. *Applied Energy*, 179:350–377, 10 2016.

A. Attarha, N. Amjady, S. Dehghan, and B. Vatani. Adaptive robust self-scheduling for a wind producer with compressed air energy storage. *IEEE Transactions on Sustainable Energy*, 9(4):1659–1671, Oct 2018. .

C. Balardy. An empirical analysis of the bid-ask spread in the german power continuous market. Technical report, CEEM, September 2018.

V. Bally and G. Pagès. A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. *Bernoulli*, 9(6):1003–1049, 2003.

A. Banshwar, N.K. Sharma, Y.R. Sood, and R. Shrivastava. Market-based participation of energy storage scheme to support renewable energy sources for the procurement of energy and spinning reserve. *Renewable Energy*, 135:326 – 344, 2019.

A. Barbry, M.F. Anjos, E. Delage, and K.R. Schell. Robust self-scheduling of a price-maker energy storage facility in the new york electricity market. *Energy Economics*, 78:629 – 646, 2019.

S. Barnett. A simple class of parametric linear programming problems. *Operations Research*, 16(6):1160–1165, 1968.

G.N. Bathurst and G. Strbac. Value of combining energy storage and wind in short-term energy and balancing markets. *Electric Power Systems Research*, 67(1):1 – 8, 2003.

L. Bayón, J.M. Grau, M.M. Ruiz, and P.M. Suárez. A comparative economic study of two configurations of hydro-wind power plants. *Energy*, 112:8 – 16, 2016.

P. Billingsley. *Probability and Measure*. Wiley Series in Probability and Statistics. Wiley, 1986.

J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer New York, 2011.

F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

Bloomberg New Energy Finance. Energy Storage Outlook 2019. Technical report, BNEF, 2019.

J. Bochnak, M. Coste, and M.F. Roy. *Real Algebraic Geometry*. A Series of Modern Surveys in Mathematics. Springer Berlin Heidelberg, 2013.

R. Bollapragada, S. Kuppusamy, and U.S. Rao. Component procurement and end product assembly in an uncertain supply and demand environment. *International Journal of Production Research*, 53(3):969–982, 2015. Production Planning.

N. Bourbaki, H. G. Eggleston, and S. Madan. *Topological Vector Spaces*. Éléments de mathématique. Springer-Verlag, 1987.

M. Branda and J. Dupačová. Approximation and contamination bounds for probabilistic programs. *Annals of Operations Research*, 193(1):3–19, Mar 2012.

British Petroleum. BP Statistical Review of World Energy 2019. Technical report, BP, 2019.

S. Bruno, S. Ahmed, A. Shapiro, and A. Street. Risk neutral and risk averse approaches to multistage renewable investment planning under uncertainty. *European Journal of Operational Research*, 250(3):979 – 989, 2016.

E.D. Castronuovo and J.A. Peças Lopes. Optimal operation and hydro storage sizing of a wind–hydro power plant. *International Journal of Electrical Power & Energy Systems*, 26(10):771 – 778, 2004.

L. Clewlow and C. Strickland. *Energy Derivatives: Pricing and Risk Management.* Lacima Publications, 2000.

Consentec. Description of load-frequency control concept and market for control reserves. Technical report, Consentec, 2014. Available at https://regelleistung.net/.

Consentec. Beschreibung systemausgleich und regelreservemärkte. Technical report, Consentec, 2020. Available at https://www.regelleistung.net/ext/static/market-information.

M. Courtillot. New methods in mathematical programming—on varying all the parameters in a linear-programming problem and sequential solution of a linear-programming problem. *Operations Research*, 10(4):471–475, 1962.

J.L. Crespo-Vazquez, C. Carrillo, E. Diaz-Dorado, J.A. Martinez-Lorenzo, and Md. Noor-E-Alam. A machine learning based stochastic optimization framework for a wind and storage power plant participating in energy pool market. *Applied Energy*, 232:341 – 357, 2018.

P. Denholm, J. Nunemaker, P. Gagnon, and W. Cole. The potential for battery energy storage to provide peaking capacity in the united states. *Renewable Energy*, 2019.

W. Dent, R. Jagannathan, and M. R. Rao. Parametric linear programming: Some special cases. *Naval Research Logistics Quarterly*, 20(4):725–728, 1973.

G. Díaz, J. Coto, and J. Gómez-Aleixandre. Optimal operation value of combined wind power and energy storage in multi-stage electricity markets. *Applied Energy*, 235: 1153 – 1168, 2019.

H. Ding, Z. Hu, and Y. Song. Stochastic optimization of the daily operation of wind farm and pumped-hydro-storage plant. *Renewable Energy*, 48:571 – 578, 2012.

H. Ding, Z. Hu, and Y. Song. Rolling optimization of wind farm and energy storage system in electricity markets. *IEEE Transactions on Power Systems*, 30(5):2676–2684, Sep. 2015. .

H. Ding, P. Pinson, Z. Hu, and Y. Song. Integrated bidding and operating strategies for wind-storage systems. *IEEE Transactions on Sustainable Energy*, 7(1):163–172, Jan 2016a. .

H. Ding, P. Pinson, Z. Hu, and Y. Song. Optimal offering and operating strategies for wind-storage systems with linear decision rules. *IEEE Transactions on Power Systems*, 31 (6):4755–4764, Nov 2016b. .

H. Ding, P. Pinson, Z. Hu, J. Wang, and Y. Song. Optimal offering and operating strategy for a large wind-storage system as a price maker. *IEEE Transactions on Power Systems*, 32(6):4904–4913, Nov 2017. .

DOE. Doe global energy storage database. Technical report, US Department of Energy, 2020. Available at https://www.energystorageexchange.org/.

J. Dong, F. Gao, X. Guan, Q. Zhai, and J. Wu. Storage sizing with peak-shaving policy for wind farm based on cyclic markov chain model. *IEEE Transactions on Sustainable Energy*, 8(3):978–989, July 2017. .

T.B. Duarte, D.M. Valladão, and A. Veiga. Asset liability management for open pension schemes using multistage stochastic programming under solvency-ii-based regulatory constraints. *Insurance: Mathematics and Economics*, 77:177 – 188, 2017.

J. Dupačová. The minimax approach to stochastic programming and an illustrative application. *Stochastics*, 20(1):73–88, 1987.

J. Dupačová. Stability and sensitivity-analysis for stochastic programming. *Annals of Operations Research*, 27(1):115–142, 1990.

E. Ela, M. Milligan, A. Bloom, A. Botterud, A. Townsend, and T. Levin. Evolution of wholesale electricity market design with increasing levels of renewable generation. Technical Report NREL/TP-5D00-61765, NREL Publications Database, September 2014.

Jonas Engels, Bert Claessens, and Geert Deconinck. Techno-economic analysis and optimal control of battery storage for frequency control services, applied to the german market. *Applied Energy*, 242:1036 – 1049, 2019. ISSN 0306-2619. . URL .

ENTSO-E, 2018. Available at https://www.entsoe.eu/.

European Energy Exchange, 2019. EEX, Available at https://www.epexspot.com/en/.

European Parliament. Understanding electricity markets in the EU. Technical report, EU, 2016.

Y. Fan, F. Schwartz, S. Voß, and D.L. Woodruff. Stochastic programming for flexible global supply chain planning. *Flexible Services and Manufacturing Journal*, 29(3): 601–633, Dec 2017. Supply Chain Management.

Fraunhofer ISE, 2020. Available at https://www.energy-charts.de/.

R. Freund. Postoptimal analysis of a linear program under simultaneous changes in matrix coefficients. In R. W. Cottle, editor, *Mathematical Programming Essays in Honor of George B. Dantzig Part I*, pages 1–13. Springer, Berlin, Heidelberg, 1985.

G. Fuchs, B. Lunz, M. Leuthold, and D.U. Saur. Technology overview on electricity storage. Technical report, Institute for Power Electronics and Electrical Drives (ISEA), June 2012.

T. Gal and H.J. Greenberg. *Advances in sensitivity analysis and parametric programming*. International series in operations research & management science. Kluwer Academic Publishers, Boston, 1997.

J. Garcia-Gonzalez, R. M. R. de la Muela, L. M. Santos, and A. M. Gonzalez. Stochastic joint optimization of wind generation and pumped-storage units in an electricity market. *IEEE Transactions on Power Systems*, 23(2):460–468, May 2008. .

H. Geman. *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*. The Wiley Finance Series. Wiley, 2009.

A. Ghasemi, S.S. Mortazavi, and E. Mashhour. Hourly demand response and battery energy storage for imbalance reduction of smart distribution company embedded with electric vehicles and wind farms. *Renewable Energy*, 85:124 – 136, 2016.

M. Ghofrani, A. Arabali, M. Etezadi-Amoli, and S. Fadali. Energy storage application for performance enhancement of wind integration. *IEEE Transactions on Power Systems*, 28:4803–4811, 11 2013.

M. Ghofrani, A. Arabali, M. Etezadi-Amoli, and M. S. Fadali. Smart scheduling and cost-benefit analysis of grid-enabled electric vehicles for wind power integration. *IEEE Transactions on Smart Grid*, 5(5):2306–2313, Sep. 2014. .

P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Applications of mathematics : stochastic modelling and applied probability. Springer, 2004.

I.L.R. Gomes, H.M.I. Pousinho, R. Melício, and V.M.F. Mendes. Stochastic coordination of joint wind and photovoltaic systems with energy storage in day-ahead market. *Energy*, 124:310 – 320, 2017.

S. Graf and H. Luschgy. *Foundations of Quantization for Probability Distributions*. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 2000.

F.J. Heredia, M.D. Cuadrado, and C. Corchero. On optimal participation in the electricity markets of wind power plants with battery energy storage systems. *Computers & Operations Research*, 96:316 – 329, 2018.

H.C. Hesse, M. Schimpe, D. Kucevic, and A. Jossen. Lithium-ion battery storage for the grid—a review of stationary battery storage system design tailored for applications in modern power grids. *Energies*, 10(12):1–42, 2017.

H. Holttinen, J. Miettinen, and S. Sillanpää. Wind power forecasting accuracy and uncertainty in Finland. Technical Report ISSN 2242-122X, VTT, 2013.

G.H. Huang and D.P. Loucks. An inexact two-stage stochastic programming model for water resources management under uncertainty. *Civil Engineering and Environmental Systems*, 17(2):95–118, 2000.

J.C. Hull. *Options, Futures, and Other Derivatives*. Pearson Education, 2014.

International Electrotechnical Commission. IEC White Paper Electrical Energy Storage. Technical report, IEC, 2011.

International Energy Agency. Key Electricity Trends 2018. Technical report, IEA, 2018.

International Energy Agency. Electricity Information 2019. Technical report, IEA, 2019a.

International Energy Agency. Tracking Energy Integration. Technical report, IEA, 2019b.

P. Jäckel. *Monte Carlo Methods in Finance.* The Wiley Finance Series. Wiley, 2002.

A. Jaramillo Duque, E.D. Castronuovo, I. Sánchez, and J. Usaola. Optimal operation of a pumped-storage hydro plant that compensates the imbalances of a wind power producer. *Electric Power Systems Research*, 81(9):1767 – 1777, 2011.

R. Jiang, J. Wang, and Y. Guan. Robust unit commitment with wind power and pumped storage hydro. *IEEE Transactions on Power Systems*, 27(2):800–810, May 2012. .

M. E. Khodayar and M. Shahidehpour. Stochastic price-based coordination of intrahour wind energy and storage in a generation company. *IEEE Transactions on Sustainable Energy*, 4(3):554–562, July 2013. .

M. E. Khodayar, L. Abreu, and M. Shahidehpour. Transmission-constrained intrahour coordination of wind and pumped-storage hydro units. *IET Generation, Transmission Distribution*, 7(7):755–765, July 2013. .

C.S. Khor, A. Elkamel, and P.L. Douglas. Stochastic refinery planning with risk management. *Petroleum Science and Technology*, 26(14):1726–1740, 2008.

A.S. Kocaman and V. Modi. Value of pumped hydro storage in a hybrid energy generation and allocation system. *Applied Energy*, 205(C):1202–1215, 2017.

M.C. Kocer, C. Cengiz, M. Gezer, D. Gunes, M.A. Cinar, B. Alboyaci, and A. Onen. Assessment of battery storage technologies for a turkish power network. *Sustainability*, 11(13):1–33, 2019.

J. Kondoh, I. Ishii, H. Yamaguchi, A. Murata, K. Otani, K. Sakuta, N. Higuchi, S. Sekine, and M. Kamimoto. Electrical energy storage systems for energy networks. *Energy Conversion and Management*, 41(17):1863 – 1874, 2000.

J. LaFrance and L.D. Barney. The envelope theorem in dynamic optimization. *Journal of Economic Dynamics and Control*, 15:355–385, 02 1991.

G. Lai, F. Margot, and N. Secomandi. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research*, 58(3):564–582, 2010.

Y. Liu, C. Jiang, J. Shen, and J. Hu. Coordination of hydro units with wind power generation using interval optimization. *IEEE Transactions on Sustainable Energy*, 6 (2):443–453, April 2015. .

Y. Liu, C. Jiang, J. Shen, J. Hu, and Y. Luo. Coordination of hydro units with wind power generation based on raroc. *Renewable Energy*, 80:783 – 792, 2015.

N. Löhndorf and A. Shapiro. Modeling time-dependent randomness in stochastic dual dynamic programming. *European Journal of Operational Research*, 273(2):650 – 661, 2019.

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

Wood Mackenzie. Global energy storage outlook 2019. Technical report, Wood Mackenzie, April 2019.

T. Malakar, S.K. Goswami, and A.K. Sinha. Impact of load management on the energy management strategy of a wind-short hydro hybrid system in frequency based pricing. *Energy Conversion and Management*, 79:200 – 212, 2014.

E. Mardan, M.S. Amalnik, and M. Rabbani. An integrated emergency ordering and production planning optimization model with demand and yield uncertainty. *International Journal of Production Research*, 53(20):6023–6039, 2015. Inventory Management.

D.H. Martin. On the continuity of the maximum in parametric linear programming. *Journal of Optimization Theory and Applications*, 17:205–210, 05 1975.

I.G. Mason. Comparative impacts of wind and photovoltaic generation on energy storage for small islanded electricity systems. *Renewable Energy*, 80:793 – 805, 2015.

B. Mauch, P.M.S. Carvalho, and J. Apt. Can a wind farm with CAES survive in the day-ahead market? *Energy Policy*, 48:584 – 593, 2012.

P. Medina, A. W. Bizuayehu, J. P. S. Catalão, E. M. G. Rodrigues, and J. Contreras.

Electrical energy storage systems: Technologies' state-of-the-art, techno-economic benefits and applications analysis. In *2014 47th Hawaii International Conference on System Sciences*, pages 2295–2304, 2014.

P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70 (2):583–601, 2002.

I.G. Moghaddam, M. Nick, F. Fallahi, M. Sanei, and S. Mortazavi. Risk-averse profit-based optimal operation strategy of a combined wind farm–cascade hydro system in an electricity market. *Renewable Energy*, 55:252 – 259, 2013.

J. Moradi, H. Shahinzadeh, A. Khandan, and M. Moazzami. A profitability investigation into the collaborative operation of wind and underwater compressed air energy storage units in the spot market. *Energy*, 141:1779 – 1794, 2017.

M.W. Murage and C.L. Anderson. Contribution of pumped hydro storage to integration of wind power in kenya: An optimal control approach. *Renewable Energy*, 63:698 – 707, 2014.

J. Márquez Angarita and J. Garcia Usaola. Combining hydro-generation and wind energy: Biddings and operation on electricity spot markets. *Electric Power Systems Research*, 77(5):393 – 400, 2007.

E. Nasrolahpour, S. J. Kazempour, H. Zareipour, and W. D. Rosehart. Strategic sizing of energy storage facilities in electricity markets. *IEEE Transactions on Sustainable Energy*, 7(4):1462–1472, Oct 2016. ISSN 1949-3037. .

Netzregelverbund, 2019a. Available at https://regelleistung.net/.

Netzregelverbund. Präqualifikationsverfahren für Regelreserveanbieter (FCR, aFRR, mFRR) in Deutschland ("PQ-Bedingungen"). Technical report, Netzregelverbund, 2019b. Available at https://www.regelleistung.net.

A.V. Ntomaris and A.G. Bakirtzis. Stochastic scheduling of hybrid power stations in insular power systems with high wind penetration. *IEEE Transactions on Power Systems*, 31:3424–3436, 2016.

M.Z. Oskouei and A.S. Yazdankhah. Scenario-based stochastic optimal operation of wind, photovoltaic, pump-storage hybrid system in frequency- based pricing. *Energy Conversion and Management*, 105:1105 – 1114, 2015.

A. Oudalov, D. Chartouni, and C. Ohler. Optimizing a battery energy storage system for primary frequency control. *IEEE Transactions on Power Systems*, 22(3):1259–1266, 2007.

S. V. Papaefthymiou, E. G. Karamanou, S. A. Papathanassiou, and M. P. Papadopoulos. A Wind-Hydro-Pumped Storage Station Leading to High RES Penetration in the Autonomous Island System of Ikaria. *IEEE Transactions on Sustainable Energy*, 1 (3):163–172, Oct 2010.

M. Parastegari, R.-A. Hooshmand, A. Khodabakhshian, and A.-H. Zare. Joint operation of wind farm, photovoltaic, pump-storage and energy storage devices in energy and reserve markets. *International Journal of Electrical Power & Energy Systems*, 64:275 – 284, 2015.

M.V.F. Pereira and L.M.V.G. Pinto. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(2):359–375, 1991.

P. Pflaum, M. Alamir, and M.Y. Lamoudi. Battery sizing for pv power plants under regulations using randomized algorithms. *Renewable Energy*, 113(C):596–607, 2017.

A. B. Philpott and Z. Guan. On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters*, 36(4):450 – 455, 2008.

D. Pollard. *A User's Guide to Measure Theoretic Probability.* Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2001.

W.B. Powell. *Approximate dynamic programming. Solving the curses of dimensionality.* Wiley, 2011.

M. Rafiee and F. Kianfar. A scenario tree approach to multi-period project selection problem using real-option valuation method. *The International Journal of Advanced Manufacturing Technology*, 56(1):411–420, Sep 2011.

R.T. Rockafellar. *Convex Analysis.* Princeton Landmarks in Mathematics and Physics. Princeton University Press, 1970.

W. Rudin. *Principles of mathematical analysis.* International series in pure and applied mathematics. McGraw-Hill, 1964.

W. Rudin. *Real and Complex Analysis.* Tata McGraw-Hill, 1987.

P.A. Samuelson. *Foundations of economic analysis.* Harvard economic studies. Harvard University Press, 1947.

A. A. Sánchez de la Nieta, J. Contreras, and J. I. Muñoz. Optimal coordinated wind-hydro bidding strategies in day-ahead markets. *IEEE Transactions on Power Systems*, 28 (2):798–809, May 2013. .

A. A. Sánchez de la Nieta, J. Contreras, J. I. Muñoz, and J. P. S. Catalão. Optimal wind reversible hydro offering strategies for midterm planning. *IEEE Transactions on Sustainable Energy*, 6(4):1356–1366, Oct 2015. .

A. A. Sánchez de la Nieta, J. Contreras, and J. P. S. Catalão. Optimal single wind hydro-pump storage bidding in day-ahead markets including bilateral contracts. *IEEE Transactions on Sustainable Energy*, 7(3):1284–1294, July 2016. .

W.-P. Schill and C. Kemfert. Modeling strategic electricity storage: The case of pumped hydro storage in germany. *The Energy Journal*, Volume 32(Number 3):59–88, 2011.

E.S. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973, 1997.

A. Shapiro. On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters*, 37(3):143 – 147, 2009.

A. Shapiro. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63 – 72, 2011.

A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory, Second Edition.* MOS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, 2014.

Z. Shu and P. Jirutitijaroen. Optimal operation strategy of energy storage system for grid-connected wind power plants. *IEEE Transactions on Sustainable Energy*, 5(1): 190–199, Jan 2014. .

R. Sioshansi. When energy storage reduces social welfare. *Energy Economics*, 41(C): 106–116, 2014.

R. Sioshansi and P. Denholm. The value of concentrating solar power and thermal energy storage. *IEEE Transactions on Sustainable Energy*, 1(3):173–183, Oct 2010. .

Iain Staffell and Mazda Rustomji. Maximising the value of electricity storage. *Journal of Energy Storage*, 8:212 – 225, 2016. ISSN 2352-152X. . URL .

C. Su, C. Cheng, P. Wang, J. Shen, and X. Wu. Optimization model for long-distance integrated transmission of wind farms and pumped-storage hydropower plants. *Applied Energy*, 242:285 – 293, 2019.

K. Sun, K.-J. Li, J. Pan, Y. Liu, and Y. Liu. An optimal combined operation scheme for pumped storage and hybrid wind-photovoltaic complementary power generation system. *Applied Energy*, 242:1155 – 1163, 2019.

A. Takayama. *Mathematical Economics*. Cambridge University Press, 1985.

Z.-F. Tan, L.-W. Ju, H.-H. Li, J.-Y. Li, and H.-J. Zhang. A two-stage scheduling optimization model and solution algorithm for wind power and energy storage system considering uncertainty and demand response. *International Journal of Electrical Power & Energy Systems*, 63:1057 – 1069, 2014.

D. Teneng. Limitations of the black-scholes model. *International Research Journal of Finance and Economics*, pages 99–102, 01 2011.

C. Tham and B. Cao. Stochastic programming methods for workload assignment in an ad hoc mobile cloud. *IEEE Transactions on Mobile Computing*, 17(7):1709–1722, July 2018.

A. A. Thatte, L. Xie, D. E. Viassolo, and S. Singh. Risk measure based robust bidding strategy for arbitrage using a wind farm and energy storage. *IEEE Transactions on Smart Grid*, 4(4):2191–2199, Dec 2013.

US Congress. H.R.776 - Energy Policy Act of 1992. *United States Congress*, 1992.

US Department of Energy. Energy Storage Technology and Cost Characterization Report. Technical report, HydroWires DOE, 2019.

D.M. Valladão, A. Veiga, and G. Veiga. A multistage linear stochastic programming model for optimal corporate debt management. *European Journal of Operational Research*, 237(1):303 – 311, 2014.

R. M. Van Slyke and R. Wets. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics*, 17(4): 638–663, 1969.

H.R. Varian. *Microeconomic Analysis*. Norton International edition. Norton, 1992.

A.K. Varkani, A. Daraeepour, and H. Monsef. A new self-scheduling strategy for integrated operation of wind and pumped-storage power plants in power markets. *Applied Energy*, 88(12):5002 – 5012, 2011.

C. Villani. *Optimal Transport: Old and New*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2008.

Q. Wang, J. Wang, and Y. Guan. Price-based unit commitment with wind power utilization constraints. *IEEE Transactions on Power Systems*, 28(3):2718–2726, Aug 2013.

J.E. Ward and R.E. Wendell. Approaches to sensitivity analysis in linear programming. *Annals of Operations Research*, 27(1):3–38, Dec 1990.

L.B. Willner. On parametric linear programming. *SIAM Journal on Applied Mathematics*, 15(5):1253–1257, 1967.

WindGuard, 2019. Available at https://www.windguard.de/.

Y. Wu. A stochastic model for production loading in a global apparel manufacturing company under uncertainty. *Production Planning & Control*, 22(3):269–281, 2011. Supply Chain Management.

U. Yıldıran and I. Kayahan. Risk-averse stochastic model predictive control-based real-time operation method for a wind energy generation system supported by a pumped hydro storage unit. *Applied Energy*, 226:631 – 643, 2018.

A. Zeh, M. Mueller, H. C. Hesse, A. Jossen, and R. Witzmann. Operating a multitasking stationary battery storage system for providing secondary control reserve on low-voltage level. In *International ETG Congress 2015; Die Energiewende - Blueprints for the new energy age*, pages 1–8, 2015.

X. Zhang, M. Prajapati, and E. Peden. A stochastic production planning model under uncertain seasonal demand and market growth. *International Journal of Production Research*, 49(7):1957–1975, 2011.

H. Zhao, Q. Wu, S. Hu, H. Xu, and C.N. Rasmussen. Review of energy storage system for wind power integration support. *Applied Energy*, 137:545 – 553, 2015.