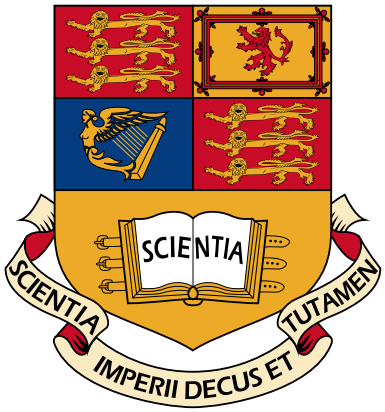
Imperial College London
Department of Computing



Decision Rule Approximations for Dynamic Optimization under Uncertainty

Phebe Theofano Vayanos

Submitted in part fulfillment of the requirements for the degree of
Doctor of Philosophy in Computing of Imperial College London and
the Diploma of Imperial College, April 2012

Abstract

Dynamic decision problems affected by uncertain data are notoriously hard to solve due to the presence of adaptive decision variables which must be modeled as functions or *decision rules* of some (or all) of the uncertain parameters. All exact solution techniques suffer from the curse of dimensionality while most solution schemes assume that the decision-maker cannot influence the sequence in which the uncertain parameters are revealed.

The main objective of this thesis is to devise tractable approximation schemes for dynamic decision-making under uncertainty. For this purpose, we develop new *decision rule approximations* whereby the adaptive decisions are approximated by finite linear combinations of prescribed basis functions.

In the first part of this thesis, we develop a tractable unifying framework for solving convex multi-stage robust optimization problems with general nonlinear dependence on the uncertain parameters. This is achieved by combining decision rule and constraint sampling approximations. The synthesis of these two methodologies provides us with a versatile data-driven framework, which circumvents the need for estimating the distribution of the uncertain parameters and offers almost complete freedom in the choice of basis functions. We obtain a-priori probabilistic guarantees on the feasibility properties of the optimal decision rule and demonstrate asymptotic consistency of the approximation.

We then investigate the problem of hedging and pricing path-dependent electricity derivatives such as swing options, which play a crucial risk management role in today's deregulated energy markets. Most of the literature on the topic assumes that a swing option can be assigned a unique fair price. This assumption nevertheless fails to hold in real-world energy markets, where the option admits a whole interval of prices consistent with those of traded instruments. We formulate two large-scale robust optimization problems whose optimal values yield the endpoints of this interval. We analyze and exploit the structure of the optimal decision rule to formulate approximate problems that can be solved efficiently with the decision rule approach discussed in the first part of the thesis.

Most of the literature on stochastic and robust optimization assumes that the sequence in which the uncertain parameters unfold is independent of the decision-maker's actions. Nevertheless, in numerous real-world decision problems, the time of information discovery can be influenced

by the decision-maker. In the last part of this thesis, we propose a decision rule-based approximation scheme for multi-stage problems with decision-dependent information discovery. We assess our approach on a problem of infrastructure and production planning in offshore oil fields.

Acknowledgements

I owe my deepest gratitude to my supervisors, Prof. Berç Rustem and Dr. Daniel Kuhn, for the astounding support and guidance they have given me throughout the course of my PhD. I find the effort and amount of time that they both invested in my academic as well as personal development truly remarkable. This thesis benefited greatly from their encouragements, dedication, enthusiasm and talent. Having them as teachers and mentors has been a true pleasure and privilege.

My sincere gratitude goes to Prof. Danilo Mandic who, with his phenomenal enthusiasm, sparked my interest in research. I would also like to thank my thesis examiners, Dr. Houyuan Jiang and Prof. Stratos Pistikopoulos, for their pertinent, challenging and inspiring comments. I am particularly grateful to Prof. Stratos Pistikopoulos for his keen interest in my research and his kind advice on my career.

I am deeply indebted to Dr. Wolfram Wiesemann who, throughout the course of my PhD, acted as a friend but also as a mentor. I benefited greatly from his insightful questions, invaluable suggestions and words of support.

In addition, I would like to thank Dr. Krysia Broda for her immense trust and understanding. My sincere gratitude also goes to Dr. Amani El-Kholy who, with her extraordinary organizational skills, helped me overcome so many obstacles.

I am grateful to my friends and colleagues, Angelos Georghiou, Iakovos Kakouris, Michalis Kapsos, Nick Baltas, Steve Zymler, Xenia Kleniati, Angelos Tsoukalas, Aris Kountouriotis, Kostas Manolarakis, Juan Jerez, Paula Rocha, Vladimir Roitch, Raquel Fonseca, Eva Kalyvianaki and Vu Ngoc Duy Luong, who made my time as a post-graduate student so enjoyable. My warmest acknowledgments go to Dimitra Bampou and Michael Hadjiyiannis for their constant support, positive feedback, excellent company and incredible humor.

I cannot begin to express my gratitude to George Elissaios for his remarkable patience, understanding and encouragements. Over the past four years, he has provided me with continuous support and advice, as much on technical as on personal matters.

Finally, I would like to thank my family for their immense support, constant encouragements and excellent advice. More than anything, I thank them for always believing in me.

Contents

Abstract **i**

Acknowledgements **iii**

1 Introduction **1**

1.1 Motivation and Objectives 1

1.2 Contributions and Structure of the Thesis 13

1.3 Statement of Originality 16

2 A Constraint Sampling Approach for Multi-Stage Robust Optimization **17**

2.1 Introduction 18

2.2 Problem Description 22

2.3 Tractable Reformulation 24

2.3.1 Decision Rule Approximation 24

2.3.2 Constraint Sampling Approximation 29

2.4 Complexity Analysis 30

2.5 Convergence Analysis 34

2.5.1 Convergence of the Decision Rule Approximation 36

2.5.2 Convergence of the Sampling Approximation 43

2.6 Examples from Inventory Management 46

2.6.1 Inventory Management with Cumulative Order Constraints 46

2.6.2 Inventory Management with Random Yield 47

2.7 Numerical Experiments 49

2.8 Appendix 53

2.8.1 Proof of Lemma 2.4 53

2.8.2 Almost Sure Density of Sample Sequences 55

3 Hedging Electricity Swing Options in Incomplete Markets 57

3.1 Introduction 58

3.2 Problem Formulation 61

3.2.1 Electricity Swing Options 62

3.2.2 Market Model 63

3.2.3 Holder's Problem 67

3.2.4 Writer's Problem 72

3.2.5 No-Arbitrage Interval 77

3.3 The Optimal Exercise Strategy 81

3.4 Tractable Reformulation 86

3.4.1 Aggregation of Time-Periods 87

3.4.2 Decision Rule Approximation 93

3.4.3 Constraint Sampling Approximation 94

3.5 Other Applications 95

3.5.1 Valuation of Mines and Oil Fields 95

3.5.2 Valuation of Power Plants and Refineries 96

3.6 Case Study 98

4 Decision Rules for Information Discovery in Stochastic Programming 101

4.1 Introduction 102

4.1.1 Motivating Examples 102

4.1.2 Literature Review 103

4.2 The Two-Stage Case 105

4.2.1 Problem Formulation 105

4.2.2 Decision Rule Approximation 107

4.2.3 MBLP Reformulation 107

4.3 The Multi-Stage Case 109

4.3.1 Problem Formulation 109

4.3.2 Decision Rule Approximation 110

4.3.3 MBLP Reformulation 113

4.4 Case Study 114

4.4.1 Problem Description 114

4.4.2 Numerical Results 116

5 Conclusion **119**

Bibliography **122**

List of Tables

2.1 Semi-analytical decision rule approximation approaches 19

2.2 Input data for the two test problems 49

4.1 Problem input parameters 116

List of Figures

1.1 Companion figure for Example 1.1 5

1.2 Companion figure for Example 1.2 8

2.1 Approximations of a nonlinear function by linear combinations of basis functions 29

2.2 Companion figure for the proof of Lemma 2.1 37

2.3 A supply chain with T stages, uncertain customer demand and yield. 48

2.4 Optimality gaps for $\mathcal{RB}_{\mathrm{COC}}$ and $\mathcal{SB}^N_{\mathrm{COC}}$ in dependence of ϵ 50

2.5 Empirical violation probabilities for problem $\mathcal{SB}^N_{\mathrm{COC}}$ in dependence of ϵ 51

2.6 Mean solver times and sample numbers for problem $\mathcal{SB}^N_{\mathrm{COC}}$ as a function of $1/\epsilon$ 52

2.7 Empirical violation probabilities for problem $\mathcal{SB}^N_{\mathrm{RY}}$ in dependence of ϵ 53

3.1 Companion figure for Example 3.1 71

3.2 Companion figure for Example 3.2 80

3.3 Companion figure for the proof of Proposition 3.6 84

3.4 Companion figure for the proof of Proposition 3.7 85

3.5 No-arbitrage interval as a function of the approximation parameters 99

3.6 Empirical violation and empirical profit/loss distributions 100

4.1 Offshore oil extraction site and expected NPV in dependence of solver time . . . 117

Chapter 1

Introduction

1.1 Motivation and Objectives

In its basic form, mathematical optimization is concerned with determining the best (in some sense) decisions amongst all decisions that satisfy a set of constraints (or requirements). Specifically, it is concerned with solving problems of the form

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & f_0\left(x, \xi\right) \\ \text{s.t.} \quad & x \in \mathcal{X}(\xi). \end{aligned} \tag{1.1}$$

Here, x denotes the vector of *decision variables* of the problem (e.g., the amount of a product to purchase or the decision to build a factory), while $\xi \in \mathbb{R}^k$ denotes the vector of *problem parameters* (e.g., the demand for a product, the availability of a resource or the capacity of a network). The function $f_0 : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}$ to be minimized is termed *objective* or *cost function*, while the set $\mathcal{X}(\xi) \subseteq \mathbb{R}^n$ is referred to as the *feasible set* of the problem, defined through

$$\mathcal{X}(\xi) := \left\{x \in \mathbb{R}^n \ : \ f_i(x, \xi) \leq 0, \ i = 1, \ldots, I\right\}.$$

The functions $f_i : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}$ are known as the *constraint functions*. The set $\mathcal{X}(\xi)$ is thus used to model constraints on the decision variables (e.g., production may not exceed capacity).

A solution $x^\star \in \mathcal{X}(\xi)$ to (1.1) is termed *optimal* if it satisfies $f_0(x^\star, \xi) \leq f_0(x, \xi)\ \forall x \in \mathcal{X}(\xi)$.

Problem (1.1) is a *convex optimization problem* if its objective and constraint functions are all convex in x. Amongst the most common classes of convex optimization problems are (in order of increasing generality) linear programs, second-order cone programs and semi-definite programs. While these do not typically admit an analytical solution, they can be solved numerically by means of interior point methods with a number of operations *polynomial* in the dimensions of the problem. We note that although linear, second-order cone and semi-definite programs are all efficiently solvable, the computational effort required per iteration for problems of similar size and structure increases with the generality of the problem class (see e.g., Alizadeh and Goldfarb [2003]). We refer the interested reader to Vandenberghe and Boyd [2004] for an in-depth review of the field of convex optimization.

Optimization problems arise naturally in a vast number of disciplines ranging from safety engineering (e.g., design of mechanical structures and buildings), energy engineering (e.g., power plant dispatch policy optimization, capacity expansion planning, electricity grid optimization), chemical engineering (e.g., protein structure prediction), aerospace engineering (e.g., aircraft and spacecraft design, air traffic control) to economics and finance (e.g., asset valuation, investment portfolio optimization). Moreover, the vast majority of decision problems encountered by practitioners are affected by uncertainty (e.g., about the prices of assets, the demand for a product or the availability of a resource). In the context of our problem formulation (1.1), this implies that some (or all) of the parameters ξ are uncertain: mathematically, they must be modeled by random variables. Consider for example the problem faced by an energy producer who needs to decide whether to build additional wind turbines so as to maximize her profits. Clearly, the profits of the producer will depend on the demand for electricity, the price of electricity and the availability of wind throughout the lifetime of the turbines, all of which are uncertain. A natural question which then arises is the following: how should one (if at all) take uncertainty into account in the decision-making process? More specifically, how can one incorporate uncertainty in problems of the form (1.1)?

Broadly speaking, the answers to these questions depend on the nature of the uncertain pa-

rameters and on the risk-preferences of the decision-maker. One can distinguish between two types of uncertain parameters: *aleatoric uncertainties* and *epistemic uncertainties* (Ellsberg [1961]). If a probability distribution can be assigned to the uncertain parameters, then one faces aleatoric uncertainty, i.e., uncertainty in view of the specific realization of the uncertainties. If the probability model is itself unknown, one faces epistemic uncertainty, i.e., uncertainty in view of the distribution of the uncertainties (this is often termed *ambiguity*). Whether one is in the face of aleatoric or epistemic uncertainty depends on numerous factors. For example, our inability to accurately measure the uncertain parameters under consideration or the lack of sufficient historical data might introduce model risk and thus the necessity to view the distributions as uncertain. At the same time, ambiguity might also arise from certain simplifications purposefully made by the modeler. Consider for example the demand for certain types of goods. This demand might be highly dependent on numerous factors which are themselves uncertain (e.g., the state of the economy). The modeler might prefer neglecting certain of these factors (potentially due to the inability to determine all of them or to fully explain their impact on the uncertainties of interest). This simplification would necessitate viewing the demand for the goods as epistemic uncertainties.

Optimization offers a vast number of versatile frameworks for taking different types of uncertainties and risk-preferences into account.

If the decision problem is affected by aleatoric uncertainties with (known) distribution $\mathbb{P}$, the resulting optimization problem can be formulated as a *stochastic optimization problem* or *stochastic program*, which takes the following form

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & \mathbb{F}[f_0(x, \xi)] \\ \text{s.t.} \quad & x \in \mathcal{X}(\xi) \quad \mathbb{P}\text{-a.s.} \end{aligned} \tag{1.2}$$

Here, $\mathbb{F}$ denotes a probability functional with respect to $\mathbb{P}$ (e.g., mean, variance, quantile or combination thereof) and may be used to model the risk-preferences of the decision maker. The constraints in (1.2) are required to hold *almost surely* (denoted "a.s."), i.e., for all possible realizations of ξ, except perhaps for sets with zero probability. Such stochastic programming

formulations date at least as far back as the papers from Markowitz [1952], Dantzig [1955] and Beale [1955]. Oftentimes, the decision-maker may tolerate violations in the constraints of (1.2). This behavior is formalized by replacing the "almost sure" constraints in (1.2) by so-called *chance constraints* of the form

$$\mathbb{P}(x \in \mathcal{X}(\xi)) \geq 1 - \epsilon,$$

where $\epsilon \in [0, 1)$ is the violation probability tolerated by the decision-maker. Chance-constrained optimization thus provides probabilistic *feasibility guarantees* with respect to the realization of ξ. Research on chance-constrained programming dates back to the works of Charnes and Cooper [1959], Miller and Wagner [1965] and Prékopa [1970]. The presence of uncertain parameters in decision problems complicates their solution considerably. We now investigate by means of a classic example drawn from the field of logistics how the explicit treatment of ξ as a random variable can help decision-makers improve the performance of their decisions.

Example 1.1 (Newsvendor problem) *Each day, a retailer receives orders from her customers for an amount ξ of a perishable product. This demand must be satisfied from the on-hand inventory of the retailer who may place advance orders with a supplier for an amount x of the product. The retailer may purchase the product at a cost £0.8 per unit product and sells it at a price £1 per unit. At the end of the day, any products left-over are worthless. The profit of the retailer can be expressed mathematically as*

$$r(x, \xi) = \min(x, \xi) - 0.8x.$$

Suppose first that the retailer knows the demand precisely, i.e., $\xi = \hat{\xi}$. Then the optimal order quantity $x^\star$ is equal to $\hat{\xi}$. Suppose to the contrary that the retailer only knows that the demand is uniformly distributed in the range $[50, 200]$. A typical approach followed by retailers is to neglect this uncertainty by effectively fixing the uncertain demand to a nominal value such as its expected value, $\mathbb{E}(\xi) = 125$. Under this strategy, the retailer would order 125 units of the product. We now investigate how the expected profit of the retailer changes in dependence of

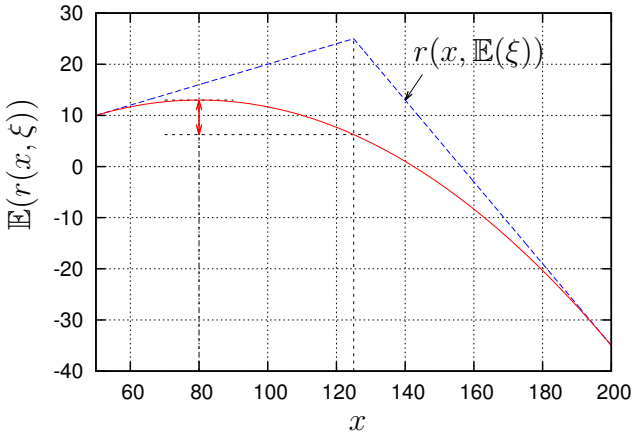


Figure 1.1: Companion figure for Example 1.1. The filled line illustrates the function $\mathbb{E}(r(x,\xi))$ whose maximum is attained for an order quantity of 80; The dashed line illustrates the function $r(x,\mathbb{E}(\xi))$ whose maximum is attained for an order quantity of 125. By taking uncertainty explicitly into account when choosing her order quantity, the retailer can increase her expected profits from £6.25 to £13. The optimality gap incurred by fixing random parameters to their expectations is illustrated by the double arrow.

the order quantity,

$$\begin{aligned}
\mathbb{E}(r(x,\xi)) &= \mathbb{E}(\min(x,\xi) - 0.8x) \\
&= \int_{-\infty}^{\infty} \min(x,\xi) - 0.8x \; \mathbb{P}(\mathrm{d}\xi) \\
&= \frac{1}{150}\left\{\int_{50}^{x} \xi - 0.8x \; \mathrm{d}\xi + \int_{x}^{200} 0.2x \; \mathrm{d}\xi\right\} \\
&= \frac{1}{150}\left\{-0.5x^2 + 80x - 1250\right\}.
\end{aligned}$$

The expected payoff of the retailer is thus a concave function attaining its maximum for an order quantity of $x^\star = 80$*, which corresponds to an expected profit of £13, see Figure 1.1. If the retailer had selected to order* $\mathbb{E}(\xi)$ *units of the product, her expected profit would be £6.25. Thus, by taking uncertainty into account explicitly, the retailer can double her expected profits. In fact, as can be seen from Figure 1.1,* $\mathbb{E}(r(x,\xi)) \leq r(x,\mathbb{E}(\xi))$ *(this statement can be derived from Jensen's inequality by noting that* $r(x,\cdot)$ *is concave for each* $x \in \mathbb{R}$*) and* $\arg\max_{x\in\mathbb{R}} \mathbb{E}(r(x,\xi)) \neq \arg\max_{x\in\mathbb{R}} r(x,\mathbb{E}(\xi))$.

If the decision problem is affected by epistemic uncertainties, broadly speaking two cases arise: the decision-maker only has (potentially partial) information on the support of the distribution

of the uncertainties or he also has additional distributional information (e.g., knowledge of certain generalized moments of ξ). In the first case, the decision problem is typically formulated as a *robust optimization problem* of the form

$$\begin{array}{ll} \min\limits_{x\in\mathbb{R}^n} & \max\limits_{\xi\in\Xi} \; f_0\left(x,\xi\right) \\ \text{s.t.} & x \in \mathcal{X}(\xi) \quad \forall \xi \in \Xi. \end{array} \tag{1.3}$$

Here, Ξ is referred to as *uncertainty set* and corresponds the set of uncertainties against which the decision-maker wishes to be *immunized*. In the latter case where the decision-maker possesses additional knowledge, distributional information may be incorporated into the decision problem by formulating a so-called *distributionally robust problem* as follows.

$$\begin{array}{ll} \min\limits_{x\in\mathbb{R}^n} & \max\limits_{\mathbb{Q}\in\mathcal{Q}} \; \mathbb{E}^{\mathbb{Q}}(f_0\left(x,\xi\right)) \\ \text{s.t.} & \mathbb{Q}(x \in \mathcal{X}(\xi)) \geq 1-\epsilon \quad \forall \mathbb{Q} \in \mathcal{Q}, \end{array} \tag{1.4}$$

where $\mathcal{Q}$ denotes the set of all distributions which conform with the partial distributional information available to the decision-maker and $\mathbb{E}^{\mathbb{Q}}(\cdot)$ denotes the expectation operator with respect to the distribution $\mathbb{Q}$. Note that if $\mathcal{Q}$ is a singleton, (1.4) reduces to (1.2). If the objective in (1.4) is deterministic, the resulting problem is typically referred to as *ambiguous chance-constrained program*. The research on distributionally robust optimization dates back to the paper of Scarf [1958] and has, since its emergence, received considerable attention (see e.g., Dupačová [1966, 1987], Kall [1988], Calafiore and El Ghaoui [2006], Erdoğan and Iyengar [2006], Shapiro [2006], Chen et al. [2007], Popescu [2007], See and Sim [2009], Delage and Ye [2010], Goh and Sim [2010] and Zymler et al. [2011]). The example below illustrates that the explicit treatment of ξ as a random variable is also necessary in a worst-case context.

Example 1.2 (Robust newsvendor problem) *We revisit Example 1.1 from the point of view of a risk-averse retailer who wishes to select the order quantity x so as to maximize her worst-case profit. We assume that the retailer wishes to be immunized against all possible realizations of ξ and thus define $\Xi := [50, 200]$. In this instance, the retailer may purchase the product at a unit cost of £0.5 and sells it at a unit price of £1. In addition the retailer now*

faces penalties in the amount of £0.2 per unit of unsatisfied demand. The profit of the retailer is thus given by

$$r(x,\xi) = \min(x,\xi) - 0.5x - 0.2\max(\xi - x, 0).$$

We now investigate how the worst-case profit of the retailer changes in dependence of the order quantity,

$$\begin{aligned}\min_{\xi\in\Xi} r(x,\xi) &= \min_{\xi\in[50,200]} \min(x,\xi) - 0.5x - 0.2\max(\xi - x, 0) \\ &= \min\left\{\min_{\xi\in[50,x]} \xi - 0.5x,\ \min_{\xi\in[x,200]} x - 0.5x - 0.2(\xi - x)\right\} \\ &= \min\left\{50 - 0.5x,\ 0.7x - 40\right\}.\end{aligned}$$

The worst-case profit of the retailer is thus a concave function attaining its maximum for an order quantity of $x^\star = 75$, which corresponds to a worst-case profit of £12.5, see Figure 1.2. If the retailer only considers the nominal demand realization, $\mathbb{E}(\xi)$, in the worst-case she makes a loss of £12.5. If the retailer only considers the minimum demand realization, $\min_{\xi\in\Xi}\xi$, in the worst-case she makes a loss of £5.

Example 1.2 illustrates the significant loss in optimality incurred by fixing the realization of the uncertain parameters in the optimization model to a single value. Of course, there are also several feasibility considerations that must be taken into account when dealing with problems incorporating robust constraints, i.e., constraints that must hold almost surely or for all $\xi \in \Xi$. We refer the interested reader to the motivating examples by Ben-Tal and Nemirovski [2000] which illustrate that, for even relatively small variations of the uncertain parameters around their nominal values, constraints may become significantly violated by the optimal solution to the nominal problem.

It is not a-priori clear when the robust optimization problem introduced in (1.3) is computationally tractable since, in its current form, it presents a potentially infinite number of constraints. Research in the field of robust optimization has burgeoned since the end of the 1990s (see e.g., El Ghaoui and Lebret [1997], Ben-Tal and Nemirovski [1998, 1999, 2000] and El Ghaoui et al. [1998]) with aim to address that specific question: when is the so-called *robust counterpart* of an

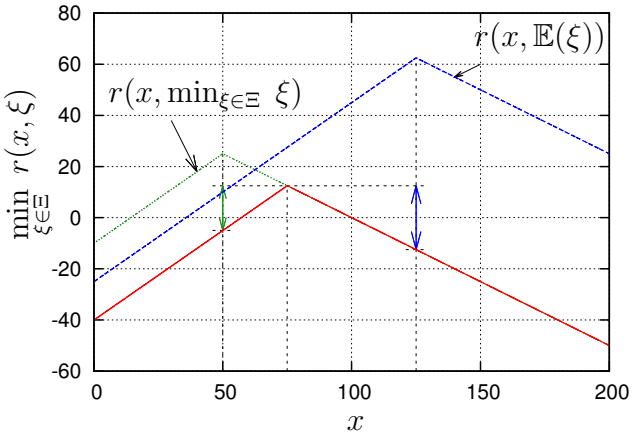


Figure 1.2: Companion figure for Example 1.2. The filled line illustrates the function $\min_{\xi\in\Xi} r(x,\xi)$ whose maximum is attained for an order quantity of 75; The dashed line illustrates the function $r(x,\mathbb{E}(\xi))$ whose maximum is attained for an order quantity of 125; The dotted line illustrates the function $r(x,\min_{\xi\in\Xi}\xi)$ whose maximum is attained for an order quantity of 50. By taking uncertainty explicitly into account when choosing her order quantity, the retailer can not only avoid making significant losses in the worst-case, but is in fact guaranteed to make profits in an amount greater than £12.5. The optimality gap incurred by fixing the demand to its expected value is illustrated by the vertical double arrow on the right; The optimality gap incurred by fixing the demand to its minimum is illustrated by the vertical double arrow on the left.

optimization problem efficiently solvable? It turns out that while robust optimization problems are generically intractable, there are numerous classes of problems of practical interest that admit equivalent tractable reformulations in the form of linear, conic or semi-definite programs. This is the case for example, under mild structural assumptions, if the nominal problem is linear and the uncertainty set is characterized through conic inequalities (Ben-Tal and Nemirovski [1999]) or if the nominal problem is a second-order cone program and the uncertainty set is a single ellipsoid (Ben-Tal and Nemirovski [1998]). Sometimes, however, the robust counterpart of a convex optimization problem does not admit an exact reformulation as a tractable optimization problem. In these cases, one must resort to conservative approximations, yielding feasible yet sub-optimal solutions. This is the case for example if the nominal problem is a semi-definite program and the uncertainty set is ellipsoidal (Ben-Tal and Nemirovski [1998], El Ghaoui et al. [1998]). A thorough review of the field of robust optimization can be found in Ben-Tal et al. [2009].

Motivated by the computational complexity of the methods discussed above, Calafiore and Campi [2005, 2006] suggest to solve robust optimization problems of the form (1.3) by con-

sidering their *scenario counterpart*. Specifically, they suggest to enforce the robust constraints in (1.3) over only a finite subset of Ξ obtained through Monte-Carlo sampling of the uncertain parameters. While the original robust problem may be intractable, the scenario counterpart can be solved efficiently under standard assumptions. Furthermore, this *constraint sampling* approach offers a number of advantages over the semi-analytical approaches discussed above. First, it results in a problem belonging to the same complexity class as the nominal problem and may thus prove computationally attractive when the semi-analytical schemes lift the problem to a higher complexity class (e.g., the robust counterpart of a linear program with ellipsoidal uncertainty set is a second-order cone program and thus more computationally expensive than the nominal problem). Moreover, the constraint sampling approach circumvents the need for estimating Ξ since it may directly employ historical samples. In addition, it remains applicable independently of the functional dependence of the objective and constraint functions on ξ and requires no structural assumptions on Ξ. Finally, it mitigates the sometimes criticized over-conservatism of mainstream robust optimization. Indeed, the scenario counterpart of a robust optimization problem constitutes a relaxation of (1.3) and thus yields super-optimal solutions that may nevertheless be infeasible for certain realizations of ξ. Calafiore and Campi [2005] and Campi and Garatti [2008] derive computationally attractive feasibility guarantees (as in chance-constrained programming) by demonstrating that $\mathcal{O}(n/\epsilon)$ samples are needed to guarantee that the solution to the scenario problem is feasible in the associated chance-constrained program. The link between the constraint sampling methodology and chance-constrained optimization is further showcased by Calafiore [2010] and Campi and Garatti [2011], who develop mechanisms for constraint discarding.

Insofar, we have discussed problems where the decision-maker takes *all* decisions *before* observing the realization of the uncertain parameters. These problems are typically referred to as *static* (or single-stage) problems in the sense that the decision-maker cannot revisit her actions once the uncertain parameters have been revealed. Thus, she may not exploit any information gained from observing the realization of ξ. These types of models are too restrictive, especially for decision problems extending well into the future (e.g., capacity expansion or financial planning, supply chain or risk management, etc). For this purpose, the stochastic programming

community has devised models capable of capturing the *dynamic* nature of the problems faced by practitioners.

The simplest form of dynamic decision problem is a *two-stage stochastic program*, see Kall and Wallace [1994], Prékopa [1995], Birge and Louveaux [2000] and Shapiro et al. [2009]. Here, the decision-maker takes a first-stage (here-and-now) decision x following which she observes all of the uncertain parameters. A second-stage or *recourse* (wait-and-see) decision y, which is allowed to depend on the realization of ξ, then comes to complement the original action. Mathematically, y is thus modeled as a function or *decision rule* of ξ. Consider for example a facility location problem, where a firm producing a single good has several clients with uncertain demand. The firm needs to decide whether and where to build new facilities (plants, warehouses, etc.) and how much to produce and stock in each location before observing the demand from the customers. Once the demand has been revealed, it may select the quantities of goods to distribute from each of the facilities that have been built to each of the customers. The goal of the firm is to maximize its expected profit (payoffs from sales less costs of building the facilities and distributing the goods). Here, the decisions to build each of the facilities correspond to first-stage actions, whereas the quantities to distribute from each facility to each customer can be viewed as recourse actions which can be adjusted to the realization of the customer demands. A two-stage stochastic program thus takes the form

$$\begin{array}{ll} \min\limits_{x\in\mathbb{R}^{n_1},\, y\in\mathcal{L}_{k,n_2}} & \mathbb{E}[g\,(x, y(\xi), \xi)] \\ \text{s.t.} & x \in \mathcal{X},\ y(\xi) \in \mathcal{Y}(x,\xi) \quad \mathbb{P}\text{-a.s.}, \end{array} \tag{1.5}$$

where $\mathcal{L}_{k,n_2}$ denotes, for example, the space of essentially bounded measurable functions from $\mathbb{R}^k$ to $\mathbb{R}^{n_2}$ and $\mathcal{Y}(x,\xi)$ corresponds to the set of feasible second-stage decisions.

Two-stage recourse problems find their natural extension in *multi-stage stochastic programs*, where uncertain parameters are revealed sequentially as time progresses. The decision-maker may then adjust his decisions whenever new information becomes available: mathematically, the decisions must thus be modeled as functions of the history of observations. The facility location problem discussed above, for example, must be viewed as a multi-stage problem if the

firm is allowed to periodically revisit its facility location decisions. In this case, the uncertain parameters correspond to the demands from the customers between consecutive revisions of the facility locations. The requirement that decisions only depend on the history of observations is referred to as *non-anticipativity* and ensures the *causality* of the decision-making process. While the incorporation of functional recourse decisions in model (1.5) is essential in order to keep the model flexible and realistic, it renders its solution (especially in the multi-stage setting) extremely challenging (Dyer and Stougie [2006], Shapiro and Nemirovski [2005]). Thus, researchers from the fields of dynamic and stochastic programming (Bertsekas [1995], Birge and Louveaux [2000]) as well as robust optimization and control (Ben-Tal et al. [2009], Dullerud and Paganini [2005]) have spent substantial research efforts on devising efficient algorithmic solution procedures.

If the distribution of the uncertain parameters is discrete, the underlying stochastic process and the associated decision process can both be modeled by means of a finite scenario tree whose nodes represent the possible realizations of ξ at each time stage. Here, the branching structure of the tree serves to model both the conditional distribution of the random process and the non-anticipative nature of the decision process (Shapiro et al. [2009]). The corresponding stochastic program can thus be solved by means of deterministic optimization, provided the number of support points of the distribution is not too large (see e.g., Shapiro and Nemirovski [2005]).

If the distribution of the uncertain parameters is continuous, a multi-stage stochastic program corresponds to an optimization problem over an infinite-dimensional function space incorporating infinitely many constraints. Thus, analytical approaches for solving problems of practical interest are rarely available and one must invariably resort to approximations combined with numerical optimization. Traditional approaches for solving such problems rely on a discretization of the underlying probability distribution and give rise to a scenario tree (Kall and Wallace [1994], Prékopa [1995], Birge and Louveaux [2000], Shapiro et al. [2009]). Most of these schemes are asymptotically *exact*, in the sense that any degree of accuracy can be achieved at the cost of increasing the number of discretization points (and thus the computational overhead). Nevertheless, for a given level of accuracy the number of discretization points needed grows *exponentially* with the number of decision stages (Shapiro and Nemirovski [2005]). Furthermore,

for small numbers of discretization points, the solution to the approximate problem may fail to be implementable in practice.

Recently, a novel solution methodology has emerged which circumvents the *curse of dimensionality* of traditional approaches by relying on a restriction of the feasible set of the problem rather than a discretization of the distribution. In the proposed scheme, pioneered by Ben-Tal et al. [2004], the decision rules are approximated by finite linear combinations of prescribed basis functions. Following this approximation, the coefficients of the linear combinations become the new decision variables of the problem. The *decision rule approximation* thus transforms the original dynamic problem to a static problem and consequently, all the machinery originally developed for single-stage problems may now be employed in a multi-stage context. We postpone a review of recent developments in the field to the introduction of Chapter 2 and only make the following observations:

1. Semi-analytical approaches for solving the problem arising from the decision rule approximation are usually only applicable when the uncertainty set and the employed decision rules have a simple algebraic structure. In addition, decision rules with a simple algebraic structure such as e.g., linear decision rules might result in overly crude approximations, i.e., in significant loss of optimality. Finally, most of these schemes necessitate knowledge or estimation of the distribution of the uncertain parameters, or at least of its support.

2. All existing decision rule approximation approaches are only applicable to problems where the decision-maker cannot influence the sequence in which the uncertain parameters are revealed. Nevertheless, in numerous real-world decision problems, the time of information discovery can be influenced by the controller who may decide whether and when to observe random parameters.

The main objective of this thesis is to devise tractable approximation schemes for solving multi-stage stochastic and robust optimization problems. More specifically, we aim:

1. To develop a unifying data-driven methodology for solving multi-stage robust optimization problems under generic nonlinear decision rules that yield arbitrarily tight approxi-

mations and to demonstrate the viability and flexibility of this approach on a large-scale real-world problem.

2. To extend the decision rule approximation paradigm to stochastic problems with decision-dependent information discovery.

1.2 Contributions and Structure of the Thesis

In this thesis, we develop solution schemes for dynamic decision-making under uncertainty. We investigate broad classes of stochastic and robust optimization problems and construct hierarchical approximations of tailored decision rules which enable their polynomial-time solution. The problems considered vary in their structure, in the required knowledge about the distribution of the uncertain parameters, and in the assumptions underlying the nature of the uncertainties. In order to illustrate the effectiveness of our approaches, we consider applications in numerous areas ranging from operations management to finance and energy engineering. Throughout the thesis, we rigorously assess both the convergence and scalability properties of our approximations and investigate their applicability in industrial-size instances.

Apart from the present chapter and the conclusions in Chapter 5, the remainder of the thesis is divided into three chapters, whose contributions can be summarized as follows.

In Chapter 2, we propose a unifying framework for solving convex, not necessarily linear, multi-stage robust optimization problems with general nonlinear dependence on the uncertain parameters. This is achieved by combining decision rule and constraint sampling approximation approaches. The synthesis of these two methodologies provides us with a versatile data-driven framework, which circumvents the need for estimating the distribution of the uncertain parameters or the uncertainty set and which remains applicable independently of their structure. Indeed, in the proposed approach, historical samples are employed directly in the optimization model. Furthermore, we introduce an axiomatic characterization of classes of decision rules that result in a tractable model and which guarantee asymptotic consistency. In fact, the use of constraint sampling techniques provides almost complete freedom in the choice of basis

functions for the decision rules. It thus enables us to exploit any a-priori knowledge about the structure of the optimal decision rule. Moreover, independently of the decision rule employed, the scenario counterpart of the static problem arising from the decision rule approximation will be of the same class as the original problem. The approach yields an algorithm parameterized in the number and type of basis functions considered and in the probability of constraint violation. These parameters serve to tune the trade-off between optimality and feasibility of the decision rules and the computational cost of the algorithm. The proposed scheme mitigates the sometimes criticized over-conservatism of mainstream robust optimization and can conveniently be tuned by the modeler to meet the risk-preferences of the decision-maker. The contents of this chapter are published in

1. P. Vayanos, D. Kuhn and B. Rustem. "A constraint sampling approach for multi-stage robust optimization". *Automatica*, 48(3):459–471, 2012.

In Chapter 3, we consider the problem of hedging and valuing path-dependent electricity derivatives such as *swing options*, which play a crucial risk-management role for utility companies in today's deregulated electricity markets. Unlike financial markets, electricity markets are *incomplete* which implies that the swing option cannot be assigned a unique fair price, but rather admits an interval of (no-arbitrage) prices consistent with those of traded instruments. We formulate two dynamic optimization problems whose optimal values yield the end-points of the no-arbitrage interval. Unfortunately, these problems, albeit accurate, involve a very large number of decision stages and decision variables and can thus not be solved by existing schemes. We propose a methodology which exploits the nature of the problem and the structure of the optimal decision rule to formulate approximate problems that can be solved efficiently with the constraint sampling approach discussed in Chapter 2. We illustrate how the violation probability parameter of the constraint sampling approximation can conveniently be used to model the negotiation process between the holder and writer of the option, who must agree on a price within the no-arbitrage interval. Our approach has natural applications to the valuation of mines, oil fields, power plants and refineries. This chapter thus develops an application of our research to problems of great interest for practitioners in energy markets and also serves to

illustrate how the modeling flexibility of the constraint sampling approach can be exploited to solve large-scale optimization problems. The contents of this chapter are partly published in

2. P. Vayanos, W. Wiesemann and D. Kuhn. "Hedging electricity swing options in incomplete markets". In *Proceedings of the 18th IFAC World Congress*, pages 846–853, 2011.

Up until the beginning of Chapter 3, the approaches discussed and investigated assume that the sequence in which the uncertain parameters are revealed is independent of the decision-maker's actions. In fact, this is the case for most of the literature on stochastic and robust optimization. Nevertheless, in numerous real-world decision problems, the time of information discovery can be influenced by the decision-maker, and uncertainties only become observable following an (often costly) investment. A classic example of this type of problem is that of infrastructure and production planning in offshore oilfields, which consist of several reservoirs of oil with uncertain volume. The goal is to decide the sequence in which to drill into each reservoir and the amount of oil to extract for production. The uncertain volume of each reservoir will only be revealed once the costly drilling process has been initiated. The drilling decisions thus control the time of information discovery in this problem. Much like in the case of traditional stochastic programming, solution approaches for solving such problems assume that the distribution of the uncertain parameters is discrete or rely on a discretization of the distribution of the uncertain parameters. This implies that they remain limited to only specific problem instances, that they yield solutions which may fail to be implementable in practice, or that they result in a combinatorial state explosion when applied to even medium sized problems. In Chapter 4, we propose a tractable, decision rule-based approximation scheme for solving such problems of great interest for practitioners. By employing the machinery of robust optimization, we develop the first scenario-free approach for solving problems with decision-dependent information discovery. The proposed scheme can be directly applied to problems with continuously distributed uncertain parameters without requiring a discretization of their distribution. Our approach conveniently captures the trade-off between sub-optimality and computational complexity in few approximation parameters. The contents of this chapter are published in

3. P. Vayanos, D. Kuhn and B. Rustem. "Decision rules for information discovery in multi-stage stochastic programming". In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pages 7368–7373, 2011.

1.3 Statement of Originality

The work presented in this thesis was conducted at Imperial College London between October 2008 and April 2012. Parts of this research are the result of collaborations with my supervisors, Prof. Berç Rustem and Dr. Daniel Kuhn, and my colleague, Dr. Wolfram Wiesemann. With this short statement, I certify that the material presented has not been submitted for the award of any degree in any other tertiary institution. This thesis is the result of my own work and no other person's work has been used without due acknowledgment.

Chapter 2

A Constraint Sampling Approach for Multi-Stage Robust Optimization

In this chapter we propose a tractable approximation scheme for convex multi-stage robust optimization problems. We approximate the adaptive decisions by finite linear combinations of prescribed basis functions and demonstrate how one can optimize over these decision rules at low computational cost through constraint randomization. We obtain a-priori probabilistic guarantees on the feasibility properties of the optimal decision rule by applying existing constraint sampling techniques to the semi-infinite problem arising from the decision rule approximation. We demonstrate that for a suitable choice of basis functions, the approximation converges as the size of the basis and the number of sampled constraints tend to infinity. The approach yields an algorithm parameterized in the basis size, the probability of constraint violation and the confidence that this probability will not be exceeded. These three parameters serve to tune the trade-off between optimality and feasibility of the decision rules and the computational cost of the algorithm. We assess the convergence and scalability properties of our approach in the context of two inventory management problems.

2.1 Introduction

Numerous real-world decision problems can be cast as *robust* dynamic optimization problems with a worst-case objective and robust constraints that must hold for all possible realizations of the uncertain problem parameters. The label dynamic indicates the presence of adaptive decisions that must be modeled as functions or decision rules of some (or all) of the uncertain parameters. One of the most recent methods to solve robust dynamic optimization problems of this type, pioneered by Ben-Tal et al. [2004], is to approximate these decision rules by linear combinations of prescribed basis functions. If, for instance, the basis functions are chosen to be the Euclidean coordinate projections, we obtain the popular class of *linear decision rules* (Ben-Tal et al. [2004, 2009]). Linear decision rules have been successfully applied in inventory management (Ben-Tal et al. [2005]), portfolio optimization (Calafiore [2008, 2009a], Rocha and Kuhn [2012]), network design (Atamtürk and Zhang [2007]), etc. The decision rule approximation transforms the original dynamic problem to a *static* robust optimization problem whose decision variables are the coefficients of the linear combinations.

While the original dynamic problem is typically intractable, the approximate static problem is sometimes equivalent to a linear, second-order conic or semidefinite program of moderate size and thus allows for an exact polynomial-time solution. This is the case, under mild structural assumptions, when the uncertainty set is characterized through conic inequalities and the employed decision rules are linear (Ben-Tal et al. [2004]), when the uncertainty set is rectangular and the decision rules are piecewise linear and separable (Ben-Tal et al. [2009], Georghiou et al. [2010]), or when the uncertainty set is ellipsoidal and the decision rules are quadratic (Ben-Tal et al. [2009]). Sometimes, however, the approximate static problem has no exact reformulation as a manifestly tractable conic program, and one has to resort to conservative approximations, yielding feasible yet sub-optimal solutions. This is the case, for instance, when the uncertainty set is semi-algebraic and the employed decision rules are polynomial (Bertsimas et al. [2011], Bampou and Kuhn [2011]), when the uncertainty set is polyhedral and the decision rules are piecewise linear (Goh and Sim [2010], Georghiou et al. [2010]), or when the uncertainty set is an intersection of concentric ellipsoids and the decision rules are quadratic (Ben-Tal et al.

[2009]). A summary of the semi-analytical decision rule approximation approaches discussed above is provided in Table 2.1.

Decision rule	Uncertainty set	Exact	References
Linear	Conic inequalities	Yes	Ben-Tal et al. [2004]
Piecewise linear and separable	Rectangular	Yes	Ben-Tal et al. [2009], Georghiou et al. [2010]
Quadratic	Ellipsoidal	Yes	Ben-Tal et al. [2009]
Polynomial	Semi-algebraic	No	Bertsimas et al. [2011], Bampou and Kuhn [2011]
Piecewise linear	Polyhedral	No	Goh and Sim [2010], Georghiou et al. [2010]
Quadratic	Intersection of concentric ellipsoids	No	Ben-Tal et al. [2009]

Table 2.1: Summary of the semi-analytical decision rule approximation approaches from the literature that admit an exact or conservative approximation as a tractable conic program.

When the uncertainty set and/or the employed decision rules have no simple algebraic structure, there are usually no semi-analytical conic programming approaches of the type described above. In this case, the static robust optimization problem emerging from the decision rule approximation is typically intractable. However, it can be solved approximately via constraint sampling, see Section 1.1.

Combining decision rule and constraint sampling techniques provides a flexible modeling framework for robust dynamic optimization problems with a general nonlinear dependence on the uncertain parameters. The sampling approach offers almost complete freedom in the choice of the basis functions for the decision rules. This allows the modeler to exploit any knowledge about the structure of the optimal solution when designing the decision rule approximation. The sampling approach is particularly attractive for problems that cannot be solved using semi-analytical schemes. However, its benefits also extend to certain problems for which semi-analytical approaches are adequate. This is the case, for instance, when the semi-analytical techniques lift the problem to a higher complexity class. Then, constraint sampling techniques prove attractive as they circumvent the added complexity, thereby yielding problems that can

potentially be solved with less computational effort. Moreover, they circumvent the need for estimating the support of the distribution of the uncertain parameters, since historical samples may be employed directly in the optimization. Finally, they bypass the sometimes criticized over-conservatism of semi-analytical approaches while providing attractive feasibility guarantees, see Campi and Garatti [2008]. There are nevertheless cases when combining decision rule and constraint sampling techniques is inadequate: this is the case in applications where even a small violation probability may be intolerable.

While the benefits of combining decision rule and sampling techniques have been explored by Calafiore and Nilim [2004], Bertsimas and Caramanis [2007], Skaf and Boyd [2009] and Lobel and Perakis [2010], several issues remain to be addressed. First, tractability results have only been provided for the case of polynomial decision rules. Furthermore, the interplay between the design parameters of the approximation has not been studied systematically. Finally, the asymptotic properties of the approximation have not been investigated rigorously.

The goal of this chapter is to present a unifying methodology for solving multi-stage convex robust optimization problems under generic nonlinear decision rules that yield arbitrarily tight approximations. In particular, we introduce an axiomatic characterization of classes of decision rules that result in a tractable scenario counterpart and guarantee asymptotic consistency. We model the degree of flexibility of the decision rules through a user specified complexity parameter that, for a fixed feasibility probability, controls the trade-off between sub-optimality and computational complexity.

The main contributions of this chapter are summarized below:

1. We consider general decision rule approximations for multi-stage robust optimization problems. The decision rules are modeled as finite linear combinations of a prescribed set of basis functions such as algebraic or trigonometric monomials, sigmoidal or Gaussian radial functions, etc. The flexibility to choose a generic set of basis functions is attractive since it enables the modeler to tailor the decision rule approximation to the problem instances if particular structural properties are known.

2. The synthesis of decision rule and constraint sampling techniques results in computationally tractable approximations of multi-stage problems. We illustrate how the underlying design parameters can be used, in a systematic fashion, to control the trade-off between sub-optimality and computational complexity of the approximation.

3. We provide a rigorous convergence proof for the decision rule approximation which applies, under mild technical assumptions, if the number of basis functions is driven to infinity. We also demonstrate almost sure convergence of the constraint randomization approach as the number of samples tends to infinity.

This chapter is organized as follows. The remainder of this section introduces the notation, while Section 2.2 presents the mathematical problem formulation. Section 2.3 is split into two parts concerned with the decision rule and constraint sampling approximations, respectively, and Section 2.4 provides a complexity analysis. The convergence properties of the approximation scheme are investigated in Section 2.5. Finally, Section 2.6 presents two problems from inventory management which are amenable to our approximation scheme, and Section 2.7 reports on numerical results.

Notation Throughout this chapter, vectors (matrices) are denoted by lowercase (uppercase) letters. For $x \in \mathbb{R}^n$, we denote the closed n-ball of radius r centered at x by $B_r(x)$, the Euclidean norm of x by $|x|$ and the i^{th} component of x by x_i. Also, for $\alpha \in \mathbb{N}_0^n$, we let $x^\alpha := \prod_{i=1}^n (x_i)^{\alpha_i}$, where $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. For any matrix $X \in \mathbb{R}^{n \times m}$ with columns $x_i \in \mathbb{R}^n$, $i = 1, \ldots, m$, we denote by $\text{vec}(X) := (x_1, x_2, \ldots, x_m) \in \mathbb{R}^{nm}$ the vector concatenation of its columns.

Uncertainty is modeled by the probability space $(\mathbb{R}^k, \mathcal{F}, \mathbb{P})$, which consists of the sample space $\mathbb{R}^k$, the Borel σ-algebra $\mathcal{F} := \mathcal{B}(\mathbb{R}^k)$ and the probability measure $\mathbb{P}$, whose support we denote by Ξ. The elements of the sample space are denoted by ξ and are assumed to possess a temporal structure in that they are representable as $\xi := (\xi_1, \xi_2, \ldots, \xi_T)$. The random vectors $\xi_t \in \mathbb{R}^{k_t}$, $t \in \mathbb{T} := \{1, 2, \ldots, T\}$, have marginal supports $\Xi_t \subseteq \mathbb{R}^{k_t}$, where $\sum_{t \in \mathbb{T}} k_t = k$. For convenience, we define combined random vectors $\xi^t := (\xi_1, \ldots, \xi_t) \in \mathbb{R}^{k^t}$, $t \in \mathbb{T}$, with marginal supports

$\Xi^t \subseteq \mathbb{R}^{k^t}$, where $k^t := \sum_{\tau=1}^{t} k_\tau$. Furthermore, for each $t \in \mathbb{T}$, we introduce the information set

$$\mathcal{F}^t := \left\{ Z^t \times \left(\bigtimes_{\tau=t+1}^{T} \mathbb{R}^{k_\tau} \right) : Z^t \in \mathcal{B}\left(\mathbb{R}^{k^t}\right) \right\}$$

which corresponds to the σ-algebra generated by ξ^t, the history of the random vector ξ up to time t. Finally, for each $t \in \mathbb{T}$ we denote by $\mathcal{C}(\Xi^t)$ the space of continuous real-valued functions on Ξ^t.

2.2 Problem Description

We consider a multi-stage decision problem under uncertainty over the finite planning horizon $\mathbb{T}$. The aim is to find a sequence of decisions $x := (x_1, x_2, \ldots, x_T)$ that minimizes a cost function $f_0(x, \xi)$ in the worst-case realization of $\xi \in \Xi$. These decisions are constrained by a set of inequalities which are required to be obeyed robustly, that is, for any realization of the uncertainties $\xi \in \Xi$. The decision $x_t \in \mathbb{R}^{n_t}$ is selected at time t after the history of realizations ξ^t has been observed but before the future outcomes $\{\xi_f\}_{f>t}$ have been revealed. This motivates us to represent x_t as an $\mathcal{F}^t$-measurable function or *decision rule* of ξ. The requirement that x_t be constant in $\{\xi_f\}_{f>t}$, which is implied by the $\mathcal{F}^t$-measurability, reflects the *non-anticipative* nature of the decision process and ensures the *causality* of the dynamic decision problem. Modeling x_t as an $\mathcal{F}^t$-measurable function of ξ is essential to keep the decision model realistic and flexible but makes it computationally challenging. Robust optimization problems of the type described here can be formulated compactly as

$$\begin{array}{ll} \inf\limits_{x \in \mathcal{N}} & \sup\limits_{\xi \in \Xi} \; f_0\left(x\left(\xi\right), \xi\right) \\ \text{s.t.} & f_i\left(x\left(\xi\right), \xi\right) \leq 0 \quad \forall \xi \in \Xi, \; i = 1, \ldots, I, \end{array} \tag{$\mathcal{R}$}$$

where $f_i : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}$, $i = 0, 1, \ldots, I$, are convex in x and continuous in (x, ξ), while

$$\mathcal{N} := \bigtimes_{t \in \mathbb{T}} L_{n_t}^{\infty}\left(\mathbb{R}^k, \mathcal{F}^t, \mathbb{P}\right)$$

denotes the space of all non-anticipative bounded decision rules from $\mathbb{R}^k$ to $\mathbb{R}^n$, $n := \sum_{t\in\mathbb{T}} n_t$. Problem $\mathcal{R}$ provides an attractive alternative to its stochastic counterpart with an expectation objective since it yields solutions that are better suited to the needs of more risk averse decision-makers.

Remark 2.1 (Role of the distribution $\mathbb{P}$) *We note at this point that the probability distribution $\mathbb{P}$ only influences problem $\mathcal{R}$ through its support Ξ. Indeed, the worst-case objective and the robust constraints are independent of probabilities. Although the distribution $\mathbb{P}$ does not play a role per se in problem $\mathcal{R}$, it will, however, assume a bivalent position when the constraint sampling approximation is introduced. Firstly, it constitutes the distribution used for sampling. Secondly, it provides the basis for establishing probabilistic guarantees on the solution quality. We remark that our approach remains applicable even if $\mathbb{P}$ is unknown. In this case, we merely require that a finite (yet "sufficiently large") number of samples from $\mathbb{P}$, e.g., in the form of past observations, be available.*

In what follows, we adopt an approach which, although aimed at approximating the robust problem $\mathcal{R}$, mitigates the sometimes criticized over-conservatism of robust optimization. In particular, we take the view of a decision-maker who may accept to be unprotected against events that have a low probability of occurrence. This idea is well known from chance-constrained programming and is often conceptually preferred to the "hard" robust paradigm since the choice in the immunization level may be tailored to the risk tolerance of the decision-maker. Indeed, sacrificing some robustness can sometimes result in a significant gain in optimality. As will become clear later on, our approach returns solutions that are, with high confidence, feasible for the following multi-stage joint chance-constrained program with probability level ϵ.

$$\begin{array}{ll} \inf\limits_{\theta\in\mathbb{R},\, x\in\mathcal{N}} & \theta \\ \text{s.t.} & \mathbb{P}\left(f_0(x(\xi),\xi)\leq\theta, f_1(x(\xi),\xi)\leq 0,\ldots,f_I(x(\xi),\xi)\leq 0\right)\geq 1-\epsilon. \end{array} \qquad (\mathcal{CC}_\epsilon)$$

Note, however, that the optimal solution of the (typically) non-convex problem $\mathcal{CC}_\epsilon$ cannot be computed efficiently even if the distribution $\mathbb{P}$ is fully known, which is not a prerequisite for our approach.

As discussed in Section 1.1, multi-stage problems of the form $\mathcal{R}$ are very difficult to solve. In what follows, we discuss two successive parametric approximations that enable us to find near-optimal solutions for $\mathcal{R}$, while controlling the quality and complexity of the approximation.

2.3 Tractable Reformulation

In this section, we construct a computationally tractable approximation for problem $\mathcal{R}$ based on a flexible decision rule and constraint sampling approach.

2.3.1 Decision Rule Approximation

The decision variables in the original robust problem $\mathcal{R}$ range over a space of non-anticipative measurable functions from $\mathbb{R}^k$ to $\mathbb{R}^n$. As a first step towards obtaining a computationally tractable model, we restrict the space of all measurable decision rules to those representable as finite linear combinations of certain prescribed basis functions. We first describe how to construct vectors of basis functions of a given complexity. We further provide conditions ensuring that the flexibility of the decision rules increases with their complexity and that the size of the basis vectors is polynomially bounded in the dimension k of the random vectors ξ. We then use the basis vectors to construct approximations for problem $\mathcal{R}$ which accommodate only a finite number of decision variables.

The construction of the basis vectors proceeds in two steps:

1. For each $t \in \mathbb{T}$, select a sequence of continuous functions $b_{t,m} : \mathbb{R}^{k^t} \to \mathbb{R}$, $m \in \mathbb{N}$.

2. Select a complexity parameter $d \in \mathbb{N}_0$. For any fixed d, choose an increasing function $s_d : \mathbb{N}_0 \to \mathbb{N}$ and define the basis vector $b_t^d : \mathbb{R}^{k^t} \to \mathbb{R}^{s_d(k^t)}$ as $b_t^d := (b_{t,1}, b_{t,2}, \ldots, b_{t,s_d(k^t)})$ for all $t \in \mathbb{T}$. We will refer to $s_d(k^t)$ as the basis size at time t.

Throughout the rest of the chapter, we assume the following conditions to hold true.

(C1) The linear hull of $\{b_{t,m}\}_{m\in\mathbb{N}}$ is dense in $\mathcal{C}(\Xi^t)$ with respect to the supremum norm.

(C2) For any fixed $d \in \mathbb{N}_0$, the function s_d is bounded above by a polynomial.

Once the basis vectors have been constructed as described above, we restrict the functional decisions x_t in $\mathcal{R}$ to be representable as linear combinations of the components of b_t^d for each $t \in \mathbb{T}$. Thus, we focus on decision rules of the form

$$x_t(\xi) = X_t\, b_t^d(\xi^t), \tag{2.1}$$

where the matrix $X_t \in \mathbb{R}^{n_t \times s_d(k^t)}$ contains the coefficients of the linear combinations, which become the new decision variables. We note that modeling b_t^d as a function of the observation history $\xi^t = (\xi_1, \ldots, \xi_t)$ up to time t preserves the non-anticipative nature of the decision rules. With the restriction (2.1), the epigraph formulation of problem $\mathcal{R}$ reduces to

$$\begin{array}{ll} \min\limits_{\theta\in\mathbb{R},\, y\in\mathbb{R}^{n_y}} & \theta \\ \text{s.t.} & \tilde{f}_0(y,\xi) - \theta \le 0 \quad \forall \xi \in \Xi \\ & \tilde{f}_i(y,\xi) \le 0 \quad \forall \xi \in \Xi, \quad i = 1, \ldots, I, \end{array}$$

where we identify y with $(\mathrm{vec}(X_1), \ldots, \mathrm{vec}(X_T))$, set $n_y := \sum_{t\in\mathbb{T}} n_t s_d(k^t)$ and define

$$\tilde{f}_i(y,\xi) := f_i\left(\left(X_1\, b_1^d\left(\xi^1\right), \ldots, X_T\, b_T^d\left(\xi^T\right)\right), \xi\right) \quad i = 0, \ldots, I.$$

The functions $\tilde{f}_i : \mathbb{R}^{n_y} \times \mathbb{R}^k \to \mathbb{R}$ are convex in y and continuous in (y, ξ). For notational convenience, we set $w := (\theta, y)$ and introduce the following equivalent robust optimization problem with basis functions

$$\begin{array}{ll} \min\limits_{w\in\mathbb{R}^{n_w}} & w_1 \\ \text{s.t.} & f(w,\xi) \le 0 \quad \forall \xi \in \Xi, \end{array} \tag{$\mathcal{RB}$}$$

where $n_w := 1 + n_y$ and $f : \mathbb{R}^{n_w} \times \mathbb{R}^k \to \mathbb{R}$ is defined through

$$f(w, \xi) := \max \left\{ \max_{1 \le i \le I} \tilde{f}_i(y, \xi), \tilde{f}_0(y, \xi) - \theta \right\}.$$

Several observations are now in order. The function f is convex in w and continuous in (w, ξ) as it is obtained by taking the maximum of $I + 1$ functions with these properties. The dimension n_w of the decision vector in $\mathcal{RB}$ is linear in the number n of functional decision variables in the original problem and polynomial in the number k of uncertainties.

Thus far, we have derived an approximation for problem $\mathcal{R}$ with a finite number of decision variables that remains polynomially bounded in n and k. In order to illustrate the flexibility of this approach, we now provide several examples of decision rules and associated basis vectors. These examples highlight the advantages of non-affine decision rules, which can result in arbitrarily tight approximations for the original problem $\mathcal{R}$ (see Section 2.5.1). An illustration of their approximation capabilities is provided in Figure 2.1.

Example 2.1 (Algebraic polynomials) *A convenient choice for b_t^d is the vector of monomials in the $\mathbb{R}$-vector space of multivariate polynomials in ξ^t with degree at most d, that is,*

$$b_t^d(\xi^t) := \left\{ \left(\zeta^t(\xi^t)\right)^\alpha : \alpha \in \mathbb{N}_0^{k^t}, \|\alpha\|_1 \le d \right\}.$$

The affine scaling function $\zeta^t : \mathbb{R}^{k^t} \to \mathbb{R}^{k^t}$ is defined through

$$\zeta^t(\xi^t) := \left(2\frac{\xi_1 - l_1}{u_1 - l_1} - 1, \ldots, 2\frac{\xi_{k^t} - l_{k^t}}{u_{k^t} - l_{k^t}} - 1 \right), \tag{2.2}$$

where $[l_s, u_s] \subseteq \mathbb{R}$ denotes an interval covering the marginal support of ξ_s. The scaling ensures that $b_t^d(\Xi^t) \subseteq [-1, 1]^{s_d(k^t)}$, which in turn implies that the decision rules are uniformly bounded on Ξ^t. This helps to avoid numerical instabilities due to poorly scaled optimization problems. By the Stone-Weierstraß theorem, the algebraic polynomials (of arbitrary finite degree) are dense in $\mathcal{C}(\Xi^t)$. Therefore, condition (C1) is satisfied. The dimension of b_t^d is given by $s_d(k^t) = \binom{k^t + d}{d} \in \mathcal{O}((k^t)^d)$, which is polynomial in k^t, as required by condition (C2). The complexity parameter

d corresponds to the degree of the polynomial decision rules.

Example 2.2 (Trigonometric polynomials) *The vector of monomials in the $\mathbb{R}$-vector space of multivariate trigonometric polynomials in ξ^t with degree at most d can be written as*

$$b_t^d(\xi^t) := \left\{\cos c^\top \zeta^t(\xi^t) : c \in \mathbb{Z}^{k^t},\ \|c\|_1 \leq d\right\} \bigcup \left\{\sin s^\top \zeta^t(\xi^t) : s \in \mathbb{Z}^{k^t},\ \|s\|_1 \leq d,\ s \neq 0\right\},$$

where the scaling function ζ^t is defined as in (2.2). Here, the scaling is needed due to the periodicity of the trigonometric polynomials. By the Stone-Weierstraß theorem, the multivariate trigonometric polynomials (of arbitrary finite degree) are dense in $\mathcal{C}(\Xi^t)$.

The total number of vectors $c \in \mathbb{N}_0^{k^t}$ with exactly $i \leq \min(d, k^t)$ non-zero elements satisfying $\|c\|_1 = j$, $i \leq j \leq d$, is $\binom{k^t}{i}\binom{j-1}{j-i}$. For each vector $c \in \mathbb{Z}^{k^t}$ with i non-zero elements, there are 2^i sign combinations which give rise to distinct cosine-type basis functions. A similar argument can be made for the sine-type basis functions. Thus, the basis size at time t is given by

$$s_d(k^t) = 1 + \sum_{i=1}^{\min(d,k^t)} \sum_{j=i}^{d} 2^{i+1} \binom{k^t}{i} \binom{j-1}{j-i} \in \mathcal{O}\left((k^t)^d\right),$$

which is polynomially bounded in k^t for any fixed d. The complexity parameter d corresponds to the degree of the trigonometric polynomial decision rules.

Example 2.3 (Sigmoidal basis functions) *Let $\sigma : \mathbb{R} \to \mathbb{R}$ be any continuous sigmoidal function such as $\sigma(t) = 1/(1+\exp(-t))$ or $\sigma(t) = \max\{0, \min\{(t+1)/2, 1\}\}$, and let $\gamma : \mathbb{Z}^{k^t+1} \to \mathbb{Q}^{k^t+1}$ be a bijection (see e.g Yu-Ting [1980] for an example of such a bijection). We can now define a basis vector b_t^d of complexity d through*

$$b_t^d\left(\xi^t\right) := \left\{\sigma\left(c^\top \zeta^t\left(\xi^t\right) + g\right) :\ (c, g) = \gamma(m),\ m \in \mathbb{Z}^{k^t+1},\ \|m\|_1 \leq d\right\},$$

where the scaling function ζ^t is defined as in (2.2). The scaling ensures that the basis functions corresponding to low values of d have significant variability within Ξ^t. By Cybenko's theorem

(Theorem 1 in Cybenko [1989]), the linear hull of all sigmoidal basis functions (for arbitrary values of d) is dense in $\mathcal{C}(\Xi^t)$.

Following a similar reasoning as in Example 2.2, it can be shown that

$$s_d\left(k^t\right) = 1 + \sum_{i=1}^{\min(d,k^t+1)} \sum_{j=i}^{d} 2^i \binom{k^t+1}{i} \binom{j-1}{j-i} \in \mathcal{O}\left((k^t)^d\right).$$

The basis size is therefore polynomially bounded in k^t *for any fixed* d. *The complexity parameter* d *corresponds to the maximum Manhattan norm of the vector* m *encoding the rationals.*

Example 2.4 (Gaussian radial basis functions) *Following a similar approach as in Example 2.3, we can construct a basis vector* b_t^d *of complexity* d *by setting*

$$b_t^d\left(\xi^t\right) := \left\{ \exp\left(-g^2 |\zeta^t\left(\xi^t\right) - c|^2\right) : \ (c,g) = \gamma\left(m\right), \ m \in \mathbb{Z}^{k^t+1}, \ \|m\|_1 \leq d \right\},$$

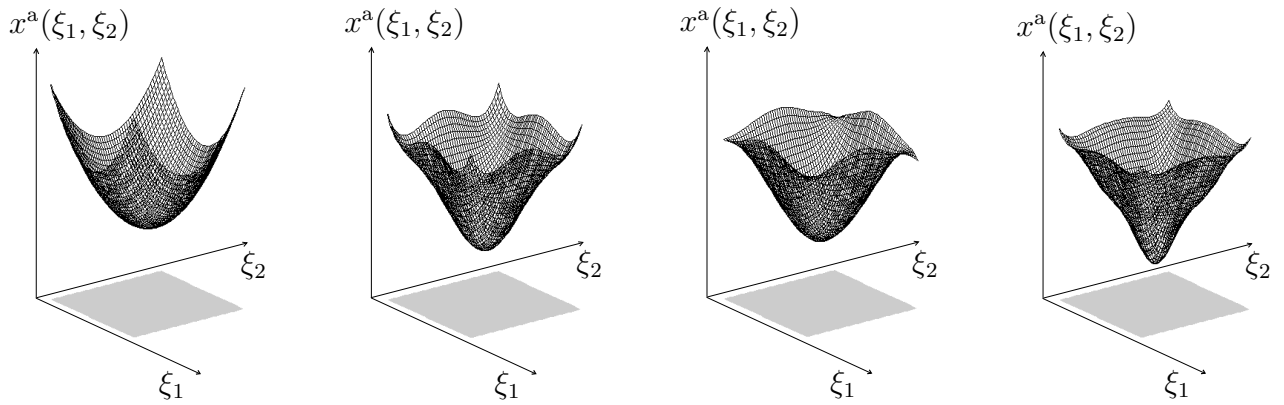
where the scaling function ζ^t *is as in* (2.2) *and the bijection* γ *is as in Example 2.3. The linear hull of all Gaussian radial basis functions (for arbitrary values of* d*) is dense in* $\mathcal{C}(\Xi^t)$, *see e.g., Proposition B.1 in Girosi and Poggio [1990].*

Following a similar approach as in Example 2.3, we obtain

$$s_d\left(k^t\right) = 1 + \sum_{i=1}^{\min(d,k^t+1)} \sum_{j=i}^{d} 2^i \binom{k^t+1}{i} \binom{j-1}{j-i} \in \mathcal{O}\left((k^t)^d\right).$$

Thus, the basis size is polynomially bounded in k^t *for any fixed* d. *As in Example 2.3, the complexity parameter* d *corresponds to the maximum Manhattan norm of the vector* m *which encodes the rationals.*

We note that the number of decision variables in problem $\mathcal{RB}$ can always be reduced by limiting the memory of the decision rules, that is, by representing x_t as a combination of functions that only depend on $(\xi_{t-m}, \ldots, \xi_t)$ for some $m \in \mathbb{N}$. Furthermore, a different basis vector may be used for each decision variable, thereby leading to a great deal of modeling flexibility.



(a) Polynomial (b) Trigonometric (c) Sigmoidal (d) Gaussian radial

Figure 2.1: Approximation of $x(\xi_1, \xi_2) := \max(\xi_1, \xi_2) + \max(-\xi_1, -\xi_2) + \max(\xi_1, -\xi_2) + \max(-\xi_1, \xi_2)$ on $[-3, 3] \times [-3, 3]$ by (a) polynomial (b) trigonometric polynomial (c) sigmoidal and (d) Gaussian radial basis functions. In all cases, the complexity parameter is set to $d = 3$.

So far, we described a methodology for reducing the number of decision variables of problem $\mathcal{R}$. However, we note that when Ξ has infinite cardinality (as is the case if $\mathbb{P}$ is continuous), the resulting problem $\mathcal{RB}$ falls in the category of semi-infinite programs, which are extremely hard to solve in general, see e.g., Blondel and Tsitsiklis [2000]. In particular, it is known that checking feasibility of a generic semi-infinite constraint $f(w, \xi) \leq 0 \; \forall \xi \in \Xi$ for a fixed $w \in \mathbb{R}^{n_w}$ is NP-hard even if Ξ is a simplex and f is indefinite quadratic in ξ. In the next section, we therefore apply a constraint sampling approach to problem $\mathcal{RB}$ to achieve computational tractability.

2.3.2 Constraint Sampling Approximation

Calafiore and Campi [2005] suggest to solve single-stage robust optimization problems that would otherwise be intractable due to their semi-infinite nature by considering their "scenario" counterparts: a finite set of N constraints is chosen at random from the typically uncountable set of constraints, and the resulting tractable problem is solved. The main result of Calafiore and Campi [2005], which was later improved by Campi and Garatti [2008], is the observation that any solution of the sampled problem will also satisfy most of the constraints of the original problem (which were not sampled).

Here we apply this approach to the decision rule approximation $\mathcal{RB}$, which has in fact the struc-

ture of a single-stage robust optimization problem. Indeed, drawing N independent samples $\xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(N)}$ distributed according to $\mathbb{P}$, we can approximate $\mathcal{RB}$ by the following scenario problem with basis functions.

$$\begin{array}{ll} \min\limits_{w \in \mathbb{R}^{n_w}} & w_1 \\ \text{s.t.} & f(w, \xi^{(s)}) \leq 0, \quad s = 1, \ldots, N. \end{array} \tag{$\mathcal{SB}^N$}$$

Since the function f is convex in w for each $\xi \in \Xi$, $\mathcal{SB}^N$ is again a convex optimization problem. We observe that since the samples $\xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(N)}$ are random variables, problem $\mathcal{SB}^N$ is itself random. However, the hope is that the variability of the solution set and the optimal value over different samples of size N is small if N is chosen sufficiently large (see Section 2.5.2).

Recall that problem $\mathcal{R}$ has infinitely many variables and constraints parameterized by $\xi \in \Xi$. The decision rule approximation described in Section 2.3.1 transformed $\mathcal{R}$ to the semi-infinite problem $\mathcal{RB}$ with finitely many variables and infinitely many constraints. Constraint sampling, in turn, yields problem $\mathcal{SB}^N$, which has finitely many variables *and* constraints. We note that all three problems are convex. While $\mathcal{RB}$ provides an upper bound on $\mathcal{R}$ due to a restriction of the feasible set, the sampled problem $\mathcal{SB}^N$ constitutes a relaxation of problem $\mathcal{RB}$, and thus its solution provides a lower bound on $\mathcal{RB}$.

A fundamental question now arising is whether these approximations break the curse of dimensionality (see e.g., Section 13.1.3 in Ben-Tal et al. [2009]) so that scalability to multi-stage problems is actually possible. In the following section, we thus investigate the trade-off between the accuracy and the tractability of the two approximations.

2.4 Complexity Analysis

As the constraints of the scenario problem $\mathcal{SB}^N$ are randomly extracted, its optimal solutions are random variables which depend on the set of extractions $\{\xi^{(s)}\}_{s=1}^N$. These solutions typically fail to satisfy all constraints of the robust problem $\mathcal{RB}$. However, they can be shown to satisfy the constraints of $\mathcal{RB}$ with high probability. The following fundamental question is addressed,

among others, by Calafiore and Campi [2005, 2006] and by Campi and Garatti [2008]: what confidence do we have that an optimal solution of $\mathcal{SB}^N$ will violate the constraints of $\mathcal{RB}$ with probability less than ϵ, where ϵ is a prescribed probability level as in chance-constrained programming (see Section 2.2)? To the best of our knowledge, the most general answer to this question is provided by Campi and Garatti [2008] who establish, for any solvable convex optimization problem and any probability distribution $\mathbb{P}$ with support Ξ, the so-called "exact feasibility" of the randomized solution. We repeat their results here, together with the definition of violation probability as introduced in Calafiore and Campi [2005].

Definition 2.1 (Calafiore and Campi [2005]) *The violation probability of a given $w \in \mathbb{R}^{n_w}$ is defined as $V(w) := \mathbb{P}(\xi \in \Xi : f(w, \xi) > 0)$.*

Theorem 2.1 (Campi and Garatti [2008]) *For any given probability level $\epsilon \in (0, 1)$ and confidence level $\beta \in (0, 1)$, let*

$$N(\epsilon, \beta) := \min \left\{ N \in \mathbb{N} : \sum_{i=0}^{n_w - 1} \binom{N}{i} \epsilon^i (1 - \epsilon)^{N-i} \leq \beta \right\}.$$

Suppose that $\mathcal{RB}$ is solvable. If problem $\mathcal{SB}^N$ is solvable and $N \geq N(\epsilon, \beta)$, then we have

$$\mathbb{P}^N \left(V(w_N^*) > \epsilon \right) \leq \beta,$$

where w_N^* denotes the (w.l.o.g.) unique optimal solution of $\mathcal{SB}^N$, and $\mathbb{P}^N := \mathbb{P} \times \mathbb{P} \times \cdots \times \mathbb{P}$ (N times) is the probability distribution of the sample $(\xi^{(1)}, \xi^{(2)}, \ldots, \xi^{(N)})$.

Remark 2.2 *Theorem 2.1 remains applicable in the case when problem $\mathcal{SB}^N$ has multiple optimal solutions, provided a suitable "tie-breaking" rule is used to systematically select a single optimal solution, see e.g., Section 4.1 in Calafiore and Campi [2005] or Section 2.1 in Campi and Garatti [2008].*

The result from Theorem 2.1 guarantees that any random solution of $\mathcal{SB}^N$ satisfies most of the original constraints with high confidence provided that the sample size N is chosen large

enough. Thus, the solution of $\mathcal{SB}^N$ is feasible in $\mathcal{CC}_\epsilon$ with high confidence $1-\beta$. We emphasize that the distribution $\mathbb{P}$ is not needed to compute this solution.

We note that by continuity of f, $\mathcal{RB}$ is always solvable when its feasible set is nonempty and bounded. A bound on the violation probability that remains valid for possibly infeasible instances of $\mathcal{RB}$ is provided by Calafiore [2010].

The following corollary is an immediate consequence of Theorem 2.1 and generalizes a result of Bertsimas and Caramanis [2007] for polynomial decision rules.

Corollary 2.1 (Complexity of problem $\mathcal{SB}^N$) *For any fixed probability level $\epsilon \in (0,1)$ and confidence level $\beta \in (0,1)$, the number of samples N needed such that the optimal solution w_N^* of $\mathcal{SB}^N$ satisfies $\mathbb{P}^N(V(w_N^*) > \epsilon) \leq \beta$ remains polynomially bounded in n and k.*

Remark 2.3 (Computational tractability) *Corollary 2.1 implies that problem $\mathcal{SB}^N$ can be solved in polynomial time with respect to the size of the input parameters, provided that for any fixed $\xi \in \mathbb{R}^k$, the set $\{w \in \mathbb{R}^n : f(w,\xi) \leq 0\}$ admits an efficient separation oracle, see Grötschel et al. [1981].*

Proof of Corollary 2.1 In Calafiore [2009b] it was shown that

$$N(\epsilon,\beta) \leq \frac{2}{\epsilon}\left(\ln\frac{1}{\beta} + n_w\right).$$

By construction of problem $\mathcal{RB}$ (see Section 2.3.1), we have that

$$n_w = 1 + \sum_{t\in\mathbb{T}} n_t s_d\left(k^t\right) \leq 1 + \sum_{t\in\mathbb{T}} n_t s_d\left(k\right) = 1 + n s_d\left(k\right),$$

where the inequality holds since s_d is increasing and $k \geq k_t$ for each $t \in \mathbb{T}$. As s_d is polynomially bounded by condition (C2), the number of decision variables n_w, and hence the number of samples $N(\epsilon,\beta)$ required for the prescribed violation probability ϵ and confidence level β, remain polynomially bounded in n and k. □

Remark 2.4 (Scalability) *Typically, n and k are both linear in T, in which case the number of samples N needed to ensure $\mathbb{P}^N(V(w_N^*) > \epsilon) \leq \beta$ is also polynomially bounded with T.*

Remark 2.5 *The basis size $s_d(k^t)$ can be exponential in the complexity parameter d for fixed values of k^t. In such cases, the number of samples required to sustain a maximum violation probability of ϵ at confidence $1-\beta$ is exponential in d. If $\mathbb{P}$ is unknown and only $\tilde{N}$ samples from $\mathbb{P}$ are available, one may opt for "simpler" decision rules, i.e., low values of d, in order to guarantee the required level of feasibility. Indeed, the maximum admissible value of d for given values of ϵ, β and $\tilde{N}$ amounts to*

$$\overline{d} := \max\left\{ d \in \mathbb{N} : \min\left\{ N \in \mathbb{N} : \sum_{i=0}^{n_w(d)-1} \binom{N}{i} \epsilon^i (1-\epsilon)^{N-i} \leq \beta \right\} \leq \tilde{N} \right\},$$

where $n_w(d) := 1 + \sum_{t\in\mathbb{T}} n_t s_d(k^t)$. In particular, any d with $s_d(k) \leq n^{-1}(\epsilon\tilde{N}/2 - \ln(1/\beta) - 1)$ satisfies $d \leq \overline{d}$ and ensures that no more than $\tilde{N}$ samples are needed to sustain a maximum violation of ϵ with confidence $1-\beta$.

For high values of d and low values of ϵ, problem $\mathcal{SB}^N$ can become computationally expensive or memory intensive. The following algorithm, which relies on the observation that only a small portion of the sampled constraints are active at optimality, provides an iterative solution approach that substantially reduces the sizes of the problems to be solved. It is specifically designed to solve $\mathcal{SB}^N$ efficiently while keeping the memory requirements of each iteration low, as it only involves the solution of subproblems with significantly fewer constraints. For ease of exposition, we assume that $\mathcal{RB}$ is feasible.

Algorithm 2.1 (Iterative constraint addition/removal procedure)

1. ***Initialization.*** *Partition the set* $\{1, \ldots, N\}$ *into* $n_{\rm it}$ *subsets, letting* $\mathcal{I}_i$ *denote the* i^{th} *subset. Set the iteration counters to* $i \leftarrow 1$ *and* $j \leftarrow 1$*. Also set* $\mathcal{I} \leftarrow \emptyset$ *and* $u \leftarrow$ *false.*

2. ***Optimization.*** *Set* $\mathcal{I} \leftarrow \mathcal{I} \cup \mathcal{I}_i$*. Solve the relaxation of* $\mathcal{SB}^N$ *involving only the constraints* $f(w, \xi^{(s)}) \leq 0$*,* $s \in \mathcal{I}$*. If the problem is solvable, denote the solution by* w_N^i *and go to Step 3, else if* $i < n_{\rm it}$*, set* $i \leftarrow i + 1$ *and repeat Step 2, else set* $u \leftarrow$ *true and go to Step 6.*

3. ***Constraint removal.*** *Set* $\mathcal{I} \leftarrow \mathcal{I} \backslash \{s \in \mathcal{I} : f(w_N^i, \xi^{(s)}) < 0\}$*. If* $i = n_{\rm it}$*, set* $\hat{w}_N^1 \leftarrow w_N^i$ *and go to Step 4, else set* $i \leftarrow i + 1$ *and go to Step 2.*

4. ***Constraint addition.*** *If* $\hat{w}_N^j$ *is feasible in* $\mathcal{SB}^N$*, go to Step 6, else set* $\mathcal{I} \leftarrow \mathcal{I} \cup \{s \in \{1, \ldots, N\} : f(\hat{w}_N^j, \xi^{(s)}) > 0\}$*.*

5. ***Re-optimization.*** *Set* $j \leftarrow j + 1$ *and solve the relaxation of* $\mathcal{SB}^N$ *involving only the constraints* $f(w, \xi^{(s)}) \leq 0$*,* $s \in \mathcal{I}$*. Denote the resulting solution by* $\hat{w}_N^j$*. Go to Step 4.*

6. ***Termination.*** *If* $u =$ *false, set* $w_N^* \leftarrow \hat{w}_N^j$ *and stop, else declare* $\mathcal{SB}^N$ *to be unbounded.*

More sophisticated algorithms for solving $\mathcal{SB}^N$ that remain applicable when $\mathcal{RB}$ is infeasible can be designed by adapting known semi-infinite optimization algorithms.

2.5 Convergence Analysis

The first part of this section analyzes the convergence of the optimal value of $\mathcal{RB}$ to the optimal value of $\mathcal{R}$ as the complexity parameter d of the basis vector tends to infinity. The second part establishes the almost sure convergence of the optimal value of the scenario problem $\mathcal{SB}^N$ to the optimal value of $\mathcal{RB}$ as the number of samples N increases. Before introducing a few technical assumptions, we define the concept of strict feasibility.

Definition 2.2 (Strict feasibility) *A decision rule $x : \mathbb{R}^k \to \mathbb{R}^n$ is said to be strictly feasible for the robust optimization problem $\mathcal{R}$ if $x \in \mathcal{N}$ and there exists $\delta > 0$ with*

$$f_i\left(x\left(\xi\right), \xi\right) \leq -\delta \quad \forall \xi \in \Xi,\ i = 1, \ldots, I.$$

If there exists a strictly feasible decision rule for $\mathcal{R}$, then the problem is said to be strictly feasible.

The subsequent convergence results are based on the following mild assumptions.

(A1) The robust problem $\mathcal{R}$ is strictly feasible.

(A2) The functions $f_i : \mathbb{R}^n \times \mathbb{R}^k \to \mathbb{R}$ are convex in x for each $\xi \in \mathbb{R}^k$ and continuous in (x, ξ), $i = 0, \ldots, I$.

(A3) There exists $R \in \mathbb{R}$ such that $|x(\xi)| \leq R$ for each $\xi \in \Xi$ and for each x feasible in $\mathcal{R}$.

(A4) An arbitrary number of independent samples from $\mathbb{P}$ can be obtained.

(A5) The uncertainty set Ξ is convex, compact, fully dimensional and rectangular in the sense that $\Xi = \bigtimes_{t \in \mathbb{T}} \Xi_t$.

Several comments are in order. Firstly, assumption (A1) is satisfied for all problems of practical interest if equality constraints are systematically eliminated by using them to reduce the number of decision variables. We remark that $\mathcal{RB}$ can be infeasible even though $\mathcal{R}$ is strictly feasible. Moreover, assumption (A3) is not restrictive since we are ultimately interested in numerical solutions for problem $\mathcal{R}$ which are necessarily bounded. The condition that Ξ be fully dimensional (meaning that there exist $\xi_0 \in \Xi$ and $\epsilon > 0$ such that $B_\epsilon(\xi_0) \subseteq \Xi$) is also non-restrictive. It can always be enforced at the cost of reducing the dimension of ξ if necessary. Finally, we note that assumption (A5) could be relaxed to require the existence of a non-anticipative homeomorphism $g : \mathbb{R}^k \to \mathbb{R}^k$, such that $\Xi = g(\Xi')$ for some set $\Xi' \subseteq \mathbb{R}^k$ that satisfies (A5). This generalization would allow us to consider even many non-convex uncertainty sets. For notational simplicity, however, we will use the simpler assumption (A5) in

the sequel. The assumptions (A1)-(A5) are always assumed to hold in the remainder of the chapter.

Throughout what follows, we denote the infimum of a problem $\mathcal{P}$ by $\inf \mathcal{P}$, the closed n-ball $B_R(0)$ by X and the Lebesgue measure on $\mathbb{R}^k$ by μ.

2.5.1 Convergence of the Decision Rule Approximation

The main result of this section is provided in Theorem 2.2. Before embarking on its proof, we provide definitions and technical background results which will facilitate our exposition.

Definition 2.3 (Mollifier) *A continuous function $\phi : \mathbb{R}^s \to \mathbb{R}$ is called an s-dimensional mollifier if*

(a) $\phi(z) \geq 0$,

(b) $\phi(z) = 0$ *if* $|z| \geq 1$ *and*

(c) $\int_{\mathbb{R}^s} \phi(z)\,\mathrm{d}z = 1$.

Let $\phi : \mathbb{R} \to \mathbb{R}$ be a unimodal one-dimensional mollifier for which there exist $\zeta \in (0,1)$ and $\kappa > 0$ such that $\phi(z) \geq \kappa$ for all z with $|z| \leq \zeta$. Let $C := \max_{z\in\mathbb{R}} \phi(z)$ and define the k-dimensional mollifiers $\phi_m : \mathbb{R}^k \to \mathbb{R}$ corresponding to ϕ through

$$\phi_m(z) := \prod_{i=1}^{k} m\phi(mz_i), \ m \in \mathbb{N}. \tag{2.3}$$

Lemma 2.1 *The mollifier ϕ_m satisfies*

$$\int_{\Xi} \phi_m(\xi - z)\,\mathrm{d}z > 0 \quad \forall \xi \in \Xi.$$

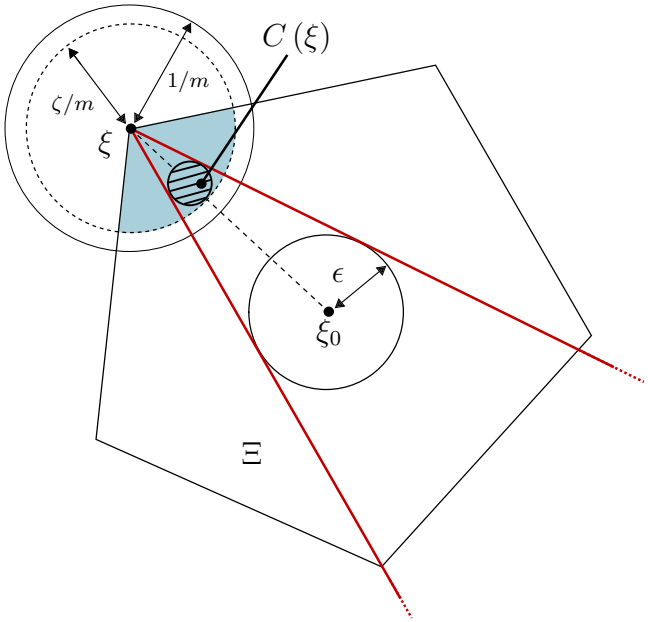


Figure 2.2: Companion figure for the proof of Lemma 2.1. The measure of $C(\xi)$ (striped area) provides a lower bound for the measure of $\Xi \cap B_{\zeta/m}(\xi)$ (shaded area).

Proof Select any $\xi \in \Xi$. By our assumption on ϕ, we have $\phi_m(z) \geq (m\kappa)^k$ for all z with $|z| \leq \zeta/m$, which implies that

$$\int_{\Xi} \phi_m(\xi - z)\,\mathrm{d}z \;\geq (m\kappa)^k\, \mu\left(\Xi \cap B_{\zeta/m}(\xi)\right). \tag{2.4}$$

Assumption (A5) guarantees that there exists $\xi_0 \in \Xi$ and $\epsilon > 0$ such that $B_\epsilon(\xi_0) \subseteq \Xi$. Next, we introduce the set

$$C(\xi) \;:=\; \lambda B_\epsilon(\xi_0) + (1-\lambda)\xi \;=\; B_{\lambda\epsilon}(\lambda\xi_0 + (1-\lambda)\xi),$$

where

$$\lambda := \min\left\{\frac{\zeta}{m(|\xi - \xi_0| + \epsilon)}, 1\right\},$$

see Figure 2.2. In the following, we show that $C(\xi) \subseteq \Xi \cap B_{\zeta/m}(\xi)$. Firstly, the set $C(\xi)$ was constructed as a convex combination of two sets contained in Ξ, and therefore $C(\xi) \subseteq \Xi$.

Secondly, note that

$$\begin{aligned} C(\xi) \subseteq B_{\zeta/m}(\xi) &\Leftrightarrow |\xi' - \xi| \leq \frac{\zeta}{m} \quad \forall \xi' \in C(\xi) \\ &\Leftrightarrow |\lambda(\xi' - \xi)| \leq \frac{\zeta}{m} \quad \forall \xi' \in B_\epsilon(\xi_0) \\ &\Leftrightarrow \begin{cases} |\xi' - \xi| \leq \frac{\zeta}{m} \quad \forall \xi' \in B_\epsilon(\xi_0) & \text{if } \lambda = 1, \\ |\xi' - \xi| \leq |\xi - \xi_0| + \epsilon \quad \forall \xi' \in B_\epsilon(\xi_0) & \text{else.} \end{cases} \end{aligned}$$

Both cases above are trivially satisfied. Therefore, $C(\xi) \subseteq B_{\zeta/m}(\xi)$. In conclusion, we have $C(\xi) \subseteq \Xi \cap B_{\zeta/m}(\xi)$ and consequently $\mu(\Xi \cap B_{\zeta/m}(\xi)) \geq \mu(C(\xi)) > 0$ since $C(\xi)$ is a k-ball of strictly positive radius. The claim now follows from (2.4). □

Lemma 2.2 *Let x be feasible for problem $\mathcal{R}$ and define the function $x^m : \mathbb{R}^k \to \mathbb{R}^n$ through*

$$x^m(\xi) := \frac{\int_\Xi x(z)\,\phi_m(\xi - z)\,\mathrm{d}z}{\int_\Xi \phi_m(\xi - z)\,\mathrm{d}z}. \tag{2.5}$$

Then,

(a) x^m is continuous on Ξ, and

(b) $x^m \in \mathcal{N}$.

Proof (a) Select $\xi' \in \Xi$. Then, we have

$$\lim_{\xi \to \xi'} \int_\Xi x(z)\,\phi_m(\xi - z)\,\mathrm{d}z \;=\; \int_\Xi \lim_{\xi \to \xi'} x(z)\,\phi_m(\xi - z)\,\mathrm{d}z \;=\; \int_\Xi x(z)\,\phi_m(\xi' - z)\,\mathrm{d}z.$$

The interchange of the limit and the integral is justified by the dominated convergence theorem, which applies because $|x(z)\,\phi_m(\xi - z)| \leq m^k C^k R$ uniformly for all $\xi, z \in \Xi$, see assumption (A3). Thus, the numerator of x^m in (2.5) is continuous on Ξ. Following similar arguments, it is possible to prove that the denominator of x^m is also continuous on Ξ. Thus, x^m is continuous on Ξ as it is a ratio of two continuous functions whose denominator is strictly positive, see Lemma 2.1.

(b) For each $t \in \mathbb{T}$, we introduce a new sequence of mollifiers $\phi_{m,t} : \mathbb{R}^{k_t} \to \mathbb{R}$, $m \in \mathbb{N}$, which are defined in terms of ϕ through

$$\phi_{m,t}(z_t) := \prod_{j=1}^{k_t} m\phi(mz_{t,j}).$$

Using this definition, ϕ_m can be written as

$$\phi_m(z) = \prod_{t=1}^{T} \phi_{m,t}(z_t), \ m \in \mathbb{N}. \tag{2.6}$$

Select an arbitrary $t \in \mathbb{T}$. As x_t is $\mathcal{F}^t$-measurable, there exists a function $\chi_t : \mathbb{R}^{k^t} \to \mathbb{R}^{n_t}$ such that $x_t(\xi) = \chi_t(\xi^t)$ for all $\xi \in \mathbb{R}^k$. Thus, we have

$$\begin{aligned} x_t^m(\xi) &= \frac{\int_\Xi x_t(z)\,\phi_m(\xi - z)\,\mathrm{d}z}{\int_\Xi \phi_m(\xi - z)\,\mathrm{d}z} \\ &= \frac{\int_{\Xi_1} \cdots \int_{\Xi_t} \chi_t(z^t) \prod_{\tau=1}^{t} \phi_{m,\tau}(\xi_\tau - z_\tau)\,\mathrm{d}z_1 \cdots \mathrm{d}z_t}{\int_{\Xi_1} \cdots \int_{\Xi_t} \prod_{\tau=1}^{t} \phi_{m,\tau}(\xi_\tau - z_\tau)\,\mathrm{d}z_1 \cdots \mathrm{d}z_t}, \end{aligned}$$

where the second equality holds due to (2.6) and since $\Xi = \bigtimes_{\tau=1}^{T} \Xi_\tau$, see assumption (A5). The last expression for x_t^m is independent of $\xi_{t+1}, \ldots, \xi_T$. Therefore x_t^m is $\mathcal{F}^t$-measurable. As the choice of $t \in \mathbb{T}$ was arbitrary, the decision rule x^m is non-anticipative. Part (a) further implies that x^m is continuous and bounded on the compact set Ξ. Therefore $x^m \in \mathcal{N}$ as postulated. □

Theorem 2.2 (Convergence of the decision rule approximation) *The infimum of the robust problem $\mathcal{RB}$ with basis functions converges to the infimum of the original problem $\mathcal{R}$ as the complexity parameter d tends to infinity.*

Proof The general strategy of our proof is to choose an arbitrary tolerance level $\epsilon > 0$ and to construct a sequence of four ϵ-optimal decision rules feasible in $\mathcal{R}$ with increasingly regular structure. The four decision rules will be (a) measurable, (b) measurable and strictly feasible, (c) continuous and strictly feasible and (d) representable as a finite linear combination of basis functions.

Select $\epsilon > 0$. Since problem $\mathcal{R}$ is feasible, there exists a feasible decision rule $x^{(1)} \in \mathcal{N}$ with

$$|\sup_{\xi\in\Xi} f_0(x^{(1)}(\xi),\xi) - \inf \mathcal{R}| \leq \frac{\epsilon}{4}. \tag{2.7}$$

Similarly, since the robust optimization problem $\mathcal{R}$ is strictly feasible, there exists a strictly feasible decision rule $x^{(2)} \in \mathcal{N}$ and $\delta^{(2)} > 0$ such that

$$f_i(x^{(2)}(\xi),\xi) \leq -\delta^{(2)} \ \ \forall \xi \in \Xi,\ i = 1,\ldots,I.$$

Set $x^\lambda := (1-\lambda)x^{(1)} + \lambda x^{(2)}$ for $\lambda \in [0,1]$. Since the functions f_i are convex in x, we have

$$\begin{aligned} f_i(x^\lambda(\xi),\xi) &\leq (1-\lambda)\, f_i(x^{(1)}(\xi),\xi) + \lambda f_i(x^{(2)}(\xi),\xi) \\ &\leq -\lambda\delta^{(2)}\ \forall \xi \in \Xi,\ i = 1,\ldots,I. \end{aligned}$$

Therefore, the causal decision rule $x^\lambda \in \mathcal{N}$ is strictly feasible in $\mathcal{R}$ for all $\lambda > 0$. The function f_0 in the objective of $\mathcal{R}$ is uniformly continuous on the compact set $X \times \Xi$. Thus, there exists $\eta > 0$ such that

$$|f_0(x,\xi) - f_0(x',\xi)| \leq \frac{\epsilon}{4} \ \ \forall x,\ x' \in X \text{ with } |x - x'| \leq \eta \text{ and } \xi \in \Xi. \tag{2.8}$$

Since $x^{(1)}$ and $x^{(2)}$ are both feasible for $\mathcal{R}$, assumption (A3) implies

$$|x^\lambda(\xi) - x^{(1)}(\xi)| = |\lambda x^{(2)}(\xi) - \lambda x^{(1)}(\xi)| \leq 2\lambda R \ \ \forall \xi \in \Xi. \tag{2.9}$$

Setting $\lambda_0 := \frac{\eta}{2R}$ and $x^{(3)} := x^{\lambda_0} \in \mathcal{N}$, we conclude via (2.8) and (2.9) that

$$|f_0(x^{(3)}(\xi),\xi) - f_0(x^{(1)}(\xi),\xi)| \leq \frac{\epsilon}{4} \ \ \forall \xi \in \Xi,$$

and therefore

$$|\sup_{\xi\in\Xi} f_0(x^{(3)}(\xi),\xi) - \sup_{\xi\in\Xi} f_0(x^{(1)}(\xi),\xi)| \leq \frac{\epsilon}{4}. \tag{2.10}$$

We note that our choice of λ_0 implies

$$f_i(x^{(3)}(\xi),\xi) \leq -\delta^{(3)} \quad \forall \xi \in \Xi,\ i = 1,\ldots,I, \tag{2.11}$$

where $\delta^{(3)} := \eta\delta^{(2)}/2R$. Thus, $x^{(3)}$ constitutes a strictly feasible, near optimal decision rule.

Now, let the mollifier sequence $\{\phi_m\}_{m\in\mathbb{N}}$ be as in (2.3) and define the mollified decision rules $\{x^m\}_{m\in\mathbb{N}}$ through

$$x^m(\xi) = \frac{\int_\Xi x^{(3)}(z)\phi_m(\xi - z)\mathrm{d}z}{\int_\Xi \phi_m(\xi - z)\mathrm{d}z}.$$

Lemma 2.2 implies that each x^m is non-anticipative and continuous on Ξ. Since the objective and constraint functions f_i, $i = 0,\ldots,I$ are convex in x, we can apply Jensen's inequality to obtain

$$f_i(x^m(\xi),\xi) \quad \leq \quad \int_\Xi f_i(x^{(3)}(z),\xi)\frac{\phi_m(\xi - z)}{\int_\Xi \phi_m(\xi - y)\mathrm{d}y}\mathrm{d}z. \tag{2.12}$$

Next, set $\delta^{(4)} := \delta^{(3)}/2$. By the uniform continuity of f_i on $X \times \Xi$, there exists $m_0 \in \mathbb{N}$ such that

$$\begin{gathered} |f_i(x^{(3)}(z),z) - f_i(x^{(3)}(z),\xi)| \leq \min\left\{\delta^{(4)}, \tfrac{\epsilon}{4}\right\} \\ \forall \xi,\ z \in \Xi \text{ with } |\xi - z| \leq \tfrac{1}{m_0},\ i = 0,\ldots,I. \end{gathered} \tag{2.13}$$

For $i = 1,\ldots,I$ (constraint functions) and $m \geq m_0$, we thus have

$$\begin{aligned} f_i(x^m(\xi),\xi) \quad &\leq \quad \int_\Xi \frac{\left(f_i(x^{(3)}(z),z) + \delta^{(4)}\right)\phi_m\left(\xi - z\right)}{\int_\Xi \phi_m\left(\xi - y\right)\mathrm{d}y}\mathrm{d}z \\ &\leq \quad \delta^{(4)} - \delta^{(3)} = -\delta^{(4)} \quad \forall \xi \in \Xi, \end{aligned}$$

where the first inequality follows from (2.12) and (2.13), while the second inequality holds because of (2.11). For $i = 0$ (objective function) and $m \geq m_0$, we can make a similar argument to obtain

$$\begin{aligned} f_0(x^m(\xi),\xi) \quad &\leq \quad \int_\Xi f_0(x^{(3)}(z),z)\frac{\phi_m(\xi - z)}{\int_\Xi \phi_m(\xi - y)\mathrm{d}y}\mathrm{d}z + \frac{\epsilon}{4} \\ &\leq \quad \sup_{z\in\Xi} f_0(x^{(3)}(z),z) + \frac{\epsilon}{4} \quad \forall \xi \in \Xi. \end{aligned} \tag{2.14}$$

Next, define $x^{(4)} := x^{m_0}$. By construction, $x^{(4)}$ is non-anticipative and continuous on Ξ. Moreover, it is near optimal and strictly feasible in $\mathcal{R}$.

By the uniform continuity of f_i on $X \times \Xi$, there exists $\eta > 0$ such that

$$|f_i(x,\xi) - f_i(x',\xi)| \leq \min\left\{\frac{\epsilon}{4}, \delta^{(4)}\right\} \ \forall x,\ x' \in X \text{ with } |x - x'| \leq \eta,\ \xi \in \Xi,\ i = 0,\ldots,I. \quad (2.15)$$

By construction of the basis vectors b_t^d and since $x^{(4)}$ is continuous on Ξ, there exists a complexity parameter $d \in \mathbb{N}$ and matrices $X_t \in \mathbb{R}^{n_t \times s_d(k^t)}$, $t \in \mathbb{T}$, such that

$$|x^{(5)}(\xi) - x^{(4)}(\xi)| \leq \eta \ \ \forall \xi \in \Xi,$$

where $x^{(5)} \in \mathcal{N}$ is defined through $x_t^{(5)}(\xi) := X_t\, b_t^d(\xi^t)$, see condition (C1). The estimate (2.15) and the strict feasibility of $x^{(4)}$ imply that

$$|\sup_{\xi\in\Xi} f_0(x^{(5)}(\xi),\xi) - \sup_{\xi\in\Xi} f_0(x^{(4)}(\xi),\xi)| \leq \frac{\epsilon}{4} \quad (2.16)$$

and

$$f_i(x^{(5)}(\xi),\xi) \leq 0 \ \ \forall \xi \in \Xi,\ i = 1,\ldots,I.$$

We have thus shown that $x^{(5)}$ is feasible for $\mathcal{R}$, and we also obtain

$$\begin{aligned} 0 \ \leq\ \sup_{\xi\in\Xi} f_0(x^{(5)}(\xi),\xi) - \inf\mathcal{R} \ &=\ \sup_{\xi\in\Xi}\ f_0(x^{(1)}(\xi),\xi) - \inf\mathcal{R} \\ &+\ \sup_{\xi\in\Xi}\ f_0(x^{(3)}(\xi),\xi) - \sup_{\xi\in\Xi}\ f_0(x^{(1)}(\xi),\xi) \\ &+\ \sup_{\xi\in\Xi}\ f_0(x^{(4)}(\xi),\xi) - \sup_{\xi\in\Xi}\ f_0(x^{(3)}(\xi),\xi) \\ &+\ \sup_{\xi\in\Xi} f_0(x^{(5)}(\xi),\xi) - \sup_{\xi\in\Xi} f_0(x^{(4)}(\xi),\xi) \leq \epsilon, \end{aligned}$$

where the first inequality follows from the feasibility of $x^{(5)}$ in $\mathcal{R}$, and the last inequality follows from (2.7), (2.10), (2.14) and (2.16). We thus conclude that the optimal value of problem $\mathcal{RB}$ differs at most by ϵ from the optimal value of $\mathcal{R}$. As the choice of ϵ was arbitrary, the claim follows. $\square$

Remark 2.6 *We presented a convergence proof for the decision rule approximation introduced in Section 2.3.1. This establishes, as special cases, convergence of the polynomial, trigonometric polynomial, Gaussian radial and sigmoidal decision rule approximations, see Examples 2.1-2.4. We note that for the exact decision rule approximations from the literature there are currently no convergence proofs. In addition, we emphasize that for the linear, quadratic or separable decision rule approximations convergence cannot be achieved.*

2.5.2 Convergence of the Sampling Approximation

In this section, we demonstrate that the optimal value of $\mathcal{SB}^N$ converges with probability 1 (w.p.1) to the optimal value of $\mathcal{RB}$ as the number of samples tends to infinity whenever problem $\mathcal{RB}$ is feasible. Thus, we will henceforth assume that $\mathcal{RB}$ is indeed feasible. Note that this can always be enforced by choosing the complexity parameter d large enough, see Theorem 2.2. The general strategy of our proof is as follows. We will first consider a fixed realization of the stochastic process $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ that is dense in Ξ. For this particular sequence of samples, we will prove that $\inf \mathcal{SB}^N$ converges to $\inf \mathcal{RB}$ as N tends to infinity. To prove the main result, we will then argue that $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ w.p.1.

For the further argumentation, we introduce the function $\psi : \mathbb{R}^{n_w} \to \mathbb{R}$, $\psi(w) := \max_{\xi\in\Xi} f(w,\xi)$ and the corresponding random function $\psi_N : \mathbb{R}^{n_w} \to \mathbb{R}$, $\psi_N(w) := \max_{\xi\in\Xi_N} f(w,\xi)$, where the random set $\Xi_N \subseteq \mathbb{R}^k$ is defined through $\Xi_N := \{\xi^{(1)}, \ldots, \xi^{(N)}\}$, $N \in \mathbb{N}$.

Remark 2.7 *By construction, the feasible set of problem $\mathcal{RB}$ is given by $\{w \in \mathbb{R}^{n_w} : \psi(w) \leq 0\}$, while the feasible set of $\mathcal{SB}^N$ is representable as $\{w \in \mathbb{R}^{n_w} : \psi_N(w) \leq 0\}$.*

Lemma 2.3 *If $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ, then ψ_N converges continuously to ψ on $\mathbb{R}^{n_w}$.*

Proof We first demonstrate that ψ_N converges pointwise to ψ on $\mathbb{R}^{n_w}$. By the continuity of f and the compactness of Ξ, $\psi(w)$ exists and is finite for each $w \in \mathbb{R}^{n_w}$. Also, $\{\psi_N(w)\}_{N\in\mathbb{N}}$ is a

non-decreasing sequence bounded above by $\psi(w)$. We thus conclude that

$$\lim_{N\to\infty} \psi_N(w) = \sup_{N\in\mathbb{N}} \max_{\xi\in\Xi_N} f(w,\xi) = \max_{\xi\in\Xi} f(w,\xi) = \psi(w),$$

where the second equality holds since f is continuous and $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ. Thus, ψ_N converges pointwise to ψ on $\mathbb{R}^{n_w}$.

We now prove the main result. Choose $w \in \mathbb{R}^{n_w}$ and $\epsilon > 0$. Since f is continuous and Ξ is compact, there exists $\delta > 0$ such that

$$|f(w,\xi) - f(w',\xi)| \leq \epsilon \;\; \forall w' \in B_\delta(w),\; \xi \in \Xi.$$

Thus, we find

$$\begin{aligned} |\psi_N(w) - \psi_N(w')| &= |\max_{\xi\in\Xi_N} f(w,\xi) - \max_{\xi\in\Xi_N} f(w',\xi)| \\ &\leq \epsilon \;\; \forall N \in \mathbb{N},\; \forall w' \in B_\delta(w). \end{aligned} \tag{2.17}$$

The sequence $\{\psi_N\}_{N\in\mathbb{N}}$ is locally bounded at w. As the choice of ϵ was arbitrary, we thus conclude from (2.17) that the sequence is equicontinuous at w, see e.g., pp.248–249 in Rockafellar and Wets [1997]. The above argument holds for all $w \in \mathbb{R}^{n_w}$, and thus $\{\psi_N\}_{N\in\mathbb{N}}$ is equicontinuous on $\mathbb{R}^{n_w}$. The pointwise convergence of ψ_N to ψ on $\mathbb{R}^{n_w}$ and Theorem 7.10 in Rockafellar and Wets [1997] ensure that ψ_N epiconverges to ψ, which implies via Theorem 7.11 in Rockafellar and Wets [1997] that ψ_N converges continuously to ψ. □

Lemma 2.4 *If $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ, then there exist $N_0 \in \mathbb{N}$ and a non-empty compact set $W^* \subseteq \mathbb{R}^{n_w}$ such that W^* contains at least one optimal solution of problem $\mathcal{SB}^N$ for each $N \geq N_0$.*

Remark 2.8 *We emphasize that W^* and N_0 may depend on the particular realization of $\{\xi^{(l)}\}_{l\in\mathbb{N}}$. We further note that Lemma 2.4 implicitly guarantees that $\mathcal{SB}^N$ is solvable (i.e., the minimum is attained) and has a finite optimal value for all $N \geq N_0$.*

Proof of Lemma 2.4 See Appendix 2.8.1. □

Lemma 2.5 *If $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ, then the infimum of the scenario problem $\mathcal{SB}^N$ converges to the infimum of $\mathcal{RB}$ as the number of samples N tends to infinity.*

The following proof is inspired by the proof of Proposition 2.2 in Pagnoncelli et al. [2009], where almost sure convergence of sample average approximations to chance-constrained programs is demonstrated.

Proof of Lemma 2.5 Since $\mathcal{SB}^N$ is a relaxation of $\mathcal{RB}$, we have

$$\limsup_{N\to\infty} \inf \mathcal{SB}^N \;\leq\; \inf \mathcal{RB}. \tag{2.18}$$

For all $N \geq N_0$, let $w_N^* \in W^*$ denote an optimal solution to $\mathcal{SB}^N$, whose existence is guaranteed by Lemma 2.4. Consider now a subsequence $\{w_{N_j}^*\}_{j\in\mathbb{N}}$ with the property that $\liminf_{N\to\infty} \mathrm{e}_1^\top w_N^* = \lim_{j\to\infty} \mathrm{e}_1^\top w_{N_j}^*$, where e_1 denotes the first canonical basis vector in $\mathbb{R}^{n_w}$. Since any w_N^* is an element of the compact set W^*, we may assume without loss of generality that $\lim_{j\to\infty} w_{N_j}^* = w^*$ for some $w^* \in W^*$. The continuous convergence of ψ_N to ψ, which is ensured by Lemma 2.3, then implies

$$\lim_{j\to\infty} \psi_N(\tilde{w}_{N_j}^*) \;=\; \psi(w^*).$$

As $\psi_{N_j}(w_{N_j}^*) \leq 0$ for all $j \in \mathbb{N}$, we conclude that $\psi(w^*) \leq 0$, that is, w^* is feasible in $\mathcal{RB}$, and thus $\mathrm{e}_1^\top w^* \geq \inf \mathcal{RB}$. Hence, we find

$$\liminf_{N\to\infty} \inf \mathcal{SB}^N = \liminf_{N\to\infty} \mathrm{e}_1^\top w_N^* = \lim_{j\to\infty} \mathrm{e}_1^\top w_{N_j}^* = \mathrm{e}_1^\top w^* \geq \inf \mathcal{RB}. \tag{2.19}$$

It follows from (2.18) and (2.19) that $\inf \mathcal{SB}^N$ converges to $\inf \mathcal{RB}$. □

Theorem 2.3 (Convergence of the sampling approximation) *The infimum of the scenario problem $\mathcal{SB}^N$ converges w.p.1 to the infimum of $\mathcal{RB}$ as the number of samples N tends to infinity.*

Proof It can be shown that $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ w.p.1, see Appendix 2.8.2. Thus, Theorem 2.3 is a simple corollary of Lemma 2.5. □

2.6 Examples from Inventory Management

We assess the convergence and scalability properties of our approach in the context of two single-echelon, multi-period supply chain models from the literature.

2.6.1 Inventory Management with Cumulative Order Constraints

The first problem considered here was originally proposed by Ben-Tal et al. [2005]. We discuss a simplified version due to Bertsimas et al. [2011], which we denote by $\mathcal{R}_{\mathrm{COC}}$.

At the beginning of each period $t \in \mathbb{T}$, a retailer receives orders for $\xi_{\mathrm{d},t}$ units of a product. This demand needs to be satisfied from the on-hand inventory, whose filling level is denoted by $s_{\mathrm{inv},t}$. In order to replenish the inventory, the retailer may place orders $x_{\mathrm{o},t}$ with a supplier, thereby incurring shipping costs d_t per unit of the product. Unsatisfied demand may be backlogged at cost p_t and inventory may be held on the premises at cost h_t per unit of the product. Furthermore, there are prescribed limits on the orders placed at each period as well as on the cumulative orders $s_{\mathrm{co},t}$ placed up to period t.

The dynamics of the inventory level and the cumulative orders are governed by

$$\begin{aligned} s_{\mathrm{inv},t+1}(\xi^{t+1}) &= s_{\mathrm{inv},t}(\xi^t) + x_{\mathrm{o},t}(\xi^t) - \xi_{\mathrm{d},t+1} \\ s_{\mathrm{co},t+1}(\xi^{t+1}) &= s_{\mathrm{co},t}(\xi^t) + x_{\mathrm{o},t}(\xi^t) \end{aligned}$$

for $t = 1, \ldots, T-1$, where $s_{\mathrm{inv},1}$ denotes the initial inventory and $s_{\mathrm{co},1}$ the initial cumulative

order level. We impose the box constraints

$$\underline{x}_{\mathrm{o},t} \;\leq\; x_{\mathrm{o},t}(\xi^t) \;\leq\; \overline{x}_{\mathrm{o},t},\; t = 1,\ldots,T-1 \tag{2.20a}$$

$$\underline{s}_{\mathrm{co},t} \;\leq\; s_{\mathrm{co},t}(\xi^t) \;\leq\; \overline{s}_{\mathrm{co},t},\; t = 1,\ldots,T, \tag{2.20b}$$

where $\underline{x}_{\mathrm{o},t}$, $\overline{x}_{\mathrm{o},t}$ and $\underline{s}_{\mathrm{co},t}$, $\overline{s}_{\mathrm{co},t}$ denote the lower and upper bounds on the instantaneous and cumulative orders, respectively.

We assume that there is no demand at time $t = 1$. The future demands are independent and uniformly distributed as $\xi_{\mathrm{d},t} \sim \mathcal{U}\left(\overline{\xi}_{\mathrm{d},t}\,(1-\rho_{\mathrm{d}})\,, \overline{\xi}_{\mathrm{d},t}\,(1+\rho_{\mathrm{d}})\right)$, where $\overline{\xi}_{\mathrm{d},t}$ denotes the nominal demand in period t, and $\rho_{\mathrm{d}} \in [0,1]$ quantifies the degree of uncertainty. The support of ξ thus corresponds to a box uncertainty set of the form $\Xi = \bigtimes_{t\in\mathbb{T}}\{\xi_{\mathrm{d},t} \in \mathbb{R} : |\xi_{\mathrm{d},t} - \overline{\xi}_{\mathrm{d},t}| \leq \rho_{\mathrm{d}}\overline{\xi}_{\mathrm{d},t}\}$.

The retailer's objective is to minimize the worst-case cumulative cost $\max_{\xi\in\Xi} \sum_{t\in\mathbb{T}} x_{\mathrm{c},t}(\xi^t)$, where the stage-wise costs $x_{\mathrm{c},t}$ satisfy

$$x_{\mathrm{c},t}(\xi^t) \geq d_t x_{\mathrm{o},t}(\xi^t) + \max\left\{h_t s_{\mathrm{inv},t}(\xi^t), -p_t s_{\mathrm{inv},t}(\xi^t)\right\}$$

for $t = 1,\ldots,T-1$ and

$$x_{\mathrm{c},T}(\xi^T) \geq \max\left\{h_T s_{\mathrm{inv},T}(\xi^T), -p_T s_{\mathrm{inv},T}(\xi^T)\right\}.$$

2.6.2 Inventory Management with Random Yield

As an extension to the basic model from Section 2.6.1, we assume now that the quality of the shipments received from the supplier is uncertain in the sense that only a fraction $\xi_{\mathrm{y},t}$ of the total number of products ordered in period t is of satisfactory quality. We denote this modified problem by $\mathcal{R}_{\mathrm{RY}}$. We assume that the yields $\xi_{\mathrm{y},t}$ are mutually independent and uniformly distributed as $\xi_{\mathrm{y},t} \sim \mathcal{U}(\overline{\xi}_{\mathrm{y},t}(1-\rho_{\mathrm{y}}), \overline{\xi}_{\mathrm{y},t}(1+\rho_{\mathrm{y}}))$, with $[\overline{\xi}_{\mathrm{y},t}(1-\rho_{\mathrm{y}}), \overline{\xi}_{\mathrm{y},t}(1+\rho_{\mathrm{y}})] \subseteq [0,1]$. We further assume that the retailer only pays for the shipments of sufficient quality and that the payment for orders placed at time t is made at time $t+1$, *after* the shipment quality has

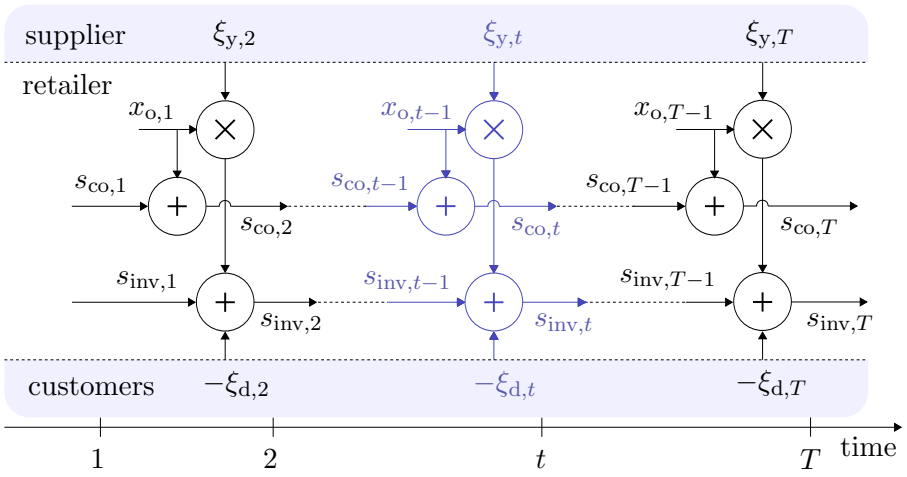


Figure 2.3: A supply chain with T stages, uncertain customer demand and yield.

been observed. These amendments to the basic model result in the following modified system dynamics

$$\begin{aligned} s_{\mathrm{inv},t+1}(\xi^{t+1}) &= s_{\mathrm{inv},t}(\xi^t) + \xi_{\mathrm{y},t+1} x_{\mathrm{o},t}(\xi^t) - \xi_{\mathrm{d},t+1} \\ s_{\mathrm{co},t+1}(\xi^{t+1}) &= s_{\mathrm{co},t}(\xi^t) + x_{\mathrm{o},t}(\xi^t) \end{aligned}$$

for $t = 1, \ldots, T-1$, and modified cost constraints

$$\begin{aligned} x_{\mathrm{c},1} &\geq \max\{h_1 s_{\mathrm{inv},1}, -p_1 s_{\mathrm{inv},1}\} \\ x_{\mathrm{c},t}(\xi^t) &\geq d_{t-1}\xi_{\mathrm{y},t} x_{\mathrm{o},t-1}(\xi^{t-1}) + \max\left\{h_t s_{\mathrm{inv},t}(\xi^t), -p_t s_{\mathrm{inv},t}(\xi^t)\right\} \end{aligned}$$

for $t = 2, \ldots, T$. An illustration of the problem flow is provided in Figure 2.3.

The extended model described here constitutes a multi-stage robust optimization problem with random recourse. Such problems are generically NP-hard even if they are solved in linear decision rules. However, in the case of linear decision rules, a tight tractable approximation can be obtained in the form of a conic-quadratic or semi-definite program (see Section 4 in Ben-Tal et al. [2004]). In what follows, we illustrate how our approach can be used to approximately solve this problem in polynomial decision rules. We remark that the problems described in Sections 2.6.1 and 2.6.2 can be brought to the form $\mathcal{R}$, without equality constraints, by eliminating the state variables $s_{\mathrm{inv},t}$ and $s_{\mathrm{co},t}$.

Table 2.2: Input data for the two test problems

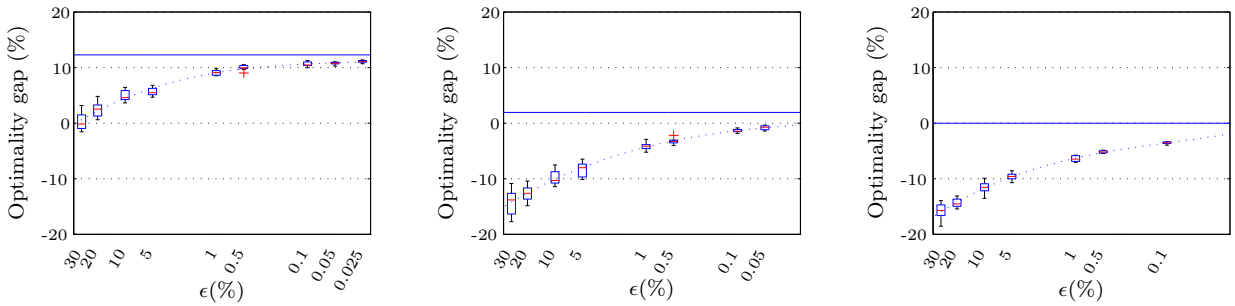
Parameter	$\mathcal{R}_{\mathrm{COC}}$	$\mathcal{R}_{\mathrm{RY}}$
T	5	7
(p_t, d_t, h_t)	$(11, 1, 10)$	$(11, 1, 10)$
$(s_{\mathrm{inv},1}, s_{\mathrm{co},1})$	$(0, 0)$	$(0, 0)$
$(\underline{x}_{\mathrm{o},t}, \overline{x}_{\mathrm{o},t})$	$(0, \infty)$	$(0, \infty)$
$(\underline{s}_{\mathrm{co},1}, \ldots, \underline{s}_{\mathrm{co},T})$	$(0, 47, 134, 188, 429)$	$(0, 32, 53, 223, 437, 437, 547)$
$(\overline{s}_{\mathrm{co},1}, \ldots, \overline{s}_{\mathrm{co},T})$	$(\infty, 94, 248, 370, 586)$	$(\infty, 98, 242, 461, 582, 838, 1068)$
$\overline{\xi}_{\mathrm{d},t}$	$100\left(1 + \frac{1}{2}\sin\left(\pi\frac{t}{6}\right)\right)$	$100\left(1 + \frac{1}{2}\sin\left(\pi\frac{t}{6}\right)\right)$
ρ_{d}	30%	30%
$\overline{\xi}_{\mathrm{y},t}$	n/a	0.90
ρ_{y}	n/a	10%

2.7 Numerical Experiments

In this section, we consider specific instances of the problems $\mathcal{R}_{\mathrm{COC}}$ and $\mathcal{R}_{\mathrm{RY}}$. The input data for the two test problems is summarized in Table 2.2.

We note that both $\mathcal{R}_{\mathrm{COC}}$ and $\mathcal{R}_{\mathrm{RY}}$ are strictly feasible. Indeed, the set $\{(x_{\mathrm{o},t})_{t\in\mathbb{T}} \in \mathbb{R}^T : \underline{x}_{\mathrm{o},t} < x_{\mathrm{o},t} < \overline{x}_{\mathrm{o},t},\ \underline{s}_{\mathrm{co},t} < \sum_{\tau=1}^{t} x_{\mathrm{o},\tau} < \overline{s}_{\mathrm{co},t}\ \forall t \in \mathbb{T}\}$ is nonempty in both cases. Therefore, there exist constant order quantities $(x_{\mathrm{o},t})_{t\in\mathbb{T}}$ satisfying the box constraints (2.20) strictly. Associated strictly feasible constant costs $(x_{\mathrm{c},t})_{t\in\mathbb{T}}$ can be found easily since all costs are unrestricted above. We report on optimality gaps, empirical violation probabilities and solver times for the corresponding approximate problems $\mathcal{SB}^N_{\mathrm{COC}}$ and $\mathcal{SB}^N_{\mathrm{RY}}$. We define the optimality gaps of the problems $\mathcal{SB}^N$ and $\mathcal{RB}$ as $(\inf \mathcal{SB}^N - \inf \mathcal{R})/\inf \mathcal{R}$ and $(\inf \mathcal{RB} - \inf \mathcal{R})/\inf \mathcal{R}$, respectively. We remark that problem $\mathcal{R}_{\mathrm{COC}}$ can be solved exactly by using a scenario tree that consists of the vertices of Ξ (see Section 3.3 in Bertsimas et al. [2011] for details). Therefore, the optimality gap for $\mathcal{SB}^N_{\mathrm{COC}}$ can be computed. To the best of our knowledge, there exists no algorithm to obtain the exact solution of problem $\mathcal{R}_{\mathrm{RY}}$ and thus the optimality gap for $\mathcal{SB}^N_{\mathrm{RY}}$ is not accessible. The empirical violation probability of a given $w \in \mathbb{R}^{n_w}$ is defined as

$$\hat{V}(w) := \frac{1}{10,000} \sum_{s=1}^{10,000} \mathbb{1}\left(f(w, \hat{\xi}^{(s)}) > 10^{-4}\right),$$



(a) $d = 1$ (b) $d = 2$ (c) $d = 3$

Figure 2.4: Optimality gaps for $\mathcal{RB}_{\mathrm{COC}}$ (solid lines) and $\mathcal{SB}^N_{\mathrm{COC}}$ (boxes and whiskers) for policies of degree 1, 2 and 3 in dependence of ϵ. The dotted curves represent a cubic fit of the average gaps as a function of $\log \epsilon$.

where $\{\hat{\xi}^{(s)}\}_{s\in\mathbb{N}}$ is a sequence of independent samples distributed according to $\mathbb{P}$. These samples are independent of the $\{\xi^{(s)}\}_{s\in\mathbb{N}}$ which are used in problem $\mathcal{SB}^N$.

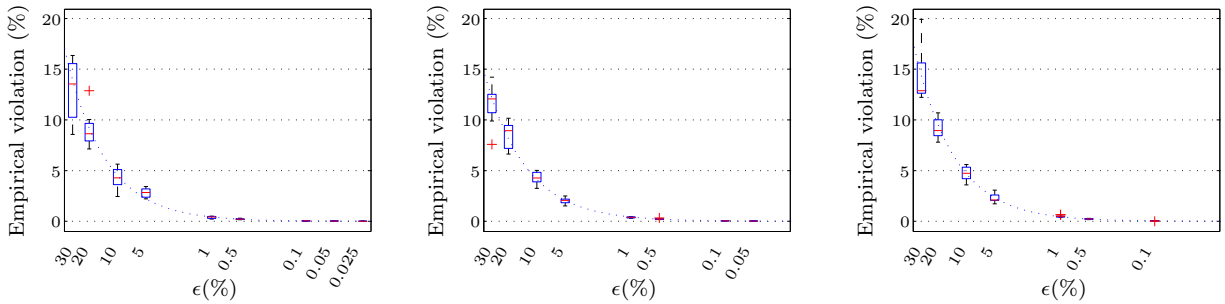
In our computational experiments we focus on polynomial decision rules, see Example 2.1, and use a fixed confidence level of $\beta = 0.1\%$ for the constraint sampling. The following procedure underlies all experiments:

- Select the degree d of the polynomial decision rules as well as the target violation probability ϵ and compute the corresponding sample size $N = N(\epsilon, \beta)$. Then, solve 100 instances of problem $\mathcal{SB}^N$, each based on a different set of N independent samples from $\mathbb{P}$. For each instance, compute the optimality gap and the empirical violation probability of the optimal solution. Moreover, record the solver run time[1].

- For each of the parameters recorded, compute statistics over the 100 problem instances.

Figure 2.4 visualizes the empirical distribution of the optimality gap of problem $\mathcal{SB}^N_{\mathrm{COC}}$ for different values of ϵ and d. The solid horizontal lines represent the optimality gap of problem $\mathcal{RB}_{\mathrm{COC}}$. Recall that the structure and the relatively small size of problem $\mathcal{R}_{\mathrm{COC}}$ enable us to compute $\inf \mathcal{R}_{\mathrm{COC}}$ exactly. We can also compute $\inf \mathcal{RB}_{\mathrm{COC}}$ exactly by solving a variant of problem $\mathcal{SB}^N_{\mathrm{COC}}$ in which the sample set coincides with the vertices of the hypercube Ξ. Thus all optimality gaps are numerically accessible. By Theorem 2.3, the optimality gap of problem

[1] All computational experiments were run on a 2.66GHz Intel Core i7-920 processor machine with 12GB RAM and all optimization problems were solved with CPLEX 12.0.



(a) $d = 1$ (b) $d = 2$ (c) $d = 3$

Figure 2.5: Empirical violation probabilities for problem $\mathcal{SB}^N_{\mathrm{COC}}$ (boxes and whiskers) for policies of degree 1, 2 and 3 in dependence of ϵ. The dotted curves represent a linear fit of the logarithm of the average empirical violation probability as a function of $\log \epsilon$.

$\mathcal{SB}^N_{\mathrm{COC}}$ converges to the optimality gap of problem $\mathcal{RB}_{\mathrm{COC}}$ as the number of samples is driven to infinity (or equivalently, as ϵ goes to 0). Our numerical results are in agreement with this convergence result, see Figure 2.4. We also observe that the variance of the optimality gap for problem $\mathcal{SB}^N_{\mathrm{COC}}$ decreases substantially as ϵ decreases and that the optimality gap of $\mathcal{RB}_{\mathrm{COC}}$ decreases rapidly as d increases. We gain about 10% in optimality when passing from linear to quadratic and about 2% when passing from quadratic to cubic policies. Note that cubic policies are indeed optimal for $\mathcal{R}_{\mathrm{COC}}$. This is not surprising since at each decision stage $t = 1, \ldots, 5$, the number of degrees of freedom offered by cubic policies exceeds the number of vertices of the uncertainty set Ξ^t.

Figure 2.5 shows the empirical distribution of the empirical violation probability in dependence of ϵ and d. As expected, it converges to 0 as ϵ decreases. We also note that the empirical violation probability is always substantially smaller than ϵ (as the confidence level β is sufficiently small, the retailer's target violation probability ϵ is never exceeded).

Our computational experiments illustrate the trade-off between optimality, feasibility and computational complexity. For less flexible policies corresponding to small values of d, ϵ can be made very small, thereby meeting the requirements of a risk averse retailer. For a retailer tolerating a higher violation probability, the costs can be significantly reduced by increasing d. For example, a retailer tolerating a violation probability of $\epsilon = 0.1\%$ can reduce the 99.9% worst-case costs from 2225.5 to 1941.4 (-12.7%) by passing from linear to cubic policies.

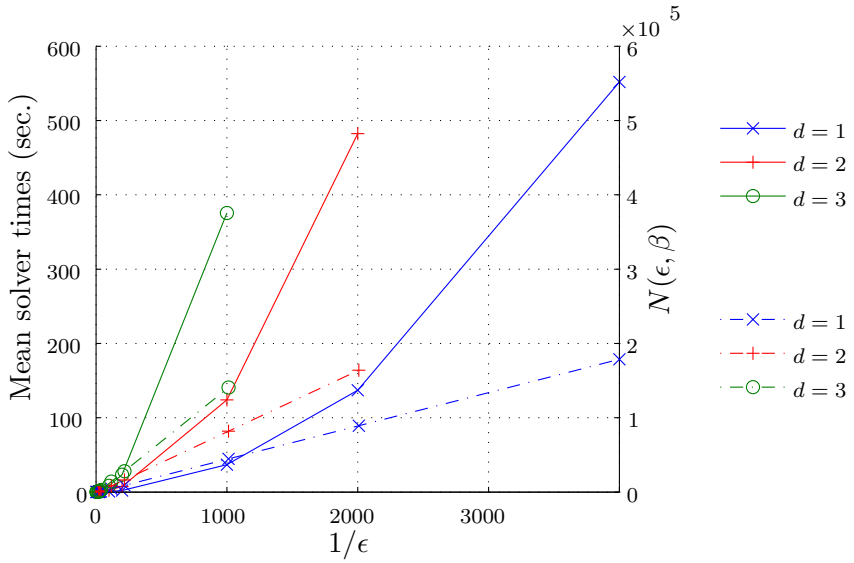
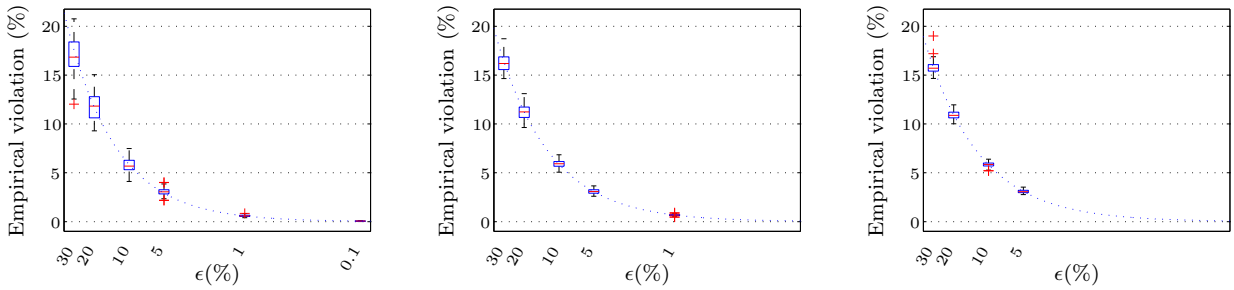


Figure 2.6: Mean solver times (solid lines) and numbers of samples (dashed lines) as a function of $1/\epsilon$ for problem $\mathcal{SB}^N_{\text{COC}}$.

Figure 2.6 displays the average solver time for problem $\mathcal{SB}^N_{\text{COC}}$ as a function of d and ϵ. As expected, the solver times for any fixed d are polynomial in the sample size $N(\epsilon, \beta)$ (indeed, Figure 2.6 reveals that $N(\epsilon, \beta)$ is linear in $1/\epsilon$, while the solver times are polynomial in $1/\epsilon$).

Next, we discuss problem $\mathcal{R}_{\text{RY}}$. Since this problem has seven stages and random recourse, the optimal value of $\mathcal{SB}^N_{\text{RY}}$ does not saturate for manageable sample sizes ($N < 180,000$). However, for any fixed violation probability ϵ, there is a substantial gain in optimality when passing from linear to quadratic decision rules. For example, the 99% worst-case cost faced by a retailer tolerating a violation probability of $\epsilon = 1\%$ amounts to 3988.8 on average when using linear decision rules and to 3173.7 when using quadratic decision rules (i.e., a reduction of about 20.4%). Figure 2.7 illustrates the convergence of the empirical violation probability as ϵ tends to 0. As before, the empirical violation probability is always substantially smaller than the target violation ϵ.

The observations above testify to the attractiveness of our approximation scheme. Firstly, at fixed target violation probability, the gain from increasing the complexity of the decision rules can be significant. Secondly, the approximation can be tailored to the risk preferences of the retailer to trade off feasibility against optimality. Finally, our framework mitigates the over-conservatism of mainstream robust optimization models, which often result in very high costs



(a) $d = 1$ (b) $d = 2$ (c) $d = 3$

Figure 2.7: Empirical violation probabilities for problem $\mathcal{SB}^N_{\mathrm{RY}}$ for policies of degree 1, 2 and 3 in dependence of ϵ. The dotted curves represent a linear fit of the logarithm of the average empirical violation probability as a function of $\log \epsilon$.

that cater for scenarios that are unlikely to materialize.

2.8 Appendix

2.8.1 Proof of Lemma 2.4

The proof of Lemma 2.4 involves the following definition.

Definition 2.4 (Linear dependence/independence on a set) *Let $S \subseteq \mathbb{R}^k$. If $c^\top b^d_t(\xi^t) = 0$ for all $\xi \in S$ implies $c = 0$, we say that the component functions of the basis vector b^d_t are linearly independent on S. Conversely, if there exists $c \neq 0$ such that $c^\top b^d_t(\xi^t) = 0$ for all $\xi \in S$, the basis functions are said to be linearly dependent on S.*

Proof of Lemma 2.4 We can assume without loss of generality that for each $t \in \mathbb{T}$ the components of b^d_t are linearly independent on Ξ. This assumption is non-restrictive since linearly dependent basis functions may always be removed without affecting the optimal value of problem $\mathcal{RB}$. Since $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ, there exists $N_0 \in \mathbb{N}$ such that the components of b^d_t are linearly independent on Ξ_N for all $t \in \mathbb{T}$ and $N \geq N_0$, see Lemma 2.6 below. Choose $N \geq N_0$. $\mathcal{SB}^N$ is feasible since it is a relaxation of $\mathcal{RB}$, which is feasible by our choice of d. Furthermore, $\mathcal{SB}^N$ has a closed feasible set and a continuous objective function, see assumption (A2). We next show that $\mathcal{SB}^N$ is bounded.

If $\mathcal{SB}^{N_0}$ is unbounded, then there is a sequence $\{w^{(i)}\}_{i\in\mathbb{N}}$ of points feasible in $\mathcal{SB}^{N_0}$ which satisfy $\lim_{i\to\infty}\theta^{(i)} = -\infty$, where $w^{(i)} = (\theta^{(i)}, \mathrm{vec}(X_1^{(i)}), \ldots, \mathrm{vec}(X_T^{(i)}))$. This is only possible if $\lim_{i\to\infty}\|X_t^{(i)}\|_2 = \infty$ for some $t \in \mathbb{T}$, where $\|\cdot\|_2$ denotes the Frobenius norm, see Section 2.3.1. However, by assumption (A3), we also have $|X_t^{(i)}\, b_t^d(\xi^t)| \leq R$ for all $\xi \in \Xi_{N_0}$, $i \in \mathbb{N}$. The bounded sequence $\{X_t^{(i)}/\|X_t^{(i)}\|_2\}_{i\in\mathbb{N}}$ is confined to the compact unit sphere and therefore has an accumulation point X_t^* with $\|X_t^*\|_2 = 1$. Let $\{X_t^{(i_j)}/\|X_t^{(i_j)}\|_2\}_{j\in\mathbb{N}}$ be a subsequence converging to X_t^*. Then, we have

$$\begin{aligned}
\lim_{j\to\infty} |X_t^{(i_j)}\, b_t^d(\xi^t)| \leq R \;\; \forall \xi \in \Xi_{N_0} \;\; &\Rightarrow \;\; \lim_{j\to\infty} \|X_t^{(i_j)}\|_2 \left| \frac{X_t^{(i_j)}}{\|X_t^{(i_j)}\|_2}\, b_t^d(\xi^t) \right| \leq R \;\; \forall \xi \in \Xi_{N_0} \\
&\Rightarrow \;\; \left| X_t^*\, b_t^d(\xi^t) \right| = 0 \;\; \forall \xi \in \Xi_{N_0} \\
&\Rightarrow \;\; X_t^* = 0,
\end{aligned}$$

where the last implication follows from the linear independence of the basis functions on Ξ_{N_0}. This contradicts our earlier result that $\|X_t^*\|_2 = 1$. We thus conclude that $\mathcal{SB}^{N_0}$ is bounded. Since $\mathcal{SB}^N$ is a restriction of $\mathcal{SB}^{N_0}$, it is also bounded. The above reasoning implies that $\mathcal{SB}^N$ is solvable and has a finite optimal value.

Moreover, the optimal solution of $\mathcal{SB}^N$ is contained in the set

$$\begin{aligned}
W^* := \Big\{ & w = (\theta, \mathrm{vec}(X_1), \ldots, \mathrm{vec}(X_T)) \in \mathbb{R}^{n_w} : \\
& \inf \mathcal{SB}^{N_0} \leq \theta \leq \inf \mathcal{RB}, |X_t\, b_t^d(\xi^t)| \leq R \; \forall \xi \in \Xi_{N_0},\; t \in \mathbb{T} \Big\},
\end{aligned}$$

which is non-empty and compact. As the choice of $N \geq N_0$ was arbitrary and the definition of W^* is independent of N, the claim follows. □

Lemma 2.6 *If the components of b_t^d are linearly independent on Ξ and $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ, then there exists $N_t \in \mathbb{N}$ such that the components of b_t^d are linearly independent on Ξ_N for all $N \geq N_t$.*

Proof Suppose that the components of b_t^d are linearly dependent on Ξ_N for all $N \in \mathbb{N}$. Thus, there exist $c_N \in \mathbb{R}^{s_d(k^t)}$ such that $|c_N| = 1$ and $c_N^\top\, b_t^d(\xi^t) = 0$ for all $\xi \in \{\xi^{(l)}\}_{l\in\mathbb{N}}$. The sequence

$\{c_N\}_{N\in\mathbb{N}}$ is confined to the compact unit sphere and therefore has an accumulation point c^* with $|c^*| = 1$. Let $\{c_{N_j}\}_{j\in\mathbb{N}}$ be a subsequence converging to c^*. Then,

$$\begin{aligned}
&\lim_{j\to\infty} c_{N_j}^\top\, b_t^d(\xi^t) = 0 \;\; \forall \xi \in \{\xi^{(l)}\}_{l\in\mathbb{N}} \\
\Rightarrow\;& c^{*\top}\, b_t^d(\xi^t) = 0 \;\; \forall \xi \in \{\xi^{(l)}\}_{l\in\mathbb{N}} \\
\Rightarrow\;& c^{*\top}\, b_t^d(\xi^t) = 0 \;\; \forall \xi \in \Xi \\
\Rightarrow\;& \text{the components of } b_t^d \text{ are linearly dependent on } \Xi.
\end{aligned}$$

The second implication follows from the continuity of the basis functions and since $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ, while the third implication holds since $c^* \neq 0$. The last implication contradicts our assumption, and thus the claim follows. □

2.8.2 Almost Sure Density of Sample Sequences

Lemma 2.7 *The sequence $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ w.p.1.*

Proof Let $\mathbb{P}^\infty := \bigtimes_{l=1}^\infty \mathbb{P}$ denote the probability distribution of the stochastic process $\{\xi^{(l)}\}_{l\in\mathbb{N}}$. Then, for $z \in \Xi$ and $\epsilon > 0$, we have

$$\mathbb{P}^\infty(B_\epsilon(z) \cap \{\xi^{(l)}\}_{l\in\mathbb{N}} = \emptyset) \;=\; \prod_{l=1}^{\infty} \mathbb{P}(\xi^{(l)} \notin B_\epsilon(z)) \;=\; 0. \tag{2.21}$$

The last equality holds since z is an element of the support of $\mathbb{P}$ and thus $\mathbb{P}(\xi^{(l)} \in B_\epsilon(z)) > 0$ independently of $l \in \mathbb{N}$. Since Ξ is convex and fully dimensional by assumption (A5), the sequence $\{\xi^{(l)}\}_{l\in\mathbb{N}}$ is dense in Ξ if and only if for every $z \in \Xi \cap \mathbb{Q}^k$ and $\epsilon \in \mathbb{Q}_+$, the set $B_\epsilon(z)$ contains at least one sample. Therefore,

$$\begin{aligned}
&\mathbb{P}^\infty(\{\xi^{(l)}\}_{l\in\mathbb{N}} \text{ is dense in } \Xi) \\
=\;& 1 - \mathbb{P}^\infty(\exists z \in \Xi \cap \mathbb{Q}^k,\; \epsilon \in \mathbb{Q}_+ \text{ with } B_\epsilon(z) \cap \{\xi^{(l)}\}_{l\in\mathbb{N}} = \emptyset) \\
\geq\;& 1 - \sum_{z\in\Xi\cap\mathbb{Q}^k} \sum_{\epsilon\in\mathbb{Q}_+} \mathbb{P}^\infty(B_\epsilon(z) \cap \{\xi^{(l)}\}_{l\in\mathbb{N}} = \emptyset) \\
=\;& 1,
\end{aligned}$$

where the second line follows from the Bonferroni inequality and the last line follows from (2.21). □

Chapter 3

Hedging Electricity Swing Options in Incomplete Markets

In Chapter 2, we proposed a tractable approximation scheme for multi-stage robust optimization problems and illustrated its convergence and scalability properties in the context of two moderately sized problems. In this chapter, we investigate a large-scale application of this approach. Specifically, we develop a methodology for hedging and valuing path-dependent electricity derivatives such as swing options in today's deregulated markets. Swing options allow the option holder to purchase electric energy from the option writer at a prescribed price during a prescribed time-period and constitute a generalized class of derivatives that contains forwards, European and American options as special cases. They enable utility companies to mitigate the risk of highly volatile spot prices and to simultaneously manage the uncertain demand from customers. Unlike financial markets, electricity markets are incomplete. Therefore, a swing option cannot be assigned a unique fair value and instead admits a whole interval of (no-arbitrage) prices consistent with those of traded instruments. In this chapter we propose to determine the option's no-arbitrage price interval by hedging its payoff stream with basic market securities (such as forward contracts) both from the perspective of the holder and the writer of the option. The end-points of the no-arbitrage interval are given by the optimal values of two robust optimization problems, which we solve with the approximation scheme developed

in Chapter 2.

3.1 Introduction

The spot price of electricity is notoriously volatile. Unpredictable demand patterns as well as the limited storability and grid-bound nature of electricity result in frequent price spikes. The associated price risk is absorbed by public utilities which buy energy at uncertain wholesale prices and sell it to end consumers at fixed retail prices. In a regulated market, the government can set prices that allow utility companies to recover their costs. In a liberalized market, however, such compensation is no longer possible.

Utility companies seek to mitigate the resulting price risks by investing in market-traded electricity derivatives such as forwards and options. An electricity forward contract is an obligation to buy electric energy at a prescribed delivery rate over a prescribed delivery period and in return for a predetermined unit price (the forward price). Both the load profile (i.e., the delivery rate in each hour of the delivery period) and the forward price are agreed at the time when the forward is issued. A European call option is the right (but not the obligation) to buy electric energy at a prescribed delivery rate over a prescribed delivery period and at a predetermined unit price (the strike price).

Unfortunately, forwards and European options constitute poor hedging instruments for utility companies. Indeed, forwards neither provide any flexibility in the volume nor in the timing of the energy delivery, both of which are essential for a utility company that is unable to control the energy consumption of its customers. Conversely, using European options to hedge against price spikes results in a costly overprotection, see Jaillet et al. [2004]. As an example, consider a utility company that wishes to hedge against up to five price spikes forecasted for a typical summer month. A perfect hedge is provided by purchasing a European option that is exercisable on each day of the month. However, since the option will be exercised at most on five days, the utility company overpays for the desired protection. Thus, utility companies often hedge spot price risks by purchasing swing options.

A swing option constitutes an agreement to purchase and/or sell electric energy during a fixed period of time and at a predetermined strike price. Like European call options, swing options offer some flexibility both in the timing and the volume of the energy delivery. However, they can only be exercised a limited number of times. For example, a typical swing option may allow its holder to purchase between 500 and 1,000 MWh of electric energy at a unit price of 60 €/MWh during the next month, while not more than 50 MWh may be bought on each day (at constant delivery rate). Swing options are popular hedging instruments in the electricity sector and are used extensively by utility companies, see Carmona and Ludkovski [2010].

In this chapter we aim to determine the monetary value of the right to exercise a swing option, that is, the premium that the option holder has to pay the writer at the time when the option is negotiated. To this end, we assume that the market of basic securities (i.e., cash as well as the forwards and European options on energy) is arbitrage-free. A market is arbitrage-free if there is no self-financing portfolio that transforms a non-positive initial investment into a nonnegative terminal wealth that is nonzero with positive probability, see Luenberger [1997]. Clearly, the market should remain arbitrage-free when the swing option is added to the existing investment opportunities. It can be shown that the option has a unique no-arbitrage price if and only if there exists a portfolio of basic securities that generates, with certainty, the same cash flow stream as the option. Such a perfect option replication is usually not possible in electricity markets due to the limited storability of electricity, the presence of jumps and spikes in the spot price, market illiquidity, and high transaction costs which preclude frequent adjustments to the replicating portfolio. In the language of finance, electricity markets are therefore *incomplete*.

In this chapter we propose a pricing framework for electricity swing options in incomplete markets. Instead of a single objective price, an incomplete market allows for a whole interval of option prices that preserve arbitrage-freeness. We determine this interval by investigating two complementary hedging portfolios of basic securities. We obtain the lower end of the no-arbitrage interval (the holder's price) by computing the maximum amount of money borrowed today that the option holder can repay through exercising the option and trading in the basic securities. Likewise, we obtain the upper end of the no-arbitrage interval (the writer's price) by calculating the minimum amount of money borrowed today that enables the option writer to

cover all obligations arising from the option by trading in the basic securities. The option price agreed by the holder and the writer emerges as the result of a negotiation process, and it must lie within the no-arbitrage interval if both contract parties act rationally. The holder's and the writer's price are representable as the optimal values of two robust optimization problems. These problems do not allow for a direct solution, and we propose an approximate solution scheme that simplifies the underlying information process and solves the resulting problems with the approach developed in Chapter 2. Our approach has natural applications to the valuation of power plants, refineries, mines and oil fields.

To the best of our knowledge, all existing pricing schemes for swing options assume either explicitly or implicitly that the underlying market is complete and that swing options can be assigned a unique price. Most approaches rely on stochastic dynamic programming to determine the optimal exercise strategy, see e.g., Thompson [1995], Jaillet et al. [2004] and Haarbrücker and Kuhn [2009]. Recently, the popular Least Squares Monte Carlo technique developed by Longstaff and Schwartz [2001] has been employed to estimate the conditional expectations arising in the stochastic dynamic programming iterations, see Ibáñez [2004]. Analytical solutions based on stochastic calculus exploit the connection between swing option pricing and multiple stopping problems, see Carmona and Touzi [2008]. A recent survey of the literature on swing option pricing is provided in Carmona and Ludkovski [2010]. The ramifications of market incompleteness have been studied extensively in the context of financial options theory. This line of research investigates strategies for super- and sub-replicating the option payoffs and evaluates the expected utility generated by the option, see e.g., King [2002], Staum [2007], Hao [2008] and the references therein. However, all studies of market incompleteness focus on standard financial options and do not easily generalize to the exotic contractual features of swing options and the peculiarities of electricity markets. In addition, emphasis is placed on a theoretical characterization of option prices, and no numerical techniques are provided to determine an option's no-arbitrage interval.

The main contributions of this chapter are summarized below:

1. We consider general incomplete market models incorporating jumps, spikes and stochastic

volatility. We formulate two robust optimization problems and demonstrate that their optimal objective values yield the end-points of the no-arbitrage interval. We highlight the existence of multiple optimal exercise strategies for the swing option and discuss how the choice of the optimal exercise strategy affects the no-arbitrage interval. In addition, we propose a methodology for selecting an exercise strategy that maximizes the utility of the holder. We consider the case where the holder is subjected to *bounded rationality* and is thus unable to commit to an optimal exercise strategy, and formulate a hedging problem for the writer in this case also.

2. We demonstrate that contrary to the case of complete markets, in incomplete markets, there does not necessarily exist an optimal bang-bang exercise strategy for the swing option.

3. We propose a tractable approximation scheme for the hedging problems and illustrate how the design parameters can be used to model the negotiation process between the holder and the writer. We demonstrate the tractability properties of our approach by hedging a swing option with a delivery period of over 900 hours (time-periods) in a market consisting of 15 basic securities.

This chapter is organized as follows. In Section 3.2 we formulate the two optimization problems that allow us to determine a swing option's no-arbitrage interval. We discuss the nature of the optimal swing option exercise strategy in Section 3.3. To ensure tractability, we propose several approximation techniques in Section 3.4. In Section 3.5, we discuss alternative applications. We illustrate our approach with a case study in Section 3.6.

3.2 Problem Formulation

We consider a finite planning horizon consisting of T time intervals (periods) $\mathbb{T} := \{1, \ldots, T\}$. Our aim is to find the value of the swing option at the beginning of the first period (i.e., today). We model security prices and dividends as random variables that are defined on a probability

space $(\Xi, \mathcal{F}, \mathbb{P})$. We assume that the elements of the sample space Ξ can be represented as $\xi = (\xi_1, \ldots, \xi_{T+1})$, where the subvector $\xi_t \in \mathbb{R}^k$, $t \in \mathbb{T}$, is observed at the beginning of period t, while $\xi_{T+1} \in \mathbb{R}^k$ is observed at the end of period T. Note that ξ_1 is known today and is therefore not random. We assume that Ξ is the smallest closed set that satisfies $\mathbb{P}(\xi \in \Xi) = 1$, and we assume that Ξ is bounded. We denote by $\xi^t = (\xi_1, \ldots, \xi_t) \in \mathbb{R}^{tk}$ the history of observations up to period t and let $\mathbb{E}(\cdot)$ denote the expectation operator with respect to $\mathbb{P}$. Finally, for any $m, n \in \mathbb{N}$, we let $\mathcal{L}_{m,n}$ denote the space of all measurable functions from $\mathbb{R}^m$ to $\mathbb{R}^n$ that are bounded on compact sets.

3.2.1 Electricity Swing Options

An electricity swing option is an agreement to purchase and/or sell electric energy at a predetermined strike price K throughout the delivery period $\mathbb{T}$. It gives its holder the right to select the sequence of exercise decisions (the exercise strategy) $e := (e_1, e_2, \ldots, e_T)$, where $e_t \in \mathbb{R}$ is selected at the beginning of period t and denotes the amount of energy to be purchased during that period. The exercise quantity e_t is subject to lower and upper limits denoted by $\underline{e}_t$ and $\overline{e}_t$, respectively. Moreover, the cumulative amount of energy received during the entire delivery period, $\sum_{t\in\mathbb{T}} e_t$, is subject to lower and upper limits $\underline{c}$ and $\overline{c}$. In practice, the exercise decision e_t is not pre-committed today. Instead, it is allowed to adapt to the information revealed up to time t and is thus modeled as a decision rule of ξ^t. For convenience, we introduce the set of all admissible exercise strategies

$$\mathcal{E} := \left\{ e \in \bigtimes_{t\in\mathbb{T}} \mathcal{L}_{kt,1} \;:\; \underline{e}_t \leq e_t(\xi^t) \leq \overline{e}_t \;\; \forall t \in \mathbb{T}, \;\; \underline{c} \leq \sum_{t\in\mathbb{T}} e_t(\xi^t) \leq \overline{c} \quad \mathbb{P}\text{-a.s.} \right\}.$$

The following comments are in order. Firstly, the delivery period of the swing option may be a strict subset of $\mathbb{T}$ given by $\{\underline{t}, \ldots, \overline{t}\}$. By setting $\underline{e}_t = \overline{e}_t = 0$ for all $t \in \mathbb{T} \setminus \{\underline{t}, \ldots, \overline{t}\}$, we can assume that $(\underline{t}, \overline{t}) = (1, T)$. Moreover, one often distinguishes base, peak and off-peak swing options. While base options can be exercised in each period of the delivery period, peak options are only exercisable in peak periods (i.e., from 8am to 8pm on working days). A peak option

can be modeled by setting $\underline{e}_t = \overline{e}_t = 0$ for all off-peak periods t. Off-peak options, which are only exercisable in off-peak periods, can be modeled accordingly. In addition, we remark that in practice, the per-period delivery quantity e_t must be chosen prior to period t so that the option writer has time to arrange the delivery. Our assumption that e_t is chosen at the beginning of period t is an acceptable idealization that facilitates a transparent exposition. Finally, some swing options permit violation of the cumulative energy limits $\underline{c}$ and $\overline{c}$, in which case penalties are imposed for each unit of shortfall or exceedance, respectively.

By our definition, an electricity forward can be regarded as a swing option with $\underline{e}_t = \overline{e}_t$ for all $t \in \mathbb{T}$ and $\underline{c} = \overline{c} = \sum_{t\in\mathbb{T}} \underline{e}_t$. Likewise, if $\underline{e}_t = 0$ for all $t \in \mathbb{T}$, $\underline{c} = 0$ and $\overline{c} = \sum_{t\in\mathbb{T}} \overline{e}_t$, then the swing option reduces to a European call option. We thus conclude that swing options constitute a generalized class of derivatives that contains forwards and European options as special cases.

3.2.2 Market Model

We consider an energy market that consists of a spot exchange and J basic securities indexed by $j \in \{1, \ldots, J\}$. We denote the average spot price of electricity over period t by $S_t(\xi_t)$. The price of security j at the beginning of period t is given by $P_{t,j}(\xi_t)$, while $P_{T+1,j}(\xi_{T+1})$ denotes the price of security j at the end of period T. Security j pays a random dividend $d_{t,j}(\xi_t)$ at the beginning of period t. We aggregate the prices and dividends to vectors $P_t := (P_{t,1}, \ldots, P_{t,J})$ and $d_t := (d_{t,1}, \ldots, d_{t,J})$.

The market we consider is incomplete. Before introducing the formal definition of a complete market, we define the concept of a portfolio strategy.

Definition 3.1 (Portfolio strategy) *A portfolio strategy is a pair* $x := (x_0, \ldots, x_T)$, $u := (u_1, \ldots, u_T)$, $x_0 \in \mathbb{R}^J$, x_t, $u_t \in \mathcal{L}_{kt,J}$ $\forall t \in \mathbb{T}$, *satisfying*

$$\left.\begin{array}{rcll} x_1(\xi^1) & = & x_0 + u_1(\xi^1) & \\ x_t(\xi^t) & = & x_{t-1}(\xi^{t-1}) + u_t(\xi^t) & \forall t \in \mathbb{T}\backslash\{1\} \end{array}\right\} \quad \mathbb{P}\text{-}a.s.$$

The variable $x_{t,j}$, $j = 1, \ldots, J$ denotes the number of basic securities of type j held within period t, while the adjustment variable $u_{t,j}$ denotes the number of securities of type j bought at the beginning of period t. The variable $x_{0,j}$ represents the number of basic securities of type j held at the beginning of period t, before any adjustments. We denote by $\mathcal{P}$ the set of all portfolio strategies. The dividend process $C^{x,u} := (C_0^{x,u}, \ldots, C_{T+1}^{x,u})$, $C_0^{x,u} \in \mathbb{R}$, $C_t^{x,u} \in \mathcal{L}_{kt,1} \forall t \in \mathbb{T} \cup \{T+1\}$, generated by the strategy (x, u) is defined through

$$\begin{aligned} C_0^{x,u} &= -P_1(\xi_1)^\top x_0 \\ C_t^{x,u}(\xi^t) &= d_t(\xi_t)^\top x_t(\xi^t) - P_t(\xi_t)^\top u_t(\xi^t) \quad \forall t \in \mathbb{T} \quad \mathbb{P}\text{-}a.s. \\ C_{T+1}^{x,u}(\xi^{T+1}) &= P_{T+1}(\xi_{T+1})^\top x_T(\xi^T) \quad \mathbb{P}\text{-}a.s. \end{aligned}$$

A portfolio strategy $(x, u) \in \mathcal{P}$ is said to be self-financing if it does not experience any exogenous cash injections, i.e.,

$$C_t^{x,u}(\xi^t) \geq 0 \quad \forall t \in \mathbb{T} \cup \{T+1\} \quad \mathbb{P}\text{-}a.s.$$

Remark 3.1 *According to standard literature convention, self-financing portfolios do not allow for exogenous cash injections or withdrawals. Nevertheless, as we work with incomplete markets and for reasons that will become clear later on, we prefer this definition.*

Definition 3.2 (Complete market) *We say that a market is complete relative to a process $f := (f_1, \ldots, f_T)$, $f_t \in \mathcal{L}_{kt,1}$ $\forall t \in \mathbb{T}$, if there exists a replicating (not necessarily self-financing) portfolio strategy $(x, u) \in \mathcal{P}$ whose dividend process $C^{x,u}$ satisfies*

$$\left.\begin{aligned} C_t^{x,u}(\xi^t) &= f_t(\xi^t) \quad \forall t \in \mathbb{T} \\ C_{T+1}^{x,u}(\xi^{T+1}) &= 0 \end{aligned}\right\} \quad \mathbb{P}\text{-}a.s.$$

The initial value $-C_0^{x,u}$ of the portfolio strategy is referred to as the no-arbitrage price of the claim (process) f. We say that a market is complete if it is complete relative to any process f.

Remark 3.2 *We note that although the replicating portfolio strategy (x, u) may not be self-financing, the portfolio obtained by purchasing (selling) the claim f and selling (purchasing) (x, u) is always self-financing (according to even the standard literature convention).*

For notational convenience, we introduce the set $\mathcal{S}(f)$ of all self-financing portfolios containing the claim $f := (f_1, \ldots, f_T)$, $f_t \in \mathcal{L}_{kt,1}$ $\forall t \in \mathbb{T}$, defined through

$$\mathcal{S}(f) := \left\{ (x,u) \in \mathcal{P} \ : \begin{array}{l} f_t(\xi^t) + C_t^{x,u}(\xi^t) \geq 0 \quad \forall t \in \mathbb{T}, \ \mathbb{P}\text{-a.s.} \\ C_{T+1}^{x,u}(\xi^{T+1}) \geq 0 \quad \mathbb{P}\text{-a.s.} \end{array} \right\}.$$

We now present the formal definition of an arbitrage opportunity.

Definition 3.3 (Arbitrage opportunity) *An arbitrage opportunity is a self-financing portfolio $(x, u) \in \mathcal{P}$ whose dividend process satisfies*

(a) $C_0^{x,u} \geq 0$, *and*

(b) $C_0^{x,u} > 0$ *or* $\mathbb{P}(C_t^{x,u}(\xi^t) > 0) > 0$ *for some* $t \in \mathbb{T} \cup \{T+1\}$.

We make the following mild assumptions, which are assumed to hold throughout the remainder of the chapter.

(M1) The market is arbitrage-free.

(M2) The market is frictionless, i.e., there are no short-sales or borrowing restrictions and no taxes or transaction costs are incurred when trading basic securities.

(M3) The market participants are price-takers, that is, their trades do not affect the market prices.

(M4) S_t, P_t and d_t are continuous functions of the random variables ξ_t observed in period t.

(M5) Security $j = 1$ is risk-less (e.g., a money market account); securities $j = 2, \dots J$ represent forwards and options on electricity: their dividends reflect the exercise costs and the revenues generated by selling the delivered electricity on the spot market.

(M6) The risk-less borrowing and lending rate is zero.

(M7) The basic market securities are traded in each period $t \in \mathbb{T}$.

Several comments are in order. Firstly, assumption (M1) is innocent: if the market was not arbitrage-free, then market participants could make infinite profits at no risk. Such arbitrage opportunities would vanish quickly. Moreover, assumption (M2) could be relaxed: our models can easily be extended to accommodate linear transaction costs and taxes or short-sales and borrowing constraints by making suitable adjustments to the dividend process of the portfolio strategy (x, u). We note that the absence of market frictions does not preclude market incompleteness. Assumption (M3) is standard in the literature. Assumption (M4) holds true for a vast number of realistic spot and forward price models in gas and electricity, which incorporate mean-reversion, jumps, spikes and stochastic volatility. Assumption (M6) is non-restrictive and can always be enforced by normalizing the prices of all securities by the price of a zero coupon bond with appropriate maturity, see e.g., Section 7 in Harrison and Kreps [1979]. Assumption (M6) is modeled by setting $P_{t,1}(\xi_t) = 1$ and $d_{t,1}(\xi_t) = 0$ for all $\xi \in \Xi$. Finally, assumption (M7) is an idealization introduced to simplify notation. Indeed, in period t some of the securities may have expired, whereas others may not yet be traded on the market.[1] When formulating our optimization problems, we will discuss how assumption (M7) may be relaxed by adding constraints to our optimization problems to avoid trading in unavailable securities.

In our market model, there are at least three sources of potential market incompleteness. Firstly, trading in the basic securities is not continuous, but only allowed at discrete time intervals. Moreover, the dividends arising from an optimal swing option exercise strategy may not be perfectly correlated with the dividend process of any portfolio strategy (note that, due to the limited storability of electricity, hedging with the underlying is not possible). Finally, we

[1] On the European Energy Exchange, for example, trading in a weekly forward starts only five weeks before its delivery period.

consider generic market models which may incorporate jumps, spikes and stochastic volatility.

We now add a swing option to the set of existing securities. Our goal is to determine the interval of swing option prices that preserve arbitrage-freeness of the market. In a complete market, this no-arbitrage interval collapses to a singleton, and the unique no-arbitrage price coincides with the initial value of a portfolio of the basic securities that replicates the swing option's dividend stream, see Definition 3.2. In an incomplete market, however, such a perfect replication is usually not possible, and the no-arbitrage interval is non-degenerate. We obtain the lower end of this interval by computing the maximum loan that the option holder can refinance through exercising the option and trading in the basic securities. This value, which we call the *holder's price*, constitutes the highest price at which all rational market participants would agree to buy the option since there is no risk involved. Any option price below the holder's price would allow the buyer to make arbitrage profits by purchasing swing options and simultaneously trading in the basic securities. Likewise, we obtain the upper end of the no-arbitrage interval by calculating the minimum loan that enables the option writer to cover all obligations arising from the option by trading in the basic securities. This value, which we call the *writer's price*, is the lowest price at which all rational market participants would agree to issue the option since there is no risk involved. Any option price above the writer's price would allow the writer to make arbitrage profits by selling swing options and simultaneously trading in the basic securities. The holder's and the writer's prices are representable as optimal values of robust optimization problems. The solutions to these problems also reveal arbitrage opportunities that emerge if the option's price falls outside the no-arbitrage interval.

3.2.3 Holder's Problem

The holder's price is given by the optimal value of the following optimization problem.

$$\begin{array}{rcll} H & = & \max\limits_{e,x,u} & C_0^{x,u} \\ & & \text{s.t.} & e \in \mathcal{E},\ (x,u) \in \mathcal{S}^{\mathrm{h}}(e), \end{array} \qquad (\mathcal{H})$$

where

$$\mathcal{S}^{\mathrm{h}}(e) := \Big\{(x,u) \in \mathcal{S}\left(\left\{e_t(\xi^t)(S_t(\xi_t)-K)\right\}_{t\in\mathbb{T}}\right)\Big\}$$

denotes the set of all self-financing portfolios containing the swing option with exercise strategy e. The decision variables are the exercise strategy e of the swing option and the portfolio strategy (x,u). We remark that the decision variables are decision rules. The first constraint in problem $\mathcal{H}$ stipulates that the exercise strategy of the option must be admissible, i.e., it must satisfy the per-period and cumulative energy limits. The second constraint stipulates that the portfolio containing the swing options must be self-financing.

The option holder aims to maximize the amount of money that can be borrowed today and repaid with certainty by exercising the swing option. Thus, the holder borrows an amount $C_0^{x,u}$ of money today by short-selling basic securities (note that $C_0^{x,u} > 0$ only if at least one position variable $x_{0,j}$ is strictly negative). Exercising the swing option at time t by buying $e_t(\xi^t)$ units of energy from the option writer at price K and selling this energy quantity immediately on the spot market at price $S_t(\xi_t)$ results in a dividend $e_t(\xi^t)(S_t(\xi_t)-K)$. Thus, the self-financing condition requires that the costs $P_t(\xi_t)^\top u_t(\xi^t)$ to rebalance the portfolio at the beginning of period t must be recovered with certainty from the dividend $e_t(\xi^t)(S_t(\xi_t)-K)$ generated by the swing option and the dividends $d_t(\xi_t)^\top x_t(\xi^t)$ received from the basic securities. It further ensures that the portfolio strategy (x,u) has non-negative terminal value with probability one and thus guarantees that the initial loan can be repaid by only using the dividends received from exercising the swing option.

We note that the set of admissible portfolio strategies (x,u) may be restricted further, for example by prohibiting short-sales, see assumption (M2). For ease of exposition, we disregard such constraints here.

Remark 3.3 *As noted in Section 3.2.2, assumption (M7) is unrealistic in electricity markets. Indeed, some of the basic securities might not be traded throughout the swing option's contract period, while the market may be closed during certain times of the day, during weekends or bank holidays. Such restrictions may be modeled in our framework by enforcing $u_{t,j} = 0$ if contract j*

is not traded at t. Similarly, some of the basic securities may not be traded anymore although they have not yet expired. If a given contract j expiring after time T is not traded beyond time T, we require $x_{T,j} = 0$. Indeed, in the absence of such a constraint, and since the swing option holder is unable to trade in the contract at time $T+1$, dividend payments might arise in the future implying that the portfolio will not necessarily be self-financing. We note that if the dividends from contract j are almost surely positive (negative), it suffices to require $x_{t,j} \geq 0$ ($x_{t,j} \leq 0$) and set $P_{T+1,j} = d_{T+1,j}$, where $d_{T+1,j}$ denotes the dividend from security j at time $T+1$.

We can now prove the correctness of the holder's problem.

Proposition 3.1 *Let $(e^\star, x^\star, u^\star)$ denote an optimal solution to the holder's problem with objective value H.*

(a) If the swing option trades at price $V < H$, then the option holder can make arbitrage profits.

(b) If the swing option trades at price $V > H$, no arbitrage can be made by buying the swing option.

Proof Assume first that the option trades at a price $V < H$. Then, buy the swing option at price V and take a loan of amount $H = C_0^{x^\star,u^\star}$, which yields an immediate payoff $H - V > 0$. Since $(x^\star, u^\star) \in \mathcal{S}^{\mathrm{h}}(e^\star)$, the loan H can be repaid almost surely by exercising the swing option according to $e^\star$ and by following the portfolio strategy $(x^\star, u^\star)$. Furthermore, the self-financing constraints implied by the requirement $(x, u) \in \mathcal{S}^{\mathrm{h}}(e^\star)$ guarantee that, with probability one, the dividend process arising from this strategy will be nonnegative. Thus, we have constructed an arbitrage.

Assume now that the swing option trades at a price $V > H$ and that arbitrage profits can be made by buying the swing option. Then, there is a swing option exercise strategy $e \in \mathcal{E}$ and a portfolio strategy $(x, u) \in \mathcal{S}^{\mathrm{h}}(e)$ which satisfies $C_0^{x,u} - V \geq 0$, see Definition 3.3, i.e., $C_0^{x,u} > H$.

Such a portfolio would correspond to a feasible solution (e, x, u) of the holder's problem with an objective value $C_0^{x,u} > H$. This contradicts the optimality of $(e^\star, x^\star, u^\star)$. Thus, no arbitrage profits can be made by buying the swing option at a price $V > H$. □

Remark 3.4 *If the option trades at price $V = H$, it is possible for the option holder to make arbitrage profits if and only if there exists a triplet (e, x, u) feasible in $\mathcal{H}$ such that for some $t \in \mathbb{T} \cup \{T+1\}$, the probability that the self-financing constraints implied by $(x, u) \in \mathcal{S}^{\mathrm{h}}(e)$ hold with strict inequality is strictly positive, see Definition 3.3. This can be seen immediately by making suitable adjustments to the second part of the proof of Proposition 3.1, see also Proposition 3.5(a).*

The holder's problem might have multiple optimal solutions (e, x, u). While each of the optimal solutions results in the same holder's price for the swing option, buyers will typically have preferences regarding the optimal strategy they wish to implement. Specifically, the option holder might wish to select the strategy which maximizes or minimizes a given function $g(e, x, u)$, for example the expected exercise profits $g(e, x, u) = \mathbb{E}[\sum_{t\in\mathbb{T}} e_t(\xi^t)(S_t(\xi_t) - K)]$. An optimal strategy with respect to the preferences of the buyer may then be obtained as the solution to the problem

$$\begin{aligned} \max_{e,x,u} \,/\, \min_{e,x,u} \quad & g(e, x, u) \\ \text{s.t.} \quad & C_0^{x,u} = H \\ & e \in \mathcal{E},\ (x, u) \in \mathcal{S}^{\mathrm{h}}(e)\,. \end{aligned} \tag{3.1}$$

We now illustrate with a simple example the possibility of multiple optimal solutions in the holder's problem and highlight how the choice of a specific strategy will depend on the utility of the holder.

Example 3.1 *Suppose $\mathbb{T} = \{1, 2, 3\}$ and consider an electricity market that consists of a spot exchange and the risk-free asset, i.e., $J = 1$. In addition, assume that $S_t(\xi_t) = \xi_t$ and that ξ is uniformly distributed on*

$$\Xi = \left\{\xi \in \mathbb{R}^3 \,:\, \xi_1 = 52,\ 47 \le \xi_2 \le 55,\ \xi_2 - 5 \le \xi_3 \le \xi_2 + 4\right\}.$$

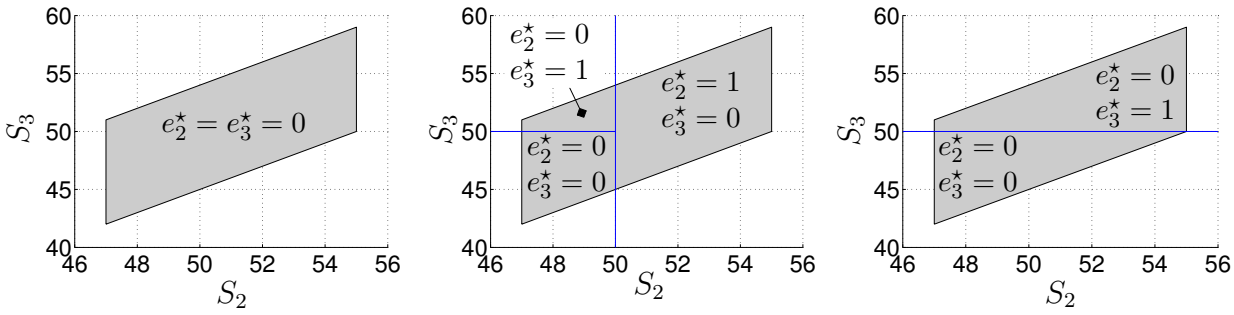


Figure 3.1: Companion figure for Example 3.1. Each sub-figure illustrates a strategy optimal in the holder's problem (all strategies satisfy $e_1^\star = 1$). No buyer with a strictly increasing utility function will exercise according to (a); The strategy maximizing the expected payoff from the swing option is shown on Figure (b); The strategy which yields the highest price for the writer is shown on Figure (c), see Section 3.2.4 and Example 3.2.

We want to determine the holder's price for a swing option with strike price $K = 50$, *per-period limits* $\underline{e}_t = 0$ *and* $\overline{e}_t = 1$ *and cumulative energy limits* $\underline{c} = 0$ *and* $\overline{c} = 2$. *Since the only basic security is the risk-free asset, the set of optimal exercise strategies in the holder's problem is equal to the set of optimal strategies of the problem*

$$\max_{e \in \mathcal{E}} \min_{\xi \in \Xi} \sum_{t \in \mathbb{T}} e_t(\xi^t)(S_t(\xi_t) - K). \tag{3.2}$$

Furthermore, the optimal objective value of (3.2) *yields the holder's price for the swing option. Since* $\underline{e}_t = 0 \ \forall t \in \mathbb{T}$, *independently of the choice of exercise strategy* $e \in \mathcal{E}$, *the realization of the spot which yields the worst-case objective is given by* $S_1 = 52$, $S_2 = 47$ *and* $S_3 = 42$, *i.e.,* $\xi = (52, 47, 42)$. *Thus every strategy* $e^\star$ *optimal in problem* (3.2) *satisfies* $e_1^\star = 1$. *Also, any feasible strategy satisfying* $e_2(\xi^2) = 0 \ \forall \xi \in \Xi : S_2(\xi_2) < 50$ *and* $e_3(\xi^3) = 0 \ \forall \xi \in \Xi : S_3(\xi_3) < 50$ *is optimal. Figure 3.1 illustrates the projection of the support of the spot price onto the hyperplane* $S_1 = 52$ *and three strategies, all optimal for the holder's problem (they all yield a holder's price for the swing option of* $H = 2$*). No buyer with a strictly increasing utility function will exercise according to the strategy from Figure 3.1(a). In this case, the strategy maximizing the expected payoff from the swing option is the strategy from Figure 3.1(b).*

An optimal exercise strategy in the holder's problem is sometimes referred to as a "ruthless" strategy. Some option holders may decide to deviate from a ruthless strategy if they face

obligations to deliver energy to third parties.

3.2.4 Writer's Problem

We now investigate the writer's price for the swing option and start by making the following assumption regarding the knowledge of the writer about the holder's exercise strategy.

(K) The swing option holder has communicated her exercise strategy $e^\star$ to the writer. Furthermore, $e^\star$ is a ruthless strategy, see Section 3.2.3.

We remark that assumption (K) is realistic in electricity markets where swing options are typically traded over-the-counter (i.e., off-exchange).

Under assumption (K), the writer's price constitutes the optimal value of the following optimization problem.

$$\begin{aligned} W \quad = \quad & \min_{x,u} \quad -C_0^{x,u} \\ & \text{s.t.} \quad (x,u) \in \mathcal{S}^{\mathrm{h}}(-e^\star) \end{aligned} \tag{$\mathcal{W}$}$$

The decision variables of this model are the elements of the portfolio strategy pair (x, u). As before, these decision variables are decision rules. Unlike the holder, the writer cannot decide on the exercise strategy e of the swing option. Instead, the holder's optimal exercise strategy $e^\star$ is a (uncertainty-affected) parameter in the writer's problem. If the option holder requests $e_t^\star(\xi^t)$ units of energy at price K, the writer has to provide this energy quantity by buying it on the spot market at price $S_t(\xi_t)$. The option writer thus receives a dividend of size $-e_t^\star(\xi^t)(S_t(\xi_t) - K)$ at time t (which is typically negative).

The writer's problem differs from the holder's problem only in the objective function and the self-financing constraint implied by the requirement that $(x, u) \in \mathcal{S}^{\mathrm{h}}(-e^\star)$. Instead of maximizing the loan that can be refinanced, the writer minimizes the initial value of a portfolio of basic securities that covers all obligations arising from the swing option. The self-financing constraint in the writer's problem stipulates that the costs $P_t(\xi_t)^\top u_t(\xi^t)$ to rebalance the portfolio at the beginning of period t are recovered from the dividends $-e_t^\star(\xi^t)(S_t(\xi_t) - K)$ arising from the

swing option and the dividends $d_t(\xi_t)^\top x_t(\xi^t)$ received from the basic securities. The dividends of the swing option cannot be influenced by the writer.

The following proposition establishes the correctness of the writer's problem.

Proposition 3.2 *Suppose that assumption (K) holds. Let $(x^\star, u^\star)$ denote an optimal solution to the writer's problem with objective value W.*

(a) If the swing option trades at price $V > W$, then the option writer can make arbitrage profits.

(b) If the swing option trades at price $V < W$, no arbitrage can be made by selling the swing option.

Proof Assume first that the option trades at a price $V > W$. Then, sell the swing option at price V and take a loan of amount $-W = C_0^{x^\star,u^\star}$, which yields an immediate payoff $V - W > 0$. Since $(x^\star, u^\star) \in \mathcal{S}^{\mathrm{h}}(-e^\star)$, the loan $-W$ can be repaid almost surely by following the portfolio strategy $(x^\star, u^\star)$ while covering all obligations arising from the swing option with exercise strategy $e^\star$. Furthermore, the self-financing constraints implied by the requirement $(x, u) \in \mathcal{S}^{\mathrm{h}}(-e^\star)$ guarantee that, with probability one, the dividend process arising from this strategy will be nonnegative. Thus, we have constructed an arbitrage.

Assume now that the swing option trades at a price $V < W$ and that arbitrage profits can be made by selling the swing option. Then, there is a portfolio strategy $(x, u) \in \mathcal{S}^{\mathrm{h}}(-e^\star)$ which satisfies $C_0^{x,u} + V \geq 0$, see Definition 3.3, i.e., $-C_0^{x,u} < W$. Such a portfolio would correspond to a feasible solution (x, u) of the writer's problem with an objective value $-C_0^{x,u} < W$. This contradicts the optimality of $(x^\star, u^\star)$. Thus, no arbitrage profits can be made by selling the swing option at a price $V < W$. □

Remark 3.5 *If the option trades at price $V = W$, it is possible for the option writer to make arbitrage profits if and only if there exists a pair (x, u) feasible in $\mathcal{W}$ such that for some*

$t \in \mathbb{T} \cup \{T+1\}$, the probability that the self-financing constraints implied by $(x,u) \in \mathcal{S}^{\mathrm{h}}(-e^\star)$ hold with strict inequality is strictly positive, see Definition 3.3.

While assumption (K) is reasonable, there may be cases where the holder may not be willing or able to commit to a ruthless exercise strategy, due to e.g., limited information, small amount of time in which she has to take a decision, expensive computational requirements of calculating an optimal solution to $\mathcal{H}$. This is closely related to the idea of "bounded rationality" coined by Simons [1955]. In such cases, it is necessary for the writer to price the swing option independently of precise a-priori knowledge about the strategy of the holder. We thus introduce the following, relaxed version of assumption (K).

($\widehat{\mathrm{K}}$) The option writer only knows that the exercise strategy e of the holder lies in the set

$$\mathcal{E}_\eta := \left\{ e \in \mathcal{E} \ : \ \exists (x,u) \in \mathcal{S}^{\mathrm{h}}(e) \ : \ C_0^{x,u} \geq H - \eta \right\},$$

for some fixed $\eta \in [0, +\infty]$. Furthermore, the holder may exercise according to any strategy $e \in \mathcal{E}_\eta$.

Under assumption ($\widehat{\mathrm{K}}$), the writer must be hedged against any exercise strategy for the swing option achieving a "satisfactory price" $H - \eta$ in the holder's problem. The writer's price thus constitutes the optimal value of the following optimization problem.

$$\begin{array}{rcll} W_\eta & = & \min\limits_{x,u} & -C_0^{x,u} \\ & & \text{s.t.} & (x,u) \in \mathcal{S}(p^e) \quad \forall e \in \mathcal{E}_\eta, \end{array} \tag{$\mathcal{W}_\eta$}$$

where the process $p^e := (p_1^e, \ldots, p_T^e)$ is defined through

$$p_t^e(\xi^t) := \begin{cases} \displaystyle\sum_{t \in \mathbb{T}} e_t(\xi^t)(K - S_t(\xi_t)) & \text{if } t = T, \\ 0 & \text{else.} \end{cases}$$

In problem $\mathcal{W}_\eta$, the swing option dividends received by the writer of the swing option are aggregated to stage T. The reason for this aggregation is the fact that the writer will not

use the same portfolio strategy independently of the exercise strategy ultimately implemented by the holder. Instead, the specific quantities of the risk-free asset traded in the portfolio strategy chosen by the writer will depend on the exercise strategy $e \in \mathcal{E}_\eta$ selected by the holder: depending on e, the writer may have to borrow (or lend) cash at some $t \in \mathbb{T}\backslash\{T\}$ to cover obligations arising from the specific strategy e chosen by the holder (or to transfer dividends across time-periods); any loans taken will be repaid with certainty before the end of period T.

We remark that the parameter η in assumption $(\widehat{\mathrm{K}})$ quantifies the deviation from optimality tolerated by the swing option holder. The higher the value of η, the more robust the hedging strategy of the writer. Specifically, we note that $\mathcal{E}_{+\infty} = \mathcal{E}$, i.e., for $\eta = +\infty$, the writer will be hedged against *any* exercise strategy, while in the setting $\eta = 0$, the writer will be hedged against *any ruthless* strategy of the holder. In the remainder of this chapter, we will refer to problem $\mathcal{W}_{+\infty}$ as the robust writer's problem.

We have the following, modified version of Proposition 3.2 which holds under assumption $(\widehat{\mathrm{K}})$.

Proposition 3.3 *Suppose that assumption $(\widehat{K})$ holds. Let $(x^\star, u^\star)$ denote an optimal solution to the writer's problem with objective value W_η.*

(a) If the swing option trades at price $V > W_\eta$, then the option writer can make arbitrage profits.

(b) If the swing option trades at price $V < W_\eta$, no arbitrage can be made by selling the swing option.

Proof Assume first that the option trades at a price $V > W_\eta$. For any $e \in \mathcal{E}_\eta$, there exists a portfolio strategy pair $(x^e, u^e) \in \mathcal{S}^{\mathrm{h}}(-e)$ satisfying $C_0^{x^e,u^e} = C_0^{x^\star,u^\star}$, which can be obtained from $(x^\star, u^\star)$ by trading in the risk-free asset. Then, sell the swing option at price V and take a loan of amount $-W_\eta = C_0^{x^e,u^e}$, which yields an immediate payoff $V - W_\eta > 0$. Since $(x^e, u^e) \in \mathcal{S}(-e)$, the loan $-W_\eta$ can be repaid almost surely by following the portfolio strategy (x^e, u^e) while covering all obligations arising from the swing option with exercise strategy e.

Furthermore, the self-financing constraints implied by $(x^e, u^e) \in \mathcal{S}^{\mathrm{h}}(-e)$ guarantee that, with probability one, the dividend process arising from this strategy will be nonnegative. Thus, we have constructed an arbitrage.

Assume now that the swing option trades at a price $V < W_\eta$ and that arbitrage profits can be made by selling the swing option. Then, for all $e \in \mathcal{E}_\eta$, there exists a portfolio strategy $(x^e, u^e) \in \mathcal{S}^{\mathrm{h}}(-e)$ such that $C_0^{x^e,u^e} + V \geq 0$, see Definition 3.3, i.e., $-C_0^{x^e,u^e} < W_\eta$. This in turn implies that there is a portfolio strategy (x, u) such that $-C_0^{x,u} < W_\eta$ and $(x, u) \in \mathcal{S}(p^e)$ $\forall e \in \mathcal{E}_\eta$. Such a portfolio would correspond to a feasible solution (x, u) of the writer's problem with an objective value $-C_0^{x,u} < W_\eta$. This contradicts the optimality of $(x^\star, u^\star)$. Thus, no arbitrage profits can be made by selling the swing option at a price $V < W_\eta$. □

We remark that for any $\eta_1, \eta_2 \in [0, +\infty]$ such that $\eta_1 < \eta_2$, it holds that

$$W \leq W_{\eta_1} \leq W_{\eta_2}.$$

We now discuss how for any $\eta \in [0, +\infty]$, an optimal solution to the writer's problem $\mathcal{W}_\eta$ can be obtained by bringing the problem to a form which, in the case $\eta = +\infty$, can be solved efficiently by using the approach discussed in Section 3.4.

Proposition 3.4 *For $\eta \in [0, +\infty]$, the writer's price W_η can be obtained as the optimal objective value of the problem*

$$\begin{aligned} \min_{x,u} \quad & -C_0^{x,u} \\ \text{s.t.} \quad & (x, u) \in \mathcal{S}(p), \end{aligned} \tag{3.3}$$

where the process $p := (p_1, \ldots, p_T)$ is defined through

$$p_t(\xi^t) := \begin{cases} -\max\limits_{e \in \mathcal{E}_\eta} \sum\limits_{t \in \mathbb{T}} e_t(\xi^t)(S_t(\xi_t) - K) & \text{if } t{=}T, \\ 0 & \text{else.} \end{cases}$$

Remark 3.6 *We note that the subproblem in the definition of p corresponds to a multi-stage robust problem over the variables $(e, x^{\mathrm{b}}, u^{\mathrm{b}})$, where $(x^{\mathrm{b}}, u^{\mathrm{b}})$ denotes the portfolio strategy used*

by the holder to hedge the dividend stream of the swing option with exercise strategy e at a satisfactory price $H - \eta$. If $\eta = +\infty$, this subproblem can be formulated as a deterministic problem.

Remark 3.7 *Proposition 3.4 implies that in order for the writer to be hedged against all (non-anticipative) exercise strategies $e^\star \in \mathcal{E}_\eta$, for $\eta \in [0, +\infty]$, it is sufficient for her to be hedged against a single* anticipative *exercise strategy $\hat{e}$ possessing the property that for any realization of $\xi \in \Xi$, there exists a non-anticipative strategy $e \in \mathcal{E}_\eta$ such that $e(\xi) = \hat{e}(\xi)$. Thus, the determination of the optimal objective value W_η of problem $\mathcal{W}_\eta$ reduces to finding an anticipative strategy $\hat{e}$ maximizing the dividend stream of the swing option for each realization of $\xi \in \Xi$ and solving a problem of the form $\mathcal{W}$, see Example 3.2.*

Proof of Proposition 3.4 The validity of the assertion can be established by noting that $\mathcal{W}_\eta$ and (3.3) are equivalent since, for any fixed $\xi \in \Xi$,

$$\begin{aligned} & P_T(\xi_T)^\top u_T(\xi^T) \;\leq\; \sum_{t\in\mathbb{T}} e_t(\xi^t)(K - S_t(\xi_t)) + d_T(\xi_t)^\top x_T(\xi^T) \quad \forall e \in \mathcal{E}_\eta \\ \Leftrightarrow \quad & P_T(\xi_T)^\top u_T(\xi^T) \;\leq\; -\max_{e\in\mathcal{E}_\eta} \sum_{t\in\mathbb{T}} e_t(\xi^t)(S_t(\xi_t) - K) + d_T(\xi_t)^\top x_T(\xi^T). \end{aligned}$$

□

In Section 3.4, we will propose a technique to solve $\mathcal{W}$ and $\mathcal{W}_{+\infty}$.

3.2.5 No-Arbitrage Interval

We now show that the holder's and the writer's prices indeed form an interval of no-arbitrage prices.

Proposition 3.5 *Let $e^\star$ denote the an optimal exercise strategy for the holder and let H denote the holder's price.*

(a) Suppose that assumption (K) holds. The optimal value W of problem $\mathcal{W}$ satisfies $H \leq W$. Equality holds if and only if the market is complete relative to the optimal dividend stream $\{e_t^\star(S_t - K)\}_{t\in\mathbb{T}}$ of the swing option, in which case $V = H = W$ is an arbitrage-free price for the option.

(b) Suppose that assumption $(\widehat{K})$ holds. The optimal value W_η of problem $\mathcal{W}_\eta$ satisfies $H \leq W_\eta$.

Remark 3.8 *We remark that the upper end-point of the no-arbitrage interval (the writer's price) depends on the assumption ((K) or $(\widehat{K})$) about writer's knowledge of the exercise strategy of the holder. Furthermore, even under assumption (K), the writer's price for the swing option may change depending on the exercise strategy $e^\star$ of the holder, see Example 3.2.*

Proof of Proposition 3.5 (a) We first show that $W - H \geq 0$. To this end, we transform the holder's problem into a minimization problem and determine $W - H$ by solving the holder's and the writer's problem simultaneously:

$$\begin{aligned} \min_{e^{\mathrm{h}},x^{\mathrm{h}},u^{\mathrm{h}},x^{\mathrm{w}},u^{\mathrm{w}}} \quad & -C_0^{x^{\mathrm{h}},u^{\mathrm{h}}} - C_0^{x^{\mathrm{w}},u^{\mathrm{w}}} \\ \text{s.t.} \quad & e^{\mathrm{h}} \in \mathcal{E} \\ & (x^{\mathrm{h}}, u^{\mathrm{h}}) \in \mathcal{S}^{\mathrm{h}}(e^{\mathrm{h}}) \\ & (x^{\mathrm{w}}, u^{\mathrm{w}}) \in \mathcal{S}^{\mathrm{h}}(-e^\star) \end{aligned} \tag{3.4}$$

In this problem, $(e^{\mathrm{h}}, x^{\mathrm{h}}, u^{\mathrm{h}})$ and $(x^{\mathrm{w}}, u^{\mathrm{w}})$ represent the decisions of the holder and the writer, respectively. We obtain a lower bound on the optimal value of this problem (and hence a lower bound on $W - H$) by aggregating (adding) the pairs of associated constraints and substituting $x = x^{\mathrm{h}} + x^{\mathrm{w}}$ as well as $u = u^{\mathrm{h}} + u^{\mathrm{w}}$:

$$\begin{aligned} \min_{e^{\mathrm{h}},x,u} \quad & -C_0^{x,u} \\ \text{s.t.} \quad & e^{\mathrm{h}} \in \mathcal{E} \\ & (x, u) \in \mathcal{S}^{\mathrm{h}}(e^{\mathrm{h}} - e^\star) \end{aligned}$$

From assumption (K), $e^{\mathrm{h}} = e^{\star}$ and the above problem simplifies to

$$\begin{aligned} \min_{e^{\mathrm{h}},x,u} \quad & -C_0^{x,u} \\ \text{s.t.} \quad & e^{\mathrm{h}} \in \mathcal{E} \\ & (x,u) \in \mathcal{S}^{\mathrm{h}}(0) \end{aligned} \tag{3.5}$$

This problem determines the minimum initial value of a self-financing portfolio (x, u) with non-negative terminal value, see Definition 3.1. Since the market of basic securities is arbitrage-free, see assumption (M1), the optimal value of this problem must be nonnegative, see Definition 3.3. We therefore conclude that $W - H$ is nonnegative as well.

Next, we show that if $H = W$, then the market of basic securities is complete relative to the optimal dividend stream of the swing option. For this purpose, we start by demonstrating that under assumption (M1) and if $H = W$ then no arbitrage profits can be made by buying or selling the swing option at $V = H = W$. Suppose that $H = W$ and that buying the swing option at price H *leads* to arbitrage opportunities. Then, under assumption (K), there exists a solution $(e^{\star}, x^{\mathrm{h}}, u^{\mathrm{h}}, x^{\mathrm{w}}, u^{\mathrm{w}})$ to (3.4) with an objective value equal to zero and such that for some $t \in \mathbb{T} \cup \{T+1\}$, one of the self-financing constraints implied by $(x^{\mathrm{h}}, u^{\mathrm{h}}) \in \mathcal{S}^{\mathrm{h}}(e^{\star})$ hold with strict inequality on a set of positive probability. This solution corresponds to a solution (x, u) to (3.5) with a zero objective value and such that either the terminal portfolio value constraint or one of the self-financing constraints implied by $(x, u) \in \mathcal{S}^{\mathrm{h}}(0)$ hold with strict inequality on a set of positive probability, which implies the presence of arbitrage opportunities in the market of basic securities. This contradicts our assumption (M1) and thus if $H = W$, no arbitrage opportunities can be made by buying the swing option at price H. A similar argumentation can be used to show that if $H = W$, no arbitrage opportunities can be made by selling the swing option at price W.

Since $H = W$ there exists a solution $(e^{\star}, x^{\mathrm{h}}, u^{\mathrm{h}}, x^{\mathrm{w}}, u^{\mathrm{w}})$ to (3.4) with an objective value equal to zero. Since no arbitrage profits can be made by buying or selling the swing option at price $V = H = W$, the self-financing constraints and the terminal portfolio value constraints in (3.4) implied by $(x^{\mathrm{w}}, u^{\mathrm{w}}) \in \mathcal{S}^{\mathrm{h}}(-e^{\star})$ must all hold with equality with probability one at this solution.

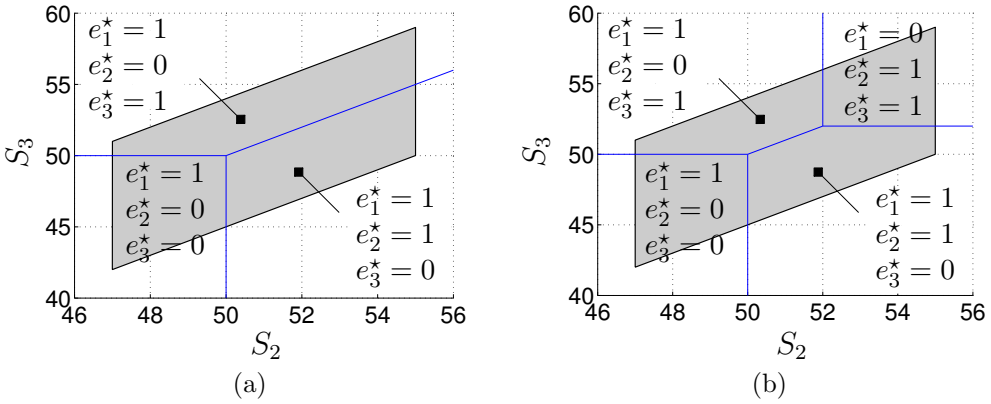


Figure 3.2: Companion figure for Example 3.2. Swing option strategy against which the writer can hedge in order to be hedged against (a) any ruthless strategy ($\eta = 0$) and (b) any admissible strategy ($\eta = \infty$).

Otherwise, positive profits could be made on a set of positive probability at zero initial cost. Thus, there exists a portfolio strategy, namely $(x^{\mathrm{w}}, u^{\mathrm{w}})$, that replicates the optimal dividend stream of the swing option with probability 1, implying that the market is complete relative to the dividend stream of the swing option, see Definition 3.2. Furthermore, the (unique) price of the swing option is given by $-C_0^{x^{\mathrm{w}},u^{\mathrm{w}}}$.

We now show that $H = W$ if the market of basic securities is complete relative to the optimal dividend stream of the swing option. Since the market is complete relative to the optimal dividend stream of the swing option, there is a portfolio strategy $(x^\star, u^\star)$ satisfying

$$\left.\begin{array}{rcll} C_t^{x^\star,u^\star}(\xi^t) & = & e_t^\star(\xi^t)(S_t(\xi_t) - K) & \forall t \in \mathbb{T} \\ C_{T+1}^{x^\star,u^\star}(\xi^{T+1}) & = & 0 & \end{array}\right\} \quad \mathbb{P}\text{-a.s.},$$

see Definition 3.2. This portfolio strategy is feasible in $\mathcal{W}$ with an objective value of $-C_0^{x^\star,u^\star}$. Furthermore, by construction, the strategy $(e^\star, -x^\star, -u^\star)$ is feasible in $\mathcal{H}$ with an objective value $C_0^{-x^\star,-u^\star} = -C_0^{x^\star,u^\star}$. We have thus constructed two solutions feasible in $\mathcal{W}$ and $\mathcal{H}$ respectively which have the same objective value, thus $H \geq W$. Equality of H and W now follows from the first part of the proof.

(b) This is a direct consequence of (a) and the fact that $W \leq W_\eta$ for all $\eta \in [0, +\infty]$. □

We now revisit Example 3.1 from the writer's viewpoint.

Example 3.2 *Consider the market and swing option from Example 3.1. We want to determine the writer's price for the swing option for the cases $\eta = 0$ and $\eta = \infty$. From Proposition 3.4, these can be obtained by solving the writer's problem for some suitably chosen anticipative exercise strategies. An anticipative strategy against which the writer can hedge in each of these two cases is illustrated on Figure 3.2. Since the market solely consists of a spot exchange and the risk-free asset, the writer's price for the swing option under a fixed (anticipative) strategy $e^\star$ can be expressed as*

$$\max_{\xi \in \Xi} \sum_{t \in \mathbb{T}} e_t^\star(\xi)(S_t(\xi_t) - K). \tag{3.6}$$

Furthermore, the optimal objective value of (3.6) *yields the writer's price for the swing option. From Figure 3.2, we obtain $W_0 = 11$ and $W_\infty = 14$. We now investigate the writer's prices for the swing option if the holder exercises according to the ruthless strategies in Figures 3.1(a)-3.1(c) and communicates his strategy to the writer. The writer's price is respectively given by 2, 7 and 11. Recall that the holder's price is $H = 2$. Thus, as expected, the market is complete relative to the strategy from Figure 3.1(a). Furthermore, the strategy from Figure 3.1(c) yields, as expected, see Proposition 3.4, a price of $W_0 = 11$ (it corresponds to one of the worst-case non-anticipative ruthless strategies from the writer's viewpoint).*

3.3 The Optimal Exercise Strategy

We now investigate the structure of the optimal swing option exercise strategy. For this purpose, we introduce the following definition.

Definition 3.4 (Bang-bang control) *An exercise strategy $e \in \mathcal{E}$ is called a bang-bang control if it satifies*

$$e_t(\xi^t) \in \left\{ \max\left(\underline{e}_t, \underline{c} - \textstyle\sum_{\tau=t+1}^{T} \overline{e}_\tau \right), \min\left(\overline{e}_t, \overline{c} - \textstyle\sum_{\tau=t+1}^{T} \underline{e}_\tau \right) \right\} \quad \mathbb{P}\text{-a.s.} \quad \forall t \in \mathbb{T}.$$

Bardou et al. [2010] have shown that under mild assumptions on the contractual specifications of the swing option, there always exists an optimal bang-bang exercise strategy that solves the

problem

$$\max_{e\in\mathcal{E}} \mathbb{E}^{\mathbb{Q}}\left[\sum_{t\in\mathbb{T}} e_t(\xi^t)(S_t(\xi_t)-K)\right],$$

where $\mathbb{Q}$ is a fixed distribution and $\mathbb{E}^{\mathbb{Q}}(\cdot)$ denotes the expectation with respect to $\mathbb{Q}$. At the same time, the second fundamental theorem of asset pricing states that a market is complete if and only if there is a *unique* risk-neutral measure $\mathbb{Q}$ (i.e., if and only if there is a unique measure under which the prices of all traded assets are martingales). Then, the risk-neutral expectation gives the unique arbitrage-free price of the swing option. Thus, the result of Bardou et al. can be stated as follows: under mild contractual constraints, in a *complete* market, there always exists an optimal bang-bang exercise strategy.

We now demonstrate that the same cannot be said if the market is *incomplete*.

Proposition 3.6 *In an incomplete market, there does not necessarily exist an optimal bang-bang exercise strategy for the swing option, even if* $\underline{e}_t = 0$, $\overline{e}_t = 1\ \forall t \in \mathbb{T}$, $\underline{c} = 0$ *and* $\overline{c} \in \mathbb{N}$.

Proof We proceed by means of an example. Suppose $\mathbb{T} = \{1, 2, 3\}$ and consider an electricity market that consists of a spot exchange and the risk-free asset. In addition, assume that $S_t(\xi_t) = \xi_t$ and that ξ is uniformly distributed on

$$\Xi := \left\{\xi \in \mathbb{R}^3 \ :\ \xi_1 = 52,\ 60 \le \xi_2 \le 61.5,\ 233 - 3\xi_2 \le \xi_3 \le \xi_2 + 3\right\}. \tag{3.7}$$

Figure 3.3(a) illustrates the projection of the support of the spot price onto the hyperplane $S_1 = 52$. We want to determine the holder's price for a swing option with strike price $K = 50$, per-period limits $\underline{e}_t = 0$ and $\overline{e}_t = 1$ and cumulative energy limits $\underline{c} = 0$ and $\overline{c} = 2$. We remark that this swing option satisfies the contractual specifications from Bardou et al. [2010].

Since the only basic security is the risk-free asset, the set of optimal exercise strategies in the holder's problem is equal to the set of strategies maximizing the worst-case payoff of the swing-option

$$\max_{e\in\mathcal{E}} \min_{\xi\in\Xi} \sum_{t\in\mathbb{T}} e_t(\xi^t)(S_t(\xi_t)-K). \tag{3.8}$$

Furthermore, the optimal objective value of (3.8) yields the holder's price for the swing option.

For $\hat{e}_1 \in [0,1]$, define the strategy $\tilde{e}$ through

$$\tilde{e}_1(\xi^1) = \hat{e}_1, \quad \tilde{e}_2(\xi^2) = 1 \quad \text{and} \quad \tilde{e}_3(\xi^3) = \begin{cases} 1 - \hat{e}_1 & \text{if } S_3(\xi_3) \geq K \\ 0 & \text{else.} \end{cases} \tag{3.9}$$

From (3.7), see also Figure 3.3(a), we have

$$S_2(\xi_2) > S_1(\xi_1) > 50 \quad \forall \xi \in \Xi \quad \text{and} \quad S_2(\xi_2) > \min_{(\xi_1,\xi_2,\tilde{\xi}_3)\in\Xi} S_3(\tilde{\xi}_3) \quad \forall \xi \in \Xi.$$

Thus, there must exist at least one strategy of the form (3.9) with $\hat{e}_1 \in [0,1]$ solving (3.8). Furthermore, if there exists an optimal bang-bang strategy for the holder's problem, $\tilde{e}$ must be optimal in (3.8) with $\hat{e}_1 = 1$. We thus investigate the behavior of the function

$$H(\hat{e}_1) = \min_{\xi\in\Xi} \ \hat{e}_1(S_1(\xi_1) - K) + S_2(\xi_2) - K + (1 - \hat{e}_1)\max\{(S_3(\xi_3) - K), 0\}$$

as $\hat{e}_1$ varies in the range $[0,1]$. Since $\min_{\xi_1,\xi_2,\tilde{\xi}_3} S_3(\tilde{\xi}_3) < 50 \ \forall \xi \in \Xi : S_2(\xi_2) > 61$, we have

$$\begin{aligned} H(\hat{e}_1) &= 2\hat{e}_1 + \min\left\{ \min_{S_2\in[60,61]} S_2 - 50 + (1-\hat{e}_1)(233 - 3S_2 - 50), \min_{S_2\in[61,61.5]} S_2 - 50 \right\} \\ &= 2\hat{e}_1 + \min\left\{(13 - 3\hat{e}_1)\mathbb{1}(\hat{e}_1 \geq 2/3) + 11\mathbb{1}(\hat{e}_1 < 2/3), 11\right\} \\ &= \min\left\{13 - \hat{e}_1, 2\hat{e}_1 + 11\right\}. \end{aligned}$$

A plot for $H(\hat{e}_1)$ is provided on Figure 3.3(b). The holder's price for the option is given by $\max_{\hat{e}_1\in[0,1]} H(\hat{e}_1) = 12\frac{1}{3}$ and is attained at $\hat{e}_1 = 2/3$ only, i.e., there does not exist an optimal bang-bang exercise strategy for this problem (the best bang-bang strategy yields a price of 12). The optimal exercise strategy obtained by maximizing $H(\hat{e}_1)$ is given by $e^\star(\xi) = (\frac{2}{3}, 1, \frac{1}{3}\mathbb{1}(S_3(\xi_3) \geq 50))$ and is illustrated on Figure 3.3(c). □

We now show that even in the case of swing options which permit violation of the cumulative energy limits, there does not necessarily exist an optimal bang-bang exercise strategy.

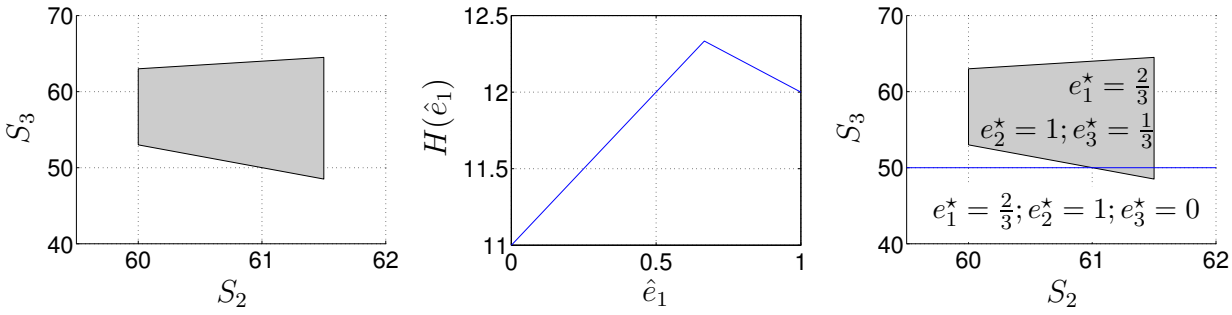


Figure 3.3: Companion figure for the proof of Proposition 3.6. The projection of the support of the spot price onto the hyperplane $S_1 = 52$ is shown on Figure (a); (b) illustrates the function $H(\hat{e}_1)$ and (c) illustrates the optimal exercise strategy $e^\star$ obtained by maximizing $H(\hat{e}_1)$.

Proposition 3.7 *In an incomplete market, there does not necessarily exist an optimal bang-bang exercise strategy for a swing option which permits violation of the cumulative energy limits (with penalties incurred for each unit of shortfall or exceedance), even if* $\underline{e}_t = 0$, $\overline{e}_t = 1\ \forall t \in \mathbb{T}$, $\underline{c}$, $\overline{c} \in \mathbb{N}$.

Proof As in the proof of Proposition 3.6, we proceed by means of an example. Suppose $\mathbb{T} = \{1, 2, 3\}$ and consider an electricity market that consists of a spot exchange and the risk-free asset. In addition, assume that $S_t(\xi_t) = \xi_t$ and that ξ is uniformly distributed on

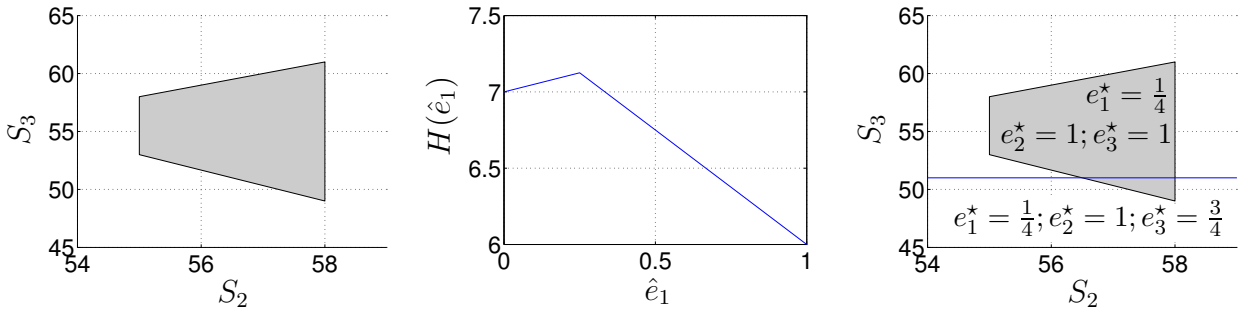
$$\Xi := \left\{\xi \in \mathbb{R}^3 \ : \ \xi_1 = 49.5,\ 55 \leq \xi_2 \leq 58,\ 126\tfrac{1}{3} - 1\tfrac{1}{3}\xi_2 \leq \xi_3 \leq \xi_2 + 3\right\}. \tag{3.10}$$

Figure 3.4(a) illustrates the projection of the support of the spot price onto the hyperplane $S_1 = 49.5$. We want to determine the holder's price for a swing option with strike price $K = 50$, per-period limits $\underline{e}_t = 0$ and $\overline{e}_t = 1$ and cumulative energy limits $\underline{c} = 2$ and $\overline{c} = 2$. A cost of 1 is incurred for each unit of shortfall or exceedance.

Since the only basic security is the risk-free asset, the set of optimal exercise strategies in the holder's problem is equal to the set of strategies maximizing the worst-case payoff of the swing-option

$$\max_{e \in \mathcal{E}} \min_{\xi \in \Xi} \sum_{t \in \mathbb{T}} e_t(\xi^t)(S_t(\xi_t) - K) - \left| 2 - \sum_{t \in \mathbb{T}} e_t(\xi^t) \right|. \tag{3.11}$$

The last term in (3.11) models the penalties incurred for over- or under-shooting the cumulative



(a) (b) (c)

Figure 3.4: Companion figure for the proof of Proposition 3.7. The projection of the support of the spot price onto the hyperplane $S_1 = 49.5$ is shown on Figure (a); (b) illustrates the function $H(\hat{e}_1)$ and (c) illustrates the optimal exercise strategy $e^\star$ obtained by maximizing $H(\hat{e}_1)$.

energy limits. We remark that the optimal objective value of (3.11) yields the holder's price for the swing option.

For $\hat{e}_1 \in [0,1]$, define the strategy $\tilde{e}$ through

$$\tilde{e}_1(\xi^1) = \hat{e}_1, \quad \tilde{e}_2(\xi^2) = 1 \quad \text{and} \quad \tilde{e}_3(\xi^3) = \begin{cases} 1 & \text{if } S_3(\xi_3) \geq K+1 \\ 1-\hat{e}_1 & \text{else.} \end{cases} \tag{3.12}$$

From (3.10), see also Figure 3.4(a), we have

$$S_2(\xi_2) > K+1 > S_1(\xi_1) \quad \forall \xi \in \Xi \quad \text{and} \quad S_2(\xi_2) > \min_{(\xi_1,\xi_2,\tilde{\xi}_3)\in\Xi} S_3(\tilde{\xi}_3) > K-1 \quad \forall \xi \in \Xi.$$

Thus, there must exist at least one strategy of the form (3.12) with $\hat{e}_1 \in [0,1]$ optimal in (3.11). Furthermore, if there exists a bang-bang strategy optimal for the holder's problem, $\tilde{e}$ must be optimal in (3.11) with $\hat{e}_1 = 1$. We thus investigate the behavior of the function

$$\begin{aligned} H(\hat{e}_1) \;=\; \min_{\xi\in\Xi} \{&\hat{e}_1(S_1(\xi_1)-K) + S_2(\xi_2) - K + [(S_3(\xi_3)-K) - \hat{e}_1]\mathbb{1}(S_3(\xi_3) \geq K+1) \\ &+(1-\hat{e}_1)(S_3(\xi_3)-K)\mathbb{1}(S_3(\xi_3) < K+1)\} \end{aligned}$$

as $\hat{e}_1$ varies in the range $[0,1]$. Since $\min_{\xi_1,\xi_2,\tilde{\xi}_3} S_3(\tilde{\xi}_3) < K+1 \;\forall \xi \in \Xi : S_2(\xi_2) > 56\frac{1}{2}$, and

$S_3(\xi) \geq K + 1 \; \forall \xi : S_2(\xi_2) \leq 56\frac{1}{2}$, we have

$$\begin{aligned} H(\hat{e}_1) &= -\frac{1}{2}\hat{e}_1 + \min\left\{ \min_{S_2 \in [55,56.5]} S_2 - 50 + 126\tfrac{1}{3} - 1\tfrac{1}{3}S_2 - 50 - \hat{e}_1, \right. \\ &\qquad \left. \min_{S_2 \in [56.5,58]} S_2 - 50 + (1-\hat{e}_1)(126\tfrac{1}{3} - 1\tfrac{1}{3}S_2 - 50) \right\} \\ &= -\frac{1}{2}\hat{e}_1 + \min\left\{ \min_{S_2 \in [55,56.5]} -\frac{S_2}{3} + 26\tfrac{1}{3} - \hat{e}_1, \min_{S_2 \in [56.5,58]} S_2\left(1\tfrac{1}{3}\hat{e}_1 - \frac{1}{3}\right) - 76\tfrac{1}{3}\hat{e}_1 + 26\tfrac{1}{3} \right\} \\ &= -\frac{1}{2}\hat{e}_1 + \min\left\{7\tfrac{1}{2} - \hat{e}_1, (7\tfrac{1}{2} - \hat{e}_1)\mathbb{1}(\hat{e}_1 \geq 1/4) + (7 + \hat{e}_1)\mathbb{1}(\hat{e}_1 < 1/4)\right\} \\ &= \min\left\{7\tfrac{1}{2} - 1\tfrac{1}{2}\hat{e}_1, 7 + \frac{1}{2}\hat{e}_1\right\}. \end{aligned}$$

A plot for $H(\hat{e}_1)$ is provided on Figure 3.4(b). The holder's price for the option is given by $\max_{\hat{e}_1 \in [0,1]} H(\hat{e}_1) = 7\frac{1}{8}$ and is attained at $\hat{e}_1 = 1/4$ only, i.e., there does not exist an optimal bang-bang exercise strategy for this problem (the best bang-bang strategy yields a price of 7). The optimal exercise strategy obtained by maximizing $H(\hat{e}_1)$ is given by

$$e^\star(\xi) = \left(\frac{1}{4}, 1, \mathbb{1}(S_3(\xi_3) \geq 51) + \frac{3}{4}\mathbb{1}(S_3(\xi_3) < 51)\right)$$

and is illustrated on Figure 3.4(c). □

3.4 Tractable Reformulation

Unfortunately, without suitable approximations, the hedging problems $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{W}_{+\infty}$ are computationally intractable for three reasons: (*i*) they involve a large number of time-periods, (*ii*) they optimize over function spaces and (*iii*) they involve infinitely many constraints.

In principle these hedging problems could be solved with the approximation discussed in Chapter 2. Unfortunately, typical swing options have delivery periods of up to one year and are exercisable in intervals of 15 min up to one day. Therefore, the planning horizon $\mathbb{T}$ may cover hundreds or even thousands of periods. Thus, although the synthesis of constraint sampling and decision rule approximations results in a tractable problem, it may not be applied to the problems $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{W}_\eta$ directly. To reduce the complexity of the hedging problems, we pro-

pose a methodology to first decrease the number of time-periods in the hedging problems before applying the approach developed in Chapter 2.

We remark that the hedging problems not only involve a very large number of time-periods, but also a large number of state variables. Indeed, consider a swing option with a delivery period of one month. The number of forward and option contracts traded in real electricity markets and with delivery period overlapping with that of the swing option is in the range of 20 to 30. This precludes the use of both dynamic programming and traditional stochastic programming solution approaches. Stochastic programming scenario tree based discretization schemes are known to suffer from limited tractability when applied to asset-allocation type problems such as $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{W}_\eta$. Indeed, in order to preclude the presence of arbitrage opportunities in the tree, the branching factor of the tree must exceed the number of nonredundant assets in the market (Geyer et al. [2010]). Otherwise, arbitrage opportunities would be built into the tree that would render the associated hedging problems unbounded. Thus, to hedge a swing option with a delivery period of month and daily adjustments to the exercise strategy using 20 basic securities would require over 10^{39} scenarios.

We now present our approximations in the context of problem $\mathcal{H}$. Our results immediately extend to problems $\mathcal{W}$ and $\mathcal{W}_\eta$.

3.4.1 Aggregation of Time-Periods

In order to obtain a problem amenable to solution via the approach discussed in Chapter 2, we propose to aggregate the time-periods in $\mathbb{T}$, which we henceforth call *micro-periods*, to fewer *macro-periods*. Formally, we denote the set of macro-periods by $\mathbb{M} = \{1, \ldots, M\}$, and we assume that there is a strictly increasing function $\tau : \mathbb{M} \cup \{M+1\} \mapsto \mathbb{T} \cup \{T+1\}$ such that $\tau(m)$ represents the first micro-period within the m^{th} macro-period. We set $\tau(1) = 1$ and $\tau(M+1) = T+1$, so that the m^{th} macro-period consists of the micro-periods $\tau(m), \ldots, \tau(m+1)-1$. With a slight abuse of notation, we also define the reduced observation histories $\xi^m = (\xi_{\tau(1)}, \ldots, \xi_{\tau(m)})$, $m \in \mathbb{M}$, consisting only of the risk factors observed at the beginning of each macro-period. As

before, $\xi^t = (\xi_1, \ldots, \xi_t)$ represents the full observation history, and we use the indices t and m to indicate whether a vector has components for each micro-period (t) or each macro-period (m).

While there does not necessarily exist an optimal bang-bang exercise strategy for the swing option, see Section 3.3, it is reasonable to assume that there exists an exercise strategy $e^\star$ optimal in $\mathcal{H}$ which is representable as a convex combination of a small (yet unknown) number of bang-bang strategies $e^{\mathrm{b}}_{t,l}$, $l = 1, \ldots, L$ of the form

$$e^{\mathrm{b}}_{t,l}(\xi^t) = \begin{cases} \min\left(\overline{e}_t, \overline{c} - \sum_{\tau=t+1}^{T} \underline{e}_\tau\right) & \text{if } S_t(\xi_t) \geq q_{t,l}(\xi^t), \\ \max\left(\underline{e}_t, \underline{c} - \sum_{\tau=t+1}^{T} \overline{e}_\tau\right) & \text{otherwise,} \end{cases}$$

for some unknown functions $q_{t,l}(\xi^t)$ which typically change slowly with t and ξ^t. These functions can conveniently be interpreted as exercise thresholds. Whenever the spot price exceeds (falls short of) $q_{t,l}$, the swing option is exercised at maximum (minimum) delivery rate. This motivates us to consider a finite number of *candidate exercise strategies* $e_{t,l}$, $l = 1, \ldots, L$, defined through

$$e_{t,l}(\xi_t) = \begin{cases} \min\left(\overline{e}_t, \overline{c} - \sum_{\tau=t+1}^{T} \underline{e}_\tau\right) & \text{if } S_t(\xi_t) \geq q_l, \\ \max\left(\underline{e}_t, \underline{c} - \sum_{\tau=t+1}^{T} \overline{e}_\tau\right) & \text{otherwise,} \end{cases}$$

for some prescribed constant exercise thresholds q_l, $l = 1, \ldots, L$. We assume that $q_1 = \infty$ and $q_L = 0$, that is, the first (last) candidate exercise strategy always exercises the swing option at minimum (maximum) delivery rate.

The basic idea to simplify problem $\mathcal{H}$ is the following. Instead of choosing an individual exercise decision $e_t(\xi^t)$ at the beginning of each micro-period, we choose, at the beginning of each macro-period, the coefficients of a convex combination of the finitely many candidate strategies. We achieve this by assigning each candidate exercise strategy $e_{t,l}$ a weight $\lambda_{m,l}(\xi^m) \geq 0$, $\sum_{l=1}^{L} \lambda_{m,l}(\xi^m) = 1$, that is held fixed in each macro-period. We then obtain the exercise strategy

$$e_t(\xi^t) = \sum_{l=1}^{L} \lambda_{m,l}(\xi^m) e_{t,l}(\xi_t) \text{ for } t = \tau(m), \ldots, \tau(m+1) - 1. \tag{3.13}$$

The weight vectors $\lambda_m(\xi^m) := (\lambda_{m,1}(\xi^m), \ldots, \lambda_{m,L}(\xi^m))$ represent the new decision variables

that replace the original variables e_t. In our numerical examples we will be able to choose $ML \ll T$, which results in a substantial complexity reduction. We remark that any other candidate exercise strategies satisfying the per-period energy limits could be used instead of the ones proposed above.

In order to reformulate problem $\mathcal{H}$ in terms of the new decision variables, we define

$$e_{m,l}(\xi) := \sum_{t=\tau(m)}^{\tau(m+1)-1} e_{t,l}(\xi_t)$$

as the cumulative energy consumption of the candidate exercise strategy $e_{t,l}$ within macro-period m and

$$\delta_{m,l}(\xi) = \sum_{t=\tau(m)}^{\tau(m+1)-1} e_{t,l}(\xi_t)\,(S_t(\xi_t) - K)$$

as the aggregate dividend earned by exercising the swing option according to $e_{t,l}$ within macro-period m. For notational convenience, we also define the vectors $e_m(\xi) := (e_{m,1}(\xi), \ldots, e_{m,L}(\xi))$ and $\delta_m(\xi) := (\delta_{m,l}(\xi), \ldots, \delta_{m,L}(\xi))$.

As for the trading strategy, there is no incentive to rebalance the portfolio of electricity derivatives if no new information is observed. Thus, we only allow portfolio adjustments $u_t(\xi^t)$ in the first micro-period of each macro-period, that is, for micro-periods $t = \tau(m)$ for some $m \in \mathbb{M}$. All other adjustment variables $u_t(\xi^t)$ are set to zero. This implies that $x_t(\xi^t) = x_{t-1}(\xi^{t-1})$ for all $\tau(m) < t < \tau(m+1)$, $m \in \mathbb{M}$ and $\xi \in \Xi$, that is, the position variables $x_t(\xi^t)$ do not change within the macro-periods. With these simplifications, we can replace the original position and adjustment variables with new decision variables $x_m(\xi^m)$ and $u_m(\xi^m)$, respectively, where $x_m(\xi^m)$ represents the portfolio positions during macro-period m and $u_m(\xi^m)$ the portfolio adjustments at the beginning of macro-period m. Finally, we define the aggregate dividend earned by holding the j^{th} basic security in macro-period m through

$$d_{m,j}(\xi) = \sum_{t=\tau(m)}^{\tau(m+1)-1} d_{t,j}(\xi_t),$$

and we set $d_m(\xi) := (d_{m,1}(\xi), \ldots, d_{m,J}(\xi))$.

We now formulate the *aggregated holder's problem*:

$$
\begin{array}{ll}
\max\limits_{\lambda,x,u} & -P_1^\top x_0 \\
\text{s.t.} & \underline{c} \leq \sum_{m\in\mathbb{M}} e_m^\top \lambda_m \leq \overline{c} \\
& \left.\begin{array}{l} \lambda_m \geq 0,\ \sum_{l=1}^{L} \lambda_{m,l} = 1 \\ x_m = x_{m-1} + u_m \\ \delta_m^\top \lambda_m + d_m^\top x_m - P_{\tau(m)}^\top u_m \geq 0 \end{array}\right\} \forall m \in \mathbb{M} \\
& P_{T+1}^\top x_M \geq 0.
\end{array}
\tag{$\mathcal{AH}$}
$$

The decision variables of this problem are $\lambda_m \in \mathcal{L}_{mt,L}$, $x_m \in \mathcal{L}_{mt,J}$ and $u_m \in \mathcal{L}_{mt,J}$, $m \in \mathbb{M}$, as well as $x_0 \in \mathbb{R}^J$. We notationally suppress dependence on ξ to avoid clutter. Thus, the decision variables are chosen once per macro-period and correspond to measurable functions of the reduced observation history ξ^m. The constraints are understood to hold almost surely.

The objective function is the same as in problem $\mathcal{H}$, while the first three constraints in $\mathcal{AH}$ ensure that the exercise strategy is admissible, and the last three constraints have the same meaning as the constraints implied by the requirement $(x,u) \in \mathcal{S}^{\mathrm{h}}(e)$ in problem $\mathcal{H}$. The first constraint ensures that the exercise strategy satisfies the cumulative energy limits. The second and third constraints ensure that λ_m is a vector of convex weights for all $m \in \mathbb{M}$. Note that the per-period energy limits are implicitly satisfied. Indeed, the candidate exercise strategies $l = 1, \ldots, L$ satisfy $\underline{e}_t \leq e_{t,l}(\xi_t) \leq \overline{e}_t$ for all $t \in \mathbb{T}$ and $\xi \in \Xi$. Since the exercise decisions (3.13) represent convex combinations of these candidate strategies, they also satisfy the per-period energy limits. In addition, since $\underline{c}_t - \sum_{\tau=t+1}^{T} \overline{e}_\tau \leq e_{t,l}(\xi_t) \leq \overline{c}_t - \sum_{\tau=t+1}^{T} \underline{e}_\tau$ for all $t \in \mathbb{T}$ and $\xi \in \Xi$, the second and third constraints guarantee, for any $t \in \mathbb{T}$ and $\xi \in \Xi$, the existence of decisions $e_\tau(\xi^\tau)$, $\tau \in \{t+1, \ldots, T\}$ satisfying the cumulative energy limits. The third constraint in $\mathcal{AH}$ ensures that the pair (x,u) is a portfolio strategy on the reduced information space, while the last two constraints ensure that the portfolio containing the swing option is self-financing.

Remark 3.9 *As previously highlighted, certain contracts in real electricity markets are not traded throughout the swing option contract period as assumed in (M7). Following the time-period aggregation, and if the number of macro-periods is small, a contract j might not be*

traded yet at the beginning of macro-period m, while it is not traded anymore at the beginning of macro-period $m+1$. In order to mitigate the loss of optimality due to the inability to trade in contract j, we suggest to allow the swing option holder (and writer) to decide at the beginning of macro-period m the amount of contract j that she wishes to buy when the contract becomes traded (at the uncertain price at which it will be traded).

We can simplify problems $\mathcal{W}$ and $\mathcal{W}_{+\infty}$ in a similar way to obtain the aggregated writer's problem $\mathcal{AW}$ and the robust aggregated writer's problem $\mathcal{AW}_{+\infty}$.

We now show that $\mathcal{AH}$, $\mathcal{AW}$ and $\mathcal{AW}_{+\infty}$ represent conservative approximations for the original problems $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{W}_{+\infty}$, respectively.

Proposition 3.8 *Assume given the optimal values*

1. *H, W and $W_{+\infty}$ of the problems $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{W}_{+\infty}$, and*
2. *AH, AW and $AW_{+\infty}$ of the problems $\mathcal{AH}$, $\mathcal{AW}$ and $\mathcal{AW}_{+\infty}$.*

These values satisfy

(a) $AH \leq H \leq W \leq AW$

(b) $AH \leq H \leq W \leq W_{+\infty} \leq AW_{+\infty}$

The aggregated problems $\mathcal{AH}$, $\mathcal{AW}$ and $\mathcal{AW}_{+\infty}$ are feasible if and only if the original problems $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{W}_{+\infty}$ are feasible.

Proof (a) We can decompose any feasible solution (λ, x, u) in problem $\mathcal{AH}$ into a feasible solution (e', x', u') in problem $\mathcal{H}$ so that $x'_0 = x_0$. This may require trading in the risk-less security so that (e', x', u') satisfies the self-financing constraint in $\mathcal{H}$ for all $t \in \mathbb{T}$. Since (λ, x, u) and (e', x', u') attain the same objective values in their respective problems, we have $AH \leq H$. A similar argument shows that $W \leq AW$ for the problems $\mathcal{W}$ and $\mathcal{AW}$. Finally, the inequality $H \leq W$ is proven in Proposition 3.5.

(b) Any solution (x, u) feasible in problem $\mathcal{W}_{+\infty}$ can be used to construct a solution (x', u') feasible in problem $\mathcal{W}$ and satisfying $x'_0 = x_0$ by trading in the risk-free asset as required. Thus $W \leq W_{+\infty}$. Following a similar argumentation as in the first part of the proof, one can also show that $W_{+\infty} \leq AW_{+\infty}$.

We now show that $\mathcal{W}$, $\mathcal{AW}$, $\mathcal{W}_{+\infty}$ and $\mathcal{AW}_{+\infty}$ are always feasible. To this end, consider problem $\mathcal{W}$. Compactness of the support Ξ and continuity of the spot price S_t imply that, for any exercise strategy $e \in \mathcal{E}$, the cumulative obligations arising from the swing option, $\sum_{t\in\mathbb{T}} e_t(S_t - K)$, are bounded above by a constant O. We thus obtain a feasible solution (x, u) to problem $\mathcal{W}$ if we take a large loan today and use this loan to cover all obligations arising from the swing option. Similar arguments show that $\mathcal{AW}$, $\mathcal{W}_{+\infty}$ and $\mathcal{AW}_{+\infty}$ are always feasible as well.

The arguments from the previous paragraph extend to problem $\mathcal{H}$. Hence, $\mathcal{H}$ is feasible if and only if there is a feasible exercise strategy e for the swing option, that is, if

$$\left[\sum_{t\in\mathbb{T}} \underline{e}_t, \sum_{t\in\mathbb{T}} \overline{e}_t\right] \cap [\underline{c}, \overline{c}] \neq \emptyset.$$

Assume that these two intervals indeed have a nonempty intersection, and let c be contained in that intersection. Then $c = \gamma \sum_{t\in\mathbb{T}} \underline{e}_t + (1-\gamma) \sum_{t\in\mathbb{T}} \overline{e}_t$ for some $\gamma \in [0, 1]$. Choose $\lambda^1_m = \gamma$, $\lambda^L_m = 1 - \gamma$ and $\lambda^l_m = 0$ for $l \notin \{1, L\}$ for each macro-period $m \in \mathbb{M}$. This choice of λ satisfies the second constraint in $\mathcal{AH}$, and we can assign λ an appropriate trading strategy (x, u) so that (λ, x, u) is feasible in $\mathcal{AH}$. □

Unfortunately, in the case of $\eta < +\infty$, deriving a conservative approximation to problem $\mathcal{W}_\eta$ is more intricate. Indeed, the exact solution of the aggregated counterpart of problem $\mathcal{W}_\eta$ necessitates the solution, for each $\xi \in \Xi$, of a problem similar to $\mathcal{H}$, see Proposition 3.4. While this subproblem could in principle be solved via the techniques described here, the solution to the resulting aggregated writer's problem with parameter η would not necessarily yield a conservative approximation to $\mathcal{W}_\eta$.

3.4.2 Decision Rule Approximation

The time-period aggregation enabled us to substantially reduce the number of time-periods in the hedging problems $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{AW}_{+\infty}$. While all approximate problems $\mathcal{AH}$, $\mathcal{AW}$ and $\mathcal{AW}_{+\infty}$ still involve infinitely many decision variables and constraints, they can now be solved efficiently with the approximation scheme developed in Chapter 2. We thus suggest to approximate the adaptive decision variables λ_m, x_m and u_m in the aggregated problems by finite linear combinations of basis functions of the reduced observation history ξ^m, see Section 2.3.1. We refer to the resulting aggregated problems with basis functions by $\mathcal{AHB}$, $\mathcal{AWB}$ and $\mathcal{AWB}_{+\infty}$, respectively.

We now show that $\mathcal{AHB}$, $\mathcal{AWB}$ and $\mathcal{AWB}_{+\infty}$ are conservative approximations for $\mathcal{AH}$, $\mathcal{AW}$ and $\mathcal{AW}_{+\infty}$, respectively.

Proposition 3.9 *Assume given the optimal values*

1. *H, W and $W_{+\infty}$ of the problems $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{W}_{+\infty}$, and*
2. *AH, AW and $AW_{+\infty}$ of the problems $\mathcal{AH}$, $\mathcal{AW}$ and $\mathcal{AW}_{+\infty}$.*
3. *AHB, AWB and $AWB_{+\infty}$ of the problems $\mathcal{AHB}$, $\mathcal{AWB}$ and $\mathcal{AWB}_{+\infty}$.*

These values satisfy the chains of inequalities

(a) $AHB \leq AH \leq H \leq W \leq AW \leq AWB$

(b) $AHB \leq AH \leq H \leq W \leq W_{+\infty} \leq AW_{+\infty} \leq AWB_{+\infty}$

The problems $\mathcal{AHB}$, $\mathcal{AWB}$ and $\mathcal{AWB}_{+\infty}$ are feasible if and only if the problems $\mathcal{H}$, $\mathcal{W}$ and $\mathcal{W}_{+\infty}$ are feasible.

Proof The chains of inequalities (a) and (b) follow from Proposition 3.8 and the fact that any feasible solution to $\mathcal{AHB}$, $\mathcal{AWB}$ and $\mathcal{AWB}_{+\infty}$ is also feasible in $\mathcal{AH}$, $\mathcal{AW}$ and $\mathcal{AW}_{+\infty}$, respectively. The proof of the last statement parallels the proof of Proposition 3.8 and can therefore be omitted. □

3.4.3 Constraint Sampling Approximation

Although $\mathcal{AHB}$, $\mathcal{AWB}$ and $\mathcal{AWB}_{+\infty}$ involve finitely many decision variables, they remain intractable since their constraints are parameterized by $\xi \in \Xi$, while Ξ typically has infinite cardinality. We thus proceed with the constraint sampling approximation discussed in Section 2.3.2 and enforce the constraints in the aggregated problems with basis functions over finite subsets of Ξ obtained through Monte-Carlo sampling of ξ. We denote these approximate problems by $\mathcal{AHB}^N$, $\mathcal{AWB}^N$ and $\mathcal{AWB}^N_{+\infty}$, respectively. As before, we denote by ϵ the target violation probability, by β the confidence that this probability will not be exceeded and by $N(\epsilon, \beta)$ the number of samples needed to guarantee that the violation probability of the solution to the sampled problem does not exceed ϵ with confidence greater than $1 - \beta$, see Theorem 2.1. We now highlight the meaning of ϵ in the context of the hedging problems. We focus our discussions on problem $\mathcal{AHB}$, but they immediately extend to $\mathcal{AWB}$ and $\mathcal{AWB}_{+\infty}$.

The parameter ϵ describes the probability that an optimal solution to $\mathcal{AHB}^N$ violates a constraint of $\mathcal{AHB}$ for a randomly chosen sample $\xi \in \Xi$, see Section 2.4. A constraint violation implies that the terminal wealth is negative, the per-period or cumulative energy constraints are violated, or that the portfolio containing the swing option and the basic securities is not self-financing. The option holder thus faces a risk when she implements an optimal solution to $\mathcal{AHB}^N$, and this risk can be controlled by choosing an appropriate value for ϵ. The choice of ϵ thus enables us to model the negotiation process between the holder and the writer of the option. As ϵ increases, the two parties take on a greater part of the risk inherent in the inability to perfectly replicate the swing option and the holder's and writer's prices for the option converge.

If we fix the complexity parameter d of the decision rules, see Section 2.3.2, as well as the violation probability ϵ and confidence β, then the required sample size $N(\epsilon, \beta)$ is bounded from above by a polynomial in the parameters k, J and M. Thus, the approximate problems $\mathcal{AHB}^N$, $\mathcal{AWB}^N$ and $\mathcal{AWB}^N_{+\infty}$ are linear programs that involve polynomially many variables and constraints and can thus be solved efficiently.

Remark 3.10 *As discussed in Section 3.2.3, the holder's problem and by extension $\mathcal{AHB}^N$ might have multiple optimal solutions. In accordance with our observations in Section 3.2.3, we henceforth select the w.l.o.g. unique strategy optimal in $\mathcal{AHB}^N$ which maximizes the expected exercise profits of the swing option, see Remark 2.2.*

3.5 Other Applications

While our approach has been developed for the purpose of hedging and valuing swing options, numerous other contracts come with contractual or physical constraints similar to those inherent in a swing contract. Our approximation scheme thus has applications for a broader class of hedging and valuation problems. In what follows, we discuss a few of these alternative applications.

3.5.1 Valuation of Mines and Oil Fields

An immediate application of our approach is in the valuation of mines and oil fields (see e.g., Brennan and Schwartz [1985], Ludkovski and Carmona [2010]). Prices of commodities such as oil, gold or copper are highly volatile and exhibit frequent spikes. Mines and oil fields can act as storage facilities which enable the transfer of the commodity from periods of low demand (low prices) to periods of high consumption. For example, the owner of a mine (or the agent holding the rights to manage the mine) may decide to withhold extraction until the price of the commodity increases further. In what follows, we discuss the valuation of oil fields, but direct parallels can be made with the case of mines.

We discuss the problem of valuing a contract which gives its holder the right to manage an oil field over a planning horizon $\mathbb{T} = \{1, \dots, T\}$. At the beginning of each period $t \in \mathbb{T}$ (e.g., each day or hour), the holder of the contract may select the amount e_t of oil she wishes to extract from the field for sale on the spot market. Extraction costs K are incurred for each unit of oil

extracted. Thus, the payoff from extracting $e_t(\xi^t)$ units of oil at t is given by

$$e_t(\xi^t)(S_t(\xi_t) - K).$$

The amount extracted during any one time-period $t \in \mathbb{T}$ cannot exceed the extraction rate $\overline{e}_t$ for that period, while we assume that no oil may be injected into the field, that is,

$$0 \leq e_t(\xi^t) \leq \overline{e}_t.$$

The cumulative amount of oil extracted from the field can never exceed the capacity $\overline{c}$ of the field, i.e.,

$$0 \leq \sum_{t \in \mathbb{T}} e_t(\xi^t) \leq \overline{c}.$$

Thus, the problem of valuing mines and oil fields can immediately be cast into our framework.

3.5.2 Valuation of Power Plants and Refineries

Power plants and refineries enable the conversion of a fuel commodity into another fuel commodity. A power plant converts for example gas, oil or coal into electricity, while a refinery converts e.g., crude oil into gasoline and natural gas into commercial or industrial fuel gas. From a financial point of view, they thus enable capitalizing on the price differential between two commodities.

Power plants and refineries are costly to build and maintain while participants in energy markets require access to the physical equipment. Indeed, numerous commodity contracts require delivery of the underlying. For this purpose, lease contracts (sometimes termed tolling agreements) were introduced, see e.g., Deng et al. [2001], Ludkovski [2005], Deng and Xia [2005]. These give their holder (i.e., the agent renting the facility) the right to plan the power plant dispatching policy, i.e., the level of production, over a finite planning horizon $\mathbb{T}$.

Our method is applicable to the valuation of such contracts which are subject to the physical characteristics of the plant or refinery. We now focus on the valuation of power plants. At the

beginning of each time-period, the contract holder may select the amount e_t of electricity to produce. This quantity may never exceed the production rate $\overline{e}_t$ of the plant,

$$0 \leq e_t(\xi^t) \leq \overline{e}_t.$$

In order to produce an amount e_t of electricity, the contract holder must purchase he_t units of the input commodity at price $G_t(\xi_t)$, where h denotes the heat rate of the plant, which we assume to be independent of the output level e_t. The heat rate thus corresponds to the number units of the input commodity needed to produce 1 MWh of electricity (the higher the heat rate, the less efficient the plant). Production costs K are incurred for each unit of electricity produced. The payoff from generating e_t units of energy at time t is thus

$$e_t(\xi^t)(S_t(\xi_t) - hG_t(\xi_t) - K).$$

Some lease contracts require the cumulative production of the plant over the lease period to lie below a certain level $\overline{c}$, that is,

$$0 \leq \sum_{t \in \mathbb{T}} e_t(\xi^t) \leq \overline{c},$$

so as to moderate plant wear.

Remark 3.11 *Certain power plants (e.g., thermal power stations) impose restrictions on the rate of change of e_t: $\rho_t^- \leq e_{t-1} - e_t \leq \rho_t^+$. These are known as ramping constraints and can be accommodated by our framework for example by constructing candidate strategies which satisfy these constraints and by keeping a sufficient number of micro-periods in between two consecutive macro-periods.*

Remark 3.12 *Certain least contract do not incorporate limitations on the cumulative amount of electricity generated. In this case, there always exists an optimal operation strategy for the power plant which is of bang-bang type with exercise threshold given by K.*

3.6 Case Study

We consider a planning horizon of five weeks subdivided into periods of one hour (i.e., $T = 958$). We hedge an hourly exercisable peak swing option with a delivery period of one month starting today (i.e., at the beginning of period 1). The per-period energy limits are $\underline{e}_t = 0$ MWh, $t \in \mathbb{T}$, $\overline{e}_t = 1$ MWh for peak hours and $\overline{e}_t = 0$ MWh for off-peak hours. The option's strike price is $K = 30$ €/MWh. The cumulative energy limits are $\underline{c} = 0$ MWh and $\overline{c} = 72$ MWh, and exceedance of the upper energy limit $\overline{c}$ is penalized by a cost of 10 €/MWh.

We hedge the swing option with a money market account and all the electricity forwards that would be traded on the European Energy Exchange[2] during the planning horizon. In particular, we consider weekly forwards for five weeks and monthly forwards for two months. In both cases, we consider base and peak contracts. In total, the market thus contains $J = 15$ basic securities. The spot price of electricity and the prices as well as dividends of the basic securities are modelled with the approach described by Haarbrücker and Kuhn [2009]. We remark, however, that any other price model can be used instead.

We partition the planning horizon into M macro-periods of (approximately) equal length, see Section 3.4.1. Moreover, we consider convex combinations of candidate exercise strategies e_l, $l = 1, \ldots, L$, whose exercise thresholds q_l are selected uniformly between the strike price K and the maximum of the spot price over the planning horizon, see Section 3.4.2. Note that exercise thresholds below the strike price K are always sub-optimal in this case since the swing option holder is not obligated to purchase electric energy ($\underline{e} = 0$ and $\underline{c} = 0$). We consider polynomial decision rules, see Example 2.1, and choose the number N of random samples $\xi \in \Xi$ such that the target violation probability ϵ is satisfied at the confidence level $\beta = 0.01\%$, see Section 3.4.3.

In the following, we denote by $\mathcal{H}$ the holder's problem, by $\mathcal{W}$ the writer's problem, and by $\mathcal{W}_{+\infty}$ the robust writer's problem, see Section 3.2.

We first investigate the impact of the approximation parameters M, L, d and ϵ on the holder's price, the writer's price and the robust writer's price of the swing option. As Figure 3.5(a)

[2]European Energy Exchange: `www.eex.com`

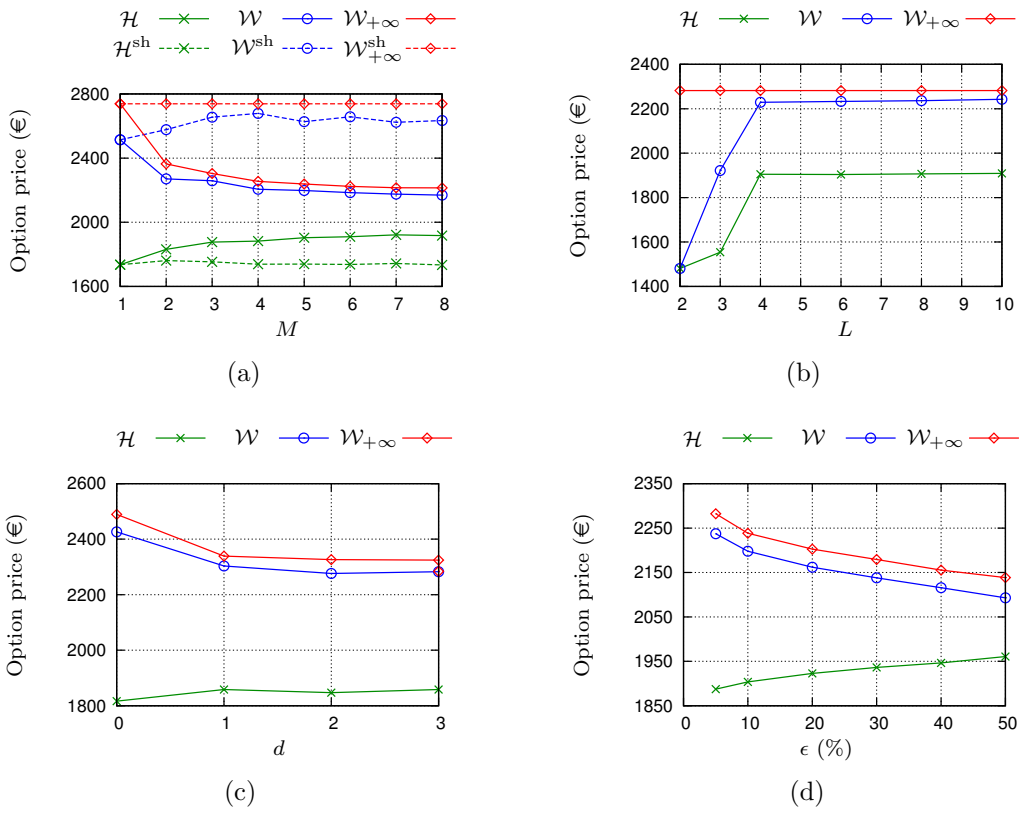


Figure 3.5: No-arbitrage interval as a function of M, L, d and ϵ. The superscript "sh" indicates static hedging. The basic parameter setting is $M = 5$, $L = 5$, $d = 1$ and $\epsilon = 5\%$.

shows, the no-arbitrage interval shrinks considerably when the number of macro-periods M increases, indicating that dynamic hedging of the swing option is essential. The no-arbitrage interval saturates when about 7 macro-periods are used. The chart also shows that one obtains severely sub-optimal no-arbitrage intervals if the holder or the writer hedges the swing option statically, that is, if the hedging portfolio can only be adjusted in the first macro-period. Figure 3.5(b) shows that a small number of candidate exercise strategies ($L \geq 4$) suffices for a good approximation of the optimal exercise strategy. Figure 3.5(c) illustrates the holder's and writer's price as the degree d of the polynomial decision rules is increased. The chart shows that the price gap reduces significantly when moving from constant to linear decision rules, while a further increase in d does not have a noticeable effect. Finally, Figure 3.5(d) illustrates the impact of the violation probability ϵ on the option's no-arbitrage interval. With larger values of ϵ, the holder and the writer accept to take on part of the risk that results from the inability to perfectly replicate the option's payoff streams. This causes the holder's and writer's price to converge as ϵ increases.

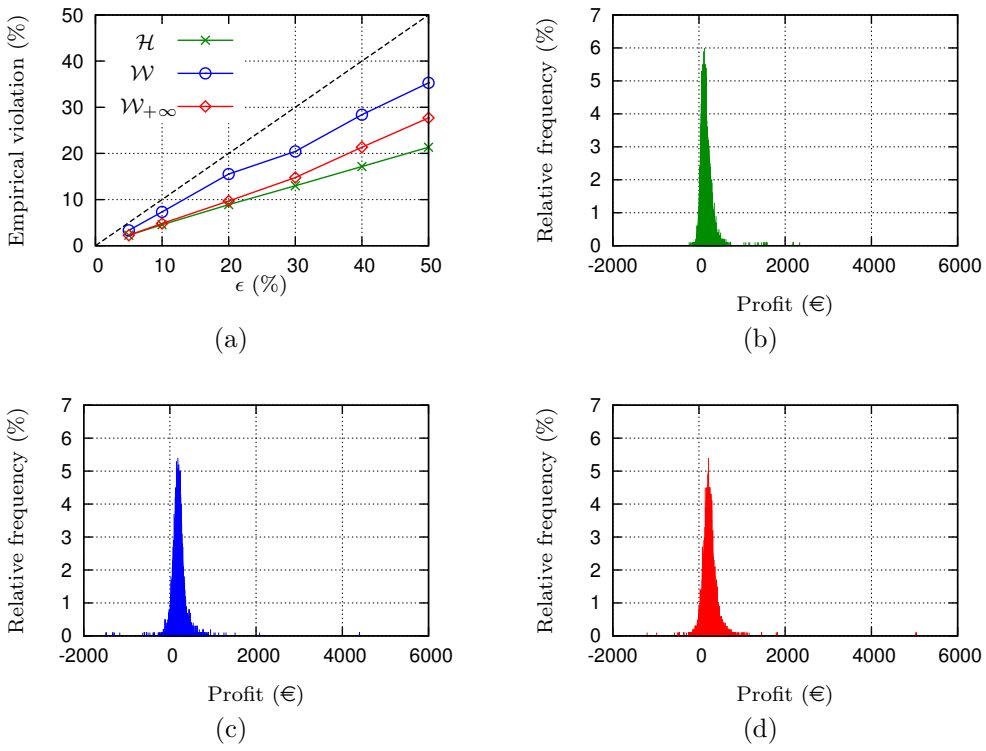


Figure 3.6: Figure (a) depicts the empirical violation probability as a function of ϵ. Figures (b), (c) and (d) illustrate the empirical profit/loss distributions for the holder, the writer and the robust writer, respectively. The basic parameter setting is $M = 5$, $L = 5$, $d = 1$ and $\epsilon = 5\%$.

We now assess the risk faced by the holder and the writer of the swing option. To this end, we draw 10,000 independent samples $\xi \in \Xi$. We define the empirical violation probability of the holder, the writer and the robust writer as the percentage of those 10,000 samples for which the solution to $\mathcal{AHB}^N$, $\mathcal{AWB}^N$ and $\mathcal{AWB}^N_{+\infty}$ violates any one of the constraints of $\mathcal{AHB}$, $\mathcal{AWB}$ and $\mathcal{AWB}_{+\infty}$, respectively. The empirical violation probabilities for different levels of (target) violation probabilities ϵ are shown in Figure 3.6(a). As expected from Theorem 2.1 and our choice of β, the empirical violation never exceeds ϵ (dashed line). The other charts in Figure 3.6 visualize the empirical profit/loss distributions. The charts show that the expected profits of both the holder and the writer are positive. The distributions have a positive skewness, indicating that the tail on the right side (profit) is longer than on the left side (loss).

We finally remark that apart from the instances considered in Figure 3.5(a), all problems were solved within 20 min on a 2.66GHz Intel Core i7-920 machine running CPLEX 12.2. The problems in Figure 3.5(a) were solved within 0.04 sec ($M = 1$) and 55 min ($M = 8$).

Chapter 4

Decision Rules for Information Discovery in Multi-Stage Stochastic Programming

Insofar, the problems discussed and investigated in this thesis assume that the order in which the uncertainties unfold is *in*dependent of the controller actions. In fact, this is true for most of the literature on dynamic decision-making under uncertainty. Nevertheless, in numerous real-world decision problems, the time of information discovery can be influenced by the decision-maker, and uncertainties only become observable following an (often costly) investment. Such problems can be formulated as mixed-binary multi-stage stochastic programs with decision-dependent non-anticipativity constraints. Unfortunately, these problems are severely computationally intractable. In this chapter we propose an approximation scheme for multi-stage problems with decision-dependent information discovery which is based on techniques commonly used in modern robust optimization. In particular, we obtain a conservative approximation in the form of a mixed-binary linear program by restricting the spaces of measurable binary and real-valued decision rules to those that are representable as piecewise constant and linear functions of the uncertain parameters, respectively. We assess our approach on a problem of infrastructure and production planning in offshore oil fields from the literature.

4.1 Introduction

The vast majority of models and algorithms for solving dynamic decision problems affected by uncertain data assume that the decision-maker cannot influence the order in which the uncertain parameters are revealed. However, this assumption fails to hold in numerous real-world decision problems, where the decision-maker can decide whether and when to observe random parameters. In order to establish a succinct terminology, Jonsbråten [1998] coined the terms of *exogenous* and *endogenous* uncertainties, which refer to parameters whose "time of revelation" is *in*dependent and dependent of the decisions, respectively. We will use this terminology throughout the remainder of this thesis. Moreover, we will refer to those decisions that trigger an information discovery as *measurement* or *observation* variables.

We highlight the practical significance of models with endogenous uncertainties by presenting several real-world decision problems in which the time of information discovery is inherently decision-dependent.

4.1.1 Motivating Examples

Oil companies spend substantial efforts on infrastructure and production planning in offshore oil fields (Goel and Grossman [2004]), which typically consist of several reservoirs with uncertain volumes. For each reservoir, one needs to determine whether and when to extract oil. The uncertain volume of a reservoir becomes observable only when an expensive well platform is built and the drilling for oil is initiated. The drilling decisions thus control the sequence of information discovery.

Pharmaceutical companies typically maintain R&D pipelines that comprise multiple candidate drugs. Before a drug can enter the marketplace, it needs to pass several costly clinical trials that may last for many years. The outcome (success/failure) of each trial is uncertain and will only be revealed once the trial is completed. Thus, pharmaceutical companies need to orchestrate the clinical trials with the goal to maximize the rate of discovering effective drugs

(Colvin and Maravelias [2008]). The decisions to proceed with different trials can thus be viewed as measurement variables which determine how the uncertainty unfolds.

A related problem is that of R&D project portfolio optimization (Solak et al. [2010]). Here, the goal is to decide how to distribute scarce resources among a number of projects with different performance characteristics. The return of any project is uncertain and will only be revealed upon the project's termination. The start times of the various projects and the resource allocations thus impact the time of information discovery.

4.1.2 Literature Review

Research on stochastic programming with endogenous uncertainties began with the works of Jonsbråten et al. [1998] and Jonsbråten [1998] in 1998. They studied decision problems in which the control actions can impact both the distribution of the uncertainties as well as the timing of their revelation. Problems with decision-dependent information discovery are perceived as particularly hard even if the distribution of the uncertainties is unaffected by the decisions, and therefore the literature on numerical solution procedures remains scarce. To the best of our knowledge, all existing algorithms rely on the assumption that the uncertain parameters follow a discrete distribution. In this case the decision process can be modeled through a finite scenario tree whose branching structure depends on the binary measurement decisions that determine the time of information discovery, see Jonsbråten [1998]. Goel and Grossman [2004] have shown that stochastic programs with discretely distributed endogenous uncertainties can be reformulated as deterministic mixed-binary programs, but unfortunately these reformulations involve an exponential number of binary variables and non-anticipativity constraints. Research efforts have consequently focused on approximation techniques that provide sub-optimal but feasible solutions to the original problem. An effective approach to complexity reduction is to require that the measurement decisions be pre-committed, that is, to approximate them by here-and-now decisions. The resulting approximate problems can be solved with an enumeration-based branch-and-bound algorithm due to Jonsbråten et al. [1998] or via decomposition techniques by Goel and Grossman [2004]. More recent branch-and-bound and branch-and-cut algorithms

truthfully account for the adaptive nature of the measurement variables, see Goel and Grossman [2006], Goel et al. [2006] as well as Colvin and Maravelias [2010], respectively. Moreover, several iterative schemes based on relaxations of the non-anticipativity constraints for the measurement variables have been proposed by Gupta and Grossman [2010] and by Colvin and Maravelias [2010].

Problems involving continuously distributed random parameters need to be discretized before any of the above solution procedures can be applied. Solak et al. [2010] propose to use a sample average approximation for this purpose. While discretization appears as a promising approach for smaller problems, it may result in a combinatorial state explosion when applied to even medium sized problems. Conversely, using only very few discretization points can result in solutions that are sub-optimal or may even fail to be implementable in practice.

In this chapter we develop a methodology for solving dynamic problems with endogenous uncertainties, which is inspired by modern robust optimization techniques, see Section 1.1. We suggest to approximate the adaptive measurement decisions by piecewise constant functions and the adaptive real-valued decisions by piecewise linear functions of the uncertainties. The resulting approximate problems are equivalent to mixed-binary linear programs (MBLP), which can be solved using standard optimization software. This decision rule approximation remains applicable when the uncertain parameters are continuously distributed, and it results in near-optimal solutions that are implementable in reality. The trade-off between the solution quality and the computational speed is controlled by the granularity of the partition of the uncertainty domain. We remark that, up to date, the benefits of decision rule techniques have only been exploited in the context of stochastic and robust optimization with *exogenous* uncertainty, see e.g., Ben-Tal et al. [2004], Shapiro and Nemirovski [2005], Goh and Sim [2010] and Kuhn et al. [2009].

This chapter is organized as follows. The remainder of this section introduces the notation, while Section 4.2 and Section 4.3 develop a new decision rule approximation for two- and multi-stage stochastic programs affected by endogenous uncertainty, respectively. The benefits of our approach are illustrated in Section 4.4 through an example in the area of infrastructure and

production planning.

Notation Throughout this chapter, uncertainty is modeled by the probability space $(\mathbb{R}^k, \mathcal{B}(\mathbb{R}^k), \mathbb{P})$, which consists of the sample space $\mathbb{R}^k$, the Borel σ-algebra $\mathcal{B}(\mathbb{R}^k)$ and the probability measure $\mathbb{P}$, whose support we denote by Ξ. We assume that Ξ is a compact polyhedral subset of $\{\xi \in \mathbb{R}^k : \xi_1 = 1\}$. This non-restrictive assumption allows us to represent affine functions of the non-degenerate uncertain parameters $(\xi_2, \ldots, \xi_k)$ in a compact way as linear functions of $\xi = (\xi_1, \ldots, \xi_k)$. We let $\mathbb{E}(\cdot)$ denote the expectation operator with respect to $\mathbb{P}$. We further denote by $\mu := \mathbb{E}(\xi)$ the first order moment vector and by $\Sigma := \mathbb{E}(\xi\xi^\top)$ the second order moment matrix of ξ under $\mathbb{P}$. For any $m, n \in \mathbb{N}$, we let $\mathcal{L}_{m,n}$ be the space of all measurable functions from $\mathbb{R}^m$ to $\mathbb{R}^n$ that are bounded on compact sets. For two vectors $x, y \in \mathbb{R}^n$, we let $x \circ y \in \mathbb{R}^n$ denote their Hadamard product and for $j \in \mathbb{N}$, we define $x_{-j} \in \mathbb{R}^{n-1}$ as $x_{-j} := (x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)$. For a square matrix $A \in \mathbb{R}^{n\times n}$, we let $\mathrm{tr}(A)$ denote the trace of A. Finally, we denote by e_1 the first canonical basis vector in $\mathbb{R}^k$.

4.2 The Two-Stage Case

4.2.1 Problem Formulation

A two-stage stochastic program with *exogenous uncertainty* is representable as

$$\begin{array}{ll} \min & c^\top x + \mathbb{E}(q(\xi)^\top y(\xi)) \\ \text{s.t.} & x \in \mathbb{R}^{n_1},\ y \in \mathcal{L}_{k,n_2} \\ & Tx + Wy(\xi) \leq h(\xi)\ \forall \xi \in \Xi, \end{array} \tag{4.1}$$

where $x \in \mathbb{R}^{n_1}$ stands for the vector of first-stage decisions and $y(\xi) \in \mathbb{R}^{n_2}$ denotes the vector of second-stage (or recourse) decisions, which may depend on the observed realization of the random vector $\xi \in \mathbb{R}^k$. Here, $c \in \mathbb{R}^{n_1}$ and $q(\xi) \in \mathbb{R}^{n_2}$ are interpreted as cost vectors, while $T \in \mathbb{R}^{m\times n_1}$ and $W \in \mathbb{R}^{m\times n_2}$ are referred to as the technology and recourse matrices, respectively. Moreover, $h(\xi) \in \mathbb{R}^m$ is termed the right hand side vector. We assume that $q(\xi) = Q\xi$ for some

$Q \in \mathbb{R}^{n_2 \times k}$ and $h(\xi) = H\xi$ for some $H \in \mathbb{R}^{m \times k}$.

The focus of this chapter is a variant of problem (4.1) in which the random vector is not necessarily observable in the second stage. Instead, any component of ξ is observed only if the decision-maker decides to observe (or measure) this particular component. A new binary decision vector $z \in \mathcal{Z} \subseteq \{0,1\}^k$ collects these measurement decisions, that is, ξ_i is observed iff $z_i = 1$. We will henceforth assume that observing random parameters incurs a cost $f^\top z$ for some $f \in \mathbb{R}^k$ and impacts the constraints through an additional term Bz for some $B \in \mathbb{R}^{m \times k}$. In this generalized model, the second stage decisions may only depend on those random parameters that have been observed, that is, they must be representable as functions of $z \circ \xi$. Note that the binary vector z "switches off" those components of ξ that remain unobserved.

A two-stage stochastic program with *endogenous uncertainty* can therefore be formalized as

$$\begin{array}{ll} \min & c^\top x + f^\top z + \mathbb{E}(q(\xi)^\top y(\xi)) \\ \text{s.t.} & x \in \mathbb{R}^{n_1},\ z \in \mathcal{Z},\ y \in \mathcal{L}_{k,n_2} \\ & \left.\begin{array}{l} Tx + Bz + Wy(\xi) \leq h(\xi) \\ y(\xi) = y(z \circ \xi) \end{array}\right\} \ \forall \xi \in \Xi. \end{array} \qquad (\mathcal{P})$$

Note that $\mathcal{Z}$ can be a strict subset of $\{0,1\}^k$, that is, it may incorporate constraints requiring that a particular component of ξ is always observed or that two components of ξ must be observed simultaneously, etc.

Problem $\mathcal{P}$ encapsulates the two-stage stochastic program (4.1) and the (deterministic) mixed-binary linear program (MBLP) as special cases, and it involves complex compositions of functional and binary decisions. Therefore, it is severely computationally intractable. In the next section we propose a conservative approximation that reduces $\mathcal{P}$ to a single-stage robust MBLP problem.

4.2.2 Decision Rule Approximation

We can substantially improve the tractability of problem $\mathcal{P}$ by reducing the space of admissible second-stage decisions to those presenting an affine data dependence, thus being representable as $y(\xi) = Y\xi$ for some $Y \in \mathbb{R}^{n_2 \times k}$. This radical but effective approach to complexity reduction was proposed in Ben-Tal et al. [2004], Goh and Sim [2010], Georghiou et al. [2010], Shapiro and Nemirovski [2005] as a means of approximating multi-stage robust and stochastic programs affected by exogenous uncertainty. Using this approach to simplify problem $\mathcal{P}$, which is affected by endogenous uncertainty, results in the following conservative (upper bound) approximation.

$$\begin{aligned} \min \quad & c^\top x + f^\top z + \operatorname{tr}\left(\Sigma Q^\top Y\right) \\ \text{s.t.} \quad & x \in \mathbb{R}^{n_1},\ z \in \mathcal{Z},\ Y \in \mathbb{R}^{n_2 \times k} \\ & Tx + Bz + WY\xi \leq h(\xi)\ \forall \xi \in \Xi \\ & |Y_{ij}| \leq M z_j \quad i = 1, \ldots, n_2,\ j = 1, \ldots, k. \end{aligned} \tag{$\mathcal{P}_u$}$$

The last constraint in $\mathcal{P}_u$ enforces non-anticipativity. It stipulates that if ξ_j was not observed in the first decision stage, then the affine decision rule $y(\xi) = Y\xi$ must be independent of ξ_j. Here, $M \in \mathbb{R}_+$ denotes a suitably chosen "big-M constant" which is large enough to guarantee that Y_{ij} is unaffected by the non-anticipativity constraint if $z_j = 1$. Problem $\mathcal{P}_u$ can be viewed as a robust MBLP, which involves semi-infinite constraints parameterized by $\xi \in \Xi$. In the following section, we reformulate $\mathcal{P}_u$ as a standard MBLP.

4.2.3 MBLP Reformulation

The key ingredient for reformulating $\mathcal{P}_u$ as an MBLP is the following proposition.

Proposition 4.1 *For any $\phi \in \mathbb{R}^k$ the following statements are equivalent:*

(a) $\phi^\top \xi \geq 0$ for all $\xi \in \Xi$;

(b) ϕ is an element of the cone dual to the cone generated by Ξ, i.e., $\phi \in \mathcal{K} := (cone(\Xi))^*$.

Proof As linear functions are positive homogeneous of degree 1, we have

$$\begin{aligned} \phi^\top \xi \geq 0 \quad \forall \xi \in \Xi \quad &\Leftrightarrow \quad \phi^\top \xi \geq 0 \quad \forall \xi \in \operatorname{cone}(\Xi) \\ &\Leftrightarrow \quad \phi \in (\operatorname{cone}(\Xi))^* \end{aligned}$$

Thus, the claim follows. □

By Proposition 4.1, $\mathcal{P}_u$ can be reformulated as

$$\begin{array}{ll} \min & c^\top x + f^\top z + \operatorname{tr}\left(\Sigma Q^\top Y\right) \\ \text{s.t.} & x \in \mathbb{R}^{n_1},\ z \in \mathcal{Z},\ Y \in \mathbb{R}^{n_2 \times k} \\ & H - (Tx + Bz)\mathrm{e}_1^\top - WY \in \mathcal{K}^m \\ & |Y_{ij}| \leq M z_j \quad i = 1, \ldots, n_2,\ j = 1, \ldots, k, \end{array} \qquad (\mathcal{P}'_u)$$

where $\mathcal{K}^m$ denotes the cone of all $m \times k$-matrices whose rows are all contained in $\mathcal{K}$. Since Ξ is a polyhedral set, $\mathcal{K}^m$ is a polyhedral cone. The conic constraint in $\mathcal{P}'_u$ therefore corresponds to a finite set of linear inequality constraints. Problem $\mathcal{P}'_u$ is thus equivalent to an MBLP involving only a finite number of decision variables and constraints. Its size grows polynomially with k, m, n_1, n_2 and the number of constraints defining the uncertainty set Ξ. The decision rule approximation thus results in a conservative approximation to $\mathcal{P}$ in the form of an MBLP whose size is polynomially bounded in the size of the original problem's inputs.

4.3 The Multi-Stage Case

4.3.1 Problem Formulation

A multi-stage stochastic program with exogenous uncertainty over the finite planning horizon $\mathbb{T} := \{1, \ldots, T\}$ is representable as

$$\begin{array}{ll} \min & \mathbb{E}\left(\sum_{t\in\mathbb{T}} c_t(\xi)^\top y_t(\xi)\right) \\ \text{s.t.} & y_t \in \mathcal{L}_{k,n_t} \ \forall t \in \mathbb{T} \\ & \left.\begin{array}{l} \sum_{\tau=1}^{t} A_{t\tau} y_\tau(\xi) \leq h_t(\xi) \\ y_t(\xi) = y_t(z_{t-1} \circ \xi) \end{array}\right\} \ \forall \xi \in \Xi,\ t \in \mathbb{T}, \end{array} \tag{4.2}$$

where $y_t(\xi) \in \mathbb{R}^{n_t}$ denotes the vector of time t decisions. The binary vector $z_t \in \{0,1\}^k$ represents the *information base* at time $t+1$, that is, it encodes the information revealed up to time t. Thus, we have $z_{t,i} = 1$ iff ξ_i has been observed at some time $\tau \in \{0, \ldots, t\}$. As information is never forgotten, we require that $z_t \geq z_{t-1}$ for all $t \in \mathbb{T}$. The last constraint in (4.2) enforces non-anticipativity by stipulating that y_t can only depend on uncertainties that have been observed up to time $t-1$.

Without much loss of generality, we assume that the problem data satisfies $c_t(\xi) = C_t\xi$ for some $C_t \in \mathbb{R}^{n_t \times k}$, $h_t(\xi) = H_t\xi$ for some $H_t \in \mathbb{R}^{m_t \times k}$ and $A_{t\tau} \in \mathbb{R}^{m_t \times n_\tau}$.

In the remainder we investigate a variant of problem (4.2) that enjoys much greater modeling power since the time of information discovery is kept flexible. We now interpret the information base $z_t(\xi) \in \mathcal{Z}_t \subseteq \{0,1\}^k$ as an adaptive decision variable, which itself depends on ξ. The set $\mathcal{Z}_t$ may incorporate constraints stipulating, for example, that a specific uncertainty can only be observed after a certain stage, etc. We assume that including ξ_i in the information base at time t, which happens iff $z_{t,i}(\xi) = 1$, incurs a cost $f_{t,i}(\xi) \in \mathbb{R}$. Moreover, $z_1(\xi), \ldots, z_t(\xi)$ also impact the time t constraints through the additional term $\sum_{\tau=1}^{t} B_{t\tau} z_\tau(\xi)$ for some $B_{t\tau} \in \mathbb{R}^{m_t \times k}$. Without much loss of generality, we assume that $f_t(\xi) = F_t\xi$ for some $F_t \in \mathbb{R}^{k \times k}$.

A multi-stage stochastic program with endogenous uncertainty can therefore be formalized as

$$
\begin{array}{ll}
\min & \mathbb{E}\left(\sum_{t\in\mathbb{T}} c_t(\xi)^\top y_t(\xi) + f_t(\xi)^\top z_t(\xi)\right) \\
\text{s.t.} & y_t \in \mathcal{L}_{k,n_t},\ z_t \in \mathcal{L}_{k,k}\ \forall t \in \mathbb{T} \\
& \left.\begin{array}{l}
\sum_{\tau=1}^{t} A_{t\tau} y_\tau(\xi) + B_{t\tau} z_\tau(\xi) \le h_t(\xi) \\
z_t(\xi) \in \mathcal{Z}_t \\
z_t(\xi) \ge z_{t-1}(\xi) \\
z_t(\xi) = z_t(z_{t-1}(\xi) \circ \xi) \\
y_t(\xi) = y_t(z_{t-1}(\xi) \circ \xi)
\end{array}\right\} \begin{array}{l} \forall \xi \in \Xi, \\ t \in \mathbb{T}. \end{array}
\end{array}
\qquad (\mathcal{MP})
$$

The fourth constraint in $\mathcal{MP}$ corresponds to an information monotonicity constraint and ensures that information is never forgotten, and the last two constraints enforce non-anticipativity of the binary and real-valued decision variables, respectively. Without loss of generality we assume that $z_0(\xi) = \mathrm{e}_1\ \forall \xi \in \Xi$, that is, only the degenerate random parameter ξ_1 is known at the beginning. Problem $\mathcal{MP}$ subsumes the multi-stage stochastic program (4.2) and it involves decision-dependent non-anticipativity constraints and binary recourse variables. It is therefore severely computationally intractable. In the next section, we propose a conservative approximation that reduces $\mathcal{MP}$ to a static robust MBLP.

4.3.2 Decision Rule Approximation

The emergence of binary recourse variables in multi-stage models of the type $\mathcal{MP}$ adds another level of complexity to the two-stage models considered in Section 4.2. Indeed, while continuous recourse variables can be approximated by linear decision rules (Ben-Tal et al. [2004], Georghiou et al. [2010], Goh and Sim [2010], Shapiro and Nemirovski [2005]), there seems to be no flexible decision rule approximation for binary recourse variables which enjoys good tractability properties. Real-valued decision rules that are piecewise constant on the subsets of an *adjustable* partition of Ξ have been studied by Bertsimas and Caramanis [2010]. However, this adjustability entails considerable complications in the presence of endogenous uncertainties. We therefore approximate the measurement decisions in problem $\mathcal{MP}$ by binary-valued decision rules that

are piecewise constant with respect to a *preselected* partition of Ξ. Similarly, we approximate all real-valued decisions in $\mathcal{MP}$ by decision rules that are piecewise linear with respect to the same partition. Without much loss of generality, we assume that all subsets of this partition are hyper-rectangles of the form

$$\Xi_s := \{\xi \in \Xi : w^i_{s_i-1} \leq \xi_i < w^i_{s_i},\ i = 1, \ldots, k\},$$

where $s \in \mathbb{S} := \bigtimes_{i=1}^{k} \{1, \ldots, r_i\} \subseteq \mathbb{N}^k$ and

$$w^i_1 < w^i_2 < \cdots < w^i_{r_i-1} \text{ for } i = 1, \ldots, k$$

represent $r_i - 1$ breakpoints along the ξ_i axis. We now approximate the binary-valued decisions in $\mathcal{MP}$ by piecewise constant decision rules of the form

$$z_t(\xi) = \sum_{s \in \mathbb{S}} \mathbb{I}_{\Xi_s}(\xi) z^s_t \tag{4.3}$$

for some $z^s_t \in \{0,1\}^k$, $s \in \mathbb{S}$, $t \in \mathbb{T}$, where $\mathbb{I}_{\Xi_s}$ denotes the indicator function of Ξ_s. Similarly, we approximate the real-valued decisions in $\mathcal{MP}$ by piecewise linear decision rules of the form

$$y_t(\xi) = \sum_{s \in \mathbb{S}} \mathbb{I}_{\Xi_s}(\xi) Y^s_t \xi \tag{4.4}$$

for some $Y^s_t \in \mathbb{R}^{n_t \times k}$, $s \in \mathbb{S}$, $t \in \mathbb{T}$.

In order to reduce the notational overhead, we henceforth suppress the domains of the variables $\xi, \xi' \in \Xi$, $t \in \mathbb{T}$, $s, s' \in \mathbb{S}$, $j, j' \in \{1, \ldots, k\}$ and $i \in \{1, \ldots, n_t\}$.

Proposition 4.2 *Under the approximations* (4.3) *and* (4.4)*, the non-anticipativity constraints in* $\mathcal{MP}$ *are equivalent to*

$$\left.\begin{array}{l} |z^s_{t,j'} - z^{s'}_{t,j'}| \leq z^s_{t-1,j} \\ |Y^s_{t,ij'} - Y^{s'}_{t,ij'}| \leq M z^s_{t-1,j}\ \forall i \end{array}\right\} \quad \begin{array}{l} \forall j, j', s, s', t : \\ s_{-j} = s'_{-j} \end{array} \tag{4.5a}$$

$$|Y^s_{t,ij}| \leq M z^s_{t-1,j}\ \forall i, j, s, t, \tag{4.5b}$$

where M is a sufficiently large big-M constant.

Proof The non-anticipativity constraints in $\mathcal{MP}$ can be re-expressed as

$$\left.\begin{array}{l} z_t(\xi) = z_t(\xi') \\ y_t(\xi) = y_t(\xi') \end{array}\right\} \forall t, \xi, \xi' : z_{t-1}(\xi) \circ \xi = z_{t-1}(\xi') \circ \xi'.$$

Substituting (4.3) and (4.4) into the above expression yields

$$\left.\begin{array}{l} z_t^s = z_t^{s'} \\ Y_t^s = Y_t^{s'} \end{array}\right\} \forall t, s, s' : z_{t-1}^s \circ s = z_{t-1}^{s'} \circ s' \tag{4.6a}$$

and

$$|Y_{t,ij}^s| \leq M z_{t-1,j}^s \ \forall i, j, s, t. \tag{4.6b}$$

Note that (4.6a) enforces non-anticipativity across distinct subsets of the partition, while (4.6b) enforces non-anticipativity for the linear decision rules within each subset and is reminiscent of the non-anticipativity constraints in $\mathcal{P}'_u$. We now demonstrate that (4.5a) and (4.6a) are equivalent.

($\Leftarrow$) Assume that (4.6a) holds, and choose some j, s, s' and t with $s_{-j} = s'_{-j}$. The information monotonicity constraint stipulated in $\mathcal{MP}$ implies that

$$z_{\tau-1,j}^s = z_{\tau-1,j}^{s'} = 1 \ \Rightarrow \ z_{\tau,j}^s = z_{\tau,j}^{s'} = 1, \tag{4.7}$$

while (4.6a) and the assumption $s_{-j} = s'_{-j}$ imply that

$$z_{\tau-1,j}^s = z_{\tau-1,j}^{s'} = 0 \ \Rightarrow z_\tau^s = z_\tau^{s'} \text{ and } Y_\tau^s = Y_\tau^{s'} \tag{4.8}$$

for all $\tau \in \{0, \ldots, t-1\}$. Since $z_0^s = z_0^{s'} = \mathrm{e}_1$, we can iteratively apply (4.7) and the first

implication in (4.8) to conclude that $z^s_{t-1,j} = z^{s'}_{t-1,j}$. Thus, (4.8) implies

$$\left.\begin{array}{l} |z^s_{t,j'} - z^{s'}_{t,j'}| \leq z^s_{t-1,j} \\ |Y^s_{t,ij'} - Y^{s'}_{t,ij'}| \leq M z^s_{t-1,j} \ \forall i \end{array}\right\} \forall j'.$$

As j, s, s' and t were chosen arbitrarily, (4.5a) follows.

($\Rightarrow$) Assume now that (4.5a) holds and choose some s, s' and t with $z^s_{t-1} \circ s = z^{s'}_{t-1} \circ s'$. As $s_j, s'_j \geq 1 \ \forall j$, we conclude that $z^s_{t-1} = z^{s'}_{t-1}$. Thus, z^s_{t-1} and $s - s'$ satisfy the complementarity condition $z^s_{t-1} \circ (s - s') = 0$. If $s = s'$, then (4.6a) holds trivially true. Next, assume that $s \neq s'$ and that there exists j' with $s_{-j'} = s'_{-j'}$, that is s and s' differ only in their j'th component. The complementarity of z^s_{t-1} and $s - s'$ then ensures that $z^s_{t-1,j'} = 0$. Together with the known identity $s_{-j'} = s'_{-j'}$, this implies via (4.5a) that $z^s_t = z^{s'}_t$ and $Y^s_t = Y^{s'}_t$. Thus, (4.6a) follows. Finally, if s and s' differ in two or more components, then (4.6a) can be established by applying the above argument iteratively. □

4.3.3 MBLP Reformulation

Substituting the decision rules (4.3) and (4.4) into $\mathcal{MP}$ and applying Proposition 4.2 yields a conservative approximation $\mathcal{MP}_u$ for $\mathcal{MP}$. We then proceed as in Section 4.2.3 to obtain the following MBLP reformulation of $\mathcal{MP}_u$,

$$\begin{array}{ll} \min & \sum_{s\in\mathbb{S}} p_s \sum_{t\in\mathbb{T}} \mu_s^\top {F_t}^\top z^s_t + \mathrm{tr}(\Sigma_s {C_t}^\top Y^s_t) \\ \text{s.t.} & z^s_t \in \mathcal{Z}_t, \ Y^s_t \in \mathbb{R}^{n_t \times k} \ \forall s, t \\ & H_t - \sum_{\tau=1}^t A_{t\tau} Y^s_\tau + B_{t\tau} z^s_\tau \mathrm{e}_1^\top \in \mathcal{K}^{m_t}_s \ \forall s, t \\ & z^s_t \geq z^s_{t-1} \forall s, t \\ & \left.\begin{array}{l} |z^s_{t,j'} - z^{s'}_{t,j'}| \leq z^s_{t-1,j} \\ |Y^s_{t,ij'} - Y^{s'}_{t,ij'}| \leq M z^s_{t-1,j} \ \forall i \end{array}\right\} \begin{array}{l} \forall s, s' : s_{-j} = s'_{-j} \\ \forall t, j, j' \end{array} \\ & |Y^s_{t,ij}| \leq M z^s_{t-1,j} \ \forall s, t, i, j, \end{array} \qquad (\mathcal{MP}'_u)$$

where $p_s := \mathbb{P}(\xi \in \Xi_s)$, $\mu_s := \mathbb{E}(\xi|\xi \in \Xi_s)$, $\Sigma_s := \mathbb{E}(\xi\xi^\top|\xi \in \Xi_s)$ and $\mathcal{K}_s := (\text{cone}(\Xi_s))^*$. Problem $\mathcal{MP}'_u$ involves only a finite number of decision variables and constraints. For a fixed number of uncertain parameters k and fixed number of breakpoints along each coordinate axis in $\mathbb{R}^k$, the size of $\mathcal{MP}'_u$ remains polynomially bounded in $m := \sum_{t\in\mathbb{T}} m_t$, $n := \sum_{t\in\mathbb{T}} n_t$ and in the number of constraints defining the uncertainty set.

4.4 Case Study

4.4.1 Problem Description

We evaluate the proposed decision rule approach on (a variant of) an infrastructure and production planning problem in offshore oil fields from the literature (Goel and Grossman [2004]). An oil company has identified an offshore oil extraction site for possible exploitation. This site comprises several oil fields (or oil reservoirs) with unknown reserves. The company is assumed to be aware of the exact locations of the individual oil fields and needs to plan the oil extraction and gas production process over a period ranging from 10 to 30 years. The objective is to maximize the expected net present value (NPV) of the oil exploitation project.

In order to extract oil from the fields, dedicated well platforms need to be installed and expanded. We denote the set of candidate well platforms (that are under consideration to be built) by $\mathcal{W}$. The oil extracted at the well platforms is sent through a network of directed pipelines to a (unique) production platform $p \in \mathcal{W}$ for gas production. The set of candidate links between well platforms is denoted by $\mathcal{L}$. For any platform $w \in \mathcal{W}$ we denote by $\mathcal{L}^+(w) \subseteq \mathcal{L}$ and $\mathcal{L}^-(w) \subseteq \mathcal{L}$ the sets of all ingoing pipelines *to* w and all outgoing pipelines *from* w, respectively. We assume that all expansion and construction decisions take immediate effect and that once a platform $w \in \mathcal{W}$ has been built, the size ξ^w of the associated oil field is revealed.

We assume that the planning horizon is subdivided into yearly intervals indexed by $t \in \mathbb{T}$. At the beginning of each year, the oil company decides which new platforms and pipelines to construct. We set $z_t^w(\xi) = 1$ if platform w exists at time t; $= 0$ otherwise. Similarly, we set

$x_t^l(\xi) = 1$ if pipeline l exists at time t; $= 0$ otherwise. We assume that platforms and pipelines cannot be decommissioned, that is $z_t^w(\xi) \geq z_{t-1}^w(\xi)$ and $x_t^l(\xi) \geq x_{t-1}^l(\xi)$.

In year t the company selects the yearly oil extraction $y_{\mathrm{e},t}^w(\xi)$ for platform w, the yearly flow $y_{\mathrm{e},t}^w(\xi)$ through pipeline l, and the amount $y_{\mathrm{c},t}^w(\xi)$ by which the capacity of platform w is increased at the start of the year. The cumulative oil extraction at a particular field can never exceed the field size,

$$\sum_{\tau=1}^{t} y_{\mathrm{e},\tau}^w(\xi) \leq \xi^w \quad \forall w \in \mathcal{W},$$

while the instantaneous oil extraction is limited by the field's maximum production rate p_w, that is,

$$0 \leq y_{\mathrm{e},t}^w(\xi) \leq p^w \quad \forall w \in \mathcal{W}.$$

The flow conservation constraints

$$y_{\mathrm{e},t}^w(\xi) + \sum_{l \in \mathcal{L}^+(w)} y_{\mathrm{f},t}^l(\xi) \geq \sum_{l \in \mathcal{L}^-(w)} y_{\mathrm{f},t}^l(\xi) \quad \forall w \in \mathcal{W}$$

ensure that no oil is created within the network, and the box-constraints

$$0 \leq y_{\mathrm{f},t}^l(\xi) \leq M x_t^l(\xi) \quad \forall l \in \mathcal{L},$$

which involve a big-M constant, force the flows through yet inexistent pipelines to vanish. Similar box constraints guarantee that yet inexistent platforms cannot be expanded.

$$0 \leq y_{\mathrm{c},t}^w(\xi) \leq M z_t^w(\xi) \quad \forall w \in \mathcal{W}.$$

The yearly amount of oil pumped into the network from a particular platform must not exceed that platform's capacity, that is,

$$\sum_{l \in \mathcal{L}^-(w)} y_{\mathrm{f},t}^l(\xi) \leq \sum_{\tau=1}^{t} y_{\mathrm{c},\tau}^w(\xi) \quad \forall w \in \mathcal{W}.$$

The company chooses the design and operating decisions with the aim to maximize the project's

expected net present value (NPV)

$$\sum_{t\in\mathbb{T}} d_t \mathbb{E}\left\{ c^p_{\mathrm{f}} \sum_{l\in\mathcal{L}^-(p)} y^l_{\mathrm{f},t}(\xi) - \sum_{l\in\mathcal{L}} c^l_{\mathrm{b}}(x^l_t(\xi) - x^l_{t-1}(\xi)) \right.$$
$$\left. - \sum_{w\in\mathcal{W}} f^w(z^w_t(\xi) - z^w_{t-1}(\xi)) + c^w_{\mathrm{c}} y^w_{\mathrm{c},t}(\xi) + c^w_{\mathrm{e}} y^w_{\mathrm{e},t}(\xi) \right\},$$

where c^p_{f} denotes the unit price for gas, while f^w and c^l_{b} denote the costs for building platform w and pipeline l, respectively. Moreover, c^w_{c} and c^w_{e} represent the unit expansion and extraction costs for platform w, and d_t denotes the discount factor for year t. Note also that $\sum_{l\in\mathcal{L}^-(p)} y^l_{\mathrm{f},t}(\xi)$ represents the total outflow from the production platform, which coincides with the yearly gas production. All decisions selected at the start of year t may depend only on $z_{t-1}(\xi)\circ\xi$. Thus, if $(x_t)_{t\in\mathbb{T}}$ are interpreted as measurement decisions for fictitious degenerate random parameters, the oil extraction problem can be brought to the form $\mathcal{MP}$.

4.4.2 Numerical Results

We consider an instance of the oil extraction problem with a 15 year horizon at the offshore site shown in Figure 4.1(a). The field sizes are mutually independent and uniformly distributed as $\xi^w \sim \mathcal{U}(0, u^w)\ \forall w \in \mathcal{W}$. The input parameters of the problem are summarized in Table 4.1.

Table 4.1: Problem input parameters

Parameter	Value	Units
$(u^w)_{w\in\mathcal{W}}$	$(10, 10, 10, 20, 20)$	$10^9\mathrm{m}^3$
$(p^w)_{w\in\mathcal{W}}$	$(0.56, 0.56, 0.56, 1.1, 1.1)$	$10^9\mathrm{m}^3$/year
c^p_{f}	1.2	EUR/$10^3\mathrm{m}^3$
$(c^l_{\mathrm{b}})_{l\in\mathcal{L}}$	$(0, 2, 2, 5, 3, 3, 2)$	MEUR
$(f^w)_{w\in\mathcal{W}}$	$(5, 2, 2, 3, 3)$	MEUR
$(c^w_{\mathrm{c}})_{w\in\mathcal{W}}$	$(0.1, 0.1, 0.1, 0.1, 0.1, 0.1)$	EUR/$10^3\mathrm{m}^3$
$(c^w_{\mathrm{e}})_{w\in\mathcal{W}}$	$(0, 0, 0, 0, 0, 0)$	EUR/$10^3\mathrm{m}^3$
d_t	$1/(1+0.01)^{t-1}$	–

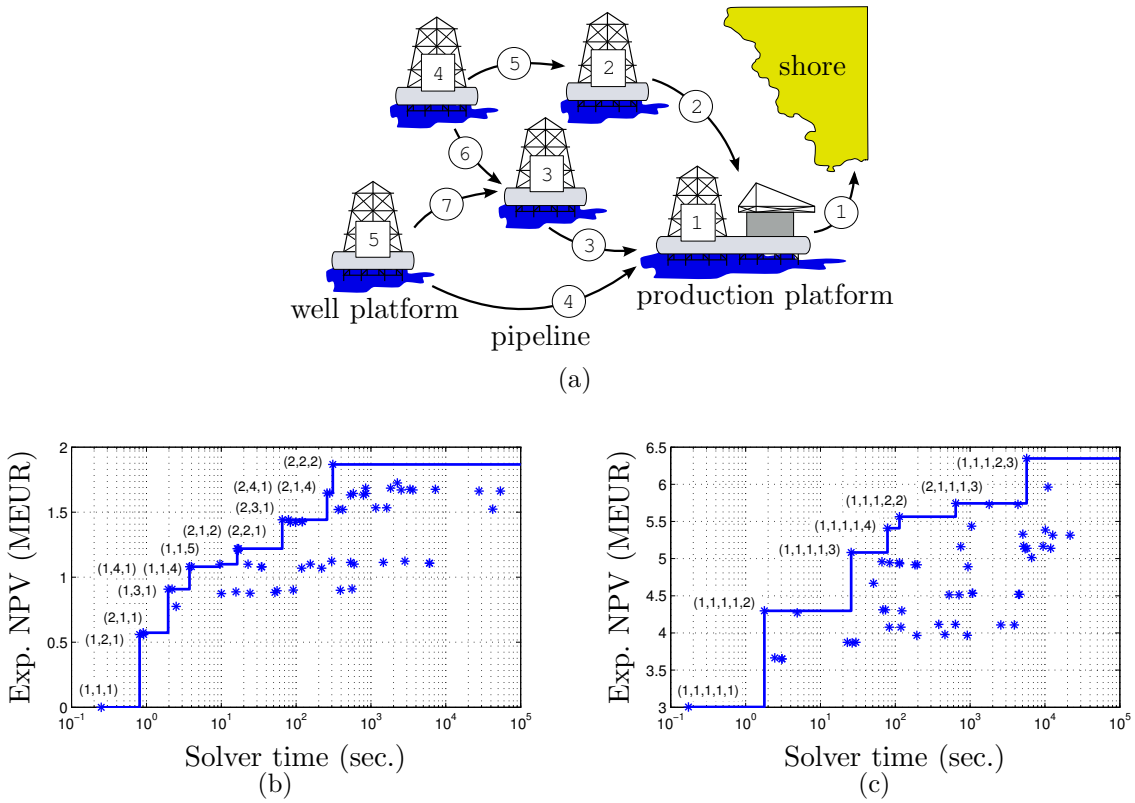


Figure 4.1: Figure (a) shows the offshore oil extraction site. The numbers in squares and circles indicate the platforms $w \in \mathcal{W}$ and the pipelines $l \in \mathcal{L}$. Figures (b) and (c) illustrate the expected NPV in dependence of solver time for projects A and B, respectively. The labels next to the markers represent the breakpoint configuration $(r_1, \ldots, r_k)$ for the problems that achieved the tightest approximation for a given time budget.

We consider two projects: project A aims at extracting oil from the fields 1 through 3, while project B considers all 5 fields. We proceed as described in Section 4.3.2 and Section 4.3.3 to obtain conservative solutions to the expected NPV maximization problems. The partitions of Ξ are constructed such that their subsets have equal probability. The results are shown on Figures 4.1(b) and 4.1(c). For project A (B), we consider all partitions with $|\mathbb{S}| \leq 12$ ($|\mathbb{S}| \leq 6$). The figures illustrate the increase in expected NPV achieved as the solver time[1] increases. For a time budget of less than 70 secs., an increase in expected NPV of more than 1.4 MEUR is achieved relative to the non-adaptive strategy which precommits the measurement decisions at time $t = 1$ and only allows for linear extraction and capacity expansion decisions. We note that for the case $|\mathbb{S}| = 1$, project A appeared not to be profitable. Finally, we remark that exploiting the full site results in a substantially higher profit.

[1] All problems were solved on a 2.66GHz Intel Core i7-920 machine running CPLEX 12.2.

Chapter 5

Conclusion

Dynamic decision problems affected by uncertainty are ubiquitous in numerous disciplines ranging from engineering to finance and economics. It is now widely recognized that disregarding this uncertainty often results in sub-optimal decisions or even in decisions that may fail to be implementable in practice. Unfortunately, all exact solution techniques suffer from the curse of dimensionality: their computational complexity grows exponentially with the number of decision stages or with the number of state variables.

In this thesis, we developed tractable approximation schemes for dynamic optimization problems affected by uncertainty. In what follows, we begin by summarizing our main contributions. We then discuss possible directions for future research.

In Chapter 2, we proposed a flexible data-driven approximation approach for generic multistage robust optimization problems which relied on a synthesis of decision rule and constraint sampling approaches. We introduced an axiomatic characterization of classes of decision rules that result in a tractable scenario counterpart and guarantee asymptotic consistency. We assessed the convergence and scalability properties of our approach in the context of two inventory management problems. In Chapter 3, we developed a dynamic hedging scheme for the valuation of path-dependent electricity derivatives, such as swing options, in incomplete markets. We formulated two robust optimization problems whose optimal values yield the end-points of the no-arbitrage interval. We considered the case where the holder is subjected to bounded

rationality and formulated a hedging problem for the writer in this case also. We analyzed and exploited the structure of the problem and the nature of the optimal decision rule to formulate approximate problems that can be solved efficiently with the approach discussed in Chapter 2. We showcased how certain approximation parameters can conveniently be used to model the negotiation process between holder and writer of the derivative contract. We illustrated the scalability properties of our approximation on a problem involving hundreds of decision stages and fifteen state variables in a market model with two risk-factors. Finally, in Chapter 4, we developed a tractable approximation scheme for solving decision problems where the decision-maker can control the sequence in which the uncertain parameters are revealed. To the best of our knowledge, this is the first scenario-free approach discussed in the literature for problems of this type. We highlighted the practical relevance of such problems and illustrated our approach on a realistic size instance of the problem of infrastructure and production planning in offshore oilfields.

During our research, we identified several interesting avenues for future work.

First, we believe that a methodology capable of assessing the loss of optimality incurred by the decision rule approximation developed in Chapter 2 would be of great interest to practitioners. Indeed, this would enable the modeler to weigh the benefits of considering more sophisticated decision rules and the computational cost of increasing the complexity of the approximation. A-posteriori methods for bounding the optimality gap have been introduced by Kuhn et al. [2009], Georghiou et al. [2010] and Bampou and Kuhn [2011] for stochastic programs subjected to linear, piecewise linear and polynomial decision rule approximations, respectively. The main challenge in extending these methods to problems where semi-analytical schemes are not applicable lies in the outer approximation entailed by the sampling. Indeed, due to the relaxation of the robust constraints, the optimal objective value of the scenario counterpart yields a lower (rather than upper) bound on the optimal objective value of the problem arising from the decision rule approximation and thus prevents us from consistently bounding the optimality gap.

Moreover, it would be interesting to develop a data-driven methodology for solving general

multi-stage distributionally robust optimization problems. Data-driven methodologies for linear single-stage distributionally robust problems have been investigated by Erdoğan and Iyengar [2006] and Wang et al. [2009] among others. Erdoğan and Iyengar [2006] develop a hybrid methodology combining constraint sampling and semi-analytical approaches and demonstrate that $\mathcal{O}(v/\epsilon)$ samples are needed to guarantee that the solution to the sampled problem is feasible in the associated distributionally robust problem (v denotes the Vapnik-Chervonenkis dimension of the maximum of the constraint functions). While their sample bounds are computationally attractive, they involve computing v which is not an easy task. In addition, extending this scheme to the multi-stage setting with general decision rule approximations cannot be readily done. Indeed, this methodology enforces the constraints robustly on balls centered about the samples and thus requires the use of semi-analytical approaches for reformulating the corresponding semi-infinite constraints. The likelihood robust optimization approach developed by Wang et al. [2009] would constitute a potentially more attractive starting point for this research.

Another interesting and promising research direction would involve analyzing the structure of the optimal decision rule in problems where the time of information discovery can be controlled by the decision-maker. The derivation of such *a-priori* results for decision problems affected by *exogenous* uncertainties is the subject of very active research, as exemplified by the recent publications of Bertsimas et al. [2010], Bertsimas and Goyal [2010, 2012] and Hadjiyiannis et al. [2011]. Nevertheless, to date, no results have been reported for the more general class of problems involving endogenous uncertainties. In order to complement this work, one could additionally develop a-posteriori bounds such as the ones described above. Furthermore, it would be interesting to investigate the convergence properties of the decision rule approximation scheme devised in Chapter 4 for problems with decision-dependent non-anticipativity constraints.

Finally, a problem which will undoubtedly receive considerable attention by the robust optimization community in the near future is that of dynamic decision-making under generic endogenous uncertainties, where the decision-maker's actions affect (in a general way) the distribution of the random parameters. These problems constitute a generalization of those

discussed in Chapter 4 in that the decision-maker does not gain *full* knowledge of the uncertain parameters following an investment, but instead, he may, for example, gain *partial* information (e.g., decrease in the variance of the distribution).

Bibliography

F. Alizadeh and D. Goldfarb. Second-order cone programming. *Mathematical Programming*, 95(1):3–51, 2003.

A. Atamtürk and M. Zhang. Two-stage robust network flow and design under demand uncertainty. *Operations Research*, 55(4):662–673, 2007.

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

O. Bardou, S. Bouthemy, and G. Pagès. When are swing options bang-bang and how to use it? *International Journal of Theoretical and Applied Finance*, 13(6):867–899, 2010.

E.M.L. Beale. On minimizing a convex function subject to linear inequalities. *Journal of the Royal Statistical Society. Series B*, 17(2):173–184, 1955.

A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

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

A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming*, 88(3):411–424, 2000.

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

A. Ben-Tal, B. Golany, A. Nemirovski, and J.-P. Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7(3):248–271, 2005.

A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, 2009.

D.P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, USA, 1995.

D. Bertsimas and C. Caramanis. Adaptability via sampling. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 4717–4722, 2007.

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

D. Bertsimas and V. Goyal. On the power of robust solutions in two-stage stochastic and adaptive optimization problems. *Operations Research*, 35(2):284–305, 2010.

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

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

D. Bertsimas, D. Iancu, and P. Parrilo. A hierarchy of near-optimal policies for multi-stage adaptive optimization. *IEEE Transactions on Automatic Control*, 56(12):2809–2824, 2011.

J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer, 2000.

V.D. Blondel and J.N. Tsitsiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9):1249–1274, 2000.

M.J. Brennan and E.S. Schwartz. Evaluating natural resource investments. *Journal of Business*, 58(2):135–157, 1985.

G.C. Calafiore. Multi-period portfolio optimization with linear control policies. *Automatica*, 44(10):2463–2473, 2008.

G.C. Calafiore. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM Journal of Control and Optimization*, 48(4):2254–2274, 2009a.

G.C. Calafiore. On the expected probability of constraint violation in sampled convex programs. *Journal of Optimization Theory and Applications*, 143(2):405–412, 2009b.

G.C. Calafiore. Random convex programs. *SIAM Journal on Optimization*, 20(6):3427–3464, 2010.

G.C. Calafiore and M.C. Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, 2005.

G.C. Calafiore and M.C. Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.

G.C. Calafiore and L. El Ghaoui. On distributionally robust chance-constrained linear programs. *Journal of Optimization Theory and Applications*, 130(1):1–22, 2006.

G.C. Calafiore and A. Nilim. Multi-stage resource allocation under uncertainty. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, volume 1, pages 417–423, 2004.

M.C. Campi and S. Garatti. The exact feasibility of randomized solutions of robust convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008.

M.C. Campi and S. Garatti. A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality. *Journal of Optimization Theory and Applications*, 148 (2):257–280, 2011.

R. Carmona and M. Ludkovski. Swing options. In R. Cont, editor, *Encyclopedia of Quantitative Finance*. Wiley, 2010.

R. Carmona and N. Touzi. Optimal multiple stopping and valuation of swing options. *Mathematical Finance*, 18(2):239–268, 2008.

A. Charnes and W.W. Cooper. Chance-constrained programming. *Management Science*, 6(1): 73–79, 1959.

X. Chen, M. Sim, and P. Sun. A robust optimization perspective of stochastic programming. *Operations Research*, 55(6):1058–1071, 2007.

M. Colvin and C.T. Maravelias. A stochastic programming approach for clinical trial planning in new drug development. *Computers & Chemical Engineering*, 32(11):2626–2642, 2008.

M. Colvin and C.T. Maravelias. Modeling methods and a branch and cut algorithm for pharmaceutical clinical trial planning using stochastic programming. *European Journal of Operational Research*, 203(1):205–215, 2010.

G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4):303–314, 1989.

G.B. Dantzig. Linear programming under uncertainty. *Management Science*, 1(3–4):197–206, 1955.

E. Delage and Y. Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, 2010.

S.-J. Deng and Z. Xia. Pricing and hedging electricity supply contracts: a case with tolling agreements. Technical report, School of Industrial and Systems Engineering, Georgia Institute of Technology, 2005.

S.-J. Deng, B. Johnson, and A. Sogomonian. Exotic electricity options and the valuation of electricity generation and transmission assets. *Decision Support Systems*, 30(3):383–392, 2001.

G.E. Dullerud and F. Paganini. *A Course in Robust Control Theory*. Springer, 2005.

J. Dupačová. On minimax solutions of stochastic linear programming problems. *Èasopis pro pìstování matematiky*, 91(4):423–430, 1966.

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

M. Dyer and L. Stougie. Computational complexity of stochastic programming problems. *Mathematical Programming*, 106(3):423–432, 2006.

L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, 18(4):1035–1064, 1997.

L. El Ghaoui, F. Oustry, and H. Lebret. Robust solutions to uncertain semidefinite programs. *SIAM Journal on Optimization*, 9(1):33–52, 1998.

D. Ellsberg. Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75 (4):643–669, 1961.

E. Erdoğan and G. Iyengar. Ambiguous chance constrained problems and robust optimization. *Mathematical Programming*, 107(1–2):37–61, 2006.

A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. Available on Optimization Online, submitted for publication, 2010.

A. Geyer, M. Hanke, and A. Weissensteiner. No-arbitrage conditions, scenario trees, and multi-asset financial optimization. *European Journal of Operational Research*, 206(3):609–613, 2010.

F. Girosi and T. Poggio. Networks and the best approximation property. *Biological Cybernetics*, 63(3):169–176, 1990.

V. Goel and I.E. Grossman. A stochastic programming approach to planning of offshore gas field developments under uncertainty in reserves. *Computers & Chemical Engineering*, 28(8): 1409–1429, 2004.

V. Goel and I.E. Grossman. A class of stochastic programs with decision dependent uncertainty. *Mathematical Programming*, 108(2):355–394, 2006.

V. Goel, I.E. Grossman, A.A. El-Bakry, and E.L. Mulkay. A novel branch and bound algorithm for optimal development of gas fields under uncertainty in reserves. *Computers & Chemical Engineering*, 30(6–7):1076–1092, 2006.

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

M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.

V. Gupta and I.E. Grossman. Solution strategies for multistage stochastic programming with endogenous uncertainties. Submitted for publication, 2010.

G. Haarbrücker and D. Kuhn. Valuation of electricity swing options by multistage stochastic programming. *Automatica*, 45(4):889–899, 2009.

M. Hadjiyiannis, P. Goulart, and D. Kuhn. A scenario approach for measuring the suboptimality of linear decision rules in two-stage robust optimization. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pages 7386–7391, 2011.

T. Hao. Option pricing and hedging bounds in incomplete markets. *Journal of Derivatives & Hedge Funds*, 14(2):78–89, 2008.

M. J. Harrison and D. M. Kreps. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, 20(3):381–408, 1979.

A. Ibáñez. Valuation by simulation of contingent claims with multiple exercise opportunities. *Mathematical Finance*, 14(2):223–248, 2004.

P. Jaillet, E. I. Ronn, and S. Tompaidis. Valuation of commodity-based swing options. *Management Science*, 50(7):909–921, 2004.

T.W. Jonsbråten. *Optimization models for petroleum field exploitation.* PhD thesis, Norwegian Shool of Economics and Business Administration, 1998.

T.W. Jonsbråten, R.J.-B. Wets, and D.L. Woodruff. A class of stochastic programs with decision dependent random elements. *Annals of Operations Research*, 82(0):83–106, 1998.

P. Kall. *Advances in Mathematical Optimization and Related Topics*, chapter Stochastic programs with recourse: upper bounds and moment problems - a review, pages 86–103. Akademie-Verlag Berlin, 1988.

P. Kall and S. Wallace. *Stochastic Programming*. Wiley-Interscience Series in Systems & Optimization. John Wiley & Sons, 2nd edition, 1994.

A.J. King. Duality and martingales: a stochastic programming perspective on contingent claims. *Mathematical Programming*, 91(3):543–562, 2002.

D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2009.

R. Lobel and G. Perakis. Dynamic pricing through sampling based optimization. Technical report, Massachusetts Institute of Technology, 2010. Submitted for publication.

F. Longstaff and E. Schwartz. Valuing american options by simulation: a simple least squares approach. *The Review of Financial Studies*, 14(1):113–147, 2001.

M. Ludkovski. *Optimal switching with applications to energy tolling agreements*. PhD thesis, Princeton university, 2005.

M. Ludkovski and R. Carmona. Valuation of energy storage: An optimal switching approach. *Quantitative Finance*, 10(4):359–374, 2010.

D. Luenberger. *Investment Science*. Oxford University Press, 1997.

H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

B.L. Miller and M.W. Wagner. Chance constrained programming with joint constraints. *Operations Research*, 13(6):930–945, 1965.

B.K. Pagnoncelli, S. Ahmed, and A. Shapiro. Sample average approximation method for chance constrained programming: theory and applications. *Journal of Optimization Theory and Applications*, 142(2):399–416, 2009.

I. Popescu. Robust mean-covariance solutions for stochastic optimization. *Operations Research*, 55(1):98–112, 2007.

A. Prékopa. On probabilistic constrained programming. In *Proceedings of the Princeton Symposium on Mathematical Programming*, pages 113–138, 1970.

A. Prékopa. *Stochastic Programming*. Kluwer Academic Publishers, 1995.

P. Rocha and D. Kuhn. Multistage stochastic portfolio optimization in deregulated electricity markets using linear decision rules. *European Journal of Operational Research*, 216(2):397–408, 2012.

R.T. Rockafellar and R.J.-B. Wets. *Variational Analysis*. Springer-Verlag, 1997.

H. Scarf. *Studies in the Mathematical Theory of Inventory and Production*, chapter A min-max solution of an inventory problem, pages 201–209. Stanford University Press, 1958.

C.-T. See and M. Sim. Robust approximation to multiperiod inventory management. *Operations Research*, 58(3):583–594, 2009.

A. Shapiro. Worst-case distribution analysis of stochastic programs. *Mathematical Programming*, 107(1):91–96, 2006.

A. Shapiro and A. Nemirovski. *Continuous Optimization: Current Trends and Applications*, chapter On complexity of stochastic programming problems, pages 111–114. Springer, 2005.

A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. MPS/SIAM Series on Optimization. SIAM, 2009.

H.A. Simons. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69 (1):99–118, 1955.

J. Skaf and S. Boyd. Nonlinear Q-design for convex stochastic control. *IEEE Transactions on Automatic Control*, 54(10):2426–2430, 2009.

S. Solak, J.-P.B. Clarke, E.L. Johnson, and E.R. Barnes. Optimization of R&D project portfolios under endogenous uncertainty. *European Journal of Operational Research*, 207(1): 420–433, 2010.

J. Staum. Incomplete markets. In J. R. Birge and V. Linetsky, editors, *Financial Engineering*, volume 15 of *Handbooks in Operations Research and Management Science*, chapter 12, pages 511–563. Elsevier, 2007.

A. Thompson. Valuation of path-dependent contingent claims with multiple exercise decisions over time: the case of take-or-pay. *Journal of Financial and Quantitative Analysis*, 30(2): 271–293, 1995.

L. Vandenberghe and S. Boyd. *Convex Optimization*. Cambridge University Press, 2004.

P. Vayanos, D. Kuhn, and B. Rustem. Decision rules for information discovery in multi-stage stochastic programming. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pages 7368–7373, 2011a.

P. Vayanos, W. Wiesemann, and D. Kuhn. Hedging electricity swing options in incomplete markets. In *Proceedings of the 18th IFAC World Congress*, pages 846–853, 2011b.

P. Vayanos, D. Kuhn, and B. Rustem. A constraint sampling approach for multi-stage robust optimization. *Automatica*, 48(3):459–471, 2012.

Z. Wang, P. Glynn, and Y. Ye. Likelihood robust optimization for data-driven newsvendor problem. Working paper, 2009.

S. Yu-Ting. A "natural" enumeration of non-negative rational numbers – an informal discussion. *The American Mathematical Monthly*, 87:25–29, 1980.

S. Zymler, D. Kuhn, and B. Rustem. Distributionally robust joint chance constraints with second-order moment information. *Mathematical Programming*, 2011. Accepted, to appear.