Technische Universität München

Lehrstuhl für Logistik und Supply Chain Management

Optimal Procurement and Inventory Control in Volatile Commodity Markets

Advances in Stochastic and Data-Driven Optimization

Christian Mandl, M.Sc.

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

Vorsitzender:		Univ.-Prof. Dr. Martin Grunow
Prüfer der Dissertation:	1.	Univ.-Prof. Dr. Stefan Minner
	2.	Univ.-Prof. Srinagesh Gavirneni, PhD Cornell University Ithaca, NY, USA

Die Dissertation wurde am 04.04.2019 bei der Technischen Universität München eingereicht und durch die Fakultät für Wirtschaftswissenschaften am 15.05.2019 angenommen.

To my parents

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my supervisor Professor Stefan Minner for his continuous support and very constructive criticism throughout my PhD studies. I also want to thank him for encouraging and supporting me to spend five months of my PhD program abroad. I am very grateful that I was given the opportunity to be part of his research group.

I also want to express my gratitude to Professor Srinagesh Gavirneni from Samuel Curtis Johnson Graduate School of Management at Cornell University for being part of the examination committee and especially for inviting me to visit Cornell from May to September 2018, which was a great experience. I also want to thank Professor Martin Grunow for being the chairman of the examination committee.

Further, I am very grateful for the support of my former and current colleagues at the chair of Logistics and Supply Chain Management at TUM: Yuka Akasaka, Szymon Albiński, Dr. Christian Bohner, Lachlan Bridges, Dr. Pirmin Fontaine, Michael Keilhacker, Dr. Miray Közen, Sebastian Malicki, Layla Martin, Santiago Nieto-Isaza, Thitinan Pholsook, Dr. Partricia Rogetzer, Dr. Martin Stößlein, Josef Svoboda, Dr. Florian Taube, Dr. Dariush Tavaghof Gigloo and Francesco Zangaro. Thank you for all the research discussions and also for the after-work activities. I thank Evelyn Gemkow for her steady support and extensive proofreading of my work.

Moreover, I would like to thank all the smart people I met during my time at Cornell for the inspiring discussions. I especially want to thank Assistant Professor Selva Nadarajah from the University of Illinois at Chicago for our joint work on Chapter 6 of this thesis.

Finally, I would like to thank my parents and my brother Stefan for their constant support and encouragement over the last years. A special thanks goes to my wonderful girlfriend Anja for her love and patience - also and especially during the time of my research stay far-off in the U.S. - for which I am deeply grateful to her.

Abstract

Volatile prices constitute a challenge for both commodity-processing and commodity-trading firms. This thesis investigates the implications of price uncertainty on the optimal operating policies in multi-period procurement and inventory control. A central contribution to the existing literature that addresses the *full information problem* is the focus on the implications of price model uncertainty, i.e., incomplete information about the underlying price process. Based on advances in stochastic and data-driven optimization, we propose mathematical models for practical decision support and test them on real data. Hence, this thesis gives guidance to managers in the digital age on how to use real-time information and *Big Data* in combination with methods from statistical learning theory (Bayesian learning, machine learning) in an optimization framework in order to improve commodity procurement and inventory management decisions.

The first problem considers operational hedging via inventory control. We show how a Bayesian belief structure can be used to express uncertainty about the price process, which is subject to switches in regimes. We prove the structure of the optimal storage policy and test its cost impact relative to several more practical but suboptimal control policies. We find that Bayesian learning yields significant cost savings.

The second problem addresses commodity procurement via forward contracting. We propose a data-driven and machine learning-enabled mixed integer linear programming model that jointly optimizes forecasts and decisions by training optimal purchase signals as functions of features related to the price. Finally, we quantify the performance loss caused by ignoring feature information in procurement.

The third problem considers optimal commodity storage from the perspective of a merchant with buying, storing and reselling opportunities. We propose several data-driven models for storage optimization. Based on empirical data of six major exchange-traded commodities, we find that optimally structured data-driven policies can outperform state-of-the-art reoptimization approaches.

Keywords: price uncertainty; procurement; inventory control; stochastic and data-driven optimization

Contents

List of Tables ix

List of Figures xi

List of Abbreviations xv

1. Introduction 1

1.1. Motivation . . . 1
1.2. Classification, Contribution and Research Questions . . . 4

2. Fundamentals of Commodity Markets 9

2.1. Commodity Price Risk . . . 9
2.2. Financial and Operational Risk Management . . . 10
2.3. Stochastic Modeling of Commodity Prices . . . 16

3. Related Literature 19

3.1. Commodity Finance . . . 19
3.2. Commodity Operations . . . 20
3.2.1. Inventory Control under Stochastic Purchase Price . . . 21
3.2.2. Inventory Control under Stochastic Purchase and Sales Price . . . 22
3.2.3. Financial Contracting under Stochastic Purchase Price . . . 24
3.3. Stochastic Optimization with Partial Information and Learning . . . 25
3.3.1. Bayesian Inventory Control under Partial Information . . . 25
3.3.2. Data-Driven and Machine Learning-Enabled Optimization . . . 26

4. Operational Hedging from a Bayesian Inventory Control Perspective 29

4.1. Introduction . . . 30
4.2. Model Formulation . . . 33
4.2.1. MRS Spot Price Model and Bayesian Updating Scheme . . . 33
4.2.2. Inventory Control Model . . . 35

4.3. Optimal Policy Structure and Monotonicity Properties 37
4.3.1. Optimality of Price(s)- and Belief-Dependent Base-Stock Policies 37
4.3.2. Monotonicity of the Optimal Base-Stock Functions 38
4.3.3. Illustration of Policy Structure and Monotonicity Properties . . . 40
4.4. Controlled Numerical Study 43
4.4.1. Setup 43
4.4.2. Cost of Price Regime Misspecification 44
4.4.3. Performance of Suboptimal Control Policies 46
4.5. Results on Empirical Data 49
4.6. Conclusion 54

5. Financial Hedging from a Data-Driven Procurement Perspective 59
5.1. Introduction 60
5.2. Model Formulation 64
5.2.1. Problem Setting 64
5.2.2. Linear Decision Rule Approximation 66
5.2.3. Data-Driven Models for Policy Parameter Optimization 68
5.2.4. Data-Driven Models under ML-Based Regularization 70
5.3. Performance Bounds and Performance Metrics 72
5.4. Controlled Numerical Study 76
5.4.1. Setup 76
5.4.2. Results 79
5.5. Results on Empirical Data 84
5.5.1. Setup 84
5.5.2. Results 87
5.6. Conclusion 90

6. Commodity Storage from a Data-Driven Merchant Perspective 91
6.1. Introduction 92
6.2. Model Formulation 97
6.2.1. Problem Setting 97
6.2.2. Optimal and Myopic Policy for Fully Flexible Storage 100
6.2.3. Optimal and Myopic Policy for Limited Flexible Storage 103
6.2.4. Sequential Reoptimization: The Rolling Intrinsic Approach 106
6.3. Data-Driven Optimization for the SCWP 108
6.3.1. Unstructured Linear Decision Rule Approach (DDA-LDR) 108

6.3.2. Optimally Structured Policy Approach (DDA-OSP) 111
6.3.3. Value Function Approximation Approach (DDA-VFA) 115
6.4. Results on Empirical Data 117
6.4.1. Setup and Descriptive Analysis 117
6.4.2. Deterministic Analysis 120
6.4.3. Stochastic Analysis: Performance of RIA 124
6.4.4. Stochastic Analysis: Performance of DDA 131
6.5. Conclusion 135

7. Conclusion **137**
7.1. Summary of Insights 137
7.2. Directions for Future Research 138

A. Appendix of Chapter 4 **141**
A.1. Proof of Theorem 1 141
A.2. Proof of Proposition 1 143
A.3. Cost of Price Regime Misspecification 145
A.4. Performance of Suboptimal Control Policies 146

B. Appendix of Chapter 5 **147**
B.1. Proof of Theorem 2 147
B.2. Model Formulation for the Best Subset Selection Problem 150
B.3. Model Formulation with Indicator Constraints 150
B.4. Computation Times 151
B.5. Results of the Controlled Numerical Study 151
B.6. Results on Empirical Data: Procurement of Natural Gas 154

C. Appendix of Chapter 6 **157**
C.1. SCWP with Fixed Costs, Storage Efficiency, Demand and Market Power 157
C.2. Empirical Forecast Error of Futures Prices 158
C.3. Deterministic Analysis under Perfect Foresight 159
C.3.1. Profit over Time 159
C.3.2. Optimal versus Myopic Performance 160
C.3.3. Impact of the Planning Horizon 162
C.4. Stochastic Analysis: Performance of RIA and DDA 164

Bibliography **167**

List of Tables

4.1. Summary of optimality and monotonicity results 40
4.2. Corn: Performance of different control policies 51
4.3. Average inventory in tons and average price paid in USD/ton 54
4.4. Potential of MRS price models in inventory control 55

5.1. General notation (Additional notation defined as required) 68
5.2. Summary of the numerical design 78
5.3. In-sample and out-of-sample VIEO in % for DDA-BD vs. SEO and DDA-ML vs. SEO-AIC under constant demand 83
5.4. Prescription error (PE): % above perfect foresight cost C^{PF} 89
5.5. Average annual savings (07-2008 to 06-2017) by DDA-ML1 compared to various benchmark policies for the real-world setting at our industry partner with annual gas demand of 10×10^6 MWh 89

6.1. Summary of policy parameter characterization 105
6.2. Commodity spot and futures price data (2000-2017) 117
6.3. MAPE, MASE and MDA of futures to predict spot prices (2000-2017) . 120
6.4. Summary of the numerical design 124
6.5. Performance of futures-based RIA in $V^{\mathrm{RIA}}/V^{\mathrm{PF}} \cdot 100\%$ across all instances from Table 6.4 125
6.6. Dominance matrix of $\mathrm{RIA_n}$ in % of instances ($\# = 1,296$) in which row performs strictly better than column 126
6.7. Performance of futures-based RIA, analyst-based RIA and AR(1)-based RIA in $V^{\mathrm{RIA}}/V^{\mathrm{PF}} \cdot 100\%$ from 2008 until 2017 across all instances (In-sample performance of AR(1)-based RIA in parenthesis) 129
6.8. Dominance Matrix of different RIA policies in % of instances ($\# = 720$) in which row performs strictly better than column 130
6.9. Benchmark approaches and feature data 131
6.10. Backtesting setup 131

6.11. Performance of DDA-OSP in $V/V^{\mathrm{PF}} \cdot 100\%$ with respect to storage flexibility 134

A.1. $\Delta\mathrm{COST}(\phi^{(1)}(\phi^{(2)}))$ in % (In gray: No speculation motives) 145
A.2. Performance of suboptimal control policies in % above optimal cost . . . 146

B.1. Historical monthly price and feature data (07-2007 to 06-2017) 154
B.2. Procurement of natural gas: Average purchase cost in Euro/MWh . . . 154
B.3. Dominance matrix: % of out-of-sample sub-periods in which a policy (row) performs strictly better than another (column) 155

C.1. Average performance (2000-2017) of futures-based RIA compared to RIA with perfect one-step directional forecasts (PDF) and RIA with perfect one-step point forecasts (PPF) in $V/V^{\mathrm{PF}} \cdot 100\%$ 164
C.2. Out-of-sample performance of futures-based IA and futures-based RIA compared to different data-driven policies measured in $V/V^{\mathrm{PF}} \cdot 100\%$ from 2002 until 2017 across 192 instances 165
C.3. Out-of-sample dominance matrix of different policies in % of instances (# = 192) in which row performs strictly better than column 166

List of Figures

1.1. Problem perspectives of this thesis and applied methodologies 5

2.1. Mean annualized price volatility σ of major commodities from 2000 to 2017 relative to corporate stocks and exchange rates (Data source: Thomson Reuters Datastream). 10
2.2. NYMEX natural gas futures curves from 01-2017 to 12-2017 (Prices in USD/mmbtu refer to closing prices at the first trading day of the corresponding month) 16

4.1. Spot price of CBOT corn [Cts per bushel] and LME zinc [USD per ton] 30
4.2. Concept of MRS in the spot market context 34
4.3. Intra-period sequence of events 35
4.4. Optimal inventory policy $S_t(p_t, \vec{\pi}_t)$ for i.i.d. regime specifications 41
4.5. Optimal inventory policy $S_t(p_t, p_{t-1}, \vec{\pi}_t)$ for Markovian regime specifications 42
4.6. Cost of correlation misspecification (Instances: $c_h = 1$, deterministic demand; Results for all other instances are reported in Table A.1; Boxplot characteristics throughout this thesis: minimum, 1st-, 2nd-, 3rd-quartile, maximum, mean ($\times$)) 44
4.7. Cost of price process misspecification (Instances: $c_h = 1$; Results for all other instances are reported in Table A.1) 45
4.8. Impact of demand volatility and speculation on cost of misspecification ($n = 2$) 46
4.9. Performance of suboptimal control policies (Instances: $c_h = 1$, $r = 0.49$; Results for all other instances are reported in Table A.2) 47
4.10. Regime switches and inventory policy implications ($c_h = \$0.4$/bu/mo) . 50
4.11. Price forecast accuracy vs. inventory performance 52
4.12. Zinc: Performance of different control policies 53

5.1. DDA embedded in the IT architecture of our industry partner 62
5.2. Training and evaluation framework 66
5.3. Out-of-sample prescription error (PE) of different procurement policies conditional on training set size T, price process type (RW, MR) and price process noise levels σ_{ϵ_t} across 100 simulations (For the results of $\sigma_{\epsilon_t} \in \{10, 20\}$, see Figure B.3 in Appendix B.5) 80
5.4. Average out-of-sample value of feature information (VFI) of DDA-BD (left), DDA-ML1 (right, solid lines) and DDA-ML2 (right, dashed lines) for different training set sizes T, price process types (RW, MR) and price process noise levels σ_{ϵ_t} across 100 simulations (Full descriptive statistics of the results in Figure B.4 of Appendix B.5) 81
5.5. TTF natural gas spot price evolution from 07-2007 until 06-2017 85
5.6. Relation between monthly TTF spot and futures contract prices [Euro/MWh] from 07-2007 to 06-2017 88

6.1. NYMEX futures curves (dashed) and realized spot prices ($\circ$) for natural gas 93
6.2. Sequence of events 98
6.3. Policy and monotonicity illustration of SCWP-FF versus SCWP-LF . . 104
6.4. RIA and MYOPIC with regard to prediction error 108
6.5. Optimization (training) and evaluation framework 110
6.6. Commodity spot prices and mean annualized volatilities σ (2000-2017) . 118
6.7. Gas storage: Performance of optimal vs. myopic policies (2000-2017) . . 121
6.8. Gas storage: Performance impact of planning horizon n (2000-2017) . . 122
6.9. Profit loss by ignoring injection and withdrawal capacity underutilization 123
6.10. Average performance of $\mathrm{RIA_n}$ in $V^{\mathrm{RIA_n}}/V^{\mathrm{PF}} \cdot 100\%$ across all instances . 125
6.11. Value of reoptimization for $n \in \{1, 3, 6, 12\}$ (light gray to dark gray) . . . 127
6.12. Average performance of RIA (black) versus IA (gray) in $V^{\mathrm{RIA}}/V^{\mathrm{PF}} \cdot 100\%$ for different time series phases of copper, gold, crude oil, natural gas, corn and soybean 127
6.13. Value of perfect one-step-ahead directional forecasts (PDF) and perfect one-step-ahead point forecasts (PPF) for RIA across all commodities and instances from 2000 until 2017 128
6.14. Forecast accuracy 2008-2017 (MAPE, MDA) of futures-based predictions (F) and analyst-based predictions (A) for horizons of 1 up to 12 months . 130

6.15. Out-of-sample performance of the different policies in $V/V^{\text{PF}} \cdot 100\%$ from 2002 until 2017 across the 192 instances . . . 132

B.1. Computation times for training of DDA-BD (Average across 100 runs) . 151
B.2. Exemplary sample paths of spot and forward prices under random walk (RW) and mean reversion (MR) assumptions and price process noise $\sigma_{\epsilon_t} = 10$ . . . 151
B.3. Out-of-sample prescription error (PE) of different procurement policies for $\sigma_{\epsilon_t} \in \{10, 20\}$ across 100 simulations conditional on training set size T and price process types (RW, MR) . . . 152
B.4. Out-of-sample value of feature information (VFI) of DDA-BD, DDA-ML1 and DDA-ML2 across 100 simulations for different training sets T, price process types (RW, MR) and noise levels σ_{ϵ_t} . . . 153

C.1. Distribution of prediction error e_t in % for copper (dark gray), gold, crude oil, natural gas, corn and soybean (light gray) (2000-2017) . . . 158
C.2. Average profit under perfect foresight for different degrees of warehouse flexibility (SCWP-FF, SCWP-LF) and frictions ($\eta = \eta^{\mathbf{i}} = \eta^{\mathbf{o}}$) . . . 159
C.3. Copper storage: Performance of optimal vs. myopic policies (2000-2017) 160
C.4. Gold storage: Performance of optimal vs. myopic policies (2000-2017) . 160
C.5. Oil storage: Performance of optimal vs. myopic policies (2000-2017) . . 161
C.6. Corn storage: Performance of optimal vs. myopic policies (2000-2017) . 161
C.7. Soybean storage: Performance of optimal vs. myopic policies (2000-2017) 161
C.8. Copper storage: Performance impact of planning horizon n (2000-2017) . 162
C.9. Gold storage: Performance impact of planning horizon n (2000-2017) . . 162
C.10. Oil storage: Performance impact of planning horizon n (2000-2017) . . . 163
C.11. Corn storage: Performance impact of planning horizon n (2000-2017) . . 163
C.12. Soybean storage: Performance impact of planning horizon n (2000-2017) 163

List of Abbreviations

AR(1)	First-Order Auto-Regressive Process
CBOT	Chicago Board Of Trade
CEC	Certainty Equivalent Control
CME	Chicago Mercantile Exchange
COMEX	Commodity Exchange
DDA	Data-Driven Approach
ERM	Empirical Risk Minimization
IA	Intrinsic Approach
LDR	Linear Decision Rule
LME	London Metal Exchange
LP	Linear Program(ming)
MAPE	Mean Absolute Percentage Error
MASE	Mean Absolute Scaled Error
MDA	Mean Directional Accuracy
MILP	Mixed Integer Linear Program(ming)
ML	Machine Learning
MO	Momentum
MR	Mean Reversion
MRS	Markov Regime Switching
NYMEX	New York Mercantile Exchange
PF	Perfect Foresight
RCM	Regime Classification Measure
RIA	Rolling Intrinsic Approach
RW	Random Walk
SCWP	Stochastic Commodity Warehouse Problem
SDP	Stochastic Dynamic Program(ming)
SEO	Separated Estimation and Optimization
SRC	Single Regime Control

Chapter 1.

Introduction

1.1. Motivation

Price uncertainty at commodity markets constitutes a significant exogenous risk factor for companies. Almost all (manufacturing) firms are exposed to commodity price risk that affects the direct costs of raw materials, packaging materials, energy consumed in operations or transportation costs. Companies that purchase material in foreign currencies face a similar risk due to exchange rate fluctuations. According to Beschaffungsmanagement (2009), commodities make up 27% of the total costs of a firm in the sector of mechanical and plant engineering, 47% in the automotive supply industry, 56% in the packaging industry and 66% in the agri-food industry.

Moreover, commodity markets have undergone some dramatic changes during the last decades: commodities exhibit rising price volatility and at the same time the liquidity of forward markets has increased, which enabled hedging activities (Geman, 2005, pp. 21). While commodity-processing firms intend to hedge against the adverse effects of price risk, commodity merchants strive to exploit inter-temporal price differentials (*purchase low, store, sell high*). However, for both purchasers and merchants, wrong decisions on managing price uncertainty can quickly lower profit margins.

Therefore, firms apply various price risk mitigation strategies, such as substituting raw materials, negotiating contractual price escalator clauses for risk sharing with both suppliers and customers, financial hedging through forward contracting or operational hedging through effective inventory control (Zsidisin and Hartley, 2012, chap. 5).

However, decision-making under price uncertainty is a complex and challenging optimization task. It requires an adequate stochastic model of the commodity price evolution (*finance discipline*) and a stochastic optimization model that determines the optimal operating policy for managing procurement and storage with respect to operational constraints, such as demand satisfaction, capacity restrictions or storage injection and withdrawal rate limits (*operations research discipline*).

Traditionally, literature treats operations and financial decisions independently. This is consistent with the Modigliani-Miller theorem (Modigliani and Miller, 1958) that states that in the absence of agency costs, taxes, bankruptcy costs, asymmetric information and market inefficiencies, the value of a firm is not affected by its financial (e.g., hedging) decisions. However, those assumptions are hardly ever met in practice, which provides rationale for risk management. Markets are not frictionless and often inefficient to some extent, taxes exist (Smith and Stulz, 1985), market participants do not necessarily have identical information (DeMarzo and Duffie, 1991), or suffer from hard budget constraints and costly external capital (Froot et al., 1993). Furthermore, firms might be risk-averse (Gaur and Seshadri, 2005). Consequently, commodity risk management can increase or decrease the value of a firm such that operational and financial risk management can be beneficial (or harmful).

A specific example for operational decision-making under price uncertainty, that refers to *Chapter 4* of this thesis, is optimal inventory management under random demand and price, i.e., how to effectively control inventory in order to avoid stock-outs and simultaneously exploit low purchase prices. If a commodity's price is anticipated to increase, an inventory manager may purchase more than required and if the price is expected to decrease, one may wait with purchasing. Stockpiling inventory is particularly prevalent for firms that want to exploit inter-temporal price differentials (e.g., commodity-trading firms) and for firms that have warehouse space available. According to the Metals Service Centers Institute, the stock of steel-processing companies increased from 2.4 to 2.7 months of inventory in November 2010 due to raising prices and the firms' anticipation of further price increases (Wall Street Journal, 2011b). Companies that actively use inventory management for price risk mitigation are, for instance, Unilever and Caterpillar (Unilever, 2016; Wall Street Journal, 2011b, p. 39).

A second example, that refers to *Chapter 5* of this thesis, is financial hedging through optimal forward contracting, i.e., the optimization of the firm's procurement position in the forward contract market that exhibits a growing liquidity during the past decades. The number of exchange-traded financial derivatives on energy, agriculture, precious metals and non-precious metals increased by 22.6% to 4.6 billion contracts between 2014 and 2015 (FIA, 2015). A large-scale empirical study by Bartram et al. (2009) shows that 50.4% of the oil-processing companies and 30.5% of the steel-processing companies have implemented some kind of commodity price risk hedging using financial contracts. Financial contracting might especially be relevant for firms that act in just-in-time environments (e.g., the automotive industry with copper or aluminum as important raw materials) or firms that purchase commodities that can hardly be stored (e.g., energy or freight capacity). A prominent example for a firm that has benefited greatly from commodity hedging is the food manufacturer General Mills, which realized hedging gains of $151 million in volatile agricultural and energy markets during the first quarter of 2008 (Wall Street Journal, 2008). On the other hand, by contractually hedging future demand, firms become inflexible to react to price declines. In 2015, the world's second-largest airline United lost $960 million, the world's third-largest airline Delta even $2.3 billion by hedging 100% of their fuel costs via long-term contracts prior to the big drop in crude oil prices (Wall Street Journal, 2016a).

A third example, that refers to *Chapter 6* of this thesis, is integrated commodity storage and physical trading[1] at storage assets that are typically owned by trading companies and serve as a link between commodity producers and commodity processors. Maximizing the merchant's profit requires the integration of trading and operational decisions since storage constraints affect the optimal trading policy. Gas storage facilities, for instance, are characterized by rate constraints that limit the amount of gas injection and withdrawal, which increases the complexity of optimal storage control.

Lately, in the course of digitalization, both academia and industry increase the focus on data-driven decision support. Companies started to routinely collect immense amounts of data and have access to real-time information from financial databases, such as Bloomberg, Quandl or Thomson Reuters. However, due to a lack of *prescriptive analytics* approaches that prescribe an optimal data-driven course of action, it is still an open question how to effectively exploit the value of data, such as economic indicators, analyst forecasts or weather information, for commodity procurement and inventory management. On the one hand, existing machine learning techniques focus on accurate predictions $Y = f(X)$ of Y based on observations $X = x$ (supervised learning), however they do not address optimal decision-making since they do not exploit the structure of the underlying optimization problem. And on the other hand, existing stochastic optimization techniques that explicitly exploit the structure of the optimization problem presume full information about the price model with a perfect out-of-sample generalization, which is not realistic[2]. In reality, procurement and inventory managers need to

[1] Throughout this thesis, we focus on physical commodity trading rather than financial trading.

[2] A *generalization error* (also referred to as the out-of-sample error) occurs, which is the prediction error over an independent test sample (Hastie et al., 2013, p. 220) and accordingly the out-of-sample cost of in-sample decisions in the operations research context (Ban and Rudin, 2019).

decide about order quantities without having complete knowledge of the price process.

Hence, it might be beneficial to integrate the fields of statistical learning and mathematical optimization (*learning-enabled optimization*) in order to identify optimal procurement and inventory plans and measuring the value of data from a decision rather than from a prediction perspective. Practical decision support is needed in the form of simple-to-use policies that are interpretable and accessible to managers in practice.

1.2. Classification, Contribution and Research Questions

Operating policies under price uncertainty, such as inventory replenishment or contracting policies for commodity-processing and commodity-trading firms are typically derived based on exogenous price models that are assumed to be fully known to the decision maker in both the price process class and the parameters (*full information problem*[3]). This typically results in the optimality of state-dependent threshold-based operating policies for inventory and procurement control.

However, the *full information optimum* is only optimal with respect to the true stochastic process that is typically not known to the decision maker. An accurate assessment of the underlying price process (and therefore the characterization of the state of the system) is difficult, especially if the commodity price time series contain structural breaks. Consequently, misspecified price models might yield unfavorable inventory replenishment and contracting decisions.

This thesis provides an integrated finance and operations perspective on commodity procurement and trading and develops practical decision support. A key contribution of this work is the focus on the operational performance implications of price model uncertainty. We rely on several fundamental operational problem settings, such as the inventory control problem under random price and demand introduced by Kalymon (1971) or the stochastic commodity warehouse problem studied by Charnes et al. (1966). The underlying problems and our proposed solution approaches are sufficiently general to allow for an application at various commodity-processing and commodity-trading firms. We test them on real data for specific commodities in the field of metals, energy and agriculturals.

[3]Following Bertsimas and Kallus (2016), we refer to the *full information problem* as optimization problems under price uncertainty with the underlying stochastic price process known in both class and parameters. Note the difference to the *perfect foresight problem* without any uncertainty.

The three main chapters of this thesis are based on three working papers. We consider two inventory control problems and one contracting problem. We furthermore distinguish between two perspectives of a risk-neutral firm: *(i)* the commodity procurement perspective under stochastic purchase prices that aims at minimizing the firm's cost and *(ii)* the commodity merchant perspective under stochastic purchase and sales prices that aims at maximizing the firm's profit. All chapters consider problems of multi-stage sequential decision-making under uncertainty that require dynamic optimization models under the consideration of the stochastic evolution of both exogenous state variables (e.g., price and demand) and endogenous state variables (e.g., inventory level or current sourcing position in the forward market). In other words, decisions taken today determine the available action space for managing price risk tomorrow and at the same time the anticipation of future inventory or contracting decisions might affect optimal decisions today. We restrict our analysis to finite-horizon problems and use dynamic programming to characterize the structure of the optimal operating policies. The absence of analytical closed-form solutions asks for numerical approaches (Williams and Wright, 1991, p. 3).

From a methodological perspective, this thesis contributes to the stochastic and data-driven optimization literature. To solve the underlying operational decision problems, methods from finance and economics (price modeling) and from the statistical learning theory (Bayesian learning, machine learning) are combined with stochastic optimization (e.g., dynamic programming and approximations, such as certainty equivalent control and decision rules). To deal with price and price model uncertainty, *Chapter 4* employs a Bayesian approach to learn a price model given price observations, while *Chapters 5* and *6* use concepts from supervised learning in order to exploit feature information.

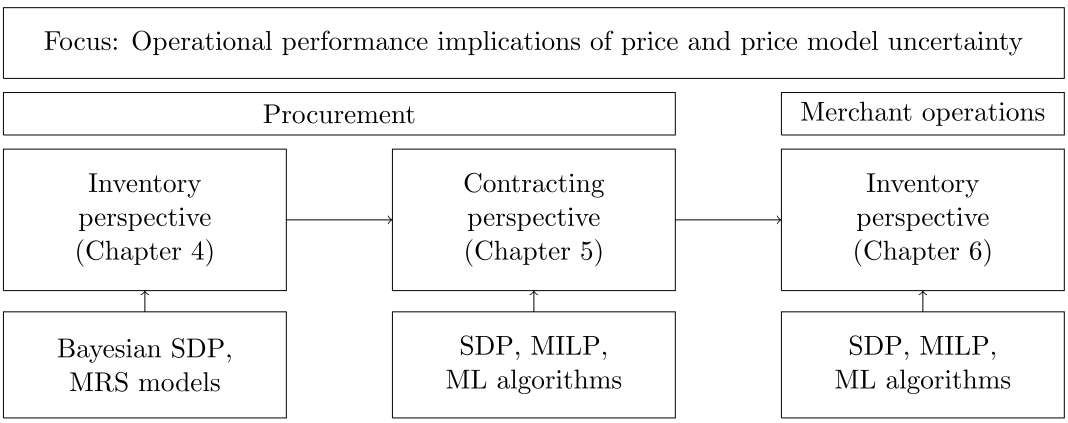

Figure 1.1.: Problem perspectives of this thesis and applied methodologies

Chapter 4 is based on Mandl and Minner (2019a) and addresses the multi-period procurement and inventory control problem under random purchase prices and demand. It aims at minimizing the sum of purchasing, holding and shortage costs. The standard literature derives procurement and inventory policies under the assumption of having full information about the stochastic price process (estimated, e.g., from historical data). However, price behavior might be non-stationary, i.e., the price process may change over time (e.g., due to regime switches), which we can infer by unsupervised learning algorithms. In order to capture uncertainty in both price and price model, we propose a Markov regime switching (MRS) approach with a Bayesian belief structure and dynamic information updates based on recent market price observations. By means of Bayesian stochastic dynamic programming (SDP), we compute optimal procurement and inventory policies under stochastic demand and purchase prices. We prove the structure of the optimal policy and find that Bayesian learning yields significant cost savings. However, computing the optimal policy parameters is difficult due to the curse of dimensionality. Therefore, we test and compare the performance of various suboptimal control policies that ignore Bayesian updates or ignore price uncertainty in general. More specifically, we answer the following research questions:

- How do partially observable spot price regimes affect the structure of the optimal inventory control policy?
- What is the cost of price regime misspecification, i.e., misspecification of the underlying price process?
- Under which conditions is it beneficial to deal with MRS price models from an inventory control perspective and when is it adequate to ignore price regimes or price uncertainty in general?
- When does a more accurate price forecast lead to better inventory decisions and is price forecast accuracy necessarily a good indicator for operational performance?

Chapter 5 is based on Mandl and Minner (2019b) and addresses the multi-period forward contracting problem under random purchase prices. Besides operational hedging via the procurement in commodity spot markets (see *Chapter 4*), the decision maker can also optimize the firm's position in the forward contract market. We present a mixed integer linear programming (MILP)-based non-parametric and data-driven approach for solving the problem. As opposed to the standard literature that assumes that the price process is fully known in its parametric form and parameters, our model relies on the machine learning (ML) principle of Empirical Risk Minimization (ERM) and trains policy parameters in the form of procurement signals directly from feature data that may have an impact on commodity prices and hence on optimal positions in the forward market. Furthermore, we combine optimization with machine learning (i.e., performance-based regularization for feature selection) in order to improve the out-of-sample performance. Based on both simulated and real data, we find that ignoring feature information can yield a significant performance loss in commodity procurement. More specifically, we answer the following research questions:

- How can firms efficiently operationalize *Big Data* for commodity procurement under price uncertainty?
- How to combine data-driven procurement with ML concepts in order to support the selection of decision-relevant (rather than prediction-relevant) features?
- What is the economic value of *Big Data* and analytics for commodity-purchasing firms?

Chapter 6 is based on Mandl et al. (2019) and addresses the multi-period inventory trading problem under random purchase and selling prices, which is referred to as the *Stochastic Commodity Warehouse Problem* (SCWP). In contrast to the previous chapters, this chapter studies commodity operations from the perspective of a merchant with purchase, storage and sales options restricted by operational constraints with regard to injection and withdrawal, which turn off the simple all-or-nothing property of the optimal storage policy. Based on six major exchange-traded commodities, we quantify the weaknesses of the state-of-the-art rolling intrinsic approach (RIA) when applied to real data. We present several learning-enabled optimization models and find that data-driven policies, if optimally structured, can significantly improve the profit of commodity storage facilities. More specifically, we answer the following research questions:

- How does the state-of-the-art RIA policy perform in backtesting settings on real data?
- How to effectively solve the SCWP in a data-driven and learning-enabled way?
- What is the out-of-sample value of data-driven storage policies for the fundamental SCWP and what is the value of exploiting known policy properties in data-driven optimization?

Chapter 2.

Fundamentals of Commodity Markets

This chapter should help those readers with a strong operations background who are not familiar with the specific characteristics of commodity markets. Readers who know the fundamentals and terminology of commodity trading can easily skip this chapter. Readers who are interested in even more details might want to take a look in the following excellent textbooks and review papers: Geman (2005) and Pirrong (2011) give a comprehensive introduction to commodity finance. Haksöz and Seshadri (2007) present literature on optimization under price risk mainly from the procurement perspective and Secomandi and Seppi (2012) focus on the commodity merchant perspective.

2.1. Commodity Price Risk

According to the Cambridge Dictionary, a commodity is *"a substance or product that can be traded, bought, or sold"*[1]. Commodities have in common that their prices fluctuate over time due to shifts in supply and demand. Based on their physical nature, the literature typically distinguishes between storable and non-storable commodities. Storable commodities include metals (precious and industrial metals) and agricultural products, while non-storable commodities are mainly those in the energy field characterized by storage limitations or storage inefficiencies (e.g., electricity). Freight capacity and weather are also often referred to as non-storable commodities, which is due to their shared characteristic of being traded on both spot and forward markets in the form of freight forward agreements and weather derivatives (Pirrong, 2011, p. 5).

Figure 2.1 shows the annualized volatilities[2] at the main commodity spot markets.

[1] `https://dictionary.cambridge.org/dictionary/english/commodity`

[2] Standardized measure based on the standard deviation of daily spot price returns. For the exact mathematical definition, we refer to Geman (2005, p. 60).

We observe that commodities are significantly more volatile than exchange rates and stocks (Geman, 2005, p. 60). Consequently, commodity operations incorporate a huge risk potential for commodity-processing and commodity-trading firms. The particularly high volatility of gas prices can be explained, besides other factors, by high storage costs, storage restrictions as well as by its strong relationship to electricity, which is the most volatile commodity of all (Geman, 2005, p. 59).

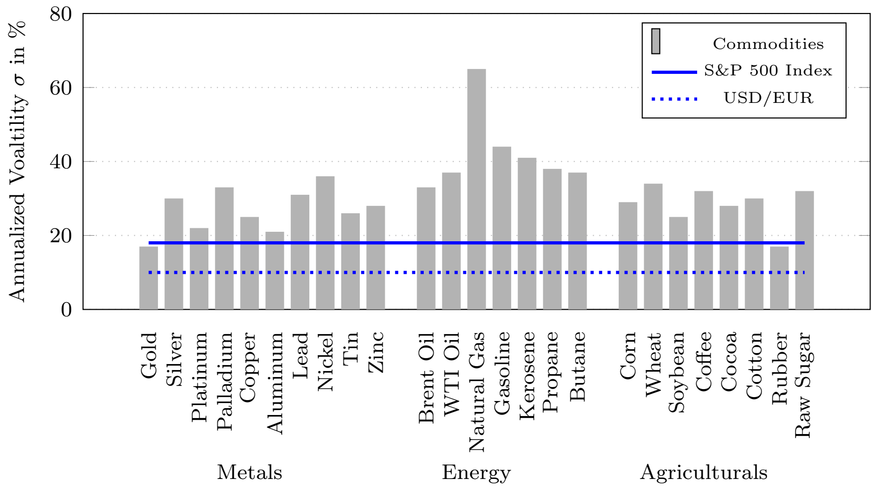

Figure 2.1.: Mean annualized price volatility σ of major commodities from 2000 to 2017 relative to corporate stocks and exchange rates (Data source: Thomson Reuters Datastream).

2.2. Financial and Operational Risk Management

In this section, we give some theoretical justification for commodity risk management of risk-neutral firms. We also point out the distinction between hedging, arbitrage and speculation and give a brief overview of financial and operational commodity risk management instruments.

Rationale for Corporate Risk Management

While risk-averse firms, per se, have an incentive to hedge (Gaur and Seshadri, 2005), risk-neutral firms are indifferent between two investments if their expected values are the same. Therefore, according to classical finance theory, a risk-neutral firm has no incentive to hedge price risk if the assumptions of Modigliani and Miller (1958) hold. According to Modigliani and Miller (1958), risk management in terms of operational or financial hedging has no effect on the firm's value as long as taxes, agency costs and bankruptcy costs are absent and information asymmetry does not exist. If at least one of those assumptions is violated, there is indeed rationale for risk management even for risk-neutral firms. For theoretical validation, we refer to Bessembinder (1991), DeMarzo and Duffie (1995), Froot et al. (1993), Smith and Stulz (1985), Stulz (1984), and Tufano (1998). Additionally, there is also empirical validation for different types of commodities (Carter et al., 2006; Haushalter, 2000; Tufano, 1996).

Furthermore, the *Efficient Market Hypothesis* (EMH) introduced by Eugene Fama in 1970 states that markets are efficient if all the available information (including past spot and forward prices) is incorporated into the current market price such that technical[3] and fundamental[4] analysis cannot lead to profits above average in the long-run. However, EMH is strongly debated in the empirical and behavioral finance literature (see, e.g., Malkiel, 2003).

Arbitrage, Speculation and Hedging

Price risk at commodity markets can either be hedged with contracts (e.g., with long-term supply contracts, risk-sharing agreements or financial derivatives, such as forwards, futures and options) or physically (e.g., with inventories). Before describing the most common operational and financial hedging instruments, we want to distinguish between the important but often misused terms *arbitrage*, *speculation* and *hedging*.

Arbitrage is the risk-free exploitation of price differentials. The common characteristics of all types of arbitrage is that the arbitrageur achieves profit with certainty and knows this profit with certainty at the time of decision. Let us assume that the arbitrageur buys a commodity at the spot market at a deterministic spot price and simultaneously sells it at the forward market at a deterministic forward price. If the marginal convenience yield is negative, the arbitrageur simultaneously buys on spot, pays for storage and sells at the forward market while earning a risk-free profit. Marginal convenience yield is defined as the cost of holding a commodity in inventory, i.e., the difference between the spot price minus the futures contract price plus inventory holding costs. To prevent arbitrage through storage, the marginal convenience yield must be greater than or equal to zero. The *no-arbitrage condition* states that, in the long-term, it is not possible to generate risk-free profit at financial markets as prices will converge. Financial markets are typically arbitrage-free or allow for arbitrage only within a very short period of time.

[3]Technical analysis refers to a chart analysis based on the historical price evolution.
[4]Fundamental analysis uses additional data (features) to explain price movements.

Chapter 4 and *Chapter 5* deal with procurement settings and do not allow for reselling and *Chapter 6* that allows for reselling does not include forward trading. Hence, cash-and-carry arbitrage situations between spot and forward markets do per se not happen.

Speculation refers to the risky exploitation of expected price differentials at different points in time, i.e., the speculator decides under uncertainty. Let us assume that the speculator buys a commodity at the forward market at a deterministic forward price in order to sell it later at the spot market at a stochastic spot price. Hence, the speculator aims at exploiting favorable price movements.

Hedging aims at limiting the risk of adverse price movements through financial contracts, such as forwards, futures, swaps or options (financial risk management) or through forward buying and stockpiling (operational risk management). We want to describe the nature of these risk management instruments in the following.

Financial Risk Management: Forward, Futures, Swaps and Options Contracts

Forward contracts are agreements made in period t to buy or sell a commodity at a pre-specified price (i.e., the forward price) at a fixed future date $\tau > t$ (often called maturity or expiration date). The forward price is paid at delivery.

Futures contracts are financial derivatives that have very similar characteristics. However, while forward contracts are over-the-counter (OTC) agreements directly between two parties, *futures contracts* are exchange-traded and highly standardized in terms of quantity, quality, delivery dates and delivery locations (Zsidisin and Hartley, 2012, p. 83). The main commodity exchanges are the New York Mercantile Exchange (NYMEX), the Chicago Mercantile Exchange (CME), the London Metal Exchange (LME), the Commodity Exchange (COMEX) as part of NYMEX for metals and the Chicago Board of Trade (CBOT) as part of CME for mainly agricultural commodities. At those exchanges, futures contracts are available for different (typically monthly) maturities, from one month (front-month contract) up to several years in the future (whereas typically only close maturities are liquidly traded). *Swaps* are generalizations of forward contracts with the agreement made over a specified period of time.

In contrast, *options contracts* give the commodity purchaser the right but not the obligation to take off a pre-specified quantity at a pre-specified price (exercise or strike price) at a future date τ (European option) or until a future date τ (American option) by paying a certain reservation price or premium. As we focus on forward and futures contracts in this thesis, we highlight these contract types in more detail in the following. For details on options and swaps, we refer the reader to, e.g., Hull (2005, 2018).

Note that, similar to Borovkova and Geman (2008, p. 10), we use the terms futures and forwards interchangeably throughout this thesis keeping in mind that forward and futures prices are identical only under non-stochastic interest rates (Geman, 2005, p. 42). This is reasonable since the interest rate risk is negligible (Geman, 2005, p. 44).

The relationship between the forward price $f_{t,\tau}$ for any future maturity $\tau \geq t$ and the spot price p_t is given by the cost-of-carry model, i.e.,

$$f_{t,\tau} = p_t e^{(r+c_h-\tilde{Y})(\tau-t)} = p_t e^{(r-Y)(\tau-t)}, \tag{2.1}$$

with $f_{\tau,\tau} \equiv p_\tau$ (convergence property). r is the risk-free interest rate and c_h are the storage costs per unit and per unit of time typically expressed as a percentage of p_t. $\tilde{Y}$ is the marginal convenience yield on the commodity, which is sometimes aggregated to $Y = \tilde{Y} - c_h$ (Geman, 2005, p. 35). The marginal convenience yield is the benefit of physically holding an additional unit of inventory, rather than a forward contract. It is an indicator for the market's expectation about the future availability of a commodity. The greater the possibility that a supply shortage might occur, the higher Y (Hull, 2018, p. 145).

Equation (2.1) is often referred to as the no-arbitrage condition as it avoids cash-and-carry arbitrage through buying at the spot market and simultaneously selling at the forward market. Equation (2.1) furthermore characterizes the forward curve $\vec{F}_t = (f_{t,\tau} : \tau > t)$ (sometimes referred to as the term structure): if $(r - Y) > 0$, then the forward curve is an increasing function of the maturity τ and the market is said to be in contango (normal market). In case of contango, speculators experience so-called negative roll yield by rolling positions across the forward curve. Purchasers have incentives to purchase commodities and put them into storage if storage is possible at a price less than the curve differential. If $(r - Y) < 0$, then the forward curve is a decreasing function of the maturity τ and the market is said to be in backwardation (inverted market). Backwardation typically happens when commodity availability is expected to be low and therefore supply is instable such that the convenience yield is large (Geman, 2005, p. 12). In case of backwardation, speculators experience positive roll yield.

The difference between spot and forward price is called the basis $B_{t,\tau} = p_t - f_{t,\tau}$ (Geman, 2005, p. 14). It can be explained by the *Theory of Storage* (Brennan, 1958; Kaldor, 1939; Telser, 1958; Working, 1949) and equals the forgone interest by a purchase in period t plus the marginal cost of storage from period t until period τ minus the marginal convenience yield (Fama and French, 1987). An important implication of the

Theory of Storage is that the price and the price volatility of commodities are both negatively correlated with the level of global inventories (Geman, 2005, p. 28).

Under the *Rational Expectations Hypothesis*, the forward price $f_{t,\tau}$ determines the best estimator of the future spot price p_τ under the risk-neutral probability measure $\mathcal{Q}$ (sometimes also referred to as the equivalent martingale measure), i.e.,

$$f_{t,\tau} = \mathbb{E}_t^{\mathcal{Q}}[p_\tau]. \tag{2.2}$$

However, statistical tests on empirical data reject the hypothesis in most cases (Geman, 2005, p. 33). If equation (2.2) is violated, then $f_{t,\tau}$ is a biased estimator of p_τ, which is typically explained by a *risk premium*, i.e., $f_{t,\tau}$ reflects both the forecast of the future spot price and the risk premium the decision maker is willing to pay to secure a fixed price in t for delivery in τ (Geman, 2005, chap. 2.4).

The difference between forward price $f_{t,\tau}$ and future spot price at maturity p_τ determines the forecast ability of futures prices to predict spot prices. However, it turns out that the forecast ability of futures prices is rather poor (Borovkova and Geman, 2008, p. 10) and might be outperformed by no-change (naïve) forecasts (Alquist and Kilian, 2010) or analyst forecasts of major financial institutions (Cortazar et al., 2018). In this thesis, we use analyst forecast data to improve inventory decisions in *Chapter 6*.

Operational Risk Management: Physical Inventories

Compared to forward contracts, spot markets provide a higher flexibility in procurement. Due to the fact that there is no equilibrium in inventory holding costs (i.e., some firms may store cheaper than the market), a firm might also operationally hedge price risk via forward buying at the spot market and carrying inventory.

In the following, we give a brief introduction to operational hedging via inventories. For an early but fundamental discussion on inventory control under stochastic raw material prices, we refer to Kingsman (1985). For a more general view of commodity storage from an economist's perspective, we refer to Williams and Wright (1991).

According to Arrow (1958), there are three motives for holding inventories: safety motives, speculation motives and transaction motives. Stochastic inventory control theory traditionally focuses on uncertain demand in order to address safety motives, i.e., keeping a safety stock to guarantee defined customer service levels. However, in volatile commodity markets, also speculation motives play a potentially important role in terms of procurement in advance and stockpiling of commodities for which one expects a price increase $p_{t+1} - p_t$ that is larger than the inventory holding costs c_h (Gavirneni and Morton, 1999). Suppose a multi-period setting with a planning horizon of n periods and deterministic period demand $d_t \;\; \forall t = 1, ..., n$ that takes place uniformly over the period and needs to be satisfied. One chooses the order quantity y_t that minimizes the expected cost C_t over the planning horizon by solving the dynamic programming equation

$$C_t(I_t, p_t) = \min_{\substack{y_t + I_t \geq d_t \\ y_t \geq 0}} \left\{ p_t y_t + c_h(y_t + I_t - \frac{1}{2} d_t) + \mathbb{E}_t[C_{t+1}(I_{t+1}, p_{t+1})] \right\} \;\; \forall t = 1, ..., n, \tag{2.3}$$

with the system state $z_t = (I_t, p_t)$ characterized by the inventory level I_t that evolves according to $I_{t+1} = I_t + y_t - d_t$ (endogenous state information) and the spot price p_t (exogenous state information). According to the seminal work by Kingsman (1969), the optimal procurement policy is of the form

$$y_t^*(z_t) = \begin{cases} (d_t - I_t)^+ & \text{if } p_t > P_2, \\ (D_k - I_t)^+ & \text{if } P_{k+1} < p_t \leq P_k \quad \forall k = 2, ..., n-1, \\ (D_n - I_t)^+ & \text{if } p_t \leq P_n. \end{cases} \tag{2.4}$$

D_k denotes the cumulative demand up to period k including the current period. The optimal policy is characterized by price thresholds P_k that are interpreted as the price at which the purchaser is indifferent between covering the cumulative demands for the next k or $k-1$ periods ahead including the present period. This is the lowest price that the purchaser expects to pay over the next $k-1$ periods for each unit of consumption in the k^{th} period ahead in time, unless it is bought now at the price offer p_t. Under the assumption of serially independent prices modeled by distribution function ϕ (Kingsman, 1969), P_k can be determined analytically by

$$P_k = \int_0^{P_{k-1}} p_{t+1} \; \phi_{t+1}(p) dp_{t+1} + \int_{P_{k-1}}^{\infty} P_{k-1} \; \phi_{t+1}(p) dp_{t+1} - c_h \;\; \forall k = 3, 4, ..., n, \tag{2.5}$$

$$P_2 = \int_0^{\infty} p_{t+1} \; \phi_{t+1}(p) dp_{t+1} - c_h. \tag{2.6}$$

However, under stochastic demand structures and more advanced and realistic price models (see Section 2.3), the inventory control problem becomes much more complicated, which is addressed in *Chapter 4* of this thesis.

2.3. Stochastic Modeling of Commodity Prices

Not only future spot prices but also future forward prices are random. While forward curves (see Figure 2.2) capture price seasonality (gas prices are expected to be higher during winter than during summer), they do not necessarily capture price uncertainty. In order to capture commodity price uncertainty in operational decision-making (e.g., inventory control), operations management typically relies on exogenous price models ϕ for dynamic modeling of spot prices p_t and the term structure $\vec{F}_t = (f_{t,\tau} : \tau > t)$. Even if we trust in the *Rational Expectations Hypothesis*, which implies that $f_{t,\tau} = \mathbb{E}_t^{\mathcal{Q}}[p_\tau]$, due to the stochastic evolution of the forward curve $\vec{F}_t$, the decision maker cannot be certain about the market's expectation in period $t+1$, i.e., about $f_{t+1,\tau} = \mathbb{E}_{t+1}[p_\tau]$ without using a price model. However, expectations in period $t+1$ affect decisions in period $t+1$, which might have an impact on the here-and-now decision in period t.

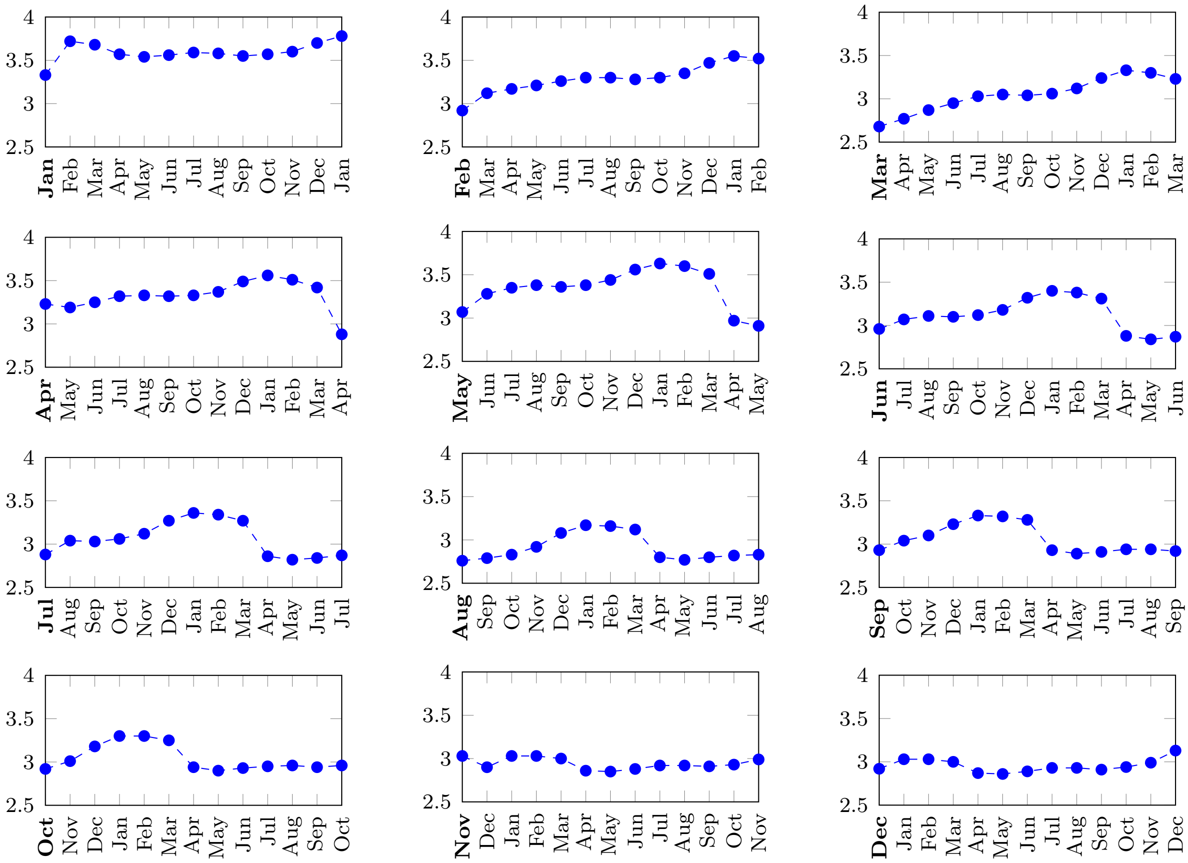

Figure 2.2.: NYMEX natural gas futures curves from 01-2017 to 12-2017 (Prices in USD/mmbtu refer to closing prices at the first trading day of the corresponding month)

In the following, we give a brief overview of prevalent commodity price models from the empirical finance and economics literature. For more details, we refer the reader to specific chapters of Geman (2005, chap. 3), Secomandi and Seppi (2012, chap. 4) and Eydeland and Wolyniec (2003, chap. 4, 5).

Basic Reduced-Form Price Models

The simplest way to capture the stochastic nature of commodity prices is a probability density function under the assumption of independent and identically distributed (i.i.d.) prices. However, if the value of the random variable (i.e., the price) evolves over time (as in case of commodity prices that are highly correlated across periods), a stochastic process, which typically is Markovian with the subsequent price depending on the current price, might be more appropriate.

The finance literature distinguishes between continuous-time and discrete-time stochastic processes for modeling commodity prices. The most basic process with Markov property is the Brownian motion that describes the basis for many others. The two most common classes of stochastic processes for modeling commodity prices are the geometric Brownian motion (GBM) and the Ornstein-Uhlenbeck process (Geman, 2005, chap. 3). GBM is a continuous-time stochastic process and can be described by the stochastic differential equation $dp_t = \mu p_t dt + \sigma p_t dW_t$, with dp_t being the price change, μ the drift of the process and $\sigma > 0$ the volatility. W_t is the Wiener process.

If drift $\mu = 0$, GBM is a martingale with the conditional expectation that the future price equals the current price. In contrast, the Ornstein-Uhlenbeck process is described by $dp_t = \kappa(\mu - p_t)dt + \sigma dW_t$ with the property of mean reversion towards the price level μ under speed $\kappa > 0$. If $p_t < \mu$, the expected price change is positive, if $p_t > \mu$, the expected price change is negative. Note that either the price or the logarithm of the price might be modeled by the mentioned stochastic processes.

For discrete-time problems, such as the multi-stage decision problems that we regard in this thesis, first-order autoregressive processes (AR(1)) are frequently used for modeling commodity price time series (see, e.g., Inderfurth et al. (2018) in the inventory control context). An AR(1) price process can be expressed by

$$p_t = \beta_0 + \beta_1 p_{t-1} + \epsilon_t, \tag{2.7}$$

with normally distributed random error term $\epsilon_t \sim N(0, \sigma_t^2)$. If $(\beta_0, \beta_1) = (0, 1)$, the AR(1) process describes a random walk (RW) without drift. A mean-reverting (MR) price process is characterized by $(\beta_0, \beta_1) = (\kappa\mu_p, 1-\kappa)$, with $\kappa \in [0,1)$ as the mean reversion speed and μ_p as the mean price level. A momentum (MO) price process with $\beta_1 > 1$ models price trends (explosive behavior) (Williams and Wright, 1991, p. 162).

Non-Linear and High-Dimensional Price Models

In reality, time series typically do not behave linearly. They are rather characterized by abrupt jumps and drops, which is due to, e.g., switches in regimes (e.g., during economic boom and bust cycles).

The state-of-the-art economic forecasting literature (see, e.g., Clements et al., 2004) presents various non-linear approaches, such as neural networks, jump diffusion models (JD), stochastic volatility models (GARCH-type models) and Markov regime switching models (MRS) that allow for time-varying parameters of the price process.

$$p_t = \begin{cases} \beta_0^{(1)} + \beta_1^{(1)} p_{t-1} + \epsilon_t^{(1)} & \text{if } s_t = 1 \\ \beta_0^{(2)} + \beta_1^{(2)} p_{t-1} + \epsilon_t^{(2)} & \text{if } s_t = 2 \end{cases} \tag{2.8}$$

In MRS models as described by equation (2.8), the price p_t is modeled by several stochastic processes with distinct parameters that depend on the current state s_t of a Markov chain. These states are typically not directly observable and must be learned from price observations (*Hidden Markov models*).

However, these methods and their implications are still not well investigated in the operational context. We investigate policy and performance implications of MRS models for an inventory control problem in *Chapter 4*.

Furthermore, *multi-factor models* that incorporate several exogenous variables (features) typically allow for a more accurate price model than one-factor models, however with a higher computational complexity (Geman, 2005, p. 69, p. 369).

The operational policy and the performance implications of multi-factor models are typically derived under the assumption of full information in terms of relevant factors and their impact. We relax full information assumptions and propose learning-enabled optimization approaches that exploit high-dimensional feature data for forward contracting in *Chapter 5* and for inventory control at commodity storage assets in *Chapter 6*.

Chapter 3.

Related Literature

This chapter presents the related literature from three relevant fields: *(i)* commodity finance with focus on explaining and modeling commodity price dynamics (Section 3.1), *(ii)* commodity operations with focus on operational decision-making (i.e., forward contracting and inventory control) in volatile commodity markets (Section 3.2) and *(iii)* methodology of stochastic and data-driven optimization under partial information and learning (Section 3.3).

3.1. Commodity Finance

Based on Cootner (1964), who describes the random character of stock prices, commodity prices are typically modeled by stochastic diffusion processes, such as geometric Brownian motions (GBM) (Samuelson, 1965), mean reverting processes (MR) (Vasicek, 1977) or jump diffusions (JD) (Merton, 1976). These processes are already well-established in stochastic inventory and procurement problems (see Section 3.2).

However, the operations literature assumes that the price process ϕ is fully known in both the class and the parameters and does not change over time. But stationarity of the price structure is unrealistic due to changing market conditions.

Markov Regime Switching Price Models

For modeling non-linearities in time series, empirical finance suggests for instance Markov regime switching (MRS) models (Hamilton, 1989, 1990). In MRS models, the time series is divided into different phases (regimes), e.g., bull and bear markets, with different underlying stochastic processes or process parameters. Regimes are latent, i.e., not directly observable, and modeled as hidden Markov chains. MRS applications in general and in the field of commodity prices in particular show that a multi-regime consideration outperforms single-regime settings in terms of accuracy for modeling spot market dynamics with a reduction of the mean absolute forecast error by 10 to 40% (for electricity see, e.g., Janczura and Weron (2012) or Yu and Sheblé (2006), for crude oil see, e.g., Vo (2009), for natural gas see, e.g., Chen and Forsyth (2010), for lumber see, e.g., Chen and Insley (2012), for copper see, e.g., Choi and Hammoudeh (2010)).

The implications of MRS price models on the optimal inventory policy structure and on the operational performance of an inventory system are studied in *Chapter 4*.

Multi-Factor and Feature-Based Price Models

Moreover, the commodity finance and econometrics literature proposes high-dimensional multi-factor models in order to explain commodity price and forward curve dynamics with a set of exogenous variables.

Gibson and Schwartz (1990), Schwartz (1997) and Schwartz and Smith (2000) suggest one-, two- and three-factor models for commodity prices that include the factors *spot price*, *convenience yield* and *interest rate*. Pindyck and Rotemberg (1990) investigate the explanatory power of variables, such as growth in industrial production, inflation and price correlation with other commodities. Geman and Nguyen (2005) empirically identify global inventory to inversely explain soybean price movements. To predict oil prices, Heath (2018) recommends measures for the economic activity and Cortazar et al. (2018) suggest a combination of futures prices and analyst forecasts. Stoll and Wiebauer (2010) propose a gas price model that considers temperature, while Brown and Yücel (2008) study natural gas prices under the explanatory variables oil price, weather, inventories and shut in production.

Chapter 5 and *Chapter 6* are motivated by the idea to explain price movements by feature data and therefore improve operating policies, such as forward contracting (*Chapter 5*) and inventory control (*Chapter 6*).

3.2. Commodity Operations

We consider three streams of literature within commodity operations. Section 3.2.1 reviews literature on optimal inventory control from a commodity purchaser's perspective and builds the foundation for *Chapter 4*. Section 3.2.2 studies optimal inventory control from a commodity merchant's perspective and builds the foundation for *Chapter 6*. Section 3.2.3 studies optimal contracting under purchase price risk and builds the foundation for *Chapter 5*. Note that there is valuable literature on commodity operations either presenting stylized two-period or deterministic models (e.g., Arnold and Minner, 2011; Arnold et al., 2011), as well as literature on commodity risk management from a supply chain perspective focusing on buyer-supplier contracting (e.g., Li and Kouvelis, 1999; Turcic et al., 2015). However, both is not the focus of this thesis and therefore not explicitly reviewed in the following.

3.2.1. Inventory Control under Stochastic Purchase Price

In their review on spot market operations, Haksöz and Seshadri (2007) recognize that price uncertainty, as opposed to demand uncertainty, has been addressed quite late.

Stochastic Price, Deterministic Demand

Starting with Fabian et al. (1959) and followed by Kingsman (1969) and Golabi (1985), inventory policies are derived with the price modeled by an i.i.d. probability density aiming at minimizing the sum of purchasing and holding costs. In this setting, the optimal inventory control policy is driven by *speculation motives* and characterized by a series of price thresholds P_k (referred to as price breaks) that, if compared to the current market price p_t, indicate how many periods k to purchase in advance (see Section 2.2).

Stochastic Price, Stochastic Demand

Under stochastic price and stochastic demand, the optimal inventory policy is driven by *speculation motives* and *safety motives* with the objective of minimizing the sum of purchasing, holding and penalty costs. Kalymon (1971) is the first to study the more complex setting with both purchase price and demand being random and modeled by a Markovian process. He proves under the assumption of non-zero setup costs, the optimality of $(s_t(p_t), S_t(p_t))$ control policies, i.e., reorder point s_t and order-up-to level S_t are both functions of the current price observation p_t. If setup costs are zero, it follows that $s_t = S_t - 1$ (base-stock policy). Yang and Xia (2009) extend Kalymon (1971) to the continuous review case. Gavirneni and Morton (1999) explicitly study the impact of speculation motives, i.e., forward buying due to an expected price increase larger than the inventory holding costs and propose effective heuristic solution procedures. Gavirneni (2004) considers changing purchasing costs due to fluctuating exchange rates modeled by a first-order Markovian approach. Inderfurth et al. (2018) study operational hedging against purchase price risk via combined inventory control in the presence of spot markets and real options contracts in the form of a capacity reservation contract. They model the purchase price as an MR process and study the cost impact of ignoring the autocorrelation between subsequent prices. They show that spot price-dependent order-up-to levels are optimal for both contract purchases and spot purchases. They furthermore show that modeling price-demand correlation has no impact on the policy structure and only a minor impact on policy parameter values and performance. Berling and Martínez-de-Albeníz (2011) and Berling and Xie (2014) consider a continuous review inventory system and model the spot price via continuous-time stochastic processes, i.e., GBM respectively an MR process. They compute optimal price-dependent base-stock policies. Berling and Martínez-de-Albeníz (2011) focus on finding efficient algorithms for the exact determination of the optimal policy parameters, whereas Berling and Xie (2014) present approximative methods, yet still under the assumption of full information about the price process ϕ, rather than on real backtesting settings.

So far, changing market environments affecting the price process have not been part of any research from an inventory perspective. According to Haksöz and Seshadri (2007), there is a lack of integrating the dynamics of information revelation of spot prices, which motivates *Chapter 4* of this thesis.

3.2.2. Inventory Control under Stochastic Purchase and Sales Price

After reviewing commodity inventory control from a procurement perspective, we now want to focus on the merchant's perspective, which is the scope of *Chapter 6*. Other than the procurement literature (Section 3.2.1), the commodity merchant and inventory trading literature typically does not incorporate an explicit demand component, but models demand indirectly by stochastic sales prices. For a more detailed review on commodity merchant operations, we additionally refer the reader to Secomandi and Seppi (2012), however with a strong focus on energy commodities.

The Stochastic Commodity Warehouse Problem (Full Flexibility Case)

The commodity trading problem dates back to the *warehouse management problem* introduced by Cahn (1948). It studies the optimal procurement, storage and sale of a single commodity under initial inventory and finite warehouse space, but under deterministic price variations and full warehouse flexibility, i.e., without injection or withdrawal capacity limits. This setting is sometimes referred to as the uncapacitated storage or fast storage setting as the warehouse can be filled and emptied within one period.

Different solution approaches, such as linear programming (Charnes and Cooper, 1955) or dynamic programming (Bellman, 1956), have been presented. Dreyfus (1957) characterizes the optimal policy under deterministic but seasonal purchase and selling price variations by recursive application of a decision rule. It consists of four regions, i.e., sell all the available inventory, buy up to the warehouse capacity, sell all the available inventory and buy up to the warehouse capacity, and do nothing.

Charnes et al. (1966) extend the problem to include uncertain purchase and sales prices. They show that the optimal policy for the *Stochastic Commodity Warehouse Problem* (SCWP) is still of a simple threshold structure (bang-bang-type) and independent of the available inventory: given the current input and output prices, it is optimal to purchase/sell (fill up/empty warehouse) and do nothing otherwise.

The Stochastic Commodity Warehouse Problem (Limited Flexibility Case)

Several variants of the SCWP have been studied in terms of high-dimensional stochastic price processes or additional operational constraints that limit the flexibility of storage operations. This case is sometimes referred to as the capacitated or slow storage case as the warehouse cannot be filled and emptied within one period due to, e.g., technical, logistical or market constraints.

Motivated by a gas storage asset, Secomandi (2010) extends Charnes et al. (1966) for commodity trading under injection and withdrawal capacity limits. The spot price is modeled by an exogenous Markov process. He establishes the optimality of a price-dependent double base-stock policy, i.e., procure-up-to and sell-down-to thresholds partition the policy into three regions (buy and inject, do nothing, withdraw and sell). If the storage asset is fully flexible and can be filled and emptied fast (within a single review period), the policy structure is the same as in Charnes et al. (1966).

However, the optimal policy parameters of the SCWP can only be derived numerically by means of stochastic dynamic programming (SDP). This yields computational intractability for real-world problem sizes in terms of, e.g., planning horizons and high-dimensional price processes (Secomandi, 2015). Hence, approximate dynamic programming (ADP) (Lai et al., 2010; Nadarajah and Secomandi, 2018; Nadarajah et al., 2015; Nascimento and Powell, 2008), approximate linear programming (ALP) (Nadarajah et al., 2015), or reoptimization heuristics, such as the rolling intrinsic policy (Lai et al., 2010; Secomandi, 2010, 2015; Wu et al., 2012), are used to efficiently solve the SCWP.

Other papers related to the SCWP deal with commodity conversion rather than storage settings (see, e.g., Devalkar et al. (2011, 2018) for soybean-to-meal and soybean-to-oil conversion or Goel and Tanrisever (2017) for corn-to-ethanol conversion). Note that even though we focus on storage settings in this thesis, the general models presented in *Chapter 6* can easily be extended to conversion settings as well.

3.2.3. Financial Contracting under Stochastic Purchase Price

Besides inventory control in the presence of volatile commodity spot markets, there is another stream of literature that optimizes procurement positions in the forwards, futures or options market, which is related to the practice-motivated setting from *Chapter 5* of this thesis.

Risk-Neutral Buyers

While risk-averse buyers, per se, have an incentive to hedge (Gaur and Seshadri, 2005), risk-neutral buyers are indifferent between two investments if their expected values are the same. Therefore, according to classical finance theory, a risk-neutral firm has no incentive to hedge price risk if the assumptions of Modigliani and Miller (1958) hold. However, due to market frictions, hedging can increase the value of a firm (Froot et al., 1993; Smith and Stulz, 1985). From Smith and Stulz (1985) follows that it is optimal to hedge fully or not at all. Froot et al. (1993) show that partial hedging can be optimal under multiple (correlated) sources of uncertainty.

Goel and Gutierrez (2011) study procurement control under the availability of both spot and forward markets and therefore incorporate contracting decisions. Secomandi and Kekre (2014) study a real options approach without storage opportunity, as it is reasonable for energy commodities. Martínez-de-Albeníz and Simchi-Levi (2005) allow for a portfolio of options contracts besides spot purchasing. Wu and Kleindorfer (2005) integrate contract and spot purchasing by using capacity options and forwards.

Risk-Averse Buyers

Even though we restrict our analysis to risk-neutral firms in this thesis, we want to briefly refer to literature on optimal risk-averse commodity procurement. Seifert et al. (2004) use a mean-variance approach for a single-period problem to determine the optimal mix of forward and spot quantities under random price and demand. Martínez-de-Albeníz and Simchi-Levi (2006) apply mean-variance analysis to support the decision between spot and options purchases. Kleindorfer and Li (2005) study a multi-period commodity contracting problem subject to value-at-risk (VaR) constraints.

3.3. Stochastic Optimization with Partial Information and Learning

Section 3.3 builds the methodological foundation for the subsequent chapters. Section 3.3.1 reviews inventory control with partial information and Bayesian learning, which refers to *Chapter 4* of this thesis. Section 3.3.2 reviews the literature on data-driven optimization and machine learning, which is particularly relevant for *Chapter 5* and *Chapter 6*.

3.3.1. Bayesian Inventory Control under Partial Information

Inventory models based only on a single stochastic process do typically not account for situations where changing information or a sharp and persistent change in the market leads to structural breaks. Therefore, inventory research proposes approaches for covering non-stationarity, however particularly addressing uncertain demand. We want to build on some of the ideas for our inventory control model in *Chapter 4*.

Markov-Modulated Demand Models with Observable States

In Markov-modulated demand inventory models, the effects of changing environments are considered, as uncertainty is not only described by one, but by various distributions characterized by the states of a Markov chain. Song and Zipkin (1993) model demand by a Poisson process with a demand rate λ_i, where i denotes the state of the world (regime). Chen and Song (2001) consider different demand states, assuming that the current state is observable. At the beginning of each period, the current state is observed, an order is placed and a shipment is received. A state-dependent base-stock policy is optimal. For the case of non-zero setup costs, Beyer and Sethi (1997) and Sethi and Cheng (1997) establish the optimality of state-dependent (s, S) policies for the full backlog case and Cheng and Sethi (1999) for the case of lost sales.

Markov-Modulated Demand Models with Unobservable States and Bayesian Learning

However, demand regimes are typically not directly observable. Hence, the random variable needs to be described by a hidden Markov model, where distributions (or its parameters) are not known with certainty and are subject to learning. Scarf (1959), Azoury (1985) and Lovejoy (1990) introduce parameter-adaptive models where the demand distribution is not entirely known and current observations are used for updating distribution parameters. Treharne and Sox (2002) model demand as a partially observable process where the distribution parameters in each period are determined by the state of a Markov chain. They use Bayesian techniques to derive probabilistic information about the current unobservable state. Treharne and Sox (2002) and more recently Wang and Mersereau (2017) show that a belief-dependent base-stock policy is optimal. For each period t, there exists an optimal base-stock $S_t(\vec{\pi}_t)$ that depends on the prior belief $\vec{\pi}_t = (\pi_t^s)$ about being in state s in period t. For the finite horizon case, Treharne and Sox (2002) compare suboptimal control policies in order to overcome computational challenges induced by the curse of dimensionality of dynamic programming.

Other papers study partially observable inventory problems where demand is censored (Bayraktar and Ludkovski, 2010) or where supply is Markov-modulated (Arifoğlu and Özekici, 2010). All contributions establish the optimality of state-dependent base-stock or (s_t, S_t) policies with the manager's prior belief $\vec{\pi}_t$ about the demand process as a state variable.

3.3.2. Data-Driven and Machine Learning-Enabled Optimization

In the emerging field of data-driven optimization, uncertainty about certain parameters is represented in a distribution-free way by historical data itself, rather than by parametric models (e.g., specific stochastic processes) (Bertsimas and Thiele, 2006). The data might be historical data about the random variable of interest (in the case of this thesis the commodity price) or feature data[1] (such as economic indicators or temperature), which is expected to have an impact on the evolution of the random variable of interest. However, data-driven optimization does not penalize model complexity per se and hence might boost overfitting issues, which favors *generalization error*. While model generalization is a central target in statistical learning theory (Vapnik, 1998), it was widely overlooked in the operations literature for a long time. Only recently, out-of-sample generalization and decision-based feature selection raise the idea of combining data-driven optimization with ML. For recent tutorials on integrating data-driven optimization and ML (referred to as *prescriptive analytics*), we want to point the reader towards Bertsimas and Kallus (2016) and Curtis and Scheinberg (2017).

As the current research output in the field of data-driven optimization is overwhelming, we mainly want to focus on papers from the fields of inventory control under demand uncertainty and portfolio optimization under return uncertainty that motivated us to

[1] often referred to as auxiliary data, covariate information or causal factors.

develop data-driven and learning-enabled models for optimizing procurement positions in commodity forward markets (*Chapter 5*) and optimizing storage policies for commodity merchants (*Chapter 6*).

Beutel and Minner (2012) and Sachs and Minner (2014) study a retail-inspired newsvendor problem with the required inventory level as a linear function of features, i.e., price, temperature and weekdays. Ban and Rudin (2019) study a similar setting with the newsvendor order quantity linearly explained by a huge amount of features (*Big Data*). They extend the data-driven newsvendor setting to regularization techniques from ML in order to select decision-relevant features and for controlling model complexity. They show that high-dimensional feature data (*Big Data*) yields significant out-of-sample cost reductions over the featureless approach (*Small Data*). Elmachtoub and Grigas (2017) study data-driven shortest path problems, assignment problems and portfolio optimization problems, using a loss function that accounts for the structure of the underlying optimization problem. Cohen et al. (2016) study data-driven dynamic pricing problems, using covariates on product characteristics and demand. Ban et al. (2018b) solve the data-driven portfolio optimization problem by means of performance-based regularization. In doing so, they improve out-of-sample generalization by obtaining more stable and less complex feature models (*Occam's razor*). Ban et al. (2018a) study a practice-motivated dynamic procurement problem with stochastic demand of new products modeled in a data-driven way by using covariates, such as the demand of similar products and product characteristics.

Chapter 4.

Operational Hedging from a Bayesian Inventory Control Perspective

Based on
Mandl, C. and S. Minner (2019a). When do commodity spot price regimes matter for inventory managers?

A large number of firms buy commodities at spot markets that are characterized by volatile prices. Due to different market regimes (e.g., bull and bear), spot price dynamics are non-stationary and only partially observable; neither the underlying stochastic price process nor its parameters are known with certainty. To capture uncertainty in both price and price model, we exploit recent spot price observations to dynamically update (*learning*) probabilistic price regime information in the context of inventory control for a storable commodity under random demand and purchase price. By means of Bayesian dynamic programming, we prove that, if prices evolve according to doubly embedded stochastic processes described by a hidden Markov regime switching (MRS) model, price(s)- and regime-belief-dependent base-stock policies, rather than price-dependent policies, are optimal. We distinguish between independent and Markovian price processes and demonstrate the difference concerning optimal base-stock functions and monotonicity properties. We find that ignoring regime shifts leads to suboptimal inventory decisions and we quantify the cost of misspecifying the spot price model. We find that Bayesian learning can yield significant cost savings that are particularly high when demand volatility and inventory holding costs are low and regime persistence is high. With regard to the curse of dimensionality at computing the optimal state-dependent base-stock levels, we propose and test different simpler heuristics (e.g., certainty equivalent control or naïve control policies) and evaluate their effectiveness.

4.1. Introduction

A significant number of firms buy raw materials at spot markets, rather than signing long-term fixed-price supply contracts (Wall Street Journal, 2002). There are several reasons: flexibility, speculation opportunities or independence from single suppliers and therefore a decrease in supply disruption risk. However, commodity spot prices possess annualized volatilities of up to 40%, with a tendency to grow (Geman, 2005, p. 2). For instance, between 2000 and 2014, price volatilities of aluminum and wheat increased on average by 8% and 6% p.a., respectively, which translates into a soaring risk for commodity-processing companies and a challenge for inventory managers.

To derive optimal inventory decisions, it is intuitively necessary to accurately model the stochastic nature of commodity prices. However, history shows that price dynamics of spot-traded commodities are not stationary over time. Changing market supply and demand due to economic boom and bust cycles, legal and political risk, innovations, weather events or commodity buffering strategies of national economies can lead to an abrupt and sharp change in price. Especially in metal and agricultural markets, purchasers may have to cope with price shifts in the future since almost half of the new mining projects and 80% of the available arable land are located in regions with a high political risk (McKinsey & Company, 2013). A raise in frequency of extreme weather events like hurricanes and floods reinforces these expectations. Furthermore, strongly increasing volatilities in commodity prices are expected in the next couple of years (Bloomberg, 2015), intensified by the observation that prices at commodity exchanges do represent market actors' anticipations of the future, rather than the economic equilibrium between today's supply and demand.

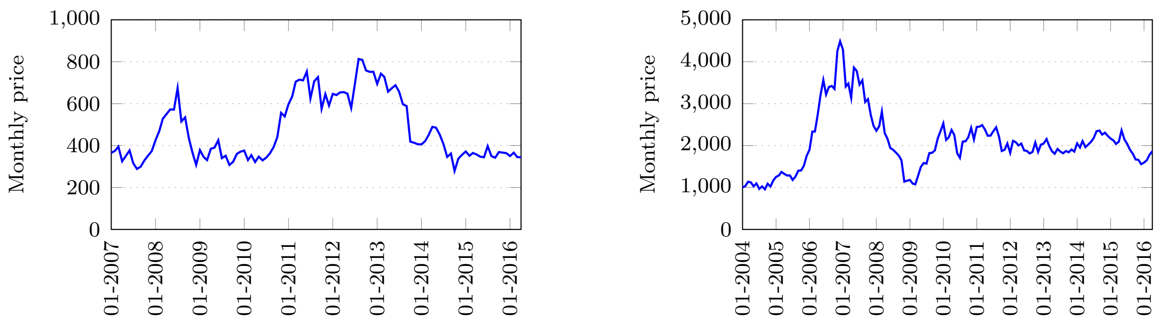

Figure 4.1.: Spot price of CBOT corn [Cts per bushel] and LME zinc [USD per ton]

Figure 4.1 shows the prices of corn at the CBOT and zinc as traded at the LME. Several structural breaks in the time series are obvious in terms of price level and/or volatility: in 2007/2008, the world food price crisis, along with an announcement by the U.S. government to push ethanol as a renewable fuel, increased corn prices dramatically; in early 2011, a mixture of heat, drought and flooding in the U.S. corn belt, which accounts for more than one third of the worldwide production, led to a price jump of almost 100%; in summer 2013, rain and favorable temperatures in the Midwest of the U.S. led to a persistent price drop. In contrast, prices for zinc were affected by metal deficits in 2006 and by the financial crisis in 2007/2008. An AR(1) coefficient of $\beta_1 = 0.96$ supports a RW regime of zinc from Jan-04 to Jul-09, whereas $\beta_1 = 0.69$ rather indicates an underlying MR price process from Aug-09 to Mar-16. Unit root tests confirm that there is evidence of different states of the price process that make predictions harder and less reliable, such that *The Wall Street Journal* even proclaimed "The End of Economic Forecasting" (Wall Street Journal, 2016b).

Structural changes in the underlying stochastic price process can be considered by Markov regime switching (MRS) models with each regime itself characterized by a specific stochastic process (Hamilton, 1989, 1990), which allows for changing economic environments. Even though these non-linear doubly stochastic time series models are well-established in empirical finance and have been shown to provide significant forecast improvements (see Section 3.1), the inventory control literature (see Section 3.2.1) still assumes full information about the purchase price process and its parameters (e.g., in the form of a RW or a MR process) and hence ignores regimes and the ability of learning based on new information that becomes available. We fill this gap and investigate the implications of the price process (and its misspecification) on the performance of an inventory system in more detail. The additional evaluation of parametric stochastic price models on empirical spot price data provides insights into the context of *misleading speculation*, resulting in an undesired excess of stocks (*misspeculative inventory*). More specifically, the following questions remain unanswered in the growing literature at the interface of finance and operations: *(Q1)* How do partially observable spot price regimes and information revelation via Bayesian learning affect the structure of the optimal inventory policy? *(Q2)* What is the cost of price regime misspecification, i.e., misspecification of the price process? *(Q3)* If there is evidence for different regimes in commodity prices, under which conditions is it particularly beneficial to use Bayesian models and when is it adequate to use simpler control policies that ignore learning (e.g., certainty equivalent control), ignore regime switches (single regime control) or ignore price uncertainty in general (naïve control)? *(Q4)* Do more accurate price forecasts per se lead to better inventory decisions and is price forecast accuracy (e.g., MAPE) necessarily a good indicator for operational performance?

We study a single-item, multi-period, discrete-time, periodic-review inventory problem under random demand and purchase price and use an MRS model and a Bayesian updating scheme to deal with incomplete information about the underlying stochastic price process.

Our *methodological contribution* is the generalization and extension of the seminal work by Kingsman (1969) (for i.i.d. price processes) and Kalymon (1971) (for Markovian price processes) to partially observable price regimes with dynamic market information updates (*learning*) via Bayesian dynamic programming. In addition to recent spot price information, the manager may have an initial belief (which may be right or wrong) about the price dynamics. We prove that a state-dependent base-stock policy is optimal with the base-stock level S_t, in addition to the current spot price p_t, depending on the manager's prior regime belief $\vec{\pi}_t$ for i.i.d. processes respectively on his posterior belief $\vec{\pi}_{t+1}$ for Markovian processes. We further analyze the violation of monotonicity properties in this context, i.e., under which regime and belief conditions to order more at higher prices and less at lower prices.

As a *managerial contribution*, we find that the evaluation of price models based only on forecast accuracy measures is insufficient and that operational implications need to be considered. On the basis of both a controlled numerical study and empirical spot market data, we sensitize managers to how structural changes in commodity prices affect the optimal inventory policy and that traditional policies that ignore price information updates (no *learning*) can lead to an increase in expected cost by up to 13%. We show under which circumstances Bayesian MRS control through $S_t(p_t, \vec{\pi}_t)$ is beneficial and when price model misspecification plays a minor role from an inventory perspective. For the latter case, we examine suboptimal control policies that are more practical with regard to the curse of dimensionality, such as certainty equivalent control without Bayesian learning ($S_t(p_t, \vec{\pi}_0)$), single regime control ($S_t(p_t)$) and naïve control (S_t).

In general, there are two reasons for divergent order decisions under different price processes: *(i) safety motives*, i.e., increased orders as protection against shortage penalty cost and *(ii) speculation motives*, i.e., hedging price risk through forward buying, expecting a price increase that is larger than the inventory holding cost. Dealing with market frictions, e.g., imperfect capital markets that imply that there is no homogeneous interest rate and hence no equilibrium in inventory holding cost (disabled separation property of financial and operational decisions according to Modigliani and Miller (1958)), hedging via financial contracts or, as in this chapter, physical inventories can increase (but also decrease) a firm's value (Froot et al., 1993; Smith and Stulz, 1985). There is also empirical evidence that firms use stockpiling for price risk mitigation (Wall Street Journal, 2011a). E.g., Unilever and Caterpillar actively manage price risk through forward buying (Unilever 2016, p. 39, Wall Street Journal 2011b). Accordingly, we explicitly study both motives and show that *speculation motives* dominate *safety motives* in their impact on the cost of price process misspecification and the cost of ignoring price regimes. The cost of price process misspecification furthermore strongly depends on the degree of demand uncertainty (rather low than high) and the type of price regimes (rather level than volatility). However, on real data, *misleading speculation* induced by stochastic models can, under certain conditions, even favor a naïve policy, which is also shown in recent literature from finance (e.g., Wang et al., 2015).

Section 4.2 presents the MRS price model, the Bayesian updating scheme and the inventory control model. Section 4.3 characterizes and illustrates the structure of the optimal inventory policy and monotonicity of its parameters. In a controlled numerical study, Section 4.4 quantifies both the cost of price regime misspecification and the value of the Bayesian MRS approach relative to several suboptimal control policies. Section 4.5 tests the policies on empirical spot market data. Section 4.6 highlights managerial insights and concludes. Appendix A includes proofs of the structural results from Section 4.3 and extensive tables of the numerical results from Section 4.4.

4.2. Model Formulation

The MRS price model including the Bayesian updating scheme is described in Section 4.2.1 and the inventory control problem is characterized in Section 4.2.2.

4.2.1. MRS Spot Price Model and Bayesian Updating Scheme

Hidden Markov Regime Switching Model

The discrete time MRS model consists of a finite number of m unobservable states, i.e., price regimes $s = (1, ..., m)$, defined by distinct underlying time-homogeneous stochastic processes $\phi^s = (p_t)_{t=1,...,n}$ to model the evolution of the commodity spot price p_t. The regimes are not directly observable, i.e., at a certain point in time t, it is not known with certainty whether the price follows, for instance, a RW or an MR process. However, the current spot price p_t is observable. Hence, by using Bayesian statistics, the probability $\mathbb{P}(s_t|p_t)$ of being in a certain regime s_t can be deduced from price observation p_t and the prior regime belief $\pi_t = (\pi_t^s)_{s \in M}$ (for the discrete case, belief vector $\vec{\pi}_t = (\pi_t^1, ..., \pi_t^m)$) that defines a probability distribution over the regime space $M = \{1, ..., m\}$. Hence, the MRS model describes a doubly embedded stochastic process where the regimes switch in a Markovian way and each regime $s \in M$ emits a stochastic price process ϕ^s.

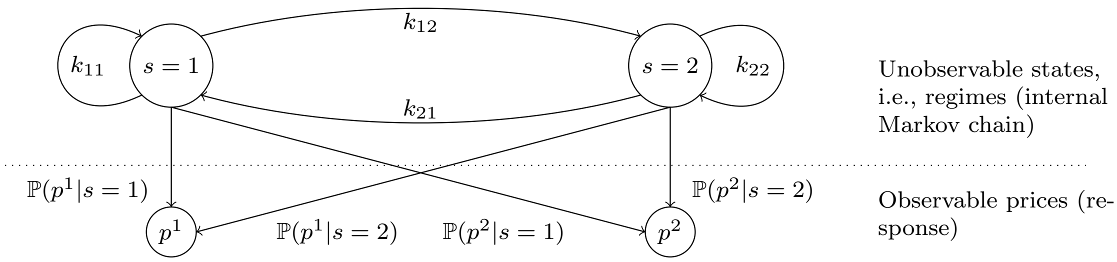

Figure 4.2.: Concept of MRS in the spot market context

Figure 4.2 gives a simplified example with regime space $M = \{1, 2\}$, two potential prices p^1 and p^2 and transition probabilities $k_{ij} = \mathbb{P}(s_{t+1} = j|s_t = i)$, i.e., the probabilities of switching from regime i at time t to regime j at time $t + 1$ for $(i, j) \in M$. $\mathbb{P}(p|s)$ is the probability of spot price p under regime s (state emission probability) and is specified by the underlying stochastic regime process ϕ^s. Since it is common in the MRS literature (Hamilton, 1989, 1990), we assume k_{ij} to be time-invariant.

Regime Estimation

The MRS parameter set $(\phi^s, \pi_t^s, k_{ij})$ is either estimated based on experts' knowledge or past time series data $(p_1, ..., p_t)$ using recursive filtering techniques such as the Baum-Welch algorithm (Baum et al., 1970), a specific instance of the expectation-maximization (EM) algorithm. This unsupervised learning method is based on maximum likelihood estimation extended to the case of incomplete data. The main goal is to achieve a good fit between regimes s and observations $(p_1, ..., p_t)$ in order to predict the actual hidden state price process from the known sequence of observed parameters. Applying the Baum-Welch algorithm, transition probabilities k_{ij} can be estimated and the initial (first prior) belief $\vec{\pi}_0 = (\pi_0^j)$ about being in price regime j can be derived by solving a set of linear equations $\pi_0^j = \sum_{i=1}^m k_{ij} \, \pi_0^i \quad \forall j \in M$ with $\sum_{i=1}^m \pi_0^i = 1$.

Bayesian Regime Belief Updates

After the current spot price $p_t = p^1$ has been observed, the prior (prior to observation p_t) belief π_t^j about being in spot price regime j is dynamically updated according to the

Bayes' theorem, which defines the learning process.

Definition 1 (Learning Process)**.** *The posterior belief* π_{t+1}^j *about being in regime* j *in the next period* $t+1$ *given* $p_t = p^1$ *and* $p_{t-1} = p^2$*, is defined for all* $j \in M$ *by*

$$\pi_{t+1}^j = \mathbb{P}(s_{t+1} = j | (p_t = p^1, p_{t-1} = p^2)) = \frac{\sum\limits_{i=1}^{m} \pi_t^i \, k_{ij} \, \mathbb{P}((p_t = p^1 \cap p_{t-1} = p^2) | s_t = i)}{\sum\limits_{i=1}^{m} \pi_t^i \, \mathbb{P}((p_t = p^1 \cap p_{t-1} = p^2) | s_t = i)}. \tag{4.1}$$

For first-order Markovian price processes $\phi^s(p_{t+1}|p_t)$ inside the regimes s, e.g., of the structure $p_t = \beta_0 + \beta_1 p_{t-1} + \epsilon_t$ (AR(1)), information about p_{t-1} is inevitably needed for updating π_t^s. For i.i.d. price processes $\phi^s(p_{t+1})$, the information about the current price observation p_t is sufficient for updating, i.e., $\pi_{t+1}^j = \mathbb{P}(s_{t+1} = j | (p_t = p^1, p_{t-1} = p^2)) = \mathbb{P}(s_{t+1} = j | p_t = p^1)$. If a certain price $p_t = p^1$ can be exclusively assigned to a regime $s_t = i'$ and a price $p_t = p^1$ is observed, then, since $\mathbb{P}(p_t = p^1 | s_t = i') = 1$, equation (4.1) reduces to $\pi_{t+1}^j = k_{i'j}$ for any $j \in M$.

4.2.2. Inventory Control Model

We consider the single-item, single-echelon, discrete-time, finite-horizon, periodic-review inventory problem as described by Kalymon (1971) with the spot market as single procurement option and without reselling (no trading). Future prices p_{t+1} follow a stochastic process $\phi(p_{t+1})$. $\phi(p_{t+1})$ is described by a convex combination of the regime processes $\phi^s(p_{t+1})$ weighted by the posterior regime belief π_{t+1}^s about being in a certain price regime s in $t+1$, i.e., $\phi(p_{t+1}) = \sum_{s=1}^m \pi_{t+1}^s \, \phi^s(p_{t+1})$. The regime processes ϕ^s can either be i.i.d. (hereafter denoted as $\phi^s(p_{t+1})$) or Markovian (hereafter denoted as $\phi^s(p_{t+1}|p_t)$). We assume the following sequential structure of the decision problem in every period $t = 1, ..., n$ (Figure 4.3).

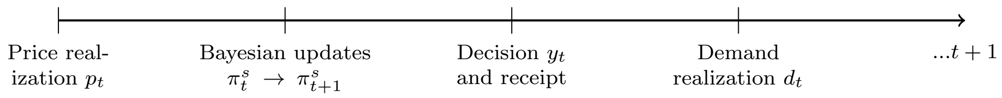

Figure 4.3.: Intra-period sequence of events

After observing price p_t, the prior belief π_t^s about being in price regime $s \in M = \{1, ..., m\}$ in period $t+1$ is updated according to equation (4.1). Then, commodity amount y_t is ordered to satisfy random demand d_t (*safety motives*) and to exploit potential forward buying benefits (*speculation motives*). Demand for the finished product is assumed to be i.i.d. stochastic with cumulative distribution function F. Without loss of generality, we assume lost sales (see, e.g., Gavirneni, 2004; Turcic et al., 2015). This is reasonable supposing the customer can substitute the finished product or buy the product elsewhere. In line with commodity operations literature (see, e.g., Berling and Martínez-de-Albeníz, 2011; Goel and Gutierrez, 2011), we assume independence of spot prices and the finished product's demand. This is realistic if the firm cannot pass higher or lower input prices on to the customer. The inventory manager is a price taker and risk-neutral. Unit inventory holding cost c_h incurs per time unit and shortages are penalized with a unit shortage cost c_p. There are no setup costs and the purchase costs equal the spot price p_t. Lead time is zero, which is reasonable for spot markets (Goel and Gutierrez, 2011). Without loss of generality, the discount factor for expected future costs is set to $\alpha = 1$. I_t is the inventory before and I_t^* is the inventory after ordering but before demand realization d_t. The order quantity is $y_t = (I_t^* - I_t)^+$ and inventory balance follows $I_{t+1} = (I_t^* - d_t)^+$. Being in state $z_t \in Z_t$, y_t minimizes the total expected cost C_t over an n-period horizon that consists of ordering costs $p_t \cdot y_t$, expected holding and shortage costs $L(I_t^*) = \mathbb{E}_{d_t}[c_h(I_t^* - d_t)^+ + c_p(d_t - I_t^*)^+]$ and cost-to-go C_{t+1}[1]:

$$C_t(z_t) = \min_{I_t^* \geq I_t} \left\{ p_t y_t + L(I_t^*) + \int_0^\infty \int_0^\infty C_{t+1}(z_{t+1}) d\phi(p_{t+1}) dF(d_{t+1}) \right\} \quad \forall t = 1, ..., n. \quad (4.2)$$

Since π_{t+1}^s is a function of $(p_t, \vec{\pi}_t)$ for i.i.d. price processes and a function of $(p_t, p_{t-1}, \vec{\pi}_t)$ for first-order Markovian (AR(1)) price processes and in order not to violate the Markov property, the state z_t of the inventory problem is described by a multi-dimensional state space $z_t = (I_t, p_t, \vec{\pi}_t)$ for i.i.d. processes $\phi^s(p_{t+1})$ and $z_t = (I_t, p_t, \underbrace{p_{t-1}, \vec{\pi}_t}_{\vec{\pi}_{t+1}})$ for AR(1) processes $\phi^s(p_{t+1}|p_t)$.

[1] One might model the salvage value on final inventories by $C_{n+1} \equiv -\mathbb{E}[p_{n+1}] \cdot I_{n+1}$ or $C_{n+1} \equiv -(\mathbb{E}[p_n] - c_h) \cdot I_{n+1}$. Without loss of generality, we set $C_{n+1} = 0$ in our numerical experiments, i.e., inventory at the end of the planning horizon is lost, which ensures that one does not order in advance for periods beyond the planning horizon. We argue in Section 4.5 why this assumption does not affect the general results.

4.3. Optimal Policy Structure and Monotonicity Properties

We characterize the optimal inventory control policy in Section 4.3.1 and present monotonicity properties of the policy parameters in Section 4.3.2, i.e., when do increasing spot prices result in increasing orders. We numerically illustrate both policy structure and monotonicity properties in Section 4.3.3.

4.3.1. Optimality of Price(s)- and Belief-Dependent Base-Stock Policies

From Kalymon (1971) follows that no matter whether prices evolve according to an i.i.d. process $\phi(p_{t+1})$ or a Markovian process $\phi(p_{t+1}|p_t)$, a state-dependent base-stock policy $S_t(p_t)$ is optimal for zero setup costs. This is also true if the price process is non-stationary but fully known, which is equivalent to multi-regime settings with observable regimes (full information case). We instead prove that if prices follow a doubly embedded stochastic process with unobservable regimes and learning (partial information case), the optimal base-stock is belief-dependent.

Theorem 1 (Optimal Policy Structure)**.**

(i) Under partially observable price regimes with i.i.d. regime processes $\phi^s(p_{t+1})$, there exists an optimal order quantity y_t^*, such that for all $1 \leq t \leq n$,

$$y_t^*(z_t) = y_t^*(I_t, p_t, \vec{\pi}_t) = \left[S_t(p_t, \vec{\pi}_t) - I_t\right]^+, \tag{4.3}$$

*i.e., base-stock level S_t is fully characterized by price p_t and **prior** regime belief $\vec{\pi}_t$.*

(ii) Under partially observable price regimes with Markovian regime processes $\phi^s(p_{t+1}|p_t)$, there exists an optimal order quantity y_t^* such that for all $1 \leq t \leq n$,

$$y_t^*(z_t) = y_t^*(I_t, p_t, p_{t-1}, \vec{\pi}_t) = \left[S_t(p_t, p_{t-1}, \vec{\pi}_t) - I_t\right]^+ = \left[S_t(p_t, \vec{\pi}_{t+1}) - I_t\right]^+, \tag{4.4}$$

*i.e., base-stock level S_t is fully characterized by price p_t and **posterior** regime belief $\vec{\pi}_{t+1}$ or equivalently by price p_t, previous price p_{t-1} and prior regime belief $\vec{\pi}_t$.*

Proof. see A.1 □

Consequently, the optimal inventory policy under partial price process information is a generalization of the policy of the full information discussed by Kalymon (1971). The practical reasoning for the dependence on p_{t-1} in the Markovian case is that in order to statistically infer the conditional probability (posterior belief $\vec{\pi}_{t+1}$) about Markovian (AR(1)) regimes, the relationship between subsequent prices ($p_{t-1} \rightarrow p_t$) is needed.

4.3.2. Monotonicity of the Optimal Base-Stock Functions

We want to investigate under which regime and belief conditions increasing prices result in decreasing base-stocks and when the reverse, more counter-intuitive, is the case where monotonicity is violated. We translate and extend the findings for the full information case by Gavirneni (2004) and Yang and Xia (2009) to the case of partial price process information with learning. The base-stock level S_t is non-increasing in the current spot price p_t, i.e., $S_t(p_t') \leq S_t(p_t)$ for $p_t' > p_t$, if the following sufficient, but not necessary, conditions *(a1)* and/or *(a2.1)+(a2.2)* are satisfied:

No-Autocorrelation Condition (a1). The no-autocorrelation condition states that future price p_{t+1} (as well as demand d_t and expected holding and shortage cost function $L(I^*)$) is independent of the current price p_t, i.e., $\phi(p_{t+1}|p_t) = \phi(p_{t+1})$. Intuitively, even though a huge price jump is expected, as long as the magnitude of the price jump is independent of the current price observation, a procurement manager would not order more at higher prices.

Mean Reversion Condition (a2.1) <u>and</u> Time Continuity Condition (a2.2). The mean reversion condition *(a2.1)* states that the expected one-period price increase at a higher price p_t' is at most of the magnitude of the expected one-period price increase at a lower price p_t, i.e., $\bar{p'}_{t+1} - p_t' \leq \bar{p}_{t+1} - p_t$, where $\bar{p}_{t+1}$ denotes the expected price subsequent to p_t and $\bar{p'}_{t+1}$ denotes the expected price subsequent to p_t'. Transferred to the case of multiple unobservable price regimes with Bayesian learning, *(a2.1)* is fulfilled if

$$\sum_{s=1}^{m} \bar{p}_{t+1}^{s}\, \pi_{t+1}'^{s} - p_t' \leq \sum_{s=1}^{m} \bar{p}_{t+1}^{s}\, \pi_{t+1}^{s} - p_t \qquad \forall\, p_t, \tag{4.5}$$

where $\bar{p}_{t+1}^{s}$ is the expected price resulting from the corresponding regime process $\phi^s(p_{t+1})$ and π_{t+1}^{s} (respectively $\pi_{t+1}'^{s}$) is the posterior regime belief after price observation p_t (respectively p_t'). This condition is violated by momentum (MO) price processes, which are the counterparts of MR price processes and indicate that small prices lead to small (or even smaller) prices and high prices lead to high (or even higher) prices.

The time continuity condition *(a2.2)* states that the price is not expected to change drastically in a short period of time. We define a set of prices $p^1 < p^2 < p^3 < ... < p^L$ and e_{ij} as the price transition probabilities from p_t^i to p_{t+1}^j. Then the probability that the price of the next period is at most p_{t+1}^j, based on the current price p_t^i, is $\sum_{l=1}^{j} e_{il}$. To satisfy condition *(a2.2)*, for all values of j, it must be ensured that

$$\sum_{l=1}^{j} e_{il} \geq \sum_{l=1}^{j} e_{i+1,l} \quad \forall\, 1 < i < L-1. \tag{4.6}$$

In the regime framework, $e_{ij} \coloneqq \sum_{s=1}^{m} \pi_{t+1}^s \, \mathbb{P}((p_{t+1}^j \cap p_t^i)|s)$. Note that if *(a1)* is fulfilled and there is no Bayesian learning based on price observations p_t, then conditions *(a2.1)* and *(a2.2)* are also fulfilled by definition as $\pi_{t+1}'^s = \pi_{t+1}^s$.

Accordingly, for multi-regime settings, we distinguish between four monotonicity cases.

Proposition 1 (Monotonicity Properties)**.**

(i) *If regime processes $\phi^s(p_{t+1})$ are i.i.d. and regime beliefs $\vec{\pi}_t$ are not dynamically updated (learned) based on price information p_t, S_t is non-increasing in p_t.*

(ii) *If regime processes $\phi^s(p_{t+1})$ are i.i.d. and regime beliefs $\vec{\pi}_t$ are updated (learned) based on price information p_t, S_t is non-increasing in p_t under sufficient, but not necessary, conditions (a2.1) and (a2.2). If regimes differ in variance but not in mean (i.e., volatility regimes), (a2.1) is not violated and S_t is non-increasing in p_t if (a2.2) is satisfied. Furthermore, more frequent regime changes, i.e., a decreasing r, where $k_{ij} = \left(\begin{smallmatrix} 0.5+r & 0.5-r \\ 0.5-r & 0.5+r \end{smallmatrix}\right)$, support the monotonicity behavior of S_t in p_t.*

(iii) *If regime processes $\phi^s(p_{t+1})$ are Markovian and regime beliefs $\vec{\pi}_t$ are not updated (learned) based on price information p_{t-1} and p_t, S_t is non-increasing in p_t under sufficient, but not necessary, conditions (a2.1) and (a2.2).*

(iv) *If regime processes $\phi^s(p_{t+1})$ are Markovian and regime beliefs $\vec{\pi}_t$ are updated (learned) based on price information p_{t-1} and p_t, S_t is non-increasing in p_t if $\phi(p_{t+1})$ fulfills the sufficient, but not necessary, conditions (a2.1) and (a2.2). If all regime processes $\phi^s(p_{t+1})$ satisfy the monotonicity conditions, monotonicity per se is given for the MRS case.*

Proof. see A.2 □

Table 4.1.: Summary of optimality and monotonicity results

Regime Observability		**Stochastic Regime Process** **i.i.d.**	**Markovian**
Observable (Full Information)	Optimality	$S_t(p_t)$	$S_t(p_t)$
	Monotonicity	S_t non-increasing in p_t	S_t non-increasing in p_t if the observed process $\phi^s(p_{t+1})$ fulfills *(a2.1)* and *(a2.2)*
Unobservable w/o Bayesian Updates (w/o Learning)	Optimality	$S_t(p_t, \vec{\pi}_0)$	$S_t(p_t, \vec{\pi}_0)$
	Monotonicity	S_t non-increasing in p_t	S_t is non-increasing in p_t if $\phi(p_{t+1})$ fulfills *(a2.1)* and *(a2.2)*
Unobservable w/ Bayesian Updates (w/ Learning)	Optimality	$S_t(p_t, \vec{\pi}_t)$	$S_t(p_t, \vec{\pi}_{t+1}) \triangleq S_t(p_t, p_{t-1}, \vec{\pi}_t)$
	Monotonicity	S_t non-increasing in p_t if $\phi(p_{t+1})$ fulfills *(a2.1)* and *(a2.2)*	S_t non-increasing in p_t if $\phi(p_{t+1})$ fulfills *(a2.1)* and *(a2.2)*

Table 4.1 summarizes that a base-stock policy is still optimal, but the base-stock functions S_t are multi-dimensional for the MRS case. Furthermore, for i.i.d. price processes, S_t is no longer necessarily non-increasing in the current spot price p_t.

4.3.3. Illustration of Policy Structure and Monotonicity Properties

Numerical Setup

We suppose discretized gamma distributed prices on a price set $p \in \{10, 15, 20, 25, 30\}$ for two types of i.i.d. regime settings ($m = 2$): *(i)* high-level-low-level (HL-LL) regimes ($\mu_{\text{HL}} = 25$, $\mu_{\text{LL}} = 15$, $\sigma_{\text{HL}} = \sigma_{\text{LL}} = 3$) and *(ii)* low-volatility-high-volatility (LV-HV) regimes ($\mu_{\text{LV}} = \mu_{\text{HV}} = 20$, $\sigma_{\text{LV}} = 2$, $\sigma_{\text{HV}} = 10$). Furthermore, we consider three common types of Markovian price processes $\phi(p_{t+1}|p_t)$ with empirical evidence for modeling commodity prices: random walk (RW), mean reversion (MR) and momentum (MO). Following Gavirneni (2004), the price transition matrices are given by

$$\text{RW}: \begin{pmatrix} .5 & .5 & 0 & 0 & 0 \\ .5 & 0 & .5 & 0 & 0 \\ 0 & .5 & 0 & .5 & 0 \\ 0 & 0 & .5 & 0 & .5 \\ 0 & 0 & 0 & .5 & .5 \end{pmatrix}, \text{MR}: \begin{pmatrix} .2 & .8 & 0 & 0 & 0 \\ .1 & .3 & .6 & 0 & 0 \\ 0 & .1 & .8 & .1 & 0 \\ 0 & 0 & .6 & .3 & .1 \\ 0 & 0 & 0 & .8 & .2 \end{pmatrix}, \text{MO}: \begin{pmatrix} .8 & .2 & 0 & 0 & 0 \\ .6 & .3 & .1 & 0 & 0 \\ 0 & .5 & 0 & .5 & 0 \\ 0 & 0 & .1 & .3 & .6 \\ 0 & 0 & 0 & .2 & .8 \end{pmatrix}.$$

Shortage cost is $c_p = 40$ and holding cost is $c_h = 1$. For illustration purposes, a planning horizon of $n = 2$ ($t = 1$ with deterministic price and random demand plus $t = 2$ with random price and random demand) is chosen and $C_{n+1} = 0$. Demand follows a discretized and rescaled normal distribution with $\mu = 5$ and $\sigma = 1$. Regime transition probabilities are $k_{ij} = \left(\begin{smallmatrix} .99 & .01 \\ .01 & .99 \end{smallmatrix}\right)$, which is reasonable for real market data (see Section 4.5).

Policy Illustration

If prices follow i.i.d. stochastic processes within the regimes, then the base-stock level S_t is a function of the spot price p_t and the prior regime belief $\vec{\pi}_t$ (Figure 4.4).

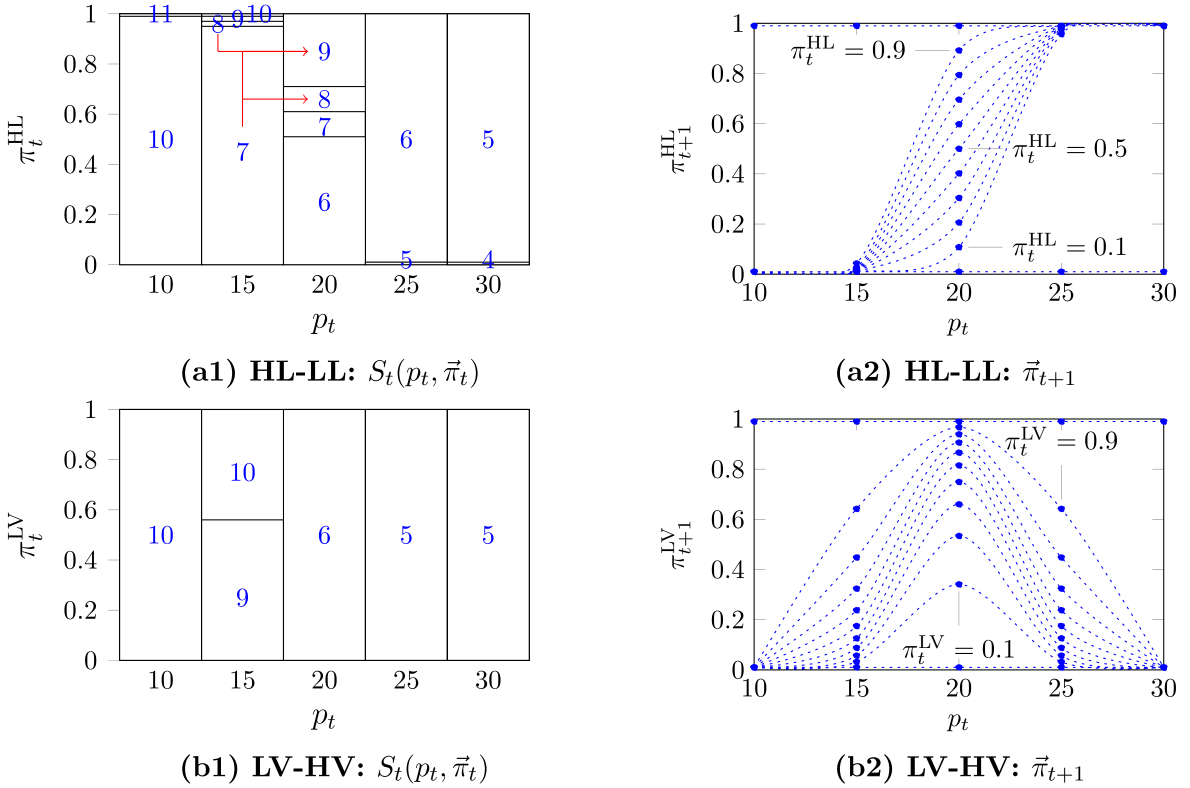

(a1) HL-LL: $S_t(p_t, \vec{\pi}_t)$ **(a2) HL-LL:** $\vec{\pi}_{t+1}$

(b1) LV-HV: $S_t(p_t, \vec{\pi}_t)$ **(b2) LV-HV:** $\vec{\pi}_{t+1}$

Figure 4.4.: Optimal inventory policy $S_t(p_t, \vec{\pi}_t)$ for i.i.d. regime specifications

Note. (a1) and (b1) illustrate S_t as a function of p_t and $\vec{\pi}_t$. Since $m = 2$, $\pi_t^{\mathrm{LL}} = 1 - \pi_t^{\mathrm{HL}}$ and $\pi_t^{\mathrm{HV}} = 1 - \pi_t^{\mathrm{LV}}$. Red arrows highlight settings where S_t increases in p_t due to the violation of monotonicity conditions. (a2) and (b2) illustrate the posterior beliefs π_{t+1}^s for different p_t given prior beliefs π_t^s.

In Figure 4.4a, regimes differ in their price level. A high (low) price p_t increases (decreases) the posterior belief π_{t+1}^{HL} of being in the HL regime (Figure 4.4a2). The higher the prior belief π_t^{HL}, the higher S_t anticipating the prices not to decrease (Figure 4.4a1). S_t is non-decreasing in π_t^{HL} for all p_t and increasing in p_t for $\pi_t^{\mathrm{HL}} \in [0.6; 0.97]$. The latter is because the price increase from $p_t = 15$ to $p_t = 20$ is small compared to the belief jump of being in HL in $t+1$ ($\pi_{t+1}^{\mathrm{HL}} | (p_t = 15)$ vs. $\pi_{t+1}^{\mathrm{HL}} | (p_t = 20)$). Hence, the mean reversion condition of monotonicity *(a2.1)* is violated.

In Figure 4.4b, regimes differ in price volatility. Even though the expected price is the same under both regimes, S_t can vary given the same price (see at $p_t = 15$) depending on π_t^s: if p_t is smaller than the expected price and one has a strong belief about being in the LV regime, i.e., π_t^{LV} is relatively large, then one will tend to order more than under a strong belief about being in HV, where there is a relatively higher chance that the price will further decrease. At rather low prices, a price process with higher volatility bears more chance than risk compared to a process with lower volatility. Consequently, at a given price level, a higher price uncertainty reduces the optimal base-stock level, whereas demand uncertainty yields the opposite effect.

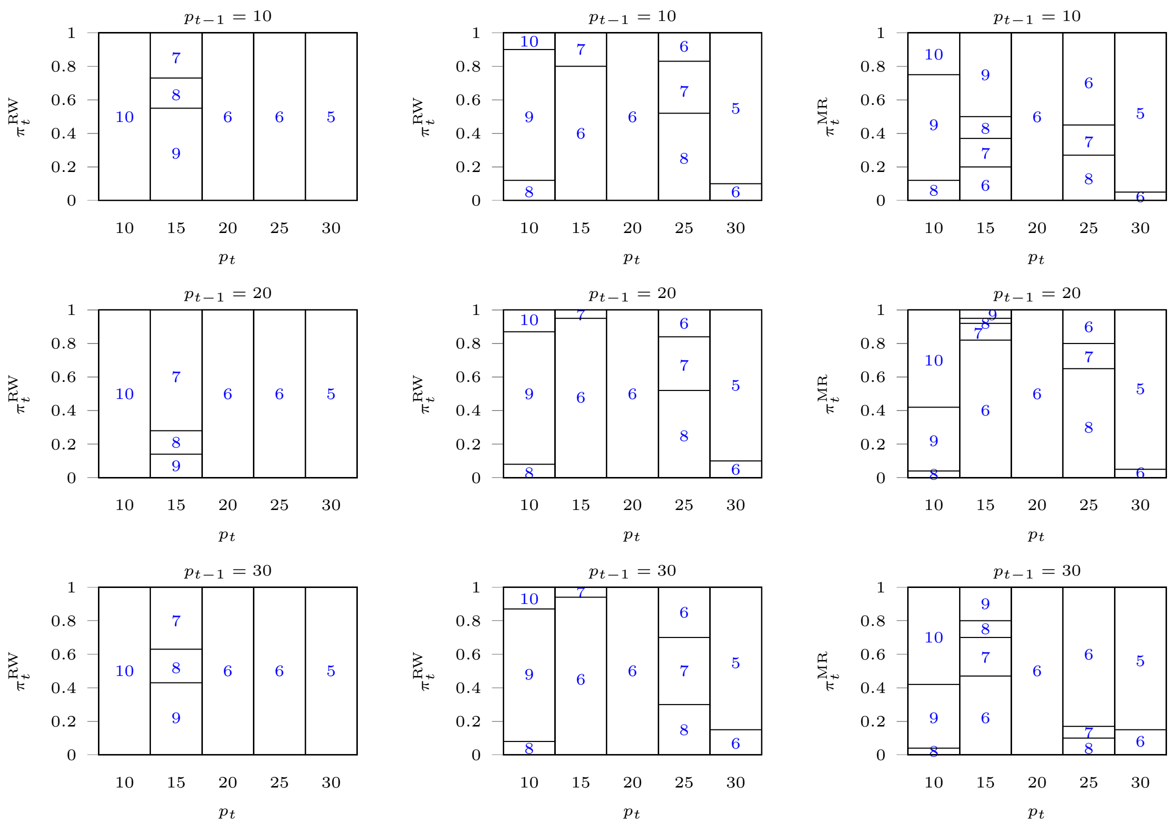

Figure 4.5.: Optimal inventory policy $S_t(p_t, p_{t-1}, \vec{\pi}_t)$ for Markovian regime specifications

Note. Figure 4.5 illustrates S_t as a function of p_t, p_{t-1} and $\vec{\pi}_t$ (equivalently, we could illustrate S_t as a function of p_t and π_{t+1}^s). Since $m = 2$, for (a) $\pi_t^{\text{MR}} = 1 - \pi_t^{\text{RW}}$, for (b) $\pi_t^{\text{MO}} = 1 - \pi_t^{\text{RW}}$ and for (c) $\pi_t^{\text{MO}} = 1 - \pi_t^{MR}$. Inventory policies for the instances $p_{t-1} \in \{15, 25\}$ are not explicitly shown here.

If the regime processes ϕ^s are AR(1), then S_t is a function of p_t and the posterior regime belief $\vec{\pi}_{t+1}$ or equivalently of p_{t-1}, p_t and the prior belief $\vec{\pi}_t$. This is demonstrated in Figure 4.5 for the regime combinations of RW, MR and MO. S_t is monotone in $\vec{\pi}_t$ for all

p_t and p_{t-1}. Since both price processes, RW and MR, satisfy the monotonicity conditions *(a2.1)* and *(a2.2)*, for the setting RW-MR, S_t is non-increasing in p_t for any p_{t-1} and $\vec{\pi}_t$ as the resulting MRS price process is a convex combination of the distinct regime processes. MO violates *(a2.1)*. Hence, for RW-MO and MR-MO, S_t is not necessarily non-increasing in p_t for any p_{t-1} and $\vec{\pi}_t$. With a high impact of the MO regime, i.e., with a sufficiently low π_t^{RW} or π_t^{MR}, S_t increases in p_t, since one would order more at high prices, expecting the price continue to increase.

4.4. Controlled Numerical Study

In order to answer our research questions *(Q2)-(Q4)*, we examine the cost impact in numerical experiments. Due to the curse of dimensionality, we test several more practical inventory control policies (see Section 4.4.3). All SDPs were solved with Matlab2016a on an Intel(R) Core(TM) i7-3770, 3.4 GHz processor with 16 GB RAM.

4.4.1. Setup

We consider the same regime setups HL-LL, LV-HV, RW-MR, RW-MO and MR-MO within a full factorial design: we vary the initial prior regime beliefs $\pi_t^s \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, with $\pi_t^s = 0.5$ representing maximum uncertainty about regime s. We distinguish between three demand settings: high volatility (uniformly distributed on $[0; 30]$), medium volatility (normally distributed with $\mu = 15$ and $\sigma^2 = 3^2$) and no volatility (deterministic with a value of 15). The initial inventory level is zero. The planning horizon is $n = 4$. We vary the regime transition probabilities $k_{ij}(r) = \left(\begin{smallmatrix} 0.5+r & 0.5-r \\ 0.5-r & 0.5+r \end{smallmatrix}\right)$ with $r \in \{0.25, 0.4, 0.49\}$. For $r \to 0$, Bayesian learning is less effective and price p_t implies less confidence for regime predictions. To study the impact of *speculation motives* and *safety motives*, we vary $c_h \in \{1, 6, 15\}$. For $c_h = 15$ (for HL-LL and LV-HV) and $c_h \in \{6, 15\}$ (for RW-MR, RW-MO, MR-MO), equation (4.7) holds for all instances, i.e., inventory decisions are driven exclusively by *safety motives*.

Definition 2 (Non-Speculative Condition)**.** *The expected one-period price increase per unit is less than or equal to the inventory holding costs per unit and unit of time, i.e.,*

$$\sum_{s=1}^{m} \bar{p}_{t+1}^s \, \pi_{t+1}^s - p_t \leq c_h, \tag{4.7}$$

with $\bar{p}_{t+1}^s$ being the expected price for $t+1$ under regime s induced by price p_t.

4.4.2. Cost of Price Regime Misspecification

Prior to evaluating the performance of MRS relative to several suboptimal control policies in Section 4.4.3, we measure the performance loss from misspecifying the price process ϕ^s by calculating the relative cost deviation

$$\Delta\text{COST}(\phi^{(1)}(\phi^{(2)})) := \left(\frac{\text{COST}^{\phi^{(2)}} - \text{COST}^{\phi^{(1)}}}{\text{COST}^{\phi^{(1)}}}\right) \cdot 100\%, \tag{4.8}$$

where $\phi^{(1)}$ is the true price process and $\phi^{(2)}$ is the supposed price process (supposed by the inventory manager). $\text{COST}^{\phi^{(1)}}$ is the minimum total expected cost over the planning horizon $t = 1, ..., n$ induced by optimal first-stage purchase decision $y_1^{(1)}$, and $\text{COST}^{\phi^{(2)}}$ is the total expected cost induced by a (potentially) suboptimal first-stage decision $y_1^{(2)}$ (which is optimal for $\phi^{(2)}$), evaluated based on $\phi^{(1)}$. $\Delta\text{COST}(\phi^{(1)}(\phi^{(2)}))$ is an upper bound for cost savings of MRS under stochastic regime switches (Section 4.4.3). Therefore, ΔCOST is an indicator of the potential to use MRS.

Table A.1 (see Appendix A) summarizes the computational results across all instances. By deciding based on the high-level regime in a low-level environment (LL(HL)), a firm may end up 35.00% above the optimal total expected cost. Supposing that the price reverts to its long-run mean but actually being in a momentum environment (MO(MR)) yields a cost increase of up to 26.29%. Table A.1 shows that the largest ΔCOST occur under $c_h = 1$ (i.e., strong *speculation motives*), low demand volatility and level regimes or Markovian regimes with opposing characteristics of the price process (MR versus MO). Under $c_h = 6$, the *speculation motive* is eliminated for the Markovian regime settings (non-speculative condition (4.7) holds) and strongly reduced for i.i.d. regime settings. Accordingly, the cost saving potential (ΔCOST) decreases or even disappears. From Table A.1, we can deduce the following four main conclusions.

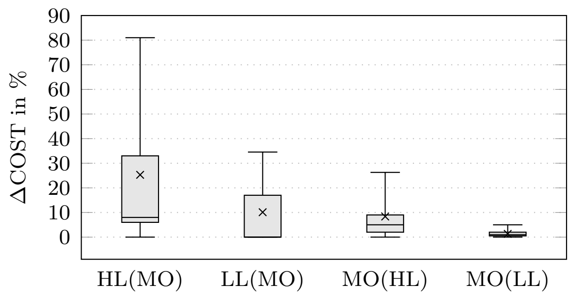

Figure 4.6.: Cost of correlation misspecification (Instances: $c_h = 1$, deterministic demand; Results for all other instances are reported in Table A.1; Boxplot characteristics throughout this thesis: minimum, 1st-, 2nd-, 3rd-quartile, maximum, mean ($\times$))

(i) Misspecifying price autocorrelation yields significant performance losses (Figure 4.6). HL(MO) and LL(MO) represent settings where autocorrelation is supposed but violated by, e.g., unexpected price jumps and drops. In the case of MO(HL) and MO(LL), existing autocorrelation is ignored. The maximum of ΔCOST = 80.78% occurs for HL(MO) at a price of $p_t = 10$ where the inventory manager ignores the price jump and orders $y_1^{\text{MO}} = 15$, rather than $y_1^{\text{HL}} = 60$. For MO(LL) at a price $p_t = 10$, $y_1^{\text{LL}} = 60$, whereas $y_1^{\text{MO}} = 15$ would have been optimal. However, as for MO the price is expected to remain low and c_h is low, the misspecification yields a cost increase of only 5.00%.

(ii) ΔCOST is particularly high if the expected prices under the regimes differ, which is true for level and Markovian regimes. In volatility regimes (HV, LV) with equal price expectations, ΔCOST is at most 2.61% and therefore rather negligible (Figure 4.7).

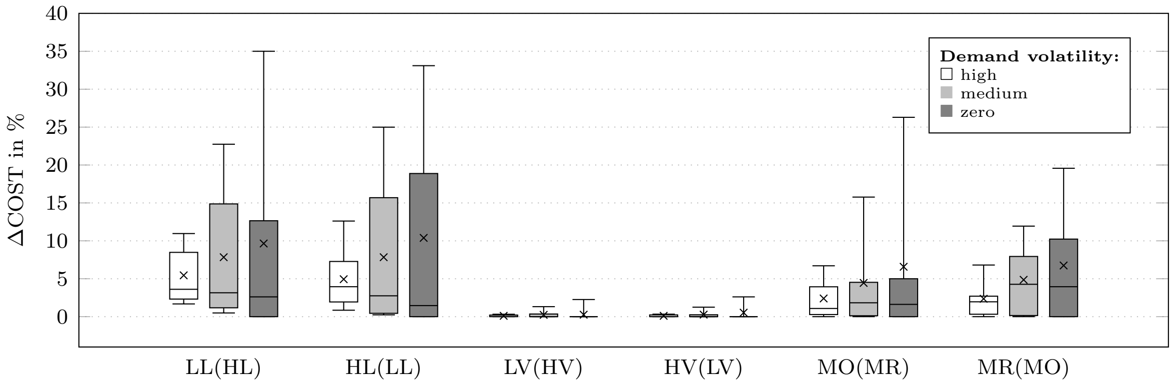

Figure 4.7.: Cost of price process misspecification (Instances: $c_h = 1$; Results for all other instances are reported in Table A.1)

(iii) The lower the demand volatility, the higher the cost of regime misspecification (Figure 4.7) - as long as *speculation motives* dominate. Figure 4.8 shows in a more illustrative way that under *speculation motives*, ΔCOST decreases with increasing demand volatility, which is not the case under pure *safety motives*, due to the additional effect that with an increasing demand volatility, the importance of the *safety motive* increases.

(iv) Speculation drives the cost of price process misspecification. Under non-speculative condition (4.7), ΔCOST is considerably smaller and converging to zero as demand uncertainty approaches zero (Figure 4.8). This is because both *speculation motives* and *safety motives* diminish. If demand is known (no *safety motive*) and equation (4.7) holds (no *speculation motive*), the price process does not affect order decision y_t. Figure 4.8 furthermore shows that the results are not very sensitive to shortage costs c_p that do not affect the dominating *speculation motive*, but solely the inferior *safety motive*.

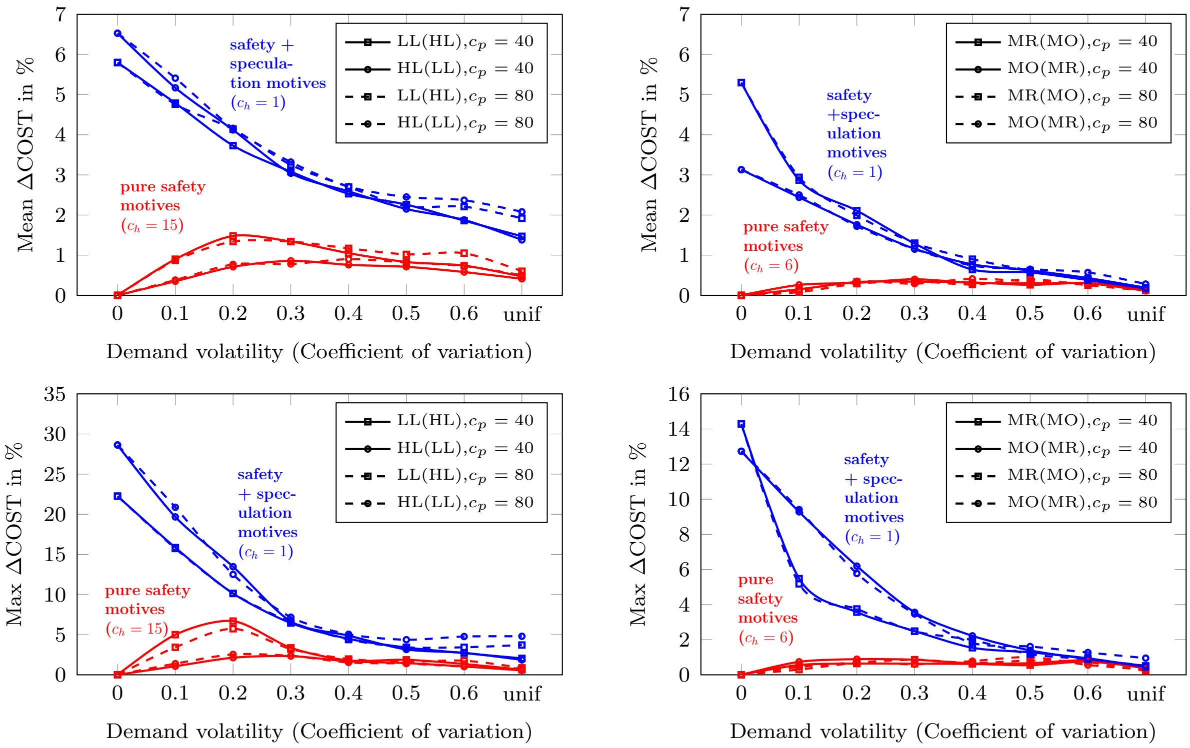

Figure 4.8.: Impact of demand volatility and speculation on cost of misspecification ($n = 2$)

4.4.3. Performance of Suboptimal Control Policies

The Bayesian mechanism of the learning process in equation (4.1) makes the state space $z_t \in Z_t$ grow exponentially in the number of periods n. Each combination of $z_t = (p_t, p_{t-1}, \vec{\pi}_t)$ yields a posterior belief $\vec{\pi}_{t+1}$ representing a next period's prior in the state space. This makes the inventory problem numerically difficult to solve and asks for practical alternatives. Therefore, we compare the optimal MRS policy $S_t(p_t, p_{t-1}, \vec{\pi}_t)$ with the following simpler (but suboptimal) policies that *(i)* ignore price uncertainty (NAIVE), *(ii)* ignore regime switches (SRC) or *(iii)* ignore learning (CEC):

(i) *NAIVE Control (S_t):* The NAIVE policy considers demand uncertainty but neglects price uncertainty and assumes that $\mathbb{E}_t[p_{t+1}] = p_t$. It therefore ignores *speculation motives* and is equivalent to just-in-time procurement without forward buying. This yields price-independent base-stock levels S_t. If demand d_t is deterministic, the NAIVE order quantity is simply given by $y_t = \max\{d_t - I_t, 0\}$.

(ii) *Single Regime Control SRC-R1/SRC-R2* ($S_t(p_t)$)*:* For SRC-R1 (respectively SRC-R2), inventory decisions are based on price process $\phi^{(1)}$ (respectively $\phi^{(2)}$) neglecting switches to price process $\phi^{(2)}$ (respectively $\phi^{(1)}$), which yields price-dependent

base-stock levels $S_t(p_t)$. SRC is equivalent to the policy proposed by Kalymon (1971) for the *full information problem*.

(iii) *Certainty Equivalent Control CEC* ($S_t(p_t, \vec{\pi}_0)$)*:* CEC replaces the doubly stochastic nature of p_{t+1} by a single stochastic counterpart in the form of price process estimate $\hat{\phi}(p_{t+1})$ based on the manager's first prior regime belief $\vec{\pi}_0$. Other than MRS, CEC does not use feedback in terms of dynamic information updates, i.e., $\pi^s_{t+\Delta} = \pi^s_0 \;\; \forall \Delta = 1, ..., n-t$. Therefore, the price process is characterized by $\hat{\phi}(p_{t+1}) = \sum_{s=1}^{m} \pi^s_0 \; \phi^s(p_{t+1})$. Since $\vec{\pi}_0$ does not evolve conditionally on observations of p_t, the size of the state space of CEC equals SRC.

Results

In the following, we benchmark optimal inventory control under regimes (MRS), i.e., $S_t(p_t, p_{t-1}, \vec{\pi}_t)$ respectively $S_t(p_t, \vec{\pi}_t)$, by computing the relative cost deviation if we decide *(i)* based on NAIVE inventory control (S_t), *(ii)* based on a specific regime SRC-R1/SRC-R2 ($S_t(p_t)$) or *(iii)* without market information updates, i.e., CEC ($S_t(p_t, \vec{\pi}_0)$). Note that for *(ii)* and *(iii)*, ΔCOST from Table A.1 defines upper bounds for the percentage above optimal (i.e., MRS) cost. We use the same setup as presented in Section 4.4.1. We distinguish between stochastic (uniform) and deterministic demand. By varying c_h, we again study strong and reduced (for i.i.d. regime processes) respectively eliminated (for Markovian regime processes) *speculation motives*. Additionally, we analyze the effects of regime persistence by varying the switching parameter r in the way that was described in Section 4.4.1. A representative instance is presented in Figure 4.9. The detailed results across all instances are summarized in Table A.2 (see Appendix A).

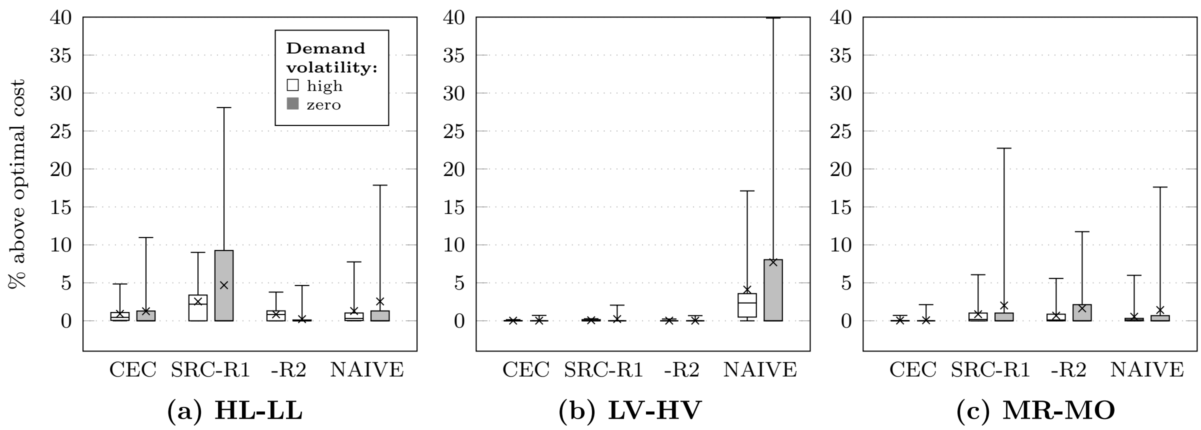

Figure 4.9.: Performance of suboptimal control policies (Instances: $c_h = 1$, $r = 0.49$; Results for all other instances are reported in Table A.2)

(i) NAIVE Control. The NAIVE heuristic performs comparatively poor across all settings. Remarkably, NAIVE can lead to significant losses in a volatility regime environment where the other heuristics (CEC, SRC-R1, SRC-R2) with price-dependent base-stock levels perform close to optimum (see Figure 4.9b). The reason is that NAIVE does not assess the current market price p_t as high or low. Therefore, the lowest imaginable market price is not even classified as a comparatively low price and the highest imaginable market price is not classified as a rather high price. We furthermore observe that with more frequent regime switches, i.e., with increasing uncertainty in the price process, NAIVE tends to perform slightly worse as it does not capture price stochasticity. However, the cost of ignoring price uncertainty and thus the performance of the MRS approach relative to the NAIVE policy decreases significantly with decreasing *speculation motives* (see Table A.2: $c_h = 1$ vs. $c_h = 6$).

(ii) Single-Regime Control (SRC). Ignoring a multi-regime framework and instead control inventory based on a specific single regime (SRC-R1, SRC-R2) yields a cost increase by up to 28.09%. Setting HL-LL implies that, in case of doubt, inventory managers should tend to underestimate the expected price (SRC-R2), rather than overestimating it (SRC-R1) to avoid speculative forward buying. In line with Section 4.4.2, the cost of ignoring regimes increases with decreasing demand volatility (under dominating *speculation motives*). The impact of *speculation motives* is in line with the results of the NAIVE policy. Concerning the impact of the expected frequency of regime switches, there are two contrary effects: on the one hand, with an increasing frequency, i.e., with decreasing r, the price observation gives less indication about the regime in $t+1$ and therefore, the Bayesian mechanism of MRS gets less effective. On the other hand, a lower r leads to higher uncertainty and more room for (mis)speculation. However, we find for most instances (especially for those with high potential performance losses), that relative benefits of Bayesian MRS increase with decreasing frequency of regime switches.

(iii) Certainty Equivalent Control (CEC). The CEC heuristic on average performs very well especially in Markovian and volatility regimes (see also Figure 4.9). The Bayesian MRS approach relative to the CEC heuristic has its highest potential under level regimes and zero demand volatility. In the worst case, ignoring dynamic information updates (CEC) yields 13.33% higher cost (Table A.2). The results furthermore demonstrate that if the demand uncertainty is high, it becomes less important to update the price process beliefs based on new market information received. Under highly volatile demand, a firm is at most 4.89% above optimal cost, while it can lose significantly more (13.33%) by not updating the price forecast as demand is less volatile. The values in parentheses

in Table A.2 show that the cost impact and thus the potential of Bayesian updating is mainly driven by *speculation motives*, i.e., updating the belief leads to speculation (i.e., forward buys), while not updating avoids speculation or (and mostly) vice versa. The opposing effects of the frequency of regime switches (explored by switching parameter r) for *SRC* are also observed for *CEC*.

Comparing the overall results of our controlled numerical study, we find that inventory decisions based on the manager's prior belief (CEC) describe an impressively effective heuristic in Markovian and volatility price regimes (at most 2.87% above optimal cost), while decisions based on any of the regimes (SRC-R1/SRC-R2) perform far worse (22.73% and 11.74%, respectively). NAIVE control can yield significant higher expected cost even and especially in volatility regimes, whereas CEC, SRC-R1 and SRC-R2 perform close to optimum here.

4.5. Results on Empirical Data

Even though an extensive empirical test of different stochastic price models for various spot-traded commodities is not within the scope of this chapter, we want to exemplarily illustrate and test the MRS approach based on real spot market data of corn (HL-LL regimes) and zinc (RW-MR regimes) (see Figure 4.1). We quantify the value of perfect spot price information for inventory control, investigate the relationship between price forecast accuracy and operational performance and point out the crucial role of *misspeculation* induced by stochastic price models in uncontrolled empirical environments.

Procurement of Corn (01-2007 until 04-2016)

Using `MS_Regress` by Perlin (2015) for the corn price from 01-2007 to 04-2016, we identify a high-level ($\mu = 651.76$, $\sigma = 78.97$) and a low-level ($\mu = 370.03$, $\sigma = 46.09$) spot price regime with the transition matrix $k_{ij} = \left(\begin{smallmatrix} 0.95 & 0.05 \\ 0.03 & 0.97 \end{smallmatrix}\right)$. The regime classification measure (RCM) (Ang and Bekaert, 2002) gives an ex-post indication for the goodness of regime specification. It is defined as $\text{RCM} := 100 \cdot m^2 \, \frac{1}{T} \sum_{t=1}^{T} \left(\prod_{s=1}^{m} \hat{\pi}_t^s\right)$, where T denotes the number of historical estimation points (periods) and m the number of regimes. RCM converging to zero means perfect classification, i.e., all data points can clearly be assigned to one of the regimes. For the underlying regime estimation, we obtain RCM $= 1.33$, which indicates good but not perfect classification (Ang and Bekaert, 2002). As an initial prior belief, we use $\vec{\pi}_0 = (0.375, 0.625)$ obtained by solving the system of linear equations from Section 4.2.1. We face a constant deterministic monthly demand of $d_t = 15,000$ bu (bushel) and hence focus on *speculation motives* only. Restricted by practical capacity limits of grain bins, we consider a horizon of $n = 4$ months, i.e., $S_t \in \{d_t, 2d_t, 3d_t, 4d_t\}$. The first-stage order quantity is restricted by $y_1 \leq (\sum_{t=1}^{4} d_t - I_1)^+$. Note that in this capacitated setting with deterministic demand, the terminal valuation of inventory C_{n+1} does not affect the first-stage decisions and hence without effect on the solution can be set to, e.g., $C_{n+1} = 0$ or $C_{n+1} \equiv -\mathbb{E}_n[p_{n+1}] \cdot I_{n+1}$. Shortage penalty costs are $c_p = \$10$/bu and inventory holding costs (in \$/bu/mo) are $c_h \in \{0.01, 0.10, 0.40\}$. A high variation in c_h is reasonable for perishable commodities that require energy-intensive drying processes and are characterized by yield losses through storage.

We compare the following inventory control policies: *(i)* ex-ante optimal policy MRS, *(ii)* CEC and *(iii)* NAIVE. Furthermore, we compare the performance of each policy to the theoretical ex-post optimal policy under perfect spot price foresight (PF). We calculate the mean absolute percentage error (FC-MAPE) for three-step-ahead forecasts (due to three potential forward buying periods) in order to study how price forecast accuracy affects the operational performance of the inventory system.

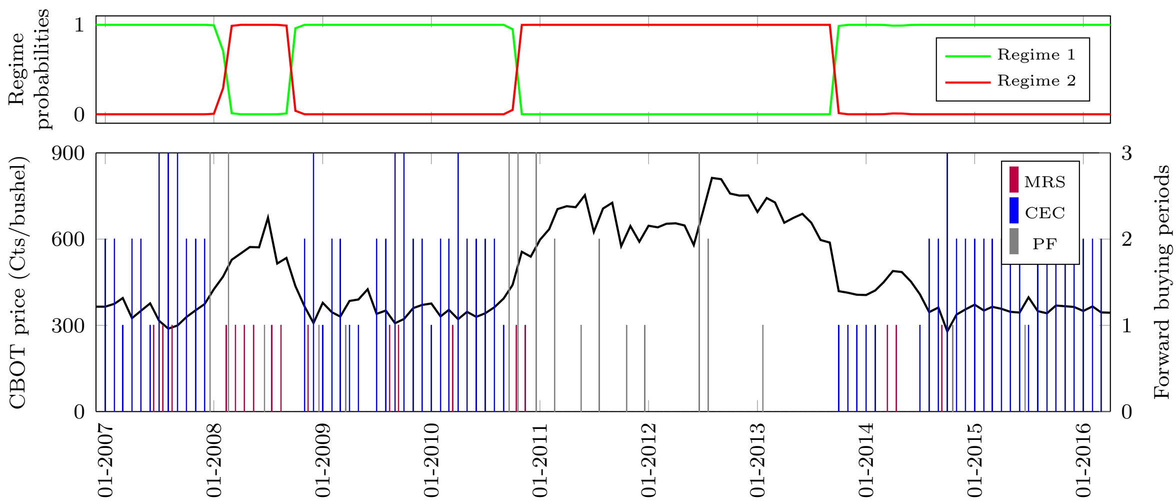

Figure 4.10.: Regime switches and inventory policy implications ($c_h = \$0.4$/bu/mo)

Figure 4.10 illustrates *(i)* the regime switching behavior of the spot price estimated by the standard MATLAB MRS package (see Perlin, 2015) and *(ii)* the corresponding operational order patterns (measured in forward buying periods) according to the different control approaches. Table 4.2 reports the forecast accuracy and operational performance for different phases within the time series. Figure 4.10 demonstrates that even at high c_h, CEC leads to massive forward buying and overstocking at low prices, expecting them to rise. This leads to a significant cost increase (+6.1% or $486,078) relative to MRS that detects (*learns*) the low-price regime and avoids forward buying (Table 4.2). Under reasonably small c_h, i.e., with larger speculation opportunities and a smaller cost of *misspeculation*, MRS clearly outperforms NAIVE. At high c_h, NAIVE, that does not build up any stocks, performs slightly better than MRS and close to PF. We conclude that in an empirical environment, NAIVE can be reasonable if c_h is high. By deciding based on stochastic price models, at high c_h, the negative consequences of speculating wrong (and build up unnecessary stock) are more severe than the positive consequences concerning the average price paid (Table 4.2). We interpret *misspeculative inventory* as the inventory surplus compared to PF.

Table 4.2.: Corn: Performance of different control policies

	c_h = \$0.01/bu/mo				c_h = \$0.1/bu/mo				c_h = \$0.4/bu/mo			
	PF	MRS	CEC	NAIVE	PF	MRS	CEC	NAIVE	PF	MRS	CEC	NAIVE
01-2007 to 12-2008 (transition phase)												
Total cost in 1,000$	1,360	1,503	**1,464**	1,516	1,416	1,589	1,525	**1,516**	1,499	1,558	1,626	**1,516**
% above PF cost	0	10.6	7.6	11.5	0	12.5	7.7	7.0	0	3.9	8.4	1.1
Av. price paid in $/bu	3.76	4.15	4.05	4.21	3.79	4.25	4.07	4.21	3.96	4.18	4.08	4.21
Av. inventory in 1,000 bu	23.1	37.5	23.8	0.0	21.3	24.4	24.4	0.0	7.5	5.6	16.3	0.0
FC-MAPE in %	0	19.78	28.06	16.13	0	19.78	28.06	16.13	0	19.78	28.06	16.13
01-2009 to 07-2010 (low-level regime)												
Total cost in 1,000$	937	**981**	999	1,007	965	**1,002**	1,057	1,007	1,001	1,013	1,168	**1,007**
% above PF cost	0	4.7	6.6	7.4	0	3.9	9.6	4.3	0	1.2	16.7	0.6
Av. price paid in $/bu	3.28	3.42	3.48	3.53	3.29	3.41	3.48	3.53	3.49	3.49	3.43	3.53
Av. inventory in 1,000 bu	18.2	32.4	35.5	0.0	14.2	16.6	33.9	0.0	0.8	2.4	25.3	0.0
FC-MAPE in %	0	9.60	35.53	9.83	0	9.60	35.53	9.83	0	9.60	35.53	9.83
08-2010 to 06-2011 (transition phase)												
Total cost in 1,000$	799	**878**	917	961	832	931	**925**	961	929	967	963	**961**
% above PF cost	0	9.9	14.8	20.3	0	11.9	11.2	15.5	0	4.1	3.7	3.5
Av. price paid in $/bu	4.82	5.31	5.55	5.83	4.82	5.55	5.55	5.83	5.01	5.79	5.73	5.83
Av. inventory in 1,000 bu	35.5	20.5	8.2	0.0	32.7	13.6	8.2	0.0	23.2	2.7	4.1	0.0
FC-MAPE in %	0	12.92	24.31	12.84	0	12.92	24.31	12.84	0	12.92	24.31	12.84
07-2011 to 09-2013 (high-level regime)												
Total cost in 1,000$	2,565	**2,711**	2,745	2,745	2,608	**2,717**	2,745	2,745	2,679	2,745	2,745	**2,745**
% above PF cost	0	5.7	7.0	7.0	0	4.2	5.3	5.3	0	2.5	2.5	2.5
Av. price paid in $/bu	6.32	6.69	6.78	6.78	6.34	6.69	6.78	6.78	6.44	6.78	6,78	6.78
Av. inventory in 1,000 bu	18.9	2.2	0.0	0.0	14.4	2.2	0.0	0.0	6.7	0.0	0.0	0.0
FC-MAPE in %	0	8.41	29.18	8.17	0	8.41	29.18	8.17	0	8.41	29.18	8.17
10-2013 to 03-2016 (transition phase)												
Total cost in 1,000$	1,637	1,752	1,723	**1,719**	1,669	1,746	1,809	**1,719**	1,714	1,734	1,960	**1,719**
% above PF cost	0	7.0	5.3	5.0	0	4.6	8.4	3.0	0	1.1	14.3	0.3
Av. price paid in $/bu	3.62	3.87	3.81	3.82	3.65	3.79	3.81	3.82	3.78	3.81	3.78	3.82
Av. inventory in 1,000 bu	18.5	31.0	33.5	0.0	9.0	13.0	31.0	0.0	1.0	1.5	21.5	0.0
FC-MAPE in %	0	14.31	26.51	10.20	0	14.31	26.51	10.20	0	14.31	26.51	10.20
01-2007 to 03-2016 (overall)												
Total cost in 1,000$	7,267	**7,838**	7,838	7,948	7,473	7,995	8,059	**7,948**	7,821	8,012	8,498	**7,948**
% above PF cost	0	7.9	7.9	9.4	0	7.0	7.8	6.4	0	2.4	8.7	1.6
Av. price paid in $/bu	4.35	4.69	4.69	4.77	4.37	4.71	4.69	4.77	4.53	4.75	4.70	4.77
Av. inventory in 1,000 bu	22.7	26.5	22.7	0.0	17.4	14.5	21.9	0.0	6.1	2.4	15.3	0.0
FC-MAPE in %	0	12.99	28.83	11.10	0	12.99	28.83	11.10	0	12.99	28.83	11.10

On average, MRS is a good compromise between *(i)* CEC as MRS avoids overstocking (exploitation of low inventories) and *(ii)* NAIVE as MRS still captures price uncertainty and allows for forward buys (exploitation of low prices).

Taking a closer look into the performance within the time series, we observe that MRS strictly outperforms CEC inside the regime phases (01-2009 to 07-2010, 07-2011 to 09-2013). That supports that MRS is mainly beneficial under fewer switches in regimes, i.e., under high regime retention times. During transition phases, no dominant order policy is observable. Unsurprisingly, none of the policies can predict sudden switches in the price behavior as the high percentage cost gap to PF confirms, e.g., for 08-2010 to 06-2011 (high value of perfect spot price information). MRS is unable to forecast regime switches, but is able to detect them through its Bayesian learning scheme.

Confirming the empirical finance literature, CEC is strictly dominated by MRS in terms of forecast accuracy. However, the quality of inventory decisions cannot simply be explained by the quality of price predictions (Figure 4.11). Over all data points from Table 4.2, FC-MAPE and the percentage cost gap to PF are positively but not perfectly correlated with Pearson correlation coefficient of $\rho = .35$ ($p < .01$). While c_h, and therefore the degree of speculation, does not impact FC-MAPE, it strongly impacts operational performance. If c_h is low (high), one speculates more (less) in terms of forward buying and therefore relies more (less) on forecast accuracy.

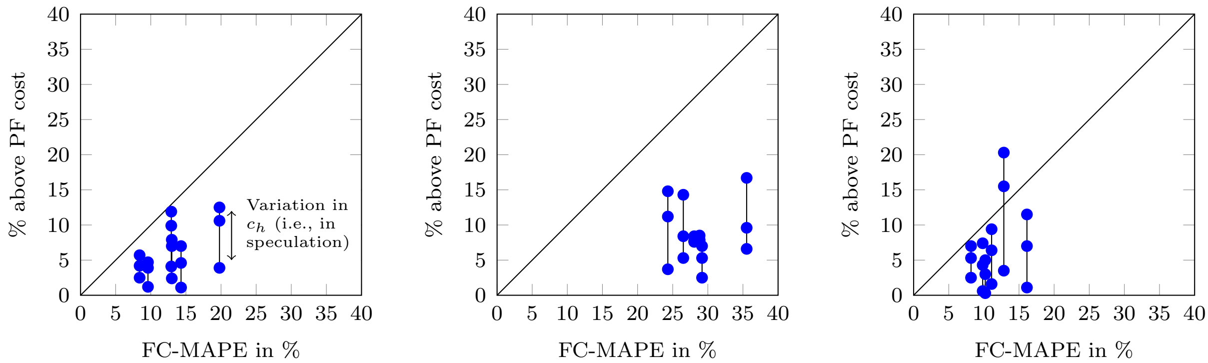

Figure 4.11.: Price forecast accuracy vs. inventory performance

Procurement of Zinc (01-2004 until 04-2016)

By estimating zinc price regimes from 01-2004 to 04-2016, we identify a random walk (RW) regime with AR(1) coefficient $\beta_1 = 0.96$ and a mean reversion regime (MR) with $\beta_1 = 0.69$. Switching occurs with transition probabilities $k_{ij} = \left(\begin{smallmatrix} 0.99 & 0.01 \\ 0.01 & 0.99 \end{smallmatrix}\right)$. As an initial prior belief, we again use the steady-state probabilities $\vec{\pi}_0 = (0.5, 0.5)$. Price transition probability matrices are generated from the estimated AR(1) processes using the method proposed by Tauchen (1986). We assume a constant deterministic monthly demand of 1 ton. Capacity limits restrict the planning horizon to $n = 4$ periods. Shortage penalty costs are $c_p = \$6000$/ton. Moreover, we vary the inventory holding cost c_h as it may have a major impact on inventory performance. Beside the control policies MRS, CEC and NAIVE, we additionally investigate the consequences deciding based on the RW regime (SRC-RW) that characterizes the zinc price evolution from 2004 to 2009.

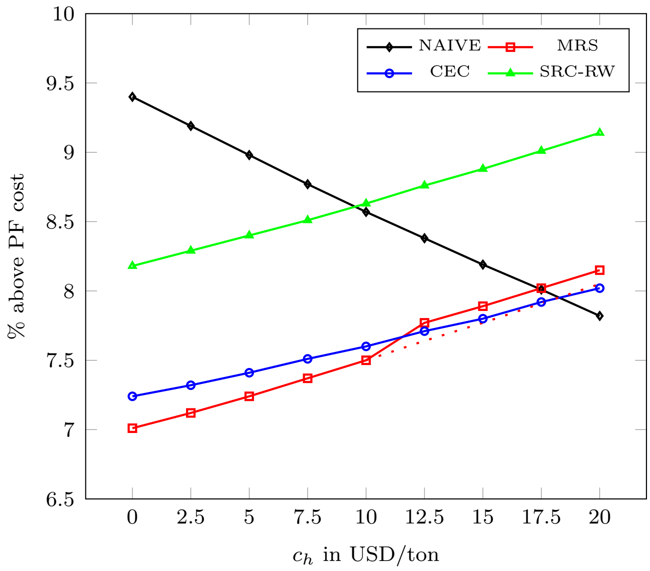

Figure 4.12.: Zinc: Performance of different control policies

While we observe that MRS (FC-MAPE: 9.21%) again clearly outperforms CEC (14.38%) with regard to price forecast accuracy, Figure 4.12 demonstrates the ambiguous character of speculation in the inventory control context. Applied to real market data, stochastic price models (monotonic increasing curves) from time to time inevitably lead to *misspeculation*. At low c_h, the cost of inventory from *misleading speculation* is less than the benefits of exploiting low purchasing prices (see Table 4.3) and therefore inventory control based on stochastic price models clearly outperforms NAIVE order control. With increasing c_h, inventory from *misleading speculation* gets more expensive and hence, stochastic price models perform worse relative to NAIVE without speculative inventory.

We furthermore see that pure random walk assumptions (i.e., not updating/learning the price process) lead to significant performance losses relative to MRS that realizes the

Table 4.3.: Average inventory in tons and average price paid in USD/ton

	c_h in USD/ton								
Control policy	0	2.5	5	7.5	10	12.5	15	17.5	20
PF	1.48	1.47	1.45	1.41	1.41	1.39	1.39	1.36	1.34
MRS	2.48	2.48	2.48	2.48	**2.48**	**2.39**	2.39	2.39	2.39
CEC	2.24	2.24	2.24	2.24	2.24	2.24	2.24	2.24	2.24
SRC-RW	2.42	2.42	2.42	2.42	2.42	2.42	2.42	2.42	2.42
NAIVE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
PF	1,890.1	1,890.1	1,890.3	1,890.4	1,890.5	1,890.5	1,890.5	1,890.8	1,890.9
MRS	2,022.5	2,022.5	2,022.5	2,022.5	**2,022.5**	**2,026.4**	2,026.4	2,026.4	2,026.4
CEC	2,026.9	2,026.9	2,026.9	2,026.9	2,026.9	2,026.9	2,026.9	2,026.9	2,026.9
SRC-RW	2,044.7	2,044.7	2,044.7	2,044.7	2,044.7	2,044.7	2,044.7	2,044.7	2,044.7
NAIVE	2,067.8	2,067.8	2,067.8	2,067.8	2,067.8	2,067.8	2,067.8	2,067.8	2,067.8

best purchase prices across all control policies. However, CEC that similar to SRC-RW does not consider price history (p_{t-1}), but considers both regimes with equal weight (due to $\vec{\pi}_0$), seems to be a good approximation in this Markovian setting. At $c_h = \$12.5$/ton, MRS reacts to the increasing cost of inventory and the base-stock levels are adjusted (i.e., reduced) in order to reduce forward buying. Nevertheless, applied to the real-market data, reduced average inventory goes along with a higher average purchase price such that in this specific case it would have been better to stick to the more speculative MRS control policy from $c_h \in \{0, ..., 10\}$ as the dashed line illustrates.

Both case studies give insights about the factor of *misleading speculation* induced by model and calibration errors of stochastic price models that are typically neglected in controlled experiments. In empirical and speculative environments, the costs of *misleading speculation* (in terms of *misspeculative inventory*) need to be carefully traded off against the price benefits from speculation in order to decide whether it might even be better to ignore price uncertainty and order NAIVE. Furthermore, the case studies demonstrate that better price forecasts (FC-MAPE) do not necessarily result in better inventory performance that strongly depends on the c_h-structure and hence the degree of speculation opportunities determined by equation (4.7).

4.6. Conclusion

We present managerial implications and give a summary and research outlook.

Managerial Implications

Our results that are summarized in Table 4.4 show under which conditions it is worth it for inventory managers to deal with spot price models and consider the price process under regimes and Bayesian learning.

Table 4.4.: Potential of MRS price models in inventory control

Setting	High potential	Low potential
Regime classification	strong	weak
Demand volatility	low	high
Regime type	level	volatility
Non-speculation condition (Holding cost)	violated (low)	fulfilled (high)
Frequency of regime switches	low	high

Impact of Regime Classification. There needs to be evidence for different price regimes that can be identified by, e.g., the Baum-Welch algorithm (Baum et al., 1970). The regime classification measure (RCM) (Ang and Bekaert, 2002) gives an ex-post indication for the quality of regime specification. The higher the RCM, the higher the impact of transition phases, which can hardly be captured by any price model as demonstrated in Section 4.5.

Impact of Demand Volatility. Under high demand uncertainty, the operational impact of the price process is smaller than in case of a low demand uncertainty. Therefore, it is more promising to consider price regimes and Bayesian learning when demand is less volatile, except when there is a high impact of the *safety motive* relative to the *speculation motive*. We demonstrated that supposing the wrong regime in a regime switching framework under deterministic demand can lead to 35.00% higher cost, while supposing the same regime under highly volatile demand in contrast only results in 10.96% higher cost.

Impact of Price Regime Setting. Structural breaks affecting price levels are more important to consider than breaks in price volatility. In our numerical tests, MRS volatility models did not have a significant cost impact in inventory control. Supposing the wrong volatility regime did not exceed values of 2.61% above the optimal cost, while supposing the wrong level or Markovian regime led to higher cost of 35.00% and 26.29%, respectively. Even though decisions based on a random regime (SRC-R1, SRC-R2) or on the manager's prior belief (CEC, i.e., no *learning*) seems to be acceptable in case of volatility regimes, NAIVE control can lead to significant performance losses.

Impact of Speculation. Speculation opportunities give an indication about when it may be worth considering MRS price models. Our numerical results show that it is significantly more beneficial to consider a regime switching framework when the non-speculative condition (4.7) is violated. While in the regime setting MR-MO under violated non-speculative condition, the relative cost increase supposing the wrong price regime (i.e., MR) is 6.71%, 15.77% and 26.29%, respectively, the increase is 0.87%, 0.44% and 0.00% under non-violation. The operational benefit of MRS, and of spot price models in general, is even zero as soon as demand is deterministic (no *safety motive*) and the non-speculative condition holds (no *speculation motive*). Nevertheless, in an uncontrolled empirical environment under speculation opportunities but relatively high holding cost, the impact of *misspeculation* induced by stochastic price models is that high that the NAIVE policy performs quite well (and in some cases even better than sophisticated stochastic price models). However, low holding costs recommend MRS price models for two reasons: *(i)* an increase in speculation motives that increases the potential value of a more accurate price model (in terms of price forecasts) and *(ii)* a decrease in the cost of *misspeculative inventory* that, due to inevitable model and calibration errors, is inherent in stochastic price models. Contrary to holding costs (that influence *speculation motives*), shortage penalty costs (that influence *safety motives*) have a minor impact on the cost of price process misspecification.

Impact of Regime Persistence. The performance of the Bayesian MRS approach is influenced by the expected frequency in regime switches. Counterintuitively, less frequent switches, i.e., higher regime retention times, support the performance of MRS. If a high frequency of price process changes is expected, then the current spot price provides less information about being in a specific regime in the future. Therefore, learning-based MRS is generally more beneficial if the expected time spent in a specific price cycle is large. This is also confirmed by the empirical results of the case study since the MRS approach dominates especially within regime phases.

Summary and Outlook

Spot-traded commodities are characterized by switches in price regimes that are only partially observable. Based on experimental and empirical data, we demonstrated in an inventory control context how partially observable spot price processes can be considered by using non-linear hidden Markov regime switching techniques with learning via dynamic Bayesian information updates. We showed that considering random regime shifts can lead to significantly different order patterns compared to expecting the price to remain in the current business cycle (SRC) or compared to inventory control without dynamic Bayesian information updates (CEC). Our findings show that, if the regimes are observable, similar to Kalymon (1971), a price-dependent order-up-to policy is optimal. If the regimes are partially observable, an order policy that depends only on the current spot price is no longer sufficient and prior respectively posterior regime beliefs need to be considered. We presented regime conditions under which the optimal base-stock level is non-increasing in prices and illustrated regime settings where the monotonicity of the base-stock level in the spot price is violated. Furthermore, we empirically showed that an MRS approach can lead to significant price forecast improvements and cost savings compared to inventory control based on state-of-the-art price models that neglect price regimes and wrongly presume full price process information. Conditions under which the price process matters give managers an indication when it is worth considering sophisticated price forecasting with stochastic spot price models and when it might be adequate to make naïve inventory decisions. An imperfect correlation between the quality of price predictions and the quality of inventory decisions implies that finance (spot price models) and operations (inventory control) cannot be separated and that the suitability of price models should be evaluated with regard to operational cost implications, rather than merely by forecast accuracy measures such as MAPE.

Further research opportunities are, e.g., a continuous time analysis of the stochastic control problem and an extension to an infinite planning horizon that requires partially observable Markov decision processes (POMDP) (see, e.g., Monahan, 1982). Furthermore, it might be interesting to investigate the operational value of price models if the assumption of uncorrelated commodity spot price and finished product's demand is relaxed. One could also extend the procurement problem to a trader's setting with reselling (to the spot market) opportunities or multiple procurement channels, including financial derivatives such as forwards, futures and options contracts. Testing the performance of further suboptimal control policies such as open-loop feedback control or limited look-ahead policies (see, e.g., Bertsekas, 1995) to overcome the curse of dimensionality of the partially observable decision problem could also be insightful. Moreover, more extensive empirical tests of spot price models on a broad range of commodities are needed in the operational context.

Chapter 5.

Financial Hedging from a Data-Driven Procurement Perspective

Based on
Mandl, C. and S. Minner (2019b). Data-driven optimization for commodity procurement under price uncertainty.

Volatile commodity prices require efficient risk management for purchasing. We study a practice-motivated multi-period stochastic commodity procurement problem with forward and spot purchase options. Existing approaches for optimizing positions in the forward contract market are based on parametric price models, which inevitably involve *price model misspecification* and *generalization error*. We propose a non-parametric, data-driven approach (DDA) that is consistent with the optimal procurement policy structure but without requiring the a-priori specification and estimation of stochastic processes. In addition to historical prices, DDA is able to leverage real-time feature data (*Big Data*), such as economic indicators, in solving the problem. This chapter provides a framework for prescriptive analytics in dynamic commodity procurement, with optimal purchase signals directly learned from data as functions of features, via mixed integer linear programming (MILP) under cost minimization objectives. Furthermore, we combine optimization with performance-based regularization from machine learning (ML) to extract decision-relevant data from noise. Based on controlled numerical experiments and empirical data, we show that there is a significant value of feature data for commodity procurement. However, overfitting deteriorates the performance of data-driven solutions, which asks for ML extensions that improve out-of-sample generalization significantly. Our approach is embedded in the IT architecture of an industry partner used for natural gas procurement. Compared to the firm's best practice benchmark, DDA saves on average 9.1 million Euro p.a. (4.33%) for ten years of backtesting.

5.1. Introduction

Substantial price volatility at commodity markets, along with violated assumptions underlying the Modigliani-Miller theorem (Modigliani and Miller, 1958), i.e., the existence of market frictions such as transaction costs, taxes, financial distress costs or information asymmetry, provide rationale for corporate risk management for both risk-neutral and risk-averse commodity-processing firms (Froot et al., 1993; Smith and Stulz, 1985). Besides operational hedging through speculative stockpiling (*Chapter 4*), which is often physically restricted, hedging price risk via contracts plays an increasing role in practice; especially with growing liquidity of commodity derivatives markets. The number of exchange-traded financial derivatives on energy, agriculture, precious metals and non-precious metals increased by 22.6% to 4.6 billion contracts between 2014 and 2015 (FIA, 2015). A large-scale empirical study by Bartram et al. (2009) shows that 50.4% of the oil-processing companies and 30.5% of the steel-processing companies have implemented some kind of commodity price risk hedging through financial contracts such as forwards, futures, swaps and call options (Hull, 2005). A prominent example for a firm that has benefited greatly from commodity hedging is the food manufacturer General Mills, which realized hedging gains of $151 million in volatile agricultural and energy markets during the first quarter of 2008 (Wall Street Journal, 2008). On the other hand, by contractually hedging future demand, firms become too inflexible to react to price declines. In 2015, the world's second-largest airline United lost $960 million, the world's third-largest airline Delta even $2.3 billion by hedging 100% of its fuel costs via forward contracts prior to the big drop in oil prices (Wall Street Journal, 2016a).

However, in theory and practice, there is still no consensus on the optimal hedging strategy (Wang et al., 2015). There is no common answer to the question if a contractually offered forward price in period t is sufficiently low to hedge demand of $t+2$ or better to wait and satisfy the demand from the forward market in $t+1$ or the spot market in $t+2$ in order to minimize the total cost of purchase.

To support operational decision-making in the multi-period context, stochastic dynamic programming (SDP) is used, which supposes full knowledge of the underlying commodity price process (e.g., Geometric Brownian motions or mean-reverting Ornstein-Uhlenbeck processes). However, solving stochastic control problems by means of dynamic programming is subject to two major drawbacks: *(i)* due to the *curse of dimensionality*, SDP is highly impracticable for real-world problems and high-dimensional spot and forward price models, as they are widely used in commodity finance (see Section

3.1). (*ii*) SDP relies on the assumption of being able to fully characterize the underlying stochastic processes and their parameters (*full information problem*), which is not true in practice. It therefore ignores the impact of two errors that inevitably affect the out-of-sample performance: *(a) price model error* (the error that occurs from model misspecification) and *(b) generalization error* (the out-of-sample error that occurs from in-sample overfitting) (Wang et al., 2015).

In addition, low-dimensional price models that are still manageable by SDP might not fully exploit all information available. Microeconomics and empirical finance (e.g., Pindyck, 2004; Pirrong, 2011) provide evidence of the impact of exogenous variables on commodity prices such as economic climate, temperature, interdependence with other commodities, inventories, convenience yield or exchange rates (*fundamental analysis*).

In this context, Arthur D. Little (2014) states: *"Suppose your organization procures plastic pellets. If your database is linked with market price data for crude oil and macroeconomic forecast data, a Big Data solution can constantly discover new opportunities and alert your organization to act, e.g., to renegotiate contracts as soon as there is a significant decrease in the price of crude oil."*

In practice, there is indeed a growing interest in *Big Data analytics.* Service providers (e.g., Quandl or Thomson Reuters) enter the market, offering real-time economic and financial data specific to commodity purchasers' needs (Boston Consulting Group, 2017). Others provide satellite imagery to forecast crop yields (CME Group, 2014) or to track ships and hence follow the movement of seaborne commodities (Reuters, 2016b), both of which can give an early indication for price movements.

Motivation from Practice

This chapter is motivated by a collaboration with a large chemical company that runs its own gas-fired power plants to generate steam and electricity for its energy-intensive production processes and for the power market. To hedge price risk in purchasing natural gas, the firm has forward procurement options at the European TTF gas market, however no storage capacity. Prior to our collaboration, the firm decided without optimization-based decision support by distributing its significant gas demand (10×10^6 MWh p.a.) equally among forward contracts of different maturities (typically 1- to 4-months-ahead futures contracts). However, having access to historical and real-time price and feature data, the purchase team wondered *(i)* how to efficiently operationalize data for optimization-based purchasing and *(ii)* whether analytics-based decisions yield significant reduction in total cost of purchase. Therefore, we developed a pragmatic and computationally tractable data-driven approach (hereafter, DDA) for dynamic decision-support that can easily be embedded in the firm's existing IT architecture of online databases such as Thomson Reuters Datastream and Eikon and the firm's existing data management system (Figure 5.1).

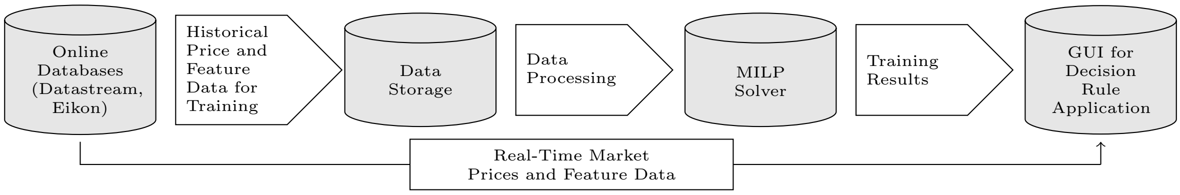

Figure 5.1.: DDA embedded in the IT architecture of our industry partner

Even though this research is motivated by an industry application, this chapter adds both methodological and managerial contribution to the existing literature.

Methodological Contribution

We present a generic data-driven approach to compute the parameters of the optimal policy in multi-period procurement problems, which yields optimally structured policy rules in the sense of the corresponding full information SDP under full knowledge of the underlying price process(es). This is fundamentally different from the existing data-driven literature that trains decisions in single-period problems (see, e.g., Ban and Rudin, 2019), which is inappropriate for constrained multi-stage problems (Bertsimas and Kallus, 2016). DDA does not require the specification and estimation of stochastic price processes or to make any statistical assumptions a-priori to optimization (*distribution-free*). It does not demand an explicit analysis of commodity spot and term structure dynamics and their existing relationship (see Geman, 2005, p. 73). Instead, DDA learns the underlying stochastic processes from data via Empirical Risk Minimization (ERM) (Vapnik, 1998, p. 32) in a robust-regression-like way within a MILP. DDA is general, as it comprises a variety of standard price models, and is also able to exploit high-dimensional real-time feature data for dynamic spot and forward procurement decisions. Therefore, DDA exhibits several benefits for generating procurement plans: *(i)* It prevents *price model misspecification* and avoids the *curse of dimensionality* of SDP from modeling future spot and forward prices by high-dimensional price models. *(ii)* DDA avoids *discretization errors*, e.g., from lattice approximations of continuous-time stochastic processes. *(iii)* DDA works with the true loss function (cost of purchase), rather than with intermediate loss functions (e.g., least squares). Hence, DDA does not separate prediction and optimization and therefore explicitly captures the interplay between finance and operations in the sense of price modeling and decision-making (e.g., *(a)* if demand is high, the cost of price misspecification is higher than if demand is low or *(b)* prediction models with coefficient of determination $R^2 \ll 1$ are sufficient for hindsight optimal decisions). *(iv)* Modeling the problem as a MILP offers additional flexibility in terms of problem-specific operational constraints.

Moreover, we focus on implications of in-sample optimization on out-of-sample procurement performance. Generalization is a central target in statistical learning theory (Hastie et al., 2013; Vapnik, 1998), however widely overlooked in the operations literature (Section 3.2). The *generalization error* is affected by overfitting or underfitting, i.e., by using feature data too extensively or too little. We combine data-driven optimization with ML-based regularization for the selection of decision-relevant (rather than prediction-relevant) features. By these means, the prescription problem is perturbed to reduce the *generalization error* by adding bias to the estimator and improving variance.

Managerial Contribution

Even though this research is motivated by the procurement of natural gas, our models are fairly general and applicable to many different commodity settings. The decision rules from DDA are easy to interpret and easy to operationalize and allow for real-time decision-making along the commodity forward curve.

A primary emphasis of the present chapter is to quantify the economic value of high-dimensional feature data in a multi-period procurement context. Therefore, we introduce several prescriptive performance measures (e.g., the *Prescription Error* and the *Value of Feature Information*) that evaluate prediction quality based on the cost of decision, rather than based on prediction error such as least squares loss. We backtest our models on empirical data for the European gas market TTF. The results that we obtained in close collaboration with our industry partner demonstrate significant out-of-sample cost reductions compared to the firm's status quo and various established benchmarks from the literature (e.g., reoptimization, AR(1) models and the featureless approach). We show that some of the best-in-practice benchmarks (e.g., REO and AR(1)) are special cases of DDA. Additionally, we sensitize the *Big Data*-driven firm towards the *generalization error*, i.e., decisions that perform well in-sample can work quite poorly out-of-sample. We demonstrate that even (or especially) in a *Big Data* environment, the firm needs to investigate carefully which data to use (*Smart Data*) and we show how interpretable performance-based machine learning algorithms can provide support.

Research Questions and Organization

We address the following research questions: *(Q1)* How can firms efficiently operationalize high-dimensional feature data for commodity procurement under price uncertainty? *(Q2)* How to combine data-driven procurement with ML in order to support the selection of decision-relevant (rather than prediction-relevant) features with the objective of reducing the out-of-sample *generalization error*? *(Q3)* What is the economic value of feature data and analytics for commodity-purchasing firms?

Section 5.2 formalizes the problem and presents the DDA models. Section 5.3 introduces performance bounds and several prescriptive performance metrics. Section 5.4 illustrates the main effects in controlled numerical experiments based on the specification of a true but unknown underlying multivariate price model. Section 5.5 tests the empirical performance for the procurement of natural gas. Section 5.6 concludes.

5.2. Model Formulation

Even though we have developed data-driven models for all major contract types (forwards/futures, swaps, European and American call options), we focus on forwards and futures contracts in the following. The other models are available upon request.

5.2.1. Problem Setting

We consider a multi-period, discrete-time, periodic-review procurement problem. In each period t (e.g., month), the firm decides on forward procurement quantities y_t^τ for delivery period $t+\tau$ by signing a forward/futures contract with time to maturity $\tau \in \mathcal{F}^+ = \mathcal{F}\backslash\{0\}$ at a nominal unit price p_t^τ quoted at the beginning of period t. Our firm is a price taker and has no access to storage capacity, which is reasonable for commodities, such as energy or for just-in-time production environments. Period demand $d_t \in \mathcal{D}$ is known (accurately forecasted) and needs to be satisfied (forced compliance) at the latest through the spot market ($\tau = 0$) at a spot price p_t^0. Market capacity is infinite, i.e., markets are supposed to be sufficiently liquid. Delivery lead time at the spot market is zero. Excess quantities cannot be resold to the market (no trading).

The problem can be formulated as a standard SDP under the Bellman equation

$$C_t(\vec{I}_t, \vec{F}_t, x_t) = \min_{\substack{y_t^\tau \geq 0 \\ I_t^0 + y_t^0 \geq d_t}} \left\{ \sum_{\tau \in \mathcal{F}} p_t^\tau \; y_t^\tau + \mathbb{E}_t\Big[C_{t+1}(\vec{I}_{t+1}, \vec{F}_{t+1}, x_{t+1})\Big] \right\} \qquad \forall t = 0, ..., n. \quad (5.1)$$

C_{t+1} denotes the cost-to-go affecting the here-and-now purchase decisions y_t^τ based on the stochastic evolution of the forward curve $\vec{F}_t = (p_t^\tau : \tau \geq 0)$. The state space $z_t \in Z_t$ of the SDP is characterized by $z_t = (\vec{I}_t, \vec{F}_t, x_t)$. $\vec{I}_t = (I_t^\tau)$ is the position of the firm in the forward market at the beginning of period t (*endogenous state information*) that evolves according to the balance equation $I_t^\tau + y_t^\tau = I_{t+1}^{\tau-1}$ with I_t^0 being the forward procurement quantity for delivery period t that has been hedged prior to t. $I_t^0 + y_t^0 \geq d_t$ ensures that d_t needs to be satisfied latest through spot purchases y_t^0. $\vec{F}_t$ is the deterministic currently-quoted forward curve including the current spot price p_t^0 (*exogenous state information*). Future forward curves $\vec{F}_{t+1}$, including future spot prices p_{t+1}^0, are stochastic and traditionally modeled by exogenous (high-dimensional) price models $\phi(p_{t+1})$ that serve as input for the SDP. x_t denotes the (unknown) state variables (i.e., features) that drive the stochastic evolution of $\vec{F}_t$ by the (unknown) function $\vec{F}_{t+1} = \phi(\vec{F}_t, x_t)$.

Theorem 2 (Structure of the Optimal Procurement Policy)**.** *The optimal procurement policy is characterized by state-dependent price thresholds P_t^τ. For all $\tau \in \mathcal{F}^+$,*

$$y_t^{\tau*}(x_t) = \begin{cases} [d_{t+\tau} - I_t^\tau]^+ & \text{if } p_t^\tau \leq P_t^\tau(x_t), \\ 0 & \text{if } p_t^\tau > P_t^\tau(x_t), \end{cases} \tag{5.2}$$

i.e., if the currently quoted forward price p_t^τ is smaller than or equal to $P_t^\tau(x_t)$, then the unhedged demand $d_{t+\tau}$ is purchased via a forward contract with time to maturity τ.

Proof. see B.1 □

However, computing $P_t^\tau(x_t)$ via dynamic programming raises two issues: *(i)* The *curse of dimensionality* in the case of multi-factor price models that result in a high-dimensional state space $z_t \in Z_t$ and *(ii) price model error*, i.e., we actually cannot characterize the exogenous feature part x_t of the state space $z_t \in Z_t$ (and consequently $P_t^\tau(x_t)$) without full information about the true price model ϕ that is typically not known (e.g., in terms of relevant price features) since mostly past data is available.

Therefore, we propose an entirely data-driven approach (DDA) without statistical assumptions when capturing uncertainty. Rather than pre-determining x_t and therefore the state space $z_t \in Z_t$, DDA exploits the optimal state space characterization $z_t^* \in Z_t$ by using historical data from periods $t = 1, ..., T \subseteq \mathbb{N}_{>0}$ to train the model for out-of-sample periods $t = T, ..., T'$, with $T' \geq T$. The firm's database $\mathfrak{D} = \{(\vec{F}_t, d_t, X_{it})\}_{t=1,...,T, i=1,...,N}$ includes historical forward curves $\vec{F}_t = (p_t^\tau : \tau \geq 0)$, demand time series $d_t \in \mathcal{D}$ and time series for feature realizations $X_{it} \in \mathcal{X} \subset \mathbb{R}^{N \times T}$ of features $i = 1, ..., N$, with $\mathcal{X}$ defining the feature space. Potential features might include lagged prices (ARIMA models), analyst forecasts or economic and market indicators. Thus, correlation structures are considered without explicitly being a-priori modeled. Feature realizations $X_{it} \in \mathcal{X}$ and prices along the forward curve $\vec{F}_t$ are observable prior to decision-making in t, and decision stages t correspond to contract maturities, which have e.g. monthly occurrence.

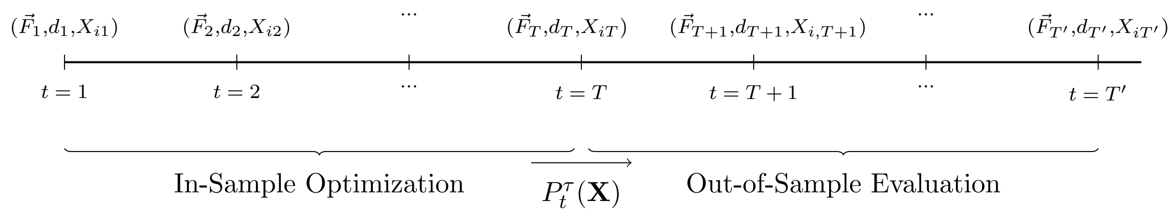

Figure 5.2.: Training and evaluation framework

We consider two stages: At stage 1, the parameters of the optimal procurement policy are trained in-sample ($t = 1, ..., T$). For each forward contract $\tau \in \mathcal{F}^+$, the training provides a state-dependent price threshold $P_t^\tau(\mathbf{X})$ as a function of the feature vector $\mathbf{X} = (X_{1t}, X_{2t}, ..., X_{Nt})$. $P_t^\tau(\mathbf{X})$ gives a dynamic purchase signal indicating whether or not to lock in the currently-quoted forward price p_t^τ. $P_t^\tau(\mathbf{X})$ may change in response to changing information about market states. At stage 2, $P_t^\tau(\mathbf{X})$ is evaluated out-of-sample (e.g., $t = T, ..., T'$, with $T' \geq T$) conditional on new feature realizations X_{it}.

5.2.2. Linear Decision Rule Approximation

We postulate an affine (i.e., linear plus a constant) decision rule approximation that is *(i)* consistent with the optimal procurement policy structure in the sense of the corresponding SDP, *(ii)* computationally tractable and *(iii)* easy to operationalize.

Definition 3 (Linear Decision Rule)**.** *For all available forward contracts $\tau \in \mathcal{F}^+$ on the forward curve $\vec{F}_t$, the firm procures the unhedged demand of period $t+\tau$ via contract τ if the following conditions are satisfied, otherwise procurement is postponed (real option):*

(i) Condition I: There are no forward contracts owned in period t for future period $t+\tau$, i.e., the demand of the corresponding future period is not fully purchased yet.

(ii) Condition II: The linear decision rule (linear as a function of the features) gives a condition-based purchase signal if

$$p_t^\tau \leq P_t^\tau(\boldsymbol{X}) := \sum_{i=0}^{N} \beta_i^\tau X_{it}. \tag{5.3}$$

Note that this approach is fundamentally different from the existing data-driven literature (e.g., Ban and Rudin, 2019) that trains decisions y_t^τ rather than policy parameters P_t^τ as a linear function of features. The standard LDR approach with $y_t^\tau := \sum_{i=0}^N \beta_i^\tau X_{it}$ would yield inconsistent and potentially suboptimal decisions if $y_t^\tau > d_{t+\tau} - I_t^\tau$ (overage).

In the regression-like equation (5.3), $\beta_i^\tau \in \mathcal{B} \subset \mathbb{R}$ are feature coefficients that are unknown to the decision maker and must be learned. β_0^τ represents the feature-independent intercept term, i.e., $X_{0t} = 1\ \forall t = 1, ..., T$. Even though the state-dependent price threshold $P_t^\tau(\mathbf{X})$ is a linear combination of features, this is not restrictive. Non-linearities can be considered by interaction terms (e.g., $P_t^\tau(\mathbf{X}) := \beta_0^\tau + \beta_1^\tau X_{1t} + \beta_2^\tau X_{2t} + \beta_3^\tau X_{1t}X_{2t}$) or by polynomial terms (e.g., $P_t^\tau(\mathbf{X}) := \beta_0^\tau + \beta_1^\tau X_{1t}^2$). Lagged observations (e.g., ARIMA models) consider correlation across time periods and offer additional flexibility (e.g., $P_t^\tau(\mathbf{X}) := \beta_0^\tau + \beta_1^\tau X_{1,t} + \beta_2^\tau X_{1,t-1} + \beta_3^\tau X_{1,t-2}$). In this regard, equation (5.3) includes many time series models as special cases, e.g., linear first-order autoregressive price models (AR(1)) of the form $p_t = \beta_0 + \beta_1 p_{t-1} + \epsilon_t$ with random error term $\epsilon_t \sim N(0, \sigma_t^2)$. AR(1) models that are widely used in the commodity procurement literature (see Section 3.2) specify typical commodity price behavior, such as random walk ($p_t = p_{t-1} + \epsilon_t$), mean reversion ($p_t = \kappa\mu_p + (1-\kappa)p_{t-1} + \epsilon_t$ with $\kappa \in [0, 1)$ as the mean-reversion speed and μ_p as the mean reversion level) or momentum ($\beta_1 > 1$).

However, rather than by a feedforward mechanism (prediction-based procurement decisions based on price forecasts) via ordinary least squares (OLS) or similar standard regression techniques, $\beta_i^\tau \in \mathcal{B}$ of features $i = 0, ..., N$ for each forward contract $\tau \in \mathcal{F}^+$ are trained via a feedback mechanism (prescription-based procurement decisions) solving an MILP based on the statistical learning theory principle of Empirical Risk Minimization (ERM; Vapnik 1998, pp. 32), i.e., $\min_{\beta_i^\tau \in \mathcal{B}} \left\{ \frac{1}{T} \sum_{t=1}^T \ell_{\mathrm{DDA}}(\hat{C}_t, C_t^{\mathrm{PF}}) \right\}$, with ℓ_{DDA} being the loss function of DDA minimizing the empirical cost $\hat{C}$ of DDA relative to the theoretical *perfect foresight policy* PF ("oracle problem"). Minimizing the loss with respect to the nominal optimization problem is equivalent to setting the coefficients $\beta_i^\tau \in \mathcal{B}$ of the linear threshold functions $P_t^\tau(\mathbf{X})$ such that the feature-conditional cost of purchase is minimized, i.e.,

$$\min_{\beta_i^\tau \in \mathcal{B}} \left\{ \frac{1}{T} \sum_{t=1}^{T} \sum_{\tau \in \mathcal{F}} \hat{C}_t(p_t^\tau, \beta_i^\tau) | \mathbf{X} \right\}. \tag{5.4}$$

Notation ˆ emphasizes costs that are estimated from data via the ERM principle, rather than the expected cost under the *full information problem.*

5.2.3. Data-Driven Models for Policy Parameter Optimization

The data-driven nature of the optimization problem (5.4) is similar to a two-stage stochastic program (see Shapiro et al., 2009), with $\beta_i^\tau \in \mathcal{B}$ being time-independent (i.e., scenario-independent) first-stage decisions and all other decision variables being second-stage scenario decisions (recourse decisions that are pre-determined by β_i^τ and feature data X_{it}). Each data observation $t = 1, ..., T$ is per se equally weighted. However, seasonality in $d_t \in \mathcal{D}$ can give greater weight to periods with high demand, which is not captured by sequential approaches with price prediction prior to optimization.

Table 5.1.: General notation (Additional notation defined as required)

Sets	
$t = 1, ..., T$	In-sample periods (training)
$\tau \in \mathcal{F}^+ = \mathcal{F}\backslash\{0\}$	Forward contract with time to maturity τ ($\tau = 0$ denotes the spot market option)
$i = 0, ..., N$	Features ($i = 0$: feature-independent intercept)
Parameters	
p_t^τ	Forward price in period t with contract maturity $t+\tau$
X_{it}	Realization of feature i in period t
d_t	Demand in period t
Decision variables	
q_t^τ	Binary purchase indicator for forward contract τ in period t
β_i^τ	Regression-like coefficient for feature i under forward contract τ

Additionally, due to regime switches and structural breaks, commodity prices typically do not come from the same data-generating process, which may yield misleading OLS estimates. Instead, we apply a least absolute deviation (LAD) regression-like approach that is robust to violations of the underlying OLS regression assumptions (i.e., particularly homoscedasticity and that outliers are rare) (Andersen, 2008, pp. 47).

The following *Big Data* model (DDA-BD) with the notation given in Table 5.1 trains the state-dependent purchase signals $P_t^\tau(\mathbf{X}) \coloneqq \sum_{i=0}^{N} \beta_i^\tau X_{it}$ of the LDR by exploiting all available feature information $\{(X_{it})\}_{t=1,...,T,i=0,...,N}$.

DDA-BD:

$$\min_{\beta_i^\tau \in \mathcal{B}} \quad \hat{C}^{\text{BD}} = \frac{1}{T}\sum_{t=1}^{T} \sum_{\tau\in\mathcal{F}|\tau\leq T-t} \left[p_t^\tau \; d_{t+\tau} \; q_t^\tau\right] \tag{5.5}$$

$$\text{s.t.} \quad \sum_{\tau\in\mathcal{F}|\tau\leq t-1} q_{t-\tau}^\tau = 1 \qquad \forall t = 1, ..., T \tag{5.6}$$

$$- M(1 - q_t^\tau) \leq \sum_{i=0}^{N} \beta_i^\tau X_{it} - p_t^\tau \qquad \forall t = 1, ..., T - \tau; \; \tau \in \mathcal{F}^+ \tag{5.7}$$

$$M\Bigg(q_t^\tau + \sum_{a \in \mathcal{F} | a \leq t+\tau-1 \cap a > \tau} q_{t+\tau-a}^a\Bigg) > \sum_{i=0}^{N} \beta_i^\tau X_{it} - p_t^\tau \qquad \forall t = 1, ..., T - \tau; \; \tau \in \mathcal{F}^+ \tag{5.8}$$

$$q_t^\tau \in \{0, 1\}, \beta_i^\tau \in \mathbb{R} \qquad \forall t = 1, ..., T; \tau \in \mathcal{F}; \; i = 0, ..., N \tag{5.9}$$

The objective (5.5) minimizes the historical average period and demand-weighted cost of spot and forward procurement quantities by means of feature data with respect to the linear decision rule framework from Definition 3. Because of the all-or-nothing property of the optimal policy, the ordering part only includes the binary variable q_t^τ. Constraint (5.6) guarantees that the demand $d_t \in \mathcal{D}$ is either satisfied by spot purchases ($\tau = 0$) or by a forward contract $\tau \in \mathcal{F}^+$ signed in period $t - \tau$. Constraints (5.7)-(5.8) control the execution of the threshold rule: First, for all forward contracts $\mathcal{F}^+$, $q_t^\tau = 0$ if $p_t^\tau > \sum_{i=0}^{N} \beta_i^\tau X_{it}$. Second, for all contracts, $q_t^\tau = 1$ if $p_t^\tau \leq \sum_{i=0}^{N} \beta_i^\tau X_{it}$, unless the demand has not already been satisfied in a previous period, which is captured by $q_{t+\tau-a}^a$. Hence, constraint (5.8) disables the execution of the threshold rule if there has already been an earlier purchase with regard to the period under consideration. Note that there is no price threshold estimated for the spot market ($\tau = 0$), which is the latest procurement option.

As a special case of DDA-BD, the *Small Data* model (DDA-SD) corresponds to the degenerate featureless intercept-only approach ($N = 0$) that offers a time- and feature-independent constant price threshold $P^\tau := \beta_0^\tau$ for each forward contract $\tau \in \mathcal{F}^+$.

To find the model specification with the highest (in-sample) explanatory power for a given dimension $\bar{N}$, we refer to the *Best Subset Selection Problem* (DDA-BSSP) presented in Appendix B.2.

As $P_t^\tau(\mathbf{X}) := \sum_{i=0}^{N} \beta_i^\tau X_{it}$ is unbounded within the MILP (as prediction and optimization is done simultaneously), we cannot derive bounds for *Big M*. However, one may set $M \geq \hat{p}$, with $\hat{p}$ as upper price limit, determined by, e.g., the price of substitutes. To avoid *Big M* and potentially resulting instability issues, we use *indicator constraints* (see Appendix B.3) for all our numerical tests in Section 5.4 and Section 5.5.

Remark 1 (Non-Uniqueness of the β-Solution)**.** *In linear OLS regression, the function* $\min_{\beta \in \mathcal{B}} \sum_{t=1}^{T} (p_{t+\Delta t} - \sum_{i=0}^{N} \beta_i X_{it})^2$ *has a single local minimum if* $T \geq N$*, i.e., no other linear model yields more accurate predictions. In least absolute deviation (LAD) regression,* $\min_{\beta \in \mathcal{B}} \sum_{t=1}^{T} |p_{t+\Delta t} - \sum_{i=0}^{N} \beta_i X_{it}|$ *that can be solved by an LP may possess infinitely many optimal solutions. Thus, DDA-BD and DDA-SD, similar to LAD, can lead to alternate optimal solutions with regard to* β_i^τ*, i.e., different values for* β_i^τ *can yield the same optimal solution* q_t^τ*, which is due to the threshold structure of the decision problem.*

Even though Remark 1 gives interesting insights about the relationship between prediction and prescription (decision), as it implies that differently accurate price forecasts still might yield optimal procurement decisions, it might not be desirable with regard to model generalization. Therefore, we extend DDA to performance-based regularization.

5.2.4. Data-Driven Models under ML-Based Regularization

The solutions to the DDA formulation from Section 5.2.3 can be highly unstable with poor generalization capability (see, e.g., Ban et al. (2018b) in the context of portfolio optimization). To avoid that the model fits the noise in the data rather than the underlying function, we apply ℓ_1-regularization for decision-based feature selection. Regularization penalizes non-zero coefficients in order to keep the model from relying too heavily on individual data points. We consider two ML-based regularization methods, however with the objective of high-quality decisions rather than high-quality predictions: *(i) Lasso regression* that sets certain coefficients $\beta_i^\tau \in \mathcal{B}$ to zero and yields sparse solutions and *(ii) ridge regression* that shrinks the size of the coefficients $\beta_i^\tau \in \mathcal{B}$ towards zero.

Performance-Based Lasso Regression

We use ℓ_1-norm regularization $R(w) = ||w||_1 = \sum_{i=1}^{N} w_i^\tau$, where w_i^τ is equal to 1 if $|\beta_i^\tau| > 0$ and 0 otherwise. We employ ℓ_1-norm regularization instead of regularization by higher norms (e.g., ℓ_2 with $R(w) = ||w||_2^2 = \sum_{i=1}^{N} \left(w_i^\tau\right)^2$) for robustness reasons and in order to avoid non-linear (e.g., quadratic) terms in the objective function. Furthermore, ℓ_1-norm is a sparse predictor, i.e., it tends to produce sparse non-zero coefficients $\beta_i^\tau \in \mathcal{B}$ and hence has feature selection property (Hastie et al., 2013, p. 141). Additionally, it also reduces multicollinearity by removing independent variables.

Lasso regression controls overfitting by selecting a subset of decision-relevant features rather than selecting prediction-relevant features via model selection criteria, such as AIC or BIC. The overall objective is to simultaneously minimize cost and model complexity, which is achieved by the following Lasso-regularized optimization problem.

DDA-ML1:

$$\min_{\beta_i^\tau \in \mathcal{B}} \quad \underbrace{\frac{1}{T}\sum_{t=1}^{T}\sum_{\tau\in\mathcal{F}|\tau\leq T-t}\left[p_t^\tau \; d_{t+\tau} \; q_t^\tau\right]}_{:=\hat{C}^{\mathrm{ML1}}} \quad + \lambda \underbrace{\sum_{i=1}^{N}\sum_{\tau\in\mathcal{F}^+} w_i^\tau}_{\text{Model complexity}} \tag{5.10}$$

$$\text{s.t.} \quad (5.6)-(5.9) \tag{5.11}$$

$$Mw_i^\tau \geq \beta_i^\tau \qquad \forall i=1,...,N;\tau\in\mathcal{F}^+ \tag{5.12}$$

$$-Mw_i^\tau \leq \beta_i^\tau \qquad \forall i=1,...,N;\tau\in\mathcal{F}^+ \tag{5.13}$$

$$w_i^\tau \in \{0,1\} \qquad \forall i=1,...,N;\tau\in\mathcal{F}^+ \tag{5.14}$$

$\lambda \geq 0$ is a parameter to control regularization. λ is typically calibrated via cross-validation (see Mohri et al., 2012, p. 28). Constraints (5.12) and (5.13) ensure for all $i = 1,...,N$ that $w_i^\tau = 1$ if $|\beta_i^\tau| > 0$ and zero otherwise. The intercept β_0^τ is not regularized, which avoids that $P_t^\tau(\mathbf{X}) = 0$ for large λ, which would lead to permanent spot procurement. For $\lambda = 0$, DDA-ML1 reduces to DDA-BD. For $\lambda \to \infty$, DDA-ML1 converges to DDA-SD as for all $i = 1,...,N$, β_i^τ is set to zero.

Performance-Based Ridge Regression

In ridge regression, the size of the coefficients $\beta_i^\tau \in \mathcal{B}$, rather than the number of non-zero coefficients, is trimmed. The objective is to minimize both cost and model complexity, which is achieved by the following regularized optimization problem where $\beta_i^{\tau,\mathbf{abs}}$ denotes the absolute value of β_i^τ.

DDA-ML2:

$$\min_{\beta_i^\tau \in \mathcal{B}} \quad \underbrace{\frac{1}{T}\sum_{t=1}^{T}\sum_{\tau\in\mathcal{F}|\tau\leq T-t}\left[p_t^\tau \; d_{t+\tau} \; q_t^\tau\right]}_{:=\hat{C}^{\mathrm{ML2}}} \quad + \lambda \underbrace{\sum_{i=1}^{N}\sum_{\tau\in\mathcal{F}^+} \beta_i^{\tau,\mathbf{abs}}}_{\text{Model complexity}} \tag{5.15}$$

$$\text{s.t.} \quad (5.6)-(5.9) \tag{5.16}$$

$$\beta_i^{\tau,\mathbf{abs}} \geq \beta_i^\tau \qquad \forall i=1,...,N;\tau\in\mathcal{F}^+ \tag{5.17}$$

$$\beta_i^{\tau,\mathbf{abs}} \geq -\beta_i^\tau \qquad \forall i=1,...,N;\tau\in\mathcal{F}^+ \tag{5.18}$$

$$\beta_i^{\tau,\mathbf{abs}} \in \mathbb{R} \qquad \forall i=1,...,N;\tau\in\mathcal{F}^+ \tag{5.19}$$

In order to ensure that the objective function for ridge regression works properly, regularization requires a-priori feature scaling to standardize the magnitude of feature data (Hastie et al., 2013, p. 63). If features are not comparable in their magnitude, $\beta_i^\tau \in \mathcal{B}$ might be of different magnitude as well, which influences shrinkage. For feature scaling, the ML literature (Hastie et al., 2013, p. 400) typically applies z-scores, with X'_{it} as the standardized feature value and σ as the standard deviation:

$$X'_{it} = \frac{X_{it} - \frac{1}{T}\sum_{t=1}^{T} X_{it}}{\sigma_{X_{it}}} \tag{5.20}$$

5.3. Performance Bounds and Performance Metrics

In this section, we derive bounds on the in-sample performance and introduce different performance measures for our data-driven prescriptions that are later used to interpret the numerical results.

We can derive in-sample cost bounds for DDA-BD, DDA-SD and DDA-ML that enhance our general understanding of the relationship between the different data-driven model formulations. An important reference value is the lower cost bound C^{PF} achieved by the truly optimal policy under perfect (deterministic) price information ("oracle"), i.e., forward quantity decisions y_t^τ based on perfect predictions of p_t^τ $\forall \tau \in \mathcal{F}, t = 1, ..., T$. Note that perfect price predictions are a sufficient, but not necessary, condition to achieve C^{PF}. C^{PF} can be determined by the following LP:

$$C^{\mathrm{PF}} = \min \frac{1}{T} \sum_{t=1}^{T} \sum_{\tau \in \mathcal{F}} \left[p_t^\tau \; y_t^\tau \right] \tag{5.21}$$

$$\text{s.t.} \quad \sum_{\tau \in F | \tau \leq t-1} y_{t-\tau}^\tau = d_t \qquad \forall t = 1, ..., T \tag{5.22}$$

$$y_t^\tau \geq 0 \qquad \forall t = 1, ..., T; \tau \in \mathcal{F} \tag{5.23}$$

Remember that the *perfect foresight problem* (5.21)-(5.23) is not equivalent to the *full information problem* where the decision maker has full information about the price process $\phi(p_{t+1})$ and its parameters but future prices are still uncertain.

Proposition 2 (In-Sample Relationships).

(i) $C^{\mathrm{UB}} \geq \hat{C}^{\mathrm{BD}} \geq C^{\mathrm{PF}}$. *$C^{\mathrm{UB}}$ describes the upper cost bound (worst case procurement policy) as obtained by solving the optimization model (5.21)-(5.23) as a maximization problem.*

(ii) $\hat{C}^{\mathrm{BD}} \to C^{\mathrm{PF}}$ *for* $N \to \infty$. *DDA is trained with respect to the loss function of the prescription problem, i.e.,* $\ell_{DDA} = |\hat{C} - C^{\mathrm{PF}}|$. *If there is no multicollinearity between features* $i = 1, ..., N$, *then with increasing number of features* N, *the data-driven solution converges to the perfect foresight optimum (PF) as the model can fit the underlying function more accurately.*

(iii) *For* $N \geq T$, $\hat{C}^{\mathrm{BD}} = C^{\mathrm{PF}}$. *If the number of features* N *is greater than or equal to the number of demand periods* T, *the model is able to give an individual price threshold* $P_t^{\tau}(\boldsymbol{X})$ *to each in-sample period* t *(overfitting). Consequently,* $\hat{C}^{\mathrm{BD}}$ *is equal to* C^{PF}. *Again, this is only true if the multicollinearity between features* $i = 1, ..., N$ *is sufficiently small.*

(iv) $\hat{C}^{\mathrm{BD}} \leq \hat{C}^{\mathrm{SD}}$. *This follows from (ii).*

(v) $\hat{C}^{\mathrm{BD}} \leq \hat{C}^{\mathrm{ML}}$. *This is true as* $\lambda \geq 0$ *and for all* $i = 1, ..., N$ *and* $\tau \in \mathcal{F}^+$, $w_i^{\tau} \geq 0$ *(DDA-ML1) respectively* $|\beta_i^{\tau}| \geq 0$ *(DDA-ML2). The idea of ML in data-driven optimization is to overcome overfitting and therefore accept an in-sample performance loss.*

(vi) $\hat{C}^{\mathrm{BD}} \leq C^{\mathrm{REO}}$. *The popular reoptimization approach (REO) (see Section 5.5) uses deterministic forward curve information from* t, *i.e.,* $\vec{F}_t$, *as predictor for future periods. REO is a special case of DDA-BD if* p_t^{τ} $\forall \tau \in \mathcal{F}$ *is used as feature information within equation (5.3).*

Based on the performance bounds, we introduce several prescriptive performance measures that measure the value of data from a decision rather than a prediction perspective.

Prescription Error (PE)

In order to evaluate and compare the data-driven models and support model selection, we introduce the *Prescription Error (PE)*, which is the prescriptive equivalent to prediction error measures, such as the mean absolute percentage error (MAPE).

Definition 4 (Prescription Error)**.**

$$\text{PE} \coloneqq \left(\frac{\hat{C} - C^{\text{PF}}}{C^{\text{PF}}} \right) \cdot 100\% \geq 0$$

PE is the percentage cost deviation from the theoretical cost C^{PF} *under perfect price information, i.e., under perfect foresight predictions of* $p_t^\tau \quad \forall \tau \in \mathcal{F}, t = 1, ..., T$.

PE can be interpreted as the maximum amount a firm should pay for perfect price forecasts. However, note that perfect price predictions are a sufficient but not necessary condition for perfect prescriptions (decisions), i.e., for PE = 0. If PE = 0, the feature set explains the price evolution as accurately such that the procurement performance cannot be improved (not even by more accurate price models).

Value of Feature Information (VFI)

In order to assess the overall prescriptive content of the feature data $\{(X_{it})\}_{t=1,...,T,i=1,...,N}$, we adopt the coefficient of prescriptiveness introduced by Bertsimas and Kallus (2016).

Definition 5 (Value of Feature Information)**.**

$$\text{VFI} \coloneqq 1 - \frac{\hat{C}^{\text{BD}} - C^{\text{PF}}}{\hat{C}^{\text{SD}} - C^{\text{PF}}} = 1 - \frac{\text{PE}^{\text{BD}}}{\text{PE}^{\text{SD}}} \leq 1$$

VFI *is unitless and determines the value of DDA-BD compared to DDA-SD.*

VFI is the prescriptive equivalent to R^2 from predictive analytics. If VFI is small (low prescriptiveness), then feature data X_{it} provides little information for prescribing an optimal decision, or DDA-BD is unable to effectively use the information in X_{it}. If VFI is large (high prescriptiveness), then feature data X_{it} provides valuable information to significantly reduce purchasing cost. VFI can be used as an in-sample and out-of-sample performance indicator: In-sample, VFI is non-decreasing in the number of features N and bounded by $0 \leq$ VFI ≤ 1. Out-of-sample, VFI is bounded by $-\infty \leq$ VFI ≤ 1 as X_{it} might provide useless or even disadvantageous (noisy) information for purchase decisions (i.e., VFI < 0).

Value of Integrating Estimation and Optimization (VIEO)

A major benefit of the integrated data-driven approach (DDA), if compared to the sequential separated estimation and optimization approach (SEO), is that there is no need to conduct regression forecasts for all contract maturities τ on the forward curves. Instead, the firm only needs to solve a single MILP that can easily be extended to ML techniques for decision-based feature selection (Section 5.2.4).

However, there might also be a performance value in integrating prediction and optimization. SEO follows a feedforward mechanism that determines procurement decisions predictively, based on price forecasts that were made without consideration of cost objectives. As opposed to SEO, DDA follows a feedback mechanism that determines procurement decisions prescriptively, by considering cost implications. E.g., price forecasts are more important in periods of high demand (e.g., during winter for natural gas) and less important in periods of low demand. In order to assess the value of integrating prediction and optimization, we introduce the performance measure VIEO.

Definition 6 (Value of Integrating Estimation and Optimization)**.**

$$\text{VIEO} \coloneqq \left(\frac{\hat{C}^{\text{SEO}} - \hat{C}^{\text{DDA}}}{\hat{C}^{\text{DDA}}} \right) \cdot 100\%$$

$\hat{C}^{\text{SEO}}$ *is the associated purchase cost if* $\beta_i^\tau \in \mathcal{B}$ *(and hence* $P_t^\tau(\boldsymbol{X})$*) is estimated predictively under OLS or LAD objectives (dependent variable* $p_{t+\tau}^0$*, independent variables* X_{it}*, normally distributed error term).* $\hat{C}^{\text{DDA}}$ *is the purchase cost if* $\beta_i^\tau \in \mathcal{B}$ *is estimated prescriptively under cost objectives.*

There are mainly three reasons for different purchase decisions based on prediction and prescription: *(i)* Different objectives of predictive and prescriptive approaches. While the predictive regression approach aims at minimizing the sum of the squared residuals (OLS) or absolute residuals (LAD), the prescriptive approach aims at cost minimization. *(ii)* Predicting the price thresholds $P_t^\tau(\mathbf{X})$ in a sequential manner does not leverage any problem structure. E.g., if a future period's demand is already hedged, the firm may not hedge again, even if the current forward price is lower than the expected future price. *(iii)* Other than SEO, DDA considers finance-operations interdependencies by minimizing cost, rather than the sum of the squared residuals. E.g., price predictions need to be more accurate for high-demand periods and less accurate for low-demand periods.

5.4. Controlled Numerical Study

Prior to using real data in Section 5.5, we illustrate the value of DDA in controlled numerical experiments based on simulation. All solutions were obtained without optimality gap with the Xpress-MP solver (version 7.6) on an Intel(R) Core(TM) i7-3770, 3.4 GHz processor with 16 GB RAM, which is a rather conservative choice compared to what is available in industry. Computation times are reported in Appendix B.4.

5.4.1. Setup

Inspired by the problem setting at our industry partner, we construct a simulation experiment with monthly procurement decisions. We simulate time series split into training periods ($t = 1, ..., T$) and evaluation periods ($t = T, ..., T'$). To study sample size effects, we vary the training set size ($T \in \{24, 48, 72\}$) (e.g., 2 years, 4 years, 6 years). The size of the evaluation period (test set size) is fixed to $T' - T + 1 = 48$ (for exemplary price sample paths, see Figure B.2 of Appendix B.5). All experiments are based on Monte Carlo simulation. We evaluate every instance on 100 independent simulation runs. To guarantee fair comparisons, we use common random numbers.

(i) Procurement Options. We assume that the firm has two procurement options: spot purchases and one-period-ahead forward purchases (front-month contract) with the relationship between forward price and spot price described by $p_t^1 = p_t^0 + \varepsilon_t$, with $\varepsilon_t \sim N(0, (\frac{p_t^0}{100})^2)$. If $\varepsilon_t > 0$, the market is in contango (upward sloping forward curve), if $\varepsilon_t < 0$, the market is in backwardation (downward sloping forward curve). $\varepsilon_t = 0$ is reasonable for markets under the spot-future parity condition $p_t^\tau = p_t^0 \cdot e^{(r-Y)\tau}$ with interest rate $r = 0$ and convenience yield $Y = 0$ (Geman, 2005, p.35). For ease of illustration, we only consider one maturity (one-period-ahead forward contracts). However, in the case study in Section 5.5, we allow for various maturities. The state-dependent price threshold of the linear decision rule that gives a signal about forward purchases is defined as $P_t(\mathbf{X}) := \beta_0 + \sum_{i=1}^{10} \beta_i X_{it}$, i.e., the firm has access to a database with 10 features available for decision-making.

(ii) Simulation of Price Trajectories. We suppose that the true but unknown underlying multivariate stochastic spot price process $\phi(p_{t+1}^0)$ (*full information problem*) has the following form:

$$p_{t+1}^0 = \beta_0^{\text{true}} + \underbrace{\beta_1^{\text{true}} p_t^0}_{\text{AR(1)}} + \underbrace{\beta_2^{\text{true}} X_{2t}}_{\text{Additional price feature}} + \epsilon_t. \tag{5.24}$$

Neither coefficients β_0^{true}, β_1^{true} and β_2^{true}, nor the two price-relevant features from the available feature set $i = 1, ..., 10$ are known and both must be learned from data. Feature $i = 1$ is assumed to be the current spot price observation p_t^0, while $i = 2$ is an unspecified additional price-relevant feature (e.g., an analyst forecast respectively forecast adjustment or related commodity prices that have an impact on p_{t+1}^0). The noise term is supposed to be $\epsilon_t \overset{\text{i.i.d.}}{\sim} N(0, \sigma_{\epsilon_t}^2)$. Note that if $\sigma_{\epsilon_t}^2$ is small, learning is expected to be more effective. Furthermore, note that $\sigma_{\epsilon_t}^2$ accounts for the difference between the *full information problem* and the *perfect foresight problem*. For the *full information problem*, price model (5.24) is known but there is still uncertainty, i.e., $\sigma_{\epsilon_t}^2 > 0$. For the *perfect foresight problem* with deterministic future prices, ϵ_t is known to the firm.

The true but unknown parameters of price process ϕ over the time horizon $t = 1, ..., T'$ are assumed to be $(\beta_0^{\text{true}}, \beta_1^{\text{true}}) = (0, 1)$ and $(\beta_0^{\text{true}}, \beta_1^{\text{true}}) = (100, 0.5)$, respectively. The first setting describes a random walk (RW) price process $p_{t+1} = p_t + \epsilon_t$ and the second setting describes a mean reverting (MR) price process $p_{t+1} = \kappa\mu_p + (1 - \kappa)p_t + \epsilon_t$ with long-run mean price level of $\mu_p = 200$ and mean reversion speed $\kappa = 0.5$. As there is no consensus in the commodity finance literature whether commodity prices (e.g. of natural gas) follow a random walk or rather are mean-reverting (see, e.g., Andersson, 2007; Geman, 2007), DDA does not make any a-priori assumptions but learns the threshold price over time. The impact of the additional price-relevant feature $i = 2$ is determined by $\beta_2^{\text{true}} = 1$. In the following, we also investigate the loss by ignoring this additional price feature and deciding based on pure random walk or mean reversion assumptions, respectively, i.e., price-dependent price threshold $P_t(p_t^0) := \beta_0 + \beta_1 p_t^0$. This order-1-autoregressive approach is denoted as DDA-AR1 in the following. We furthermore distinguish between four different levels of noise in the data, i.e., $\sigma_{\epsilon_t} \in \{5, 10, 20, 30\}$. We generate price trajectories using the following procedure: Feature data X_{2t} is drawn from $N(0, 15^2)$ for all $t = 1, ..., T'$. Initial spot price is set to $p_1^0 = 200$. Consequently, p_{t+1}^0 can be generated according to equation (5.24) for all $t = 1, ..., T' - 1$. Price simulation paths that include negative prices (which occurred for RW in minor cases) are ignored. For exemplary price paths of RW and MR, see Figure B.2 of Appendix B.5. Over all simulations, the average monthly spot price change is of the order $6 - 16\%$, which is reasonable for the gas price time series (9%) that we analyze in Section 5.5.

(iii) Simulation of Feature Trajectories. Feature data relevant for price predictions (i.e., p_t^0 and X_{2t}) is generated according to the previously described procedure. All other feature data X_{it} $\forall i = 3, ..., 10$ without effect on prices is generated by sampling randomly from $N(10i, (2i)^2)$.

(iv) Simulation of Demand Trajectories. For the first part of our experiments, we assume constant demand, as this is the situation at our industry partner. For interpretation purposes, we set $d_t = 1$ for all $t = 1, ..., T'$. In the second part, which studies the *Value of Integrating Estimation and Optimization (VIEO)*, we additionally suppose a seasonal demand pattern modeled by the following process with mean demand level $\mu_d = 1$ and amplitude $\mu_d/2$:

$$d_t = \left[1 + \frac{1}{2}\sin\left(\frac{\pi\,(t-2)}{6}\right)\mu_d\right]. \tag{5.25}$$

(v) Procurement Policies. We compare the *Big Data* model (DDA-BD) that has access to all feature data ($i = 1, ..., 10$) with the featureless approach (DDA-SD), DDA-AR1 that only learns auto-regression in prices (i.e., feature p_t^0) but ignores further feature data $i = 2, ...10$, and regularization-based ML extensions (DDA-ML1, DDA-ML2). We furthermore consider the pure spot procurement policy (P-SPOT) and the pure forward procurement policy (P-M1). In the second part of the experiments, we additionally consider the separated estimation and optimization approach (SEO), which we describe in detail in the corresponding section.

(vi) Choice of Further Model Parameters. To ensure that the ML models (DDA-ML1, DDA-ML2) work properly, feature data X_{it} is standardized according to equation (5.20). We optimize the regularization parameter λ for the ML models (DDA-ML1, DDA-ML2) with an accuracy of 10^{-2} within a cross-validation procedure splitting the in-sample data equally into training and validation sets.

Table 5.2.: Summary of the numerical design

In-sample periods (training)	$t = 1, ..., T$ with $T \in \{24, 48, 72\}$
Out-of-sample periods (evaluation)	$t = T, ..., T'$ with $T' - T + 1 = 48$
General price process	Random walk (RW), mean reversion (MR)
Price process noise σ_{ϵ_t}	5, 10, 20, 30
Demand process	Constant, seasonal
Procurement policies	DDA-BD, -SD, -AR1, -ML1, -ML2, P-SPOT, P-M1, SEO
Performance measures	PE, VFI, VIEO

In the following, we answer research question (*Q3*). In particular, we want to exploit the drivers of *Prescription Error (PE)*, *Value of Feature Information (VFI)* and *Value of Integrating Estimation and Optimization (VIEO)* in commodity procurement.

5.4.2. Results

All results are reported across the 100 independent simulations that we run for each parameter configuration of Table 5.2. Boxplots report the minimum, 1st-, 2nd-, 3rd-quartile, maximum and the mean (×) value across the 100 simulations. Note that all results within a graph allow for fair comparisons as we evaluate on the same out-of-sample data (common random numbers).

Prescription Error (PE)

Figure 5.3 summarizes the out-of-sample PE of the different procurement policies for different parameter configurations in terms of training set size T, underlying price behavior (RW, MR) and level of price process noise σ_{ϵ_t}. It shows the rather intuitive result that PE of DDA-BD decreases with increasing training set size T (sample size effect) and increases with increasing noise in the price process σ_{ϵ_t}.

Regularization-based ML extensions (DDA-ML1, DDA-ML2) that decrease the in-sample performance yield less complex feature models and reduce overfitting, which increases the out-of-sample performance significantly (both average and worst case) except for the case with RW price behavior and price process noise of large size ($\sigma_{\epsilon_t} = 30$), in which the performance improvements through regularization are still positive but rather small. If T is sufficiently large and σ_{ϵ_t} is sufficiently small, regularization-based ML extensions generate less additional value relative to DDA-BD as there is less overfitting of the threshold function $P_t(\mathbf{X})$ generated by DDA-BD.

Figure 5.3 furthermore demonstrates that by ignoring additional feature information and deciding only based on AR(1) assumptions (DDA-AR1), purchase costs are significantly higher. The same is true for decisions based on a time-invariant price threshold (DDA-SD). However, there is a difference in the relative performance of DDA-SD with regard to the price process ϕ: If commodity prices are mean-reverting without trend, the performance loss of DDA-SD relative to DDA-BD, DDA-ML1 and DDA-ML2 is smaller than under RW assumptions where DDA-SD cannot capture the changing price levels (trends). If in addition price process noise σ_{ϵ_t} is large (Figure 5.3b2), then DDA-SD might even outperform DDA-BD, which misspecifies the price threshold $P_t(\mathbf{X})$ by overfitting. For the pure policies (P-SPOT and P-M1), the opposite is true: their performance loss relative to the other policies is higher under MR than under RW due to longer periods of upward and downward trends under the RW regime, for which P-M1 (for bull markets) or P-SPOT (for bear markets), respectively, is beneficial.

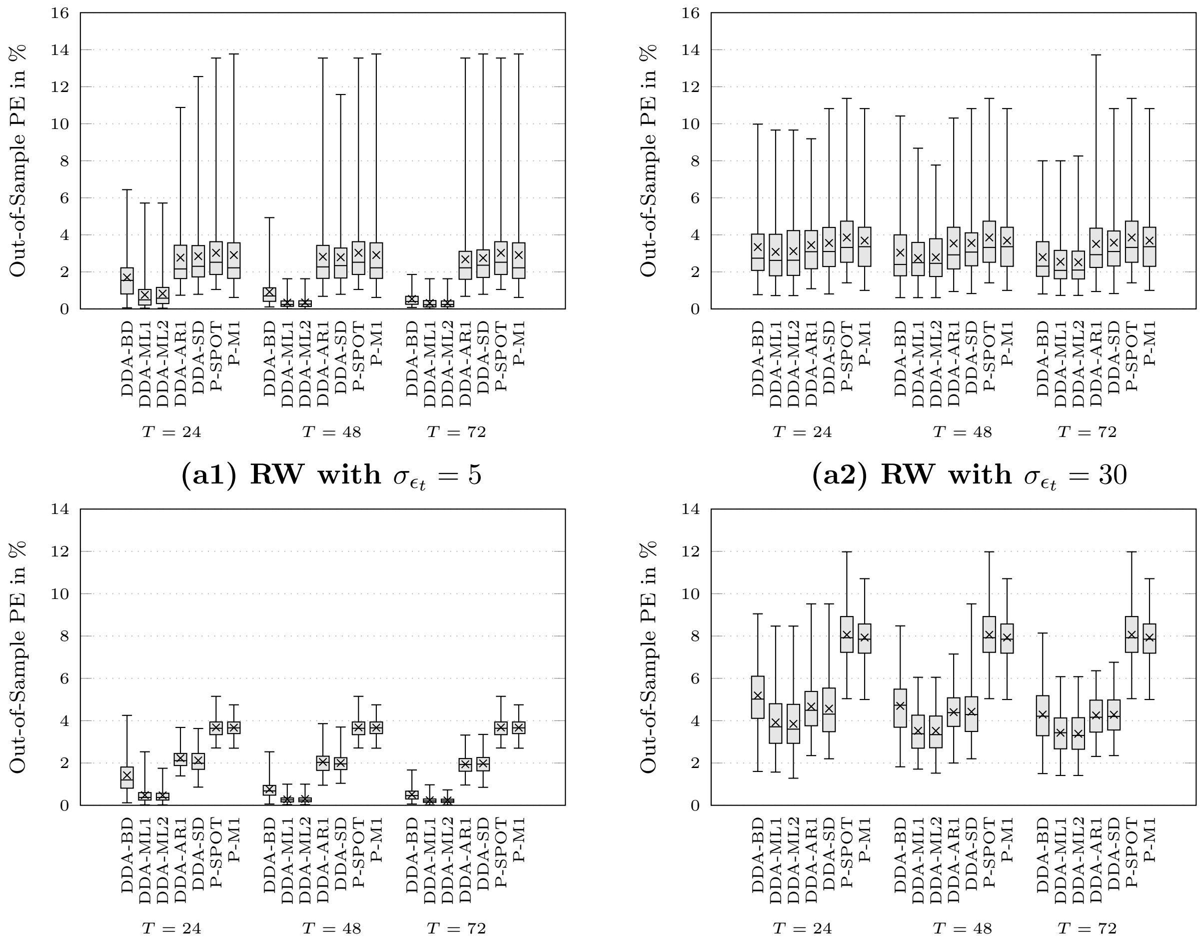

(a1) RW with $\sigma_{\epsilon_t} = 5$

(a2) RW with $\sigma_{\epsilon_t} = 30$

(b1) MR with $\sigma_{\epsilon_t} = 5$

(b2) MR with $\sigma_{\epsilon_t} = 30$

Figure 5.3.: Out-of-sample prescription error (PE) of different procurement policies conditional on training set size T, price process type (RW, MR) and price process noise levels σ_{ϵ_t} across 100 simulations (For the results of $\sigma_{\epsilon_t} \in \{10, 20\}$, see Figure B.3 in Appendix B.5)

In general, we observe that regularization-based DDA-ML1 and DDA-ML2 perform similarly well, and yield the best average out-of-sample procurement performance for all price processes, training set sizes and noise levels. If σ_{ϵ_t} is small, DDA-ML1 and DDA-ML2 come close to PF even for short training cycles. Across the 2,400 out-of-sample comparisons (2,400 = 2 price process types × 4 noise levels × 3 training set sizes × 100 simulation runs), DDA-ML1 (DDA-ML2) strictly outperforms DDA-BD in 76.08% (75.79%) of the settings, DDA-AR1 in 91.71% (91.63%), DDA-SD in 92.04% (92.08%), P-SPOT in 98.04% (97.88%) and P-M1 in 95.54% (95.38%) of the settings. The minimum dominance of DDA-ML1 (DDA-ML2) against DDA-BD is 57.00% (55.00%) for parameter configuration (RW, $T = 24$, $\sigma_{\epsilon_t} = 30$), against DDA-AR1 62.00% (62.00%)

for the same parameter configuration, against DDA-SD 64.00% (62.00%) for parameter configuration (MR, $T = 24$, $\sigma_{\epsilon_t} = 30$), against P-SPOT 85.00% (87.00%) and against P-M1 73.00% (74.00%) both for parameter configuration (RW, $T = 24$, $\sigma_{\epsilon_t} = 30$) - all instances of short training cycles and high level of noise.

Value of Feature Information (VFI)

Figure 5.4 reports the average out-of-sample results of the VFI for our simulation setup. We observe that VFI increases with increasing training sample size and decreases with increasing price process noise σ_{ϵ_t}.

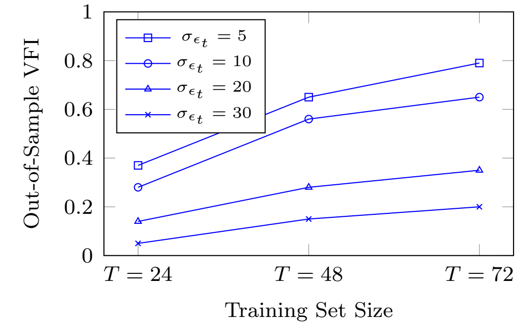

(a1) RW: DDA-BD

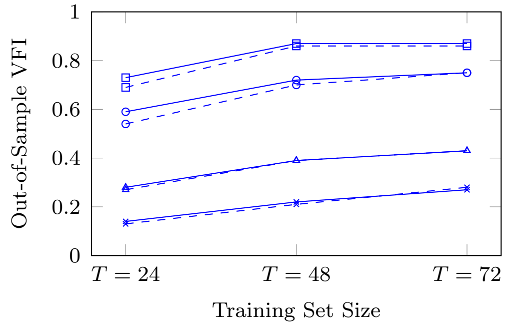

(a2) RW: DDA-ML1/-ML2

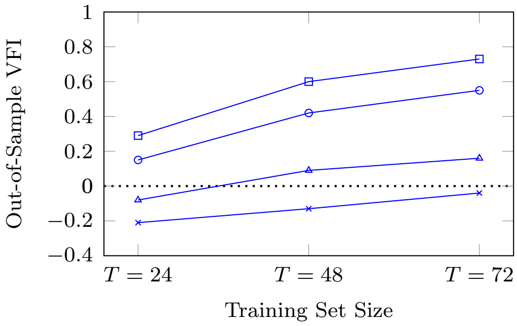

(b1) MR: DDA-BD

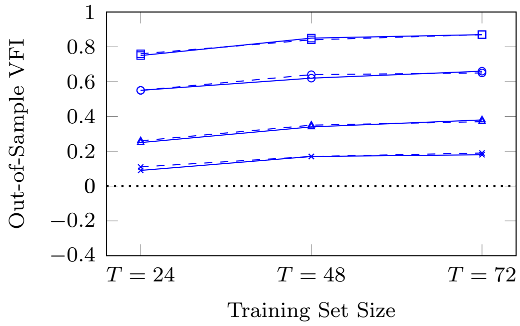

(b2) MR: DDA-ML1/-ML2

Figure 5.4.: Average out-of-sample value of feature information (VFI) of DDA-BD (left), DDA-ML1 (right, solid lines) and DDA-ML2 (right, dashed lines) for different training set sizes T, price process types (RW, MR) and price process noise levels σ_{ϵ_t} across 100 simulations (Full descriptive statistics of the results in Figure B.4 of Appendix B.5)

Regularization-based ML extensions that extract the relevant features from noise reduce overfitting and increase VFI significantly compared to DDA-BD for both price process types RW and MR (plots on the right hand-side of Figure 5.4 compared to plots on the left hand-side). Furthermore, we notice a concave shape, i.e., more additional value is generated by increasing the sample size from $T = 24$ to $T = 48$ rather than from $T = 48$ to $T = 72$. Especially for the ML approaches under small price process noise σ_{ϵ_t}, the additional value relative to the featureless approach generated from increasing the training set size from $T = 48$ to $T = 72$ is negligible.

We also observe for MR that the average out-of-sample VFI of DDA-BD is negative if σ_{ϵ_t} is sufficiently large and T is sufficiently small (Figure 5.4b1). In this case, overfitted DDA-BD generalizes worse than decisions based on constant feature-independent purchase signals (DDA-SD), i.e., feature data yields disadvantageous information. Consequently, ML extensions that extract decision-relevant features from noise are mandatory to generate value from covariate data (Figure 5.4b2).

Value of Integrating Estimation and Optimization (VIEO)

Even though we believe that the advantage of DDA compared to SEO comes, in particular, from reduced effort, the built-in cost-based backtesting opportunity (cost implications of predictions) and the easy to interpret and easy to operationalize decision rule framework, we still want to study whether there is a significant performance impact from integrating prediction and optimization.

Therefore, we compare the performance of DDA with the performance of SEO. To obtain price estimates for SEO, we apply OLS regression without (SEO) and with (SEO-AIC) forward-backward stepwise regression under the complexity criterion AIC. The OLS equivalent to DDA-BD is SEO, which does not use any complexity criterion for feature selection. Hence, both DDA-BD and SEO are heavily affected by overfitting, while DDA-ML and SEO-AIC aim at minimizing overfitting. Price models are estimated with the software `R` (`lm`-function for linear OLS models and `stepAIC`-function for stepwise forward-backward regression under the AIC model selection criterion). As we consider a setting with only one forward contract available, the price threshold $P_t^1(\mathbf{X})$ simply describes the price expectation $\mathbb{E}_t[p_{t+1}^0]$. This allows us to evaluate the cost-based out-of-sample performance of SEO by evaluating prediction-based $\beta_i \;\; \forall i = 0, ..., 10$ from OLS regression with the MILP of DDA-BD. For fair evaluation, we compare DDA-BD with SEO and DDA-ML with SEO-AIC based on the same 2,400 experiments (same 100 simulation runs as before for each of the parameter configurations with 2 price process types, 4 price process noise levels, 3 training set sizes). Therefore, we need to run 2,400 OLS regressions (SEO) and 2,400 OLS-AIC regressions (SEO-AIC) in total. Even though we are typically interested in out-of-sample implications, we also report in-sample results, which provides valuable additional information in this paragraph. Both the in-sample and the out-of-sample results across the 2,400 experiments are reported in Table 5.3 for different noise levels σ_{ϵ_t}.

Table 5.3.: In-sample and out-of-sample VIEO in % for DDA-BD vs. SEO and DDA-ML vs. SEO-AIC under constant demand

	DDA-BD vs. SEO							
	In-Sample				Out-of-Sample			
σ_{ϵ_t}	5	10	20	30	5	10	20	30
Mean	0.11	0.37	0.92	1.45	-0.73	-0.66	-0.47	-0.52
Min	0.00	0.00	0.00	0.00	-5.56	-10.36	-4.66	-5.02
25% Quantile	0.05	0.21	0.55	0.80	-0.99	-0.95	-0.88	-1.19
50% Quantile	0.09	0.33	0.82	1.35	-0.51	-0.46	-0.37	-0.35
75% Quantile	0.14	0.47	1.18	1.94	-0.22	-0.18	0.03	0.16
Max	0.59	2.08	9.22	4.82	1.72	1.32	3.19	4.52
DDA Dominance	100%	100%	100%	100%	4.8%	11.3%	27.2%	32.2%
	DDA-ML vs. SEO-AIC							
	In-Sample				Out-of-Sample			
σ_{ϵ_t}	5	10	20	30	5	10	20	30
Mean	0.06	0.16	0.28	0.35	-0.17	-0.12	0.10	0.24
Min	-0.12	-0.51	-1.20	-1.71	-5.13	-3.56	-2.46	-2.03
25% Quantile	0.02	0.05	0.02	-0.01	-0.24	-0.31	-0.33	-0.36
50% Quantile	0.05	0.13	0.23	0.30	-0.07	-0.07	0.08	0.10
75% Quantile	0.09	0.24	0.48	0.71	0.02	0.15	0.45	0.76
Max	0.53	1.34	5.87	2.69	1.72	2.85	4.55	5.33
DDA Dominance	95.7%	89.2%	78.5%	74.3%	31.5%	42.5%	57.0%	60.5%

Note. Due to similar performance of DDA-ML1 and DDA-ML2, we use DDA-ML1 for evaluation; results aggregated across price process types (RW,MR) and training set sizes T. While in-sample VIEO ≥ 0 for DDA-BD, the in-sample VIEO of DDA-ML can also be negative due to cost-complexity objectives rather than pure cost objectives.

DDA-BD versus SEO. The in-sample results demonstrate that fitting the model using OLS regression rather than fitting the model within the optimization problem yields on average $0.11 - 1.45\%$ higher cost (with a maximum of 9.22%). This percentage quantifies the value of a robust regression approach that leverages the optimization problem structure by minimizing the cost delta between DDA and PF, rather than minimizing squared residuals between price forecasts and price realizations. However, we observe that the in-sample benefits of DDA-BD over SEO are overlaid by the poor generalization of DDA-BD due to massive overfitting and alternate optimal solutions. Consequently, it is not surprising that, on average, SEO outperforms DDA-BD out-of-sample (VIEO < 0).

DDA-ML versus SEO-AIC. We already observed that regularization-based ML extensions (DDA-ML) increase out-of-sample generalization significantly. However, the comparison of DDA-ML with SEO would not be fair as DDA-ML uses feature selection while SEO does not. Therefore, we compare DDA-ML with SEO-AIC. DDA-ML benefits from leveraging the structure of the optimization problem and, at the same time, reduces the *generalization error* that deteriorates the performance of DDA-BD. We observe that DDA-ML outperforms SEO-AIC (in mean and median) when σ_{ϵ_t} is rather large ($\sigma_{\epsilon_t} \in \{20, 30\}$). In these settings SEO-AIC can yield a 5.33% higher cost. While with increasing σ_{ϵ_t}, the predictive power of both DDA-ML and SEO-AIC decreases, it seems that the performance of SEO-AIC is more strongly affected by this issue. Higher noise might cause more severe outliers which fosters robust regression-based DDA-ML compared to OLS-based SEO that gives more weight to outliers.

Effect of Demand Seasonality. Even though we do not address demand patterns in detail (as our industry partner faces constant gas demand over the year), demand seasonality might reinforce the gap between DDA and SEO, and hence the value of leveraging the optimization problem structure for price predictions. When the 2,400 in-sample experiments are repeated under the seasonal demand pattern modeled by equation (5.25), VIEO increases on average by 2.5% (for maximum instances by 27.3%). This is because DDA, which aims at minimizing costs, takes demand variations into account, while price predictions of SEO remain unaffected. Hence, SEO does not capture that the cost of price misspecification is larger in periods of high demand than in periods of low demand.

5.5. Results on Empirical Data

We perform backtests of our data-driven approaches for the multi-period procurement of natural gas at our industry partner and compare them to several benchmark policies. All models are solved to optimality. Without loss of generality, we set $d_t = 1$ for all periods since the gas demand of our industry partner is non-seasonal.

5.5.1. Setup

Natural gas is characterized by local markets, storage limitations and particularly high price volatility compared to other commodities. While for a long time, the European market has been characterized by long-term contracts, this fundamentally changed with its liberalization and the emergence of short-term financial hedging instruments (Geman, 2005, p. 237). For our case study, we use ten years of data from 07-2007 to 06-2017. In addition to the spot option, we consider monthly futures contracts of different time to maturity ($\tau \in \{1, 2, 3, 4\}$), i.e., M1-M4, that describe liquid financial products at the European TTF hub from which our industry partner purchases its gas.

Figure 5.5 shows the historical spot prices capturing various market trends, including the financial crisis. All price and feature data is retrieved from the databases Thomson Reuters Datastream and Eikon. Prices refer to closing prices on the last trading day of the month. Based on the results of empirical studies (see Section 3.1) and exten-

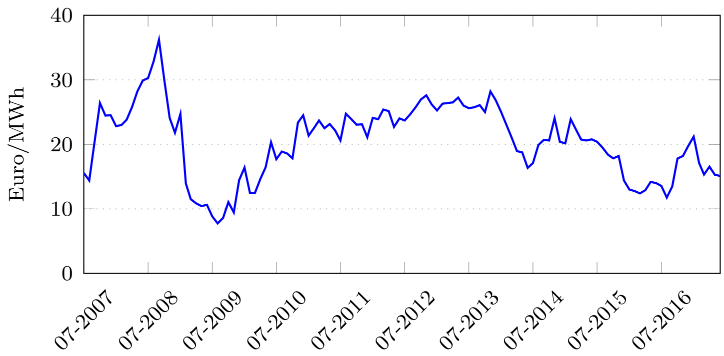

Figure 5.5.: TTF natural gas spot price evolution from 07-2007 until 06-2017

sive discussions with our industry partner's experienced gas purchase team, we consider the following feature categories: price history (e.g., spot price returns and lagged spot prices), related commodity prices (e.g., coal, Brent oil, Henry Hub gas), gas production (e.g., production volumes in Germany), gas consumption (e.g., gas consumption in Germany), economic indicators (e.g., EUR/USD, Bloomberg Commodity Index, Producer Price Index Energy) and environmental factors (e.g., temperature). In total, $N = 20$ features are pre-selected and standardized according to equation (5.20). For a detailed list of data sources, see Table B.1 of Appendix B.6. Similar to Secomandi et al. (2015), we split the dataset into rolling sub-periods for training and evaluation. Training (in-sample) and evaluation (out-of-sample) consist of 12 months each. Consequently, we perform 9 in-sample optimizations (07/2007-06/2008, 07/2008-06/2009,..., 07/2015-06/2016) and 9 out-of-sample evaluations (07/2008-06/2009, 07/2009 - 06/2010, ..., 07/2016 - 06/2017). The regularization parameter λ is calibrated with an accuracy of 10^{-2} within a cross-validation procedure, with a 50/50 split of in-sample data into training and validation sets, starting with $\lambda = 0$ (DDA-BD) and stopping as soon as $\beta_i^\tau = 0 \;\; \forall i = 1, ..., N, \tau \in \mathcal{F}^+$ (i.e., with convergence to DDA-SD).

Benchmark Policies and Policy Evaluation

As reference values, we use lower and upper performance bounds of the *perfect foresight problem* (5.21)-(5.23). Remember that C^{PF} is the theoretical optimal cost from the ex-post optimal policy (PE = 0%) under perfect price predictions. C^{UB} is the worst case performance solving the problem (5.21)-(5.23) under maximization objective. C^{PF} and C^{UB} describe a bandwidth for the hedging potential. If the hedging potential HP $:= \left(\frac{C^{\text{UB}}-C^{\text{PF}}}{C^{\text{PF}}}\right) \cdot 100\%$ is large (low), hedging decisions might strongly (hardly) affect procurement performance.

Beside DDA-BD, DDA-SD and DDA-ML, we test DDA-AR1, which uses the spot price p_t^0 as a single explanatory variable. Remember that DDA-AR1 exploits an AR(1) process with spot price-dependent thresholds $P_t^\tau(p_t^0) = \beta_0^\tau + \beta_1^\tau p_t^0$. Depending on the estimated values for β_0^τ and β_1^τ, DDA-AR1 characterizes RW, MR or MO. We furthermore compare our data-driven models (DDA) with the following procurement policies:

(i) *P-SPOT Policy:* The pure spot policy (P-SPOT) with $y_t^0 = d_t$ purchases under the current market prices, i.e., in the absence of risk management with the resulting purchase cost $C^{\text{P}-\text{SPOT}} = \sum_{t=1}^{T} p_t^0\, d_t$.

(ii) *P-τ Policy:* The pure forward policy (P-τ) with $y_t^\tau = d_{t+\tau}$ hedges demand fully via forward contracts of a specific time to maturity τ and resulting cost $C^{\text{P}-\tau} = \sum_{t=1|t\leq\tau}^{T} p_t^0\, d_t + \sum_{t=1|t>\tau}^{T} p_{t-\tau}^\tau\, d_t$ (Note: The beginning of the horizon requires spot purchases. As this is applied to all policies, it still ensures a fair comparison).

(iii) *1/N Policy (Status Quo):* 1/N portfolios allocate demand equally between the available procurement options (including the spot market). Note that this was the best-practice policy of our industry partner prior to our collaboration.

(iv) *REO Policy:* The reoptimization approach (REO) uses the deterministic forward curve $\vec{F}_t$ as predictor for $t + \Delta t$ with $\Delta t = 1, ..., T - t$ and considers forward curve updates by periodic reoptimization. REO avoids *model error* but ignores the extrinsic value from the stochastic evolution of the forward curve. REO is related to the rolling intrinsic policy primarily applied to natural gas storage, i.e., buying under random future input price, storage and selling under random future output price (see *Chapter 6*). In this context, the forward curve $\vec{F}_t$ is used for spot price forecasts, i.e., $p_t^\tau = \mathbb{E}_t^{Q}[p_{t+\tau}^0]$, which is true under the risk-neutral probability measure $\mathcal{Q}$. Consequently, $p_t^\tau = \mathbb{E}_t^{\mathcal{Q}}[p_{t+\tau-1}^{\tau-1}] \quad \forall \tau \geq 2$. However, this is only of interest if storage opportunities exist. For financial hedging, the rolling intrinsic

policy under $\mathcal{Q}$ would lead to the conclusion that financial hedging decisions do not matter. Therefore, we modify the rolling intrinsic policy for our problem setting: REO supposes that the future is equal to the present (naïve forecasts), i.e., the forward curves of all future periods equal the current forward curve. In this case, $p_t^\tau = \mathbb{E}_t[p_{t+\Delta t}^\tau]$ for all $\tau \in \mathcal{F}$, with $\Delta t = 1, ..., T - t$.
Example. *If the forward curve in period* t *is given by* $\vec{F}_t = (p_t^0, p_t^1, p_t^2) = (5, 3, 4)$, *then hedging of* d_{t+2} *in period* t *at a price* $p_t^2 = 4$ *is postponed to period* $t + 1$ *since* $\mathbb{E}_t[p_{t+1}^1] = 3 < 4$.

Reoptimization has been shown to perform impressively in the inventory context of gas storage valuation (Lai et al., 2010; Nadarajah et al., 2015; Secomandi, 2010, 2015; Wu et al., 2012) and is part of commercial software (Lacima, 2018; MathWorks, 2018). Similar to DDA-AR1, REO is a special case of DDA-BD if the forward curve $\vec{F}_t = (p_t^\tau : \tau \geq 0)$ is used as feature and the price threshold is estimated as $P_t^{\tau'}(\mathbf{X}) = \min_{\tau \in \mathcal{F}|\tau < \tau'} \{p_t^\tau\}\ \forall \tau' \in \mathcal{F}^+$.

Algorithm 1 Reoptimization Approach (REO)

for $t = 1, ..., T$ **do**
  for $\tau' \in \mathcal{F}^+ | t + \tau' \leq T$ **do**
    if $d_{t+\tau'}$ *is not fully hedged yet* **then**
      if $p_t^{\tau'} \leq p_t^\tau$ *for all* τ *with* $\tau \in \mathcal{F}|\tau < \tau'$ *feasible to hedge* $d_{t+\tau'}$ *in* $t + 1, ..., t + \tau'$ **then**
        set $y_t^{\tau'} = d_{t+\tau'}$
      else
        set $y_t^{\tau'} = 0$
      end if
    else
      set $y_t^{\tau'} = 0$
    end if
  end for
  if d_t *is not fully hedged yet* **then**
    set $y_t^0 = d_t$
  else
    set $y_t^0 = 0$
  end if
end for

5.5.2. Results

Descriptive Statistics

Figure 5.6 illustrates the spot-futures price relations, i.e., the relationship between p_t^0 and $p_t^\tau, \tau > 0$ (upper plots) and the relationship between $p_{t+\tau}^0$ and $p_t^\tau, \tau > 0$ (lower plots).

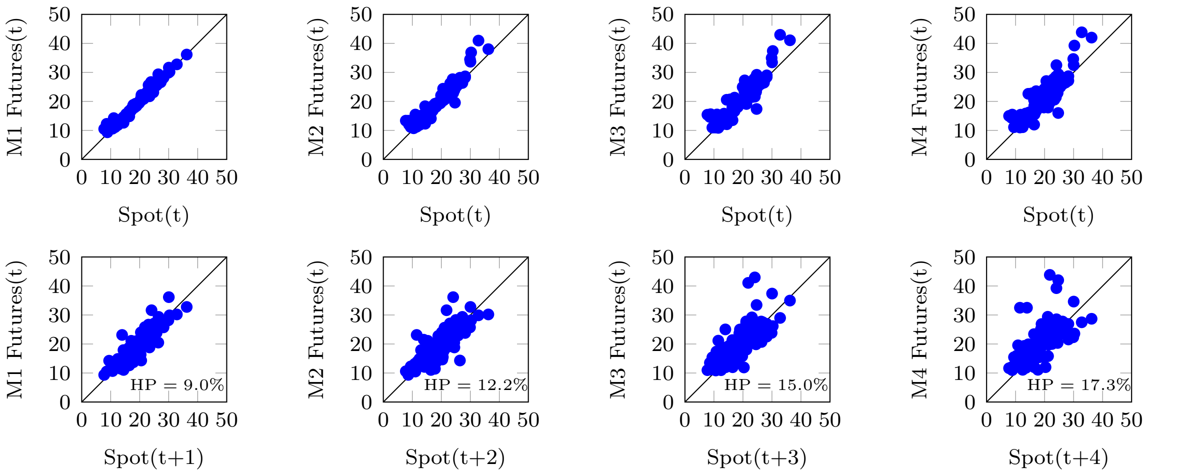

Figure 5.6.: Relation between monthly TTF spot and futures contract prices [Euro/MWh] from 07-2007 to 06-2017

In the upper plots, values above the diagonal indicate a normal futures curve (upward sloping, contango), while values below the diagonal indicate an inverted futures curve (downward sloping, backwardation). The plots below illustrate that, in some periods, it would have been better to buy at the spot, while in others hedging with futures is advantageous. The hedging potential HP shows that the larger the contract maturity, the higher the potential hedging gains/losses. Allowing the use of all hedging instruments simultaneously (Spot, M1, M2, M3, M4), then HP = 24.6% with an average purchase cost of $C^{\mathrm{PF}} = 18.86$ Euro/MWh (best case) and $C^{\mathrm{UB}} = 23.50$ Euro/MWh (worst case).

Performance of Procurement Policies

Table 5.4 presents the performance of the different procurement policies measured by the *Prescription Error* (PE), i.e., the cost deviation from the *perfect foresight problem* (PF). PE = 0.82% (DDA-AR1) and PE = 1.23% (DDA-SD) indicate that both AR(1) processes and constant purchase signals are not able to fully prescribe the problem not even in-sample. Even though there is a positive out-of-sample value of feature information VFI = 0.26, the performance of DDA-BD suffers from overfitting and poor generalization (on average 2.42 effective features per maturity excluding intercept). Decision-based feature selection by DDA-ML1 (1.64 features per maturity on average) and DDA-ML2 (1.72 features per maturity on average) adds regularization bias to the in-sample decision in order to reduce overfitting issues, which improves out-of-sample generalization significantly and increases VFI to 0.40 (DDA-ML1) and 0.38 (DDA-ML2), respectively.

Table 5.4.: Prescription error (PE): % above perfect foresight cost C^{PF}

	In-Sample				Out-of-Sample			
	Mean	StDev	Min	Max	Mean	StDev	Min	Max
DDA-BD	0.00	0.00	0.00	0.00	7.07	6.42	0.93	17.81
DDA-SD	1.23	1.26	0.20	3.59	9.52	5.47	3.47	19.92
DDA-AR1	0.82	1.12	0.00	3.58	9.42	5.25	3.07	17.94
DDA-ML1	0.57	1.19	0.00	3.59	5.68	4.34	0.93	13.09
DDA-ML2	0.59	1.18	0.00	3.59	5.93	4.23	0.93	13.09
REO	6.22	4.53	2.02	15.91	6.88	5.16	2.02	15.91
1/N	10.41	6.93	3.31	25.91	10.47	6.92	3.31	25.91
P-SPOT	6.86	7.02	0.93	21.68	6.14	6.40	0.93	21.68
P-M1	7.73	4.65	3.63	18.46	7.67	4.67	3.63	18.46
P-M2	10.36	7.14	3.02	27.15	10.69	7.05	3.02	27.15
P-M3	12.74	10.02	3.64	37.77	13.04	9.92	3.64	37.77
P-M4	14.59	13.53	2.69	48.78	14.56	13.59	2.69	48.78
UB	21.68	14.81	6.92	54.95	21.53	14.80	6.92	54.95

The most relevant features for DDA-ML1 over all backtests are spot price returns, the NYMEX front-month gas contract price and the producer price index (PPI) of energy in Germany. For two sub-periods (i.e., 07-2008 to 06-2009 and 07-2015 to 06-2016), DDA-ML1 yields, equivalently to DDA-SD, intercept-only purchase signals.

Furthermore, note that based on our analyzed dataset, DDA-ML (neither ML1 nor ML2) is never dominated by DDA-BD or DDA-SD in any of the out-of-sample sub-periods due to underfitting (DDA-SD) and overfitting (DDA-BD). Note that the pure spot policy (P-SPOT) is comparatively strong due to abrupt price declines, especially in 2008 (financial crisis), that cannot fully and immediately be captured by the other policies. More detailed results on cost and policy dominance are given in Table B.2 and B.3 of Appendix B.6.

Table 5.5.: Average annual savings (07-2008 to 06-2017) by DDA-ML1 compared to various benchmark policies for the real-world setting at our industry partner with annual gas demand of 10×10^6 MWh

	DDA-BD	DDA-SD	DDA-AR1	DDA-ML2	REO	1/N
in Mio Euro	2.634	7.295	7.099	0.400	2.278	9.087
in %	1.30	3.51	3.42	0.20	1.12	4.33
	P-SPOT	P-M1	P-M2	P-M3	P-M4	UB
in Mio Euro	0.880	3.780	9.514	13.981	16.852	30.084
in %	0.44	1.85	4.53	6.52	7.75	13.04

Supposing an annual demand of 10×10^6 MWh, as in the case of our industry partner, the data-driven approach with feature selection (DDA-ML1) yields significant savings in procurement cost compared to the firm's current policy 1/N (9.087 Mio Euro p.a., 4.33%) as well as compared to other benchmark policies (Table 5.5).

5.6. Conclusion

This chapter arose from collaboration with a chemical company that asked for analytics support in the field of dynamic natural gas procurement with both spot and forward purchase options. Optimizing the firm's position in the forward contract market describes a challenging multi-stage stochastic optimization problem under price and price model uncertainty. In addition, covariates related to prices, such as economic indicators, are available to the firm and might be worth exploiting for optimal decision-making.

In this chapter, we demonstrate how optimization in combination with machine learning can be used to compute the parameters of the optimal procurement policy. We propose a novel prescriptive analytics approach that avoids *price model error*, and is able to control *generalization error* by means of performance-based regularization. It incorporates prediction and optimization rather than executing these tasks sequentially. Uncertainty in future spot and forward prices is not expressed analytically but by historical (covariate) data. Our approach does not require *(i)* the a-priori specification and estimation of the underlying price processes, *(ii)* an explicit analysis of correlation structures or *(iii)* an a-priori selection of relevant price features. The associated decision rules are easy to interpret and easy to operationalize. We show that, in a *Big Data* environment, the firm needs to investigate carefully which data (*Smart Data*) to use, and that the performance of data-driven approaches depends significantly on model specification. Thus, we combine data-driven optimization with regularization methods from machine learning for a built-in selection of decision-relevant (rather than prediction-relevant) features, which increases out-of-sample generalization significantly.

The data-driven approach developed has already been implemented and is in use at the gas purchasing department of our industry partner (see Figure 5.1). Nevertheless, several extensions to this approach could be considered. Future research could incorporate the opportunity of keeping inventory and random demand, which may lead to partial hedging (see Froot et al., 1993). An extension to a merchant's setting with reselling opportunities might also be interesting. Furthermore, it would be worth exploring different types of sparsity penalties for the proposed ML models. Alternative regression techniques, such as kernel regression, or deep learning (non-linear kernels), can be an extension, although they might deteriorate interpretability.

Chapter 6.

Commodity Storage from a Data-Driven Merchant Perspective

Based on
Mandl, C., S. Nadarajah, S. Minner and S. Gavirneni (2019). Structured data-driven operating policies for commodity storage.

Storage assets are critical for the temporal trading of commodities under volatile prices. State-of-the-art methods for managing storage, such as the reoptimization-based rolling intrinsic approach (RIA), approximate a stochastic dynamic program (SDP) assuming full information regarding the state and the stochastic commodity price process and hence suffer from price modeling and forecasting errors, respectively, which are instances of *generalization error*. Based on extensive backtests, we find that this kind of error can lead to significantly suboptimal RIA policies, contrary to their known near-optimality in the full information setting. We develop two classes of non-parametric data-driven approaches (DDAs) that leverage machine learning and mathematical programming to overcome *generalization error*. The first class trains parameters of bang-bang and base-stock type policies, respectively, by solving linear and mixed integer programs, thereby extending known DDAs that parameterize decisions as functions of features without enforcing a policy structure. The second class trains value function parameters and encodes a base-stock policy structure via a linear program that can be solved to compute decisions. We backtest their performance on six major exchange-traded commodities from 2000 until 2017 with features constructed using Thomson Reuters and Bloomberg data streams. We find that DDA can improve the profits obtained by RIA significantly, with a base-stock policy structure needed to realize this improvement. Overall, our research advances the state-of-the-art for storage operations and suggests ways of restructuring commercial storage software to handle *generalization error*.

6.1. Introduction

Inter-temporal storage plays a fundamental role in the commodity industry. Examples for storage facilities are gas storage caverns, grain silos or metal warehouses. E.g., in 2017, more than 600 warehouses across 14 countries were approved by the London Metal Exchange (LME), one of the major marketplaces for physical metal trading. These warehouses are typically owned by trading companies and serve as a link between commodity producers and commodity processors. In 2015, 2.23 million tons of metal were delivered to and 3.77 million tons were delivered from LME warehouses (LME, 2017). But also other market places, such as the Chicago Merchantile Exchange (CME), encourage warehouse companies to expand their storage network (Reuters, 2016a). At commodity warehouses, storage volumes tie up a significant amount of money. For example, CME-registered warehouses in Salt Lake City held $131,774$ tons of copper in January 2018 (Reuters, 2018). As evaluated by the average COMEX copper price for January 2018 (3.19 USD/lb), this translates into an inventory value of 927 million USD.

Besides its operational use as a buffer against unexpected changes in supply and demand, storage also has commercial use by taking advantage of positive price differentials over time (Williams and Wright, 1991, p.24). Even though most academic papers assume that markets are efficient (frictionless), this is not true in practice due to, e.g., transaction costs and information asymmetry. Furthermore, as there is no equilibrium in inventory holding costs, some warehouses might provide cheaper storage than others. For instance, during the last couple of years, merchants increasingly shifted commodities from LME warehouses to cheaper private non-LME warehouses, whose storage rates are often around a tenth of the LME rates (Wall Street Journal, 2015) or to floating storage, i.e., offshore storage at vessels (Financial Times, 2015; Reuters, 2009).

In this chapter, we consider the well-known *Stochastic Commodity Warehouse Problem* (SCWP), sometimes simply called *Warehouse Management Problem* or *Warehousing Problem* that studies the optimal operating and physical trading policy over a finite horizon under the objective to *buy low, store, and sell high.* It optimizes the timing and quantity of purchase and sale at a storage facility with initial inventory, stochastic variability in prices, finite warehouse space and potential injection and withdrawal rate constraints. The optimization of injection and withdrawals thus critically depends on the movement of commodity prices in the future, which is difficult to predict.

Our analysis of commodity price data over 18 years (2000-2017) for instance shows a MAPE of 14-31% to predict the corresponding 6-months-ahead spot price by using the futures contract price (which is the expected spot price under the risk-neutral measure) such that even naïve (no-change) forecasts can perform better (see also Alquist and Kilian, 2010).

Figure 6.1 exemplarily illustrates the gap between futures contract prices and realized spot prices for natural gas at the New York Merchantile Exchange (NYMEX). It shows that futures prices might significantly fail in predicting future spot prices (especially for more distant spot prices).

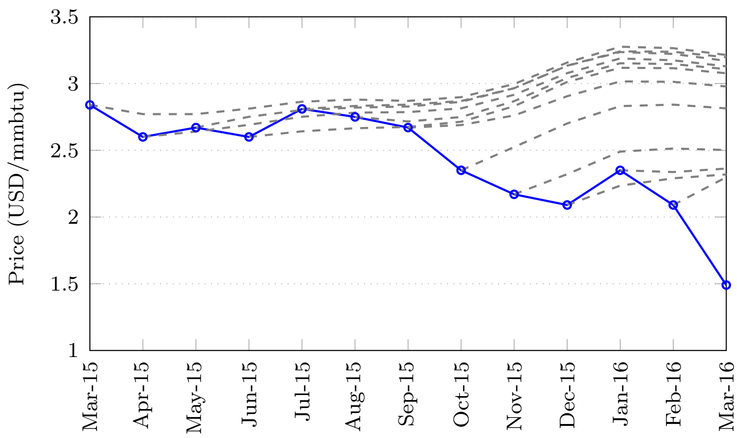

Figure 6.1.: NYMEX futures curves (dashed) and realized spot prices (∘) for natural gas

Optimal storage decisions for the SCWP are naturally obtained by solving a stochastic dynamic program (SDP) with the state described by on-hand inventory (endogenous state variable) and a price model to describe the evolution of the spot price (exogenous state variable). The latter is typically high-dimensional when the number of features that drive the commodity price is large. This makes the SDP intractable to solve and asks for alternatives. Two fundamental strategies exist: distribution-forecast based approaches and point-forecast based approaches. The first is related to approximate dynamic programming (ADP) techniques that approximate the SDP but are still based on distribution forecasts and a calibration of a stochastic process using historical data. The second strategy relies on replacing the uncertainty by a point forecast obtained through market information (e.g., from futures contract markets; see Figure 6.1).

For storage operations, popular examples of distribution-forecast based and point-forecast based approaches, respectively, are the least square Monte Carlo approach (LSMCA) and the rolling intrinsic approach (RIA), both of which are part of commercial storage management software (Energy Quants, 2018; Kyos, 2018; Lacima, 2018; MathWorks, 2018) describing state-of-the-art methods in natural gas storage practice

(Breslin et al., 2008, 2009; Gray and Khandelwal, 2004a,b). For natural gas storage valuation, reoptimization-based RIA has been shown to perform near-optimal in simulation settings under a calibrated price model ϕ (see, e.g., Lai et al., 2010; Nadarajah and Secomandi, 2018; Secomandi, 2010, 2015; Wu et al., 2012).

Doing so, the existing stochastic optimization literature aims at finding a storage policy π that minimizes the expected loss function

$$\min_{\pi}\left\{\mathbb{E}_{\mathbf{X}^{\mathrm{TE}}}\left[V^{\pi^*}(\mathbf{X}^{\mathrm{TE}}) - V^{\pi}(\mathbf{X}^{\mathrm{TE}})|\mathbf{X}^{\mathrm{TR}}\right]\right\} \tag{6.1}$$

on test data $\mathbf{X}^{\mathrm{TE}}$ given training data $\mathbf{X}^{\mathrm{TR}}$, with V^{π^*} being the expected profit resulting from the *full information optimum* π^*. The expectation about $\mathbf{X}^{\mathrm{TE}}$ is taken with respect to the calibrated price model ϕ (*full information problem*).

However, to ascertain the true performance of RIA, one needs to consider the inherent *generalization error* from price model and forecast error, which can be significant when employing RIA in practice. While the machine learning literature considers *generalization error*, the stochastic optimization literature at most addresses *price model error* due to price model misspecification (Secomandi et al., 2015). Even though zero price model error yields the *full information optimum*, only a sufficiently small *generalization error* yields the *perfect foresight optimum*. Hence, under explicit consideration of the *generalization error*, the truly desired policy π minimizes the expected loss function

$$\min_{\pi}\left\{\mathbb{E}_{\mathbf{X}^{\mathrm{TE}}}\left[V^{\mathrm{PF}}(\mathbf{X}^{\mathrm{TE}}) - V^{\pi}(\mathbf{X}^{\mathrm{TE}})|\mathbf{X}^{\mathrm{TR}}\right]\right\}, \tag{6.2}$$

with V^{PF} being the theoretically achievable oracle profit from storage under perfect foresight. Considering *generalization error*, the performance upper bound is not the performance of the *ex ante* optimal policy in the sense of the SDP under the assumption of full information about the price model (*full information problem*), but the theoretical *ex post* optimal performance under perfect foresight (*perfect foresight problem*). While under loss function (6.1), RIA is near-optimal, it is not known whether it is also near-optimal under loss function (6.2), i.e., once *generalization error* is accounted for.

Contribution

In this paper, we perform extensive backtesting on publicly available commodity price data. We show that RIA policies based on forward curves can be highly suboptimal, contrary to their known near-optimality in the full information setting. We observe that some main insights about the performance of RIA change fundamentally when evaluating its performance on real data affected by *generalization error* (which is addressed by the ML community), rather than under the full information paradigm of the stochastic optimization literature:

- While RIA has been shown to capture on average more than 99% of the value of the *full information optimum* (Lai et al., 2010; Secomandi, 2010), we find that due to *generalization error*, RIA yields an average value of 11.0% of the *perfect foresight optimum*.
- While RIA generates a positive (expected) value for the *full information problem*, we find that under *generalization error*, RIA yields negative profits for 30.1% of our backtest instances.
- While the value of reoptimization based on forward curve updates is positive for the *full information problem* (Secomandi, 2015), we find that under *generalization error*, it is negative for 37.0% of our backtest instances.
- While ignoring far-ahead forward price information might not matter for optimal first-stage decisions (Cruise et al., 2019), ignoring information (e.g. myopic decisions) can even be beneficial under *generalization error*.

Motivated by the surprising performance results of RIA under *generalization error*, we propose learning-enabled data-driven approaches (hereafter, DDA) for the SCWP in data-rich contexts. There is evidence in empirical commodity finance that the spot price forecast ability of futures prices that do not include estimates of the risk premium might be poor and outperformed by naïve forecasts (Alquist and Kilian, 2010), by a combination of futures prices and analyst forecasts (Cortazar et al., 2018) or by backward-looking ARMA models (see, e.g., discussions in Williams and Wright 1991, chap. 7 and references therein such as Fama and French 1987). Furthermore, macroeconomic features can improve predictions (see, e.g., Heath, 2018) as there is no consensus on the amount of new information already captured by quoted market prices (Cortazar et al., 2018). Therefore, commodity merchants recently try to take advantage of ML techniques in order to recognize patterns across commodities to anticipate demand and supply and its impact on prices (Oliver Wyman, 2017). Hence, while for RIA, the state space of the inventory system is fully characterized by the inventory level and the currently quoted forward curve, DDA learns the optimal state space representation directly from the data, which might include the aforementioned additional features, such as analyst forecasts, that are increasingly available in real-time having access to databases like Thomson Reuters Datastream, Bloomberg or Quandl.

For this purpose, we adapt the standard linear decision rule approach (DDA-LDR) from the literature (see, e.g., Ban and Rudin, 2019) that trains decisions as linear functions of features. However, this is not necessarily consistent with the optimal inventory policy structure. Instead of decisions, our novel approaches train optimal inventory policies (DDA-OSP) and value functions (DDA-VFA), both of which exploit well-known structural properties with regard to the optimal policy of the SCWP. Consequently, based on a fundamental multi-stage decision problem, we are the first to exploit the value of adding structure to data-driven solutions, which is not required for single-period problems such as the well-studied data-driven newsvendor (e.g., Ban and Rudin, 2019). To reduce *generalization error*, we combine DDA with ML-inspired interpretable regularization methods (i.e., lasso regression) for performance-based model selection from a set of candidate feature models.

We evaluate DDA in backtests based on six major commodities (i.e., copper, gold, crude oil, natural gas, corn, soybean) by using three feature sets (spot prices, futures prices and Bloomberg analyst forecasts). We observe that DDA generates a median profit of 26.7% of the perfect foresight value and hence improves the RIA profits (median of 12.0%) significantly. Furthermore, there is a considerable value in being consistent with optimal policy structures in multi-stage data-driven optimization. Over all instances and commodities, the unstructured state-of-the-art linear decision rule approach yields a median profit of only 2.4% of the perfect foresight profit compared to a value of 26.7% obtained by structured DDA policies that outperform the unstructured approach in 77.8% out of 1,152 out-of-sample backtesting instances. Our major performance results are statistically significant at the 1% level.

Research Questions and Organization

We address the following research questions: *(Q1)* How does the state-of-the-art software solution RIA that has been shown to perform near-optimal relative to the *full information optimum* perform in backtesting settings on real data affected by *generalization error*? *(Q2)* How to effectively solve the SCWP in a data-driven and learning-enabled way that addresses the adverse effects of *generalization error*? *(Q3)* What is the out-of-sample value of data-driven policies for the fundamental SCWP relative to RIA and can the exploitation of known policy properties improve data-driven solutions?

Section 6.2 presents the formal model of the SCWP and summarizes structural properties of the optimal and myopic inventory trading policies. It furthermore presents RIA as the state-of-the-art solution approach in practice. Section 6.3 presents data-driven models based on decision rule and value function approximations. Section 6.4 illustrates the weakness of RIA policies in extensive backtests and compares the results of the DDA models for various commodities. Section 6.5 concludes.

6.2. Model Formulation

We build on the stochastic variant (SCWP) of the fundamental *Warehouse Problem* under full operational flexibility (SCWP-FF) as formulated by Charnes et al. (1966) and under limited operational flexibility (SCWP-LF) as studied by Secomandi (2010). We regard the SCWP from the warehouse operator's perspective (optimal storage operating decisions), rather than from the investor's perspective (storage valuation).

6.2.1. Problem Setting

We consider a single-item, multi-period, discrete-time, periodic-review inventory replenishment problem at a single commodity storage facility with a finite planning horizon n. Periods $t = 0, 1, ..., n$ equal decision stages and might correspond to hours, days, weeks or months. The state of the warehouse is described by I_t and denotes the commodity amount in storage at the beginning of period t. I_t is bounded by lower and upper inventory levels $\underline{I}_t$ and $\bar{I}_t$, i.e., $\underline{I}_t \leq I_t \leq \bar{I}_t$. Without loss of generality, we assume that $\underline{I}_t = 0$ and $\bar{I}_t = C$, with C being the warehouse capacity. The per unit of inventory and unit of time holding costs are denoted as $c_h \geq 0$. $y_t^{\mathbf{i}} \geq 0$ defines the periodic storage injection quantity and $y_t^{\mathbf{o}} \geq 0$ the periodic storage withdrawal quantity.

Definition 7 (Operational Flexibility)**.** *It is distinguished between fully flexible facilities (SCWP-FF) and facilities with limited flexibility (SCWP-LF). SCWP-FF allows for full injection and withdrawal within one inventory review period, while SCWP-LF is restricted by injection rate limit $G^{\boldsymbol{i}} \in (0, C)$ and/or withdrawal rate limit $G^{\boldsymbol{o}} \in (0, C)$, i.e.,*

$$y_t^{\boldsymbol{i}} \leq \min\left\{C - I_t + y_t^{\boldsymbol{o}}, G^{\boldsymbol{i}}\right\}, \qquad y_t^{\boldsymbol{o}} \leq \min\left\{I_t, G^{\boldsymbol{o}}\right\}.$$

Limited flexibility ($G^{\mathbf{i}}, G^{\mathbf{o}} < C$) can be due to technical, logistical or market restrictions. E.g., certain gas storage facilities require 300 days to be filled (Secomandi, 2010).

Following the standard literature (Secomandi, 2010; Wu et al., 2012), we assume that the merchant is risk-neutral, has access to the spot market only (physical rather than financial trading via futures contracts) and is a price taker, i.e., storage operations do not affect the spot price. This is reasonable if the warehouse capacity C is restricted to such an extent that buying and selling has no effect on the spot market equilibrium. Furthermore, as it is standard in the merchant operations literature (see Section 3.2.2), we do not use an explicit demand component due to physical trading in spot markets where prices determine the demand. The unit nominal commodity spot price in period t is denoted as p_t.

Definition 8 (Frictions)**.** *The friction-adjusted purchase and selling prices are given by*

$$p_t^{\boldsymbol{i}} = \frac{1}{\eta^{\boldsymbol{i}}}\, p_t + c^{\boldsymbol{i}}, \qquad\qquad p_t^{\boldsymbol{o}} = \eta^{\boldsymbol{o}}\, p_t - c^{\boldsymbol{o}},$$

with $c^{\boldsymbol{i}} \geq 0$ and $c^{\boldsymbol{o}} \geq 0$ as the marginal injection and withdrawal cost and $\eta^{\boldsymbol{i}} \in (0,1]$ and $\eta^{\boldsymbol{o}} \in (0,1]$ as the injection and withdrawal loss.

These can be operational frictions (e.g., gas loss during the injection process due to the use of gas for running injection pumps) or transaction costs that are linear in the market price. Secomandi (2010) argues that $(\eta^{\mathbf{i}}, \eta^{\mathbf{o}}, c^{\mathbf{i}}, c^{\mathbf{o}})$=(0.99,1,\$0.02/mmbtu,\$0.02/mmbtu) is reasonable for gas storage facilities. In the following, we distinguish between the friction case with $p_t^{\mathbf{i}} > p_t^{\mathbf{o}}$ and the frictionless case with $p_t = p_t^{\mathbf{i}} = p_t^{\mathbf{o}}$.

The commodity merchant has precise information about past and present prices but not about prices in the future. Future market prices p_{t+1} are random and follow an unknown price process $\phi(p_{t+1})$.

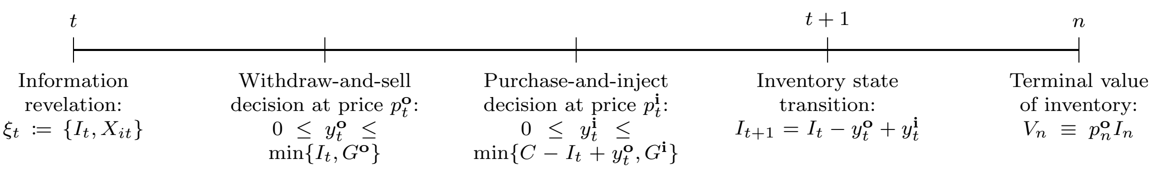

Figure 6.2.: Sequence of events

Let $\xi_t := \{I_t, X_{it}\}$ denote all information available to the merchant at the beginning of period $0 \leq t \leq n$ where $\xi_1 \subseteq \xi_2 \subseteq ... \subseteq \xi_n$. the inventory level I_t is endogenous information, $X_{it} \in \mathcal{X}$ is exogenous information, i.e., realizations of features $i = 1, ..., N$. This includes the current market price p_t, the current forward curve $\vec{F}_t$ and other features such as lagged prices or price forecasts.

In every period $0 \leq t \leq n$, the merchant decides about purchase-and-inject quantities $y_t^{\mathbf{i}}$ and withdraw-and-sell quantities $y_t^{\mathbf{o}}$ based on the following sequence: *(i)* Merchant information revelation ξ_t. *(ii)* Merchant's withdraw-and-sell decision $y_t^{\mathbf{o}} \geq 0$ at price $p_t^{\mathbf{o}}$ with $y_t^{\mathbf{o}} \leq \min\{I_t, G^{\mathbf{o}}\}$. *(iii)* Merchant's purchase-and-inject decision $y_t^{\mathbf{i}} \geq 0$ at price $p_t^{\mathbf{i}}$ with $y_t^{\mathbf{i}} \leq \min\{C - I_t + y_t^{\mathbf{o}}, G^{\mathbf{i}}\}$. Hence, the amount purchased in t is not sold in t (no-arbitrage assumption). Inventory state transitions follow $I_{t+1} = I_t - y_t^{\mathbf{o}} + y_t^{\mathbf{i}}$. The initial inventory I_0 is known to the merchant. Terminal inventory is valuated by the market price, i.e., $V_n(I_n, p_n) \equiv p_n^{\mathbf{o}} I_n$, which yields clearance of available inventory in the last period n as long as $p_n^{\mathbf{o}} > 0$. This assumption is reasonable if we assume that forward prices (and hence price forecasts) beyond n are not available or decreasingly liquid. Furthermore, we demonstrate in our numerical results from Section 6.4.2 that optimal first-stage decisions are only affected by a limited future horizon. Without loss of generality, we assume a discount factor of $\alpha = 1$. This simplification is uncritical as for all $t = 0, ..., n$, prices $p_t^{\mathbf{i}}$ and $p_t^{\mathbf{o}}$ can be interpreted as already discounted.

Accordingly, the objective of the SCWP is to maximize profit V_0 of the inventory system over the horizon of n future periods. The corresponding optimization problem is given by

$$\max_{y_t^{\mathbf{i}}, y_t^{\mathbf{o}}} \quad V_0 = \sum_{t=0}^{n} \mathbb{E}\Big[\big(p_t^{\mathbf{o}}\, y_t^{\mathbf{o}} - p_t^{\mathbf{i}}\, y_t^{\mathbf{i}} - c_h\, I_t\big) | X_{it}\Big] \tag{6.3}$$

$$\text{s.t.} \quad I_{t+1} = I_t - y_t^{\mathbf{o}} + y_t^{\mathbf{i}} \qquad \forall t = 0, 1, ..., n \tag{6.4}$$

$$0 \leq y_t^{\mathbf{i}} \leq \min\{C - I_t + y_t^{\mathbf{o}}, G^{\mathbf{i}}\} \qquad \forall t = 0, 1, ..., n \tag{6.5}$$

$$0 \leq y_t^{\mathbf{o}} \leq \min\{I_t, G^{\mathbf{o}}\} \qquad \forall t = 0, 1, ..., n. \tag{6.6}$$

Note that for capturing storage loss (e.g., in terms of agriculturals or electricity), storage efficiency $\rho \in (0, 1]$ as the fraction of stored commodity maintained over one period can be considered by modifying constraint (6.4), i.e., $I_{t+1} = \rho(I_t - y_t^{\mathbf{o}} + y_t^{\mathbf{i}})$. As this does not affect the structural results, we set $\rho = 1$ in the following. For model extensions with regard to positive fixed costs, storage inefficiency, explicit demand component and market power (price setter), we refer to Appendix C.1.

The corresponding dynamic programming formulation facilitates the analysis of the optimal policy structure. The problem can be formulated by the following Bellman equation:

$$V_t(z_t) = \max_{\substack{0 \le y_t^{\mathbf{i}} \le \min\{C - I_t + y_t^{\mathbf{o}}, G^{\mathbf{i}}\} \\ 0 \le y_t^{\mathbf{o}} \le \min\{I_t, G^{\mathbf{o}}\}}} \left\{ p_t^{\mathbf{o}}\; y_t^{\mathbf{o}} - p_t^{\mathbf{i}}\; y_t^{\mathbf{i}} - c_h\; I_t + \mathbb{E}_t \Big[V_{t+1}(z_{t+1}) \Big] \right\} \quad \forall\, 0 \le t < n \tag{6.7}$$

The state $z_t \in Z_t$ of the SDP includes both endogenous and exogenous information. Endogenous information describes the current operating conditions, i.e., I_t. Exogenous information represents the current market conditions X_{it}. For convenience in notation, we denote the feature vector as $\mathbf{X} = (X_{1t}, X_{2t}, ..., X_{Nt})$ in the following.

Within the existing merchant operations literature (see Section 3.2.2), the price model $\phi(p_{t+1})$ is assumed to be fully known in both the price process class and the process parameters (*full information problem*). The choice of the price model determines the state of the inventory model $z_t \in Z_t$. If the price follows a Markovian price process $\phi(p_{t+1}|p_t)$ (see, e.g., Secomandi, 2010), then the system state z_t is fully characterized by $z_t = (I_t, p_t)$, if the price follows a forward price model $\phi(p_{t+1}|\vec{F}_t)$ (see, e.g., Lai et al., 2010), then z_t is fully characterized by $z_t = (I_t, \vec{F}_t)$. Intractability of solving the SDP is mainly due to high-dimensional price models.

In contrast, if the underlying price process $\phi(p_{t+1})$ is not known, the state space $z_t \in Z_t$ cannot be specified a-priori. Hence, our distribution-free DDA models must identify the optimal state representation $z_t \in Z_t$ (which might include lagged prices or other features) directly from the data.

In the following, we distinguish between *(i)* SCWP-FF, i.e., fully flexible storage (Charnes et al., 1966) and *(ii)* SCWP-LF, i.e., storage with limited flexibility due to, e.g., injection and withdrawal constraints (Secomandi, 2010) or logistical respectively market constraints. SCWP-FF and SCWP-LF yield significantly different policy implications (Proposition 3-6). We make use of these propositions when formulating our data-driven optimization models in Section 6.3.

6.2.2. Optimal and Myopic Policy for Fully Flexible Storage

Proposition 3 (Optimal Policy under Full Operational Flexibility). *The optimal policy under full operational flexibility (FF) is of a bang-bang type. Depending on a state-dependent price threshold $P_t(\boldsymbol{X})$, for every stage $0 \le t < n$, the optimal inventory inflow and outflow decisions (y_t^{i}, y_t^{o}) only take either the maximum or the minimum. This splits the policy into three regions: (i) Do nothing, (ii) Fill the warehouse, or (iii) Empty the* *warehouse. The optimal inventory trading policy for SCWP-FF is:*

$$(y_t^{\boldsymbol{i}^*}, y_t^{\boldsymbol{o}^*}) = \begin{cases} (0,0) & \textit{if } p_t^{\boldsymbol{i}} \geq P_t(\boldsymbol{X}) \wedge p_t^{\boldsymbol{o}} < P_t(\boldsymbol{X}) \\ (C - I_t, 0) & \textit{if } p_t^{\boldsymbol{i}} < P_t(\boldsymbol{X}) \\ (0, I_t) & \textit{if } p_t^{\boldsymbol{i}} \geq P_t(\boldsymbol{X}) \wedge p_t^{\boldsymbol{o}} \geq P_t(\boldsymbol{X}) \end{cases}$$

The proof by backward induction follows Charnes et al. (1966) with the difference that the case

$$(y_t^{\mathbf{i}^*}, y_t^{\mathbf{o}^*}) = (C, I_t) \quad \textit{if } p_t^{\mathbf{i}} < P_t(\mathbf{X}) \wedge p_t^{\mathbf{o}} > p_t^{\mathbf{i}},$$

i.e., empty and refill the warehouse, is not reachable in our storage setting as by definition $p_t^{\mathbf{i}} \geq p_t^{\mathbf{o}}$ if $\eta^{\mathbf{i}}, \eta^{\mathbf{o}} \in (0,1]$ and $c^{\mathbf{i}}, c^{\mathbf{o}} \geq 0$. Hence, this case might only be relevant for commodity conversion settings with different input and output commodities.

Proposition 3 demonstrates that the optimal policy is of a state-dependent order-up-to-C-sell-down-to-0 type and can be fully characterized by a single price threshold $P_t(\mathbf{X})$. To decide about inventory trading, the merchant compares the current friction-adjusted spot prices $p_t^{\mathbf{i}}$ and $p_t^{\mathbf{o}}$ with the threshold value $P_t(\mathbf{X})$ regardless of the current inventory level I_t.

$P_t(\mathbf{X})$ can be interpreted as an inventory evaluator, i.e., the expected value of holding a unit of inventory at the beginning of period $t+1$. If we compute $P_t(\mathbf{X})$ via SDP, the recursive numerical evaluation relies on the assumption that we can fully characterize the underlying price process ϕ:

$$P_t(\mathbf{X}) = \mathbb{E}_\phi \left[\max \left\{ p_{t+1}^{\mathbf{o}}, \min[p_{t+1}^{\mathbf{i}}, P_{t+1}(p_{t+1}^{\mathbf{i}}, p_{t+1}^{\mathbf{o}}, \phi(p_{t+2}))] \right\} \right] - c_h \tag{6.8}$$

Therefore, we can derive lower and upper bounds on the inventory evaluator P_t, i.e., $P_t \geq \mathbb{E}_t[p_{t+1}^{\mathbf{o}}] - c_h$ and $P_t \leq \mathbb{E}_t[\max\{p_{t+1}^{\mathbf{o}}, P_{t+1}\}] - c_h$.

Example 1 (Recursive Solution)**.** *Suppose a deterministic 3-period SCWP-FF setting with initial inventory $I_t = 2$, capacity $C = 4$ and holding costs $c_h = 0$. Prices are $(p_t^{\boldsymbol{i}}, p_{t+1}^{\boldsymbol{i}}, p_{t+2}^{\boldsymbol{i}}) = (5.50, 6.00, 5.50)$ and $(p_t^{\boldsymbol{o}}, p_{t+1}^{\boldsymbol{o}}, p_{t+2}^{\boldsymbol{o}}) = (4.50, 4.00, 5.00)$. $P_t = \mathbb{E}_t[p_{t+1}^{\boldsymbol{o}}]$ and $P_t = \mathbb{E}_t[p_{t+1}^{\boldsymbol{i}}]$ yield the suboptimal decisions $(y_t^{\boldsymbol{i}}, y_t^{\boldsymbol{o}}) = (0,2)$ and $(y_t^{\boldsymbol{i}}, y_t^{\boldsymbol{o}}) = (2,0)$ respectively. The recursive character of P_t is required to evaluate the value of inventory at the beginning of period $t+2$, which yields the optimal decision $(y_t^{\boldsymbol{i}}, y_t^{\boldsymbol{o}}) = (0,0)$.*

Even though this simple numerical example already implies suboptimality of myopic decision-making, we want to close the gap in literature and formalize myopic policies as potential heuristics for the SCWP in the following.

Definition 9 (Myopic Policy Under Full Operational Flexibility)**.** *The myopic policy (MYOPIC) with a limited look-ahead of* $\Delta t = 1$ *is fully characterized by the price threshold* $P_t(\boldsymbol{X}) = \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h$. *This leads to the following myopic policy in closed-form:*

$$(y^{\boldsymbol{i}}_t, y^{\boldsymbol{o}}_t) = \begin{cases} (0,0) & \textit{if } p^{\boldsymbol{i}}_t \geq \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \wedge p^{\boldsymbol{o}}_t < \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \\ (C - I_t, 0) & \textit{if } p^{\boldsymbol{i}}_t < \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \\ (0, I_t) & \textit{if } p^{\boldsymbol{i}}_t \geq \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \wedge p^{\boldsymbol{o}}_t \geq \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \end{cases}$$

Other than for the optimal policy (OPT), MYOPIC is not characterized by $\mathbb{E}_t[p^{\mathbf{i}}_{t+1}]$ and $\mathbb{E}_t[P_{t+1}]$. If we assume that the marginal friction amounts $(\eta^{\mathbf{i}}, \eta^{\mathbf{o}}, c^{\mathbf{i}}, c^{\mathbf{o}})$ do not change over time, we can further simplify MYOPIC knowing that $p^{\mathbf{i}}_t = \frac{1}{\eta^{\mathbf{i}}} p_t + c^{\mathbf{i}}$, $\mathbb{E}_t[p^{\mathbf{o}}_{t+1}] = \eta^{\mathbf{o}}\ \mathbb{E}_t[p_{t+1}] - c^{\mathbf{o}}$ and $p^{\mathbf{o}}_t = \eta^{\mathbf{o}}\ p_t - c^{\mathbf{o}}$:

$$(y^{\mathbf{i}}_t, y^{\mathbf{o}}_t) = \begin{cases} (0,0) & \text{if } p_t \geq (\eta^{\mathbf{o}}\ \mathbb{E}_t[p_{t+1}] - c^{\mathbf{o}} - c^{\mathbf{i}} - c_h)\ \eta^{\mathbf{i}} \wedge p_t < \mathbb{E}_t[p_{t+1}] - \frac{c_h}{\eta^{\mathbf{o}}} \\ (C - I_t, 0) & \text{if } p_t < (\eta^{\mathbf{o}}\ \mathbb{E}_t[p_{t+1}] - c^{\mathbf{o}} - c^{\mathbf{i}} - c_h)\ \eta^{\mathbf{i}} \\ (0, I_t) & \text{if } p_t \geq (\eta^{\mathbf{o}}\ \mathbb{E}_t[p_{t+1}] - c^{\mathbf{o}} - c^{\mathbf{i}} - c_h)\ \eta^{\mathbf{i}} \wedge p_t \geq \mathbb{E}_t[p_{t+1}] - \frac{c_h}{\eta^{\mathbf{o}}} \end{cases}$$

Note that MYOPIC does not even require a precise one-step-ahead price prediction $\mathbb{E}_t[p_{t+1}]$ but allows for a certain prediction error $e_t > 0$ without affecting myopic inventory decisions. The allowable size of e_t depends on operational parameters such as injection and withdrawal costs $(c^{\mathbf{i}}, c^{\mathbf{o}})$, holding costs (c_h) and injection and withdrawal loss $(\eta^i, \eta^{\mathbf{o}})$:

$$(y^{\mathbf{i}}_t, y^{\mathbf{o}}_t) = \begin{cases} (0,0) & \text{if } p_t + \frac{c_h}{\eta^{\mathbf{o}}} < \mathbb{E}_t[p_{t+1}] \leq \frac{\frac{p_t}{\eta^{\mathbf{i}}} + c^{\mathbf{o}} + c^{\mathbf{i}} + c_h}{\eta^{\mathbf{o}}} \\ (C - I_t, 0) & \text{if } \frac{\frac{p_t}{\eta^{\mathbf{i}}} + c^{\mathbf{o}} + c^{\mathbf{i}} + c_h}{\eta^{\mathbf{o}}} < \mathbb{E}_t[p_{t+1}] \\ (0, I_t) & \text{if } \mathbb{E}_t[p_{t+1}] \leq \min\left\{p_t + \frac{c_h}{\eta^{\mathbf{o}}}, \frac{\frac{p_t}{\eta^{\mathbf{i}}} + c^{\mathbf{o}} + c^{\mathbf{i}} + c_h}{\eta^{\mathbf{o}}}\right\} \end{cases}$$

Proposition 4 (Optimality of MYOPIC for SCWP-FF without Frictions)**.** *For the frictionless case, MYOPIC is optimal and fully characterized in closed-form by* $P_t(\boldsymbol{X}) = \mathbb{E}_t[p_{t+1}] - c_h$, *i.e., price forecasting and inventory optimization can be fully decoupled, which results in the following trivial policy:*

$$(y^{\boldsymbol{i}^*}_t, y^{\boldsymbol{o}^*}_t) = \begin{cases} (C - I_t, 0) & \textit{if } p_t < \mathbb{E}_t[p_{t+1}] - c_h \\ (0, I_t) & \textit{if } p_t \geq \mathbb{E}_t[p_{t+1}] - c_h \end{cases}$$

This is because, for the frictionless case, $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 1$ and $c^{\mathbf{i}} = c^{\mathbf{o}} = 0$. Consequently, $p_t = p_t^{\mathbf{i}} = p_t^{\mathbf{o}}$ and $\mathbb{E}_t[p_{t+1}] = \mathbb{E}_t[p_{t+1}^{\mathbf{i}}] = \mathbb{E}_t[p_{t+1}^{\mathbf{o}}]$. Hence, as equation (6.8) implies that $P_t \geq \mathbb{E}_t[p_{t+1}^{\mathbf{o}}] - c_h$ and $P_t \leq \mathbb{E}_t[\max\{p_{t+1}^{\mathbf{o}}, p_{t+1}^{\mathbf{i}}\}] - c_h$, $P_t = \mathbb{E}_t[p_{t+1}] - c_h$ for the frictionless case. Consequently, the optimal policy from Proposition 3 reduces to the myopic policy.

Note that, under the risk-neutral probability measure $\mathcal{Q}$, the period t spot price expectation for future period τ equals the quoted forward price $f_{t,\tau}$. Consequently, for the frictionless case $P_t(\mathbf{X}) = f_{t,t+1} - c_h$. Hence, SDP, MYOPIC and RIA lead to the same storage decisions since stochastic changes in the forward curve $\vec{F}_t$ (which is considered by SDP but not for MYOPIC and RIA) are not relevant (Note that this is only the case for the frictionless SCWP-FF).

In Section 6.3, we want to compute the inventory evaluator $P_t(\mathbf{X})$ in a data-driven way by exploiting feature information $X_{it} \in \mathcal{X}$, rather than by SDP under full characterization of $\phi(p_{t+1})$ by, e.g., high-dimensional spot price models.

6.2.3. Optimal and Myopic Policy for Limited Flexible Storage

Proposition 5 (Optimal Policy Under Limited Operational Flexibility)**.** *Under limited operational flexibility (LF), for every stage $0 \leq t < n$, there exists a state-dependent purchase-up-to base-stock level $S_t^{\boldsymbol{i}}(\boldsymbol{X})$ and a state-dependent sell-down-to base-stock level $S_t^{\boldsymbol{o}}(\boldsymbol{X})$ with $S_t^{\boldsymbol{i}}(\boldsymbol{X}) \leq S_t^{\boldsymbol{o}}(\boldsymbol{X})$ that split the optimal policy into three regions: (i) Do nothing, (ii) Purchase-and-inject, or (iii) Withdraw-and-sell. The optimal inventory trading policy for SCWP-LF is:*

$$(y_t^{\boldsymbol{i}^*}, y_t^{\boldsymbol{o}^*}) = \begin{cases} (\min\{S_t^{\boldsymbol{i}}(\boldsymbol{X}) - I_t, G^{\boldsymbol{i}}\}, 0) & \textit{if } I_t < S_t^{\boldsymbol{i}}(\boldsymbol{X}) \\ (0, 0) & \textit{if } S_t^{\boldsymbol{i}}(\boldsymbol{X}) \leq I_t \leq S_t^{\boldsymbol{o}}(\boldsymbol{X}) \\ (0, \min\{I_t - S_t^{\boldsymbol{o}}(\boldsymbol{X}), G^{\boldsymbol{o}}\}) & \textit{if } I_t > S_t^{\boldsymbol{o}}(\boldsymbol{X}) \end{cases}$$

The proof follows Secomandi (2010).

Proposition 6 (Monotonicity of Base-Stock Levels)**.** *In every stage $0 \leq t < n$, the optimal base-stock levels $S_t^{\boldsymbol{i}}$ and $S_t^{\boldsymbol{o}}$ decrease in the spot price p_t (i.e., the optimal purchase amount $y_t^{\boldsymbol{i}}$ decreases and the optimal sales amount $y_t^{\boldsymbol{o}}$ increases in p_t) if the following two conditions hold for the underlying spot price process ϕ in every stage $0 \leq t < n$:*

(i) The distribution function of random variable p_{t+1} conditional on the spot price p_t stochastically increases in p_t, i.e., the one-period-ahead expected spot price increases in the current price,

(ii) The function $\mathbb{E}_t[p^{\boldsymbol{i}}_{t+1}|p_t] - p^{\boldsymbol{o}}_t$ decreases in p_t, i.e., the expected spot price in the next stage increases at a slower rate than the current spot price.

From Proposition 6, that has been proven by Secomandi (2010), follows that monotonicity of $S^{\mathbf{i}}_t$ and $S^{\mathbf{o}}_t$ strongly depends on the underlying spot price process. A RW process satisfies conditions *(i)* and *(ii)*, while a MR process only satisfies *(i)*. However, we do not rely on parametric price models in this chapter and hence do not determine monotonicity a-priori for DDA.

Figure 6.3 illustrates the optimal policy and monotonicity structure of SCWP-FF and SCWP-LF (under the assumptions of Proposition 6).

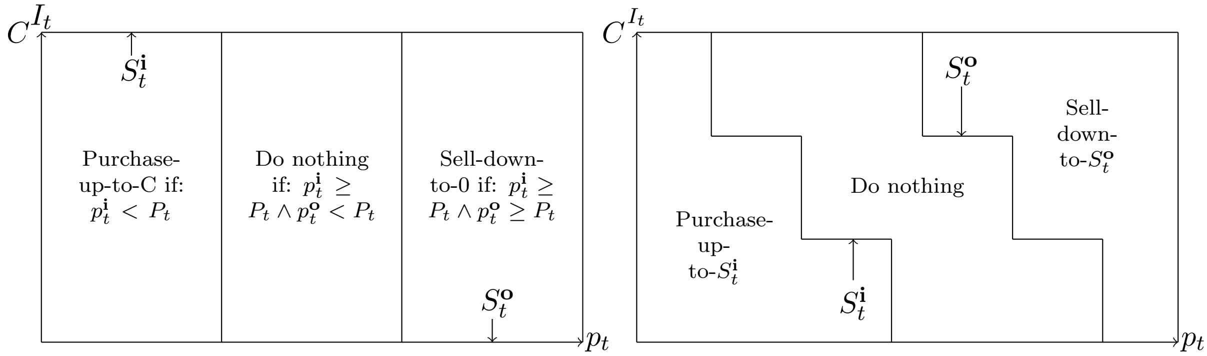

(a) SCWP-FF **(b) SCWP-LF**

Figure 6.3.: Policy and monotonicity illustration of SCWP-FF versus SCWP-LF

Definition 10 (Myopic Policy Under Limited Operational Flexibility)**.** *Equivalently to the fully flexible storage case, the myopic policy for the limited flexibility case is fully characterized by a single price threshold $P_t(\boldsymbol{X}) = \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h$, rather than by base-stock levels $S^{\boldsymbol{i}}_t(\boldsymbol{X})$ and $S^{\boldsymbol{o}}_t(\boldsymbol{X})$. This leads to the following myopic policy in closed-form, which is similar to the myopic policy of fully flexible warehouses from Definition 9:*

$$(y^{\boldsymbol{i}}_t, y^{\boldsymbol{o}}_t) = \begin{cases} (0,0) & \text{if } p^{\boldsymbol{i}}_t \geq \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \wedge p^{\boldsymbol{o}}_t < \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \\ (\min\{C - I_t, G^{\boldsymbol{i}}\}, 0) & \text{if } p^{\boldsymbol{i}}_t < \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \\ (0, \min\{I_t, G^{\boldsymbol{o}}\}) & \text{if } p^{\boldsymbol{i}}_t \geq \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \wedge p^{\boldsymbol{o}}_t \geq \mathbb{E}_t[p^{\boldsymbol{o}}_{t+1}] - c_h \end{cases}$$

Proposition 7 (Suboptimality of MYOPIC for SCWP-LF)**.** *Other than for the fully flexible storage case, for storage facilities with limited flexibility, MYOPIC is not optimal for the frictionless case ($\eta^{\boldsymbol{i}} = \eta^{\boldsymbol{o}} = 1$, $c^{\boldsymbol{i}} = c^{\boldsymbol{o}} = 0$) with $p_t = p^{\boldsymbol{i}}_t = p^{\boldsymbol{o}}_t$.*

This is because the myopic policy with a limited look-ahead of $\Delta t = 1$ is characterized by a price threshold $P_t(\mathbf{X})$, rather than by a double base-stock structure $(S_t^{\mathbf{i}}(\mathbf{X}), S_t^{\mathbf{o}}(\mathbf{X}))$, and leads to all-or-nothing decisions ignoring injection and withdrawal rate underutilization except for the trivial case where $y_t^{\mathbf{i}} = 0$ or $y_t^{\mathbf{o}} = 0$.

Example 2 (Capacity Underutilization). *Suppose the following deterministic three-period setting of the SCWP-LF without frictions: Spot prices are* $(p_t, p_{t+1}, p_{t+2}) = (2, 1, 3)$*, warehouse capacity is* $C = 10$*, storage cost* $c_h = 0$*, injection and withdrawal limits are* $G^{i} = 6$ *and* $G^{o} = 9$*, respectively. If* $I_t = 0$*, then* $y_t^{\boldsymbol{i,MYOPIC}} = 0$*, while* $y_t^{\boldsymbol{i,OPT}} = 3 < \min\{C - I_t, G^{i}\}$ *(injection rate underutilization). If* $I_t = 10$*, then* $y_t^{\boldsymbol{o,MYOPIC}} = 10$*, while* $y_t^{\boldsymbol{o,OPT}} = 7 < \min\{I_t, G^{o}\}$ *(withdrawal rate underutilization).*

In Section 6.3, the base-stock levels $S_t^{\mathbf{i}}$ and $S_t^{\mathbf{o}}$ of the optimal policy should be derived in a data-driven way by exploiting feature information $X_{it} \in \mathcal{X}$ of features $i = 1, ..., N$, rather than by SDP under full characterization of $\phi(p_{t+1})$ by, e.g., high-dimensional spot price models.

Table 6.1.: Summary of policy parameter characterization

		SCWP-FF	SCWP-LF
w/o frictions	Optimal	$P_t(\mathbf{X}) = \mathbb{E}_t[p_{t+1}]$	$S_t^{\mathbf{i}}(\mathbf{X}), S_t^{\mathbf{o}}(\mathbf{X})$
$(p_t^{\mathbf{i}} = p_t^{\mathbf{o}})$	Myopic	$P_t(\mathbf{X}) = \mathbb{E}_t[p_{t+1}]$	$P_t(\mathbf{X}) = \mathbb{E}_t[p_{t+1}]$
w/ frictions	Optimal	$P_t(\mathbf{X})$	$S_t^{\mathbf{i}}(\mathbf{X}), S_t^{\mathbf{o}}(\mathbf{X})$
$(p_t^{\mathbf{i}} > p_t^{\mathbf{o}})$	Myopic	$P_t(\mathbf{X}) = \mathbb{E}_t[p_{t+1}^{\mathbf{o}}]$	$P_t(\mathbf{X}) = \mathbb{E}_t[p_{t+1}^{\mathbf{o}}]$

Table 6.1 summarizes the policy results. The state-dependent threshold price $P_t(\mathbf{X})$ gives a signal regarding when to order, while state-dependent base-stock levels $S_t^{\mathbf{i}}(\mathbf{X})$ and $S_t^{\mathbf{o}}(\mathbf{X})$ give signals how much to order.

In our numerical analysis, we additionally test MYOPIC as a heuristic for solving the SCWP. Even though it is suboptimal (except for the case of fully flexible storage facilities without frictions), MYOPIC, if applied in a rolling horizon manner, might be a promising trading policy in the presence of *generalization error* since it only uses forecast information about period $t + 1$, which might be less affected by prediction errors than forecasts beyond $t + 1$.

6.2.4. Sequential Reoptimization: The Rolling Intrinsic Approach

Since optimal policy parameters are typically too complicated to compute via SDP (Eydeland and Wolyniec, 2003, pp. 351), we summarize the state-of-the-art reoptimization heuristic RIA. Being part of storage management software, such as Lacima (2018), Kyos (2018), Energy Quants (2018) or MathWorks (2018), RIA is widely used in industry practice for the valuation of gas storage assets (see, e.g., Breslin et al., 2008, 2009; Gray and Khandelwal, 2004a,b), where it has been shown to yield near-optimal results relative to the full information SDP. For instance, Secomandi (2010) reports that RIA is able to capture on average 99.81% of the value of the *full information optimum*.

RIA, which is sometimes referred to as forward dynamic optimization (Eydeland and Wolyniec, 2003, p. 355), is a sequential reoptimization heuristic and a type of certainty equivalent control (CEC) (Bertsekas, 1995) that makes decisions about $y^{\mathbf{i}}_t$ and $y^{\mathbf{o}}_t$ by optimizing a tractable LP based on period t point estimates for future spot prices, i.e., p_τ with $\tau > t$.

Stochastic future spot prices p_τ with $\tau > t$ are approximated, if available, by the (discounted) deterministic forward curve $\vec{F}_t = (f_{t,\tau} : \tau \in \mathcal{F} = \{t, t+1, ..., t+n\})$ with $f_{t,\tau}$ being the quoted forward price in t with contract maturity τ and n denoting the planning horizon. This is reasonable under rational expectation and the risk neutral measure $\mathcal{Q}$, i.e., futures prices equal expected spot prices with $f_{t,\tau} = \mathbb{E}^{\mathcal{Q}}_t[f_{s,\tau}], t < s \leq \tau$ (standard no-arbitrage assumption). Note that $f_{t,t}$ describes the spot price p_t. Equivalently to spot prices p_t, forward prices $f_{t,\tau}$ need to be adjusted to frictions, i.e., $f^{\mathbf{i}}_{t,\tau} = \frac{1}{\eta^{\mathbf{i}}}\, f_{t,\tau} + c^{\mathbf{i}}$ and $f^{\mathbf{o}}_{t,\tau} = \eta^{\mathbf{o}}\, f_{t,\tau} - c^{\mathbf{o}}$.

RIA:

$$\max_{y^{\mathbf{i}}_\tau, y^{\mathbf{o}}_\tau} \quad V^{\mathrm{I}}_t(I_t, \vec{F}_t) = \sum_{\tau \in \mathcal{F}} \left[f^{\mathbf{o}}_{t,\tau}\, y^{\mathbf{o}}_\tau - f^{\mathbf{i}}_{t,\tau}\, y^{\mathbf{i}}_\tau - c_h\, I_\tau \right] \tag{6.9}$$

$$\text{s.t.} \quad \underbrace{I_t + \sum_{\tau'=t}^{\tau} (y^{\mathbf{i}}_{\tau'} - y^{\mathbf{o}}_{\tau'})}_{I_{\tau+1}} = I_\tau - y^{\mathbf{o}}_\tau + y^{\mathbf{i}}_\tau \qquad \forall \tau \in \mathcal{F} = \{t, ..., t+n\} \tag{6.10}$$

$$0 \leq y^{\mathbf{i}}_\tau \leq C - I_\tau + y^{\mathbf{o}}_\tau \qquad \forall \tau \in \mathcal{F} \tag{6.11}$$

$$0 \leq y^{\mathbf{o}}_\tau \leq I_\tau \qquad \forall \tau \in \mathcal{F} \tag{6.12}$$

$$y^{\mathbf{i}}_\tau \leq G^{\mathbf{i}} \qquad \forall \tau \in \mathcal{F} \tag{6.13}$$

$$y^{\mathbf{o}}_\tau \leq G^{\mathbf{o}} \qquad \forall \tau \in \mathcal{F} \tag{6.14}$$

The LP ((6.9)-(6.14)) is solved in the current period t based on the available deterministic price information $\vec{F}_t = (f_{t,\tau} : \tau \in \mathcal{F}, \tau \geq t)$ and the inventory state information I_t maximizing the intrinsic value $V_t^I(I_t, \vec{F}_t)$ of storage regardless of the random evolution of forward curve $\vec{F}_t$ over time. Hence, $V_t^I(I_t, \vec{F}_t)$ comes from the seasonality of the forward curve (see Figure 6.1). However, other than for the intrinsic policy IA, for RIA only decisions of the current stage and state $y^{\mathbf{i}}_{\tau=t}$ and $y^{\mathbf{o}}_{\tau=t}$ are implemented and the intrinsic LP is resolved in $t+1$ for the updated inventory state information $I_{t+1} = I_t - y^{\mathbf{o}}_t + y^{\mathbf{i}}_t$ and the updated forward curve $\vec{F}_{t+1} = (f_{t+1,\tau} : \tau \in \mathcal{F}, \tau \geq t+1)$. This determines the rolling intrinsic value of storage.

RIA is consistent with the deterministic dynamic programming model and hence with the optimal policy structure. Also note that the myopic policy is a special case of RIA with $n = 1$, i.e., $\mathcal{F} = \{t, t+1\}$. A major benefit of RIA is its computational attractiveness. It avoids the curse of dimensionality of dynamic programming. However, thereby RIA disregards the stochastic evolution of the forward curve $\vec{F}_t$, i.e., expectations about $\vec{F}_{t+1}$ do not affect $y^{\mathbf{i}}_t$ and $y^{\mathbf{o}}_t$. Consequently, RIA ignores the extrinsic value of storage $V_t^E = V^*(I_t, \vec{F}_t) - V^I(I_t, \vec{F}_t)$, with V^* being the expected profit according to equation (6.7).

Moreover, we identify two further issues of RIA that are not addressed by the current literature in the SCWP context: *(i)* Are forward prices $f_{t,\tau}$ good estimates for spot prices p_τ? *(ii)* How to set the planning horizon n? I.e., is the consideration of all available forward price information $\vec{F}_t = (f_{t,\tau} : \tau \in \mathcal{F}, \tau \geq t)$ reasonable in the light of *generalization error*? Cruise et al. (2019) show that optimal first-stage decisions only depend on a short future time horizon n, which is determined by the parameters of the warehouse (e.g., frictions). However, they ignore *generalization error*, even though there is empirical evidence that prediction error increases with maturity τ (see, e.g., Cortazar et al. (2018) or our analysis in Section 6.4). Hence, large look-ahead horizons might not only be irrelevant (Cruise et al., 2019), but even disadvantageous.

Example 3 (RIA vs. MYOPIC)**.** *We consider fully flexible storage (SCWP-FF) with frictions $c^{\boldsymbol{i}} = c^{\boldsymbol{o}} = 0.9$ ($\eta^{\boldsymbol{i}} = \eta^{\boldsymbol{o}} = 1$). The warehouse with capacity $C = 1$ is empty ($I_t = 0$). Storage costs are neglected ($c_h = 0$). Available forward prices $f_{t,\tau}$ describe spot price expectations for both approaches RIA and MYOPIC, which is reasonable under the risk-neutral probability measure $\mathcal{Q}$. We assume that the market is in contango (i.e., upward sloping forward curves). Given $\vec{F}_t$ from Figure 6.4, MYOPIC yields $y^{\boldsymbol{i}}_t = y^{\boldsymbol{o}}_t = 0$, while RIA yields $y^{\boldsymbol{i}}_t = C$ for expected withdraw-and-sell in $t+5$. However, RIA is highly affected by prediction errors due to changing forward curves. If price realizations for*

$t+1, ..., t+5$ are smaller than 11.8 (= 10+0.9+0.9), RIA's period t decision is not even profitable. In contrast, MYOPIC is less affected by prediction errors since expectations about $t+2, ..., t+5$ do not affect period t decisions.

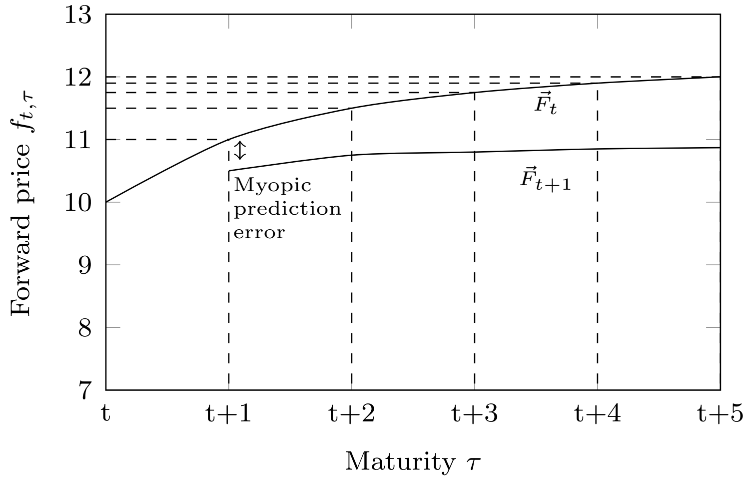

Figure 6.4.: RIA and MYOPIC with regard to prediction error

To overcome the limitations of RIA (i.e., *(i)* forward prices might not provide accurate spot price estimates and *(ii)* it is an open question how much forward price information to consider), we present several learning-enabled data-driven approaches in the following.

6.3. Data-Driven Optimization for the SCWP

In this section, we propose data-driven approaches (DDA) for solving the SCWP under the relaxation of *full information problem* assumptions, i.e., full characterization of a price model $\phi(p_{t+1})$ and therefore of the state space $z_t \in Z_t$. The DDA models are based on decision rule approximations (Section 6.3.1), policy rule approximations (Section 6.3.2) and value function approximations (Section 6.3.3). DDA-LDR trains decisions as functions of features, DDA-OSP trains policy parameters as functions of features and DDA-VFA trains the value-to-go function based on features.

6.3.1. Unstructured Linear Decision Rule Approach (DDA-LDR)

DDA-LDR trains decisions via decision rule approximations as applied in the data-driven and non-parametric stochastic optimization literature (see Section 3.3.2 and particularly Ban and Rudin (2019) for the newsvendor problem). Decision rules map observations of random variables to decisions such as the newsvendor order quantity. In the following, we focus on affine (intercept plus slope) linear decision rules (see Garstka and Wets (1974) and Georghiou et al. (2005) for comprehensive tutorials on decision rule approximations).

Definition 11 (Linear Decision Rule Approximation for the SCWP)**.** *The injection and withdrawal quantities $y_t^{\mathbf{i}}$ and $y_t^{\mathbf{o}}$ are approximated by affine linear decision rules (LDR) of the form*

$$y_t^{\boldsymbol{i}}(\boldsymbol{X}) := \sum_{i=0}^{N} \beta_i^{\boldsymbol{i}} X_{it}, \qquad y_t^{\boldsymbol{o}}(\boldsymbol{X}) := \sum_{i=0}^{N} \beta_i^{\boldsymbol{o}} X_{it}.$$

In this regression-like equation, $\beta_i^{\mathbf{i}} \in \mathcal{B} \subset \mathbb{R}$ and $\beta_i^{\mathbf{o}} \in \mathcal{B} \subset \mathbb{R}$ are feature coefficients that are unknown to the merchant and must be learned from historical time series data $t \in \mathcal{T} = \{1, ..., T\}$ of length T. To allow for a feature-independent intercept term, we set $X_{0t} = 1 \; \forall t \in \mathcal{T}$.

Even though $y_t^{\mathbf{i}}$ and $y_t^{\mathbf{o}}$ are linear in feature observations X_{it}, this is not restrictive. Non-linearities can be considered by interaction terms (e.g., $y_t^{\mathbf{i}}(\mathbf{X}) := \beta_0^{\mathbf{i}} + \beta_1^{\mathbf{i}} X_{1t} + \beta_2^{\mathbf{i}} X_{2t} + \beta_3^{\mathbf{i}} X_{1t} X_{2t}$) or by polynomial terms (e.g., $y_t^{\mathbf{i}}(\mathbf{X}) := \beta_0^{\mathbf{i}} + \beta_1^{\mathbf{i}} X_{1t}^2$). Lagged observations offer additional flexibility (e.g., $y_t^{\mathbf{i}}(\mathbf{X}) := \beta_0^{\mathbf{i}} + \beta_1^{\mathbf{i}} X_{1t} + \beta_2^{\mathbf{i}} X_{1,t-1}$).
The coefficients $\beta_i^{\mathbf{i}}, \beta_i^{\mathbf{o}} \in \mathcal{B}$ are determined based on the statistical learning theory principle of ERM (Vapnik, 1998, pp. 32) through $\min_{\beta_i^{\mathbf{i}}, \beta_i^{\mathbf{o}} \in \mathcal{B}} \frac{1}{T} \sum_{t=1}^{T} \ell_{\mathrm{DDA}}(\hat{V}_t, V_t^{\mathrm{PF}})$ with ℓ_{DDA} being the profit-based loss function of the data-driven approach relative to the theoretical profit V^{PF} of the *perfect foresight problem.* Notation $\hat{}$ emphasizes that the profit is estimated from data via the ERM principle, rather than under the *full information problem.* Minimizing the loss with respect to the nominal optimization problem is equivalent to setting $\beta_i^{\mathbf{i}}, \beta_i^{\mathbf{o}} \in \mathcal{B}$ of the LDR such that the average period feature-conditional (Big-Data-conditional) inventory trading profit is maximized over the training set $t = 1, ..., T$:

$$\max_{\beta_i^{\mathbf{i}}, \beta_i^{\mathbf{o}} \in \mathcal{B}} \left\{ \frac{1}{T} \sum_{t=1}^{T} \hat{V}_t(I_t, p_t) | \mathbf{X} \right\} \tag{6.15}$$

Figure 6.5 illustrates the framework of in-sample optimization (training) and out-of-sample evaluation of feature coefficients $\beta_i^{\mathbf{i}}, \beta_i^{\mathbf{o}} \in \mathcal{B}$.

Once the coefficients $\beta_i^{\mathbf{i}}, \beta_i^{\mathbf{o}} \in \mathcal{B}$ are trained via DDA-LDR, injection and withdrawal quantities for $t = T^{\mathrm{s}}, ..., T'$ are determined by $y_t^{\mathbf{i}} = \max\{\min\{C - I_t + y_t^{\mathbf{o}}, G^{\mathbf{i}}, \sum_{i=0}^{N} \beta_i^{\mathbf{i}} X_{it}\}, 0\}$ and $y_t^{\mathbf{o}} = \max\{\min\{I_t, G^{\mathbf{o}}, \sum_{i=0}^{N} \beta_i^{\mathbf{o}} X_{it}\}, 0\}$ respectively. However, the training is different to the standard ERM approach in the data-driven optimization literature that only

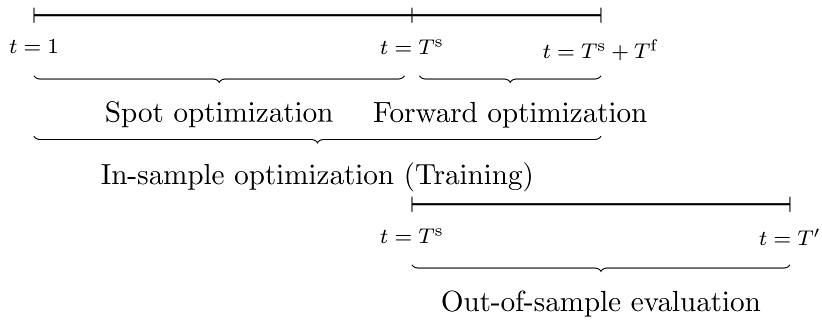

Figure 6.5.: Optimization (training) and evaluation framework

optimizes the in-sample profit/cost over historical periods (i.e., $t = 1, ..., T^{\mathrm{s}}$) and does not exploit available forward-looking information during the training. However, for commodities, estimates for future periods $t = T^{\mathrm{s}} + 1, ..., T^{\mathrm{s}} + T^{\mathrm{f}}$ are available at $t = T^{\mathrm{s}}$ in the form of the quoted forward curve $\vec{F}_t$ that might be worth considering for the training of the feature coefficients (forward optimization). Therefore, the ERM objective is extended to capture the current best-estimate of future profits by considering available forward prices. Injection and withdrawal decisions are trained via the following LP.

DDA-LDR:

$$\begin{aligned}\max_{\substack{y_t^{\mathbf{i}}(\mathbf{X})\\ y_t^{\mathbf{o}}(\mathbf{X})}} \hat{V} = & \frac{1}{T^{\mathrm{s}} + T^{\mathrm{f}}} \Bigg[\underbrace{\sum_{t=1}^{T^{\mathrm{s}}} (p_t^{\mathbf{o}} y_t^{\mathbf{o}} - p_t^{\mathbf{i}} y_t^{\mathbf{i}} - c_h I_t)}_{\text{Standard ERM}} \\ & + \underbrace{\sum_{t=1}^{T^{\mathrm{f}}} (f_{T^{\mathrm{s}}, T^{\mathrm{s}}+t}^{\mathbf{o}} \, y_{T^{\mathrm{s}}+t}^{\mathbf{o}} - f_{T^{\mathrm{s}}, T^{\mathrm{s}}+t}^{\mathbf{i}} \, y_{T^{\mathrm{s}}+t}^{\mathbf{i}} - c_h I_{T^{\mathrm{s}}+t})}_{\text{Forward optimization: Estimate of future profit}} \Bigg]\end{aligned} \tag{6.16}$$

$$\text{s.t.} \quad \underbrace{I_1 + \sum_{t'=1}^{t} (y_{t'}^{\mathbf{i}} - y_{t'}^{\mathbf{o}})}_{I_{t+1}} = I_t - y_t^{\mathbf{o}}(\mathbf{X}) + y_t^{\mathbf{i}}(\mathbf{X}) \qquad \forall t \in \mathcal{T} = \{1, ..., T^{\mathrm{s}} + T^{\mathrm{f}}\} \tag{6.17}$$

$$0 \leq y_t^{\mathbf{i}}(\mathbf{X}) \leq C - I_t + y_t^{\mathbf{o}}(\mathbf{X}) \qquad \forall t \in \mathcal{T} \tag{6.18}$$

$$0 \leq y_t^{\mathbf{o}}(\mathbf{X}) \leq I_t \qquad \forall t \in \mathcal{T} \tag{6.19}$$

$$y_t^{\mathbf{i}}(\mathbf{X}) \leq G^{\mathbf{i}} \qquad \forall t \in \mathcal{T} \tag{6.20}$$

$$y_t^{\mathbf{o}}(\mathbf{X}) \leq G^{\mathbf{o}} \qquad \forall t \in \mathcal{T} \tag{6.21}$$

Note that $\beta_i^{\mathbf{i}}, \beta_i^{\mathbf{o}} \in \mathcal{B} \subset \mathbb{R}$ are decision variables. If p_t is used as a feature, a price-dependent base-stock policy can be estimated by the feature models $y_t^{\mathbf{i}} = \beta^{\mathbf{i}} \cdot \frac{1}{p_t} - I_t$ and $y_t^{\mathbf{o}} = I_t - \beta^{\mathbf{o}} \cdot \frac{1}{p_t}$. However, $y_t^{\mathbf{i}} = (S_t^{\mathbf{i}} - I_t)^+$ and $y_t^{\mathbf{o}} = (I_t - S_t^{\mathbf{o}})^+$ are nonlinear equations, i.e., there is no foundation to assume that a base-stock policy is close to an LDR. Consequently, DDA-LDR is not necessarily consistent with the optimal policy structure (e.g., it is feasible to simultaneously purchase and sell). One can measure the suboptimality of LDRs by solving the dual problem (Georghiou et al., 2005). The dual variable of constraint (6.17) denotes the value of having an additional unit of inventory in $t+1$. Especially when decisions are constrained, LDRs might be highly inappropriate (Bertsimas and Kallus, 2016). While this is not an issue for newsvendor-like problems (see, e.g., Ban and Rudin, 2019), multi-period problems exhibit structural properties that might be valuable to exploit in data-driven optimization. Therefore, we extend DDA-LDR in the next section in order to guarantee consistency with the optimal policy structure from Section 6.2.

6.3.2. Optimally Structured Policy Approach (DDA-OSP)

In this section, we explicitly use the knowledge about structural properties for data-driven optimization. DDA-OSP trains policy parameters, rather than decisions, by mapping feature observations X_{it} to inventory control parameters. This approach ensures policy consistency by construction and is related to Iyer and Schrage (1992), who directly compute the (s, S) parameters that would have been optimal for the original deterministic demand stream. However, they focus on the inventory problem under uncertain demand rather than price and also do not consider feature-based learning.

Due to different policy characterizations (see Section 6.2), we distinguish between full operational flexibility (SCWP-FF) and limited operational flexibility (SCWP-LF) being aware that SCWP-FF is a special case of SCWP-LF.

DDA-OSP-FF for SCWP-FF

Definition 12 (Linear Policy Rule Approximation for SCWP-FF)**.** *The inventory evaluator $P_t(\boldsymbol{X})$ from Proposition 3 that gives purchase-and-inject and withdraw-and-sell signals is approximated by an affine policy rule of the form*

$$P_t(\boldsymbol{X}) \coloneqq \sum_{i=0}^{N} \beta_i X_{it}.$$

The following MILP models are consistent with the optimal policy structure in the sense of the full information SDP under full knowledge about the price process (Proposition 3). The goal is to train inventory evaluator $P_t(\mathbf{X})$ as a function of feature vector $\mathbf{X}$.

DDA-OSP-FF:

$$\begin{aligned}\max_{P_t(\mathbf{X})} \hat{V} =& \frac{1}{T^{\mathrm{s}}+T^{\mathrm{f}}}\Bigg[\sum_{t=1}^{T^{\mathrm{s}}}(p_t^{\mathbf{o}}y_t^{\mathbf{o}} - p_t^{\mathbf{i}}y_t^{\mathbf{i}} - c_h I_t) \\ &+ \sum_{t=1}^{T^{\mathrm{f}}}(f_{T^{\mathrm{s}},T^{\mathrm{s}}+t}^{\mathbf{o}}\, y_{T^{\mathrm{s}}+t}^{\mathbf{o}} - f_{T^{\mathrm{s}},T^{\mathrm{s}}+t}^{\mathbf{i}}\, y_{T^{\mathrm{s}}+t}^{\mathbf{i}} - c_h I_{T^{\mathrm{s}}+t})\Bigg]\end{aligned} \tag{6.22}$$

$$\begin{aligned}
\text{s.t.}\quad & I_1 + \sum_{t'=1}^{t}(y_{t'}^{\mathbf{i}} - y_{t'}^{\mathbf{o}}) = I_t - y_t^{\mathbf{o}} + y_t^{\mathbf{i}} && \forall t \in \mathcal{T} = \{1, ..., T^{\mathrm{s}}+T^{\mathrm{f}}\} && (6.23)\\
& 0 \le y_t^{\mathbf{i}} \le C - I_t + y_t^{\mathbf{o}} && \forall t \in \mathcal{T} && (6.24)\\
& 0 \le y_t^{\mathbf{o}} \le I_t && \forall t \in \mathcal{T} && (6.25)\\
& Mq_t \ge P_t(\mathbf{X}) - p_t^{\mathbf{i}} && \forall t = 1, ..., T^{\mathrm{s}} && (6.26)\\
& -M(1-q_t) < P_t(\mathbf{X}) - p_t^{\mathbf{i}} && \forall t = 1, ..., T^{\mathrm{s}} && (6.27)\\
& Mr_t > p_t^{\mathbf{o}} - P_t(\mathbf{X}) && \forall t = 1, ..., T^{\mathrm{s}} && (6.28)\\
& -M(1-r_t) \le p_t^{\mathbf{o}} - P_t(\mathbf{X}) && \forall t = 1, ..., T^{\mathrm{s}} && (6.29)\\
& Mq_t \ge P_t(\mathbf{X}) - f_{T^{\mathrm{s}},t}^{\mathbf{i}} && \forall t = T^{\mathrm{s}}+1, ..., T^{\mathrm{s}}+T^{\mathrm{f}} && (6.30)\\
& -M(1-q_t) < P_t(\mathbf{X}) - f_{T^{\mathrm{s}},t}^{\mathbf{i}} && \forall t = T^{\mathrm{s}}+1, ..., T^{\mathrm{s}}+T^{\mathrm{f}} && (6.31)\\
& Mr_t > f_{T^{\mathrm{s}},t}^{\mathbf{o}} - P_t(\mathbf{X}) && \forall t = T^{\mathrm{s}}+1, ..., T^{\mathrm{s}}+T^{\mathrm{f}} && (6.32)\\
& -M(1-r_t) \le f_{T^{\mathrm{s}},t}^{\mathbf{o}} - P_t(\mathbf{X}) && \forall t = T^{\mathrm{s}}+1, ..., T^{\mathrm{s}}+T^{\mathrm{f}} && (6.33)\\
& y_t^{\mathbf{i}} \le Cq_t && \forall t \in \mathcal{T} && (6.34)\\
& y_t^{\mathbf{i}} \ge C - I_t - C(1-q_t) && \forall t \in \mathcal{T} && (6.35)\\
& y_t^{\mathbf{o}} \le Cr_t && \forall t \in \mathcal{T} && (6.36)\\
& y_t^{\mathbf{o}} \ge I_t - C(1-r_t) && \forall t \in \mathcal{T} && (6.37)\\
& P_t(\mathbf{X}) \in \mathbb{R}, q_t, r_t \in \{0,1\} && \forall t \in \mathcal{T}; i = 0, ..., N && (6.38)
\end{aligned}$$

Besides the deterministic formulation of the SCWP ((6.22)-(6.25)), additional constraints (6.26)-(6.38) are required to preserve the optimal policy structure of SCWP-FF (see Proposition 3) and to derive $P_t(\mathbf{X}) \coloneqq \sum_{i=0}^{N} \beta_i X_{it}$ accordingly. Constraints (6.26)-

(6.27) and (6.28)-(6.29) determine the relationship between $p_t^{\mathbf{i}}$ and $P_t(\mathbf{X})$ respectively between $p_t^{\mathbf{o}}$ and $P_t(\mathbf{X})$, i.e., $q_t = 0$ if $p_t^{\mathbf{i}} \geq P_t(\mathbf{X})$ and $q_t = 1$ otherwise, respectively $r_t = 0$ if $P_t(\mathbf{X}) > p_t^{\mathbf{o}}$ and $r_t = 1$ otherwise. Constraints (6.30)-(6.33) are duplications of constraints (6.26)-(6.29) that are required for the forward optimization extension. Constraint (6.34) ensures that $y_t^{\mathbf{i}} = 0$, if $p_t^{\mathbf{i}} \geq P_t(\mathbf{X})$ (i.e., $q_t = 0$) independent of $p_t^{\mathbf{o}}$. Constraint (6.35) ensures that $y_t^{\mathbf{i}} = C - I_t$ if $p_t^{\mathbf{i}} < P_t(\mathbf{X})$ (i.e., $q_t = 1$). Constraint (6.36) ensures that $y_t^{\mathbf{o}} = 0$ if $p_t^{\mathbf{o}} < P_t(\mathbf{X})$ (i.e., $r_t = 0$). Constraint (6.37) ensures that $y_t^{\mathbf{o}} = I_t$ if $p_t^{\mathbf{o}} \geq P_t(\mathbf{X})$ (i.e., $r_t = 1$).

In order to improve *model generalization*, we extend our DDA formulations to regularization. We illustrate this procedure for DDA-OSP-FF in the following.

We apply ℓ_1-norm regularization for feature selection to avoid that the model fits the noise in the data, rather than the underlying functions (over-fitting). We use *Lasso regression* that penalizes non-zero coefficients $\beta_i \in \mathcal{B}$ in order to keep the model from relying too heavily on individual data points. The objective are high-quality decisions, rather than high-quality predictions, which is a benefit over model selection criteria such as AIC or BIC in standard regression that aim at selecting prediction-relevant features.

DDA-OSP-FF with regularization:

$$\max_{\beta_i \in \mathcal{B}} \quad \hat{V} - \lambda \underbrace{\sum_{i=1}^{N} w_i}_{\text{Model complexity}} \tag{6.39}$$

$$\text{s.t.} \quad (6.23) - (6.38) \tag{6.40}$$

$$Mw_i \geq \beta_i \qquad \forall i = 1, ..., N \tag{6.41}$$

$$-Mw_i \leq \beta_i \qquad \forall i = 1, ..., N \tag{6.42}$$

$$w_i \in \{0, 1\} \qquad \forall i = 1, ..., N \tag{6.43}$$

$\lambda \geq 0$ controls regularization and is typically calibrated by n-fold cross-validation (see Mohri et al., 2012, p. 28). For $\lambda = 0$, DDA-OSP-FF with regularization reduces to the basic formulation without penalizing complexity. For $\lambda \to \infty$, DDA-OSP-FF with regularization converges to a solution that estimates P_t as a time-invariant constant. The intercept $\beta_0 \in \mathcal{B}$ is not regularized, which avoids that $P_t(\mathbf{X}) = 0$ for a large λ, which yields $(y_t^{\mathbf{i}^*}, y_t^{\mathbf{o}^*}) = (0, I_t)$, i.e., immediate clearance sale of all available inventory in period t and no replenishment in subsequent periods.

DDA-OSP-LF for SCWP-LF

For SCWP-LF, the estimation of a single price threshold is not sufficient since at the same market price p_t, different purchase-and-inject respectively withdraw-and-sell decisions $y_t^{\mathbf{i}}$ and $y_t^{\mathbf{o}}$ can be optimal depending on the current inventory level I_t.

Definition 13 (Linear Policy Rule Approximation for SCWP-LF)**.** *The base-stock levels $S_t^{\boldsymbol{i}}(\boldsymbol{X})$ and $S_t^{\boldsymbol{o}}(\boldsymbol{X})$ of the optimal policy from Proposition 5 are approximated by*

$$S_t^{\boldsymbol{i}}(\boldsymbol{X}) := \max\left\{0, \sum_{i=0}^{N} \beta_i^{\boldsymbol{i}} X_{it}\right\}, \ \ S_t^{\Delta}(\boldsymbol{X}) := \max\left\{0, \sum_{i=0}^{N} \beta_i^{\Delta} X_{it}\right\},$$

with $S_t^{\boldsymbol{o}}(\boldsymbol{X}) := S_t^{\boldsymbol{i}}(\boldsymbol{X}) + S_t^{\Delta}(\boldsymbol{X})$.

Note that this incremental formulation is required to ensure that $S_t^{\mathbf{i}} \leq S_t^{\mathbf{o}}$. Furthermore, note that $\beta_i^{\mathbf{i}}, \beta_i^{\Delta} \in \mathbb{R}$ are decision variables in the following MILP. For computational performance reasons, we use a fractional formulation with all quantities $y_t^{\mathbf{i}}$, $y_t^{\mathbf{o}}$, $S_t^{\mathbf{i}}(\mathbf{X})$, $S_t^{\mathbf{o}}(\mathbf{X})$, $S_t^{\Delta}(\mathbf{X})$, I_t, $G^{\mathbf{i}}$ and $G^{\mathbf{o}}$ defined as fractions of the warehouse capacity C in order to avoid *Big M* notation.

DDA-OSP-LF:

$$\begin{aligned} \max_{\substack{S_t^{\mathbf{i}}(\mathbf{X}) \\ S_t^{\mathbf{o}}(\mathbf{X})}} \hat{V} = & \frac{C}{T^{\mathrm{s}} + T^{\mathrm{f}}} \Bigg[\sum_{t=1}^{T^{\mathrm{s}}} (p_t^{\mathbf{o}} y_t^{\mathbf{o}} - p_t^{\mathbf{i}} y_t^{\mathbf{i}} - c_h I_t) \\ & + \sum_{t=1}^{T^{\mathrm{f}}} (f_{T^{\mathrm{s}},T^{\mathrm{s}}+t}^{\mathbf{o}} \, y_{T^{\mathrm{s}}+t}^{\mathbf{o}} - f_{T^{\mathrm{s}},T^{\mathrm{s}}+t}^{\mathbf{i}} \, y_{T^{\mathrm{s}}+t}^{\mathbf{i}} - c_h I_{T^{\mathrm{s}}+t}) \Bigg] \end{aligned} \tag{6.44}$$

$$\text{s.t.} \quad I_1 + \sum_{t'=1}^{t} (y_{t'}^{\mathbf{i}} - y_{t'}^{\mathbf{o}}) = I_t - y_t^{\mathbf{o}} + y_t^{\mathbf{i}} \qquad \forall t \in \mathcal{T} = \{1, ..., T^{\mathrm{s}} + T^{\mathrm{f}}\} \tag{6.45}$$

$$S_t^{\mathbf{i}}(\mathbf{X}) - u_t^{\mathbf{i}} \leq I_1 + \sum_{t'=1}^{t} (y_{t'}^{\mathbf{i}} - y_{t'}^{\mathbf{o}}) \leq S_t^{\mathbf{o}}(\mathbf{X}) + u_t^{\mathbf{o}} \quad \forall t \in \mathcal{T} \tag{6.46}$$

$$I_t + q_t^{\mathbf{i}} \geq S_t^{\mathbf{i}}(\mathbf{X}), \quad I_t - q_t^{\mathbf{o}} \leq S_t^{\mathbf{o}}(\mathbf{X}) \qquad \forall t \in \mathcal{T} \tag{6.47}$$

$$I_1 + \sum_{t'=1}^{t} (y_{t'}^{\mathbf{i}} - y_{t'}^{\mathbf{o}}) \leq S_t^{\mathbf{i}}(\mathbf{X}) + (1 - q_t^{\mathbf{i}}) \qquad \forall t \in \mathcal{T} \tag{6.48}$$

$$I_1 + \sum_{t'=1}^{t} (y_{t'}^{\mathbf{i}} - y_{t'}^{\mathbf{o}}) \geq S_t^{\mathbf{i}}(\mathbf{X}) - (1 - q_t^{\mathbf{i}}) - u_t^{\mathbf{i}} \qquad \forall t \in \mathcal{T} \tag{6.49}$$

$$I_1 + \sum_{t'=1}^{t} (y^{\mathbf{i}}_{t'} - y^{\mathbf{o}}_{t'}) \geq S^{\mathbf{o}}_t(\mathbf{X}) - (1 - q^{\mathbf{o}}_t) \qquad \forall t \in \mathcal{T} \tag{6.50}$$

$$I_1 + \sum_{t'=1}^{t} (y^{\mathbf{i}}_{t'} - y^{\mathbf{o}}_{t'}) \leq S^{\mathbf{o}}_t(\mathbf{X}) + (1 - q^{\mathbf{o}}_t) + u^{\mathbf{o}}_t \qquad \forall t \in \mathcal{T} \tag{6.51}$$

$$y^{\mathbf{i}}_t \geq G^{\mathbf{i}} - (1 - u^{\mathbf{i}}_t), \quad y^{\mathbf{o}}_t \geq G^{\mathbf{o}} - (1 - u^{\mathbf{o}}_t) \qquad \forall t \in \mathcal{T} \tag{6.52}$$

$$y^{\mathbf{i}}_t \leq G^{\mathbf{i}}, \quad y^{\mathbf{o}}_t \leq G^{\mathbf{o}} \qquad \forall t \in \mathcal{T} \tag{6.53}$$

$$y^{\mathbf{i}}_t \leq q^{\mathbf{i}}_t, \quad y^{\mathbf{o}}_t \leq q^{\mathbf{o}}_t \qquad \forall t \in \mathcal{T} \tag{6.54}$$

$$u^{\mathbf{i}}_t \leq q^{\mathbf{i}}_t, \quad u^{\mathbf{o}}_t \leq q^{\mathbf{o}}_t \qquad \forall t \in \mathcal{T} \tag{6.55}$$

$$0 \leq y^{\mathbf{i}}_t \leq 1 - I_t + y^{\mathbf{o}}_t, \quad 0 \leq y^{\mathbf{o}}_t \leq I_t \qquad \forall t \in \mathcal{T} \tag{6.56}$$

$$u^{\mathbf{i}}_t, u^{\mathbf{o}}_t, q^{\mathbf{i}}_t, q^{\mathbf{o}}_t \in \{0, 1\} \qquad \forall t \in \mathcal{T} \tag{6.57}$$

Constraint (6.45) ensures inventory balance. Constraint (6.46) ensures that $S^{\mathbf{i}}_t \leq S^{\mathbf{o}}_t$. Constraint (6.47) ensures that $q^{\mathbf{i}}_t = 1$ if $I_t < S^{\mathbf{i}}_t(\mathbf{X})$ and $q^{\mathbf{o}}_t = 1$ if $I_t > S^{\mathbf{o}}_t(\mathbf{X})$. Constraints (6.48)-(6.49) ensure to purchase-up-to $S^{\mathbf{i}}_t(\mathbf{X})$ if $q^{\mathbf{i}}_t = 1$ and $u^{\mathbf{i}}_t = 0$. Constraints (6.50)-(6.51) ensure to sell-down-to $S^{\mathbf{o}}_t(\mathbf{X})$ if $q^{\mathbf{o}}_t = 1$ and $u^{\mathbf{o}}_t = 0$. Constraints (6.52)-(6.53) ensure to purchase $G^{\mathbf{i}}$ if $u^{\mathbf{i}}_t = 1$ respectively sell $G^{\mathbf{o}}$ if $u^{\mathbf{o}}_t = 1$. Constraints (6.54)-(6.55) ensure not to buy or sell if $q^{\mathbf{i}}_t = u^{\mathbf{i}}_t = 0$ respectively $q^{\mathbf{o}}_t = u^{\mathbf{o}}_t = 0$. The auxiliary variables $u^{\mathbf{i}}_t$, $u^{\mathbf{o}}_t$, $q^{\mathbf{i}}_t$ and $q^{\mathbf{o}}_t$ can be interpreted as follows: if $q^{\mathbf{i}}_t = u^{\mathbf{i}}_t = 0$ ($q^{\mathbf{o}}_t = u^{\mathbf{o}}_t = 0$), then $y^{\mathbf{i}}_t = 0$ ($y^{\mathbf{o}}_t = 0$) since $S^{\mathbf{i}}_t \leq I_t$ ($S^{\mathbf{o}}_t \geq I_t$). If $q^{\mathbf{i}}_t = 1 \wedge u^{\mathbf{i}}_t = 0$ respectively $q^{\mathbf{o}}_t = 1 \wedge u^{\mathbf{o}}_t = 0$, then buy-up to $S^{\mathbf{i}}_t$ respectively sell-down to $S^{\mathbf{o}}_t$. If $q^{\mathbf{i}}_t = 1 \wedge u^{\mathbf{i}}_t = 1$ respectively $q^{\mathbf{o}}_t = 1 \wedge u^{\mathbf{o}}_t = 1$, then buy $G^{\mathbf{i}}$ respectively sell $G^{\mathbf{o}}$.

6.3.3. Value Function Approximation Approach (DDA-VFA)

The basic idea of DDA-VFA is similar to Wu et al. (2012), who manipulate the forward curve $\vec{F}_t$ a-priori such that it includes estimates of the extrinsic value generated by its stochastic evolution. Their goal is to close the gap between RIA and the optimal solution of the *full information problem*. However, while Wu et al. (2012) focus on compensating the RIA drawback of ignoring the extrinsic value of storage, we focus on compensating the RIA drawback of ignoring *generalization error*, i.e., we aim at closing the gap to the theoretical perfect foresight solution.

We know that the value function V_t is piecewise linear concave in the inventory level I_t. The idea of DDA-VFA is to approximate the value-to-go function V_{t+1} using feature information by mapping feature data to spot price predictions $\hat{f}_{t,\tau}$ with $\tau > t$ (i.e., prediction in period t for period τ), rather than to decisions (DDA-LDR) or to policy parameters (DDA-OSP). The predictions $\hat{f}_{t,\tau}$ can be obtained by, e.g., a regression model of the following form:

$$\hat{f}_{t,\tau}(\mathbf{X}) \coloneqq \sum_{i=0}^{N} \beta_i^{\tau} X_{it} \qquad \forall \tau = t+1, ..., t+n$$

The estimates $\hat{f}_{t,\tau}$ are thereafter used as input for an LP. Equivalently to RIA, the LP is solved in a rolling horizon fashion to determine $y_t^{\mathbf{i}}$ and $y_t^{\mathbf{o}}$. The first part of the objective function (6.58) is the immediate profit, the second part is the deterministic value-to-go function approximation with approximated spot price estimates $\hat{f}_{t,\tau}$ adjusted to frictions.

DDA-VFA:

$$\max_{y_t^{\mathbf{i}}, y_t^{\mathbf{o}}} \quad \hat{V} = p_t^{\mathbf{o}}\, y_t^{\mathbf{o}} - p_t^{\mathbf{i}}\, y_t^{\mathbf{i}} - c_h\, I_t + \sum_{\tau \in \mathcal{F} | \tau > t} \left[\hat{f}_{t,\tau}^{\mathbf{o}}(\mathbf{X})\, y_\tau^{\mathbf{o}} - \hat{f}_{t,\tau}^{\mathbf{i}}(\mathbf{X})\, y_\tau^{\mathbf{i}} - c_h\, I_\tau \right] \tag{6.58}$$

$$\text{s.t.} \quad I_t + \sum_{\tau'=t}^{\tau} (y_{\tau'}^{\mathbf{i}} - y_{\tau'}^{\mathbf{o}}) = I_\tau - y_\tau^{\mathbf{o}} + y_\tau^{\mathbf{i}} \qquad \forall \tau \in \mathcal{F} = \{t, ..., t+n\} \tag{6.59}$$

$$0 \leq y_\tau^{\mathbf{i}} \leq C - I_\tau + y_\tau^{\mathbf{o}} \qquad \forall \tau \in \mathcal{F} \tag{6.60}$$

$$0 \leq y_\tau^{\mathbf{o}} \leq I_\tau \qquad \forall \tau \in \mathcal{F} \tag{6.61}$$

$$y_\tau^{\mathbf{i}} \leq G^{\mathbf{i}}, \quad y_\tau^{\mathbf{o}} \leq G^{\mathbf{o}} \qquad \forall \tau \in \mathcal{F} \tag{6.62}$$

A benefit of DDA-VFA is that it results in simple LPs as the uncertain parameters p_τ with $\tau > t$ exclusively appear linearly in the objective function. This ensures that we search for policies that belong to the same family as the optimal policy. Furthermore, unlike DDA-LDR and DDA-OSP, DDA-VFA is not necessarily based on ERM and all kinds of linear and non-linear ML methods, including local learning (e.g., kernel regression or random forest), can be employed to obtain the price estimates. In contrast to ERM, local learning makes predictions based on past data that is similar to the data within the recent state. Therefore, equation (6.15) changes to $\max \sum_{t=1}^{T} w_{T,t}(\mathbf{X}) \hat{V}_t(I_t, p_t)$ with $w_{T,t}$ re-weighting the observed data (Bertsimas and Kallus, 2016).

However, just like RIA, DDA-VFA requires the decision maker to define a planning horizon n, which might affect the performance of DDA-VFA significantly. Furthermore, since $\beta_i^{\tau} \in \mathcal{B}$ is determined a-priori, DDA-VFA separates prediction and optimization, which might not yield decision-optimal predictions.

6.4. Results on Empirical Data

Based on six major commodities, we quantify the backtest performance of the presented solution approaches for the SCWP conditional on warehouse characteristics. Section 6.4.1 describes the setup and data used and quantifies the empirical forecast error of futures prices to predict spot prices. In Section 6.4.2, we analyze the sensitivity in a deterministic setting, which enhances our understanding about the myopic performance, the required planning horizon and the value of capacity underutilization. In Section 6.4.3, we demonstrate the weaknesses of RIA in stochastic settings that inevitably involve *generalization error*. In Section 6.4.4, we compare the performance of the different DDA policies relative to RIA. All LP and MILP solutions were obtained by using the XPRESS-MP solver (version 7.6) on an Intel(R) Core(TM) i7-3770, 3.4 GHz processor with 16 GB RAM.

6.4.1. Setup and Descriptive Analysis

We use monthly price data and assume monthly inventory review periods, i.e., each review period coincides with a futures contract maturity (Lai et al., 2010; Secomandi, 2010, 2015). This allows us to evaluate RIA based on futures prices without further assumptions (e.g., price models). Recall that futures prices are used by RIA as a proxy for the market's spot price expectation (futures-based forecasts).

Table 6.2.: Commodity spot and futures price data (2000-2017)

Commodity	Price quotation	Spot market (Data source)	Futures market (Data source)
Metals			
Copper	USD/lb	Nevada Copper (Datastream: NCUCASH)	COMEX (Eikon: HGc1-12)
Gold	USD/ounce	Gold (Eikon: XAU=)	COMEX (Eikon: GCc1-12)
Energy			
Crude oil	USD/bbl	West Texas Intermediate (Datastream: CRUDOIL)	NYMEX (Eikon: CLc1-12)
Natural gas	USD/mmbtu	Henry Hub Natural Gas (Datastream: NATLGAS)	NYMEX (Eikon: NGc1-12)
Agricultural			
Corn	USc/bushel	No.2 Yellow Corn (Datastream: CORNUS2)	CBOT (Eikon: Cc1-12)
Soybean	USc/bushel	No.1 Yellow Soybean (Datastream: SOYBEAN)	CBOT (Eikon: Sc1-12)

Our data refers to monthly closing prices at the first trading day of the corresponding month. Futures contracts for metals and energy are traded at the NYMEX, for agricultural commodities at the CBOT. COMEX is the metal division of NYMEX. We use futures data for the first 12 maturities, i.e., 1- to 12-months-ahead contracts. Even though contracts beyond one year are available for various commodities, these markets are typically highly illiquid with only very few contracts traded, which implies that the predictive content for future spot prices might be low (Alquist and Kilian, 2010). CBOT corn futures mature in March, May, July, September and December. CBOT soybean futures mature in January, March, May, July, August, September and November (www.cmegroup.com). Therefore, corn and soybean prices in-between are linearly interpolated, which is a frequently used and reasonable approximation (Guthrie, 2009).

Note also that futures markets might be more liquid than spot markets that are often thinly traded (Geman and Smith, 2013). We additionally tested the models based on trading in the futures market with closest expiry (the so-called front-month contract) as a proxy for the spot price. As the results are similar, we do not explicitly report them.

Figure 6.6 shows the commodity price evolutions and annualized volatilities σ for the considered time frame.

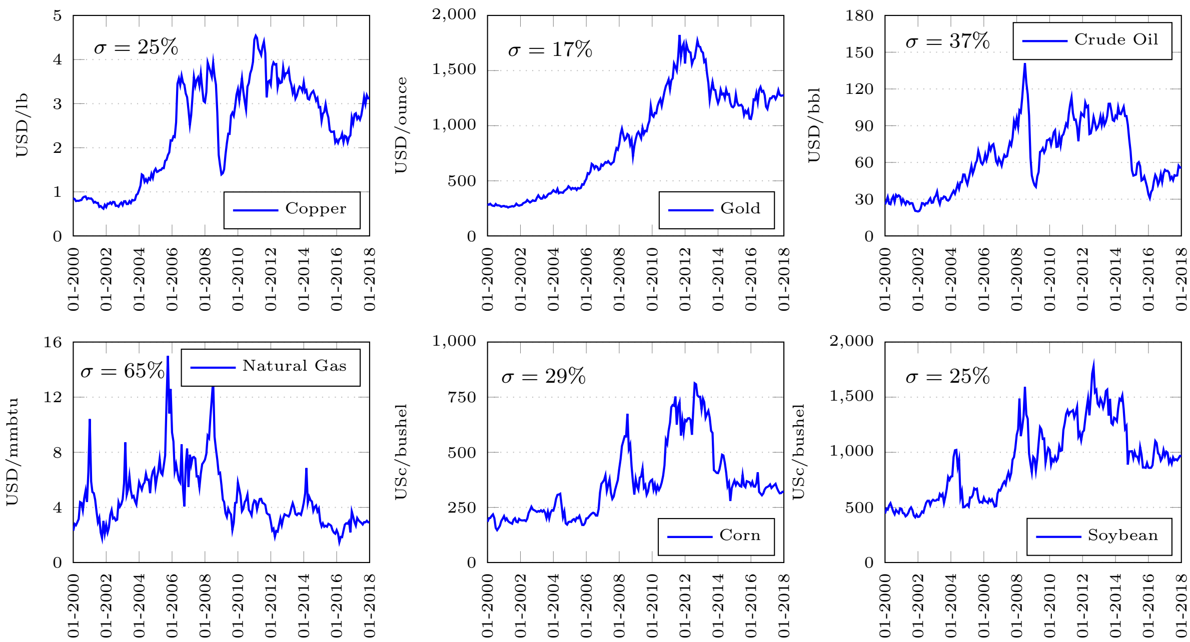

Figure 6.6.: Commodity spot prices and mean annualized volatilities σ (2000-2017)

Empirical Forecast Error of Futures Prices

Under the risk-neutral probability measure $\mathcal{Q}$, deterministic futures prices $f_{t,\tau}$ represent expected future spot prices p_τ with $\tau > t$. Hence, RIA uses the futures curve $\vec{F}_t = (f_{t,\tau} : \tau \in \mathcal{F}, \tau \geq t)$ for inventory trading decisions in t. However, futures prices are not free from prediction error e_t. Based on the commodity data from 2000-2017, we calculate the mean absolute percentage forecasting error (MAPE) of futures prices of different maturities τ to predict spot prices:

$$\text{MAPE}_\tau := \frac{100\%}{N}\sum_{t=1}^{N}\left|\frac{e_t}{p_t}\right| = \frac{100\%}{N}\sum_{t=1}^{N}\left|\frac{p_t - f_{t-\tau,t}}{p_t}\right|$$

As a scale-variant measure, MAPE is not appropriate to compare forecast accuracy across different time series (commodities). Therefore, we additionally calculate the mean absolute scaled error (MASE) that compares the forecast accuracy of $f_{t-\tau,t}$ with the forecast accuracy of naïve forecasts under the random walk model without drift, i.e., $\mathbb{E}_t[p_{t+\Delta t}] = p_t$ for all $\Delta t > 0$ (Hyndman and Koehler, 2006):

$$\text{MASE}_\tau := \frac{1}{N}\sum_{t=1}^{N}\frac{|e_t|}{\frac{1}{N-\tau}\sum\limits_{t=\tau+1}^{N}|p_t - p_{t-\tau}|}$$

Consequently, MASE < 1 indicates that futures-based forecasts outperform naïve forecasts, while MASE > 1 indicates that a naïve forecast is better on average (for which there is evidence in the literature; see Alquist and Kilian 2010).

To measure the ability to detect the correct direction of price changes (classification accuracy), we furthermore calculate the mean directional accuracy (MDA):

$$\text{MDA}_\tau := \frac{100\%}{N}\sum_{t=1}^{N}1_{sign(p_t-p_{t-\tau})==sign(f_{t-\tau,t}-p_{t-\tau})}$$

with *sign* being the sign function and 1 being the indicator function.

Table 6.3 shows that the forecast error in terms of MAPE typically increases with increasing time to maturity. This is in line with Cortazar et al. (2018) and raises the question whether to use all available futures price information for inventory trading decisions, or to use a myopic approach as a series of simple-to-solve two-period problems, which might be less affected by errors.

Note that there is no systematic overestimation ($e_t < 0$) or underestimation ($e_t > 0$) of p_t by $f_{t-\tau,t}$ (see Appendix C.2). However, while MAPE might lead to the conclusion that futures prices are a bad choice for forecasting gas spot prices, MASE indicates their usefulness (compared to naïve forecasts). The MDA results show that $\vec{F}_t$ is able to correctly classify the direction (upward or downward) of front-month spot prices in less than 50% of the periods for copper (46.0%), gold (46.5%) and crude oil (49.3%).

Table 6.3.: MAPE, MASE and MDA of futures to predict spot prices (2000-2017)

	Months to Maturity											
	1	2	3	4	5	6	7	8	9	10	11	12
	MAPE in %											
COMEX Copper	6.2	8.6	11.5	13.9	15.8	17.4	18.6	19.7	20.9	21.8	22.9	24.2
COMEX Gold	3.9	5.1	6.0	6.7	7.5	8.2	9.0	9.8	10.4	11.0	11.5	12.3
NYMEX Crude Oil	8.2	11.6	14.2	16.6	18.9	21.0	22.2	23.4	24.1	25.0	25.7	26.9
NYMEX Natural Gas	13.5	18.2	21.7	25.4	28.4	31.4	33.7	35.7	37.2	38.4	39.6	40.9
CBOT Corn	8.1	10.8	13.7	16.1	18.1	19.2	20.2	21.0	22.0	22.3	22.9	24.2
CBOT Soybean	6.6	8.7	10.2	11.6	13.8	14.2	15.1	15.6	16.4	15.9	16.2	16.4
	MASE											
COMEX Copper	1.09	1.02	1.03	1.01	1.02	1.00	1.01	1.01	1.02	1.03	1.04	1.05
COMEX Gold	1.00	1.00	0.99	0.99	0.98	0.97	0.96	0.95	0.94	0.94	0.94	0.93
NYMEX Crude Oil	1.03	0.99	0.98	0.97	0.96	0.97	0.97	0.97	0.96	0.96	0.96	0.94
NYMEX Natural Gas	0.95	0.95	0.96	0.98	0.94	0.97	0.95	0.94	0.93	0.95	0.96	0.98
CBOT Corn	1.08	1.01	1.02	1.02	1.01	1.01	0.99	1.00	1.01	0.99	0.98	0.98
CBOT Soybean	0.98	0.94	0.89	0.87	0.87	0.85	0.86	0.84	0.85	0.83	0.84	0.82
	MDA in %											
COMEX Copper	46.0	54.7	54.5	54.7	51.2	53.8	55.5	57.7	55.1	54.4	54.1	53.9
COMEX Gold	46.5	56.1	56.8	61.8	63.5	66.7	69.4	69.7	71.5	70.4	72.2	73.5
NYMEX Crude Oil	49.3	50.5	53.5	52.4	55.0	56.7	56.9	56.7	56.0	58.7	61.5	59.8
NYMEX Natural Gas	60.9	53.3	56.3	52.4	53.6	54.3	58.4	57.7	61.8	61.2	64.9	64.7
CBOT Corn	60.9	57.9	54.9	59.0	55.0	57.1	54.1	54.8	56.0	59.7	58.5	58.8
CBOT Soybean	57.7	60.7	67.1	65.1	59.7	61.0	63.6	62.5	60.4	68.4	67.8	69.1

6.4.2. Deterministic Analysis

We start with a deterministic analysis of the SCWP with prices known in advance.

Loss Through Myopic Inventory Decisions

In the first experiment, we test the performance of myopic inventory decisions (MYOPIC) relative to the optimal inventory decisions (OPT) based on the six commodity time series in a rolling horizon fashion from 2000 to 2017. We aim at identifying operational drivers of the performance of MYOPIC (i.e., the simplest limited-look-ahead approach). As in Lai et al. (2010) and Secomandi (2015), we normalize the warehouse capacity to $C = 1$, initial inventory is $I_0 = 0$ (in January 2000) and frictions occur due to injection and withdrawal loss $\eta^{\mathbf{i}}, \eta^{\mathbf{o}} \in \{0.95, 0.955, 0.96, ..., 1\}$. Following the standard literature (Lai et al., 2010; Secomandi, 2010), marginal storage costs are $c_h = 0$ (however, we observed that plausible numbers for c_h lead to similar results). There are no injection and withdrawal costs ($c^{\mathbf{i}} = c^{\mathbf{o}} = 0$). Furthermore, we distinguish between fully flexible storage (SCWP-FF) with $G^{\mathbf{i}} = G^{\mathbf{o}} = C$ and limited flexible storage (SCWP-LF) with $G^{\mathbf{i}} = G^{\mathbf{o}} = \frac{1}{2}C$. The following results show what a merchant with perfect foresight price information might lose through (i) limited operational flexibility, (ii) frictions and (iii) myopic decision-making. The corresponding optimal profits under perfect foresight for the six commodities and for different time series phases are presented in Appendix C.3.1. They describe upper performance bounds.

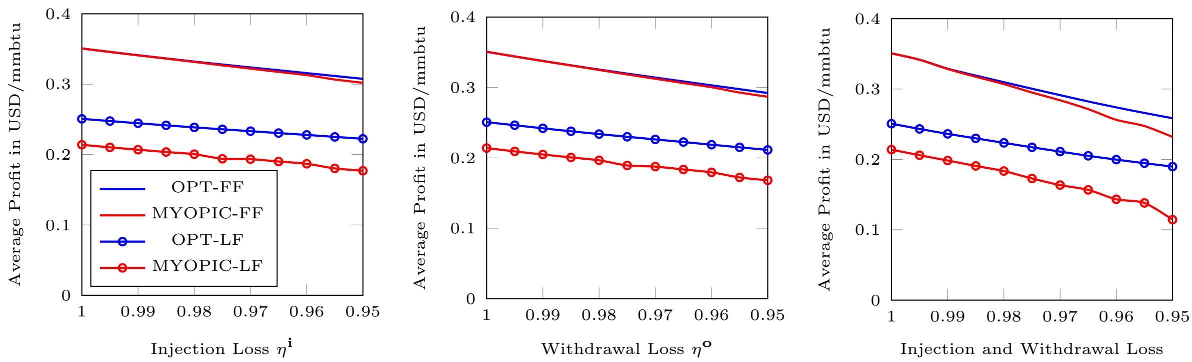

Figure 6.7.: Gas storage: Performance of optimal vs. myopic policies (2000-2017)

Figure 6.7 shows that frictions decrease the profit of gas storage facilities significantly. For the fully flexible storage case (e.g., salt caverns), a 1% injection and withdrawal loss translates into 6.22% profit decline. However, we observe for the fully flexible storage case that, if frictions are realistically small (1%; see Secomandi 2010), the profit decline by myopic decision-making is approximately 0.2%. This is an interesting observation as MYOPIC is less affected by prediction errors (Note that we would not reach OPT anyway due to *generalization error*). Limited storage flexibility (e.g., at gas aquifers) decreases the merchant's profit significantly and MYOPIC deviates from the optimum by 14.7% even for the frictionless case. We make similar observations for all other analyzed commodities (see Appendix C.3.2). Our analysis on real data shows that the myopic policy, which does not require any optimization and is less affected by prediction errors since it is based on one-step-ahead price forecasts (that even do not have to be very adequate) is optimal for full operational flexibility without frictions and near-optimal for plausible frictions. However, for limited storage flexibility (that is mainly relevant for natural gas), myopic policies are far from optimal ($> 10\%$).

Impact of the Planning Horizon

In order to quantify the impact of the planning horizon, we solve the RIA-LP under perfect spot price foresight for different planning horizons $n \in \{1, 2, 3, ..., 12\}$ and identical warehouse characteristics as before ($C = 1$, $c_h = 0$, SCWP-FF versus SCWP-LF with $G^{\mathbf{i}} = G^{\mathbf{o}} = 0.5$, frictions with $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} \in \{1, 0.995, 0.99, 0.985, ..., 0.95\}$). We evaluate RIA-LP in a rolling horizon fashion based on the spot price data from 2000 to 2017.

Consequently, for each instance, we sequentially solve 216 LPs based on the updated inventory state I_{t+1} and the updated futures curve $\vec{F}_{t+1}$ (that equals the vector of perfect spot price forecasts in this deterministic setting). Figure 6.8 shows the profit exemplarily for gas storage measured as percentage of the theoretical optimal profit under perfect foresight with the planning horizon limited by the last data point (December 2017). The results for the other commodities are qualitatively similar and reported in Appendix C.3.3.

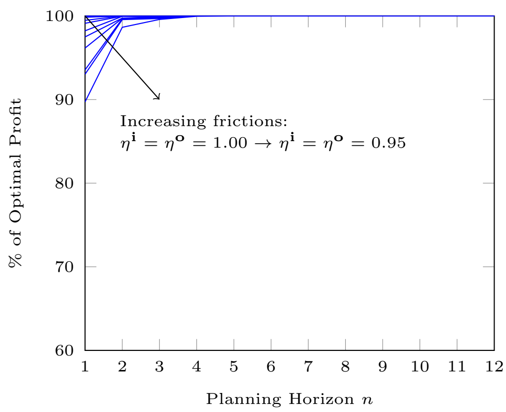

(a) SCWP-FF

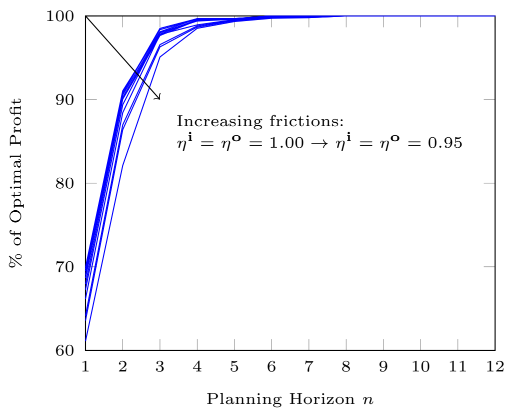

(b) SCWP-LF

Figure 6.8.: Gas storage: Performance impact of planning horizon n (2000-2017)

Figure 6.8 shows that the required planning horizon increases with increasing frictions and with decreasing operational flexibility. For plausible frictions $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.99$, a planning horizon of $n = 1$ generates 99.84% (SCWP-FF) respectively 69.49% (SCWP-LF) of the potential value. $n = 2$ generates 100% (SCWP-FF) respectively 90.83% (SCWP-LF) and $n = 3$ generates 100% (SCWP-FF) respectively 98.36% (SCWP-LF) of the potential value. A planning horizon of $n = 2$ (SCWP-FF) respectively $n = 6$ (SCWP-LF) is sufficient to generate 100% of the value. Beyond 6 future periods, it is not necessary to look at for optimal first-stage decisions. That is related to the definition of forecast horizons, i.e., data beyond that period do not affect optimal first-stage actions (Chand et al., 2002). We observe that the forecast horizon at gas storage facilities compared to, e.g., gold (Appendix C.3.3) is smaller due to a higher market volatility. If the price varies more often, then shorter planning horizons are sufficient for optimal decisions.

Value of Injection and Withdrawal Capacity Underutilization

In this paragraph, we quantify the loss by ignoring injection and withdrawal rate constraints $G^{\mathbf{i}} < C$ and $G^{\mathbf{o}} < C$. This is insofar interesting to study as a single price threshold P_t for the SCWP-FF (i.e., under the absence of capacity limits) might be faster to train than base-stock levels $S_t^{\mathbf{i}}$ and $S_t^{\mathbf{o}}$ for the SCWP-LF (i.e., under the presence of capacity limits). The difference is that $S_t^{\mathbf{i}}$ and $S_t^{\mathbf{o}}$ allow for capacity underutilization, while P_t does not. If the value of capacity underutilization is already low in the deterministic analysis under perfect foresight, then training P_t might be a good approximation for training $S_t^{\mathbf{i}}$ and $S_t^{\mathbf{o}}$ in the stochastic setting.

We calculate the value of considering capacity underutilization by calculating the loss if we restrict OPT-LF to purchase and selling decisions whenever OPT-FF purchases and sells (timing), but limit the buying and selling quantity $y_t^{\mathbf{i}}$ and $y_t^{\mathbf{o}}$ ex-post by the capacity functions $G^{\mathbf{i}}$ and $G^{\mathbf{o}}$. In other words: we solve OPT-FF (i.e., without capacity limits) to obtain when to buy/to sell under full operational flexibility and restrict the purchase and selling decisions in a second step a-posteriori by $G^{\mathbf{i}}$ and $G^{\mathbf{o}}$. The resulting profit is compared to the profit of OPT-LF (i.e., with a-priori consideration of $G^{\mathbf{i}}$ and $G^{\mathbf{o}}$) within the optimization. The profit gap quantifies the loss caused by ignoring capacity underutilization and is plotted in Figure 6.9 for the different commodities over the entire evaluation horizon (2000-2017) with $c_h = 0$, $C = 1$, $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.99$ and an empty warehouse at the beginning of the year 2000.

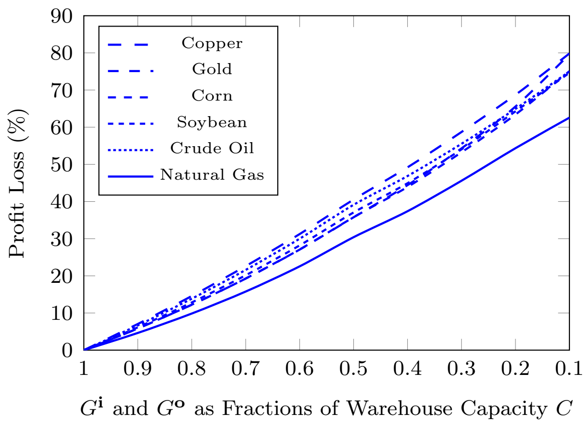

Figure 6.9.: Profit loss by ignoring injection and withdrawal capacity underutilization

Figure 6.9 implies that ignoring the double base-stock structure ($S_t^{\mathbf{i}}$, $S_t^{\mathbf{o}}$), i.e., ignoring the option of capacity underutilization and estimating a price threshold P_t instead, reduces the merchant's profit significantly as soon as capacity restrictions are tight (i.e., small $G^{\mathbf{i}}$ and $G^{\mathbf{o}}$). For $G^{\mathbf{i}} = G^{\mathbf{o}} = C$, the storage facility is fully flexible and capacity underutilization is never optimal. If $G^{\mathbf{i}}$ and $G^{\mathbf{o}}$ are close to C, then SCWP-LF might be approximated by SCWP-FF, i.e., the double base-stock structure ($S_t^{\mathbf{i}}$,$S_t^{\mathbf{o}}$) can be approximated by a simple price threshold policy (P_t) without huge losses.

6.4.3. Stochastic Analysis: Performance of RIA

The deterministic analysis under perfect foresight from Section 6.4.2 provides an upper bound for the profit but disregards uncertainty about future prices and hence ignores the implications of *generalization error*. It is an open question how close the real warehouse operator can come to the theoretical optimal profit under perfect foresight.

Futures-Based RIA

In the following, we evaluate the performance of RIA that uses $\vec{F}_t = (f_{t,\tau} : \tau \in \mathcal{F} = \{t, t+1, ..., t+n\})$ to predict spot prices p_τ. Note that RIA with $n = 1$ is equivalent to the myopic policy. Table 6.4 summarizes the parameters that we vary to obtain $4 \times 2 \times 3 \times 9 = 216$ instances per commodity (i.e., 1,296 instances in total) to perform our RIA backtests. We use the same commodity spot price data (2000-2017) and the identical storage setting ($C = 1$, $c_h = 0$, $I_0 = 0$) as for the deterministic analysis.

Table 6.4.: Summary of the numerical design

Planning horizon	$n \in \{1, 3, 6, 12\}$
Storage flexibility	$G^{\mathbf{i}} = G^{\mathbf{o}} \in \{0.5, 1\}$
Frictions	$\eta^{\mathbf{i}} = \eta^{\mathbf{o}} \in \{1, 0.995, 0.99\}$
Sub-periods	$2000-2001$, $2002-2003$, ..., $2016-2017$

For each instance, we solve the RIA-LPs sequentially over the evaluation horizon from 2000 to 2017 based on the updated inventory state information $I_{t+1} = I_t - y_t^{\mathbf{o}} + y_t^{\mathbf{i}}$ and the updated futures curve $\vec{F}_{t+1} = (f_{t+1,\tau} : \tau \in \mathcal{F}, \tau \geq t+1)$. The performance is measured by $\frac{V^{\text{RIA}}}{V^{\text{PF}}} \cdot 100\%$, with V^{RIA} being the achieved profit according to RIA and V^{PF} being the theoretical profit under perfect foresight. We make the following observations.

Observation 1. *Under generalization error, RIA can yield unprofitable operations.*

The results in Table 6.5 show that RIA achieves on average 11.0% of the perfect foresight profit. This finding is in strong contrast to the literature that establishes the near-optimality of RIA in the full information setting, mainly for natural gas storage (Lai et al., 2010; Secomandi, 2010). RIA capturing only a small percentage of the perfect foresight bound is intriguing but even more concerning is the significant occurrence of negative RIA profits in our backtests (in 30.1% of the instances). Negative profits for instance occur whenever the warehouse operator expects decreasing prices and therefore would sell the available inventory to the market disregarding sunk cost.

Table 6.5.: Performance of futures-based RIA in $V^{\mathrm{RIA}}/V^{\mathrm{PF}} \cdot 100\%$ across all instances from Table 6.4

	Mean	Min	25%-Q	50%-Q	75%-Q	Max	Mean*	Mean**
Copper	1.3	-58.5	-12.2	0.0	10.6	70.0	4.1	1.5
Gold	10.0	-107.5	0.0	0.0	23.2	78.7	9.4	11.1
Crude Oil	-5.1	-171.9	-30.1	7.1	26.3	70.5	-3.3	7.6
Natural Gas	18.6	-125.9	6.0	25.6	44.8	79.5	21.2	28.5
Corn	16.7	-44.3	-5.8	14.5	42.2	64.9	20.0	19.4
Soybean	24.5	-31.6	-1.0	20.8	45.6	79.8	28.1	27.3
Overall	11.0	-171.9	-4.6	9.5	34.9	79.8	13.3	15.9

Note. * Mean w/o 2008-2009 (financial crisis), ** Mean w/o 2014-2015 (oil price drop).

We also observe a strong impact of the oil price drop during 2015 on the performance of energy storage systems. If we exclude 2014-2015 from the dataset, the average performance of RIA for crude oil increases from -5.1% (unprofitable storage) to 7.6% (profitable storage) with a worst case performance of -64.2%, rather than -171.9%.

Furthermore, the performance also varies with the operational setting $(G^{\mathbf{i}}, G^{\mathbf{o}}, \eta^{\mathbf{i}}, \eta^{\mathbf{o}})$ (see Table C.1 in Appendix C.4): identical price forecasts can yield both profitable and unprofitable storage. While for a certain operational setting (e.g., copper, SCWP-FF, $n = 12$, $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.995$), a forecast is beneficial, it can be disadvantageous for another setting (e.g., copper, SCWP-FF, $n = 12$, $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.99$). Moreover, if frictions are large (i.e., small $\eta^{\mathbf{i}}$ and $\eta^{\mathbf{o}}$) relative to the (expected) price changes, the warehouse slows down its activity (see for some instances of gold in Appendix C.4).

Observation 2. *Under generalization error, ignoring information can be beneficial.*

We investigate the performance effect of the planning horizon n and refer to a version of RIA restricted to a certain planning horizon $n \in \{1, 3, 6, 12\}$ as $\mathrm{RIA_n}$.

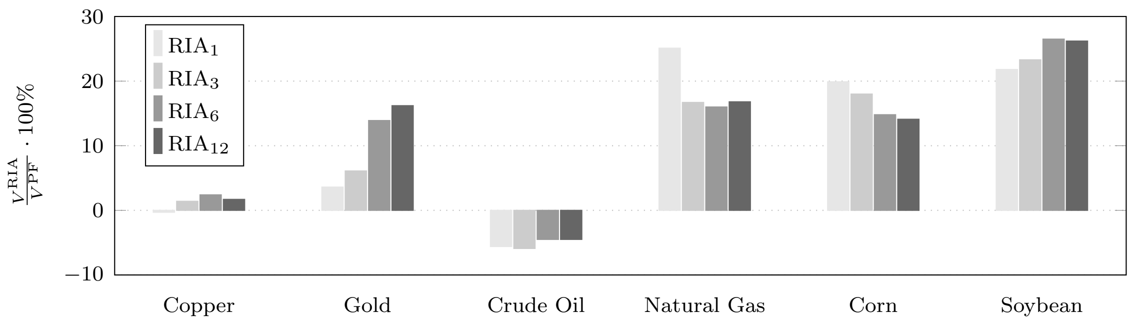

Figure 6.10.: Average performance of $\mathrm{RIA_n}$ in $V^{\mathrm{RIA_n}}/V^{\mathrm{PF}} \cdot 100\%$ across all instances

Figure 6.10 shows that the performance of RIA is highly sensitive to the planning horizon n (except for SCWP-FF instances without frictions where myopic policies are optimal; see Table C.1 of Appendix C.4). For a variety of settings, it is beneficial to exclude futures prices from its objective that correspond to maturities further out in the planning horizon, which often have lower liquidity and hence may provide a poorer indication of the spot price at its maturity. Furthermore, we observe that futures prices up to 12 months do not provide much additional valuable information for RIA compared to futures price information up to 6 months.

Table 6.6.: Dominance matrix of RIA_n in % of instances ($\# = 1,296$) in which row performs strictly better than column

	RIA_1	RIA_3	RIA_6	RIA_{12}
RIA_1	-	28.7	28.4	29.9
RIA_3	31.2	-	9.6	14.8
RIA_6	38.9	21.9	-	8.0
RIA_{12}	41.4	27.5	12.0	-

In total, RIA_1 (myopic) strictly outperforms RIA_{12} in 29.9% of the instances and, in $100 - 41.4 = 58.6\%$ of the instances, myopic does at least as well as RIA_{12} (see Table 6.6). Further, RIA_6 is equal to or better than RIA_{12} in $100 - 12.0 = 88.0\%$ of the instances, which corroborates our intuition that futures prices with later maturities adversely affect the RIA_{12} operating policy. This finding is different in nature from forecast horizon results in operations (Chand et al., 2002) including the RIA literature (Cruise et al., 2019), where it is shown that using a shorter planning horizon results in same optimal decisions in the full information setting. Instead, the suggestion to reduce the planning horizon stems from futures prices being poor forecasts for longer horizons.

Observation 3. *Under generalization error, the value of reoptimization is not necessarily positive.*

The preceding results also bring to question whether the reoptimization embedded in the definition of the RIA policy adds value over the static intrinsic policy computed using the futures curve available at the initial stage. For the full information setting, we know that reoptimization adds significant value (Lai et al., 2010; Secomandi, 2015).

We define the value of reoptimization as $\text{VReO} := \left(\frac{V^{\text{RIA}} - V^{\text{IA}}}{V^{\text{PF}}}\right) \cdot 100\%$, with V^{IA} as the corresponding profit of the intrinsic approach (IA). While RIA revises the inventory plan in each period based on new futures price information available, IA determines the plan in t for all remaining periods $t, t{+}1, ..., t{+}n$ based on $\vec{F_t} = (f_{t,\tau} : \tau \in \mathcal{F} = \{t, t{+}1, ..., t{+}n\})$.

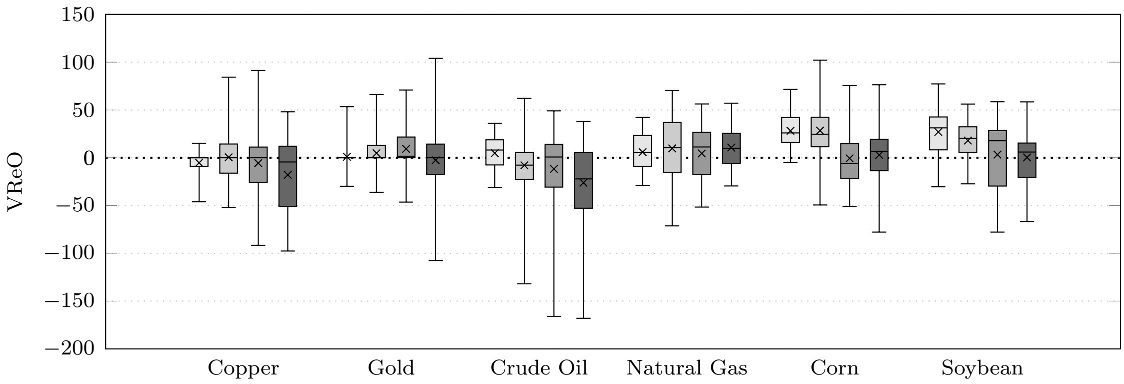

Figure 6.11.: Value of reoptimization for $n \in \{1, 3, 6, 12\}$ (light gray to dark gray)

Figure 6.11 summarizes the backtesting results across all instances with respect to the planning horizon n. We observe that particularly for copper and crude oil, IA performs surprisingly well if compared to RIA. In total over all commodities, the value of reoptimization is negative in 37.0% of the 1,296 instances, which fundamentally contradicts the results of the literature (Lai et al., 2010; Secomandi, 2015) that compares IA and RIA for the *full information problem*.

Figure 6.12 presents when reoptimization would have generated value and when not. We observe that IA outperforms RIA most notably (but not exclusively) in phases of sharp price jumps or sharp price drops.

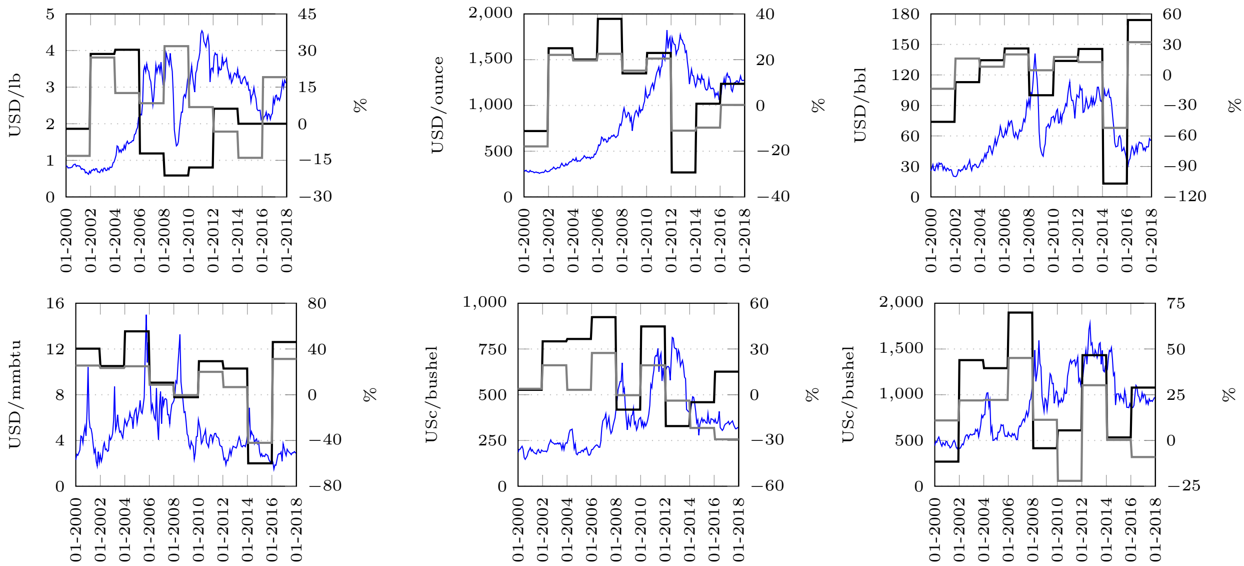

Figure 6.12.: Average performance of RIA (black) versus IA (gray) in $V^{\mathrm{RIA}}/V^{\mathrm{PF}} \cdot 100\%$ for different time series phases of copper, gold, crude oil, natural gas, corn and soybean

Observation 4. *Correctly specifying the direction of one-step-ahead spot prices (upward or downward) generates significant additional value compared to RIA. The value generated by estimating the correct forecast direction can hardly be improved by perfect one-step-ahead point forecasts.*

Finally, we investigate the nature of "high fidelity" information that RIA would need to perform well. Therefore, everything else being equal, we manipulate the front-month price predictions $\hat{p}_{t+1}$. We distinguish between *(i)* perfect one-step ahead point forecasts (PPF) with $\hat{p}_{t+1} = p_{t+1}$ and *(ii)* perfect one-step ahead directional forecasts (PDF) with $\hat{p}_{t+1} = p_t \cdot (1 + \bar{\delta})$ if $p_{t+1} > p_t$ and $\hat{p}_{t+1} = p_t \cdot (1 - \bar{\delta})$ if $p_{t+1} \leq p_t$. We set $\bar{\delta} > 0$ equal to the average monthly percentage spot price change of the corresponding commodity over the previous two years.

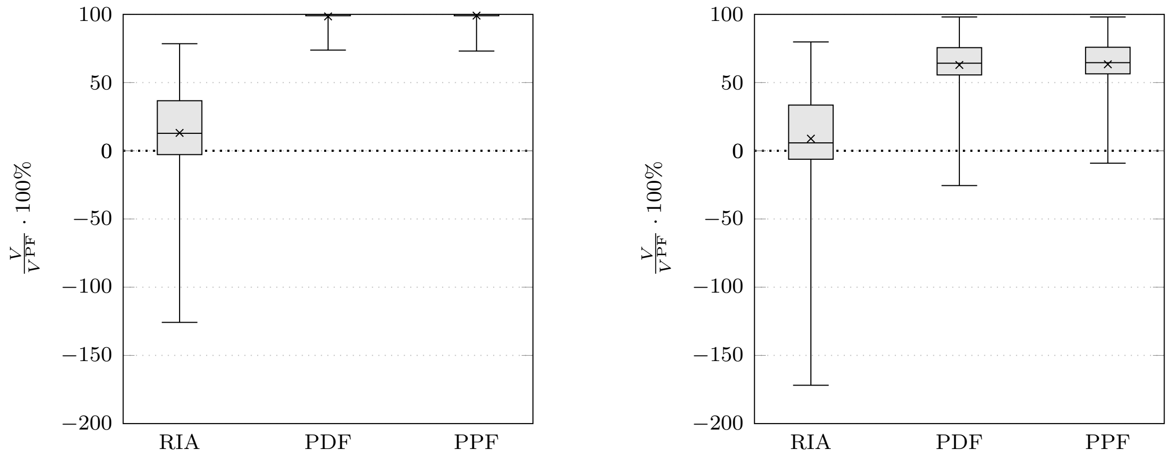

(a) SCWP-FF **(b) SCWP-LF**

Figure 6.13.: Value of perfect one-step-ahead directional forecasts (PDF) and perfect one-step-ahead point forecasts (PPF) for RIA across all commodities and instances from 2000 until 2017

Figure 6.13 shows that appropriate one-step-ahead price estimates generate significant additional profit compared to RIA that misspecifies both future prices (see the MAPE in Table 6.3) and price trends (see the MDA in Table 6.3). The perfect foresight value for SCWP-FF can almost fully (on average 98.5%) be captured by correctly classifying the direction of one-step-ahead price movements. Appropriate point forecasts aiming at a small MAPE do not generate significant additional value (on average 99.1%) over directional forecasts that aim at large MDA. Detailed results on the performance of PDF and PPF relative to RIA are presented in Table C.1 of Appendix C.4.

RIA Based on Analyst Forecasts

Cortazar et al. (2018) argue that there is no consensus on the amount of new information already captured in market prices and that futures-based forecasts do not incorporate explicit information about the current risk premium. They find that analyst forecasts can outperform futures-based forecasts in predicting future spot prices.

Hence, we exploit the value of the *Bloomberg's Analysts' Median Composite Forecasts (CPFC)* that report the medians composed of the quarterly price forecasts offered by several major financial institutions. While individual expert forecasts may exhibit high prediction errors, by using the median forecast over a variety of well-established financial institutions (up to 31) we expect some error diversification (Cortazar et al., 2018). Based on the quarterly median forecasts, we generate monthly analyst forecast curves $\vec{A}_t = (a_{t,\tau} : \tau \in \mathcal{A} = \{t, t+1, ..., t+n\})$ for the six commodities for planning horizons up to 12 months. Due to restricted data availability (2008 until 2017), we can only evaluate 5, rather than 9, sub-periods, which yields 120 instances per commodity.

Besides analyst-based RIA, we also evaluate RIA based on deterministic price forecasts obtained from AR(1) models, each estimated on the previous two years of spot price data.

Table 6.7.: Performance of futures-based RIA, analyst-based RIA and AR(1)-based RIA in $V^{\text{RIA}}/V^{\text{PF}} \cdot 100\%$ from 2008 until 2017 across all instances (In-sample performance of AR(1)-based RIA in parenthesis)

	Mean	Min	25%-Q	50%-Q	75%-Q	Max
	Futures-based RIA					
Copper	-6.6	-45.6	-16.4	0.0	0.0	25.8
Gold	3.6	-107.5	0.0	0.0	18.6	61.3
Crude Oil	-6.7	-171.9	-30.1	9.7	35.2	70.5
Natural Gas	7.3	-125.9	-15.6	21.7	39.4	64.7
Corn	4.9	-44.3	-18.9	-2.1	25.6	62.1
Soybean	15.7	-18.6	-2.8	10.4	32.4	76.3
	Analyst-based RIA					
Copper	-22.2	-242.1	-29.3	-4.8	9.8	47.9
Gold	5.9	-110.1	-12.8	17.5	33.9	64.5
Crude Oil	-37.6	-338.1	-75.4	16.0	45.3	77.7
Natural Gas	-21.0	-103.5	-58.5	-49.8	33.9	60.2
Corn	18.0	-45.9	-3.3	15.4	40.9	64.2
Soybean	27.3	-3.0	14.8	26.3	34.1	78.2
	AR(1)-based RIA					
Copper	-36.7 (29.3)	-412.9 (-3.0)	0.0 (3.5)	0.0 (29.1)	39.0 (49.4)	57.4 (77.2)
Gold	5.5 (31.7)	-62.8 (-20.3)	0.0 (0.0)	0.0 (34.2)	28.8 (57.2)	69.1 (92.1)
Crude Oil	-36.4 (36.1)	-262.0 (0.0)	-20.2 (14.6)	0.0 (31.2)	7.7 (60.0)	48.0 (85.4)
Natural Gas	-4.2 (22.7)	-79.0 (-40.7)	-32.3 (-3.2)	-21.4 (34.5)	37.2 (59.4)	60.7 (66.6)
Corn	6.9 (24.5)	-59.1 (0.0)	0.0 (7.1)	9.6 (19.2)	27.1 (33.1)	38.5 (78.8)
Soybean	-0.5 (42.7)	-91.4 (0.0)	-18.8 (25.8)	0.0 (36.8)	26.7 (60.7)	61.9 (77.8)

Table 6.7 summarizes the results. We observe that analyst forecast-based RIA outperforms both futures-based and AR(1)-based RIA for the commodities gold, corn and soybean in both average and median values (Table 6.7). The performance loss of AR(1)-based RIA from in-sample to out-of-sample evaluation indicates a weak generalization of AR(1) price models.

We furthermore observe that analyst-based RIA strictly outperforms both futures-based RIA and AR(1)-based RIA for the majority of the 720 instances (see Table 6.8). However, there is no significant dominance of one approach over all others.

Table 6.8.: Dominance Matrix of different RIA policies in % of instances (# = 720) in which row performs strictly better than column

	Futures-based RIA	Analyst-based RIA	AR(1)-based RIA
Futures-based RIA	-	44.6	50.0
Analyst-based RIA	54.4	-	51.1
AR(1)-based RIA	41.5	46.8	-

Our analysis also shows the importance of the directional forecast accuracy for optimal inventory decisions: even though the MAPE of futures-based forecasts is smaller than the MAPE of analyst forecasts for almost all commodities and forecast horizons (see Figure 6.14 on the left), analyst-based RIA outperforms futures-based RIA for gold, corn and soybean due to a higher directional accuracy (MDA), particularly in the more important close maturities (see Figure 6.14 on the right).

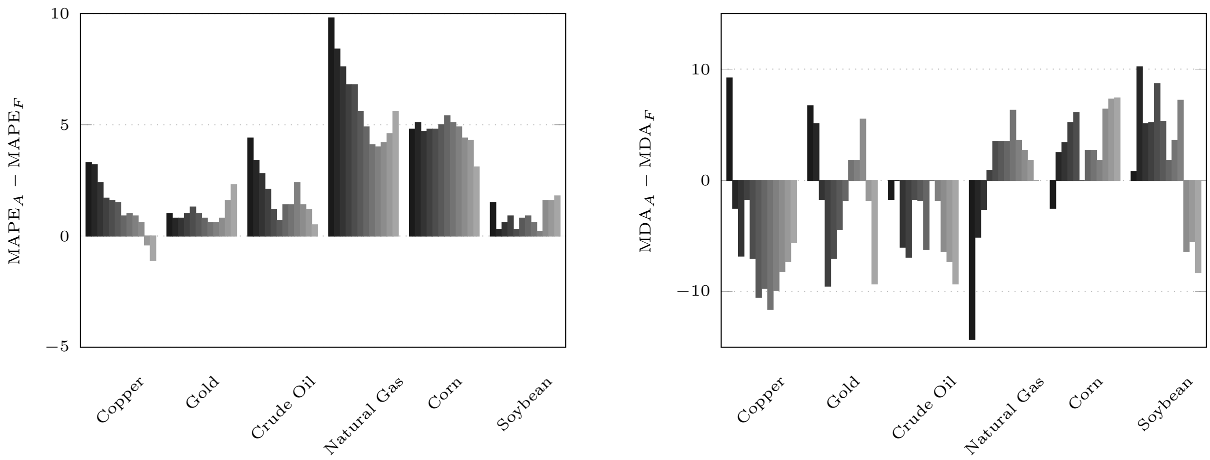

(a) MAPE **(b) MDA**

Figure 6.14.: Forecast accuracy 2008-2017 (MAPE, MDA) of futures-based predictions (F) and analyst-based predictions (A) for horizons of 1 up to 12 months

6.4.4. Stochastic Analysis: Performance of DDA

From observations (1)-(4), we conclude that futures-based RIA yields several severe drawbacks when applied to real data. In this section, we want to exploit whether feature-based DDA that is capable of combining all available information can outperform RIA that is solely based on futures price data. Furthermore, we want to find out which of the DDA approaches (DDA-LDR, DDA-OSP or DDA-VFA) can exploit the same available dataset most effectively with regard to reducing inventory-related costs.

Table 6.9.: Benchmark approaches and feature data

Approach	Computation	Policy consistency	Futures prices	Spot history	Analyst forecasts
IA	LP	yes	✓	✗	✗
RIA	LP	yes	✓	✗	✗
DDA-LDR	LP	no	✓	✓	✓*
DDA-OSP	MILP	yes	✓	✓	✓*
DDA-VFA	Regression + LP	yes	✓	✓	✓*

* Due to data availability, analyst forecasts are used starting from January 2008.

Table 6.9 summarizes the approaches that we compare and the data that the approaches exploit. Besides futures prices ($f_{t,t+1}, f_{t,t+2}, ..., f_{t,t+12}$), we additionally consider recent and past spot prices ($p_t, p_{t-1}, p_{t-2}, p_{t-3}$) (see Sioshansi et al. (2009) with application to electricity storage) and quarterly analyst forecasts ($a_{t,q}, a_{t,q+1}, a_{t,q+2}, a_{t,q+3}$) (see Cortazar et al. (2018) with application to oil price prediction).

Table 6.10.: Backtesting setup

Sample	In-Sample	Out-of-Sample	Sample	In-Sample	Out-of-Sample
1	2000-2001	2002-2003	5	2008-2009	2010-2011
2	2002-2003	2004-2005	6	2010-2011	2012-2013
3	2004-2005	2006-2007	7	2012-2013	2014-2015
4	2006-2007	2008-2009	8	2014-2015	2016-2017

Similar to Secomandi et al. (2015), we split the data into sub-periods (Table 6.10). We use the identical storage setup as described in Table 6.4. This yields $2 \times 3 \times 8 \times 4 = 192$ instances per commodity (Note that the performance of DDA-LDR and DDA-OSP is not sensitive to the planning horizon n). We test DDA-LDR and DDA-OSP with and without forward optimization. For forward optimization, we use futures curves $\vec{F}_t$ with the 12 closest monthly maturities ($T^{\mathrm{f}} = 12$). To avoid overfitting and to enable feature selection, we apply Lasso regression (see equations (6.39)-(6.43)) to regularize DDA-LDR and DDA-OSP. We optimize the regularization parameter $\lambda \in \{0, 10^{-4}, 10^{-3}, ..., 10^{3}\}$ within a cross-validation procedure splitting the in-sample data equally into training and validation sets. For the a-priori predictions of DDA-VFA, we use OLS regression with the `stepAIC`-function in R (`MASS` package) for stepwise forward-backward regression under the AIC model selection criterion for feature selection.

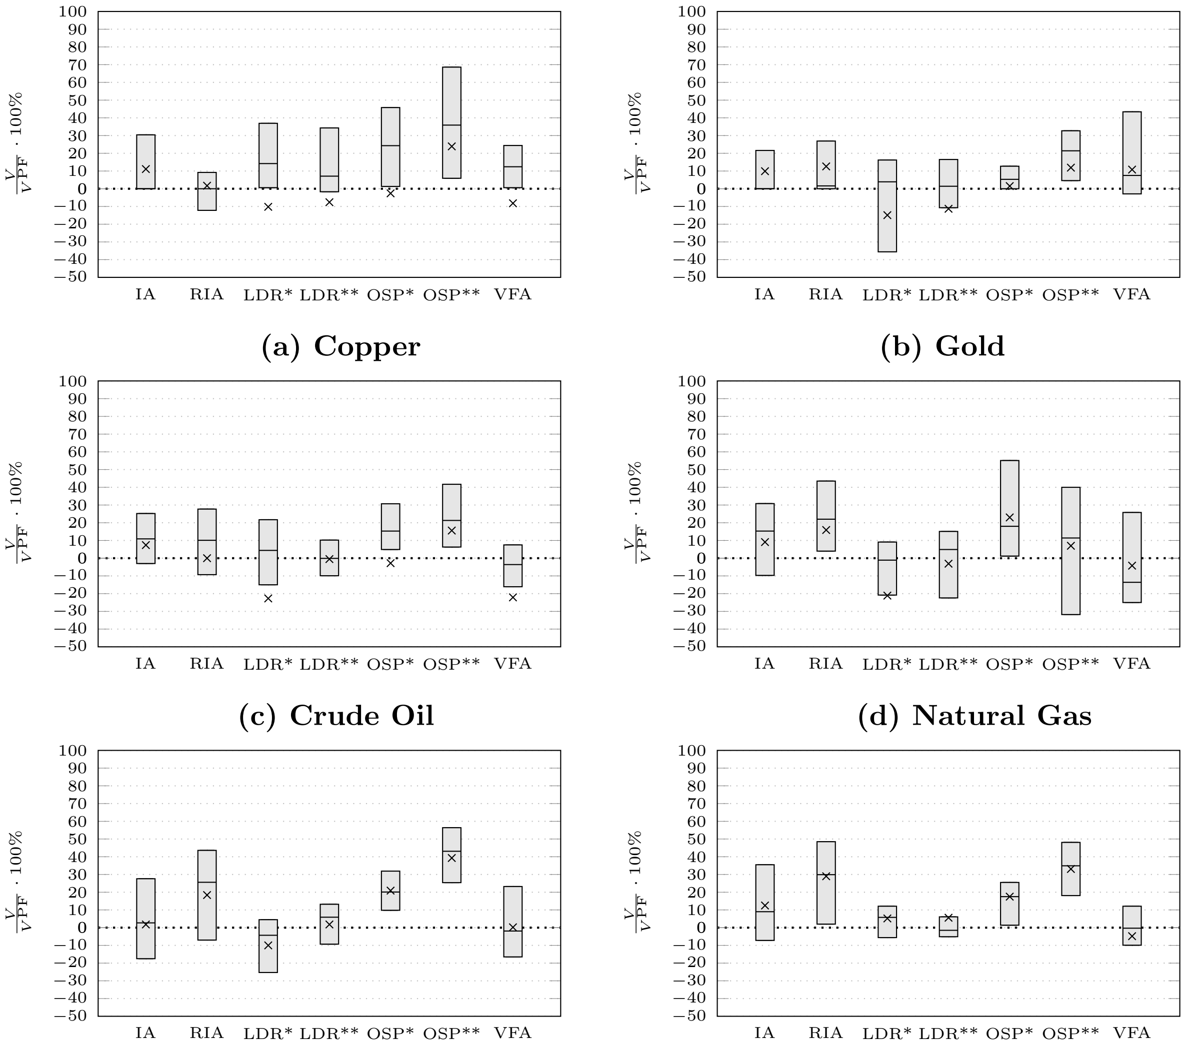

(a) Copper **(b) Gold**

(c) Crude Oil **(d) Natural Gas**

(e) Corn **(f) Soybean**

Figure 6.15.: Out-of-sample performance of the different policies in $V/V^{\mathrm{PF}} \cdot 100\%$ from 2002 until 2017 across the 192 instances

Note. * w/o forward optimization, ** w/ forward optimization. Boxplots characteristics: 1st-, 2nd-, 3rd-quartile, mean (×). For a better graphical comparability of the quartiles, we do not explicitly show the whiskers in these plots. The corresponding minimum and maximum values are reported in Table C.2 of Appendix C.4.

Figure 6.15 summarizes the out-of-sample performance of the different storage policies. More detailed numbers are reported in Table C.2 and Table C.3 of Appendix C.4.

DDA versus RIA

We observe for the majority of the commodities that structured data-driven policies (DDA-OSP) can outperform RIA both significantly (see Table C.2) and consistently (see Table C.3) on our analyzed dataset: over all commodities, DDA-OSP with forward optimization strictly dominates RIA in 63.7% of the out-of-sample instances (Table C.3) with a median performance of 26.7% of the perfect foresight profit, while RIA achieves a median performance of 12.0% (Table C.2). For copper, crude oil and corn, the profit improvements relative to RIA are statistically significant at the 1% significance level.

For natural gas, RIA seems to be a comparatively strong policy that can hardly be improved by feature-based DDA policies. This is reasonable in markets that are particularly efficient, which is the case for the natural gas market compared to less efficient metal and agricultural markets (see, e.g., Kristoufek and Vosvrda, 2013). In highly efficient markets, all available information is already included in the futures prices. Another reasonable explanation is the extraordinary high volatility of gas prices ($\sigma = 65\%$) if compared to the other commodities ($\sigma = 17 - 37\%$) under consideration (see Figure 6.6). This might favor periodic reoptimization (RIA) that exclusively uses forward-looking information, while feature-based DDA is sensitive to structural breaks in the training history. However, we see that DDA policies can compete with RIA also for natural gas.

DDA with Forward Optimization versus DDA w/o Forward Optimization

We observe a positive value of considering forward optimization for the DDA policies. E.g., for DDA-OSP, forward optimization is valuable in 63.9% of the instances (Table C.3). Except for natural gas, the profit improvement of DDA-OSP through forward optimization is statistically significant at the 1% level.

DDA-OSP versus DDA-LDR

We further observe that the state-of-the-art LDR approach from the data-driven operations literature, which does not ensure policy consistency, performs poor for our constrained multi-stage stochastic optimization problem. DDA-OSP, which respects policy structure, strictly outperforms DDA-LDR in 77.8% (with forward optimization) and

73.6% (without forward optimization) of the instances respectively (Table C.3). The profit gain from DDA-LDR** to DDA-OSP** is statistically significant at the 1% level for all commodities under consideration, which indicates that there is a positive value of policy structure for data-driven optimization for our analyzed dataset.

DDA-OSP versus DDA-VFA

OLS-based DDA-VFA policies perform comparatively poor on the dataset considered. Besides the valid question whether an OLS model is the most appropriate one, another reasonable explanation seems to be that DDA-VFA, other than DDA-OSP, separates price prediction and optimization. Therefore, operational parameters, such as injection and withdrawal rates or frictions do not affect predictions. However, this may be important as the RIA results from Table C.1 of Appendix C.4 imply: identical forecasts can yield both unprofitable and profitable storage depending on the warehouse setting, which suggests to integrate prediction and optimization.

DDA-OSP-FF versus DDA-OSP-LF

Table 6.11 shows the disaggregated performance of DDA-OSP** from Figure 6.15 with respect to storage flexibility. Similar to the RIA results (Table C.1 of Appendix C.4), the results of DDA-OSP relative to the perfect foresight bound are not fundamentally different between fully flexible storage assets (SCWP-FF) and limited flexible storage assets (SCWP-LF), i.e., DDA-OSP performs well for both of the storage settings.

Table 6.11.: Performance of DDA-OSP in $V/V^{\mathrm{PF}} \cdot 100\%$ with respect to storage flexibility

	Mean	Min	25%-Q	50%-Q	75%-Q	Max
	SCWP-FF					
Copper	24.1	-115.0	12.1	35.9	58.4	78.6
Gold	17.4	-65.1	8.6	21.4	32.3	65.3
Crude Oil	13.3	-129.5	7.4	21.3	39.9	66.6
Natural Gas	9.0	-69.3	-21.4	17.7	40.0	59.5
Corn	35.0	-6.6	27.8	39.3	49.6	64.8
Soybean	28.8	0.9	18.1	30.2	38.2	52.9
Overall	**21.3**	**-129.5**	**8.9**	**26.6**	**42.9**	**78.6**
	SCWP-LF					
Copper	23.7	-188.2	5.9	33.7	70.2	83.6
Gold	6.3	-136.6	2.3	17.0	32.8	80.5
Crude Oil	18.0	-159.1	4.7	22.7	45.3	75.5
Natural Gas	5.0	-68.0	-34.5	1.9	39.7	71.5
Corn	43.6	-3.7	24.8	55.2	60.3	71.9
Soybean	37.5	0.0	19.2	45.5	53.6	71.1
Overall	**22.4**	**-188.2**	**7.1**	**12.0**	**36.9**	**83.6**

However, we observe for SCWP-FF that it is more effective to train a price threshold P_t (DDA-OSP-FF), rather than the more general double base-stock structure $(S_t^{\mathbf{i}}, S_t^{\mathbf{i}})$ (DDA-OSP-LF). While DDA-OSP-FF yields an average (median) performance of 21.3% (26.6%), DDA-OSP-LF for solving SCWP-FF yields 11.9% (22.8%). Furthermore, DDA-OSP-FF is more efficient and reduces computation times of DDA-OSP-LF (above $3,600$ seconds) by on average almost 90% over all backtest instances.

Regarding feature selection of DDA-OSP, for each commodity, features from all of the three feature categories (past spot prices, futures prices and analyst forecasts) were relevant for data-driven storage decisions for at least one sub-period - mostly in combination. This is in line with, e.g., Cortazar et al. (2018), who show that the combination of market prices and analyst forecasts is valuable in terms of forecast accuracy.

6.5. Conclusion

This chapter studies the well-known *Stochastic Commodity Warehouse Problem* (SCWP). For solving the problem, the reoptimization-based rolling intrinsic approach (RIA) is widely used in academia and practice due to its computational attractiveness and its near-optimality with respect to the *full information problem*.

However, we show that this is misleading if applied for real storage decisions in practice due to the inherent *generalization error* driven by the quality of futures prices in predicting spot prices. We make four observations that were not addressed by the literature, even though being crucial for storage managers: Under *generalization error*, *(i)* RIA can yield unprofitable storage operations, *(ii)* ignoring futures price information can be beneficial, *(iii)* the value of reoptimization is not necessarily positive, and *(iv)* the correct prediction of the direction (upward or downward) of the one-step-ahead price is essential for the performance of RIA.

To mitigate the adverse effects of *generalization error*, we propose data-driven and ML-based policies that systematically explore feature data. We find that these policies can outperform RIA without requiring to solve an LP at every decision stage and without determining the planning horizon. We furthermore show that the linear decision rule approach from the data-driven optimization literature is not effective for our multi-stage decision problem where structural results with regard to the optimal policy structure should be respected. Moreover, the standard ERM principle should be extended to include forward-looking information, which can improve the performance of DDA policies significantly.

Our results are valuable for both commodity storage managers in terms of challenging their current RIA-based decisions and data scientists in terms of sensitizing them that structural results are still valuable in data-rich contexts.

There are several limitations of our research: similar to, e.g., Secomandi (2010) and Wu et al. (2012), we do not analyze combined inventory trading at spot and forward markets. However, it might be interesting to analyze performance effects of different policies if the warehouse can trade simultaneously on both spot and futures markets. Furthermore, following the standard literature (e.g., Lai et al., 2010; Secomandi, 2010), we assume monthly inventory review periods. It might be interesting to analyze the performance effects on empirical data if injection and withdrawal decisions can be taken on a higher granularity (e.g., weekly or daily). Also the effects of periodic reoptimization of the training-based DDA policies might be worth investigating especially for natural gas where RIA performs comparatively well.

Chapter 7.

Conclusion

7.1. Summary of Insights

In many industries, volatile commodity prices constitute a significant exogenous risk factor. This thesis aimed at optimizing procurement and inventory policies in the presence of stochastic price fluctuations. As opposed to the existing literature, we relax the unrealistic assumption of full information about the underlying commodity price process.

First, we proved the optimal inventory control policy when prices and demand are uncertain and the prices follow a Markov regime switching process (MRS). The regime belief is learned based on recent price observations via a Bayesian updating scheme. The resulting state-dependent and learning-enabled base-stock policies are complex. Therefore, we tested various simpler but suboptimal control policies that ignore regime switches, learning or price uncertainty in general. Our numerical results suggest that inventory managers should care about sophisticated price models, such as MRS, especially when demand uncertainty is low and expected price changes are large relative to the inventory holding costs. In this case, ignoring Bayesian updating of regime beliefs can yield 13% higher costs.

Next, we studied a multi-stage forward contracting problem with the optimal purchase signals derived in a data-driven way as functions of price features. Employing the principle of ERM from statistical learning theory, we trained the threshold parameters of the procurement policy based on historical price and feature data with a MILP model. To support performance-based feature selection and to avoid overfitting, we extended our math programs to include regularization. Our results with regard to reduced procurement costs show that there is a significant value in combining interpretable machine learning with problem-specific mathematical optimization, in order to extract decision-relevant data (i.e., *Smart Data*) from noise. We offer decision rules that are simple to operationalize and simple to interpret, as well as effective and easily accessible to managers in practice.

Last, we considered optimal commodity storage from a merchant's perspective with buying and selling opportunities. In practice, reoptimization via RIA is the state-of-the-art for solving this optimization problem. However, in extensive backtests based on six major commodities, we find that *generalization error* can lead to significantly suboptimal RIA policies, contrary to their known near-optimality in the full information setting. To reduce *generalization error*, we developed different classes of non-parametric data-driven approaches that leverage machine learning and optimization and yield improved commodity storage policies based on structured decision rules.

7.2. Directions for Future Research

In all of our problem settings, the firm (a commodity-purchasing firm in *Chapters 4* and *5* and a commodity-trading firm in *Chapter 6*) is assumed to be *(i)* risk-neutral and *(ii)* a price taker, i.e., prices are taken as exogenous. It might be worth investigating both the impact of risk aversion and the impact of market power on policy results and performance. We also assume that our firm is a stand-alone firm. It might be interesting to study the effects of price risk in supply chain or competition settings.

Furthermore, we either study operational hedging of commodity price risk via optimized inventory control (*Chapters 4 and 6*) or via financial hedging through optimized forward contracting decisions (*Chapter 5*). It might be interesting to study settings where both options exist simultaneously.

Moreover, following the standard literature, we do not explicitly consider the effects of correlation structures between price and demand in any of the chapters. However, this is only reasonable if *(i)* commodity prices make up only a small amount of the final product of the firm and the firm needs the input commodity regardless of the price or *(ii)* the firm cannot pass higher or lower input prices on to the customer. If a firm can pass its commodity prices on to the customer, there is a customer's demand reaction on prices. Therefore, it might be valuable to study the effects of price-demand correlation on both policy structure and performance.

Furthermore, the results on empirical data in *Chapters 4* and *5* are based on single case studies rather than extensive empirical investigations. They need justification on a broader range of commodities.

Appendix

Appendix A.

Appendix of Chapter 4

A.1. Proof of Theorem 1

In this section, we present the proof for the more general Markovian case. The proof under i.i.d. price processes is straightforward to the Markovian case, however without p_{t-1} as a state variable in the dynamic programming equation (4.2) as π^s_{t+1} is a function of $(p_t, \vec{\pi}_t)$, rather than $(p_t, p_{t-1}, \vec{\pi}_t)$ according to equation (4.1). Consequently, for the i.i.d. case, S_t is fully characterized by p_t and the prior regime belief $\vec{\pi}_t$.

Lemma 1. *$C_t(z_t)$ is a convex function of I_t for all p_t, p_{t-1} and $\vec{\pi}_t$ respectively p_t and $\vec{\pi}_{t+1}$.*

Proof. Equivalently to equation (4.2),

$$C_t(z_t) = \min_{I_t^* \geq I_t} \left\{ g_t\big((I_t^* - I_t)|p_t\big) + \mathbb{E}_{d_{t+1}, p_{t+1}|(p_t, \underbrace{p_{t-1}, \vec{\pi}_t)}_{\vec{\pi}_{t+1}}} \big[C_{t+1}(z_{t+1})\big] \right\}, \; t = 1, ..., n \quad \text{(A.1)}$$

where

$$g_t\big((I_t^* - I_t)|p_t\big) = p_t(I_t^* - I_t)^+ + \mathbb{E}_{d_t}\big[c_h(I_t^* - d_t)^+ + c_p(d_t - I_t^*)^+\big]. \quad \text{(A.2)}$$

The optimality equation (A.1) can be rewritten as

$$C_t(z_t) = \min_{I_t^* \geq I_t} \left\{ G_t\big(I_t^*|(p_t, \underbrace{p_{t-1}, \vec{\pi}_t}_{\vec{\pi}_{t+1}})\big) - p_t I_t \right\} = -p_t I_t + \min_{I_t^* \geq I_t} \left\{ G_t\big(I_t^*|(p_t, \underbrace{p_{t-1}, \vec{\pi}_t}_{\vec{\pi}_{t+1}})\big) \right\} \quad \text{(A.3)}$$

where

$$G_t\big(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t)\big) := p_t I_t^* + \mathbb{E}_{d_t}\big[c_h(I_t^* - d_t)^+ + c_p(d_t - I_t^*)^+\big] + \mathbb{E}_{d_{t+1},p_{t+1}|(p_t,p_{t-1},\vec{\pi}_t)}\big[C_{t+1}(z_{t+1})\big]. \tag{A.4}$$

Thus, the optimal decision starting with inventory I_t in period t is found by minimizing $G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t))$ over $\{I_t^*|I_t^* \geq I_t\}$. Because $C_t(z_t) = -p_t I_t + \min_{I_t^* \geq I_t}\{G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t))\}$, we need to proof that $G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t))$ is convex in I_t^*. Because $G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t))$ is the sum of three functions, $G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t))$ is convex in I_t^* if the three functions are convex in I_t^*. The first function $p_t \cdot I_t^*$ is linear and therefore convex in I_t^*. As proven by Porteus (2002, pp. 67), the second and third terms are convex, too. Therefore, $G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t))$ is convex in I_t^* and hence $C_t(z_t)$ is convex in I_t^*. Note: From equation (4.1) follows that $G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t)) = G_t(I_t^*|(p_t, \vec{\pi}_{t+1}))$ as p_{t-1} is solely required for determining $\vec{\pi}_{t+1}$. □

Lemma 2. *Any minimizer of G_t is an optimal base-stock level.*

Proof. Let $S_t(p_t, p_{t-1}, \vec{\pi}_t)$ denote the value over all real I_t^* that minimize $G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t))$. If $I_t < S_t(p_t, p_{t-1}, \vec{\pi}_t)$, then the optimal $I_t^* \geq I_t$ is at $I_t^* = S_t(p_t, p_{t-1}, \vec{\pi}_t)$, i.e., an amount $y_t = S_t(p_t, p_{t-1}, \vec{\pi}_t) - I_t$ is ordered in period t and the expected total cost $C_t(z_t)$ is written as

$$\begin{aligned} C_t(z_t) =& p_t S_t(p_t, p_{t-1}, \vec{\pi}_t) \\ &+ \mathbb{E}_{d_t}\Big[c_h(S_t(p_t, p_{t-1}, \vec{\pi}_t) - d_t)^+ + c_p\big(d_t - S_t(p_t, p_{t-1}, \vec{\pi}_t)\big)^+\Big] \\ &+ \mathbb{E}_{d_{t+1},p_{t+1}|(p_t,p_{t-1},\vec{\pi}_t)}\big[C_{t+1}(z_{t+1})\big]. \end{aligned} \tag{A.5}$$

If $I_t \geq S_t(p_t, p_{t-1}, \vec{\pi}_t)$, then $G_t(I_t^*|(p_t, p_{t-1}, \vec{\pi}_t))$ is non-decreasing to the right of I_t (by convexity) and thus the optimal $I_t^* \geq I_t$ is $I_t^* = I_t$, i.e., we will not order in period t and consequently the expected total cost $C_t(z_t)$ is written as

$$C_t(z_t) = \mathbb{E}_{d_t}\big[c_h(I_t - d_t)^+ + c_p(d_t - I_t)^+\big] + \mathbb{E}_{d_{t+1},p_{t+1}|(p_t,p_{t-1},\vec{\pi}_t)}\big[C_{t+1}(z_{t+1})\big]. \tag{A.6}$$

$C_t(z_t)$ is convex for all I_t since it is convex in I_t for $I_t \geq S_t(p_t, p_{t-1}, \vec{\pi}_t)$ (see equation (A.6)) and it is convex in I_t for $I_t < S_t(p_t, p_{t-1}, \vec{\pi}_t)$ (see equation (A.5)). Note: From equation (4.1) follows that $S_t(p_t, p_{t-1}, \vec{\pi}_t) = S_t(p_t, \vec{\pi}_{t+1})$, i.e., S_t is fully characterized by p_t and the posterior regime belief $\vec{\pi}_{t+1}$. □

By means of Lemma 1 and Lemma 2, we can prove Theorem 1.

Proof. The base-stock policy from equation (4.4) is optimal since $C_t(z_t)$, as defined in equation (A.1), is a convex function of I_t (see also Porteus 2002, pp. 67). $\square$

A.2. Proof of Proposition 1

Part(i)

Proof. The proof follows Kalymon (1971). If p_{t+1} follows an i.i.d. price process $\phi^s(p_{t+1}|p_t) = \phi^s(p_{t+1}) \quad \forall s \in M$ and the regime belief $\vec{\pi}_t$ is not updated based on p_t, then $\phi(p_{t+1}|p_t) = \phi(p_{t+1})$. Therefore, equation (4.2) reduces to

$$C_t(z_t) = \min_{I_t^* \geq I_t} \left\{ p_t y_t + L(I_t^*) + \int_0^\infty \int_0^\infty C_{t+1}\big((I_{t+1}^* - d_{t+1})^+, p_{t+1}\big) d\phi(p_{t+1}) dF(d_{t+1}) \right\}. \quad \text{(A.7)}$$

We define

$$Q(I_t^*) \coloneqq L(I_t^*) + \int_0^\infty \int_0^\infty C_{t+1}\big((I_{t+1}^* - d_{t+1})^+, p_{t+1}\big) d\phi(p_{t+1}) dF(d_{t+1}). \quad \text{(A.8)}$$

Therefore,

$$G(I_t^*|p_t) = p_t I_t^* + Q(I_t^*). \quad \text{(A.9)}$$

If $p_t' > p_t$, then

$$G(I_t^*|p_t') = p_t' I_t^* + Q(I_t^*) = p_t I_t^* + Q(I_t^*) + (p_t' - p_t) I_t^* = G(I_t^*|p_t) + (p_t' - p_t) I_t^*. \quad \text{(A.10)}$$

Suppose that $S_t(p_t') > S_t(p_t)$. Then from $G(S_t|p_t) = \min_{I_t^*}\{G(I_t^*|p_t)\}$, it follows that:

$$\begin{aligned} G\big(S_t(p_t')|p_t\big) + (p_t' - p_t) S_t(p_t') &= G\big(S_t(p_t')|p_t'\big) \\ \leq G\big(S_t(p_t)|p_t'\big) &= G\big(S_t(p_t)|p_t\big) + (p_t' - p_t) S_t(p_t). \end{aligned} \quad \text{(A.11)}$$

Since $(p_t' - p_t) S_t(p_t') > (p_t' - p_t) S_t(p_t)$, we get $G(S_t(p_t')|p_t) < G(S_t(p_t)|p_t)$, which contradicts the optimality of $S_t(p_t)$. Thus, if $p_t' > p_t$, then $S_t(p_t') \leq S_t(p_t)$. $\square$

Part(ii)

Proof. For the unobservable case with dynamic information updates (*learning*), due to equation (4.1), $\phi(p_{t+1}|p_t) \neq \phi(p_{t+1})$ and hence *(a1)* is violated. If the Bayesian relationship between p_t and p_{t+1} does not violate either *(a2.1)* or *(a2.2)* or both, S_t is non-increasing in p_t. The proof is analogous to the observable case discussed, e.g., by Gavirneni (2004). However, Bayes' theorem (4.1) can violate *(a2.1)* if the predicted price increase at higher prices is greater than at lower prices. A high price leads to an increase in the belief about being in the high price regime in subsequent periods. In volatility regimes $(i, j) \in \{1, ..., m\}$ with $\mathbb{E}_i[p_{t+1}] = \mathbb{E}_j[p_{t+1}] \quad \forall (i, j) \in \{1, ..., m\}$ and $Var_i(p_{t+1}) \neq Var_j(p_{t+1})$, *(a2.1)* is not violated as the expected price is equal in each regime and hence not expected to increase at higher prices.

Let the price in each regime s follow a distribution with mean μ_s and variance σ_s^2. Mean and variance of the mixture $\phi(p_{t+1})$ are defined as $\mathbb{E}[p_{t+1}] = \sum_{s=1}^{m} \pi_{t+1}^s \mu_s$ and $Var(p_{t+1}) = \mathbb{E}[p_{t+1}^2] - (\mathbb{E}[p_{t+1}])^2$. For $m = 2$, $Var(p_{t+1}) = \pi_{t+1}^{(1)} \sigma_1^2 + \pi_{t+1}^{(2)} \sigma_2^2 + \pi_{t+1}^{(1)} \pi_{t+1}^{(2)} (\mu_1 - \mu_2)^2$. If $\mu_1 = \mu_2$ (unimodal mixture), the variance is a linear combination of the regime variances. Consequently, if all ϕ^s fulfill the monotonicity conditions, this is also true for the unobservable case and S_t is non-increasing in p_t. In an m-regime setting, as $k_{ii} \to 1/m$, according to equation (4.1), $\pi_{t+1}^i \to 1/m$ and $\pi_{t+1}'^i \to 1/m$. With convergence, equation (4.5) reduces to $-p_t' \leq -p_t$, which is true by definition. Therefore, *(a2.1)* holds. The reasoning is that since $k_{ii} \to 1/m$, information p_t is worth less than if $k_{ii} \to 1$ or $k_{ii} \to 0$ because the conditional regime forecast based on p_t is less reliable. Hence, one would not order more at higher prices speculating to be in the high price regime in the future. □

Part(iii)

Proof. As the price processes in the regimes ϕ^s are Markovian, condition *(a1)* is violated by definition since $\phi^s(p_{t+1}|p_t) \neq \phi^s(p_{t+1}) \quad \forall s \in M$. S_t is non-increasing in p_t if $\phi(p_{t+1}) = \sum_{s=1}^{m} \pi_t^s \phi^s(p_{t+1})$ satisfies both *(a2.1)* and *(a2.2)*. For the detailed proof, see, e.g., Gavirneni (2004). □

Part(iv)

Proof. Condition *(a1)* is violated by definition since $\phi^s(p_{t+1}|p_t) \neq \phi^s(p_{t+1}) \quad \forall s \in M$. S_t is non-increasing in p_t if $\phi(p_{t+1}) = \sum_{s=1}^{m} \pi_{t+1}^s \phi^s(p_{t+1})$ satisfies *(a2.1)* and *(a2.2)*. The proof is analogous to Proposition 1(iii). It is sufficient to consider exclusively π_{t+1}^s, disregarding π_{t+l}^s with $l > 1$, since one would only buy more at higher prices if p_{t+1} is expected to be greater than $p_t + c_h$, independent of any expected price increase in periods after the next one (i.e., $l > 1$). □

A.3. Cost of Price Regime Misspecification

Table A.1.: $\Delta\text{COST}(\phi^{(1)}(\phi^{(2)}))$ in % (In gray: No speculation motives)

			Regime type																	
			i.i.d.				Markovian						Combination							
c_h	Demand volatility		LL(HL)	HL(LL)	LV(HV)	HV(LV)	RW(MR)	MR(RW)	RW(MO)	MO(RW)	MR(MO)	MO(MR)	HL(MR)	MR(HL)	LL(MR)	MR(LL)	HL(MO)	MO(HL)	LL(MO)	MO(LL)
1	high	Mean	5.45	4.94	0.11	0.11	0.59	0.53	0.70	0.48	2.36	2.40	1.63	1.09	1.57	0.93	10.17	5.26	5.07	1.92
		StDev	3.53	3.88	0.13	0.14	0.84	0.63	0.93	0.61	2.73	2.87	1.44	1.02	0.94	0.46	13.58	5.94	4.73	1.38
		Min	1.68	0.85	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.02	0.18	0.18	0.55	0.40	0.19	0.24
		Max	**10.96**	**12.61**	**0.32**	**0.33**	**2.00**	**1.39**	**2.20**	**1.50**	**6.81**	**6.71**	**3.01**	**2.53**	**2.54**	**1.39**	**32.24**	**14.15**	**11.95**	**3.64**
	medium	Mean	7.87	7.84	0.24	0.27	0.94	0.89	1.55	0.91	4.86	4.45	1.85	1.19	1.83	0.96	17.42	6.65	8.25	1.35
		StDev	8.44	8.99	0.44	0.46	1.78	1.51	2.01	1.24	5.14	6.59	2.86	1.19	2.71	1.46	21.06	8.03	11.83	1.35
		Min	0.49	0.24	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.44	0.17	0.24	0.12	0.31	0.20
		Max	**22.74**	**24.99**	**1.32**	**1.25**	**4.11**	**3.56**	**4.22**	**2.82**	**11.94**	**15.77**	**6.79**	**4.56**	**6.66**	**3.56**	**49.91**	**20.32**	**28.23**	**3.15**
	zero	Mean	9.65	10.39	0.26	0.53	1.46	0.79	2.29	1.32	6.75	6.58	1.66	1.62	2.62	0.79	25.30	8.37	10.11	1.32
		StDev	12.71	14.04	0.75	1.06	3.27	1.77	3.28	2.17	8.30	11.21	3.72	3.62	5.86	1.77	33.68	10.59	15.32	2.17
		Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
		Max	**35.00**	**33.10**	**2.26**	**2.61**	**7.32**	**3.95**	**7.07**	**5.00**	**19.57**	**26.29**	**8.32**	**8.09**	**13.11**	**3.95**	**80.78**	**26.29**	**34.55**	**5.00**
6	high	Mean	2.90	1.89	0.01	0.03	0.08	0.08	0.07	0.07	0.26	0.33	1.62	1.77	0.47	0.26	2.35	4.30	1.25	0.62
		StDev	1.88	1.59	0.02	0.04	0.10	0.08	0.05	0.06	0.21	0.35	2.54	2.24	0.36	0.17	3.21	5.59	1.32	0.51
		Min	0.63	0.31	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.02	0.04	0.00	0.13	0.17	0.07	0.08
		Max	**5.67**	**5.31**	**0.05**	**0.10**	**0.24**	**0.21**	**0.14**	**0.17**	**0.51**	**0.87**	**6.06**	**5.16**	**0.96**	**0.42**	**7.61**	**12.06**	**3.01**	**1.39**
	medium	Mean	5.24	2.98	0.01	0.01	0.04	0.06	0.04	0.06	0.16	0.16	3.54	3.23	0.12	0.09	4.08	9.13	0.60	0.37
		StDev	5.26	4.23	0.02	0.02	0.06	0.09	0.04	0.08	0.13	0.18	5.65	4.48	0.12	0.06	6.44	13.01	0.63	0.33
		Min	0.27	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.11	0.01	0.03	0.15
		Max	**12.29**	**12.10**	**0.04**	**0.05**	**0.15**	**0.20**	**0.11**	**0.19**	**0.34**	**0.44**	**13.14**	**10.23**	**0.27**	**0.15**	**15.02**	**28.84**	**1.57**	**0.94**
	zero	Mean	6.56	3.39	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	4.20	3.43	0.00	0.00	4.20	10.42	0.00	0.00
		StDev	8.41	5.57	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	7.09	5.35	0.00	0.00	7.09	15.86	0.00	0.00
		Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
		Max	**22.03**	**16.09**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**16.36**	**12.19**	**0.00**	**0.00**	**16.36**	**35.84**	**0.00**	**0.00**
15	high	Mean	0.69	0.55	0.00	0.14	0.02	0.03	0.01	0.02	0.07	0.10	0.34	0.44	0.16	0.13	0.56	0.94	0.34	0.26
		StDev	0.30	0.26	0.01	0.08	0.02	0.02	0.02	0.03	0.07	0.09	0.44	0.52	0.14	0.10	0.68	1.16	0.37	0.22
		Min	0.27	0.18	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.03	0.03	0.03	0.07
		Max	**1.12**	**1.05**	**0.02**	**0.21**	**0.06**	**0.06**	**0.04**	**0.07**	**0.16**	**0.24**	**1.06**	**1.24**	**0.31**	**0.23**	**1.51**	**2.45**	**0.90**	**0.60**
	medium	Mean	0.83	0.32	0.00	0.00	0.00	0.00	0.01	0.01	0.03	0.01	0.26	0.99	0.12	0.03	0.39	1.77	0.16	0.07
		StDev	1.27	0.33	0.00	0.00	0.00	0.00	0.03	0.02	0.07	0.02	0.41	1.96	0.11	0.03	0.69	3.27	0.09	0.07
		Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
		Max	**4.08**	**0.90**	**0.00**	**0.00**	**0.00**	**0.00**	**0.06**	**0.05**	**0.15**	**0.05**	**0.93**	**4.48**	**0.22**	**0.06**	**1.59**	**7.53**	**0.23**	**0.18**
	zero	Mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
		StDev	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
		Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
		Max	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**

Note. MO(MR) denotes that MO is the true underlying price process, whereas MR is the supposed price process. Example: Inventory decisions based on MR in an MO regime in a speculation-friendly environment ($c_h = 1$) with deterministic demand yields a cost increase by up to 26.29%.

A.4. Performance of Suboptimal Control Policies

Table A.2.: Performance of suboptimal control policies in % above optimal cost

				$c_h = 1$			$c_h = 6$		
Heuristic	Regime setting	Demand volatility		r=0.25	r=0.40	r=0.49	r=0.25	r=0.40	r=0.49
NAIVE	HL-LL	high	Mean	2.78	1.81	1.29	0.43	0.26	0.20
			StDev	4.44	3.26	2.38	0.49	0.34	0.26
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**14.21**	**10.51**	**7.76**	**1.59**	**1.08**	**0.81**
		zero	Mean	5.93	3.93	2.55	0.27	0.02	0.00
			StDev	11.37	8.32	5.60	0.78	0.05	0.00
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**34.83**	**26.10**	**17.86**	**2.44**	**0.17**	**0.00**
	LV-HV	high	Mean	4.30	4.18	4.11	0.85	0.84	0.83
			StDev	5.59	5.39	5.24	0.85	0.83	0.81
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**18.12**	**17.52**	**17.11**	**2.90**	**2.80**	**2.73**
		zero	Mean	8.28	7.94	7.71	0.92	0.92	0.92
			StDev	14.06	13.47	13.06	1.98	1.99	2.00
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**42.61**	**40.99**	**39.87**	**6.14**	**6.16**	**6.17**
	MR-MO	high	Mean	0.56	0.54	0.54	0.08	0.08	0.08
			StDev	1.15	1.19	1.21	0.10	0.10	0.10
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**5.42**	**5.78**	**5.99**	**0.47**	**0.47**	**0.47**
		zero	Mean	1.52	1.48	1.42	0.00	0.00	0.00
			StDev	3.60	3.73	3.70	0.00	0.00	0.00
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**16.32**	**17.54**	**17.61**	**0.00**	**0.00**	**0.00**
SRC-R1/-R2	HL-LL	high	Mean	1.43/0.86	1.88/0.77	2.54/0.85	1.12/0.25	1.58/0.29	2.00/0.34
			StDev	1.36/0.34	2.12/0.68	2.93/0.95	1.11/0.16	1.75/0.30	2.30/0.37
			Min	0.13/0.30	0.01/0.07	0.00/0.00	0.04/0.06	0.00/0.00	0.00/0.00
			Max	**5.53/1.64**	**7.34/2.80**	**9.01/3.78**	**2.90/0.65**	**4.31/0.95**	**5.51/1.17**
		zero	Mean	2.49/1.26	3.18/0.51	4.69/0.22	4.69/0.27	4.56/0.02	6.09/0.00
			StDev	4.30/2.57	5.59/1.12	7.55/0.74	3.69/0.78	5.66/0.05	7.59/0.00
			Min	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00
			Max	**17.67/7.99**	**22.69/3.45**	**28.09/4.65**	**9.93/2.44**	**15.92/0.17**	**21.31/0.00**
	LV-HV	high	Mean	0.04/0.02	0.06/0.02	0.07/0.02	0.01/0.00	0.02/0.00	0.02/0.00
			StDev	0.05/0.03	0.07/0.04	0.09/0.05	0.02/0.01	0.03/0.01	0.04/0.01
			Min	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00
			Max	**0.15/0.13**	**0.21/0.19**	**0.27/0.24**	**0.05/0.02**	**0.08/0.03**	**0.09/0.03**
		zero	Mean	0.17/0.00	0.20/0.00	0.23/0.02	0.00/0.00	0.00/0.00	0.00/0.00
			StDev	0.38/0.00	0.47/0.00	0.56/0.10	0.00/0.00	0.00/0.00	0.00/0.00
			Min	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00
			Max	**1.64/0.00**	**1.90/0.00**	**2.06/0.68**	**0.00/0.00**	**0.00/0.00**	**0.00/0.00**
	MR-MO	high	Mean	0.68/0.69	0.76/0.68	0.87/0.68	0.10/0.07	0.11/0.08	0.12/0.08
			StDev	1.07/1.07	1.24/1.14	1.42/1.18	0.14/0.09	0.16/0.10	0.18/0.12
			Min	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00
			Max	**5.17/5.02**	**5.30/5.37**	**6.07/5.58**	**0.70/0.40**	**0.70/0.41**	**0.81/0.46**
		zero	Mean	1.61/1.70	1.70/1.62	2.02/1.64	0.00/0.00	0.00/0.00	0.00/0.00
			StDev	3.94/2.66	4.37/2.79	4.88/3.01	0.00/0.00	0.00/0.00	0.00/0.00
			Min	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00	0.00/0.00
			Max	**18.51/10.45**	**19.19/11.67**	**22.73/11.74**	**0.00/0.00**	**0.00/0.00**	**0.00/0.00**
CEC	HL-LL	high	Mean	0.33 (0.59)	0.55 (1.44)	0.92 (2.10)	0.28 (1.05)	0.47 (1.62)	0.65 (2.02)
			StDev	0.40	0.73	1.18	0.52	0.86	1.15
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**1.55**	**2.80**	**4.85**	**2.51**	**3.79**	**4.89**
		zero	Mean	0.40 (1.32)	0.56 (3.05)	1.26 (4.35)	0.54 (3.17)	0.90 (5.32)	1.34 (5.49)
			StDev	1.00	1.48	2.68	1.46	2.24	3.04
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**4.79**	**6.03**	**10.97**	**5.50**	**8.49**	**13.33**
	LV-HV	high	Mean	0.01 (-)	0.01 (-)	0.02 (-)	0.00 (-)	0.00 (-)	0.01 (-)
			StDev	0.01	0.02	0.03	0.01	0.02	0.02
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**0.08**	**0.12**	**0.15**	**0.05**	**0.08**	**0.09**
		zero	Mean	0.01 (-)	0.01 (-)	0.02 (-)	0.00 (-)	0.00 (-)	0.00 (-)
			StDev	0.05	0.08	0.11	0.00	0.00	0.00
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**0.32**	**0.54**	**0.71**	**0.00**	**0.00**	**0.00**
	MR-MO	high	Mean	0.05 (0.36)	0.03 (0.27)	0.04 (0.34)	0.00 (-)	0.01 (-)	0.01 (-)
			StDev	0.14	0.10	0.11	0.01	0.02	0.02
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**1.19**	**0.88**	**0.71**	**0.06**	**0.10**	**0.10**
		zero	Mean	0.11 (0.72)	0.05 (0.77)	0.07 (1.02)	0.00	0.00	0.00
			StDev	0.39	0.22	0.28	0.00	0.00	0.00
			Min	0.00	0.00	0.00	0.00	0.00	0.00
			Max	**2.87**	**1.99**	**2.13**	**0.00**	**0.00**	**0.00**

Note. The given numbers demonstrate the value of a MRS price model relative to simple models in the controlled numerical study from Section 4.4. In parentheses: Instances with speculation under MRS and non-speculation under CEC or vice versa (Non-speculation as defined in equation (4.7)).

Appendix B.

Appendix of Chapter 5

B.1. Proof of Theorem 2

In the following, we show that a price threshold $P_t^\tau(x_t)$ fully characterizes the optimal procurement policy, which is of a bang-bang type using a particular option in a period, i.e., we either procure all uncovered demand or nothing with a contract $\tau \in \mathcal{F}^+$.

However, the state $z_t \in Z_t$ is not fully known. We only know that z_t contains the firm's position in the forward market $\vec{I} = (I_t^\tau)$ and the current forward curve $\vec{F}_t = (p_t^\tau : \tau \geq 0)$. However, without price model specifications, we do not know the drivers of the evolution of $\vec{F}$. Therefore, we introduce $x_t \in \mathcal{X}$ as the unknown parts (features) of the state space $z_t = (\vec{I}_t, \vec{F}_t, x_t)$ that drives the evolution of $\vec{F}$ and must be learned from the data $\mathbf{X}$ for which we formulate DDA models. In order to provide the policy structure, we formulate the problem as a standard SDP with the endogenous state transition $I_{t+1}^{\tau-1} = I_t^\tau + y_t^\tau$ and the exogenous price evolution $\vec{F}_{t+1} = \phi(\vec{F}_t, x_t)$. This (unknown) exogenous price transition replaces typical stochastic price processes in models where they are assumed.

$$C_t(\vec{I}_t, \vec{F}_t, x_t) = \min_{\substack{y_t^\tau \geq 0 \\ I_t^0 + y_t^0 \geq d_t}} \left\{ \sum_{\tau \in \mathcal{F}} p_t^\tau y_t^\tau + \mathbb{E}_t \left[C_{t+1}(\vec{I}_{t+1}, \vec{F}_{t+1}, x_{t+1}) \right] \right\} \qquad \forall t = 0, ..., n. \tag{B.1}$$

For every period t, we prove that for all $\tau \in \mathcal{F}^+$,

$$y_t^\tau(x_t) = \begin{cases} d_{t+\tau} & \text{if } p_t^\tau \leq P_t^\tau(x_t) \text{ and } I_t^\tau = 0, \\ 0 & \text{if } p_t^\tau > P_t^\tau(x_t). \end{cases} \tag{B.2}$$

The proof exploits two properties: The value function of a given period is separable in the procurement instruments and linear with regard to the quantities y_t^τ.

Period t=n.

We start with the last period $t = n$, which contains a non-stochastic purchase decision y_n^0 with

$$C_n(\vec{I}_n, \vec{F}_n, x_n) = C_n(\vec{I}_n, p_n^0) = p_n^0 y_n^0, \tag{B.3}$$

that yields

$$y_n^0 = [d_n - I_n^0]^+. \tag{B.4}$$

Period t=n-1.

At the second to last stage $t = n-1$, the problem becomes stochastic, since procurement decisions y_{n-1}^τ affect the cost-to-go C_n of period $t = n$:

$$C_{n-1}(\vec{I}_{n-1}, \vec{F}_{n-1}, x_{n-1}) = \min_{\substack{y_{n-1}^1 \geq 0 \\ I_{n-1}^0 + y_{n-1}^0 \geq d_{n-1}}} \left\{ p_{n-1}^0 y_{n-1}^0 + p_{n-1}^1 y_{n-1}^1 + (d_n - y_{n-1}^1 - I_{n-1}^1)\mathbb{E}_{n-1}[p_n^0 | x_{n-1}] \right\}. \tag{B.5}$$

This function is linear and separable in y_{n-1}^0 and y_{n-1}^1. The two decisions to be taken in $t = n - 1$ are *(i)* the spot purchase decision

$$y_{n-1}^0 = [d_{n-1} - I_{n-1}^0]^+ \tag{B.6}$$

and *(ii)* the forward purchase decision

$$y_{n-1}^1(x_{n-1}) = \begin{cases} [d_n - I_{n-1}^1]^+ & \text{if } p_{n-1}^1 \leq \mathbb{E}_{n-1}[p_n^0 | x_{n-1}], \\ 0 & \text{if } p_{n-1}^1 > \mathbb{E}_{n-1}[p_n^0 | x_{n-1}], \end{cases} \tag{B.7}$$

which is driven by whether the derivative with regard to y_{n-1}^1 is positive or negative. Consequently, the implied threshold is

$$P_{n-1}^1(x_{n-1}) = \mathbb{E}_{n-1}[p_n^0 | x_{n-1}], \tag{B.8}$$

and the functional value is again separable and linear in the remaining demands:

$$\begin{aligned} C_{n-1}(\vec{I}_{n-1}, \vec{F}_{n-1}, x_{n-1}) &= (d_{n-1} - I_{n-1}^0)p_{n-1}^0 \\ &+ (d_n - I_{n-1}^1)p_{n-1}^1 \cdot 1_{p_{n-1}^1 \leq P_{n-1}^1} \\ &+ (d_n - I_{n-1}^1)\mathbb{E}_{n-1}[p_n^0 | x_{n-1}] \cdot 1_{p_{n-1}^1 > P_{n-1}^1} \end{aligned} \tag{B.9}$$

with 1 as the indicator function.

Periods t=0,...,n-2.

For all stages $t = 0, ..., n-2$, the linearity of the period cost and the cost-to-go in the decision variables reinforces the all-or-nothing decisions for using a procurement instrument depending on the state-dependent thresholds $P_t^\tau(x_t)$ and again results in the linearity of the resulting value function.

$$P_{n-1}^1(x_{n-1}) = \mathbb{E}_{n-1}[p_n^0|x_{n-1}] \tag{B.10}$$

The functional value is again separable and linear in the remaining uncovered demands:

$$\begin{aligned} C_{n-2}(\vec{I}_{n-2}, \vec{F}_{n-2}, x_{n-2}) = (d_{n-2} - I_{n-2}^0)p_{n-2}^0 \\ +(d_{n-1} - I_{n-2}^1)p_{n-2}^1 \cdot 1_{p_{n-2}^1 \leq P_{n-2}^1} \\ +(d_{n-1} - I_{n-2}^1)\mathbb{E}_{n-2}[p_{n-1}^0|x_{n-2}] \cdot 1_{p_{n-2}^1 > P_{n-2}^1} \\ +(d_n - I_{n-2}^2)p_{n-2}^2 \cdot 1_{p_{n-2}^2 \leq P_{n-2}^2} \\ +(d_n - I_{n-2}^2)\mathbb{E}_{n-2}[p_{n-1}^1|x_{n-2}] \cdot 1_{p_{n-2}^2 > P_{n-2}^2, p_{n-1}^1 \leq P_{n-1}^1} \\ +(d_n - I_{n-2}^2)\mathbb{E}_{n-2}[p_n^0|x_{n-2}] \cdot 1_{p_{n-2}^2 > P_{n-2}^2, p_{n-1}^1 > P_{n-1}^1}. \end{aligned} \tag{B.11}$$

The two decisions to be taken in $t = 0, ..., n-2$ are again *(i)* the spot purchase decision

$$y_t^0 = [d_t - I_t^0]^+ \tag{B.12}$$

and *(ii)* the forward purchase decision

$$y_t^\tau(x_t) = \begin{cases} [d_{t+\tau} - I_t^\tau]^+ & \text{if } p_t^\tau \leq P_t^\tau(x_t), \\ 0 & \text{if } p_t^\tau > P_t^\tau(x_t). \end{cases} \tag{B.13}$$

The resulting policy is a bang-bang-type policy. The optimal hedging policy uses the expected cheapest source, which is in accordance with Smith and Stulz (1985), who show that partial hedging that leaves the hedger exposed to some residual price risk is not cost-optimal.

B.2. Model Formulation for the Best Subset Selection Problem

DDA-BSSP:

$$\min_{\beta_i^\tau \in \mathcal{B}} \quad \hat{C}^{\text{BSSP}} = \frac{1}{T}\sum_{t=1}^{T} \sum_{\tau \in \mathcal{F} | \tau \leq T-t} \left[p_t^\tau \; d_{t+\tau} \; q_t^\tau \right] \tag{B.14}$$

$$\text{s.t.} \quad (5.6) - (5.9) \tag{B.15}$$

$$Mw_i^\tau \geq \beta_i^\tau \qquad \forall i = 1, ..., N; \tau \in \mathcal{F}^+ \tag{B.16}$$

$$- Mw_i^\tau \leq \beta_i^\tau \qquad \forall i = 1, ..., N; \tau \in \mathcal{F}^+ \tag{B.17}$$

$$\sum_{i=1}^{N} w_i^\tau \leq \bar{N} \qquad \forall \tau \in \mathcal{F}^+ \tag{B.18}$$

$$w_i^\tau \in \{0, 1\} \qquad \forall i = 1, ..., N; \tau \in \mathcal{F}^+ \tag{B.19}$$

B.3. Model Formulation with Indicator Constraints

DDA-BD (Indicator Constraints):

$$\min_{\beta_i^\tau \in \mathcal{B}} \hat{C}^{\text{BD}} = \frac{1}{T}\sum_{t=1}^{T} \sum_{\tau \in \mathcal{F} | \tau \leq T-t} [p_t^\tau \; d_{t+\tau} \; q_t^\tau] \tag{B.20}$$

$$\text{s.t.} \sum_{\tau \in \mathcal{F} | \tau \leq t-1} q_{t-\tau}^\tau = 1 \qquad \forall t = 1, ..., T \tag{B.21}$$

$$\text{Indicator ct:} \;\; aux_t^\tau = 0 \;\; \text{if} \;\; \sum_{i=0}^{N} \beta_i^\tau X_{it} < p_t^\tau \qquad \forall \tau \in \mathcal{F}^+; t = 1, ..., T-\tau \tag{B.22}$$

$$\text{Indicator ct:} \;\; aux_t^\tau = 1 \;\; \text{if} \;\; \sum_{i=0}^{N} \beta_i^\tau X_{it} \geq p_t^\tau \qquad \forall \tau \in \mathcal{F}^+; t = 1, ..., T-\tau \tag{B.23}$$

$$q_t^\tau \leq aux_t^\tau \qquad \forall \tau \in \mathcal{F}^+; t = 1, ..., T-\tau \tag{B.24}$$

$$q_t^\tau \geq aux_t^\tau - \sum_{a \in \mathcal{F} | a \leq t+\tau-1 \cap a > \tau} q_{t+\tau-a}^a \qquad \forall \tau \in \mathcal{F}^+; t = 1, ..., T-\tau \tag{B.25}$$

$$q_t^\tau, aux_t^\tau \in \{0, 1\}, \beta_i^\tau \in \mathbb{R} \qquad \forall t = 1, ..., T; i = 0, ..., N; \tau \in \mathcal{F} \tag{B.26}$$

B.4. Computation Times

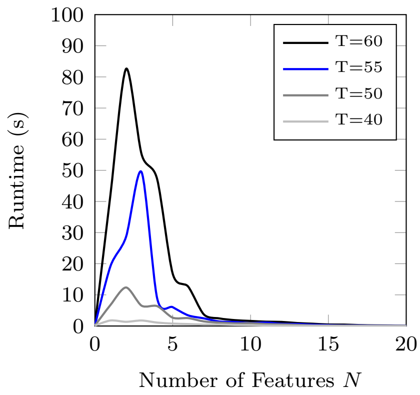

Figure B.1.: Computation times for training of DDA-BD (Average across 100 runs)

Note. Spot and forward ($\tau \in \{1, 2, 3, 4\}$) prices randomly sampled from $N(100, 20)$ $\forall t = 1, ..., T$. Features $i = 1, ..., N$ randomly sampled from $N(20, 10)$. Main observation: Runtime-overfitting trade-off, i.e., by strongly increasing the number of features N (overfitting), computation times decrease (alternate optimal solutions). The computation times refer to the implementation from B.3 with indicator constraints.

B.5. Results of the Controlled Numerical Study

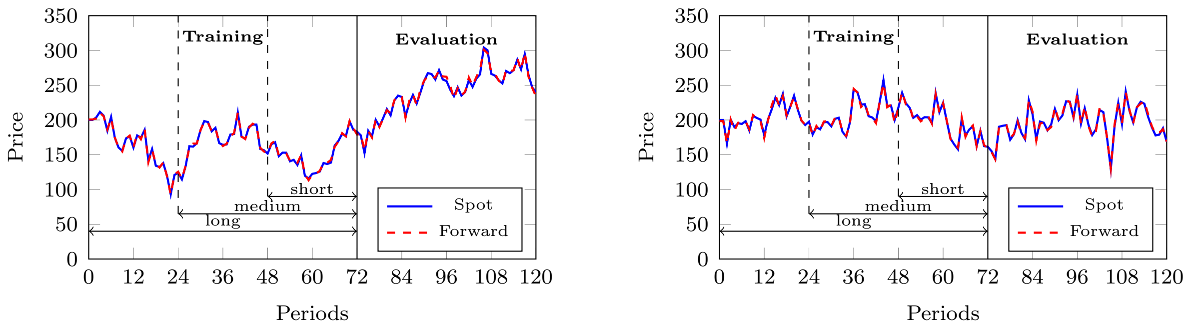

(a) RW **(b) MR**

Figure B.2.: Exemplary sample paths of spot and forward prices under random walk (RW) and mean reversion (MR) assumptions and price process noise $\sigma_{\epsilon_t} = 10$

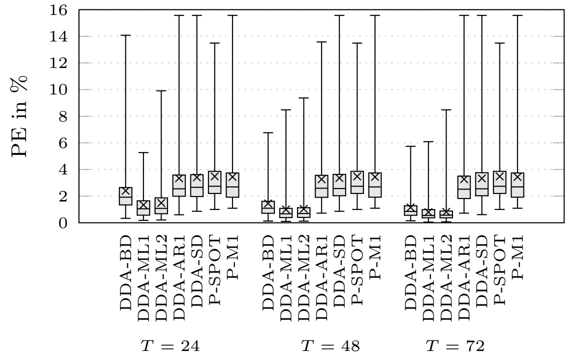

(a1) RW with $\sigma_{\epsilon_t} = 10$

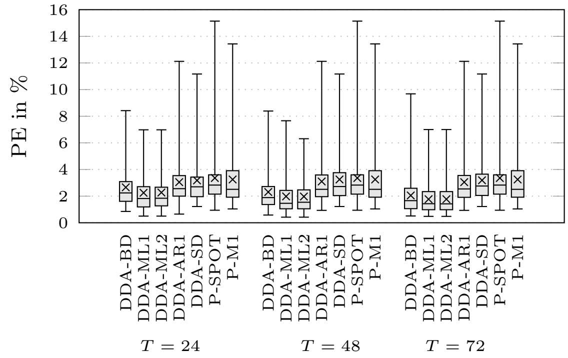

(a2) RW with $\sigma_{\epsilon_t} = 20$

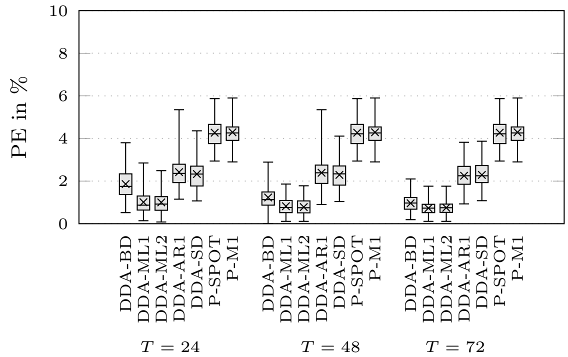

(b1) MR with $\sigma_{\epsilon_t} = 10$

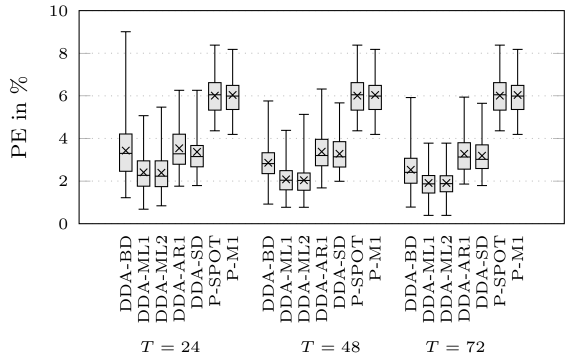

(b2) MR with $\sigma_{\epsilon_t} = 20$

Figure B.3.: Out-of-sample prescription error (PE) of different procurement policies for $\sigma_{\epsilon_t} \in \{10, 20\}$ across 100 simulations conditional on training set size T and price process types (RW, MR)

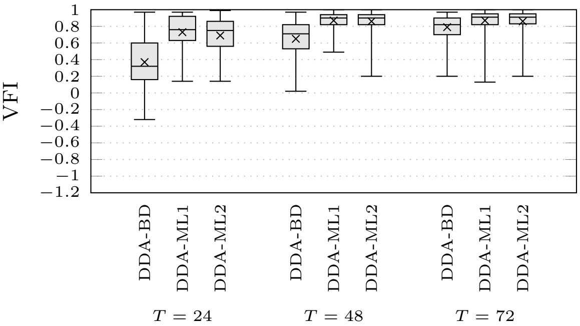

(a1) RW with $\sigma_{\epsilon_t} = 5$

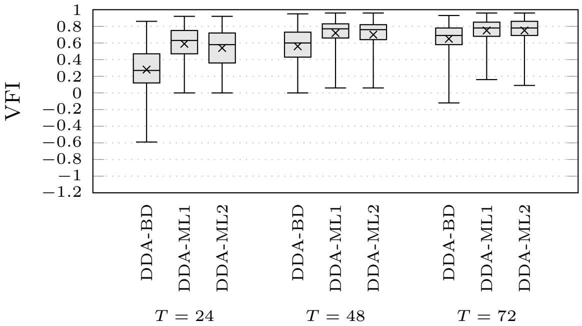

(a2) RW with $\sigma_{\epsilon_t} = 10$

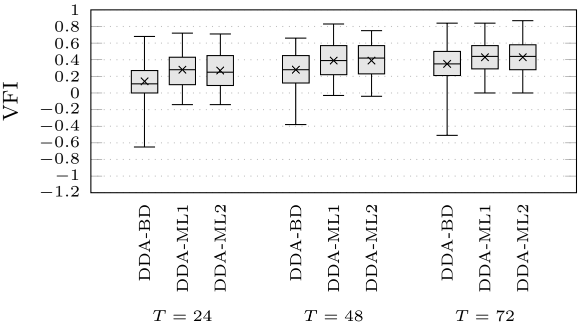

(a3) RW with $\sigma_{\epsilon_t} = 20$

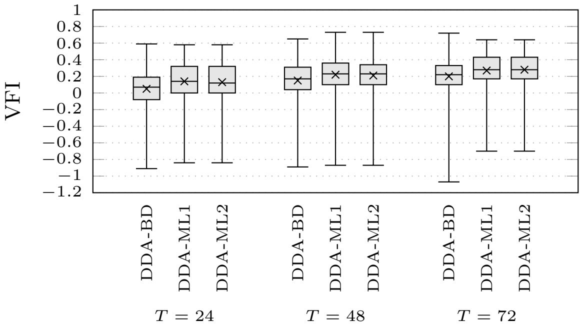

(a4) RW with $\sigma_{\epsilon_t} = 30$

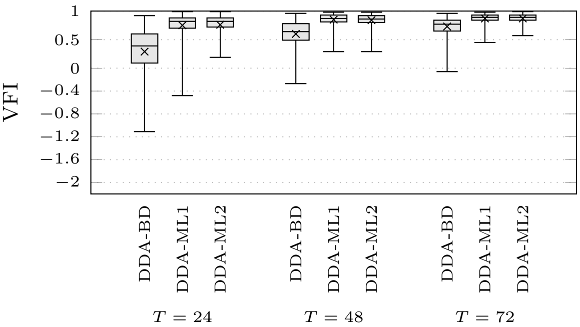

(b1) MR with $\sigma_{\epsilon_t} = 5$

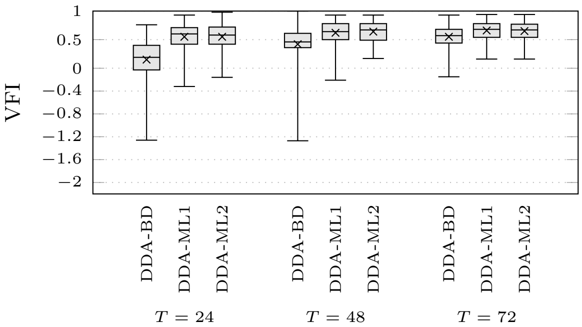

(b2) MR with $\sigma_{\epsilon_t} = 10$

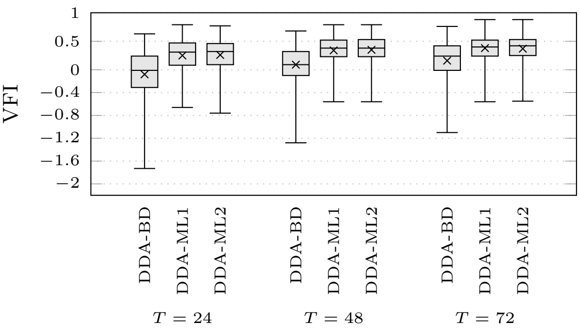

(b3) MR with $\sigma_{\epsilon_t} = 20$

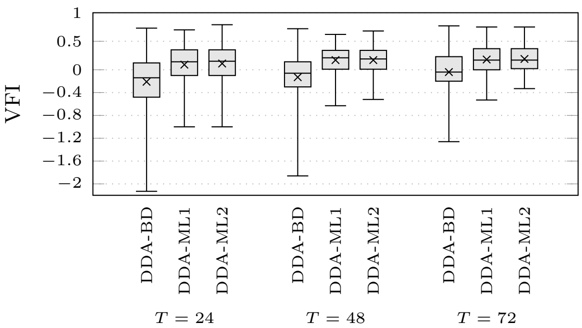

(b4) MR with $\sigma_{\epsilon_t} = 30$

Figure B.4.: Out-of-sample value of feature information (VFI) of DDA-BD, DDA-ML1 and DDA-ML2 across 100 simulations for different training sets T, price process types (RW, MR) and noise levels σ_{ϵ_t}

B.6. Results on Empirical Data: Procurement of Natural Gas

Table B.1.: Historical monthly price and feature data (07-2007 to 06-2017)

Time Series	Unit	Data Source
Prices		
TTF Spot Price (Day-Ahead)	Euro/MWh	Eikon: TRNLTTFDA
TTF Futures Prices	Euro/MWh	Eikon: TRNLTTFMc1-4
Features		
Spot Price Returns	%	Eikon: TRNLTTFDA
TTF Spot Price Lags (t-1,t-2,t-3)	Euro/MWh	Eikon: TRNLTTFDA
Gaspool Spot Price (Day-ahead)	Euro/MWh	Eikon: TRDEGSPD1
Henry Hub Spot Price	USD/mmbtu	Datastream: NATLGAS
Henry Hub NYMEX Front-Month	USD/mmbtu	Eikon: NGc1
Coal Price Front-Month	USD/tonne	Eikon: TRAPI2Mc1
Brent Oil Spot Price	USD/barrel	Datastream: EIACRBR
Gas Production Germany	TWh	Datastream: BDELPRNGP
Domestic Gas Demand Germany	Mio Tons of Oil Equivalent	Datastream: BDXDGAS.P
EUR/USD	-	Eikon: EUR=
EUR/GBP	-	Eikon: UKEURSP
USD Index	-	Datastream: BOEUSA$
Federal Fund Rate	-	Datastream: USFDFUND
S&P 500 Index	-	Datastream: S&PCOMP
Bloomberg Commodity Index	USD	Datastream: DJUBSTR
PPI Energy Germany	-	Datastream: BDENERGYF
PPI Energy UK	-	Datastream: UKOPIEN2F
Temperature (Paris)	Celsius	Eikon: PARIS-OBS

Table B.2.: Procurement of natural gas: Average purchase cost in Euro/MWh

	In-Sample (07-2007 to 06-2016)				**Out-of-Sample (07-2008 to 06-2017)**			
	Mean	StDev	Min	Max	Mean	StDev	Min	Max
DDA-BD	19.58	4.41	10.47	25.17	20.32	4.13	12.31	25.65
DDA-SD	19.82	4.39	10.76	25.22	20.79	4.78	11.45	26.05
DDA-AR1	19.74	4.42	10.60	25.17	20.77	4.58	12.34	25.95
DDA-ML1	19.69	4.44	10.47	25.17	20.06	4.41	11.45	25.65
DDA-ML2	19.70	4.45	10.47	25.17	20.10	4.42	11.45	25.65
REO	20.80	4.48	11.59	25.68	20.28	4.69	11.59	25.68
1/N	21.62	4.57	12.41	26.02	20.97	4.92	12.41	26.02
P-SPOT	20.93	4.18	12.74	26.05	20.14	4.33	12.74	26.05
P-M1	21.09	4.32	12.40	26.09	20.43	4.60	12.40	26.09
P-M2	21.61	4.68	12.27	26.28	21.01	5.03	12.27	26.28
P-M3	22.08	5.03	12.24	28.47	21.46	5.43	12.24	28.47
P-M4	22.44	5.54	11.75	30.75	21.74	5.94	11.75	30.75
PF (=LB)	19.58	4.41	10.47	25.17	18.98	4.71	10.47	25.17
UB	23.83	5.12	14.46	32.02	23.07	5.47	14.46	32.02

Table B.3.: Dominance matrix: % of out-of-sample sub-periods in which a policy (row) performs strictly better than another (column)

	DDA-BD	DDA-SD	DDA-AR1	DDA-ML1	DDA-ML2	REO	1/N	P-SPOT	P-M1	P-M2	P-M3	P-M4
DDA-BD	-	44.4	66.7	0.0	0.0	55.6	66.7	22.2	55.6	55.6	55.6	55.6
DDA-SD	33.3	-	44.4	0.0	0.0	33.3	44.4	33.3	44.4	44.4	66.7	66.7
DDA-AR1	22.2	44.4	-	11.1	11.1	22.2	55.6	33.3	33.3	33.3	55.6	66.7
DDA-ML1	44.4	55.6	66.7	-	11.1	66.7	66.7	44.4	77.8	66.7	66.7	77.8
DDA-ML2	44.4	44.4	66.7	0.0	-	66.7	66.7	44.4	66.7	66.7	66.7	77.8
REO	44.4	66.7	77.8	33.3	33.3	-	66.7	55.6	66.7	77.8	77.8	77.8
1/N	33.3	55.6	44.4	33.3	33.3	33.3	-	44.4	33.3	66.7	66.7	66.7
P-SPOT	33.3	33.3	55.6	11.1	11.1	44.4	55.6	-	55.6	55.6	66.7	66.7
P-M1	44.4	55.6	66.7	22.2	33.3	33.3	66.7	44.4	-	55.6	66.7	66.7
P-M2	44.4	55.6	66.7	33.3	33.3	22.2	33.3	44.4	44.4	-	66.7	66.7
P-M3	44.4	33.3	44.4	33.3	33.3	22.2	33.3	33.3	33.3	33.3	-	66.7
P-M4	44.4	33.3	33.3	22.2	22.2	22.2	33.3	33.3	33.3	33.3	33.3	-

Appendix C.

Appendix of Chapter 6

C.1. SCWP with Fixed Costs, Storage Efficiency, Demand and Market Power

$$\max_{y_t^{\mathbf{i}}, y_t^{\mathbf{o}}} \quad V_0 = \sum_{t=0}^{n} \mathbb{E}\Big[\big(p_t^{\mathbf{o}}\, y_t^{\mathbf{o}} - p_t^{\mathbf{i}}\, y_t^{\mathbf{i}} - K^{\mathbf{i}}\, z_t^{\mathbf{i}} - K^{\mathbf{o}}\, z_t^{\mathbf{o}} - c_h\, I_t\big)|X_{it}\Big] \tag{C.1}$$

s.t.

$$I_{t+1} = \rho(I_t - y_t^{\mathbf{o}} + y_t^{\mathbf{i}}) \qquad \forall t = 0, 1, ..., n \tag{C.2}$$

$$0 \leq y_t^{\mathbf{o}} \leq \min\{I_t, d_t\} \qquad \forall t = 0, 1, ..., n \tag{C.3}$$

$$0 \leq y_t^{\mathbf{i}} \leq C - I_t + y_t^{\mathbf{o}} \qquad \forall t = 0, 1, ..., n \tag{C.4}$$

$$y_t^{\mathbf{i}} \leq C z_t^{\mathbf{i}} \qquad \forall t = 0, 1, ..., n \tag{C.5}$$

$$y_t^{\mathbf{o}} \leq C z_t^{\mathbf{o}} \qquad \forall t = 0, 1, ..., n \tag{C.6}$$

$$z_t^{\mathbf{i}}, z_t^{\mathbf{o}} \in \{0, 1\} \qquad \forall t = 0, 1, ..., n \tag{C.7}$$

with $z_t^{\mathbf{i}}$ and $z_t^{\mathbf{o}}$ being indicator variables for purchase and sale. $K^{\mathbf{i}}$ and $K^{\mathbf{o}}$ are fixed costs for purchase and sale, respectively. $0 \leq \rho \leq 1$ is the storage efficiency. We can also consider a demand component d_t that limits $y_t^{\mathbf{o}}$.

If the warehouse has market impact (price setter), then p_{t+1} is a non-decreasing linear function $B - y_t^{\mathbf{o}} + y_t^{\mathbf{i}}$ with B being a constant. Large $y_t^{\mathbf{o}}$ decrease the price p_{t+1}, while large $y_t^{\mathbf{i}}$ increase the price p_{t+1}. While under price taker assumptions, the problem can be formulated as an LP, under price setter assumptions, due to nonlinearities when inventory decisions affect prices, the optimization becomes a quadratic program. The models can be solved using quadratic solvers.

C.2. Empirical Forecast Error of Futures Prices

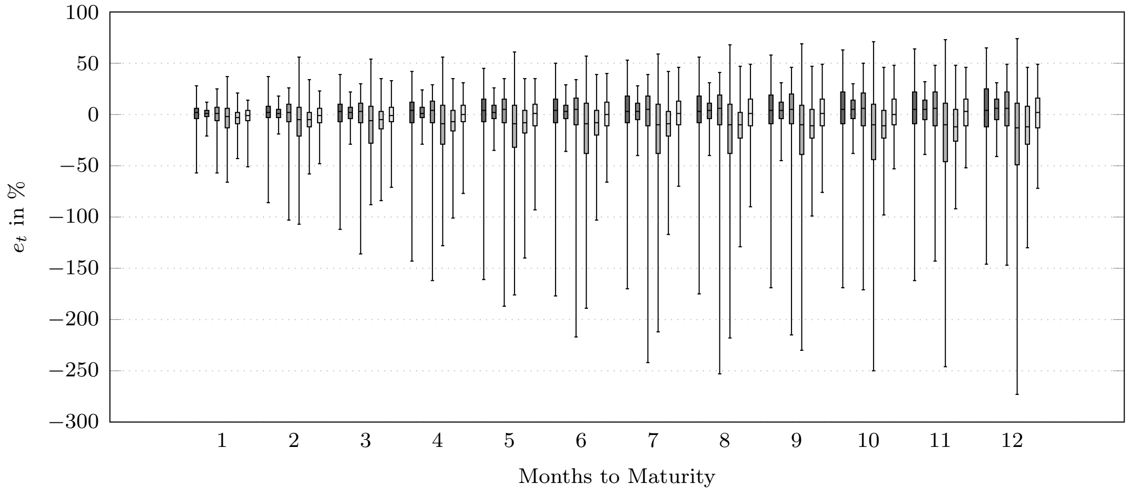

Figure C.1.: Distribution of prediction error e_t in % for copper (dark gray), gold, crude oil, natural gas, corn and soybean (light gray) (2000-2017)

Note. Boxplots: minimum, 1st-, 2nd-, 3rd-quartile, maximum. Some abrupt price declines (e.g., for copper, crude oil and natural gas during the financial crisis) lead to a significant overestimation of spot prices by futures prices and therefore to a strongly negative prediction error e_t.

C.3. Deterministic Analysis under Perfect Foresight

C.3.1. Profit over Time

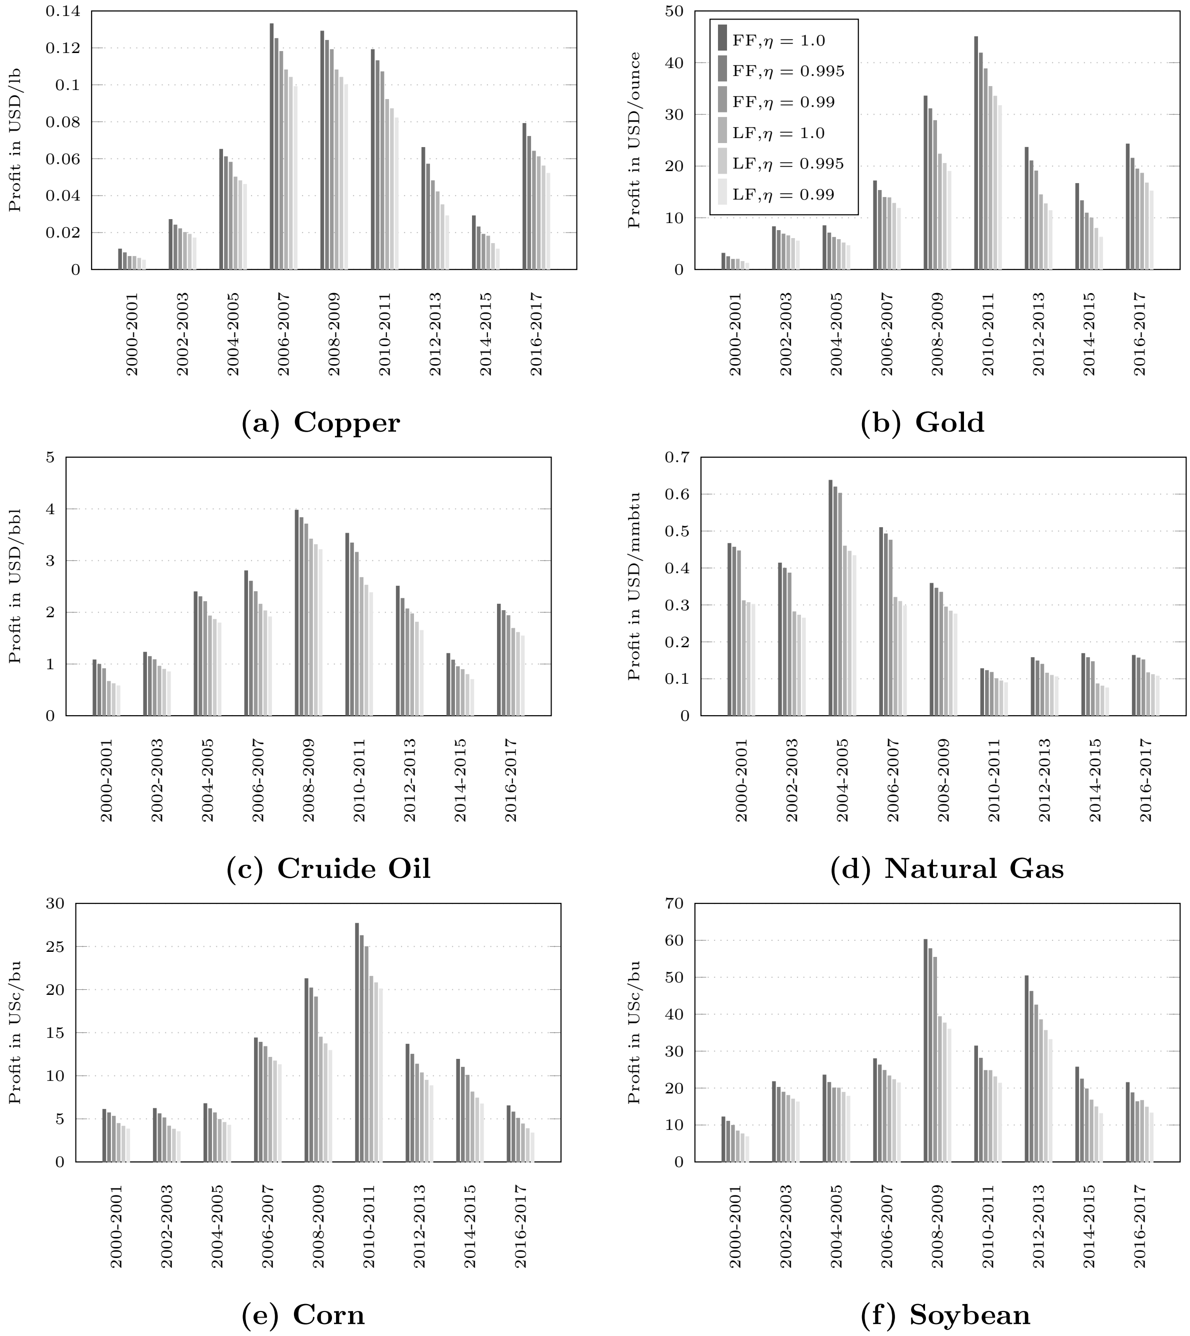

(a) Copper **(b) Gold**

(c) Cruide Oil **(d) Natural Gas**

(e) Corn **(f) Soybean**

Figure C.2.: Average profit under perfect foresight for different degrees of warehouse flexibility (SCWP-FF, SCWP-LF) and frictions ($\eta = \eta^{\mathbf{i}} = \eta^{\mathbf{o}}$)

C.3.2. Optimal versus Myopic Performance

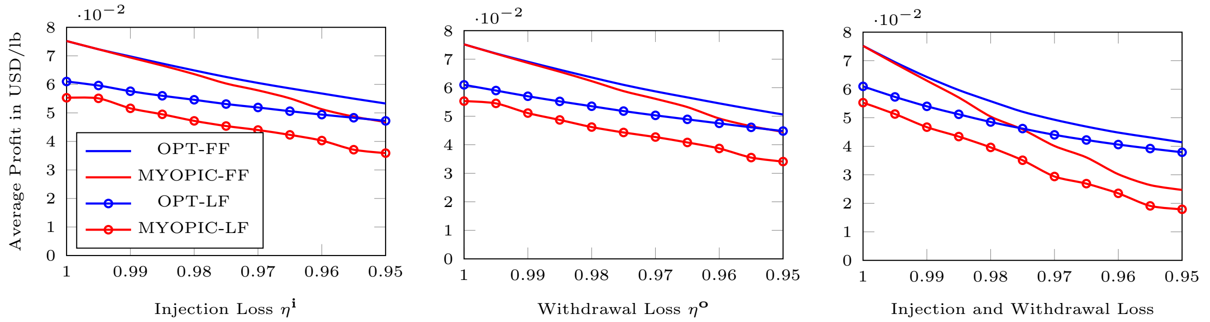

Figure C.3.: Copper storage: Performance of optimal vs. myopic policies (2000-2017)

Note. For fully flexible copper storage, a 1% injection and withdrawal loss yields a profit decline of 14.48%. If frictions are realistically small (1%), the profit decline by myopic decision-making is approximately 2.0%.

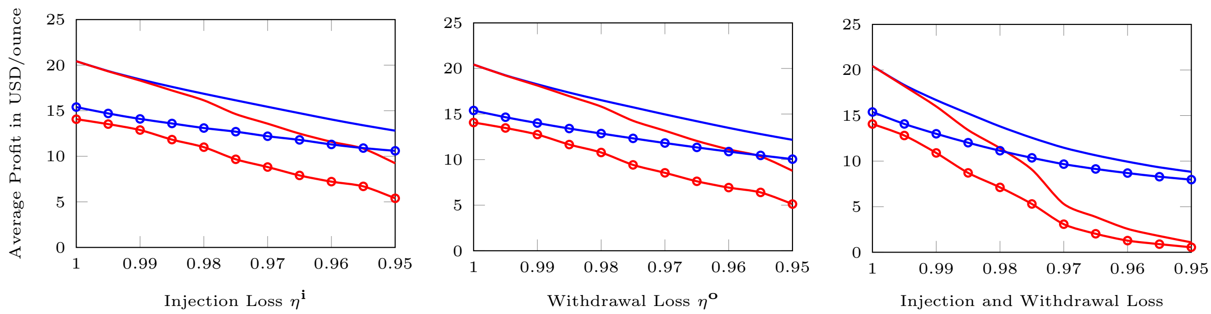

Figure C.4.: Gold storage: Performance of optimal vs. myopic policies (2000-2017)

Note. For fully flexible gold storage, a 1% injection and withdrawal loss yields a profit decline of 18.26%. If frictions are realistically small (1%), the profit decline by myopic decision-making is approximately 4.2%.

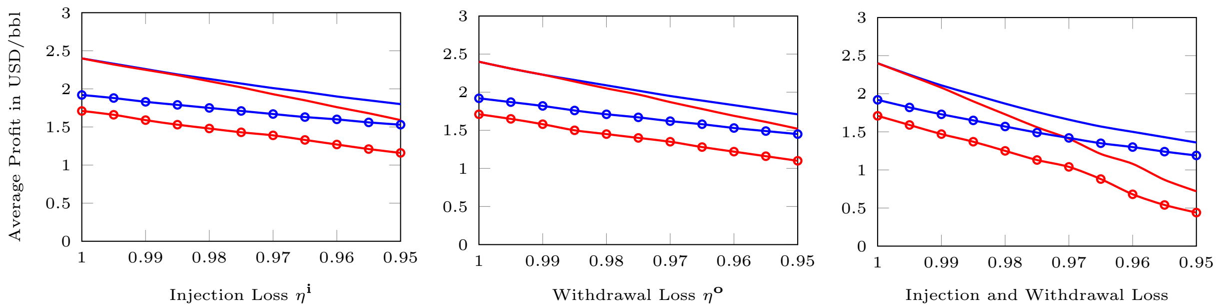

Figure C.5.: Oil storage: Performance of optimal vs. myopic policies (2000-2017)

Note. For fully flexible oil storage, a 1% injection and withdrawal loss yields a profit decline of 11.94%. If frictions are realistically small (1%), the profit decline by myopic decision-making is approximately 1.6%.

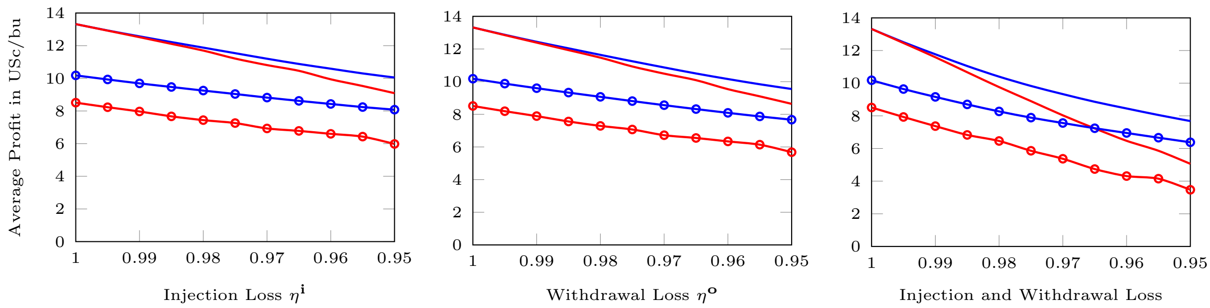

Figure C.6.: Corn storage: Performance of optimal vs. myopic policies (2000-2017)

Note. For corn silos, a 1% injection and withdrawal loss leads to a profit decline of 11.66%. If frictions are small (1%), profit decline by myopic decision-making is approximately 1.5%.

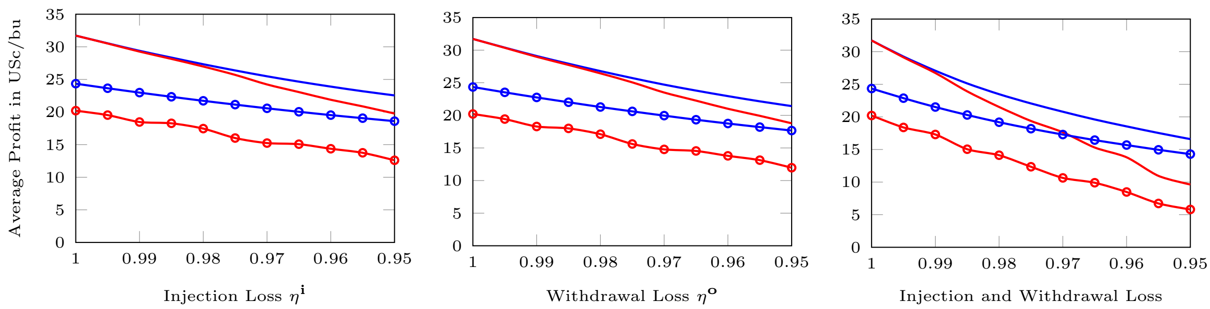

Figure C.7.: Soybean storage: Performance of optimal vs. myopic policies (2000-2017)

Note. For soybean storage, a 1% injection and withdrawal loss leads to a profit decline of 14.62%. If frictions are small (1%), losses by myopic decision-making are approximately 1.3%.

C.3.3. Impact of the Planning Horizon

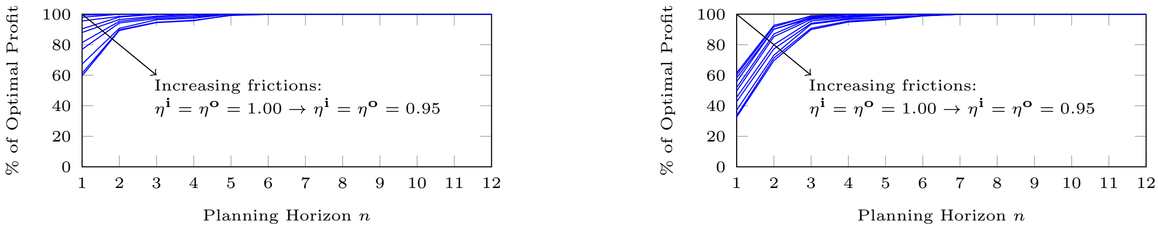

(a) SCWP-FF **(b) SCWP-LF**

Figure C.8.: Copper storage: Performance impact of planning horizon n (2000-2017)

Note. For plausible frictions $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.99$, a planning horizon of 1/2/3 periods generates 97.98%/99.97%/99.97% (SCWP-FF) and 58.27%/90.13%/97.99% (SCWP-LF) of the potential value. $n = 4$ (SCWP-FF, SCWP-LF) is sufficient to generate 100% of the potential value.

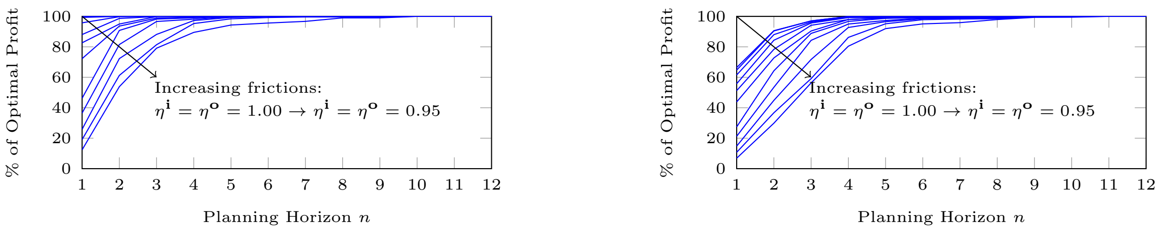

(a) SCWP-FF **(b) SCWP-LF**

Figure C.9.: Gold storage: Performance impact of planning horizon n (2000-2017)

Note. For plausible frictions $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.99$, a planning horizon of 1/2/3 periods generates 95.77%/99.78%/100% (SCWP-FF) and 61.53%/88.01%/96.26% (SCWP-LF) of the potential value. $n = 3$ (SCWP-FF) respectively $n = 8$ (SCWP-LF) is sufficient to generate 100% of the potential value.

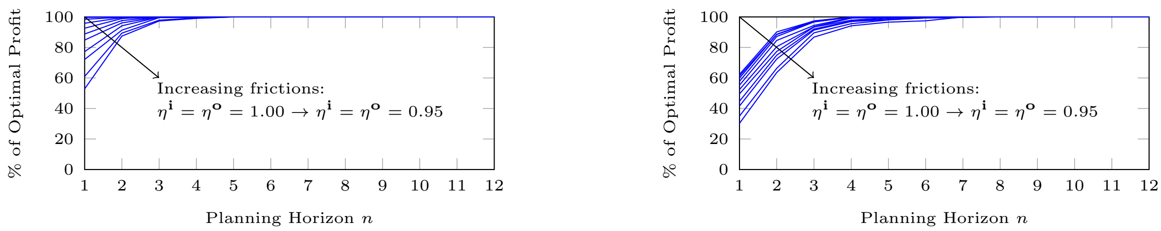

(a) SCWP-FF **(b) SCWP-LF**

Figure C.10.: Oil storage: Performance impact of planning horizon n (2000-2017)

Note. For plausible frictions $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.99$, a planning horizon of 1/2/3 periods generates 98.42%/99.80%/100% (SCWP-FF) and 59.92%/88.64%/97.37% (SCWP-LF) of the potential value. $n = 3$ (SCWP-FF) respectively $n = 4$ (SCWP-LF) is sufficient to generate 100% of the potential value.

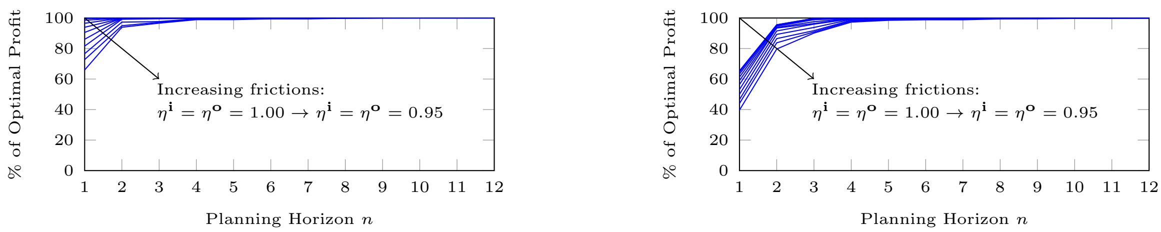

(a) SCWP-FF **(b) SCWP-LF**

Figure C.11.: Corn storage: Performance impact of planning horizon n (2000-2017)

Note. For plausible frictions $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.99$, a planning horizon of 1/2/3 periods generates 98.51%/99.94%/100% (SCWP-FF) and 63.26%/94.98%/99.20% (SCWP-LF) of the potential value. $n = 3$ (SCWP-FF) respectively $n = 6$ (SCWP-LF) is sufficient to generate 100% of the potential value.

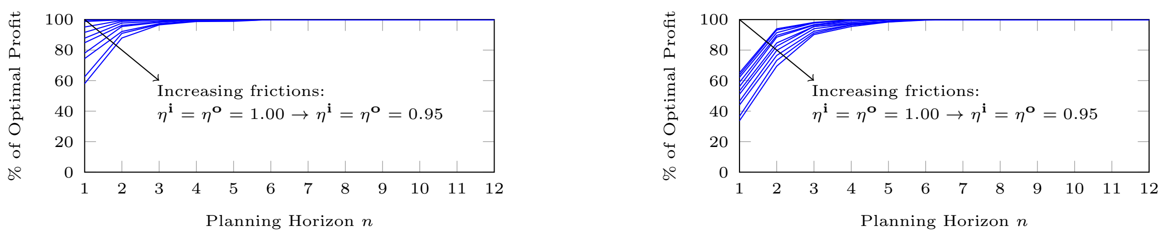

(a) SCWP-FF **(b) SCWP-LF**

Figure C.12.: Soybean storage: Performance impact of planning horizon n (2000-2017)

Note. For plausible frictions $\eta^{\mathbf{i}} = \eta^{\mathbf{o}} = 0.99$, a planning horizon of 1/2/3 periods generates 99.58%/100%/100% (SCWP-FF) and 63.69%/93.17%/97.79% (SCWP-LF) of the potential value. $n = 2$ (SCWP-FF) respectively $n = 5$ (SCWP-LF) is sufficient to generate 100% of the potential value.

C.4. Stochastic Analysis: Performance of RIA and DDA

Table C.1.: Average performance (2000-2017) of futures-based RIA compared to RIA with perfect one-step directional forecasts (PDF) and RIA with perfect one-step point forecasts (PPF) in $V/V^{\mathrm{PF}} \cdot 100\%$

	$\eta^{\mathrm{i}} = \eta^{\mathrm{o}} = 1$			$\eta^{\mathrm{i}} = \eta^{\mathrm{o}} = 0.995$			$\eta^{\mathrm{i}} = \eta^{\mathrm{o}} = 0.99$		
Planning horizon n	RIA	PDF	PPF	RIA	PDF	PPF	RIA	PDF	PPF
SCWP-FF ($G^{\mathrm{i}} = G^{\mathrm{o}} = C$)									
Copper									
1 (myopic)	13.2	100.0	100.0	-7.9	98.1	99.3	-7.3	93.0	96.7
3	13.2	100.0	100.0	-0.7	98.1	99.0	-1.6	93.0	97.6
6	13.2	100.0	100.0	5.6	98.1	99.3	-3.8	93.0	98.3
12	13.2	100.0	100.0	5.6	98.1	99.3	-3.8	93.0	98.3
Gold									
1 (myopic)	13.3	100.0	100.0	0.0	97.8	99.3	0.0	90.7	94.2
3	13.3	100.0	100.0	-0.8	97.8	99.2	0.0	90.7	96.2
6	13.3	100.0	100.0	13.3	97.8	99.0	8.2	90.7	94.2
12	13.3	100.0	100.0	13.5	97.8	99.2	10.0	90.7	94.9
Crude Oil									
1 (myopic)	-3.6	100.0	100.0	-4.9	99.8	99.6	-10.2	98.8	98.0
3	-3.6	100.0	100.0	-4.3	99.8	99.6	-10.8	98.8	98.6
6	-3.6	100.0	100.0	-4.3	99.8	99.6	-6.6	98.8	98.4
12	-3.6	100.0	100.0	-4.3	99.8	99.6	-6.6	98.8	98.4
Natural Gas									
1 (myopic)	27.2	100.0	100.0	34.7	99.9	99.9	28.4	99.4	99.7
3	27.2	100.0	100.0	16.3	99.9	99.9	20.5	99.4	99.3
6	27.2	100.0	100.0	16.3	99.9	99.9	19.9	99.4	99.2
12	27.2	100.0	100.0	16.3	99.9	99.9	19.9	99.4	99.2
Corn									
1 (myopic)	24.7	100.0	100.0	21.7	99.5	99.7	25.3	98.2	99.0
3	24.7	100.0	100.0	20.3	99.5	99.6	19.2	98.2	98.8
6	24.7	100.0	100.0	17.7	99.5	99.6	16.7	98.2	98.8
12	25.6	100.0	100.0	18.7	99.5	98.8	16.7	98.2	98.8
Soybean									
1 (myopic)	29.4	100.0	100.0	27.8	99.7	99.4	25.0	98.1	98.1
3	29.4	100.0	100.0	28.4	99.7	99.4	21.1	98.1	97.6
6	29.4	100.0	100.0	29.2	99.7	99.4	28.5	98.1	98.4
12	29.4	100.0	100.0	30.7	99.7	98.6	28.5	98.1	98.4
SCWP-LF ($G^{\mathrm{i}} = G^{\mathrm{o}} = 0.5C$)									
Copper									
1 (myopic)	9.4	68.1	68.1	-4.8	66.3	67.3	-4.4	62.3	65.4
3	7.7	61.0	61.0	-8.8	59.0	59.2	-1.3	53.9	56.2
6	7.7	62.8	62.8	-3.5	56.8	57.9	-4.8	46.1	53.0
12	7.7	62.8	62.8	-6.4	54.6	55.3	-5.9	50.1	53.5
Gold									
1 (myopic)	8.4	71.9	71.9	0.0	69.6	70.7	0.0	64.3	67.3
3	24.7	72.8	72.8	-0.7	48.3	50.3	0.0	40.6	43.9
6	24.8	73.6	73.6	17.3	48.8	56.0	6.6	33.8	45.9
12	24.8	73.6	73.6	18.3	61.9	62.6	17.1	41.5	48.8
Crude Oil									
1 (myopic)	-3.6	65.8	65.8	-3.8	65.1	65.0	-7.6	64.2	63.6
3	-7.3	58.9	58.9	-3.9	57.6	57.4	-5.5	61.1	57.4
6	-7.3	64.2	64.2	-5.0	59.3	59.1	-0.2	60.0	59.9
12	-7.3	64.2	64.2	-5.0	59.3	59.1	-0.2	60.7	60.4
Natural Gas									
1 (myopic)	18.2	73.1	73.1	23.4	73.0	73.1	18.7	72.6	72.8
3	13.9	66.6	66.6	8.8	69.9	69.8	13.4	67.5	68.8
6	13.9	67.8	67.8	11.4	66.7	66.6	7.2	65.9	65.7
12	13.9	67.8	67.8	11.4	67.1	66.9	11.9	65.9	65.7
Corn									
1 (myopic)	16.7	69.1	69.1	14.6	68.7	68.8	16.7	67.6	68.1
3	17.4	67.5	67.5	14.3	65.2	65.5	12.1	62.3	63.3
6	17.4	67.5	67.5	7.5	64.2	65.7	4.9	58.2	60.5
12	10.6	62.1	62.9	8.1	60.5	60.9	4.9	58.2	58.0
Soybean									
1 (myopic)	18.5	66.5	66.5	16.5	65.3	65.1	14.0	63.4	63.4
3	26.3	70.4	70.4	24.6	66.9	64.9	10.4	60.7	55.9
6	26.3	70.8	70.7	25.8	68.5	66.6	19.7	61.8	57.5
12	25.9	66.6	63.6	23.2	64.6	61.3	19.7	61.8	58.1

Table C.2.: Out-of-sample performance of futures-based IA and futures-based RIA compared to different data-driven policies measured in $V/V^{\mathrm{PF}} \cdot 100\%$ from 2002 until 2017 across 192 instances

	Mean	Min	25%-Q	50%-Q	75%-Q	Max
			IA			
Copper	11.1	-91.3	0.0	0.0	30.4	79.4
Gold	9.9	-100.5	0.0	0.0	21.6	79.1
Crude Oil	7.4	-121.6	-3.0	10.9	25.2	65.3
Natural Gas	9.1	-96.4	-9.7	15.3	30.8	66.3
Corn	1.8	-123.5	-17.5	2.7	27.6	67.6
Soybean	12.5	-64.3	-7.2	9.0	35.5	68.0
Overall	**8.6**	**-123.5**	**-4.0**	**5.9**	**29.4**	**79.4**
			RIA			
Copper	1.7	-45.6	-12.2	0.0	9.2	70.0
Gold	12.6	-107.5	0.0	1.6	26.9	78.7
Crude Oil	0.0	-171.9	-9.3	10.1	27.7	70.5
Natural Gas	15.9	-125.9	4.0	22.0	43.5	79.5
Corn	18.4	-44.3	-7.0	25.6	43.6	64.9
Soybean	29.0	-18.6	2.0	29.9	48.5	79.8
Overall	**12.9**	**-171.9**	**-2.7**	**12.0**	**36.9**	**79.8**
			DDA-LDR w/o Forward Optimization			
Copper	-10.2	-460.6	0.6	14.2	36.9	67.2
Gold	-14.9	-236.6	-35.6	3.9	16.2	67.2
Crude Oil	-22.7	-344.8	-15.0	4.4	21.7	57.0
Natural Gas	-21.1	-208.7	-20.8	-1.1	9.1	22.0
Corn	-10.0	-72.1	-25.3	-4.3	4.5	31.3
Soybean	5.2	-41.3	-5.6	5.8	12.1	60.4
Overall	**-12.3**	**-460.6**	**-16.3**	**2.1**	**16.5**	**67.2**
			DDA-LDR w/ Forward Optimization			
Copper	-7.6	-237.6	-1.7	7.1	34.3	68.8
Gold	-11.3	-129.4	-10.7	1.4	16.5	47.8
Crude Oil	-0.5	-50.0	-9.9	-0.1	10.2	48.7
Natural Gas	-3.1	-68.9	-22.4	4.9	15.1	44.8
Corn	1.8	-49.3	-9.3	5.9%	13.2	34.3
Soybean	5.6	-20.7	-5.1	-1.5	6.1	58.3
Overall	**-2.5**	**-237.6**	**-8.4**	**2.4**	**15.5**	**68.8**
			DDA-OSP w/o Forward Optimization			
Copper	-2.6	-325.5	1.3	24.3	45.8	72.3
Gold	1.5	-121.1	0.0	5.3	12.7	85.1
Crude Oil	-2.8	-209.1	4.9	15.3	30.7	60.9
Natural Gas	23.0	-42.9	1.2	18.0	55.1	69.0
Corn	20.9	-33.6	9.8	20.1	31.9	60.8
Soybean	17.5	-5.6	1.4	17.5	25.5	54.8
Overall	**9.6**	**-325.5**	**0.7**	**15.7**	**31.8**	**85.1**
			DDA-OSP w/ Forward Optimization			
Copper	23.9	-188.2	5.9	35.9	68.6	83.6
Gold	11.9	-136.6	4.6	21.4	32.7	80.5
Crude Oil	15.6	-159.1	6.3	21.3	41.7	75.5
Natural Gas	7.0	-69.3	-31.8	11.4	40.0	71.5
Corn	39.3	-6.6	25.4	43.1	56.4	71.9
Soybean	33.1	0.0	18.1	34.9	48.1	71.1
Overall	**21.8**	**-188.2**	**7.6**	**26.7**	**49.4**	**83.6**
			DDA-VFA			
Copper	-8.2	-303.5	0.6	12.4	24.4	82.4
Gold	10.8	-78.1	-2.9	7.5	43.4	75.0
Crude Oil	-22.1	-239.1	-16.1	-3.6	7.5	34.5
Natural Gas	-4.2	-77.6	-25.0	-13.6	25.8	59.0
Corn	0.2	-76.9	-16.5	-1.9	23.2	41.0
Soybean	-4.8	-96.7	-9.9	-0.3	12.1	36.4
Overall	**-4.7**	**-303.5**	**-15.2**	**1.4**	**21.1**	**82.4**

Table C.3.: Out-of-sample dominance matrix of different policies in % of instances (# = 192) in which row performs strictly better than column

	IA	RIA	DDA-LDR*	DDA-LDR**	DDA-OSP*	DDA-OSP**	DDA-VFA
				Copper			
IA	-	42.7	41.7	47.4	37.5	27.6	51.0
RIA	30.2	-	31.8	34.9	22.9	18.8	39.6
DDA-LDR*	57.3	68.2	-	54.2	45.8	18.8	52.1
DDA-LDR**	52.6	65.1	45.8	-	35.4	18.8	47.9
DDA-OSP*	58.3	77.1	52.1	64.6	-	16.7	50.0
DDA-OSP**	68.2	81.3	79.2	81.3	70.8	-	68.8
DDA-VFA	49.0	60.4	47.9	52.1	50.0	31.3	-
				Gold			
IA	-	27.1	53.6	60.9	42.2	37.5	43.8
RIA	34.4	-	61.5	62.5	52.6	42.7	46.4
DDA-LDR*	44.3	36.5	-	56.3	45.8	27.1	16.7
DDA-LDR**	39.1	37.5	43.8	-	45.8	31.3	22.9
DDA-OSP*	48.4	38.0	50.0	54.2	-	25.0	41.7
DDA-OSP**	62.5	57.3	72.9	68.8	60.4	-	56.3
DDA-VFA	51.0	47.4	83.3	77.1	58.3	41.7	-
				Crude Oil			
IA	-	50.0	62.0	65.6	52.1	33.9	79.2
RIA	50.0	-	69.3	59.9	47.9	33.9	71.9
DDA-LDR*	38.0	30.7	-	50.0	18.8	14.6	56.3
DDA-LDR**	34.4	40.1	50.0	-	31.3	31.3	58.3
DDA-OSP*	46.9	52.1	81.3	68.8	-	27.1	77.1
DDA-OSP**	65.6	66.1	85.4	68.8	66.7	-	87.5
DDA-VFA	20.8	28.1	43.8	41.7	22.9	12.5	-
				Natural Gas			
IA	-	38.5	83.9	64.6	35.4	45.3	59.4
RIA	61.5	-	93.8	81.3	33.9	53.6	69.3
DDA-LDR*	16.1	6.3	-	33.3	2.1	27.1	37.5
DDA-LDR**	35.4	18.8	66.7	-	20.8	37.5	52.1
DDA-OSP*	64.6	66.1	97.9	79.2	-	64.6	68.8
DDA-OSP**	54.7	46.4	72.9	62.5	35.4	-	54.2
DDA-VFA	40.6	30.7	62.5	47.9	31.3	45.8	-
				Corn			
IA	-	27.6	67.7	57.8	29.2	6.8	50.5
RIA	72.4	-	87.5	70.3	51.0	22.9	79.7
DDA-LDR*	32.3	12.5	-	25.0	12.5	0.00	35.4
DDA-LDR**	42.2	29.7	75.0	-	18.8	2.1	45.8
DDA-OSP*	70.8	44.8	87.5	81.3	-	29.2	77.1
DDA-OSP**	93.2	75.5	100.0	97.9	70.8	-	91.7
DDA-VFA	49.5	19.8	64.6	54.2	18.8	8.3	-
				Soybean			
IA	-	20.3	57.8	58.9	47.4	30.2	60.9
RIA	79.7	-	79.7	74.0	62.5	44.3	83.3
DDA-LDR*	42.2	20.3	-	52.1	22.9	10.4	60.4
DDA-LDR**	41.1	26.0	47.9	-	20.8	12.5	64.6
DDA-OSP*	50.5	37.5	72.9	79.2	-	14.6	79.2
DDA-OSP**	69.8	55.7	87.5	87.5	79.2	-	91.7
DDA-VFA	39.1	16.7	35.4	35.4	4.2	4.2	-
				Overall			
IA	-	34.4	61.1	59.2	40.6	30.2	57.5
RIA	54.7	-	70.6	63.8	45.1	36.0	65.0
DDA-LDR*	38.4	29.1	-	45.1	24.7	16.3	43.1
DDA-LDR**	40.8	36.2	54.9	-	28.8	22.2	48.6
DDA-OSP*	56.6	52.6	73.6	71.2	-	29.5	65.6
DDA-OSP**	69.0	63.7	83.0	77.8	63.9	-	75.0
DDA-VFA	41.7	33.9	56.3	51.4	30.9	24.0	-

* w/o forward optimization, ** w/ forward optimization

Bibliography

Alquist, R. and L. Kilian (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics* 25 (1), pp. 539–573.

Andersen, R. (2008). Modern methods for robust regression. *Series: Quantitative Applications in the Social Sciences*. Ed. by T.F. Liao. 1st ed. Sage Publications.

Andersson, H. (2007). Are commodity prices mean reverting? *Applied Financial Economics* 17 (10), pp. 769–783.

Ang, A. and G. Bekaert (2002). Regime switches in interest rates. *Journal of Business & Economic Statistics* 20 (2), pp. 163–182.

Arifoğlu, K. and S. Özekici (2010). Optimal policies for inventory systems with finite capacity and partially observed Markov-modulated demand and supply processes. *European Journal of Operational Research* 204 (3), pp. 421–438.

Arnold, J. and S. Minner (2011). Financial and operational instruments for commodity procurement in quantity competition. *International Journal of Production Economics* 131 (1), pp. 96–106.

Arnold, J., S. Minner, and B. Eidam (2011). Raw material procurement with fluctuating prices. *International Journal of Production Economics* 121 (2), pp. 353–364.

Arrow, K.A. (1958). Historical background. *Studies in the Mathematical Theory of Inventory and Production*. Ed. by K.J. Arrow, S. Karlin, and H.E. Scarf. 1st ed. Stanford University Press.

Arthur D. Little (2014). Big data in procurement. *Arthur D. Little Publications*.

Azoury, K.S. (1985). Bayes solution to dynamic inventory models under unknown demand distribution. *Management Science* 31 (9), pp. 1150–1160.

Ban, G.-Y. and C. Rudin (2019). The big data newsvendor: Practical insights from machine learning. *Operations Research* 67 (1), pp. 90–108.

Ban, G.-Y., J. Gallien, and A. Mersereau (2018a). Dynamic procurement of new products with covariate information: The residual tree method. *Forthcoming in Manufacturing & Service Operations Management*.

Ban, G.-Y., N. El Karoui, and A. Lim (2018b). Machine learning and portfolio optimization. *Management Science* 64 (3), pp. 1136–1154.

Bartram, S.M., G.W. Brown, and F.R. Fehle (2009). International evidence on financial derivatives usage. *Financial Management* 38 (1), pp. 185–206.

Baum, L.E., T. Petrie, G. Soules, and N. Weiss (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics* 41 (1), pp. 164–171.

Bayraktar, E. and M. Ludkovski (2010). Inventory management with partially observed nonstationary demand. *Annals of Operations Research* 176 (1), pp. 7–39.

Bellman, R. (1956). On the theory of dynamic programming - A warehousing problem. *Management Science* 2 (3), pp. 272–275.

Berling, P. and V. Martínez-de-Albeníz (2011). Optimal inventory policies when purchase price and demand are stochastic. *Operations Research* 59 (1), pp. 109–124.

Berling, P. and Z. Xie (2014). Approximation algorithms for optimal purchase/inventory policy when purchase price and demand are stochastic. *OR Spectrum* 36 (4), pp. 1077–1095.

Bertsekas, D. (1995). *Dynamic programming and optimal control*. 1st ed. Athena Scientific.

Bertsimas, D. and N. Kallus (2016). From predictive to prescriptive analytics. *Working Paper*.

Bertsimas, D. and A. Thiele (2006). Robust and data-driven optimization: Modern decision making under uncertainty. *INFORMS Tutorials in Operations Research*, pp. 95–122.

Beschaffungsmanagement (2009). Rohstoffkosten-Management: Einkaufsexperten fordern jetzt aktives Management von Rohstoffkosten für schlechtere Zeiten. *Beschaffungsmanagement 04/2009*, pp. 12–13.

Bessembinder, H. (1991). Forward contracts and firm value: Incentives and contracting effects. *Journal of Financial and Quantitative Analysis* 26 (1), pp. 491–532.

Beutel, A-L. and S. Minner (2012). Safety stock planning under causal demand forecasting. *International Journal of Production Economics* 140 (2), pp. 637–645.

Beyer, D. and S.P. Sethi (1997). Average cost optimality in inventory models with Markovian demands. *Journal of Optimization Theory and Applications* 92 (3), pp. 497–526.

Bloomberg (2015). Market volatility has changed immensely. Retrieved from `http://www.bloomberg.com/news/articles/2015-09-08/market-volatility-has-changed-immensely`.

Borovkova, S. and H. Geman (2008). Forward curve modeling in commodity markets. *Risk management in commodity markets - From shipping to agriculturals and energy*. Ed. by H. Geman. 1st ed. Wiley Finance. Chap. 2, pp. 9–32.

Boston Consulting Group (2017). Capturing commodity trading's $70 billion prize - How digitalization is changing commodity trading. Retrieved from `https://www.bcg.com/publications/2017/commodity-trading-risk-management-energy-environment-capturing-commodity-trading-billion-prize.aspx`.

Brennan, M.J. (1958). The supply of storage. *The American Economic Review* 48 (1), pp. 50–72.

Breslin, J., L. Clewlow, T. Elbert, C. Kwok, and C. Strickland (2008). Gas storage: Overview and static valuation. *Energy Risk November 2008*, pp. 62–68.

— (2009). Gas storage: Rolling intrinsic valuation. *Energy Risk January 2009*, pp. 61–65.

Brown, S. and K. Yücel (2008). What drives natural gas prices? *The Energy Journal* 29 (2), pp. 43–58.

Cahn, A.S. (1948). The warehouse problem. *Bulletin of the American Mathematical Society* 54 (11), p. 1073.

Carter, D.A., D.A. Rogers, and B.J. Simkins (2006). Does hedging affect firm value? Evidence from the U.S. airline industry. *Financial Management* 35 (1), pp. 53–86.

Chand, S., V.N. Hsu, and S. Sethi (2002). Forecast, solution, and rolling horizons in operations management problems: A classified bibliography. *Manufacturing and Service Operations Management* 4 (1), pp. 25–43.

Charnes, A. and W.W. Cooper (1955). Generalization of the warehousing problem. *Operations Research Quarterly* 6 (4), pp. 131–172.

Charnes, A., M. Drèze, and M. Miller (1966). Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica* 34 (2), pp. 307–330.

Chen, F. and J.-S. Song (2001). Optimal policies for multiechelon inventory problems with Markov-modulated demand. *Operations Research* 49 (2), pp. 226–234.

Chen, S. and M. Insley (2012). Regime switching in stochastic models of commodity prices: An application to an optimal tree harvesting problem. *Journal of Economic Dynamics and Control* 36 (2), pp. 201–219.

Chen, Z. and P.A. Forsyth (2010). Implications of a regime-switching model on natural gas storage valuation and optimal operation. *Quantitative Finance* 10 (2), pp. 159–176.

Cheng, F. and S.P. Sethi (1999). Optimality of state-dependent (s,S) policies in inventory models with Markov-modulated demand and lost sales. *Production and Operations Management* 8 (2), pp. 183–192.

Choi, K. and S. Hammoudeh (2010). Volatility behavior of oil, industrial commodity and stock markets in a regime-switching environment. *Energy Policy* 38 (8), pp. 4388–4399.

Clements, M.P., P.H. Franses, and N.R. Swanson (2004). Forecasting economic and financial time-series with non-linear models. *International Journal of Forecasting* 20 (2), pp. 169–183.

CME Group (2014). How satellite technology is predicting crop yields. Retrieved from `http://openmarkets.cmegroup.com/9625/how-satellite-technology-is-predicting-crop-yields`.

Cohen, M.C., I. Lobel, and R.P. Leme (2016). Feature-based dynamic pricing. *Working Paper*.

Cootner, P.H. (1964). *The random character of stock market prices*. 1st ed. MIT Press.

Cortazar, G., C. Millard, H. Ortega, and E. Schwartz (2018). Commodity price forecasts, futures prices and pricing models. *Forthcoming in Management Science*.

Cruise, J., L. Flatley, R. Gibbens, and Zachary S. (2019). Control of energy storage with market impact: Lagrangian approach and horizons. *Operations Research* 67 (1), pp. 1–9.

Curtis, F.E. and K. Scheinberg (2017). Optimization methods for supervised machine learning: From linear models to deep learning. *INFORMS Tutorials in Operations Research*, pp. 89–113.

DeMarzo, P.M. and D. Duffie (1991). Corporate financial hedging with proprietary information. *Journal of Economic Theory* 53 (1), pp. 261–286.

— (1995). Corporate incentives for hedging and hedge accounting. *Review of Financial Studies* 8 (3), pp. 743–772.

Devalkar, S.K., R. Anupindi, and A. Sinha (2011). Integrated optimization of procurement, processing, and trade of commodities. *Operations Research* 59 (6), pp. 1369–1381.

— (2018). Dynamic risk management of commodity operations: Model and analysis. *Manufacturing and Service Operations Management* 20 (2), pp. 317–332.

Dreyfus, S.E. (1957). An analytic solution of the warehouse problem. *Management Science* 4 (1), pp. 99–104.

Elmachtoub, A.N. and P. Grigas (2017). Smart "predict, then optimize". *Working Paper*.

Energy Quants (2018). Energy Quants - Software. Retrieved from `http://www.energyquants.com/software/`.

Eydeland, A. and K. Wolyniec (2003). *Energy and power risk management - New developments in modeling, pricing, and hedging*. 1st ed. Wiley Finance.

Fabian, T., J.L. Fisher, M.W. Sasieni, and A. Yardeni (1959). Purchasing raw material on a fluctuating market. *Operations Research* 7 (1), pp. 107–122.

Fama, E.F. and K.R. French (1987). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business* 60 (1), pp. 55–73.

FIA (2015). 2015 annual survey: Global derivatives volume. *Market Voice - The Magazine of the Global Futures, Options and Cleared Swaps Markets*.

Financial Times (2015). Oil traders eye floating storage options. Retrieved from `https://www.ft.com/content/8be5347a-9a32-11e4-9602-00144feabdc0`.

Froot, K.A., D.S. Scharfstein, and J.C. Stein (1993). Risk management: Coordinating corporate investment and financing policies. *The Journal of Finance* 47 (5), pp. 1629–1658.

Garstka, S.J. and R. Wets (1974). On decision rules in stochastic programming. *Mathematical Programming* 7 (1), pp. 117–143.

Gaur, V. and S. Seshadri (2005). Hedging inventory risk through market instruments. *Manufacturing and Service Operations Management* 7 (2), pp. 103–120.

Gavirneni, S. (2004). Periodic review inventory control with fluctuating purchasing costs. *Operations Research Letters* 32 (4), pp. 374–379.

Gavirneni, S. and T.E. Morton (1999). Inventory control under speculation: Myopic heuristics and exact procedures. *European Journal of Operational Research* 117 (2), pp. 211–221.

Geman, H. (2005). *Commodities and commodity derivatives*. 1st ed. Wiley Finance.

— (2007). Mean reversion versus random walk in oil and natural gas prices. *Advances in Mathematical Finance*. Ed. by M.C. Fu, R.A. Jarrow, J.-Y. Yen, and R.J. Elliott. 1st ed. Springer, pp. 219–228.

Geman, H. and V. Nguyen (2005). Soybean inventory and forward curve dynamics. *Management Science* 51 (7), pp. 1076–1091.

Geman, H. and W. Smith (2013). Theory of storage, inventory and volatility in the LME base metals. *Resources Policy* 38 (1), pp. 18–28.

Georghiou, A., W. Wiesemann, and D. Kuhn (2005). The decision rule approach to optimisation under uncertainty: Methodology and applications in operations management. *Working Paper.*

Gibson, R. and E.S. Schwartz (1990). Stochastic convencience yield and the pricing of oil contingent claims. *Journal of Finance* 45 (3), pp. 959–976.

Goel, A. and G.J. Gutierrez (2011). Multiechelon procurement and distribution policies for traded commodities. *Management Science* 57 (12), pp. 2228–2244.

Goel, A. and F. Tanrisever (2017). Financial hedging and optimal procurement policies under correlated price and demand. *Production and Operations Management* 26 (10), pp. 1924–1945.

Golabi, K. (1985). Optimal inventory policies when ordering prices are random. *Operations Research* 33 (3), pp. 575–588.

Gray, J. and P. Khandelwal (2004a). Towards a realistic gas storage model. *Commodities Now June 2004*, pp. 1–4.

— (2004b). Realistic gas storage models II: Trading strategies. *Commodities Now September 2004*, pp. 1–5.

Guthrie, G. (2009). *Real options in theory and practice.* 1st ed. Oxford University Press.

Haksöz, Ç. and S. Seshadri (2007). Supply chain operations in the presence of a spot market: A review with discussion. *Journal of the Operational Research Society* 58 (11), pp. 1412–1429.

Hamilton, J.D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57 (2), pp. 357–384.

— (1990). Analysis of time series subject to changes in regimes. *Journal of Econometrics* 45 (1-2), pp. 39–70.

Hastie, T., R. Tibshirani, and J. Friedman (2013). *The elements of statistical learning: Data mining, inference, and prediction.* 1st ed. Springer Series in Statistics.

Haushalter, D. (2000). Financing policy, basis risk, and corporate hedging: Evidence from oil and gas producers. *The Journal of Finance* 55 (1), pp. 107–152.

Heath, D. (2018). Macroeconomic factors in oil futures markets. *Forthcoming in Management Science.*

Hull, J.C. (2005). *Fundamentals of futures and options markets.* 5th ed. Pearson Prentice Hall.

— (2018). *Options, futures, and other derivatives.* 9th ed. Pearson Prentice Hall.

Hyndman, R. and A. Koehler (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting* 22 (4), pp. 679–688.

Inderfurth, K., P. Kelle, and R. Kleber (2018). Inventory control in dual sourcing commodity procurement with price correlation. *Central European Journal of Operationsl Research* 26 (1), pp. 93–119.

Iyer, A.V. and L.E. Schrage (1992). Analysis of the deterministic (s,S) inventory problem. *Management Science* 38 (9), pp. 1299–1313.

Janczura, J. and R. Weron (2012). Efficient estimation of Markov regime-switching models: An application to electricity spot prices. *Advances in Statistical Analysis* 96 (3), pp. 385–407.

Kaldor, N. (1939). Speculation and economic stability. *The Review of Economic Studies* 7 (1), pp. 1–27.

Kalymon, B.A. (1971). Stochastic prices in a single-item inventory purchasing model. *Operations Research* 19 (6), pp. 1434–1458.

Kingsman, B.G. (1969). Commodity purchasing. *Operational Research Quarterly* 20 (1), pp. 59–79.

— (1985). *Raw materials purchasing: An operational research approach*. 1st ed. Pergamon Press.

Kleindorfer, P. and L. Li (2005). Multi-period VaR-constrained portfolio optimization with applications to the electric power sector. *Energy Journal* 26 (1), pp. 1–26.

Kristoufek, L. and M. Vosvrda (2013). Commodity futures and market efficiency. *Working Paper*.

Kyos (2018). KyStore - Gas storage optimization software. Retrieved from `https://www.kyos.com/energy-asset-optimization/gas-storage-optimization/`.

Lacima (2018). Lacima Analytics - Valuation and optimization suite. Retrieved from `https://www.lacimagroup.com/page18759/ValuationampOptimisationSuite.aspx`.

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

Li, C.-L. and P. Kouvelis (1999). Flexible and risk-sharing contracts under price uncertainty. *Management Science* 45 (10), pp. 1378–1398.

LME (2017). A detailed guide to the London metal exchange. Retrieved from `https://www.lme.com/-/media/Files/Brochures/Detailed-Guide-to-the-LME.pdf?la=en-GB`.

Lovejoy, W.S. (1990). Myopic policies for some inventory models with uncertain demand distributions. *Management Science* 36 (6), pp. 724–738.

Malkiel, B.G. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives* 17 (1), pp. 59–82.

Mandl, C. and S. Minner (2019a). *When do commodity spot price regimes matter for inventory managers?* Working paper. TUM School of Management. Technical University of Munich.

— (2019b). *Data-driven optimization for commodity procurement under price uncertainty.* Working paper. TUM School of Management. Technical University of Munich.

Mandl, C., S. Nadarajah, S. Minner, and S. Gavirneni (2019). *Structured data-driven operating policies for commodity storage.* Working paper. TUM School of Management. Technical University of Munich.

Martínez-de-Albeníz, V. and D. Simchi-Levi (2005). A portfolio approach to procurement contracts. *Production and Operations Management* 14 (1), pp. 90–114.

— (2006). Mean-variance trade-offs in supply contracts. *Naval Research Logistics* 53 (1), pp. 603–616.

MathWorks (2018). Natural gas storage valuation.

McKinsey & Company (2013). Resource revolution: Tracking global commodity markets - Trends survey 2013.

Merton, R.C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3 (1-2), pp. 125–144.

Modigliani, F. and M.H. Miller (1958). The cost of capital, corporation finance and the theory of investment. *The American Economic Review* 48 (3), pp. 261–297.

Mohri, M., A. Rostamizadeh, and A. Talwalkar (2012). *Foundations of machine learning.* 1st ed. MIT Press.

Monahan, G.E. (1982). A survey of partially observable Markov decision processes: Theory, models, and algorithms. *Management Science* 28 (1), pp. 1–16.

Nadarajah, S. and N. Secomandi (2018). Merchant energy trading in a network. *Operations Research* 66 (12), pp. 1304–1320.

Nadarajah, S., F. Margot, and N. Secomandi (2015). Relaxations of approximate linear programs for the real option management of commodity storage. *Management Science* 61 (12), pp. 3054–3076.

Nascimento, J.M. and W.B. Powell (2008). Optimal approximate dynamic programming algorithms for a general class of storage problems. *Princeton University, Department of Operations Research and Financial Engineering.*

Oliver Wyman (2017). The endgame for commodity traders - Why only the biggest and digitally advanced traders will thrive. Retrieved from `http://www.oliverwyman.com/our-expertise/insights/2017/nov/the-endgame-for-commodity-traders.html`.

Perlin, M. (2015). MS_Regress - The Matlab package for Markov regime switching models. Retrieved from `http://ssrn.com/abstract=1714016`.

Pindyck, R.S. (2004). Volatility and commodity price dynamics. *The Journal of Futures Markets* 24 (11), pp. 1029–1047.

Pindyck, R.S. and J.J. Rotemberg (1990). The excess co-movement of commodity prices. *The Economic Journal* 100 (403), pp. 1173–1189.

Pirrong, G. (2011). *Commodity price dynamics - A structural approach.* 1st ed. Cambridge University Press.

Porteus, E.L. (2002). *Foundations of stochastic inventory theory.* 1st ed. Stanford Business Books.

Reuters (2009). Recession delivers boom in commodity storage. Retrieved from `https://www.reuters.com/article/columns-us-column-storage/recession-delivers-boom-in-commodity-storage-john-kemp-idUSTRE54E4S020090515`.

— (2016a). CME bids to boost its metal storage network, challenge LME. Retrieved from `https://www.reuters.com/article/us-cme-lme-warehouses/exclusive-cme-bids-to-boost-its-metal-storage-network-challenge-lme-idUSKCN0YE0GS`.

— (2016b). Israeli shipping data firm plans push into commodities markets. Retrieved from `https://www.reuters.com/article/us-tech-shipping-windward-idUSKCN0WI2RZ`.

— (2018). U.S. copper stocks sprint to record high, reflecting transport costs, weak demand. Retrieved from `https://www.reuters.com/article/us-comex-copper-stocks/u-s-copper-stocks-sprint-to-record-high-reflecting-transport-costs-weak-demand-idUSKCN1G61JU`.

Sachs, A-L. and S. Minner (2014). The data-driven newsvendor with censored demand observations. *International Journal of Production Economics* 149 (3), pp. 28–36.

Samuelson, P.A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review* 6 (2), pp. 41–49.

Scarf, H. (1959). Bayes solutions of the statistical inventory problem. *The Annals of Mathematical Statistics* 30 (2), pp. 490–508.

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

Schwartz, E. and J. Smith (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science* 46 (7), pp. 893–911.

Secomandi, N. (2010). Optimal commodity trading with a capacitated storage asset. *Management Science* 56 (3), pp. 449–467.

— (2015). Merchant commodity storage practice revisited. *Operations Research* 63 (5), pp. 1131–1143.

Secomandi, N. and S. Kekre (2014). Optimal energy procurement in spot and forward markets. *Manufacturing and Service Operations Management* 16 (2), pp. 270–282.

Secomandi, N. and D.J. Seppi (2012). Real options and merchant operations of energy and other commodities. *Foundations and Trends in Technology, Information and Operations Management* 6 (3-4), pp. 161–331.

Secomandi, N., G. Lai, F. Margot, A. Scheller-Wolf, and D.J. Seppi (2015). Merchant commodity storage and term structure model error. *Manufacturing and Service Operations Management* 17 (3), pp. 302–320.

Seifert, R.W., U.W. Thonemann, and W.H. Hausman (2004). Optimal procurement strategies for online spot markets. *European Journal of Operational Research* 152 (3), pp. 781–799.

Sethi, S.P. and F. Cheng (1997). Optimality of (s,S) policies in inventory models with Markovian demand. *Operations Research* 45 (6), pp. 931–939.

Shapiro, A., D. Dentcheva, and A. Ruszczyński (2009). *Lectures on stochastic programming*. 1st ed. Society for Industrial and Applied Mathematics (SIAM).

Sioshansi, R., P. Denholm, T. Jenkin, and J. Weiss (2009). Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics* 31 (2), pp. 269–277.

Smith, C.W. and R.M. Stulz (1985). The determinants of firms' hedging policies. *Journal of Financial and Quantitative Analysis* 20 (4), pp. 391–405.

Song, J.-S. and P. Zipkin (1993). Inventory control in a fluctuating demand environment. *Operations Research* 41 (2), pp. 351–370.

Stoll, S.O. and K. Wiebauer (2010). A spot price model for natural gas considering temperature as an exogenous factor with applications. *Journal of Energy Markets* 3 (3), pp. 113–128.

Stulz, R.M. (1984). Optimal hedging policies. *Journal of Financial and Quantitative Analysis* 19 (2), pp. 127–140.

Tauchen, G. (1986). Finite-state Markov chain approximations to univariate and vector autoregressions. *Economics Letters* 20 (2), pp. 177–181.

Telser, L.G. (1958). Futures trading and the storage of cotton and wheat. *Journal of Political Economy* 66 (3), pp. 233–255.

Treharne, J.T. and C.R. Sox (2002). Adaptive inventory control for nonstationary demand and partial information. *Management Science* 48 (5), pp. 607–624.

Tufano, P. (1996). Who manages risk? An empirical examination of risk management practices in the gold mining industry. *Journal of Finance* 51 (4), pp. 1097–1137.

— (1998). The agency costs of corporate risk management. *Financial Management* 27 (1), pp. 67–77.

Turcic, D., P. Kouvelis, and E. Bolandifar (2015). Hedging commodity procurement in a bilateral supply chain. *Manufacturing and Service Operations Management* 17 (2), pp. 221–235.

Unilever (2016). 2016 Annual report and accounts. Retrieved from `https://www.unilever.com/investor-relations/annual-report-and-accounts/archive-of-unilever-annual-report-and-accounts.html`.

Vapnik, V.N. (1998). *Statistical learning theory.* 1st ed. Wiley.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics* 5 (2), pp. 177–188.

Vo, M.T. (2009). Regime-switching stochastic volatility: Evidence from the crude oil market. *Energy Economics* 31 (5), pp. 779–788.

Wall Street Journal (2002). Deadline scramble - A new hazard for recovery. *The Wall Street Journal* 25.

— (2008). Against the grain: Food firms hedge costs. Retrieved from `http://www.wsj.com/articles/SB120632251149658521`.

— (2011a). Companies stock up as commodities prices rise. Retrieved from `http://www.wsj.com/articles/SB10001424052748704775604576120533736097682`.

— (2011b). Steel-price increases creep into supply chain. Retrieved from `https://www.wsj.com/articles/SB10001424052748704775604576120382801078352`.

— (2015). Unseen stashing pays off for metal traders. Retrieved from `https://www.wsj.com/articles/unseen-stashing-pays-off-for-metal-traders-1432881202`.

Wall Street Journal (2016a). Airlines pull back on hedging fuel costs. Retrieved from `http://www.wsj.com/articles/airlines-pull-back-on-hedging-fuel-costs-1458514901`.

— (2016b). The end of economic forecasting. Retrieved from `https://www.wsj.com/articles/the-end-of-economic-forecasting-1466723149`.

Wang, Y., C. Wu, and L. Yang (2015). Hedging with futures: Does anything beat the naïve hedging strategy? *Management Science* 61 (12), pp. 2870–2889.

Wang, Z.F. and A.J. Mersereau (2017). Bayesian inventory management with potential change-points in demand. *Production and Operations Management* 26 (2), pp. 341–359.

Williams, J.C. and B.D. Wright (1991). *Storage and commodity markets.* 1st ed. Cambridge University Press.

Working, H. (1949). The theory of price of storage. *The American Economic Review* 39 (6), pp. 1254–1262.

Wu, D.J. and P.R. Kleindorfer (2005). Competitive options, supply contracting, and electronic markets. *Management Science* 51 (3), pp. 452–466.

Wu, O., D. Wang, and Z. Qin (2012). Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management* 14 (3), pp. 455–471.

Yang, X. and Y. Xia (2009). Acquisition management under fluctuating raw material prices. *Production and Operations Management* 18 (2), pp. 212–225.

Yu, W. and G.B. Sheblé (2006). Modeling electricity markets with hidden Markov model. *Electric Power Systems Research* 76 (6-7), pp. 445–451.

Zsidisin, G.A. and J.L. Hartley (2012). *Managing commodity price risk - A supply chain perspective.* 1st ed. Business Expert Press.