

MODELLING AND CONTROLLING RISK IN ENERGY SYSTEMS

A THESIS SUBMITTED TO THE UNIVERSITY OF MANCHESTER
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
IN THE FACULTY OF ENGINEERING AND PHYSICAL SCIENCES

2015

Jhonny Gonzalez
School of Mathematics

Contents

Abstract	15
Declaration	16
Copyright Statement	17
Acknowledgements	18
Dedication	19
1 Introduction	20
1.1 The Autonomic Power System	20
1.2 Risk in energy markets	22
1.3 Outline of Thesis	24
2 Background	29
2.1 Risk-sensitive optimal switching	29
2.2 Risk-sensitive optimal control	32
2.3 Overview of risk-sensitive control and related work	34
2.4 Energy price risk	36
2.5 Convex and coherent measures of risk	37
2.5.1 Value at risk	39
2.5.2 Conditional value at risk	40
2.5.3 Entropic value at risk	41
2.5.4 Entropic risk measure	42
2.6 Risk-sensitive criterion versus risk measures	42
2.7 Bayesian inference for multi-factor models	44
2.7.1 Bayesian statistical inference	45

2.7.2	Basic theory on Markov chains	47
2.8	Markov Chain Monte Carlo algorithms	49
3	Risk-sensitive optimal switching	57
3.1	Preliminaries	58
3.2	Problem formulation	59
3.3	Iterative optimal stopping	63
3.4	Main properties of Snell envelopes	65
3.5	Properties of the auxiliary processes	67
3.6	Characterisation of optimal solutions	74
3.7	Special cases	88
4	RS-LSM algorithm	91
4.1	Discrete time approximation	92
4.1.1	Time discretisation error	96
4.2	A least-squares Monte Carlo approach	97
4.2.1	Extension of the algorithm for control-dependent drift coefficient	100
4.3	The RS-LSM algorithm	102
4.3.1	Switching regions	104
4.4	Choice of basis functions	105
4.5	Variance reduction techniques	110
5	Bayesian inference on mixture of OU processes	115
5.1	Stylised facts of energy prices	116
5.2	Stochastic modelling of energy markets	122
5.2.1	A multi-factor model for energy spot prices	123
5.3	Literature review	126
5.4	Bayesian inference for a sum of two OU processes	127
5.4.1	Data augmentation based on a marked Poisson process	130
5.4.2	A partially non-centred parameterisation	131
5.5	MCMC implementation	134
5.5.1	Update step for the latent process Φ	137
5.6	Bayesian inference for a sum of three OU processes	140

5.7	Analysis of algorithms on simulated data	142
5.7.1	2-OU model	142
5.7.2	3-OU model	148
5.7.3	Diagnostics	152
5.8	Application to real data	154
5.9	APXUK index for electricity spot prices	155
5.9.1	Deterministic part	155
5.9.2	Stochastic part, 2-OU model	157
5.9.3	Stochastic part, 3-OU model	159
5.9.4	Diagnostics	164
5.10	EEX Phelix index for electricity spot prices	169
5.10.1	Deterministic part	169
5.10.2	Stochastic part, 2- and 3-OU models	171
5.10.3	Diagnostics	176
5.10.4	Prior sensitivity analysis	177
5.11	National grid SAP for gas spot prices	179
5.11.1	Deterministic part	180
5.11.2	Stochastic part, 2- and 3-OU models	182
5.11.3	Diagnostics	183
6	Numerical examples	186
6.1	Energy tolling agreements	188
6.1.1	Example 1. An on/off optimal switching problem	190
6.2	Energy storage	198
6.2.1	Example 2. Without seasonality	202
6.2.2	Example 3. A depleted gas reservoir on the UK market	208
6.3	Flexible district energy systems	215
6.3.1	Example 4. District energy systems	216
7	Conclusions and future work	229
	Bibliography	232

List of Tables

4.1	Convergence of the Monte Carlo standard error (SE) for the value function of Example 1 in Section 6.1.1 using the RS-LSM and the RN-LSM algorithms. The results show the sample mean of $V(t_0, X_0, i_0)$ across 60 separate runs of each algorithm as the number of sample paths N_p increases (top to bottom) and when using different random number generation techniques (left to right).	113
5.1	True value of parameters and the prior distributions used to asses the performance of the MCMC algorithm on the 2-OU model.	143
5.2	For different values of the true intensity rate η , the sample mean and standard deviation of the MCMC posterior means across 60 simulated paths. The posterior means were obtained by running the MCMC algorithm two million iterations, removing the first half-million as burn-in and thinning the chain one every 100.	145
5.3	The true value and prior distribution of each parameter of the 3-OU model, alongside the properties of the posterior means (sample mean and standard deviation) across 60 samples paths. For each path the MCMC algorithm was run two million iterations, the first half-million were removed as burn-in and the resulting chain was thinned one every 100.	149
5.4	Fit of the seasonal trend parameters in equation (5.21) to the APXUK time series.	156

5.5	Fit of the autocorrelation function of the APXUK time series using superpositions of OU processes. $n + 1$ denotes the number of OU components, the ω_i are the estimated weights of the exponential functions, the λ_i are the estimated times to mean reversion and SS is the sum of squares resulting from the procedure.	157
5.6	Prior distributions and posterior properties obtained for the APXUK dataset when fitting the 2-OU model. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the chain one every 100. The latent process Φ was updated five times per update of the model parameters.	158
5.7	Summary of posterior properties for each parameter when fitting the 2- and 3-OU models to the APXUK series. The results were obtained by running two million MCMC iterations, removing 500000 as burn-in and thinning the resulting chain one every 100. The latent process Φ was updated five times per update of the remaining parameters.	163
5.8	Summary statistics of the APXUK time series and simulated data generated with the estimated 2- and 3-OU models. The results for simulated data are based on the grand average across 500 simulated sample paths. The standard deviation is provided in parenthesis and Q_i denotes the i -th quartile.	168
5.9	Fit of the seasonal trend parameters in equation (5.21) to the EEX time series.	170
5.10	Fit of the autocorrelation function of the EEX time series using superpositions of OU processes. n denotes the number of OU components, the ω_i are the estimated weights of the exponential functions, the λ_i are the estimated times to mean reversion and SS is the sum of squares resulting from the procedure.	171

5.11	Summary of posterior properties for each parameters when fitting the 2- and 3-OU models to the EEX series. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100. The latent process Φ was updated 5 times per update of the remaining parameters. The process Y_2 is used to model negative jumps.	172
5.12	Summary statistics of the EEX time series and simulated data generated with the estimated 2- and 3-OU models. The results for simulated data are based on the grand average across 500 simulated sample paths. The standard deviation is provided in parenthesis and Q_i denotes the i -th quartile.	178
5.13	Summary of posterior properties for each parameters when fitting the 2- and 3-OU models to the EEX series. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100. The latent process Φ was updated five times per update of the remaining parameters. The process Y_2 is used to model negative jumps.	180
5.14	Fit of the autocorrelation function of the SAP time series using superpositions of OU processes. $n+1$ denotes the number of OU components, the ω_i are the estimated weights of the exponential functions, the λ_i are the estimated times to mean reversion and SS is the sum of squares resulting from the procedure.	181
5.15	Summary of posterior properties for each parameter when fitting the 2- and 3-OU model to the gas SAP series. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100. The latent process Φ was updated 5 times per update of the remaining parameters. The process Y_2 is used to model negative jumps.	183

5.16	Summary statistics of the SAP time series and simulated data generated with the estimated 2- and 3-OU models. The results for simulated data are based on the grand average across 500 simulated sample paths. The standard deviation is provided in parenthesis and Q_i denotes the i -th quartile.	185
6.1	Main characteristics of the examples considered in this chapter.	187
6.2	Parameter estimates in hourly units of the spot price model parameters in Example 4.	220
6.3	Estimated coefficients of the mean price level functions f^e and f^g in Example 4.	220

List of Figures

4.1	The estimated continuation values obtained with different sets of basis functions. The risk-sensitive value function must be strictly negative. The positive continuation values obtained by regressing against monomial functions (top panel) violate this assumption. The piecewise polynomials (bottom panel) provide a better fit and respect the negativity assumption.	107
4.2	The computational time versus the Monte Carlo standard error (SE) obtained with different random number generation techniques. The results are for the value function $V(t_0, X_0, i_0)$ of Example 1 in Section 6.1.1, using $\rho = 0.5$. The standard error (blue and orange lines) are given with respect to left y -axis, while the computational time (green line) is given with respect to the right y -axis.	114
5.1	Typical patterns of price jumps in the APXUK and EEX indexes for daily average electricity spot prices. The x -axis is given in days. . . .	118
5.2	Trace plots of η corresponding to simulated datasets from the 2-OU model with different true values of η , $\eta_{true} = \{0.05, 0.1, 0.2, 0.3\}$	146
5.3	Autocorrelation function of the Markov chain of η for sample paths simulated with $\eta_{true} = \{0.05, 0.1, 0.2, 0.3\}$. The results were obtained by running the MCMC algorithm two million iterations, removing 500000 as burn-in and thinning the chain one every 100.	147
5.4	The ACF of η_1 when fitting the 3-OU model to simulated data, using one (blue line) and five (orange line) update(s) of the latent process Φ per update of the remaining parameters.	150

5.5	The figure shows the MCMC chain (left), the posterior distribution (middle) and the ACF (right) of the model parameters (top to bottom) when fitting the 3-OU model to simulated data. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100. . . .	151
5.6	Posterior of p -values obtained when testing the estimated noise process of Y_0 for Normality using the Kolmogorov-Smirnov test at the 5% significance level. The results were obtained when fitting the 3-OU model to simulated data.	154
5.7	The daily average APXUK electricity spot log-prices along with fitted seasonal trend (top graph), and the resulting deseasonalised time series (bottom graph).	157
5.8	The figure shows the MCMC chain (left), the posterior distribution (middle) and the ACF (right) of the model parameters (top to bottom) when fitting the 2-OU model to the APXUK dataset. The results were obtained by running two million MCMC iterations, removing a burn-in period of 500000 and thinning the resulting chain one every 100. The latent process Φ was updated five times per update of the model parameters.	160
5.9	The figure shows the MCMC chain (left), the posterior distribution (middle) and the ACF (right) of the model parameters (top to bottom) when fitting the 3-OU model to the APXUK dataset. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100. . . .	162
5.10	The deseasonalised APXUK time series superimposed to a MCMC estimate of the path of the Gaussian OU process, and the corresponding estimate of the latent process Φ for the 2-OU model and the 3-OU model.	165
5.11	Posterior of p -values obtained when testing the estimated noise process of Y_0 for Normality using the Kolmogorov-Smirnov test at the 5% significance level.	167

5.12	The daily average APXUK electricity spot log-prices along with a fitted seasonal trend (top graph), and the resulting deseasonalised time series (bottom graph).	170
5.13	The figure shows the MCMC chain (left), the posterior distribution (middle) and the ACF (right) of the model parameters (top to bottom) when fitting the 3-OU model to the EEX dataset. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the chain one every 100.	174
5.14	The deseasonalised EEX time series superimposed to a MCMC estimate of the path of the Gaussian OU process, and the corresponding estimate of the latent process Φ for the 2-OU model and the unrestricted 3-OU model.	175
5.15	Posterior of p -values obtained when testing the estimated noise process of Y_0 for Normality using the Kolmogorov-Smirnov test at the 5% significance level.	177
5.16	The daily average SAP series of gas spot log-prices along with fitted seasonal trend (top graph), and the resulting deseasonalised time series (bottom graph).	182
5.17	Posterior of p -values obtained when testing the estimated noise process of Y_0 for Normality using the Kolmogorov-Smirnov test at the 5% significance level.	184
6.1	Top panel: The switching boundaries of $V(t, X_t, i)$ for Example 1 obtained with different degrees of risk-sensitivity ρ , superimposed to a sample path of the spark spread process X_t . Bottom panel: The corresponding cumulative revenue of the power plant operator for the given path of the spark spread and each value of ρ	192
6.2	Top panel: The switching boundaries of $V(t, X_t, i)$ for Example 1 obtained with different degrees of risk-sensitivity ρ , superimposed to a sample path of the spark spread process X_t . Bottom panel: The corresponding cumulative revenue of the power plant operator for the given path of the spark spread and each value of ρ	194

6.3	For the total revenue function of Example 1, its mean, standard deviation (SD) and risk measures VaR, CVaR and EVaR all at the 95% confidence level, as a function of the degree of risk-sensitivity ρ . In this example the greater the risk measures the smaller the risk faced by the plant operator. The results are based on 16000 sample paths of the spark spread process (X_t)	196
6.4	For Example 1, the net change with respect to the risk-neutral case $\rho = 0$ of the mean and risk measures of the total revenue, as the risk-sensitivity ρ increases. A negative value indicates a decrease with respect to the risk-neutral case. The results are based on 16000 sample paths of the spark spread process (X_t)	197
6.5	For the total revenue function of Example 2, its mean, standard deviation (SD) and risk measures VaR, CVaR and EVaR all at the 99% confidence level, as a function of the degree of risk-sensitivity ρ . In this example the greater the risk measures the smaller the risk faced by the plant operator. The results are based on 2000 sample paths.	205
6.6	For Example 2, the net change with respect to the risk-neutral case $\rho = 0$ of the mean and risk measures of the total revenue, as the risk-sensitivity ρ increases. A negative value indicates a decrease with respect to the risk-neutral case. The results are based on 2000 sample paths of the gas price model.	205
6.7	Switching regions from regime store obtained for different values of the risk-sensitivity parameter ρ . The marker colours in each graph indicate the optimal operating regime for each pair (G_t, C_t) of gas-inventory values: red for regime inject, black for regime store and blue for regime withdraw. The results are based on 2000 simulations of the gas price process.	207

6.8	Custom box plots of the total revenue function of Example 3. The two lower whiskers below each box indicate the risk measures $\text{CVaR}_{95\%}$ and $\text{CVaR}_{99\%}$. The red line is the median, the edges (top and bottom) of the box are the 25th and 75th percentiles, the upper whisker is the highest datum that is not an outlier, i.e. it is still within 1.5IQR (interquartile) of the upper quartile.	212
6.9	Change with respect to the risk-neutral case $\rho = 0$ in the expected total revenue (mean), and risk measures $\text{VaR}_{99\%}$ and $\text{CVaR}_{99\%}$ as ρ increases. A negative value means that the corresponding index has decreased. . .	212
6.10	For Example 3, the switching regions from regime 'store' as function of time and gas price when the gas inventory level is fixed and equal to 30 Bcf. The top panel presents the risk-neutral case, whereas the bottom panel presents the risk-sensitive case with $\rho = 0.072$	214
6.11	The smart district energy system studied in Example 4.	217
6.12	The winter heat and electricity demand profiles of Example 4.	217
6.13	For the total cost function of Example 4, its mean, standard deviation (SD) and risk measures VaR , CVaR and EVaR all at the 99% confidence level, as a function of the degree of risk-sensitivity ρ . The smaller the value of the risk measures the smaller the associated risk. The results are based on 5000 sample paths.	222
6.14	For Example 4, the net change with respect to the risk-neutral case $\rho = 0$ of the mean and risk measures of the total revenue, as the risk-sensitivity ρ increases. A negative value indicates a decrease with respect to the risk-neutral case. The results are based on 5000 sample paths.	222
6.15	Optimal operating regimes for the gas boiler (top) and the CHP unit (bottom) at 9am on 2nd January, using risk-neutral control (left) and risk-sensitive control with $\rho = 0.15$ (right). The current level of stored heat is 4.5 MWh.	224

6.16	Contours of the value function $V(t, (Y_0^e, Y_1^e, Y_0^g, Y_1^g), C_t, i; \rho)$ for Example 4 as function of the gas price factors Y_0^g and Y_1^g when $Y_0^e = 1, Y_1^e = 0$. The time to maturity is 24 hours, $\rho = 0$, the inventory level is $C_t = 0$ and the current operating regime is $i = C_0 B_0$	226
6.17	Contours of the value function $V(t, (Y_0^e, Y_1^e, Y_0^g, Y_1^g), C_t, i; \rho)$ for Example 4 as function of the electricity price factors Y_0^e and Y_1^e when $Y_0^g = 0, Y_1^g = 0$. The time to maturity is 24 hours, $\rho = 0$, the inventory level is $C_t = 0$ and the current operating regime is $i = C_0 B_0$	228

The University of Manchester

Jhonny Gonzalez

Doctor of Philosophy

Modelling and controlling risk in energy systems

23rd June 2015

The Autonomic Power System (APS) grand challenge was a multi-disciplinary EPSRC-funded research project that examined novel techniques that would enable the transition between today's and 2050's highly uncertain and complex energy network. Being part of the APS, this thesis reports on the sub-project 'RR2: Avoiding High-Impact Low Probability events'. The goal of RR2 is to develop new algorithms for controlling risk exposure to high-impact low probability (Hi-Lo) events through the provision of appropriate risk-sensitive control strategies. Additionally, RR2 is concerned with new techniques for identifying and modelling risk in future energy networks, in particular, the risk of Hi-Lo events.

In this context, this thesis investigates two distinct problems arising from energy risk management. On the one hand, we examine the problem of finding managerial strategies for exercising the operational flexibility of energy assets. We look at this problem from a risk perspective taking into account non-linear risk preferences of energy asset managers. Our main contribution is the development of a risk-sensitive approach to the class of optimal switching problems. By recasting the problem as an iterative optimal stopping problem, we are able to characterise the optimal risk-sensitive switching strategies. As byproduct, we obtain a multiplicative dynamic programming equation for the value function, upon which we propose a numerical algorithm based on least squares Monte Carlo regression.

On the other hand, we develop tools to identify and model the risk factors faced by energy asset managers. For this, we consider a class of models consisting of superposition of Gaussian and non-Gaussian Ornstein-Uhlenbeck processes. Our main contribution is the development of a Bayesian methodology based on Markov chain Monte Carlo (MCMC) algorithms to make inference into this class of models. On extensive simulations, we demonstrate the robustness and efficiency of the algorithms to different data features. Furthermore, we construct a diagnostic tool based on Bayesian p -values to check goodness-of-fit of the models on a Bayesian framework. We apply this tool to MCMC results from fitting historical electricity and gas spot price datasets corresponding to the UK and German energy markets. Our analysis demonstrates that the MCMC-estimated models are able to capture not only long- and short-lived positive price spikes, but also short-lived negative price spikes which are typical of UK gas prices and German electricity prices.

Combining together the solutions to the two problems above, we strive to capture the interplay between risk, uncertainty, flexibility and performance in various applications to energy systems. In these applications, which include power stations, energy storage and district energy systems, we consistently show that our risk management methodology offers a tradeoff between maximising average performance and minimising risk, while accounting for the jump dynamics of energy prices. Moreover, the tradeoff is achieved in such way that the benefits in terms of risk reduction outweigh the loss in average performance.

Declaration

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright Statement

- i. The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the “Copyright”) and s/he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.
- ii. Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made **only** in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.
- iii. The ownership of certain Copyright, patents, designs, trade marks and other intellectual property (the “Intellectual Property”) and any reproductions of copyright works in the thesis, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.
- iv. Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see <http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=487>), in any relevant Thesis restriction declarations deposited in the University Library, The University Library’s regulations (see <http://www.manchester.ac.uk/library/aboutus/regulations>) and in The University’s Policy on Presentation of Theses.

Acknowledgements

It has been a pleasure to work with my supervisor Dr John Moriarty. I thank him for his careful guidance, great support, helpful advice, and for getting me interested in the exciting world of energy markets. This work is the result of the many meetings and conversations we have had over the last three years.

Many thanks also to Dr Jan Palczewski for his expertise and collaboration over the past year and a half. I am grateful to the investigators in the Autonomic Power System project, in particular, to Dr Joseph Mutale, Dr Pierluigi Mancarella, Prof Janusz Bialek and Dr Chris Dent. I am also grateful to Dr Christiana Charalambous and Dr Alexandros Beskos for agreeing to be my examining committee and carefully reading this thesis.

The School of Mathematics at Manchester University has been a quite inspiring place to work in. Thanks to all other PhD students, post-docs and academics for making a very pleasant environment. I am thankful to my colleagues and friends Steven Falconer and Jonathan Schachter for always fruitful academic and non-academic conversations. Special thanks also go out to Randall Martyr for providing useful suggestions to improve Chapter 3.

I would also like to thank my family, who have supported me in so many ways during my studies. Infinitely many thanks to Mildred for her invaluable love and patience. I know many dreams have come true only because you exist in my life.

Finally, I would like to thank the School of Mathematics at Manchester University and EPSRC grant EP/I031650 for their generous support.

Dedication

To Mildred

Chapter 1

Introduction

1.1 The Autonomic Power System

The research conducted in this thesis was part of the Autonomic Power System (APS)¹ Grand Challenge (McArthur et al., 2012), a multi-disciplinary multi-institutional research project between University of Cambridge, Durham University, Imperial College London, King’s College London, Manchester University, Newcastle University, University of Strathclyde and University of Sussex, which was funded by EPSRC and had the participation of the industrial partners Accenture, Agilent, E.On, IBM, KEMA, Mott MacDonald, PB Power, National Grid and SSE.

The focus of the APS project is the electricity network of 2050. The vision is that this network will be shaped by numerous key drivers (including, among others, increasing energy prices, electric vehicles functioning as roving loads and energy storage, massive use of renewable energy and increased customer participation) which will bring a significant increased level of uncertainty and complexity into the power system. The challenge is to deliver the fundamental changes in the power system that will facilitate the transition between today’s and the future’s energy network. By analogy with the biological autonomic nervous system, the autonomic power system is intended to be a fully self-managing system, which optimises its operational decisions in real-time by using a decentralised configuration. Based on high-level goals, the APS will control itself, optimise itself, protect itself and heal itself (McArthur et al., 2012). This type of architectural frameworks for complex systems was started by IBM in the context of

¹<http://autonomicpowersystem.org.uk/>

computer science (Kephart and Chess, 2003).

To achieve this, the APS project identified four key, interrelated areas where major research and advancement is necessitated:

1. Self* Network Operation and Control. This workstream analyses and evaluates the autonomic control philosophy, determining the level of decentralisation that can be obtained and examining the stability properties of an APS in a complex and uncertain energy network.
2. Autonomic Economics. This workstream focuses on economic decision-making problems on the operational (short-term) and investment (long-term) time scale, developing new bespoke frameworks and methodologies for the APS.
3. Resilience and Risk management. The aim of this workstream is to study the resilience and risk associated with the operation of an APS, while developing new techniques to quantify and manage risk and taking into account information available in real-time and the high-degree of uncertainty in the system.
4. Active Participation of Consumers. This workstream investigates key aspects such as behavioural patterns of energy consumption and customers preferences, which will enable active participation of customers and small energy producers in an APS.

This thesis was part of the Resilience and Risk management (RR) workstream, within the project ‘RR2: Avoiding High-Impact Low Probability (Hi-Lo) events’. This project is motivated by the assumption that risk, uncertainty, flexibility and performance will play key roles in the future electrical network envisioned by the APS grand challenge. RR2 pays special attention to a class of risk events so-called Hi-Lo events. As their name indicate, these risk events have a relatively small probability of occurrence, but once they happen, their impact on the system can be catastrophic. From a economic perspective, Hi-Lo events can be extremely large variations in energy prices due to shocks on energy demand or supply. From a power system security perspective, they can be cascading failures (Dobson et al., 2007) leading to blackouts.

In any case, the aim of RR2 is to model these risk factors and develop appropriate techniques to minimise their probability, and thus control risk exposure, while

accounting for uncertainty, flexibility and average performance of the system. For expositional clarity, this thesis is focused on economic risk. Nonetheless, the tools and methodologies developed here are widely applicable outside this context, by appropriately defining the type of risk of interest and the underlying risk factors to be modelled.

1.2 Risk in energy markets

The worldwide liberalisation of energy markets, which started in Chile in the early 80s and was followed by the UK and many other countries in Europe and the USA over the last 25 years, has significantly transformed the way in which the energy industry operates (Weron, 2007). The implications of this liberalisation process are numerous. On the one hand, this process has given place to the introduction of competitive energy markets, which have replaced the traditional deterministic way of pricing energy. In these markets, energy prices are stochastic, being set at the temporary balance between supply and demand arising from selling and buying bids of producers and wholesale consumers. On the other hand, liberalised energy markets have considerably increased the level of risk faced by active market players. For economical reasons, their decision-making process must now be based on the relative movements of the market, in sympathy with the physical constraints of their energy assets and uncontrollable factors such as weather, which greatly influence both production and demand. Finally, being in a very capital intensive business, energy market players must take informed operational (and planning) decisions and adopt suitable risk management strategies in this tough environment. Thus, the energy industry depends heavily on appropriate quantitative tools for modelling risk factors such as volatile energy prices, volume production, and demand levels.

These implications are aggravated by taking into account the strong effect that large shocks in either energy supply or demand may have on energy prices. In particular, the strong variability and uncertainty of weather conditions, alongside generation outages and interruptions in transmission, pose a major risk to both producers and retailers (Benth, 2013). For example, demand for electricity increases during the cold season in the Nordic area due to extended use of household heating. In the US, it is the

relatively high temperature over summer that increases demand, owing to excessive use of air conditioning. Additionally, generation of hydropower is strongly influenced by seasonal changes in reservoir levels because of rainfall and snow melting. In most energy markets, these sudden increase in demand or decrease in supply translate into high and volatile energy prices, which in turn may change according to short-term weather conditions and demand levels.

In this economical landscape, active market participants face many types of risk including market risk, weather risk, credit risk, and operational risk, among others (Simkins and Simkins, 2013). Of particular importance it is price risk, i.e. the risk arising from the uncertainty associated with movements in energy prices. As mentioned earlier, these price movements can have large variations due to different factors affecting supply or demand. For instance, power prices may increase 10-fold in a matter of hours or a few days; also, they tend to exhibit a extreme daily volatility that can be up to 50%, while other commodities (such natural gas and crude oil) and financial instruments have a daily volatility that goes only up to 5% (Weron, 2007).

Under these conditions, energy risk management has come to play a key role in the decision-making process of energy assets' managers, in particular, in regards to the operational strategies for power generation plants and energy storage facilities (see for example Thompson et al. (2004), Carmona and Ludkovski (2008) and Felix (2012), among others). Risk management can be defined as “the process and tools used for evaluating, measuring and managing the various risks within a company’s portfolio of financial, commodity and other assets” (Pilipovic, 2007). It is useful to think of energy price risk management as involving a number of steps, including (Simkins and Simkins, 2013):

1. Identification of risk factors through price modelling.
2. Assessment of risk impact and likelihood (quantification of risk exposure).
3. Finding options to mitigate risk.
4. Implementation of mitigation risk options.

The main purpose of this thesis is to develop new tools to tackle the first and third steps for the risk management of energy assets. More specifically, the goal of this

thesis is twofold. On the one hand, our goal is to propose new price risk mitigation options for energy asset managers. To do this, we develop a methodology for finding optimal control strategies that exploit the operational flexibility of energy assets. Our approach takes into account the price risk associated with energy markets and strikes a balance between profitability (or operating costs) and risk. Our main tool for doing this is the theory of stochastic risk-sensitive optimal control (Whittle, 1990), and its connections with dynamic measures of risks (Barrieu and El Karoui, 2008).

On the other hand, since any operational strategy must account for the actual dynamics of risk factors, in our case, of energy prices, our second goal is to study and develop models to capture the key features exhibited by energy prices (both gas and electricity). In these premises, we aim to provide an efficient Bayesian method for making inference into multi-factor models for energy spot prices. The models are based on mixture of Gaussian and non-Gaussian Ornstein-Uhlenbeck processes and appropriate modelling of seasonal components.

By combining the proposed methodologies, we strive to capture the interplay between risk, uncertainty, flexibility and performance in various applications to energy systems. In these examples, we illustrate the implementation of our risk mitigation approach and quantify the resulting risk exposure with sophisticated, modern measures of risk. We notice that these are the second and fourth risk management steps given above. Implementing our risk management approach and quantifying risk in these numerical applications allows us to illustrate the potential benefits of the tools develop in this thesis.

1.3 Outline of Thesis

This thesis develops methods to tackle two distinct but related mathematical problems lying at the interface between the theory of stochastic optimal control and energy markets. The backbone of the thesis is composed of Chapter 3 (theoretical analysis of the control problem), Chapter 4 (numerical algorithm for the solution of the control problem) and Chapter 5 (Bayesian inference of multi-factor models and applications to simulated and real data). In Chapter 6, we put together the results of these previous chapters in various examples related to the operational control of energy assets. For

the reader convenience, the following paragraphs provide an outline of the thesis and the main research contributions of our work.

Chapter 2

The general purpose of Chapter 2 is to provide an introduction to stochastic risk-sensitive optimal control, the modern theory of measures of risk and Bayesian inference based on Markov Chain Monte Carlo (MCMC) methods. Also,

1. We establish a connection between risk-sensitive control and convex measures of risk.
2. We show that in special cases minimising the classical risk-sensitive optimisation criterion based on exponential functions is equivalent to minimising coherent measures of risk.

Chapter 3

In Chapter 3, we develop a risk-sensitive approach to optimal switching problems. We start by reviewing the notion of Snell envelope of processes and setting up our stochastic control problem. In addition,

1. We re-formulate the risk-sensitive optimal switching problem as an iterative optimal stopping problem for controlled diffusions.
2. We introduce suitable auxiliary processes which aid to solve the control problem, and prove in Proposition 3.7 a number of properties which they satisfy.
3. In Proposition 3.8, we characterise the optimal control strategy for a particular case of the general problem.
4. In Proposition 3.9, we provide a recursive characterisation of the optimal control for the general risk-sensitive optimal switching problem.
5. As byproduct, we obtain a multiplicative dynamic programming (MDP) equation for the risk-sensitive optimal switching problem, equations (3.10) and (3.23).

6. We conclude the chapter by showing that the risk-sensitive optimal switching formulation generalises two particular control problems typically encountered in the literature.

Chapter 4

The purpose of Chapter 4 is to provide a simulation-based algorithm for the solution of the risk-sensitive optimal switching problem. Here

1. We re-write the MDP equation and propose two recursive methods for its numerical solution.
2. We provide an algorithm based on least squares Monte Carlo (LSM) methods to approximate the conditional expectations in the MDP equation.
3. We also propose a suitable choice of basis functions for the numerical algorithm. The basis functions adapt to the distribution of the underlying stochastic process and ensure the sign of the value function of the control problem is respected.
4. We demonstrate the potential benefits of combining the least squares methodology with Quasi-Monte Carlo techniques to improve the rate of converge of the proposed algorithm.

Chapter 5

In Chapter 5, we make Bayesian inference into a mixture of Gaussian and non-Gaussian Ornstein-Uhlenbeck (OU) processes for energy spot price modelling. In addition to reviewing key features of energy spot prices, we provide empirical evidence of the existence of different patterns of energy price spikes, which call for using mixture of jump models. Also,

1. We propose a missing data methodology for making Bayesian inference into a general class of models based on Gaussian and non-Gaussian OU processes.
2. We develop efficient MCMC algorithms for tackling the inference problem.
3. We run extensive simulations in order to test the proposed algorithms, and show their robustness to different data features.

4. We develop quantitative diagnostics tools for model checking based on Bayesian p -values.
5. We then apply the algorithms to real datasets consisting of electricity and gas spot prices in the UK and German energy markets.
6. Using the proposed diagnostic tool, we suggest a suitable minimum number of components in the mixture model for the real datasets under study. In particular, we demonstrate the need to use different jump components to better capture not only positive price jumps but also negative ones.

Chapter 6

The main goal of Chapter 6 is to illustrate by example the potential benefits of our risk management approach to the operation of energy assets. This chapter brings together the novel risk-sensitive algorithm of Chapter 4 with the multi-factor jump-diffusion models of Chapter 5. The illustrations are composed of thought and real world examples of increasing level of complexity. In particular, we study the risk-sensitive stochastic optimal control of power stations, energy storage and district energy systems. Additionally,

1. We compare our risk-sensitive approach with the classical risk-neutral approach to solving optimal switching problems.
2. We propose an intuitive, general way to choose a ‘best’ degree of risk-sensitivity on each of the examples.
3. We illustrate risk-averse attitudes from part of energy asset controllers, and show the effect of price seasonality on their managerial decisions.
4. In all the examples, we consistently show that our methodology strikes a balance between average performance and risk (as measured by different risk indices). The balance is struck in such way that the decrease in risk index is greater than the loss in average performance.
5. In our final and most complex example, we apply our risk-sensitive control algorithm to district energy systems using a 4-dimensional jump-diffusion model

for gas and electricity prices. Here, we successfully demonstrate that our methodology is able to learn the jump dynamics of energy prices and produce appropriate control strategies that exploit the operational flexibility of the system while balancing average performance and risk.

Let us turn now to a brief introduction of the stochastic control problem addressed in this thesis.

Chapter 2

Background

2.1 Risk-sensitive optimal switching

In the first part of the thesis we investigate a stochastic control problem known as optimal switching. It is convenient to introduce this problem with an example. To do this, let us consider a power plant whose electricity output is sold in a liberalised energy market. As mentioned earlier, in these markets electricity prices are subject to random fluctuations due to different factors affecting power demand and supply. These factors can be e.g. weather, power outages, statistical correlation with other energy commodities (such as gas and oil) and governmental policies related to renewable generation. As a consequence, future electricity prices are unknown, albeit statistical models can be constructed using available market information. Moreover, since currently electricity cannot be directly and economically stored in large amounts, it must be consumed almost at the same time it is produced.

We suppose that the power plant has available a finite number of operating regimes. For example, the unit can be either off-line, working at full capacity, or working at some given intermediate modes (e.g. half capacity, $3/4$ capacity, etc). These operating regimes represent the operational flexibility of the power plant, or equivalently the operational options available to the manager of the plant. By exercising one of these options, the manager receives the market price of electricity proportional to the chosen working capacity. We assume however that changing the operating modes of the power plant is not free. These costs cover the extra fuel required to change regimes (e.g. startup, ramp-up and shutdown costs) and various overhead costs. It is also natural

to assume that there are operating expenses associated with any regime, for example, due to ongoing operation and maintenance (O&M) costs.

Under these assumptions, the manager aims to find an ‘optimal’ management strategy for the power plant. In other words, as time evolves, the manager needs to decide (i) when it is optimal to switch from the current operating regime to another one, and (ii) to which new operating regime she should switch, provided a decision to change regimes has been taken.

This managerial strategy can be seen as a double sequence $u := (\tau_n, \xi_n)_{n \geq 0}$ of increasing switching times τ_n , and the corresponding new operating regimes ξ_n . Obviously, the strategy must take into account the current and future dynamics of electricity prices. The problem of finding this type of control strategies is known as *optimal switching* (between the different operating regimes).

The optimality of the management strategy is defined depending on the optimisation objective of the manager. One possible objective is to choose a switching control strategy maximising the average total reward of the plant manager over a given horizon $[0, T]$. Let us denote by $P(u; [0, T])$ the total reward of the manager when implementing the control policy u over this horizon. Further, assume that the initial regime of the unit at τ_0 is $\xi_0 = i$. Informally, the aim is to find a managerial strategy u that maximises the risk-neutral criterion

$$\mathbb{E}(P(u; [0, T])).$$

This problem is known as a *risk-neutral* optimal switching problem (Djehiche et al., 2009). An alternative objective is to minimise the ‘risk’ associated with operating the unit, or to find a tradeoff between the expected total reward and risk, in the sense that for a reduce level of risk the manager is willing to sacrifice some of the expected profit. In this case, the optimisation problem can be written as

$$\min_u \kappa(-P(u; [0, T])),$$

where κ defines some measure of risk. This type of switching control problems are the main focus of the thesis and we shall refer to them as *risk-sensitive* optimal switching problems. More specifically, we are interested in switching control problems where

$\kappa(\cdot)$ is a nonlinear measure of risk of the form

$$\kappa_\rho(X) = \frac{1}{\rho} \log \mathbb{E}(e^{\rho X}), \rho \in \mathbb{R} \setminus \{0\}.$$

We notice that the risk-neutral objective falls in this second case (the risk minimisation case) by taking $\kappa(X) = \mathbb{E}(-X)$. At this point we will not make any formal assumptions and leave the mathematical definition of measures of risk until Section 2.5. Risk-sensitive optimal switching problems are formally studied in Chapter 3.

Why do we want to control risk?

In the above example, controlling energy price risk, for instance, is motivated by the high levels of price volatility typically observed in power markets, which can be up to 50% higher than in other commodity or financial markets (Weron, 2007). These extreme volatility levels are mainly due to the non-storability of electricity and changing weather conditions, which affect both electricity demand and supply. More details about this are given in Section 5.1.

Under these premises, control policies that only maximise total average reward (with the risk-neutral optimisation criterion above) may not be appropriate to manage price risk. In fact, as we shall illustrate in various examples in Chapter 6, these risk-neutral control policies may lead optimisers to a large amount of risk exposure. It is also worth mentioning that when optimising the risk-neutral criterion the actual realised total profit of the manager can be far from the average value. Hence indicators such as the median of the distribution of the total reward (if the distribution is skew) and its variance are also relevant in the optimisation objective. Moreover, although large profits relative to the mean are fine (indeed, any rational operator does or should welcome this scenario), unusually low total profits (say, extremely below the average, and even negative ones) are a matter of concern to the power plant operator, and so she may want to avoid these scenarios and/or limit its exposure to price risk. By using a risk optimisation objective, we aim to find suitable control policies to manage the downside risk faced by the manager of the power plant.

We remark that those scenarios giving unusually low total profits may be classified as *high-impact low-probability* events. Intuitively, once they happen, although with a

small probability, the impact they have can be catastrophic from the manager perspective. For example, these events may lead to unbearable losses, credit defaults, and ultimately, put the operator of the power plant out of business. In this thesis, we will be concerned with the price risk associated with operating energy assets in liberalised energy markets. However, the results of the thesis are also applicable to other cases calling for risk management.

2.2 Stochastic risk-sensitive optimal control

As mentioned earlier, in the risk-neutral approach to solving stochastic control problems the aim is to optimise the expected value of a reward functional over a specified time horizon. For expositional clarity, in the following discussion we choose the equivalent problem of minimising the expected value of a cost functional. Accordingly, let us denote by $C = C(u)$ the specified cost functional, where u is the control being applied. Then the risk-neutral optimal control problem is to find an admissible control that minimises

$$\mathbb{E}(C), \tag{2.1}$$

while an alternative is to minimise the criterion

$$\gamma(\rho) = \rho^{-1} \log \mathbb{E}(e^{\rho C}), \tag{2.2}$$

where ρ is a scalar. In general, the problem of finding an optimal control that minimises the criterion (2.2) is known as a stochastic risk-sensitive optimal control problem (Whittle, 1990, Fleming and Soner, 2006). This class of optimisation problems was first introduced by Jacobson (1973). Hereafter we refer to the parameter ρ as the *risk-sensitivity parameter*, and call the criterion (2.2) *the risk-sensitive criterion*. As we shall see in our discussion below, and in fact throughout the thesis, the parameter ρ represents an extra degree of freedom that the optimiser can tune in order to indicate their preference to risk.

Some intuition behind the optimisation criterion (2.2) is obtained by finding its Taylor series expansion around the point $\rho = 0$. Let us denote by $\text{Var}(C)$ the variance of C and assume that $\rho \text{Var}(C)$ is sufficiently small. Then it is straightforward to show

that (Whittle, 1981)

$$\rho^{-1} \log \mathbb{E}(e^{\rho C}) = \mathbb{E}(C) + \frac{\rho}{2} \text{Var}(C) + O(\rho^2). \quad (2.3)$$

From this equation we can readily see that the relative values of ρ around zero define different local properties of the optimiser. For a fixed value of $\mathbb{E}(C)$, when $\rho < 0$ the criterion prefers controls producing large variability of the cost C around its expected value, as this improves the criterion. Conversely, when $\rho > 0$ the criterion favours controls producing small variability of C about $\mathbb{E}(C)$. Moreover, when $\rho \rightarrow 0$ it follows that $\gamma(\rho) \rightarrow \mathbb{E}(C)$, i.e. in this case we recover the risk-neutral criterion (2.1). Interestingly, this latter case implies that risk-sensitive optimal control can be seen as a general approach that includes risk-neutral optimal control as a limiting case.

The conceptual intuition of the risk-sensitive criterion extends to the case when ρ is large, by considering the nonlinear relation that the exponential function puts to the smaller and larger values of C . For example, when $\rho > 0$, the aim is to minimise the expectation of $e^{\rho C}$, which is a convex decreasing function of the cost functional C (note that $\rho > 0$ and the logarithmic function is increasing, so we can either minimise $\mathbb{E}(e^{\rho C})$ or $\gamma(\rho)$). With such criterion the larger (smaller) values of C are given relatively more (less) weight with increasing ρ , when compared to the linear risk-neutral objective. Therefore, as ρ increases the optimiser seeks to find controls that minimise C , but she puts more and more concern on controlling the unusually large values of C (the worst cases) than its potential small values (the best cases). Further details on the interpretation of the risk-sensitive criterion can be found in Ploeg (1984) and Whittle (1990).

Traditionally, the cases $\rho < 0$, $\rho > 0$ and $\rho = 0$ are said to indicate a risk-seeking, risk-averse and risk-neutral attitude from part of the controller (Whittle, 1990), owing to the differentiated preferences of cost variability of the optimiser. In this thesis we are particularly interested in the risk-averse case and its relation with the risk-neutral one. We would like to clarify, however, that in applications to economics before the 1990's it was usual to quantify 'risk' in terms of the variance of random variables (see for example Markowitz (1952) and Brennan and Schwartz (1985)). The current paradigm for measuring risk goes beyond the variance, and is much more sophisticated

and suitable for financial applications. We will mention this point again in Section 2.5, where we formally review the modern concept of measures of risk.

2.3 Overview of risk-sensitive control and related work

The generalisation of stochastic optimal control theory by means of the risk-sensitive criterion was pioneered by Jacobson (1973). In his seminal work Jacobson was mainly interested in a discrete-time setting of the so-called linear quadratic Gaussian (LQG) systems, where one assumes linear dynamics, quadratic costs and Gaussian noise. The possibility to incorporate risk sensitivity in control problems influenced several extensions, until the general solution to the linear exponential quadratic Gaussian (LEQG) model was given by Whittle (1981). An advantage of the LEQG control problem is that it can be solved analytically, and thus allows a direct interpretation of the optimal control solutions, see for example Whittle (1990) and Ploeg (1984). The solution to the continuous-time version of the LEQG model is due to Bensoussan and van Schuppen (1985).

Over the last three decades many authors have further developed and analysed the theory of risk-sensitive optimal control, including different research areas such as:

- Connections with major topics such as robust control, differential games and backward stochastic differential equations (BSDEs) (Fleming and McEneaney, 1992, James, 1992, Bensoussan et al., 1998, Dupuis et al., 2000, El-Karoui and Hamadène, 2003).
- Characterisation of the value function and optimal control policy, by means of partial differential equations (PDEs) and probabilistic methods (Nagai, 1996, Bensoussan et al., 1998, Bensoussan and Nagai, 2000).
- Applications to risk management within mathematical finance, and connections with measures of risk (Bielecki and Pliska, 1999, Nagai, 2003, Hata and Iida, 2006, Hata and Sekine, 2010, Davis and Lleo, 2011, Barrieu and El Karoui, 2008).

Relevant to this thesis are the aforementioned second and third areas of research.

During the 1990s many authors worked on deriving solutions to the non-linear case using PDE methods. In this line of research, Nagai (1996) seems to be the first to link the solution of risk-sensitive control problems to the existence of the solution of a Bellman equation (a parabolic PDE in the dynamic programming formulation of the problem). This author considers a risk-sensitive control problem whereby the system dynamics are governed by a controlled diffusion process and the cost is the exponential of non-quadratic functions. Under quite restrictive assumptions on the range of values of the risk-sensitivity parameter, he proves that the solution of the Bellman equation coincides with the value function of the control problem.

To relax some of those assumptions Bensoussan et al. (1998) use more advanced PDE methods, and also prove a uniqueness result for the value function of the risk-sensitive control problem. Their results allow to solve the problem in a much more general framework. The more recent paper by Bensoussan and Nagai (2000) relaxes further these assumptions, so larger values of ρ can be used without breaking stability conditions.

Nonetheless, improving the assumptions on ρ has come at the cost of using several sophisticated methods from the theory of PDEs. Moreover, the main focus of the above papers has been on finding explicit solutions of the control problem. For realistic applications however the conditions are quite restrictive. The more recent paper by Hata and Sekine (2010) deals also with the non-linear case but limits the analysis to analytical solutions. See also Fleming and Soner (2006) and the recent overview on risk-sensitive control by Nagai (2014).

Recently, the particular class of optimal switching problems (illustrated at the beginning of this chapter) has become popular since they provide a powerful framework to capture the interplay between flexibility and uncertainty in many problems arising from applications to the energy sector. See for example Carmona and Ludkovski (2008, 2010), Aïd et al. (2014) and Kitapbayev et al. (2015). Optimal switching problems have long been studied starting with Brennan and Schwartz (1985) and then Brekke and Øksendal (1994). More recently, Hamadène and Jeanblanc (2007) solved completely the simplest case of these problems, namely, the case of two operating regimes. Extensions to the multi-regime case and further advances in these area are due to

Djehiche et al. (2009), Asri and Hamadene (2009), Pham et al. (2009), Hamadène and Zhang (2010), among others. All these authors deal with the risk-neutral formulation of the stochastic control problem. However, little attention has been paid to the risk-sensitive approach to optimal switching problems and their numerical solution.

One of the goals of this thesis is to generalise stochastic risk-sensitive optimal control to the case of optimal switching problems. For this, we use probabilistic methods based on the notion of Snell envelope of stochastic processes. In our setting the state dynamics are governed by a controlled diffusion, as in the early PDE papers. Our main contribution in this area is to recast the problem as a multiple optimal stopping problem, which then allows to characterise the optimal control policy recursively. In Chapter 3, we shall see that as byproduct of this characterisation we obtain a multiplicative version of the dynamic programming principle. Additionally, we suggest an efficient method for the numerical solution of the risk-sensitive optimal switching problem and demonstrate in various applications the potential benefits of risk-sensitive control when compared to the risk-neutral approach for the risk management of energy systems.

2.4 Energy price risk

A rigorous mathematical treatment of the theory of risk goes back to 1783 with the work of Daniel Bernoulli (1954). Over the last century the mean-variance framework developed within modern portfolio was very popular in economics. The variance, as a measure of risk, however, does not account for the asymmetry proper of financial variables, where one is mainly interested in downside risk (Föllmer and Schied, 2002b).

It was not until the stock market crisis of 1987 that a suitable and consistent way to quantify risk was more carefully considered. The axiomatic framework under which the theory of risk lies today is due to Artzner et al. (1999). Since the introduction of this work, the concept of risk and appropriate methods to quantify it has been the object of study by economists and mathematicians. See Föllmer and Schied (2010) for a recent overview of this subject.

In this thesis we adopt this axiomatic framework, and refer to risk as the uncertainty related with the future value of a (financial) position. In particular, we shall be

interested in energy price risk, namely, the uncertainty arising from changes in energy prices over time. As mentioned earlier, the volatility of wholesale electricity prices can be extreme, in some cases doubling the price volatility of other commodity and financial markets. Thus controlling and managing energy price risk is of fundamental importance for market participants operating energy assets in deregulated markets.

Our focus on energy risk management lies; on the one hand, in the identification of risk factors through price modelling (which we discuss in Chapter 5); and, on the other hand, in the development of risk mitigation options by designing algorithms to produce appropriate managerial strategies for energy assets. As illustrated in the power plant example at the beginning of this chapter, managerial strategies define the future and uncertain value of our objective function (which, in the example, refers to striking a balance between total reward and risk exposure). Thus, in general, we shall consider the price risk of adopting a specified managerial strategy, whether it is risk-neutral or risk-sensitive, for a given energy system.

Since the mean-variance framework is not appropriate for the financial applications we study in this thesis, an important question is then how to suitably quantify price risk. We address this question next by reviewing the modern axiomatic framework of risk measures, which we will use throughout the thesis.

2.5 Convex and coherent measures of risk

In this section, we briefly present some of the most popular risk measures used in the literature and review their properties and suitability for quantifying risk in a financial context. These risk measures shall be used in the numerical examples of Chapter 6 to measure the associated risk of following a given control policy, and thus to compare risk-neutral and risk-sensitive control strategies.

Loosely speaking, a risk measure is a mapping that assigns a real number to a random outcome or risk position. The random outcome may represent, for instance, the accumulated wealth or the operating costs associated with running a system over a specified period of time. The assigned number given through the risk measure summarises the amount of risk associated with the given outcome. In this thesis, we are concerned with the class of coherent risk measures, first introduced by Artzner

et al. (1999), and its extension to the class of convex risk measures by Föllmer and Schied (2002a).

To be precise, let us fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω denotes the sample space, \mathcal{F} is a σ -algebra of events of Ω and \mathbb{P} is probability measure defined on \mathcal{F} . Also, let us denote by $\mathbf{L}(\Omega, \mathcal{F}, \mathbb{P})$ the set of all Borel measurable functions and by $\mathbf{L}_\infty(\Omega, \mathcal{F}, \mathbb{P})$ the set of all bounded Borel functions including constants. We will write \mathbf{L} and \mathbf{L}_∞ if no confusion is possible, and assume in the definition below that $X \in \mathbf{L}_\infty$ represents a loss functional.

Definition 2.1. The mapping $\kappa : \mathbf{L}_\infty(\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ is a convex risk measure, if for every $X_1, X_2 \in \mathbf{L}_\infty$ and $c \in \mathbb{R}$, $m \in [0, 1]$ it satisfies the following properties:

1. Monotonicity: If $X_1 \leq X_2$, then $\kappa(X_1) \leq \kappa(X_2)$.
2. Translation invariance: $\kappa(X + c) = \kappa(X) + c$.
3. Convexity: $\kappa(mX_1 + (1 - m)X_2) \leq m\kappa(X_1) + (1 - m)\kappa(X_2)$.

A convex risk measure is called a coherent risk measure if it also satisfies

1. Positive homogeneity: if $c \geq 0$, then $\kappa(cX) = c\kappa(X)$.

We note that under the assumption of positive homogeneity, convexity is equivalent to the subadditivity property:

1. Subadditivity: For every $X_1, X_2 \in \mathbf{L}_\infty$, $\kappa(X_1 + X_2) \leq \kappa(X_1) + \kappa(X_2)$.

These properties admit intuitive interpretations in economical terms (Artzner, 1999, Föllmer and Schied, 2002a, 2010), which can be easily translated to energy systems, considering for instance a portfolio of energy assets consisting of two power plants. In these premises, the convexity property makes precise the notion of diversification: the sum of the individual risks involved with the operation of each unit should not increase by considering the overall risk of their combined operation. In fact, this global risk can be limited by controlling the separate risk associated with each power plant. On the other hand, the positive homogeneity implies that the risk increases linearly with the size of the financial position of a power plant. Although quite intuitive,

the positive homogeneity property has been replaced with the convexity property in order to account for nonlinear impacts of the size of risky positions on the risk measure (Föllmer and Schied, 2002a, Barrieu and El Karoui, 2008). Thus convex risk measures can be seen as a generalisation of coherent risk measures, which nevertheless remain popular in the literature and practical applications in risk management.

Elementary examples

Some elementary examples of risk measures are in order.

- The expectation operator defines a coherent risk measure,

$$\kappa_E(X) := \mathbb{E}(X).$$

- Also, the mapping defined by the essential supremum of random variables,

$$\kappa_{max}(X) := \text{ess sup } X,$$

is a coherent risk measure. This risk measure can be interpreted as the worst-case scenario implied by the realisations of X . See Lemma 3.3 for a definition of the essential supremum of a random variable.

- The mean-variance risk measure introduced by Markowitz is related to the risk-sensitive criterion through equation (2.3),

$$\kappa_{ms}(X) := \mathbb{E}(X) + \gamma \sqrt{\text{Var}(X)}, \gamma > 0.$$

While this risk measure satisfies properties 2. and 3., it does not satisfy the monotonicity property, and therefore it is neither convex nor coherent.

Other popular and more sophisticated choices of risk measures are as follows.

2.5.1 Value at risk

The value at risk (VaR) measure is one of the most popular and widely used risk measures. For a confidence level $1 - \alpha$, the VaR of X is defined as the $(1 - \alpha)$ -percentile

of the distribution of X ,

$$\text{VaR}_{1-\alpha}(X) := \inf_z \{z : \mathbb{P}(X \leq z) \geq 1 - \alpha\}. \quad (2.4)$$

Typical values of α are in the range $[0.01, 0.1]$. The VaR provides an intuitive interpretation of the risk position X . Roughly, with a confidence level of $1 - \alpha$, the VaR specifies that the potential losses (or costs) that may be obtained by following a given control strategy will not exceed the value indicated by VaR. We note however that VaR does not give any indications about the potential losses in the case that the VaR is exceeded.

Although the VaR satisfies properties 1, 3, 4, it does not satisfy the subadditivity property in general. For instance, this property is not satisfied when X_1 and X_2 are Normally distributed and $\alpha > 0.5$. Hence the risk measure VaR is neither coherent nor convex. The implications that this has on risk management are acute (Artzner et al., 1999, p. 218). For example, VaR may discourage or even prevent diversification, in the sense that controlling the individual risk positions X_1 and X_2 (even if they are independent) may not limit the risk involved with their combination $X_1 + X_2$.

2.5.2 Conditional value at risk

The conditional value at risk (CVaR) is a coherent risk measure (Rockafellar and Uryasev, 2000) that has been proposed as a suitable alternative to overcome the weaknesses of VaR. For a confidence level $1 - \alpha$, the CVaR is related to the VaR by the formula

$$\text{CVaR}_{1-\alpha}(X) := \frac{1}{\alpha} \int_0^\alpha \text{VaR}_{1-t}(X) dt. \quad (2.5)$$

The CVaR may be interpreted as the average loss that can be incurred provided the VaR has been exceeded. Intuitively, the CVaR says “how bad is bad”. For this reason, CVaR is also known as expected shortfall or tail conditional expectation to make explicit its interpretation.

The CVaR measure is more conservative or risk-averse than VaR, in the sense that for a given random outcome CVaR is in general larger than the VaR at the same confidence level. Since the end of the 1990’s, the CVaR risk measure has become more

and more popular thanks to the fact that it is coherent and straightforward to interpret. Furthermore, in stochastic optimisation problems requiring limiting risk, CVaR has been widely used as an optimisation criterion. However, except in special cases, numerical calculations of the CVaR tend to be computationally expensive (Ahmadi-Javid, 2012b) and thus other coherent risk measures have been proposed such as the entropic value at risk.

2.5.3 Entropic value at risk

The entropic value at risk (EVaR) has been recently introduced and studied by Ahmadi-Javid (2012b). The EVaR measure provides the tightest upper bound for both VaR and CVaR that can be obtained from the Chernoff inequality (Chernoff, 1952). The EVaR of $X \in \mathbf{L}_\infty$ at the confidence level $1 - \alpha$ is defined by the quantity

$$\text{EVaR}_{1-\alpha}(X) := \inf_{z>0} \{z^{-1} \log(\alpha^{-1} M_X(z))\}, \quad (2.6)$$

where M_X is the moment-generating function of X . As shown in Ahmadi-Javid (2012b), (see also Ahmadi-Javid (2012a)), EVaR is a coherent risk measure for every $\alpha \in (0, 1]$. There are important connections between EVaR and some of the risk measures that we have presented earlier. The EVaR is a more conservative risk measure than CVaR. More specifically, it can be shown that

$$\mathbb{E}(X) \leq \text{VaR}_{1-\alpha}(X) \leq \text{CVaR}_{1-\alpha}(X) \leq \text{EVaR}_{1-\alpha}(X) \leq \text{ess sup}(X).$$

Additionally, the following equalities hold true

$$\mathbb{E}(X) = \text{VaR}_{0.5}(X) = \text{EVaR}_0(X), \quad \lim_{\alpha \rightarrow 0} \text{EVaR}_{1-\alpha}(X) = \text{ess sup}(X).$$

We notice that by setting α close to zero, the latter equality allows to explore worst-case scenarios, that is, to quantify the risk of high-impact low-probability events (recall our discussion at the beginning of this chapter). Proofs of the above expressions and further details about the EVaR measure can be found in Ahmadi-Javid (2012b).

2.5.4 Entropic risk measure

The entropic risk measure is one of the most important examples of a convex risk measure that is not coherent. The entropic risk measure with risk parameter $\theta \in \mathbb{R}$ is defined by (Kupper and Schachermayer, 2009)

$$\kappa_{er}^\theta(X) := \begin{cases} \frac{1}{\theta} \log \mathbb{E}(e^{\theta X}) & \text{if } \theta \in \mathbb{R} \setminus \{0\} \\ \mathbb{E}(X) & \text{if } \theta = 0 \end{cases}, \quad X \in \mathbf{L}_\infty.$$

From this definition, the connection between the risk-sensitive criterion in (2.2) and convex measures of risk is clear. Furthermore, the entropic risk measure highlights that our risk-sensitive optimal switching problem (recall Section 2.1) does not consist of optimising a mean-variance criterion as in (2.3). Instead we aim at optimising a convex measure of risk which enjoys the properties defined on page 38.

Apart from the properties given for the risk-sensitive criterion, the entropic risk measure has other interesting properties from a risk management perspective.

For fixed $X \in \mathbf{L}_\infty$, it is straightforward to show that the function $\theta \rightarrow \kappa_{er}^\theta(X)$ is increasing. This property means that increasing θ makes the entropic risk measure more conservative, i.e. the perceived measured risk increases. In the limit case of infinite risk-aversion we have that (Kupper and Schachermayer, 2009)

$$\lim_{\theta \rightarrow \infty} \kappa_{er}^\theta(X) = \text{ess sup}(X),$$

that is, by increasing θ the optimiser protects increasingly against worst-case scenarios. We notice that this is similar to using EVaR with decreasing α .

2.6 Risk-sensitive criterion versus risk measures

As mentioned earlier, one of our aims in this thesis is to optimise a risk-sensitive criterion, or equivalently the entropic risk measure, in problems of optimal switching. From a practitioner point of view, however, one may ask whether minimising the convex entropic risk measure is equivalent to minimising coherent risk measures such as CVaR and EVaR, or even the risk measure VaR. In general the answer is no.

Nevertheless, it is possible to obtain this equivalence by carefully choosing the risk-sensitivity parameter ρ .

Let us illustrate this point with an example we have introduced in Gonzalez and Moriarty (2014). For the purpose of illustration, let us assume momentarily that the cost functional C is Normally distributed with mean μ and variance σ^2 , and that $1 - \alpha$ is a fixed confidence level. In this case, all the VaR, CVaR and EVaR can be written in the general form

$$\kappa(C) = \mu + a\sigma, \quad (2.7)$$

with different scalars a , while the risk-sensitive criterion is given by

$$\gamma(\rho) = \mu + \rho^2\sigma. \quad (2.8)$$

For each measure VaR, CVaR and EVaR, a takes the value z_α , $\Phi(z_\alpha)/\alpha$ and $\sqrt{-2\log \alpha}$, respectively, where z_α and $\Phi(\cdot)$ denote the $(1 - \alpha)$ -percentile and the probability density function of the standard Normal distribution. We note that, for fixed ρ and $\gamma(\rho)$, equation (2.8) makes precise the tradeoff between mean and variance for Normal costs: under the given assumptions, an increase in mean must be compensated with a decrease in standard deviation, and conversely, a decrease in mean must be accompanied by an increase in standard deviation. Furthermore, combining equations (2.7) and (2.8), it follows that the risk κ and the performance γ are related by

$$\kappa(C) = \gamma(\rho) - \rho\sigma^2/2 + a\sigma, \quad (2.9)$$

and by keeping $\gamma(\rho) = P$ fixed, for some constant $P \in \mathbb{R}$, we have that κ is a quadratic function of σ attaining its maximum value at $\rho^* = a\rho^{-1}$. Let us focus on the results for CVaR. In this case, we find that when $\sigma < \rho^* = \Phi(\alpha)(\alpha\rho)^{-1}$, σ has an increasing relationship with CVaR, and when $\sigma > \Phi(\alpha)(\alpha\rho)^{-1}$ the relationship reverses, i.e. increasing the variance of C actually decreases its CVaR. We notice also that, in this specialised example, equation (2.9) implies that for a fixed value of σ , decreasing the risk-sensitivity ρ increases the risk measure κ .

Thus by choosing ρ sufficiently small such that $\sigma < \rho^*$, it follows that minimising the risk-sensitive criterion amounts to minimising the coherent risk measure CVaR. We

notice that similar results are obtained for VaR and EVaR by taking the appropriate value of a given earlier.

It is worth mentioning that if ρ is not small enough then the equivalence in the minimisation problems above no longer holds, highlighting the need to carefully choose ρ . In addition, we note that similar results are obtained for other families of distributions for the cost C . For instance, Gonzalez and Moriarty (2014) consider the case of asymmetric costs having a Weibull distribution. In Chapter 6, these analytical results will be generalised in various applications of risk-sensitive stochastic control to energy systems. There, we also explain how to choose appropriate values of ρ so that we can indirectly minimise alternative risk measures by minimising the risk-sensitive criterion.

2.7 Bayesian inference for multi-factor models

In the second part of the thesis we focus on another of the core capabilities of risk management, namely, the identification of risk factors through price modelling. This is motivated by the fact that in order to perform appropriate price risk management we first need to understand and model the dynamics of energy prices and the risk they pose to market participants.

For this, we will review in Chapter 5 key features of energy spot price dynamics. Moreover, we will focus on a class of stochastic multi-factor price models recently introduced by Benth et al. (2007). The specific problem we solve there refers to making inference in this class of models from historical data and analysing the minimum number of factors needed to suitably account for the key features of energy prices. We will postpone until Chapter 5 a thorough discussion of the problem and the context in which we work.

The main tool we use for making inference in this class of models is Bayesian statistical inference aided by a range of simulation-based methods known as Markov chain Monte Carlo algorithms. In the following we provide a review of this subject and cover the necessary background required in the thesis. Throughout, we use the customary convention of denoting by $\pi(Y)$ both the distribution and density of a random variate Y .

2.7.1 Bayesian statistical inference

In statistical inference we deal with a collection of observations $X = \{x_1, x_2, \dots, x_n\}$ having a parametric probability distribution $\pi(X|\Theta)$, where the set of parameters $\Theta = (\theta_1, \dots, \theta_k)$ is assumed unknown. The aim is to make an inference on Θ (i.e. to learn some information about Θ) from the observed data. For example, one may want to find point estimates of the parameters that better describe the observed data. In Bayesian statistics, the parameter Θ is not only assumed to be an unknown quantity, but it is treated as a random variable and assigned a probability distribution. Inference on Θ is then made based on a probability distribution which is assumed to contain all the relevant statistical information about the parameters.

More precisely, a Bayesian statistical model specifies two main ingredients:

1. A statistical model for the observed data, conditioned on the unknown parameter Θ ,

$$X \sim \pi(X|\Theta).$$

Equivalently, the model may be specified through the likelihood of the data $\ell(\Theta|X) \propto \pi(X|\Theta)$ regarded as a function of Θ .

2. A marginal distribution $\pi(\Theta)$ for Θ , which is known as the prior distribution of Θ ,

$$\Theta \sim \pi(\Theta).$$

The prior distribution (or simply the prior) incorporates our ‘beliefs’ about Θ before any observations are actually made. The parameters of the prior are called the hyperparameters.

Bayesian statistical inference combines our prior beliefs about Θ with the actual observed data in order to update the statistical information about the parameters. Such combination leads to the (joint) posterior distribution of Θ , $\pi(\Theta|X)$, the posterior for short. The posterior is linked together with the data model and the parameter model

through Bayes' theorem:

$$\pi(\Theta | X) = \frac{\ell(\Theta|X)\pi(\Theta)}{\int \ell(\Theta|X)\pi(\Theta)d\Theta} \quad (2.10)$$

$$\propto \ell(\Theta|X)\pi(\Theta) \quad (2.11)$$

i.e. Posterior \propto Likelihood \times prior.

Note that $\pi(\Theta|X)$ can be seen as the conditional distribution of Θ given the observed data, so intuitively, once observations are available the posterior reflects this new statistical information about Θ . Bayesian inference on Θ is made in terms of the posterior distribution, from which we can compute quantities of interest such as the mean, the variance and other summary statistics. It is worth mentioning that the normalising constant $1/\int \ell(\Theta|X)\pi(\Theta)d\Theta$ is in general difficult to compute, and so Bayes' theorem is typically used in the form (2.11). For a general view on the Bayesian paradigm the reader is referred to Robert (2007) or Bernardo and Smith (2009).

The following concepts are widely used in Bayesian statistics. Let us denote by θ_{-i} the set of parameters $(\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_k)$, where we have omitted the i -th component. The full conditional (posterior) distribution of θ_j is defined as the distribution of θ_j conditional on the remaining parameters and the observed data

$$\pi(\theta_j | \theta_{-j}, X).$$

It is easy to see that this distribution is proportional to the posterior distribution. In the particular case when all the full conditional distributions are known, the Gibbs sampler (described later) provides an iterative algorithm to sample from the posterior of Θ .

Another useful concept is that of a conjugate prior. A family of probability distributions $\pi(\Theta)$ parameterised by Θ is called the conjugate for a likelihood function $\ell(\Theta|X)$, if the posterior $\pi(\Theta | X)$ also belongs to that family for every $\pi(\Theta)$. In this case the posterior and the prior are said to be conjugate distributions. Conjugate

priors make up one of the most widely chosen class of prior distributions. An important result is that every element of the exponential family admits a conjugate prior (Robert, 2007). So when available, this class of conjugate priors is analytically and computationally convenient to use. Comprehensive lists of common conjugate prior distributions are found for example in DeGroot (2005) or Robert (2007). See also the compendium collected by Fink (1997).

2.7.2 Basic theory on Markov chains

Bayesian statistical inference is much facilitated by the use of iterative sampling algorithms known as Markov chain Monte Carlo (MCMC) algorithms. In this section, we briefly review some concepts related to Markov chains which will be useful for our presentation on MCMC methods. A formal introduction to Markov chains in general spaces can be found in Meyn and Tweedie (1993), which we follow here. See also Robert and Casella (2004).

We work with a general measurable space $(E, \mathcal{B}(E))$, for which the σ -algebra is countably generated. In this thesis the state space E will be usually a subset of \mathbb{R}^d for some $d \geq 1$ and $\mathcal{B}(E)$ the Borel σ -algebra. We also consider point processes for which the state space is non-Euclidean.

A Markov chain $Z = \{Z_0, Z_1, \dots\}$ is a special type of discrete-time stochastic process taking values Z_n in a space E . It is useful to think in terms of the sample paths of Z as lying on the countable product space $\Omega := \prod_{i=0}^{\infty} E_i$, where each E_i is a copy of E endowed with the σ -algebra $\mathcal{B}(E)$. The space is equipped with the corresponding product σ -algebra \mathcal{F} . Markov chains may be constructed by specifying their transition probability kernel.

Definition 2.2. A transition probability kernel P is a function taking values on $E \times \mathcal{B}(E)$ such that

1. For each $A \in \mathcal{B}(E)$, $P(\cdot, A)$ is a measurable function.
2. For each $x \in E$, $P(x, \cdot)$ is a probability measure.

Let μ_0 be an initial probability measure on $(E, \mathcal{B}(E))$. Then there exists (Meyn and Tweedie, 1993) a stochastic process $Z = \{Z_0, Z_1, \dots\}$ on the countable set Ω ,

measurable with respect to \mathcal{F} , and a probability measure P_μ on \mathcal{F} such that for every $B \subset \mathcal{F}$, $P_\mu(B)$ is the probability of the event that the process takes values in B , and for any n and measurable $A_i \subset E_i, i = 0, \dots, n$, we have

$$P_\mu(Z_0 \in A_0, \dots, Z_n \in A_n) = \int_{y_0 \in A_0} \dots \int_{y_{n-1} \in A_{n-1}} \mu(dy_0) P(y_0, dy_1) \dots P(y_{n-1}, A_n) \quad (2.12)$$

The stochastic process Z defined on (Ω, \mathcal{F}) is called a time-homogenous Markov chain with transition probability kernel P and initial distribution P_μ if for every n its finite dimensional distributions satisfy (2.12). Equation (2.12) is known as the Markov property, and can be interpreted by saying that given the present and past states of the process Z , its future path only depends on its present state. Note that the path of Z is defined by its one-step transitions $P(y_i, y_{i+1})$.

The concepts of invariant distribution and reversibility are important for constructing MCMC algorithms. These concepts specify some required stability conditions on the Markov chain Z . A σ -finite measure π on $(E, \mathcal{B}(E))$ is said to be invariant for the transition kernel P if for every $A \in \mathcal{B}(E)$

$$\pi(A) = \int_E P(x, A) \pi(dx),$$

A Markov chain with transition probability kernel P and initial distribution π is called reversible if and only if the equation (Tierney, 1998)

$$P(x, dy) \pi(dx) = P(y, dx) \pi(dy) \quad (2.13)$$

is satisfied. Relation (2.13) is known as the detailed balance equation, and must be understood as equality of measures defined on the product space $\mathcal{B}(E) \otimes \mathcal{B}(E)$. Roughly, detailed balance expresses the property that both the probability of a transition from y to x and the probability of a transition from x to y are equal. In other words, the direction of time does not affect the dynamics of a reversible Markov chain. The detailed balance equation is a sufficient condition for π to be invariant for P (Tierney, 1998).

2.8 Markov Chain Monte Carlo algorithms

Given an invariant distribution π , the purpose of MCMC algorithms is to construct a reversible Markov chain whose invariant distribution is π . In our case, the target distribution π is the posterior of the set of parameters Θ .

Over the last 25 years MCMC methods have revolutionised Bayesian analysis, providing an efficient approach to tackle highly dimensional simulation problems with relative ease of implementation. Today there is a impressive amount of literature on Bayesian inference based on MCMC methods, accounting for the theoretical background but also providing practical implementations of the algorithms. Useful references in this direction are Casella and George (1992), Tierney (1994), Chib and Greenberg (1995), the survey by Roberts et al. (2004b), and the books by Gilks et al. (1995), Robert and Casella (2004), and Robert and Casella (2009), among others.

From a practical perspective, MCMC methods are iterative simulation-based sampling methods giving as output a chain which converges asymptotically to the target distribution π . Typically, when using an MCMC algorithm we need to specify the initial state of the chain Z_0 , which may be random or deterministic. Since in general Z_0 is not necessarily drawn from π , the initial samples say Z_1, \dots, Z_B are typically removed because they may not be drawn from the stationary distribution π either. This period, $1, \dots, B$, is known as burn-in period, and for efficiency any inference is carried out with the samples $\{Z_{B+1}, Z_{B+2}, \dots\}$.

It is also worth mentioning that, in general, MCMC algorithms do not produce independent samples from π , and so one must bear in mind this when calculating characteristics of π which assume independence. A useful technique to improve serial correlation of the chain is to ‘thin’ the chain. In this case one takes the sequence of values $\{Z_{B+1}, Z_{B+2n}, \dots\}$ for some $n > 1$. In any case, we typically follow the steps below when using any MCMC algorithm

1. Initialise the chain Z_0 .
2. Simulate $N > 0$ samples $\{Z_1, \dots, Z_N\}$ by running an MCMC algorithm.
3. Check if N samples are enough for convergence, for example, with the aid of autocorrelation functions (ACFs) or by plotting the output chain.

4. Discard the first B samples as burn-in and thin if necessary.
5. Obtain relevant characteristics of π such as mean, median, variance and so forth, from $\{Z_{B+1}, Z_{B+2n}, \dots\}$.

In the following, we provide a brief presentation of the MCMC methods that we will use in the thesis.

Gibbs sampler

In the statistical literature, the Gibbs sampler was introduced by Geman and Geman (1984) for studying Gibbs distributions in the context of image processing. Drawing from this work, the algorithm was later generalised by Gelfand and Smith (1990) and employed for missing data problems by Tanner and Wong (1987). Since then the Gibbs sampler has become one of the most popular MCMC algorithms, due to its ease of implementation and versatility to divide a highly multi-dimensional simulation problem into various low-dimensional problems.

The idea behind the Gibbs sampler is as follows. Suppose we are interested in simulating from a multi-dimensional distribution π and that the state space E can be written as $E = E_1 \times E_2 \times \dots \times E_k$ for some $k > 1$. The subspaces E_j can be either uni- or multi-dimensional, for example, $E = \mathbb{R}^d, E_j = \mathbb{R}^{d_j}$ with $d = \sum_j d_j$. For any $z \in E$, let us write $z = (z_1, z_2, \dots, z_k)$, where $z_j \in E_j, j = 1, \dots, k$, and denote by $\pi_j(dz_j | z_{-j})$ the full conditional distributions of the j -th component. We assume that the conditional distributions exist and that we can simulate from them directly. The Gibbs sampler uses the knowledge of the full conditional distributions to construct a transition probability kernel P with stationary distribution π . Algorithmically, the Gibbs sampler can be written as follows:

Algorithm for the Gibbs sampler

1. Initialise the chain $Z_0 = (Z_0^{(1)}, \dots, Z_0^{(k)})$.
2. For each $i = 1, \dots, N$

Sample $Z_{i+1}^{(1)} \sim \pi_1(\cdot | z_i^{(2)}, \dots, z_i^{(k)})$

Sample $Z_{i+1}^{(2)} \sim \pi_2(\cdot | z_{i+1}^{(1)}, z_i^{(3)}, \dots, z_i^{(k)})$

$$\vdots$$

$$\text{Sample } Z_{i+1}^{(k)} \sim \pi_k(\cdot \mid z_{i+1}^{(1)}, \dots, z_{i+1}^{(k-1)})$$

Under mild regularity conditions, it follows that the Gibbs sampler produces a Markov chain with transition kernel P (Liu et al., 1994, 1995). In fact, P arises as the composition of k transition kernels $P = P_k P_{k-1} \dots P_1$, each associated with the conditional distribution $\pi_j, j = k, \dots, 1$ (Roberts et al., 2004b). Note that in this algorithm each component is visited in a deterministic way, i.e. the procedure cycles through component 1, component 2 and so on, repeating the cycle after visiting component k . This algorithm is known as the Gibbs sampler with systematic (or deterministic) scan. It is also possible to design a Gibbs sampler where the components are visited randomly, provided the probability of visiting each component is positive. More general, it can be shown that the Gibbs sampler converges to the specified invariant distribution under any visiting scheme, as long as each component is updated infinitely often (Gelfand and Smith, 1990, Liu et al., 1995).

One of the attractive features of the Gibbs sampler is that it transforms the simulation of a highly dimensional distribution into simulation of possibly univariate distributions. In this context, the use of conjugate priors may facilitate the implementation of Gibbs samplers by producing a posterior from which it is easy to simulate from.

Data augmentation algorithm (or Two-step Gibbs sampler)

The data augmentation algorithm of Tanner and Wong (1987) is a particular case of the Gibbs sampler that is widely used to tackle missing data problems. The algorithm is interesting in its own right because it is the simplest case of the Gibbs sampler (there are only two components), and its analysis has provided a good understanding of the theoretical properties of the Gibbs sampler. See for instance Liu et al. (1994, 1995).

As the name suggests, data augmentation consists of augmenting the observed data X by the ‘missing’ or latent data $Z^{(mis)}$. This is usually done in such a way that the augmented data $(X, Z^{(mis)})$ is easy to analyse. For example, suppose we are interested

in sampling from the posterior

$$\pi(\Theta|X) \propto \ell(X|\Theta)\pi(\Theta),$$

but for computational reasons it is not feasible to simulate directly from it (for instance, because the likelihood function is neither analytically nor numerically tractable). Instead, we can consider the augmented data $(X, Z^{(mis)})$ if it is now easy to sample from the distribution of the missing data $\pi(Z^{(mis)}|X, \Theta)$ and the posterior of Θ given the observed and missing data $\pi(\Theta|X, Z^{(mis)})$. In this case, we can design a Gibbs sampler iterating the steps:

1. Sample $Z_i^{(mis)} \sim \pi(\cdot|X, \Theta_{i-1})$
2. Sample $\Theta \sim \pi(\cdot|X, Z^{(mis)})$

An interesting result related to this algorithm is that in addition to the chain $\{Z^{(mis)}, \Theta\}$ being Markov, each of the marginal chains $\{Z_i^{(mis)}\}_i$ and $\{\Theta_i\}_i$ has also the Markov property (Liu et al., 1994). This is not true for the general Gibbs sampler presented earlier.

We remark here that missing data does not necessarily refer to data that has not been collected. In general, it can take many forms including difference schemes of data augmentation, latent variables, hidden states or processes, or any unobserved and unobservable variables (Meng, 2000).

Augmentation is typically suggested by the model under analysis. In our case, we shall see in Chapter 5 that it is not computationally efficient to calculate the likelihood $\ell(X|\Theta)$ associated with the multi-factor processes with which we work there. However, by augmenting the observed data by an appropriate latent process, the likelihood of the augmented data $\ell(X|\Theta, Z^{(mis)})$ becomes Gaussian. Thus, we can use a missing data methodology based on the Gibbs sampler to simulate from the posterior distribution $\pi(\Theta, Z^{(mis)}|X)$ as indicated earlier. This allows us to replace a rather difficult simulation problem with a number of easier problems. More details will be provided in Section 5.4.

Metropolis-Hastings algorithm

The use of Markov-chain simulation methods was pioneered by Metropolis et al. (1953), proposing the original Metropolis-Hastings (MH) algorithm to tackle problems in statistical physics. This algorithm was further generalised by Hastings (1970), who introduced general proposal functions, and then by Green (1995), who introduced reversible jump MH algorithms (RJMH for short) to address Bayesian inference problems where the dimension of the vector parameter Θ is not fixed. The MH algorithm is the most general MCMC algorithm, being widely applicable under very mild assumptions. A recent introduction to the MH algorithm is given by Chib and Greenberg (1995) and Robert and Casella (2004). Here we follow Tierney (1998) who studies the algorithm in general state spaces.

The MH algorithm is a technique for constructing a reversible Markov transition kernel $P(x, dy)$ on a measurable space $(E, \mathcal{B}(E))$, or equivalently, a Markov chain $\{Z_0, Z_1, \dots\}$ with a specified invariant distribution π . The algorithm requires a transition kernel $Q(x, dy)$, which is called the proposal kernel, and a measurable function $\alpha(x, y) : E \times E \rightarrow [0, 1]$, which is known as the acceptance probability. The basic step of the algorithm is as follows. When the current state of the chain is $Z_i = x$, the kernel $Q(x, \cdot)$ is used to propose its next state, say, y . This new state is accepted with probability $\alpha(x, y)$, in which case the new state of chain is $Z_{i+1} = y$, otherwise, the chain remains in its previous state $Z_{i+1} = x$. The MH algorithm can be summarised as follows:

MH algorithm

1. Initialise Z_0
2. For each $i = 0, \dots, N$
 - Sample $Y_{i+1} \sim Q(Z_i, \cdot)$
 - Sample $u \sim \text{Un}[0, 1]$
 - If $u \leq \alpha(Z_i, Y_{i+1})$
 - Set $Z_{i+1} = Y_{i+1}$
 - Else
 - Set $Z_{i+1} = Z_i$

The transition kernel constructed by the MH algorithm is given by

$$P(x, dy) = Q(x, dy)\alpha(x, y) + \delta_x(dy) \int (1 - \alpha(x, z))Q(x, dz), \quad (2.14)$$

where δ_x is the Dirac measure centred at x . Roughly, the first summand on the right of equation (2.14) accounts for the accepted transitions from x to candidates y belonging to a suitable measurable set $A \in \mathcal{B}(E)$, while the second summand captures all cases when the chain remains in x . Ensuring that the constructed Markov chain has π as its invariant distribution amounts to finding a suitable acceptance probability α . For a Markov chain Z with initial distribution π and transition kernel Q , let us denote by R the set of pairs (x, y) such that transitions from x to y and from y to x are both possible. Then under some suitable conditions (Tierney, 1998), there exists a density r on $(E \times E)$ such that the acceptance probability given by

$$\alpha(x, y) = \begin{cases} \min\{r(x, y), 1\}, & \text{if } (x, y) \in R \\ 0 & \text{otherwise} \end{cases},$$

makes π a stationary distribution for P . Explicit results can be obtained, for instance, when there exists a measure ν on $(E, \mathcal{B}(E))$ such that $\pi(dx) = \pi(x)\nu(dx)$ and the proposal kernel takes the form $Q(x, dy) = q(x, y)\nu(dy)$. In this case (Tierney, 1994)

$$R = \{(x, y) : \pi(x)q(x, y) > 0 \text{ and } \pi(y)q(y, x) > 0\},$$

and $\alpha(x, y)$ is given through the well-known acceptance ratio

$$r(x, y) = \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)}.$$

Note that the normalising constants (assumed independent of the current state x) of the densities π or $q(x, \cdot)$ need not be known since the algorithm only depends on the ratio of these densities.

A fundamental characteristic of the MH algorithm is realised by noting that the proposal kernel Q can be chosen rather arbitrarily. This makes the MH algorithm a generic and universal method which can be applied in a broad range of situations.

Nevertheless, some practical choices on the proposal density $q(x, y)$ have been given in the literature. The most common choices for metric spaces are based on random walks, for which the proposal takes the form $q(x, y) = q(|x - y|)$. In this case, the proposal moves symmetrically within a neighbourhood of the current state x , and the algorithm is called a random walk MH algorithm (RWMH for short) (Sherlock et al., 2010). A specific proposal is based on Gaussian distributions. When the current state is x , the proposal candidates are drawn from $Y \sim x + \xi$, where $\xi \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$ has a multivariate Gaussian distribution, i.e. $q(x, \cdot) = N(x, \sigma^2 \mathbf{I})$.

In the case where a component of the Markov chain is restricted to $E_j \subset \mathbb{R}^+$, it is useful to employ a multiplicative RWMH. The proposals may now be sampled from $Y \sim x\zeta$, where $\zeta \sim \text{LN}(0, \sigma^2)$. The acceptance ratio of this algorithm takes the form

$$r(x, y) = \frac{\pi(y)}{\pi(x)} \cdot \frac{y}{x}.$$

Note that the multiplicative RWMH is equivalent to the (additive) RWMH when a logarithmic scale is used, which gives $\log Y \sim x + \log(\zeta)$.

In any case special attention must be paid to the variance of the proposal distribution q . If the variance is extremely large, the chain will tend to move in large steps over the space, resulting in small acceptance rates and the chain getting ‘stuck’ at the same state for a large number of steps. Conversely if the variance is extremely small, the acceptance probability will tend to be large, although the chain will explore the space quite slowly, resulting in a high serial correlation and a long time to converge to its stationary distribution. Therefore some tuning is typically required using, say, pilot simulations. A practical advice is to tune the proposal variance so that the acceptance rate is in the range 20% – 40% (Roberts and Rosenthal, 2001). See also Roberts et al. (2004b) and Chapter 7 in Robert and Casella (2004). Further implementation issues and advice can be found for example in Tierney (1994), Roberts and Sahu (1997) and Gilks et al. (1995).

Hastings-within-Gibbs

The last MCMC algorithm we review is actually a hybrid MCMC algorithm resulting from the combination of the Gibbs sampler and the MH algorithm. The algorithm was

proposed by Müller (1991, 1993), and is motivated by situations where some of the full conditionals in the Gibbs algorithm on page 50 are not easy to sample from. In such cases, one can replace the difficult Gibbs steps with MH steps. The algorithm has a similar structure to the Gibbs sampler and is naturally called the Hastings-within-Gibbs algorithm/sampler. It can be shown that this hybrid algorithm constructs a Markov chain with π as its invariant distribution.

From a practical perspective, it is worth mentioning that any of the MH steps in the Hasting-within-Gibbs sampler can be run either one or several times (Robert and Casella, 2004), depending on the desired level of approximation of the components involved. However, one must keep in mind that running multiple MH steps may not necessarily lead to a better approximation of the target distribution π , and that the decision to do this depends on the particular properties of the problem being analysed. In the simulation results of Chapter 5 we shall see the advantage of using multiple MH updates in a data augmentation algorithm (see Section 5.7). For more comments on this, we kindly refer the interested reader to Chapter 10 in Robert and Casella (2004) and references therein.

Chapter 3

Risk-sensitive optimal switching

In this chapter, we develop a risk-sensitive approach to the class of optimal switching problems introduced in Section 2.1. A novel feature of our approach is that we consider a multi-dimensional controlled diffusion process in combination with a risk-sensitive criterion of exponential type. Our objective is to characterise risk-sensitive optimal switching strategies, describing when and to which new regime it is optimal to switch under a risk management perspective.

To do this, we adopt a probabilistic approach based on the concept of the Snell envelope of stochastic processes. This concept, which is reviewed in Theorem 3.6 below, is a powerful probabilistic tool that provides a supermartingale characterisation to solutions of optimal stopping problems. In this class of control problems, an optimal stopping time (if one exists) occurs at the first time when the underlying process (e.g. a revenue process) and its Snell envelope are equal almost surely.

We recast the risk-sensitive optimal switching problem as an iterative optimal stopping problem, and by a repeated application of the Snell envelope theorem, we shall be able to characterise recursively optimal solutions to risk-sensitive switching problems. As byproduct, we obtain a multiplicative version of the dynamic programming principle, upon which a numerical solution to risk-sensitive optimal switching problems is described in the next chapter.

3.1 Preliminaries

In the following we present some essential definitions and results which are needed in this chapter. Throughout we fix a finite time horizon $T > 0$ and consider a standard d -dimensional Brownian motion $W = (W_t)_{t \in [0, T]}$ defined on a given probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We denote by $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ the natural filtration generated by W and completed with the \mathbb{P} -null sets of \mathcal{F} . Thus \mathbb{F} satisfies the usual conditions of right-continuity and augmentation by null sets. We use the shorthand notation *rcll* to say that the sample paths of a given process are right continuous, with left limits. Also, an event happening almost surely with respect to \mathbb{P} will be indicated by writing \mathbb{P} -a.s. or simply a.s, if no confusion is possible.

Definition 3.1. (Stopping times). A random variable $\tau : \Omega \rightarrow [0, \infty)$ is called a *stopping time* with respect to the filtration \mathbb{F} if $\{\tau \leq t\} \in \mathcal{F}_t$ for each $t \geq 0$.

Let us denote by S the class of stopping times in $[0, T]$ with respect to the filtration \mathbb{F} , and for any $\theta \in S$, let us set $S_\theta = \{\tau \in S : \tau \geq \theta, \mathbb{P}\text{-a.s.}\}$, i.e. the set of \mathbb{F} -stopping times τ such that $\theta \leq \tau \leq T$ a.s. Also, let us denote by \mathcal{I}^2 the set of real-valued, \mathbb{F} -adapted, continuous processes G such that

$$\mathbb{E} \left[\sup_{s \in [t, T]} |G_s|^2 \right] < \infty.$$

Definition 3.2. (Class [D]). A real-valued \mathbb{F} -adapted process $G = (G_s)_{s \in [t, T]}$ is said to be of class [D] if the family of random variables $\{G_\tau : \tau \in S\}$ is uniformly integrable (UI).

We recall that a family of random variables that is bounded in \mathcal{L}^p , for any $p > 1$, is necessarily UI. We shall use this fact below when proving that a given stochastic process is of class [D]. Since random variables are defined only up to equivalence classes with \mathbb{P} -a.s equality, it does not make sense in general to consider the supremum of a family of random variables for each given $\omega \in \Omega$. The following lemma introduces the concept of essential supremum of a family of random variables which is the appropriate notion of supremum in this context.

Lemma 3.3. (Essential supremum, Lemma 1.3, Peskir and Shiryaev (2006)). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let $\{Z_\alpha : \alpha \in I\}$ be an arbitrary nonempty family

of random variables. Then there exists a countable subset \mathcal{C} of I such that the random variable $Z^* : \Omega \rightarrow \bar{\mathbb{R}}$ given by $Z^* = \sup_{\alpha \in \mathcal{C}} Z_\alpha$ is such that

1. For each $\alpha \in I$, $Z_\alpha \leq Z^*$ a.s.
2. If \bar{Y} is a random variable such that $Z_\alpha \leq \bar{Y}$ a.s. for all $\alpha \in I$, then $Z^* \leq \bar{Y}$ a.s.

Note that properties 1 and 2 in Lemma 3.3 imply that Z^* is unique up to a \mathbb{P} -null set. The random variable Z^* is called the *essential supremum* of $\{Z_\alpha : \alpha \in I\}$ and it is denoted here by $Z^* = \text{ess sup}_{\alpha \in I} Z_\alpha$.

3.2 Problem formulation

Next we formulate the finite horizon risk-sensitive optimal switching problem addressed in the thesis.

Operational regimes and switching controls

Let $\mathcal{R} := \{1, 2, \dots, m\}$ be a finite control space representing the set of all possible operational regimes that a controller can assign to a system. Roughly, a switching control is a sequence of increasing times and corresponding regimes chosen by the controller, which characterise the time periods over which the system is assigned a certain operating regime belonging to \mathcal{R} . More precisely, for a fixed $t \in [0, T]$ we define a *switching control* in $[t, T]$ as a double sequence

$$((\xi_0, \tau_0 = t), (\xi_1, \tau_1), (\xi_2, \tau_2), \dots),$$

where $(\tau_n)_{n \geq 0}$ is an increasing sequence of stopping times in S_t and for each $n \geq 0$, ξ_n belongs to \mathcal{R} and is \mathcal{F}_{τ_n} -measurable. The regime ξ_n indicates the actual regime that is applied to the system from time τ_n to τ_{n+1} . A switching control will be said to be admissible if the two following conditions hold,

- $\lim_{k \rightarrow \infty} \tau_k = T$, \mathbb{P} -a.s., the switching times are required to cover the time period $[t, T]$.
- $\mathbb{P}(\{\omega \in \Omega : \tau_n(\omega) < T, \forall n \geq 1\}) = 0$, the switching control is finite, i.e. the controller is allowed to make only a finite number of switches over $[t, T]$.

Let us denote by $\mathcal{U}(t)$ the set of admissible switching controls in $[t, T]$, and by $\mathcal{U}^i(t)$ the switching controls in $\mathcal{U}(t)$ such that the initial regime is $\xi_0 = i$. Given a control $(\xi_n, \tau_n) \in \mathcal{U}^i(t)$, there is an associated process $(u_s)_{s \in [t, T]}$ indicating the operating regime being applied at time $s \in [t, T]$. Following Pham et al. (2009) and Aïd et al. (2014), we define the regime indicator process by

$$u_s := \sum_{n \geq 0} \xi_n \mathbb{I}_{[\tau_n, \tau_{n+1})}(s),$$

where $\mathbb{I}_{[a, b)}(s) = 1$ if $s \in [a, b)$ and 0 otherwise. We notice that (u_s) is a rcll piecewise process adapted to the filtration \mathbb{F} . Additionally, there is a unique correspondence between the set of admissible switching controls $(\xi_n, \tau_n) \in \mathcal{U}^i(t)$ and the regime indicator processes. Hence we shall denote by u both an element of $\mathcal{U}^i(t)$ and its associated regime indicator process.

Controlled diffusion process

We assume that the state dynamics is described by the controlled regime-switching diffusion

$$dX_s = b(X_s, u_s)ds + \sigma(X_s, u_s)dW_s, \quad X_t = x, \quad u_t = i. \quad (3.1)$$

where $u \in \mathcal{U}^i(t)$, $x \in \mathbb{R}^d$ is the initial state, and $b : \mathbb{R}^d \times \mathcal{R} \rightarrow \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \times \mathcal{R} \rightarrow \mathbb{R}^{d' \times d}$ are two measurable functions. We call $b(\cdot, \cdot)$ the drift coefficient and $\sigma(\cdot, \cdot)$ the diffusion coefficient (or matrix) of the SDE (3.1). For each $j \in \mathcal{R}$, we assume that $b(\cdot, j)$ and $\sigma(\cdot, j)$ satisfy Lipschitz and linear growth conditions. More specifically,

Assumption 3.4. *For every $x, y \in \mathbb{R}^d$, $j \in \mathcal{R}$ and for some constant $C > 0$, the following conditions hold*

$$|b(x, j) - b(y, j)| + |\sigma(x, j) - \sigma(y, j)| \leq C|x - y|, \quad (3.2)$$

$$|b(x, j)| + |\sigma(x, j)| \leq C(1 + |x|). \quad (3.3)$$

Additionally we assume that the diffusion matrix is positive-definite, i.e. for every non-zero $x \in \mathbb{R}^d$, $j \in \mathcal{R}$, $x\sigma(x, j)^T\sigma(x, j)x^T > 0$.

Assumption 3.4 ensures that the SDE (3.1) with random coefficients has a unique

non-degenerate strong solution with initial data (t, x, i) . We denote this solution by $\{X_s^{t,x,u}, t \leq s \leq T\}$. For a review on the solution to SDEs with random coefficients see for example Appendix D in Fleming and Soner (2006) or Krylov (2008). Without loss of generality we will assume throughout this chapter that the system starts its operation at a fixed time t and a fixed initial state x .

Payoff functions and switching costs

For each operating regime $j \in \mathcal{R}$, we assume that there is an associated payoff rate function $\psi_j : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$, which represents the instantaneous reward accrued by the system controller in regime j . We make the following assumption.

Assumption 3.5. *For each $j \in \mathcal{R}$, $\psi_j : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is \mathbb{F} -measurable, jointly continuous and uniformly bounded by a constant $M > 0$, i.e.,*

$$|\psi_j(s, y)| \leq M, \quad \forall (s, y) \in [0, T] \times \mathbb{R}^d, \forall j.$$

Further, we suppose that a change of regimes, say, from i to j , $i \neq j$, implies an instantaneous positive cost denoted by $K_{i,j}(\cdot) := K(i, j, \cdot)$, and which may be dependent on the current state of the system. We assume that K is given by a measurable Lipschitz function with respect to $y \in \mathbb{R}^d$,

$$K : \mathcal{R} \times \mathcal{R} \times \mathbb{R}^d \rightarrow \mathbb{R}.$$

Further, we assume that for every $x \in \mathbb{R}^d$, $K_{i,i}(x) = 0$, and that for some $\epsilon > 0$, $K_{i,j}(x) > \epsilon > 0$ for every $j \neq i$. Additionally, for every $x \in \mathbb{R}^d$ we assume the following triangular condition holds:

$$K_{i,k}(x) < K_{i,j}(x) + K_{j,k}(x), j \neq i, k.$$

The above conditions on the switching costs are standard (see e.g. Carmona and Ludkovski (2008), Gassiat et al. (2012), Aïd et al. (2014)), and admit intuitive economic interpretations. On one hand, there is a positive cost for switching regimes, but no cost is incurred for not switching. On the other hand, a direct switching from one

regime to any other is less costly than switching through any intermediate regimes. Hereafter we remove the dependence of K_{ij} on the process $X^{t,x,u}$.

The risk-sensitive optimal switching problem

For each outcome $\omega \in \Omega$, the total revenue over $[t, T]$ associated with the switching control $u \in \mathcal{U}^i(t)$ is given by

$$P(t, x, i; u)(\omega) := \int_t^T \psi_{u_s}(s, X_s^{t,x,u}(\omega)) ds - \sum_{t \leq \tau_k < T} K_{u_{\tau_{k-1}(\omega)}, u_{\tau_k}(\omega)}. \quad (3.4)$$

The first expression on the right-hand side of (3.4) corresponds to the cumulative revenue associated with u , whilst the second expression amounts to the total switching costs generated with such control. Our aim is to maximise a function of the total revenue function as follows.

Let $\rho < 0$ be a parameter representing risk sensitivity. For fixed $(t, x, i) \in [0, T] \times \mathbb{R}^d \times \mathcal{R}$, we define the performance index of the risk-sensitive optimal switching problem by

$$J(t, x, i; u) = \mathbb{E} \left[-e^{\rho P(t,x,i;u)} \mid \mathcal{F}_t \right].$$

The risk-sensitive optimal switching problem consists of finding a switching control $u^* \in \mathcal{U}^i(t)$ such that

$$J(t, x, i; u^*) = \sup_{u \in \mathcal{U}^i(t)} J(t, x, i; u). \quad (3.5)$$

Accordingly, we define the value function of the risk-sensitive optimal switching problem by

$$\begin{aligned} V(t, x, i) &:= J(t, x, i; u^*) \\ &= \sup_{u \in \mathcal{U}^i(t)} \mathbb{E} \left[-\exp \left(\rho \left(\int_t^T \psi_{u_s}(s, X_s^{t,x,u}) ds - \sum_{t \leq \tau_k < T} K_{u_{\tau_{k-1}}, u_{\tau_k}} \right) \right) \mid \mathcal{F}_t \right]. \end{aligned} \quad (3.6)$$

For short, we call $V(t, x, i)$ the risk-sensitive (RS) value function. Notice that the analogous of the risk-sensitive criterion in equation (2.2) reads $-\rho^{-1} \log [-J(t, x, i; u)]$. Nonetheless, since the logarithmic function is injective it does not play a role in our optimisation problem, and so we focus on the optimisation problem formulated above. When necessary, e.g. in numerical applications, we shall re-scale $V(t, x, i)$ so that it is

on the original scale of $P(t, x, i; u)$.

3.3 Iterative optimal stopping

Optimal switching can be understood as a generalisation of optimal stopping problems (Peskir and Shiryaev, 2006). Indeed, in the latter the aim is to find a single stopping time, where it is optimal to take a pre-specified action. A typical example of optimal stopping time problems is the pricing of American-style financial derivatives such the American put option. Here, one seeks to find a stopping time where exercising the option gives maximum return to the buyer. On the other hand, optimal switching requires finding a sequence of stopping times and actions taken from a finite set, which gives the maximum expected revenue over the planning period. In a few words, there are more ‘options’ and decision times in optimal switching than in optimal stopping.

Nonetheless, there exists a connection between these two problems. The process of finding a sequence of optimal stopping times can be divided in a sequence of sub-problems as follows: given an initial time t and initial operating regime i , we find both the first time where it is optimal to switch after t and the corresponding optimal regime. That is, we find τ_1 and ξ_1 in $u = ((i, t), (\xi_1, \tau_1), \dots, ((\xi_n, \tau_n), \dots))$. Note that this is a standard optimal stopping time problem where the options are in the set of regimes \mathcal{R} . Now, setting this optimal switching time and regime as the new initial conditions, we proceed to find the first time after τ_1 where switching is optimal, alongside the optimal regime. This will give the values of τ_2 and ξ_2 . We can continue applying the same reasoning until the end of the planning period T .

This idea has been recently exploited in Carmona and Ludkovski (2008) and Djehiche et al. (2009) in the context of optimal switching, when the system dynamics is not affected by the control. An application to stochastic impulse control problems is considered in Djehiche et al. (2010). There, the authors tackle their impulse control problem using the theory of reflected backward SDEs.

In this thesis, we use this idea to re-formulate the risk-sensitive optimal switching problem as an iterative optimal stopping problem. Thus we extend the results of Carmona and Ludkovski (2008) and Djehiche et al. (2009) in two ways. On the one hand, we consider a situation where the system dynamics is affected by the switching control.

On the other hand, we include a risk-sensitive approach to the optimisation problem, considering a non-linear optimisation objective of exponential type. As we illustrate extensively in Chapter 6, this approach allows us to include the risk-preferences of system controllers in order to avoid situations that represent a potential high risk. Our method of proof follows a similar methodology to that presented in Djehiche et al. (2009) and Djehiche et al. (2010), together with some ideas that are specific to our context. In the following we re-cast the risk-sensitive optimal switching as an iterative optimal stopping problem.

Let $u \in \mathcal{U}(t)$ be an admissible control. For any $\theta \in S_t$ and ξ an \mathcal{F}_θ -measurable regime in \mathcal{R} , we define the switching control $u \oplus (\xi, \theta)$ by

$$(u \oplus (\xi, \theta))_s := \begin{cases} u_s, & \text{on } \{s < \theta\} \\ \xi, & \text{on } \{s \geq \theta\} \end{cases}.$$

Intuitively, the switching control $u \oplus (\xi, \theta)$ is a constant switching control after time θ , from where it takes the value ξ . Note that $u \oplus (\xi, \theta)$ is uniquely defined up to a \mathbb{P} -null set by u, ξ and θ . Furthermore, if u is admissible then $u \oplus (\xi, \theta)$ is also admissible. The operator \oplus is left associative, but for simplicity we write

$$(u \oplus (j_1, s_1)) \oplus (j_2, s_2) = u \oplus (j_1, s_1) \oplus (j_2, s_2),$$

for every $s_1 \leq s_2 \leq T$ and $j_1, j_2 \in \mathcal{R}$.

Let $u \in \mathcal{U}(t)$ be an admissible control, θ be a stopping time in S_t and ξ an \mathcal{F}_θ -measurable regime in \mathcal{R} . For $s \in [t, T]$ and each $k \geq 0$, we introduce the processes $(V_s^k(t, x, u, \xi, \theta))_s$ defined recursively by

$$V_s^0(t, x, u, \xi, \theta) := \mathbb{E} \left[-\exp \left(\rho \int_s^T \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \middle| \mathcal{F}_s \right], \quad (3.7)$$

and for $k \geq 1$,

$$\begin{aligned} V_s^k(t, x, u, \xi, \theta) \\ := \operatorname{ess\,sup}_{\tau \in S_s} \mathbb{E} \left[\exp \left(\rho \int_s^\tau \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M} V_\tau^k(t, x, u \oplus (\xi, \theta), \xi, \theta) \middle| \mathcal{F}_s \right], \end{aligned} \quad (3.8)$$

where

$$\mathcal{M} V_s^k(t, x, u, \xi, \theta) := \max_{j \neq \xi} (V_s^{k-1}(t, x, u, j, \theta) e^{-\rho K_{\xi,j} \mathbb{I}_{[s < T]}}). \quad (3.9)$$

We provide some intuition on these auxiliary processes. On the right-hand of (3.7), notice that the integrand takes the value 0 on $\{s' < \theta\}$, and that after time θ the switching control $u \oplus (\xi, \theta)$ is constant by definition. Thus, informally, the expression $V_s^0(t, x, u, \xi, \theta)$ represents the expected total revenue over $[\theta, T]$ in a risk-sensitive sense (i.e. with the exponential criterion) when the control is constant and equal to ξ over this period. Clearly, when $k = 0$ no switches are allowed on $[\theta, T]$.

The operator \mathcal{M} in (3.9) indicates the best value (in a risk-sensitive sense) that can be obtained at time s by switching immediately from ξ to j , provided k switches remain and the switching cost $K_{\xi,j}$ is paid. Hence $V_s^k(t, x, u, \xi, \theta)$ corresponds to the maximum expected total revenue over $[\theta, T]$ (again, with the exponential criterion) that can be obtained when k switches are allowed.

3.4 Main properties of Snell envelopes

In the proofs constructed in this chapter, we demonstrate and exploit a number of properties that the processes $(V_s^k(t, x, u, \xi, \theta))_s$ enjoy. Our chief tool to do this is the concept of Snell envelope of processes, whose main properties are summarised in the following theorem. For further details about this concept, we refer the interested reader to Djehiche et al. (2009), Appendix D in Karatzas and Shreve (1998a), the Appendix in Cvitanic and Karatzas (1996) and Peskir and Shiryaev (2006).

Theorem 3.6. *(The Snell envelope, Proposition 2, Djehiche et al. (2009)). Let $G = (G_s)_{s \in [t, T]}$ be an \mathbb{F} -adapted rcll (right-continuous with left limits) real-valued process of*

class $[D]$. Then there exists an \mathbb{F} -adapted rcll real-valued process of class $[D]$, $Z := (Z_s)_{s \in [t, T]}$, such that Z is the smallest supermartingale which dominates G , i.e. if $Z' := (Z'_s)_{s \in [t, T]}$ is another rcll supermartingale of class $[D]$ satisfying $G_s \leq Z'_s$, then $G_s \leq Z_s \leq Z'_s$ for all $s \in [t, T]$.

The process Z is called the Snell envelope of the process G and has the following properties:

1. For any $\theta \in S_t$ we have

$$Z_\theta = \operatorname{ess\,sup}_{\tau \in S_\theta} \mathbb{E}[G_\tau | \mathcal{F}_\theta].$$

2. The Doob-Meyer decomposition of Z implies the existence of a martingale $(M_s)_{s \in [t, T]}$ and two non-decreasing processes $(C_s)_{s \in [t, T]}$ and $(D_s)_{s \in [t, T]}$ such that $C_0 = D_0 = 0$, C is continuous, D is purely discontinuous predictable and for any $s \in [t, T]$,

$$Z_s = M_s - C_s - D_s,$$

and

$$\{\Delta_s D = D_s - D_{s-} > 0\} \subset \{\Delta_s G < 0\} \cap \{Z_{t-} = G_{s-}\}.$$

3. If G has only positive jumps, then Z is a continuous process. Moreover, for any $\theta \in S_t$ the stopping time defined by

$$\tau_\theta^* = \inf \{s \geq \theta : Z_s = G_s\} \wedge T,$$

is optimal after θ , that is,

$$Z_\theta = \mathbb{E}[Z_{\tau_\theta^*} | \mathcal{F}_\theta] = \mathbb{E}[G_{\tau_\theta^*} | \mathcal{F}_\theta] = \operatorname{ess\,sup}_{\tau \in S_\theta} \mathbb{E}[G_\tau | \mathcal{F}_\theta].$$

4. Let us denote Z^G the Snell envelope of G . If $(G^n)_{n \geq 0}$ and G are rcll processes of class $[D]$ such that $(G^n)_{n \geq 0}$ converges increasingly and pointwisely to G , then $(Z^{G^n})_{n \geq 0}$ converges increasingly and pointwisely to Z^G . Finally, $G \in \mathcal{I}^2$ implies that $Z^G \in \mathcal{I}^2$.

The Snell envelope theorem is a powerful probabilistic tool that will help us to characterise solutions to the risk-sensitive optimal switching problem. The key insight

of our proofs is the construction of suitable processes (G_s) , whose Snell envelopes can be obtained directly from the processes $(V_s^k(t, x, u, \xi, \theta))_s$. Thanks to property 3 in Theorem 3.6, we can use these processes, for example, to characterise solutions to the iterative optimal stopping problems in (3.8). Additionally, with property 4 we can prove limit properties as k increases. We proceed now to the main results of this chapter.

3.5 Properties of the auxiliary processes

In Proposition 3.7 we explore the properties of the processes $V_s^k(t, x, u, \xi, \theta)$. The key results of this proposition are, on the one hand, that these processes converge to an appropriate limit as the number of available switches k increases, and on the other hand, that the limit process satisfies equation (3.10). This equation is a multiplicative version of the dynamic programming principle (see e.g. Bertsekas (1976) and Fleming and Soner (2006)) for optimal switching problems. In Proposition 3.9 we show that the risk-sensitive value function $V(t, x, i)$ also satisfies (3.10).

Proposition 3.7. *Let θ be a stopping time in S_t , $\xi \in \mathcal{R}$ be \mathcal{F}_θ -measurable and $u \in \mathcal{U}(t)$. Then*

1. *For each $k \geq 0$, the process $V^k(t, x, u, \xi, \theta)$ belongs to \mathcal{I}^2 and verifies*

$$V_T^k(t, x, u, \xi, \theta) = -1.$$

2. *The sequence of processes $(V^k(t, x, u, \xi, \theta))_{k \geq 0}$ is increasing and converges pointwisely a.s. on $[t, T]$ to a rcll process $\bar{V}(t, x, u, \xi, \theta)$.*

3. *The limit process $\bar{V}(t, x, u, \xi, \theta)$ satisfies*

- (a) *For any $s \in [t, T]$,*

$$\begin{aligned} & \bar{V}_s(t, x, u, \xi, \theta) \\ &= \operatorname{esssup}_{\tau \in S_s} \mathbb{E} \left[\exp \left(\rho \int_s^\tau \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M} \bar{V}_\tau(t, x, u \oplus (\xi, \theta), \xi, \theta) \middle| \mathcal{F}_s \right]. \end{aligned} \tag{3.10}$$

(b) For every $\theta' \in S_\theta$ and $s \geq \theta'$,

$$\bar{V}_s(t, x, u, \xi, \theta) = \bar{V}_s(t, x, u, \xi, \theta') \text{ a.s.}$$

(c) $\mathbb{E} [\sup_{s \in [t, T]} |\bar{V}_s(t, x, u, \xi, \theta)|^2] < \infty$.

Proof. **1.** Let $\theta \in S_t$, $\xi \in \mathcal{R}$ be an \mathcal{F}_θ -measurable regime and $u \in \mathcal{U}(t)$. We show by induction that for each $k \geq 0$, $V^k(t, x, u, \xi, \theta)$ is in \mathcal{I}^2 and verifies $V_T^k(t, x, u, \xi, \theta) = -1$. Let us start with the case $k = 0$. From the definition of $V_s^0(t, x, u, \xi, \theta)$ in equation (3.7), it readily follows that $V_T^0(t, x, u, \xi, \theta) = -1$. Also, we have that

$$\begin{aligned} \mathbb{E} \left[\sup_{s \in [t, T]} |V_s^0(t, x, u, \xi, \theta)|^2 \right] &= \mathbb{E} \left[\sup_s \left| \mathbb{E} \left[-\exp \left(\rho \int_s^T \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \middle| \mathcal{F}_s \right] \right|^2 \right] \\ &\leq \mathbb{E} \left[\sup_s \mathbb{E} \left[\exp \left(2|\rho| \int_s^T |\psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)})| \mathbb{I}_{[s' \geq \theta]} ds' \right) \middle| \mathcal{F}_s \right] \right] \\ &\leq \mathbb{E} \left[\sup_s \mathbb{E} \left[\exp \left(2|\rho| \int_t^T |\psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)})| ds' \right) \middle| \mathcal{F}_s \right] \right] \\ &\leq \exp(2|\rho|M(T-t)), \end{aligned}$$

since ψ_j is uniformly bounded by assumption 3.5. Now, we notice that

$$\begin{aligned} &V_s^0(t, x, u, \xi, \theta) \\ &= \mathbb{E} \left[-\exp \left(\rho \int_s^T \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \middle| \mathcal{F}_s \right] \\ &= \mathbb{E} \left[-\exp \left(\rho \int_t^T \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \middle| \mathcal{F}_s \right] \exp \left(-\rho \int_t^s \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right), \end{aligned}$$

that is, $V_s^0(t, x, u, \xi, \theta)$ can be written as a square-integrable martingale with respect to (w.r.t.) a Brownian filtration times a continuous process. Since every martingale w.r.t to \mathbb{F} has a rcll modification, we always assume that this modification is taken. Then we can apply the martingale representation theorem (see e.g. Theorem 4.15, Chapter 3 in Karatzas and Shreve (1998b)) to conclude that $V^0(t, x, u, \xi, \theta)$ is continuous. Therefore $V^0(t, x, u, \xi, \theta) \in \mathcal{I}^2$ and $V_T^0(t, x, u, \xi, \theta) = -1$, for every $\theta \in S_t$, \mathcal{F}_θ -measurable regime $\xi \in \mathcal{R}$ and $u \in \mathcal{U}(t)$.

Assume now that these two properties are satisfied for some $k \geq 1$. Further, let us assume that for this k

$$\mathbb{E} \left[\sup_{s \in [t, T]} |V_s^k(t, x, u, \xi, \theta)|^2 \right] \leq \exp(2(k+1)|\rho|M(T-t)),$$

which holds true for $k = 0$. Then, by equation (3.8), we have that $V_T^{k+1}(t, x, u, \xi, \theta) = -1$ since

$$\mathcal{M}V_T^{k+1}(t, x, u, \xi, \theta) = \max_{j \neq \xi} (V_T^k(t, x, u, j, \theta) e^{-\rho K_{\xi, j} \mathbb{I}_{[T < T]}}) = -1.$$

Let us define the process G by

$$(G_s)_{s \in [t, T]} := \left(\exp \left(\rho \int_t^s \psi_{\xi}(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M}V_s^{k+1}(t, x, u \oplus (\xi, \theta), \xi, \theta) \right)_{s \in [t, T]}. \quad (3.11)$$

We claim that G is of class [D]. Indeed, for any $\tau \in S_t$ we have that,

$$\begin{aligned} & \mathbb{E}[|G_{\tau}|^2] \\ &= \mathbb{E} \left[\left| \exp \left(\rho \int_t^{\tau} \psi_{\xi}(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M}V_{\tau}^{k+1}(t, x, u \oplus (\xi, \theta), \xi, \theta) \right|^2 \right] \\ &\leq \exp(2|\rho|M(T-t)) \mathbb{E} [| \mathcal{M}V_{\tau}^{k+1}(t, x, u \oplus (\xi, \theta), \xi, \theta) |^2] \\ &\leq \exp(2|\rho|M(T-t)) \mathbb{E} \left[\left| \max_{j \neq \xi} (V_{\tau}^k(t, x, u \oplus (\xi, \theta), j, \theta) e^{-\rho K_{\xi, j} \mathbb{I}_{[\tau < T]}}) \right|^2 \right] \\ &\leq \exp(2|\rho|M(T-t)) \mathbb{E} \left[\left| \max_{j \neq \xi} \sup_{s \in [t, T]} (V_s^k(t, x, u \oplus (\xi, \theta), j, \theta)) \right|^2 \right] \\ &\leq \exp(2|\rho|M(T-t)) \mathbb{E} \left[\left| \max_{j \in \mathcal{R}} \sup_{s \in [t, T]} (V_s^k(t, x, u \oplus (j, \theta), j, \theta)) \right|^2 \right] \\ &\leq \exp(2|\rho|M(T-t)) \exp(2(k+1)|\rho|M(T-t)) \\ &\leq \exp(2(k+2)|\rho|M(T-t)), \end{aligned}$$

where we have used that \mathcal{R} is a finite set and the induction hypothesis. Hence, the family $\{G_\tau : \tau \in S_t\}$ is bounded in \mathcal{L}^2 and therefore it is UI. Consequently, G is of class [D] as we wanted. Note that a similar reasoning as before shows that

$$\mathbb{E} \left[\sup_{s \in [t, T]} |V_s^{k+1}(t, x, u, \xi, \theta)|^2 \right] \leq \exp(2(k+2)|\rho|M(T-t)).$$

Multiplying both sides of (3.8) by $\exp\left(\rho \int_t^s \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds'\right)$, we can conclude that

$$\left(V_s^{k+1}(t, x, u, \xi, \theta) \exp\left(\rho \int_t^s \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds'\right) \right)_{s \in [t, T]},$$

is the Snell envelope of the process $(G_s)_{s \in [t, T]}$, thanks to property 1 in Theorem 3.6.

By the induction hypothesis we know that $V^k(t, x, u, \xi, \theta)$ is continuous, but

$V_s^k(t, x, u, j, \theta) e^{-\rho K_{\xi, j}}|_{s=T} < -1$ since $K_{ij} > 0$ for all $i \neq j$ and $\rho < 0$. Hence

$$\mathcal{M}V_s^{k+1}(t, x, (u \oplus (\xi, \theta)) \oplus (j, s), \xi, \theta) = \max_{j \neq \xi} \left(V_s^k(t, x, u \oplus (\xi, \theta) \oplus (j, s), j, \theta) e^{-\rho K_{\xi, j} \mathbb{I}_{[s < T]}} \right),$$

is continuous on $[0, T)$ and has a positive jump at $s = T$. Thus, by property 3 in Theorem 3.6, $V^{k+1}(t, x, u, \xi, \theta)$ is continuous and hence it belongs to \mathcal{I}^2 . Proceeding by induction, it holds true that for each $k \geq 0$, $V^k(t, x, u, \xi, \theta)$ belongs to \mathcal{I}^2 and verifies $V_T^k(t, x, u, \xi, \theta) = -1$.

2. Next we show that $(V^k(t, x, u, \xi, \theta))_{k \geq 0}$ is an increasing sequence and that $V_s^k(t, x, u, \xi, \theta) \leq -e^{\rho M(T-s)}$ for all $s \in [t, T]$, $k \geq 0$. By definitions (3.7) and (3.8), we have that

$$\begin{aligned} & V_s^1(t, x, u, \xi, \theta) \\ &= \operatorname{ess\,sup}_{\tau \in S_s} \mathbb{E} \left[\exp\left(\rho \int_s^\tau \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds'\right) \mathcal{M}V_\tau^1(t, x, u \oplus (\xi, \theta), \xi, \theta) \middle| \mathcal{F}_s \right] \\ &\geq \mathbb{E} \left[\exp\left(\rho \int_s^T \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds'\right) \mathcal{M}V_T^1(t, x, u \oplus (\xi, \theta), \xi, \theta) \middle| \mathcal{F}_s \right] \\ &= \mathbb{E} \left[-\exp\left(\rho \int_s^T \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds'\right) \middle| \mathcal{F}_s \right] \end{aligned}$$

$$= V_s^0(t, x, u, \xi, \theta).$$

Obviously, $V_s^0(t, x, u, \xi, \theta) \leq -e^{\rho M(T-s)}$ by assumption 3.5, and hence,

$$\begin{aligned} \mathcal{M}V_\tau^1(t, x, u \oplus (\xi, \theta), \xi, \theta) &= \max_{j \neq \xi} (V_\tau^0(t, x, u \oplus (\xi, \theta), j, \theta) e^{-\rho K_{\xi, j} \mathbb{I}_{[\tau < T]}}) \\ &\leq \max_{j \neq \xi} (-e^{\rho M(T-\tau)} e^{-\rho K_{\xi, j} \mathbb{I}_{[\tau < T]}}) \\ &\leq -e^{\rho M(T-\tau)}, \end{aligned}$$

since $K_{i, j} > 0$, for all $i \neq j$. Therefore,

$$V_s^1(t, x, u, \xi, \theta) \leq -e^{\rho M(\tau-s)} e^{\rho M(T-\tau)} \leq -e^{\rho M(T-s)}.$$

Let us assume now that for some $k \geq 1$, $V_s^k(t, x, u, \xi, \theta) \leq V_s^{k+1}(t, x, u, \xi, \theta) \leq -e^{\rho M(T-s)}$. By definition this implies that \mathbb{P} -a.s.

$$\mathcal{M}V_s^k(t, x, u \oplus (\xi, \theta), \xi, \theta) \leq \mathcal{M}V_s^{k+1}(t, x, u \oplus (\xi, \theta), \xi, \theta),$$

and hence, using (3.8), we conclude that $V_s^{k+1}(t, x, u, \xi, \theta) \leq V_s^{k+2}(t, x, u, \xi, \theta)$ for every $t \leq s \leq T$. Finally, since $V_s^{k+1}(t, x, u, \xi, \theta) \leq -e^{\rho M(T-s)}$ we must have that $V_s^{k+2}(t, x, u, \xi, \theta) \leq -e^{\rho M(T-s)}$. Therefore for all $k \geq 0$, $s \in [t, T]$,

$$V_s^k(t, x, u, \xi, \theta) \leq V_s^{k+1}(t, x, u, \xi, \theta) \leq -e^{\rho M(T-s)}, \quad (3.12)$$

from where we deduce that $\lim_{k \rightarrow \infty} V_s^k(t, x, u, \xi, \theta)$ exists. Accordingly, for a fixed $(t, x) \in [0, T]$, and for every $u \in \mathcal{U}(t)$, $\theta \in S_t$, $\xi \in \mathcal{R}$, we define

$$\bar{V}_s(t, x, u, \xi, \theta) := \lim_{k \rightarrow \infty} V_s^k(t, x, u, \xi, \theta), \quad (3.13)$$

for every $s \in [t, T]$. We notice that $\bar{V}_s(t, x, u, \xi, \theta) \leq -e^{\rho M(T-s)}$ by definition.

We show next that $\bar{V}_s(t, x, u, \xi, \theta)$ is a rcll process. Recall that for $k \geq 1$,

$$\left(V_s^{k+1}(t, x, u, \xi, \theta) \exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) ds' \right) \right)_{s \in [t, T]},$$

is a continuous supermartingale, being the Snell envelope of the process in (3.11).

Then

$$\left(\bar{V}_s(t, x, u, \xi, \theta) \exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) ds' \right) \right)_{s \in [t, T]},$$

is indistinguishable from a rcll process, thanks to Theorem 18, p. 79 in Dellacherie and Meyer (1982).

3.(a) We can now prove that $\bar{V}_s(t, x, u, \xi, \theta)$ satisfies the multiplicative dynamic programming equation in (3.10). Indeed, using the results in 2., we observe that the process

$$\left(\exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M}V_s^k(t, x, u \oplus (\xi, \theta), \xi, \theta) \right)_{s \in [t, T]},$$

converges increasingly and pointwisely to the process

$$\left(\exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M}\bar{V}_s(t, x, u \oplus (\xi, \theta), \xi, \theta) \right)_{s \in [t, T]}.$$

Hence, thanks to property 4 in Theorem 3.6, we find that

$$\left(\bar{V}_s(t, x, u, \xi, \theta) \exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) ds' \right) \right)_{s \in [t, T]},$$

is the Snell envelope of

$$\left(\exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M}\bar{V}_s(t, x, u \oplus (\xi, \theta), \xi, \theta) \right)_{s \in [t, T]},$$

from where (3.10) follows.

3.(b) Let $\theta' \in S_\theta$. Then, if $s \geq \theta' \geq \theta$, we obtain that

$$\begin{aligned} V_s^0(t, x, u, \xi, \theta) &= \mathbb{E} \left[- \exp \left(\rho \int_s^T \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \middle| \mathcal{F}_s \right] \\ &= \mathbb{E} \left[- \exp \left(\rho \int_s^T \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta']} ds' \right) \middle| \mathcal{F}_s \right] \\ &= V_s^0(t, x, u, \xi, \theta'). \end{aligned}$$

Let us assume now that for some k , $V_s^k(t, x, u, \xi, \theta) = V_s^k(t, x, u, \xi, \theta')$ a.s., $s \geq \theta'$. This implies that, for any $s \geq \theta'$,

$$\mathcal{M}V_s^{k+1}(t, x, u, \xi, \theta) = \mathcal{M}V_s^{k+1}(t, x, u, \xi, \theta'),$$

and then we must have that $V_s^{k+1}(t, x, u, \xi, \theta) = V_s^{k+1}(t, x, u, \xi, \theta')$, $s \geq \theta'$. Therefore, for every $k \geq 0$, it holds true that

$$V_s^k(t, x, u, \xi, \theta) = V_s^k(t, x, u, \xi, \theta'), s \geq \theta'.$$

Since $\lim_{k \rightarrow \infty} V_s^k(t, x, u, \xi, \theta) = \bar{V}_s(t, x, u, \xi, \theta)$ we find that

$$\bar{V}_s(t, x, u, \xi, \theta) = \bar{V}_s(t, x, u, \xi, \theta'), s \geq \theta',$$

as we wanted.

3.(c) Finally, from the characterisation of

$$\left(V_s^{k+1}(t, x, u, \xi, \theta) \exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) ds' \right) \right)_{s \in [t, T]},$$

as a continuous supermartingale and Doob's maximal inequality, it follows that

$$\begin{aligned}
& \mathbb{E} \left[\sup_{s \in [t, T]} |V_s^k(t, x, u, \xi, \theta)|^2 \right] \\
&= \mathbb{E} \left[\sup_{s \in [t, T]} \left| V_s^k(t, x, u, \xi, \theta) \exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) ds' \right) \right. \right. \\
&\quad \left. \left. \exp \left(-\rho \int_t^s \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) ds' \right) \right|^2 \right] \\
&\leq \mathbb{E} \left[\sup_{s \in [t, T]} \left| V_s^k(t, x, u, \xi, \theta) \exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) ds' \right) \exp(|\rho| M(T-t)) \right|^2 \right] \\
&\leq 4 \exp(2|\rho| M(T-t)) \mathbb{E} \left[\left| V_T^k(t, x, u, \xi, \theta) \exp \left(\rho \int_t^T \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) ds' \right) \right|^2 \right] \\
&\leq 4 \exp(2|\rho| M(T-t)) \mathbb{E} \left[\left| \exp \left(\rho \int_t^T \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) ds' \right) \right|^2 \right] \\
&\leq 4 \exp(2|\rho| M(T-t)) \mathbb{E} [\exp(2|\rho| M(T-t))] \\
&\leq 4 \exp(4|\rho| M(T-t)).
\end{aligned}$$

Hence, thanks to the Lebesgue dominated convergence theorem, we conclude that

$$\mathbb{E} \left[\sup_{s \in [t, T]} |\bar{V}_s(t, x, u, \xi, \theta)|^2 \right] < \infty.$$

□

3.6 Characterisation of optimal solutions

In Proposition 3.8 we justify the construction of the processes $V^k(t, x, u, \xi, \theta)$, $k \geq 0$, given in equations (3.7) and (3.8). In particular, we show that if at time t the state is x and the regime is i , then $V_t^n(t, x, u, i, t)$ is the value function of the risk-sensitive optimal switching problem when only n switches are allowed. Additionally, we will characterise optimal switching controls in such situation, in terms of a sequence of first hitting times of suitably defined processes.

Let us call $\mathcal{U}^{i,n}(t)$ the subset of admissible controls with at most n switches, i.e.

$$\mathcal{U}^{i,n}(t) := \{u \in \mathcal{U}^i(t) : \tau_{n+1} = T \text{ a.s.}\}.$$

Proposition 3.8. *Let u be an admissible control in $\mathcal{U}^i(t)$ and for any $n \geq 0$ let us*

define the sequence of \mathbb{F} -stopping times $(\tau_k^n)_{k \geq 0}$ by $\tau_0^n = t$,

$$\tau_1^n := \inf \{s \geq t : V_s^n(t, x, u, i, t) = \mathcal{M}V_s^n(t, x, u \oplus (i, t), i, t)\} \wedge T,$$

and, for $2 \leq k \leq n$,

$$\begin{aligned} \tau_k^n : &= \inf \{s \geq \tau_{k-1}^n : V_s^{n-k+1}(t, x, u \oplus (\xi_0^n, \tau_0^n) \oplus \dots \oplus (\xi_{k-2}^n, \tau_{k-2}^n), \xi_{k-1}^n, \tau_{k-1}^n) \\ &= \mathcal{M}V_s^{n-k+1}(t, x, u \oplus (\xi_0^n, \tau_0^n) \oplus \dots \oplus (\xi_{k-1}^n, \tau_{k-1}^n), \xi_{k-1}^n, \tau_{k-1}^n)\} \wedge T, \end{aligned}$$

where, $\xi_0^n = i$,

$$\xi_1^n := \arg \max_{j \neq i} \left\{ V_{\tau_1^n}^{n-1}(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i,j} \mathbb{I}_{[\tau_1^n < T]}} \right\},$$

and, for $2 \leq k \leq n$,

$$\xi_k^n := \arg \max_{j \neq \xi_{k-1}^n} \left\{ V_{\tau_k^n}^{n-k-1}(t, x, u \oplus (\xi_0^n, \tau_0^n) \oplus \dots \oplus (\xi_{k-1}^n, \tau_{k-1}^n), j, \tau_{k-1}^n) e^{-\rho K_{\xi_{k-1}^n, j} \mathbb{I}_{[\tau_k^n < T]}} \right\},$$

Finally, let us set for any $k \geq n+1$, $\tau_k^n = T$ and $\xi_k^n = \xi_n^n$. Then

$$u^{*,n} = ((i, t), (\xi_1^n, \tau_1^n), \dots, (\xi_k^n, \tau_k^n), \dots),$$

is optimal when the supremum in (3.5) is taken over $\mathcal{U}^{i,k}$ and

$$V_t^n(t, x, u, i, t) = \sup_{v \in \mathcal{U}^{i,n}(t)} J(t, x, i; v) = J(t, x, i; u^{*,n}).$$

Proof. Let $u \in \mathcal{U}^i(t)$. By equation (3.8) we know that

$$\begin{aligned} V_t^n(t, x, u, i, t) &= \operatorname{ess\,sup}_{\tau \in S_t} \mathbb{E} \left[\exp \left(\rho \int_t^\tau \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) \mathbb{I}_{[s' \geq t]} ds' \right) \mathcal{M}V_\tau^n(t, x, u \oplus (i, t), i, t) \middle| \mathcal{F}_t \right] \\ &= \operatorname{ess\,sup}_{\tau \in S_t} \mathbb{E} \left[\exp \left(\rho \int_t^\tau \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) \mathbb{I}_{[s' \geq t]} ds' \right) \mathcal{M}V_\tau^n(t, x, u \oplus (i, t), i, t) \right], \end{aligned}$$

where the last inequality follows from noting that

$$\exp \left(\rho \int_t^\tau \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) \mathbb{I}_{[s' \geq t]} ds' \right) \mathcal{M}V_\tau^n(t, x, u \oplus (i, t), i, t),$$

is \mathcal{F}_t -measurable. Also, since the expression has only positive jumps, it follows from property 3 in Theorem 3.6 that τ_1^n is optimal. Then

$$\begin{aligned}
V_t^n(t, x, u, i, t) &= \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1^n} \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) ds' \right) \mathcal{M} V_{\tau_1^n}^n(t, x, u \oplus (i, t), i, t) \right] \\
&= \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1^n} \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) ds' \right) \max_{j \neq i} \left(V_{\tau_1^n}^{n-1}(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i, j} \mathbb{I}_{[\tau_1^n < T]}} \right) \right] \\
&= \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1^n} \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) ds' \right) V_{\tau_1^n}^{n-1}(t, x, u \oplus (i, t), \xi_1^n, t) e^{-\rho K_{i, \xi_1^n} \mathbb{I}_{[\tau_1^n < T]}} \right],
\end{aligned}$$

where we have used the definition of ξ_1^n . Now, by proposition 3.7.3(b) we have that

$$V_{\tau_1^n}^{n-1}(t, x, u \oplus (i, t), \xi_1^n, t) = V_{\tau_1^n}^{n-1}(t, x, u \oplus (i, t), \xi_1^n, \tau_1^n),$$

given that $\tau_1^n \geq t$, and hence

$$V_t^n(t, x, u, i, t) = \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1^n} \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) ds' \right) V_{\tau_1^n}^{n-1}(t, x, u \oplus (i, t), \xi_1^n, \tau_1^n) e^{-\rho K_{i, \xi_1^n} \mathbb{I}_{[\tau_1^n < T]}} \right]. \quad (3.14)$$

Let us now work out the expression $V_{\tau_1^n}^{n-1}(t, x, u \oplus (i, t), \xi_1^n, \tau_1^n)$. Following the same steps as before, we have that

$$\begin{aligned}
 & V_{\tau_1^n}^{n-1}(t, x, u \oplus (i, t), \xi_1^n, \tau_1^n) \\
 &= \operatorname{ess\,sup}_{\tau \in S_{\tau_1^n}} \mathbb{E} \left[\exp \left(\rho \int_{\tau_1^n}^{\tau} \psi_{\xi_1^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) \mathbb{I}_{[s' \geq \tau_1^n]} ds' \right) \right. \\
 & \quad \left. \mathcal{M}V_{\tau}^{n-1}(t, x, u \oplus (i, t) \oplus (\xi_1^n, \tau_1^n), \xi_1^n, \tau_1^n) \Big| \mathcal{F}_{\tau_1^n} \right] \\
 &= \mathbb{E} \left[\exp \left(\rho \int_{\tau_1^n}^{\tau_2^n} \psi_{\xi_1^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) ds' \right) \right. \\
 & \quad \left. \mathcal{M}V_{\tau_2^n}^{n-1}(t, x, u \oplus (i, t) \oplus (\xi_1^n, \tau_1^n), \xi_1^n, \tau_1^n) \Big| \mathcal{F}_{\tau_1^n} \right] \\
 &= \mathbb{E} \left[\exp \left(\rho \int_{\tau_1^n}^{\tau_2^n} \psi_{\xi_1^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) ds' \right) \right. \\
 & \quad \left. \max_{j \neq \xi_1^n} \left(V_{\tau_2^n}^{n-1}(t, x, u \oplus (i, t) \oplus (\xi_1^n, \tau_1^n), j, \tau_1^n) e^{-\rho K_{\xi_1^n, j} \mathbb{I}_{[\tau_2^n < T]}} \right) \Big| \mathcal{F}_{\tau_1^n} \right] \\
 &= \mathbb{E} \left[\exp \left(\rho \int_{\tau_1^n}^{\tau_2^n} \psi_{\xi_1^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) ds' \right) \right. \\
 & \quad \left. \left(V_{\tau_2^n}^{n-1}(t, x, u \oplus (i, t) \oplus (\xi_1^n, \tau_1^n), \xi_2^n, \tau_1^n) e^{-\rho K_{\xi_1^n, \xi_2^n} \mathbb{I}_{[\tau_2^n < T]}} \right) \Big| \mathcal{F}_{\tau_1^n} \right] \\
 &= \mathbb{E} \left[\exp \left(\rho \int_{\tau_1^n}^{\tau_2^n} \psi_{\xi_1^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) ds' \right) \right. \\
 & \quad \left. \left(V_{\tau_2^n}^{n-1}(t, x, u \oplus (i, t) \oplus (\xi_1^n, \tau_1^n), \xi_2^n, \tau_2^n) e^{-\rho K_{\xi_1^n, \xi_2^n} \mathbb{I}_{[\tau_2^n < T]}} \right) \Big| \mathcal{F}_{\tau_1^n} \right].
 \end{aligned}$$

Putting this expression inside (3.14), and using that

$$\exp \left(\rho \int_t^{\tau_1^n} \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) ds' \right) e^{-\rho K_{i, \xi_1^n} \mathbb{I}_{[\tau_1^n < T]}},$$

is $\mathcal{F}_{\tau_1^n}$ -measurable, that $[\tau_2^n < T] \subset [\tau_1^n < T]$ and that

$$\begin{aligned}
 & \exp \left(\rho \int_t^{\tau_1^n} \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) ds' \right) \exp \left(\rho \int_{\tau_1^n}^{\tau_2^n} \psi_{\xi_1^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) ds' \right) \\
 &= \exp \left(\rho \left(\int_t^{\tau_1^n} \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) ds' + \int_{\tau_1^n}^{\tau_2^n} \psi_{\xi_1^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) ds' \right) \right) \\
 &= \exp \left(\rho \int_t^{\tau_2^n} \psi_{u_s^*}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) ds' \right),
 \end{aligned}$$

we obtain that

$$\begin{aligned}
 V_t^n(t, x, u, i, t) &= \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_2^n} \psi_{u_s^{*,n}}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n)}) ds' \right) e^{-\rho(K_{i,\xi_1^n} + K_{\xi_1^n, \xi_2^n}) \mathbb{I}_{[\tau_2^n < T]}} \right. \\
 &\quad \left. V_{\tau_2^n}^{n-2}(t, x, u \oplus (i, t) \oplus (\xi_1^n, \tau_1^n), \xi_2^n, \tau_2^n) \right].
 \end{aligned}$$

We repeat this procedure as many times as necessary to get

$$\begin{aligned}
 V_t^n(t, x, u, i, t) &= \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_n^n} \psi_{u_s^{*,n}}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n) \oplus \dots \oplus (\xi_{n-1}^n, \tau_{n-1}^n)}) ds' \right) \right. \\
 &\quad \cdot e^{-\rho(K_{i,\xi_1^n} + K_{\xi_1^n, \xi_2^n} + \dots + K_{\xi_{n-1}^n, \xi_n^n}) \mathbb{I}_{[\tau_n^n < T]}} \\
 &\quad \left. \cdot V_{\tau_n^n}^0(t, x, u \oplus (i, t) \oplus (\xi_1^n, \tau_1^n) \oplus \dots \oplus (\xi_{n-1}^n, \tau_{n-1}^n), \xi_n^n, \tau_n^n) \right].
 \end{aligned}$$

However, we have that

$$\begin{aligned}
 &V_{\tau_n^n}^0(t, x, u \oplus (i, t) \oplus (\xi_1^n, \tau_1^n) \oplus \dots \oplus (\xi_{n-1}^n, \tau_{n-1}^n), \xi_n^n, \tau_n^n) \\
 &= \mathbb{E} \left[- \exp \left(\rho \int_{\tau_n^n}^T \psi_{\xi_n^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n) \oplus \dots \oplus (\xi_n^n, \tau_n^n)}) ds' \right) \middle| \mathcal{F}_{\tau_n^n} \right],
 \end{aligned}$$

and hence,

$$\begin{aligned}
 V_t^n(t, x, u, i, t) &= \mathbb{E} \left[- \exp \left(\rho \int_t^T \psi_{u_s^{*,n}}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^n, \tau_1^n) \oplus \dots \oplus (\xi_n^n, \tau_n^n)}) ds' \right) e^{-\rho(K_{i,\xi_1^n} + K_{\xi_1^n, \xi_2^n} + \dots + K_{\xi_{n-1}^n, \xi_n^n}) \mathbb{I}_{[\tau_n^n < T]}} \right] \\
 &= \mathbb{E} \left[- \exp \left(\rho \int_t^T \psi_{u_s^{*,n}}(X_{s'}^{t,x,u^*}) ds' \right) e^{-\rho \sum_{k=1}^n K_{\xi_{k-1}^n, \xi_k^n}} \right] \\
 &= \mathbb{E} \left[- \exp \left(\rho \left(\int_t^T \psi_{u_s^{*,n}}(X_{s'}^{t,x,u^*}) ds' - \sum_{k=1}^n K_{\xi_{k-1}^n, \xi_k^n} \right) \right) \right] \\
 &= J(t, x, i; u^*).
 \end{aligned}$$

Finally, we need to prove that $u^{*,n}$ is optimal over the strategies in $\mathcal{U}^{i,n}$, i.e. that for any other strategy $v \in \mathcal{U}^{i,n}$, $J(t, x, i; u^{*,n}) \geq J(t, x, i; v)$. Let $v = ((i, t), (\hat{\xi}_1^n, \hat{\tau}_1^n), \dots)$ be an strategy in $\mathcal{U}^{i,n}$. Then

$$\begin{aligned} V_t^n(t, x, u, i, t) &= \operatorname{ess\,sup}_{\tau \in S_t} \mathbb{E} \left[\exp \left(\rho \int_t^\tau \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) \mathbb{I}_{[s' \geq t]} ds' \right) \mathcal{M}V_\tau^n(t, x, u \oplus (i, t), i, t) \middle| \mathcal{F}_t \right] \\ &\geq \mathbb{E} \left[\exp \left(\rho \int_t^{\hat{\tau}_1^n} \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) ds' \right) \mathcal{M}V_{\hat{\tau}_1^n}^n(t, x, u \oplus (i, t), i, t) \right], \end{aligned}$$

since τ_1^n is optimal. Now, by proposition 3.7.3(b), we have that

$$\begin{aligned} \mathcal{M}V_{\hat{\tau}_1^n}^n(t, x, u \oplus (i, t), i, t) &= \max_{j \neq i} \left(V_{\hat{\tau}_1^n}^{n-1}(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i,j} \mathbb{I}_{[\hat{\tau}_1^n < T]}} \right) \\ &= \max_{j \neq i} \left(V_{\hat{\tau}_1^n}^{n-1}(t, x, u \oplus (i, t), j, \hat{\tau}_1^n) e^{-\rho K_{i,j} \mathbb{I}_{[\hat{\tau}_1^n < T]}} \right) \\ &\geq V_{\hat{\tau}_1^n}^{n-1}(t, x, u \oplus (i, t), \hat{\xi}_1^n, \hat{\tau}_1^n) e^{-\rho K_{i,\hat{\xi}_1^n} \mathbb{I}_{[\hat{\tau}_1^n < T]}} \end{aligned}$$

and therefore

$$\begin{aligned} V_t^n(t, x, u, i, t) & \tag{3.15} \\ &\geq \mathbb{E} \left[\exp \left(\rho \int_t^{\hat{\tau}_1^n} \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) ds' \right) V_{\hat{\tau}_1^n}^{n-1}(t, x, u \oplus (i, t), \hat{\xi}_1^n, \hat{\tau}_1^n) e^{-\rho K_{i,\hat{\xi}_1^n} \mathbb{I}_{[\hat{\tau}_1^n < T]}} \right]. \end{aligned}$$

Similarly, we have that

$$\begin{aligned} &V_{\hat{\tau}_1^n}^{n-1}(t, x, u \oplus (i, t), \hat{\xi}_1^n, t) \\ &\geq \mathbb{E} \left[\exp \left(\rho \int_{\hat{\tau}_1^n}^{\hat{\tau}_2^n} \psi_{\hat{\xi}_1^n}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\hat{\xi}_1^n, \hat{\tau}_1^n)}) ds' \right) \right. \\ &\quad \left. V_{\hat{\tau}_2^n}^{n-1}(t, x, u \oplus (i, t) \oplus (\hat{\xi}_1^n, \hat{\tau}_1^n), \hat{\xi}_2^n, \hat{\tau}_2^n) e^{-\rho K_{\hat{\xi}_1^n, \hat{\xi}_2^n} \mathbb{I}_{[\hat{\tau}_2^n < T]}} \middle| \mathcal{F}_{\hat{\tau}_1^n} \right], \end{aligned}$$

and putting this expression back into the right hand side of (3.15), we obtain that

$$\begin{aligned} V_t^n(t, x, u, i, t) &\geq \mathbb{E} \left[\exp \left(\rho \int_t^{\hat{\tau}_2^n} \psi_{v_s}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\hat{\xi}_1^n, \hat{\tau}_1^n)}) ds' \right) e^{-\rho(K_{i, \hat{\xi}_1^n} + K_{\hat{\xi}_1^n, \hat{\xi}_2^n}) \mathbb{I}_{[\hat{\tau}_2^n < T]}} \right. \\ &\quad \left. V_{\hat{\tau}_2^n}^{n-2}(t, x, u \oplus (i, t) \oplus (\hat{\xi}_1^n, \hat{\tau}_1^n), \hat{\xi}_2^n, \hat{\tau}_2^n) \right]. \end{aligned}$$

Repeating these steps as many times as necessary it follows that

$$\begin{aligned} V_t^n(t, x, u, i, t) &\geq \mathbb{E} \left[-\exp \left(\rho \left(\int_t^T \psi_{v_s}(X_{s'}^{t,x,v}) ds' - \sum_{k=1}^n K_{\hat{\xi}_{k-1}^n, \hat{\xi}_k^n} \right) \right) \right] \\ &= J(t, x, i; v), \end{aligned}$$

and therefore $V_t^n(t, x, u, i, t) = J(t, x, i; u^*) \geq J(t, x, i; v)$, as we wanted. \square

We turn now to Proposition 3.9, which is the main result of this chapter. Here we characterise optimal solutions to the risk-sensitive optimal switching problem and show that the risk-sensitive value function $V(t, x, i)$ satisfies the multiplicative dynamic equation in (3.10). Actually, our objective is to demonstrate that under the given assumptions, $\bar{V}_t(t, x, u, i, t)$ actually corresponds to $V(t, x, i)$.

We aim first at proving that the limit process $\bar{V}(t, x, u, \xi, \theta)$ is continuous and then, by mean of the characterisation (3.10) and the properties of the Snell envelope, we show the optimality of the switching control u^* defined below. As in Proposition 3.8, this optimal switching control is defined in terms of a sequence of stopping times, representing the first hitting times of suitably defined processes.

Proposition 3.9. *Let u be an admissible control in $\mathcal{U}^i(t)$. Let us define the sequence of \mathbb{F} -stopping times $(\tau_k^*)_{k \geq 0}$ by $\tau_0^* = t$,*

$$\tau_1^* := \inf \left\{ s \geq t : \bar{V}_s(t, x, u, i, t) = \mathcal{M}\bar{V}_s(t, x, u \oplus (i, t), i, t) \right\} \wedge T,$$

and, for $k \geq 2$,

$$\begin{aligned}\tau_k^* : &= \inf \left\{ s \geq \tau_{k-1}^* : \bar{V}_s(t, x, u \oplus (\xi_0^*, \tau_0^*) \oplus \dots \oplus (\xi_{k-2}^*, \tau_{k-2}^*), \xi_{k-1}^*, \tau_{k-1}^*) \right. \\ &= \mathcal{M}\bar{V}_s(t, x, u \oplus (\xi_0^*, \tau_0^*) \oplus \dots \oplus (\xi_{k-1}^*, \tau_{k-1}^*), \xi_{k-1}^*, \tau_{k-1}^*) \left. \right\} \wedge T,\end{aligned}$$

where, $\xi_0^* = i$,

$$\xi_1^* := \arg \max_{j \neq i} \left\{ \bar{V}_{\tau_1^*}(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i,j} \mathbb{I}_{[\tau_1^* < T]}} \right\},$$

and, for $k \geq 2$,

$$\xi_k^* := \arg \max_{j \neq \xi_{k-1}^*} \left\{ \bar{V}_{\tau_k^*}(t, x, u \oplus (\xi_0^*, \tau_0^*) \oplus \dots \oplus (\xi_{k-1}^*, \tau_{k-1}^*), j, \tau_{k-1}^*) e^{-\rho K_{\xi_{k-1}^*, j} \mathbb{I}_{[\tau_k^* < T]}} \right\}.$$

Then $u^* = ((i, t), (\xi_1^*, \tau_1^*), \dots, (\xi_k^*, \tau_k^*), \dots)$ is optimal for the risk-sensitive optimal switching problem. Moreover,

$$\bar{V}_t(t, x, u, i, t) = \sup_{v \in \mathcal{U}^i(t)} J(t, x, i; v) = V(t, x, i).$$

Proof. Let θ be an stopping time in S_t and ξ an \mathbb{F}_θ -measurable regime in \mathcal{R} . Let us prove first that the process $\bar{V}_t(t, x, u, \xi, \theta)$ is continuous. Note that by (3.10) we have that for every $s \in [t, T]$,

$$\begin{aligned}\bar{V}_s(t, x, u, \xi, \theta) \\ = \operatorname{ess\,sup}_{\tau \in S_s} \mathbb{E} \left[\exp \left(\rho \int_s^\tau \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M}\bar{V}_\tau(t, x, u \oplus (\xi, \theta), \xi, \theta) \middle| \mathcal{F}_s \right],\end{aligned}$$

which implies that the process

$$(Z_s(t, x, u, \xi, \theta))_{s \in [t, T]} := \left(\bar{V}_s(t, x, u, \xi, \theta) \exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t,x,u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \right)_{s \in [t, T]},$$

is the Snell envelope of

$$\begin{aligned} & (G_s(t, x, u, \xi, \theta))_{s \in [t, T]} \\ & := \left(\exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \mathcal{M} \bar{V}_s(t, x, u \oplus (\xi, \theta), \xi, \theta) \right)_{s \in [t, T]}. \end{aligned}$$

Hence, thanks to property 2 in Theorem 3.6, the process $Z(t, x, u, \xi, \theta)$ admits the following decomposition

$$Z_s(t, x, u, \xi, \theta) = M_s(t, x, u, \xi, \theta) - C_s(t, x, u, \xi, \theta) - D_s(t, x, u, \xi, \theta), \quad (3.16)$$

where $(M_s(t, x, u, \xi, \theta))_{s \in [t, T]}$ is a martingale w.r.t. to the Brownian filtration \mathbb{F} (and hence it is a continuous martingale) and $(C_s(t, x, u, \xi, \theta))_{s \in [t, T]}$ and $(D_s(t, x, u, \xi, \theta))_{s \in [t, T]}$ are two non-decreasing processes such that $C_0(t, x, u, \xi, \theta) = D_0(t, x, u, \xi, \theta) = 0$, $C(t, x, u, \xi, \theta)$ is continuous and $D(t, x, u, \xi, \theta)$ is purely discontinuous and predictable. We recall that the process

$$\left(\exp \left(\rho \int_t^s \psi_\xi(X_{s'}^{t, x, u \oplus (\xi, \theta)}) \mathbb{I}_{[s' \geq \theta]} ds' \right) \right)_{s \in [t, T]},$$

is continuous thanks to assumptions 3.4 and 3.5.

Our aim is to show that $D_\tau(t, x, u, \xi, \theta) = 0$ for all $\tau \in S_\theta$ and hence that $\bar{V}(t, x, u, \xi, \theta)$ is continuous. We will prove this by contradiction. Indeed, we will assume that the process $D(t, x, u, \xi, \theta)$ has a jump and show that this implies that $\bar{V}(t, x, u, \xi, \theta)$ does not belong to \mathcal{I}^2 .

Assume that for the control u and regime ξ , there exists a stopping time $\tau \in S_\theta$ such that $\Delta_\tau Z(t, x, u, \xi, \theta) < 0$, or equivalently that $\Delta_\tau D(t, x, u, \xi, \theta) > 0$. Note that, by equation (3.10), it is only interesting to look for jumps after θ . Also, the jumps of $Z(t, x, u, \xi, \theta)$, if any, must be negative, given that $D(t, x, u, \xi, \theta)$ is non-decreasing.

Then, by property 2 in Theorem 3.6, it holds true that $\Delta_\tau G(t, x, u, \xi, \theta) < 0$ and

$Z_{\tau-}(t, x, u, \xi, \theta) = G_{\tau-}(t, x, u, \xi, \theta)$, and consequently,

$$\Delta_\tau \mathcal{M} \bar{V}(t, x, u \oplus (\xi, \theta), \xi, \theta) < 0, \text{ and } \bar{V}_{\tau-}(t, x, u, \xi, \theta) = \mathcal{M} \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi, \theta).$$

Therefore,

$$\mathcal{M}\bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi, \theta) > \mathcal{M}\bar{V}_{\tau}(t, x, u \oplus (\xi, \theta), \xi, \theta),$$

but since \mathcal{R} is finite we have that for some ξ_1 ,

$$\begin{aligned} \mathcal{M}\bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi, \theta) &= \max_{j \neq i} \left(\bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi, \theta) e^{-\rho K_{\xi, j} \mathbb{I}_{[\tau- < T]}} \right) \\ &= \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi_1, \theta) e^{-\rho K_{\xi, \xi_1} \mathbb{I}_{[\tau < T]}}. \end{aligned}$$

Hence,

$$\bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi_1, \theta) e^{-\rho K_{\xi, \xi_1} \mathbb{I}_{[\tau < T]}} > \bar{V}_{\tau}(t, x, u \oplus (\xi, \theta), \xi_1, \theta) e^{-\rho K_{\xi, \xi_1} \mathbb{I}_{[\tau < T]}},$$

which implies that $\bar{V}(t, x, u \oplus (\xi, \theta), \xi_1, \theta)$ has a negative jump at τ . Thus, given the control u and regime ξ , there exists a regime ξ_1 such that

$$\Delta_{\tau} \bar{V}(t, x, u \oplus (\xi, \theta), \xi_1, \theta) < 0, \text{ and } \bar{V}_{\tau-}(t, x, u, \xi, \theta) = \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi_1, \theta) e^{-\rho K_{\xi, \xi_1} \mathbb{I}_{[\tau < T]}}. \quad (3.17)$$

Similarly, given the control $u \oplus (\xi, \theta)$ and the regime ξ_1 , there exists a regime ξ_2 such that

$$\begin{aligned} \Delta_{\tau} \bar{V}(t, x, u \oplus (\xi, \theta) \oplus (\xi_1, \theta), \xi_2, \theta) &< 0 \text{ and} \\ \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi_1, \theta) &= \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta) \oplus (\xi_1, \theta), \xi_2, \theta) e^{-\rho K_{\xi_1, \xi_2} \mathbb{I}_{[\tau < T]}}. \end{aligned}$$

We can repeat this argument many times to show that there exists a sequence of regimes $\xi_1, \xi_2, \dots, \xi_N$ such that (3.17) holds true and for each $2 \leq k \leq N$,

$$\Delta_{\tau} \bar{V}(t, x, u \oplus (\xi, \theta) \oplus (\xi_1, \theta) \oplus \dots \oplus (\xi_{k-1}, \theta), \xi_k, \theta) < 0, \text{ and}$$

$$\begin{aligned} \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta) \oplus (\xi_1, \theta) \oplus \dots \oplus (\xi_{k-1}, \theta), \xi_k, \theta) \\ = \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta) \oplus (\xi_1, \theta) \oplus \dots \oplus (\xi_{k-2}, \theta), \xi_{k-1}, \theta) e^{-\rho \sum_{k=1}^N K_{\xi_{k-1}, \xi_k} \mathbb{I}_{[\tau < T]}}, \end{aligned} \quad (3.18)$$

where conventionally we set $\xi_0 = \xi$. From equations (3.18) and (3.12), it then follows

that

$$\begin{aligned}
\bar{V}_{\tau-}(t, x, u, \xi, \theta) &= \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta), \xi_1, \theta) e^{-\rho K_{\xi_0, \xi_1} \mathbb{I}_{[\tau < T]}} \\
&= \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta) \oplus (\xi_1, \theta), \xi_2, \theta) e^{-\rho K_{\xi_1, \xi_2} \mathbb{I}_{[\tau < T]}} \\
&= \dots \\
&= \bar{V}_{\tau-}(t, x, u \oplus (\xi, \theta) \oplus (\xi_1, \theta) \oplus \dots \oplus (\xi_{N-1}, \theta), \xi_N, \theta) e^{-\rho \sum_{k=1}^N K_{\xi_{k-1}, \xi_k} \mathbb{I}_{[\tau < T]}} \\
&\leq -e^{\rho(T-t)M} e^{-\rho \sum_{k=1}^N K_{\xi_{k-1}, \xi_k} \mathbb{I}_{[\tau < T]}} \\
&\leq -e^{\rho(T-t)M} e^{-\rho N \inf_{i \neq j} \{K_{i,j}\}} \\
&\leq -e^{\rho(T-t)M} e^{-\rho N \epsilon},
\end{aligned}$$

where we have used our assumption that $K_{i,j}(x) > \epsilon > 0$ for every $j \neq i$ and $x \in \mathbb{R}^d$. But then, in the limit when $N \rightarrow \infty$, and recalling that $\Delta_\tau Z(t, x, u, \xi, \theta) < 0$, we get that

$$\bar{V}_\tau(t, x, u, \xi, \theta) < \bar{V}_{\tau-}(t, x, u, \xi, \theta) \leq \infty,$$

which cannot be true, since $\mathbb{E} [\sup_{s \in [t, T]} |\bar{V}_s(t, x, u, \xi, \theta)|^2] < \infty$ by Proposition 3.7.3(c). Therefore there does not exist any $\tau \in S_\theta$ for which $\Delta_\tau Z(t, x, u, \xi, \theta) < 0$. Thus, $\bar{V}(t, x, u, \xi, \theta)$ is continuous on $[t, T]$.

Next we will prove that the switching control u^* is admissible and such that

$$\bar{V}_t(t, x, u, i, t) = J(t, x, i; u^*).$$

This will be similar to what we did in Proposition 3.8. From equation (3.10), it follows that

$$\begin{aligned}
\bar{V}_t(t, x, u, i, t) &= \operatorname{ess\,sup}_{\tau \in S_t} \mathbb{E} \left[\exp \left(\rho \int_t^\tau \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) \mathbb{I}_{[s' \geq t]} ds' \right) \mathcal{M} \bar{V}_\tau(t, x, u \oplus (i, t), i, t) \middle| \mathcal{F}_t \right] \\
&= \operatorname{ess\,sup}_{\tau \in S_t} \mathbb{E} \left[\exp \left(\rho \int_t^\tau \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) ds' \right) \right. \\
&\quad \left. \max_{j \neq i} \left(\bar{V}_\tau(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i,j} \mathbb{I}_{[\tau < T]}} \right) \right].
\end{aligned}$$

Since the process $(\bar{V}_s(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i,j} \mathbb{I}_{[s < T]}})_{s \in [t, T]}$ can only have positive jumps,

we have that τ_1^* is optimal by property in 3 Theorem 3.6, and then

$$\bar{V}_t(t, x, u, i, t) = \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1^*} \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) ds' \right) \max_{j \neq i} \left(\bar{V}_{\tau_1^*}(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i, j} \mathbb{I}_{[\tau_1^* < T]}} \right) \right].$$

Now, by definition, we know that ξ_1^* is \mathcal{F}_{τ_1} -measurable. In fact, we have that ξ_k^* is \mathcal{F}_{τ_k} -measurable for all $k \geq 0$, and thanks to Proposition 3.7.3(b), we get that

$$\begin{aligned} \max_{j \neq i} \left(\bar{V}_{\tau_1^*}(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i, j} \mathbb{I}_{[\tau_1^* < T]}} \right) &= \bar{V}_{\tau_1^*}(t, x, u \oplus (i, t), \xi_1^*, t) e^{-\rho K_{i, j} \mathbb{I}_{[\tau_1^* < T]}} \\ &= \bar{V}_{\tau_1^*}(t, x, u \oplus (i, t), \xi_1^*, \tau_1^*) e^{-\rho K_{i, j} \mathbb{I}_{[\tau_1^* < T]}}. \end{aligned}$$

Hence,

$$\bar{V}_t(t, x, u, i, t) = \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1^*} \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) ds' \right) \bar{V}_{\tau_1^*}(t, x, u \oplus (i, t), \xi_1^*, \tau_1^*) e^{-\rho K_{i, j} \mathbb{I}_{[\tau_1^* < T]}} \right]. \quad (3.19)$$

Similarly, we have that

$$\begin{aligned} &\bar{V}_{\tau_1^*}(t, x, u \oplus (i, t), \xi_1^*, \tau_1^*) \\ &= \mathbb{E} \left[\exp \left(\rho \int_{\tau_1^*}^{\tau_2^*} \psi_{\xi_1^*}(X_{s'}^{t, x, u \oplus (i, t) \oplus (\xi_1^*, \tau_1^*)}) ds' \right) \bar{V}_{\tau_2^*}(t, x, u \oplus (i, t) \oplus (\xi_1^*, \tau_1^*), \xi_2^*, \tau_2^*) e^{-\rho K_{\xi_1^*, \xi_2^*} \mathbb{I}_{[\tau_2^* < T]}} \right]. \end{aligned}$$

Replacing this expression in (3.19), it follows that

$$\begin{aligned} \bar{V}_t(t, x, u, i, t) &= \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1^*} \psi_{u_s^*}(X_{s'}^{t, x, u \oplus (i, t) \oplus (\xi_1^*, \tau_1^*)}) ds' \right) \right. \\ &\quad \left. e^{-\rho(K_{i, \xi_1^*} + K_{\xi_1^*, \xi_2^*}) \mathbb{I}_{[\tau_2^* < T]}} \bar{V}_{\tau_2^*}(t, x, u \oplus (i, t) \oplus (\xi_1^*, \tau_1^*), \xi_2^*, \tau_2^*) \right], \end{aligned}$$

where we have used that

$$\exp \left(\rho \int_t^{\tau_1^*} \psi_i(X_{s'}^{t, x, u \oplus (i, t)}) ds' \right) e^{-\rho K_{i, \xi_1^*} \mathbb{I}_{[\tau_1^* < T]}},$$

is $\mathcal{F}_{\tau_1^*}$ -measurable, that $[\tau_2^* < T] \subset [\tau_1^* < T]$ and that

$$\begin{aligned} \exp \left(\rho \int_t^{\tau_1^*} \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) ds' \right) \exp \left(\rho \int_{\tau_1^*}^{\tau_2^*} \psi_{\xi_1^*}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^*, \tau_1^*)}) ds' \right) \\ = \exp \left(\rho \int_t^{\tau_2^*} \psi_{u_s^*}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^*, \tau_1^*)}) ds' \right). \end{aligned}$$

Repeating this argument as many times as necessary, we obtain

$$\begin{aligned} \bar{V}_t(t, x, u, i, t) \\ = \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_N^*} \psi_{u_s^*}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1^*, \tau_1^*) \oplus \dots \oplus (\xi_{N-1}^*, \tau_{N-1}^*)}) ds' \right) e^{-\rho(K_{i,\xi_1^*} + K_{\xi_1^*, \xi_2^*} + \dots + K_{\xi_{N-1}^*, \xi_N^*}) \mathbb{I}_{[\tau_N^* < T]}} \right. \\ \left. \bar{V}_{\tau_N^*}(t, x, u \oplus (i, t) \oplus (\xi_1^*, \tau_1^*) \oplus \dots \oplus (\xi_{N-1}^*, \tau_{N-1}^*), \xi_N^*, \tau_N^*) \right]. \quad (3.20) \end{aligned}$$

This equality will help us to show that u^* is admissible and that $\bar{V}_t(t, x, u, i, t) = J(t, x, i; u^*)$.

Indeed, suppose that $\mathbb{P}(\{\tau_n^* < T, \forall n \geq 1\}) > 0$, i.e. that u^* is not an admissible control. Hence, from equation (3.20), it follows that

$$\begin{aligned} \bar{V}_t(t, x, u, i, t) &\leq \mathbb{E} \left[-\exp(\rho M(T-t)) \exp(\rho M(T-t)) e^{-\rho(K_{i,\xi_1^*} + K_{\xi_1^*, \xi_2^*} + \dots + K_{\xi_{N-1}^*, \xi_N^*}) \mathbb{I}_{[\tau_N^* < T]}} \right] \\ &\leq -\exp(2\rho M(T-t)) \mathbb{E} \left[e^{-\rho(K_{i,\xi_1^*} + K_{\xi_1^*, \xi_2^*} + \dots + K_{\xi_{N-1}^*, \xi_N^*}) \mathbb{I}_{[\tau_N^* < T]}} \right] \\ &\leq -\exp(2\rho M(T-t)) \mathbb{E} \left[e^{-\rho N \epsilon \mathbb{I}_{[\tau_N^* < T]}} (\mathbb{I}_{\{\tau_n^* < T, \forall n \geq 1\}} + \mathbb{I}_{\{\tau_n^* < T, \forall n \geq 1\}^c}) \right] \\ &\leq -\exp(2\rho M(T-t)) e^{-\rho N \epsilon} \mathbb{P}(\{\tau_n^* < T, \forall n \geq 1\}). \end{aligned}$$

But, by taking the limit as $N \rightarrow \infty$ in the last inequality, we find that $\bar{V}_t(t, x, u, i, t) = \infty$, which contradicts Proposition 3.7.3(c). Therefore, u^* is an admissible switching control. Finally, letting $N \rightarrow \infty$ in (3.20), and using that u^* is admissible and that $\bar{V}_t(t, x, u, i, t) = \lim_{k \rightarrow \infty} V_T^k(t, x, u, i, t) = -1$, we conclude that

$$\begin{aligned} \bar{V}_t(t, x, u, i, t) &= \mathbb{E} \left[-\exp \left(\rho \int_t^T \psi_{u_s^*}(X_{s'}^{t,x,u^*}) ds' \right) e^{-\rho \sum_{k \geq 1} K_{\xi_{k-1}^*, \xi_k^*} \mathbb{I}_{[\tau_k^* < T]}} \right] \\ &= \mathbb{E} \left[-\exp \left(\rho \int_t^T \psi_{u_s^*}(X_{s'}^{t,x,u^*}) ds' \right) e^{-\rho \sum_{t \leq \tau_k^* < T} K_{u_{\tau_{k-1}^*}^*, u_{\tau_k^*}^*}} \right] \\ &= J(t, x, i; u^*). \end{aligned}$$

Finally, we show that for any other admissible control $v \in \mathcal{U}^i(t)$, $\bar{V}_t(t, x, u, i, t) \geq J(t, x, i; v)$. Let $v = ((i, t), (\xi_1, \tau_1), \dots)$ be an admissible control in $\mathcal{U}^i(t)$. Then

$$\begin{aligned} \bar{V}_t(t, x, u, i, t) &= \operatorname{ess\,sup}_{\tau \in S_t} \mathbb{E} \left[\exp \left(\rho \int_t^\tau \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) \mathbb{I}_{[s' \geq t]} ds' \right) \mathcal{M} \bar{V}_\tau(t, x, u \oplus (i, t), i, t) \middle| \mathcal{F}_t \right] \\ &\geq \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1} \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) ds' \right) \mathcal{M} \bar{V}_{\tau_1}(t, x, u \oplus (i, t), i, t) \right], \end{aligned}$$

since τ_1^* is optimal. Using proposition 3.7.3(b), we find that

$$\begin{aligned} \mathcal{M} \bar{V}_{\tau_1}(t, x, u \oplus (i, t), i, t) &= \max_{j \neq i} (\bar{V}_{\tau_1}(t, x, u \oplus (i, t), j, t) e^{-\rho K_{i,j} \mathbb{I}_{[\tau_1 < T]}}) \\ &= \max_{j \neq i} (\bar{V}_{\tau_1}(t, x, u \oplus (i, t), j, \tau_1) e^{-\rho K_{i,j} \mathbb{I}_{[\tau_1 < T]}}) \\ &\geq \bar{V}_{\tau_1}(t, x, u \oplus (i, t), \xi_1, \tau_1) e^{-\rho K_{i,\xi_1} \mathbb{I}_{[\tau_1 < T]}}, \end{aligned}$$

from where we get

$$\bar{V}_t(t, x, u, i, t) \geq \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_1} \psi_i(X_{s'}^{t,x,u \oplus (i,t)}) ds' \right) \right] \quad (3.21)$$

$$\bar{V}_{\tau_1}(t, x, u \oplus (i, t), \xi_1, \tau_1) e^{-\rho K_{i,\xi_1} \mathbb{I}_{[\tau_1 < T]}}. \quad (3.22)$$

In a similar way,

$$\begin{aligned} &\bar{V}_{\tau_1}(t, x, u \oplus (i, t), \xi_1, \tau_1) \\ &\geq \mathbb{E} \left[\exp \left(\rho \int_{\tau_1}^{\tau_2} \psi_{\xi_1}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1, \tau_1)}) ds' \right) \bar{V}_{\tau_2}(t, x, u \oplus (i, t) \oplus (\xi_1, \tau_1), \xi_2, \tau_2) e^{-\rho K_{\xi_1, \xi_2} \mathbb{I}_{[\tau_2 < T]}} \middle| \mathcal{F}_{\tau_1} \right]. \end{aligned}$$

Repeating this argument many times we get

$$\begin{aligned} &\bar{V}_t(t, x, u, i, t) \\ &\geq \mathbb{E} \left[\exp \left(\rho \int_t^{\tau_n} \psi_{v_s}(X_{s'}^{t,x,u \oplus (i,t) \oplus (\xi_1, \tau_1) \oplus \dots \oplus (\xi_{N-1}, \tau_{N-1})}) ds' \right) \right. \\ &\quad \left. \bar{V}_{\tau_N}(t, x, u \oplus (i, t) \oplus (\xi_1, \tau_1) \oplus \dots \oplus (\xi_{N-1}, \tau_{N-1}), \xi_N, \tau_N) e^{-\rho(K_{\xi_1, \xi_2} + \dots + K_{\xi_{N-1}, \xi_N}) \mathbb{I}_{[\tau_N < T]}} \right]. \end{aligned}$$

Now, using that v is admissible and that $\bar{V}_T(t, x, u, \xi, \theta) = -1$, we take the limit as

$N \rightarrow \infty$ to find that

$$\begin{aligned}\bar{V}_t(t, x, u, i, t) &\geq \mathbb{E} \left[-\exp \left(\rho \int_t^T \psi_{v_s}(X_{s'}^{t,x,v}) ds' \right) e^{-\rho \sum_{k \geq 1} K_{\xi_{k-1}, \xi_k}} \mathbb{I}_{[\tau_k < T]} \right] \\ &= J(t, x, i; v).\end{aligned}$$

Therefore, since v was chosen arbitrarily, we can conclude that

$$\bar{V}_t(t, x, u, i, t) = \sup_{v \in \mathcal{U}^i(t)} J(t, x, i; v) = J(t, x, i; u^*) = V(t, x, i). \quad (3.23)$$

□

3.7 Special cases

In Section 3.2 we presented a general optimal switching problem where the state is governed by a controlled diffusion (equation (3.1)) where both the drift coefficient and the diffusion coefficient may depend on the switching control (u_s). We now turn to two special cases of this general problem which are worth considering on their own. As described below, these cases represent a broad class of control problems associated with applications to energy systems.

Uncontrolled process

First we consider a risk-sensitive optimal switching problem where the state process is uncontrolled. Accordingly, we assume now that the drift and the diffusion coefficients of (3.1) are only functions of the system state,

$$b(x, i) = \bar{b}(x), \quad \sigma(x, i) = \bar{\sigma}(x),$$

where $\bar{b} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\bar{\sigma} : \mathbb{R}^d \rightarrow \mathbb{R}^{d' \times d}$. In this case we denote by $X^{t,x}$ the corresponding solution to the SDE (3.1). We note that the payoff functions can still be functions of the switching control. This particular case arises naturally in applications to energy systems where the state dynamics evolves according to a stochastic exogenous process. For example, $X^{t,x}$ may represent the time evolution of commodity prices when the system controller is a price taker. In this context, the control decisions taken by

the controller do not have any significant effects on the prices, in the sense that the control strategy adopted by the controller cannot move prices either up or down.

For this class of problems Proposition 3.9 and equations (3.8) and (3.9) imply that the risk-sensitive value function satisfies for each $t \leq s \leq T$ and regime $i \in \mathcal{R}$,

$$V(s, X_s^{t,x}, i) = \operatorname{ess\,sup}_{\tau \in S_s} \mathbb{E} \left[\exp \left(\rho \int_s^\tau \psi_i(X_{s'}^{t,x}) ds' \right) \mathcal{M}V(\tau, X_\tau^{t,x}, i) \middle| \mathcal{F}_s \right], \quad (3.24)$$

where

$$\mathcal{M}V(s, X_s^{t,x}, i) = \max_{j \neq i} (V(s, X_s^{t,x}, j) e^{-\rho K_{i,j} \mathbb{I}_{[s < T]}}). \quad (3.25)$$

Furthermore, we have that

$$V(s, X_s^{t,x}, i) = \mathbb{E} \left[\exp \left(\rho \int_s^{\tau^*} \psi_i(X_{s'}^{t,x}) ds' \right) \mathcal{M}V(\tau^*, X_{\tau^*}^{t,x}, i) \middle| \mathcal{F}_s \right], \quad (3.26)$$

where

$$\tau^* = \inf \{ s' \in [s, T] : V(s', X_{s'}^{t,x}, i) = \mathcal{M}V(s', X_{s'}^{t,x}, i) \} \wedge T,$$

is the first switching time after s , and that $(V(s, X_s^{t,x}, i))_s$ is a martingale on $[t, \tau^*]$ with respect to the filtration $(\mathcal{F}_s)_{s \in [t, T]}$. In the next chapter we exploit this characterisation of the value function in order to develop an efficient algorithm for the numerical solution of risk-sensitive optimal switching control problems.

Uncontrolled diffusion coefficient

We also consider a situation in which only the drift (but not the diffusion) coefficient of the SDE (3.1) is a function of the switching control. More specifically, we suppose that the process $X^{t,x,u}$ can be written in terms of two (possibly multi-dimensional) components $X_s = (P_s, C_s)$ such that

$$dX_s = \begin{bmatrix} dP_s \\ dC_s \end{bmatrix} = \begin{bmatrix} \tilde{b}(P_s) \\ a(C_s, u_s) \end{bmatrix} ds + \begin{bmatrix} \tilde{\sigma}(P_s) & 0 \\ 0 & 0 \end{bmatrix} dW_s. \quad (3.27)$$

Note that the P component does not depend on the switching control u , and hence it represents an exogenous process (e.g. the energy price process) as in the previous case. Also, the C component depends explicitly on the control u but not on P . Finally, the diffusion coefficient is independent on the switching control u .

Typical examples having the model in (3.27) originate from applications to energy (gas, electricity, heat) storage problems (Carmona and Ludkovski, 2010, van de Ven et al., 2013, Kitapbayev et al., 2015) and resource extraction problems (Brennan and Schwartz, 1985, Brekke and Øksendal, 1994, Evatt et al., 2010), where inventory levels change according to given output and input rates. In these applications, C_s represents the inventory level at time s , and $a(C_s, u_s)$ is the input/output rate, which depends on the current regime in which the system is being operated and the current inventory level. Thus, it is typically assumed that changes in inventory level are modelled by an ODE,

$$dC_s = a(C_s, u_s)ds, \quad C_t = c,$$

with $a : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ being a measurable Lipschitz function. We shall consider specific inventory models in Example 2, 3 and 4 of Chapter 6. Additionally, in the next chapter we develop a numerical method for solving risk-sensitive optimal switching problems with inventory variables.

Chapter 4

Numerical algorithm for risk-sensitive optimal switching

This chapter proposes a numerical method for solving the class of risk-sensitive optimal switching problems in Section 3.7. The method is a simulation-based technique which combines Monte Carlo simulations of the state process and a regression-based approximation of conditional expectation functions. Since our aim is to compare the risk-neutral and risk-sensitive approaches to solving optimal switching problems, we describe algorithms to solving these two cases.

Estimating the conditional expectations is a critical step of the method, specially for the risk-sensitive case. We discuss different ways to carry out this step and, in particular, how to choose basis functions for the regression step. Our choice of basis functions is independent of the payoff/cost functions and adapts to the distribution of the state process over time.

Finally, we present ways to improve the Monte Carlo error associated with these numerical techniques. Specifically, we employ Quasi-Monte Carlo techniques based on antithetic variables and low-discrepancy sequences, and implement them in an illustrative example, demonstrating the computational advantages of using these techniques in conjunction with the proposed numerical algorithms.

We start by developing a numerical method for the uncontrolled diffusion case presented of Section 3.7. Throughout this chapter we assume that the driving process $X^{t,x}$ has the Markov property.

4.1 Discrete time approximation

In practice it is convenient to consider a discrete time version of the original (continuous time) problem in order to find its numerical solution. We consider such approximation here and present some useful results which can be used to justify this approximation.

Let N be a fixed positive integer and consider the following partition of the time interval $[0, T]$,

$$\mathcal{P} = \{0 = t = t_0 < t_1 < \cdots < t_N = T\}$$

For simplicity, the time step between points in the time grid is fixed and equal to $\Delta t = T/N$. We note that the time grid need not be uniform and more general time grids are possible.

In the discrete-time version of the problem, switchings are now restricted to the time points in the grid \mathcal{P} . In this case the analogous of equation (3.24) reads,

For each $s \in \mathcal{P}$,

$$V(s, X_s^{t,x}, i) = \operatorname{ess\,sup}_{\tau \in S_s^{\mathcal{P}}} \mathbb{E} \left[\exp \left(\rho \int_s^{\tau} \psi_i(X_{s'}^{t,x}) ds' \right) \mathcal{M}V(\tau, X_{\tau}^{t,x}, i) \middle| \mathcal{F}_s \right], \quad (4.1)$$

where $S_s^{\mathcal{P}}$ is the set of stopping times taking values in $[s, T] \cap \mathcal{P}$ and

$$\mathcal{M}V(s, X_s^{t,x}, i) = \max_{j \neq i} (V(s, X_s^{t,x}, j) e^{-\rho K_{i,j} \mathbb{I}_{[s < T]}}). \quad (4.2)$$

Furthermore, we have that

$$V(s, X_s^{t,x}, i) = \mathbb{E} \left[\exp \left(\rho \int_s^{\tau^*} \psi_i(X_{s'}^{t,x}) ds' \right) \mathcal{M}V(\tau^*, X_{\tau^*}^{t,x}, i) \middle| \mathcal{F}_s \right], \quad (4.3)$$

where

$$\tau^* = \inf \{s' \in [s, T] \cap \mathcal{P} : V(s', X_{s'}^{t,x}, i) = \mathcal{M}V(s', X_{s'}^{t,x}, i)\} \wedge T,$$

is the first switching time in the grid \mathcal{P} after s , and that $(V(s, X_s^{t,x}, i))_s$ is a martingale on $[t, \tau^*]$. The integrals in (4.1) are interpreted as a summation on the points of \mathcal{P} ,

$$\int_{t_k}^{t_{\ell}} \psi_i(X_s^{t,x}) ds := \sum_{n=k}^{\ell-1} \psi_i(X_{t_n}^{t,x}) \Delta t.$$

In a discrete time formulation, the Snell envelope of $V(s, X_s^{t,x}, i)$ can be constructed via a recursive scheme (see e.g. Chapter 6 in Föllmer and Schied (2002b)), where at each time t_k one can either stay in regime i or switch immediately to another regime j . Therefore, it follows that V satisfies the following multiplicative backward dynamic programming equations:

$$V(t_N, X_{t_N}^{t,x}, i) = -1, \quad (4.4)$$

and for each $k = N - 1, \dots, 0$,

$$V(t_k, X_{t_k}^{t,x}, i) = \max \left\{ e^{\rho \int_{t_k}^{t_{k+1}} \psi_i(X_s^{t,x}) ds} \cdot \mathbb{E} \left[V(t_{k+1}, X_{t_{k+1}}^{t,x}, i) \middle| \mathcal{F}_{t_k} \right], \right. \\ \left. \max_{j \neq i} \left\{ V(t_k, X_{t_k}^{t,x}, j) e^{-\rho K_{i,j}} \right\} \right\} \quad (4.5)$$

$$= \max_j \left\{ e^{\rho \int_{t_k}^{t_{k+1}} \psi_j(X_s^{t,x}) ds} \cdot \mathbb{E} \left[V(t_{k+1}, X_{t_{k+1}}^{t,x}, j) \middle| \mathcal{F}_{t_k} \right] e^{-\rho K_{i,j}} \right\}. \quad (4.6)$$

At this point we recall that for each regime $j \in \mathcal{R}$, $K_{j,j} = 0$. The recursive equations given by (4.4) and (4.6) can be interpreted as follows. Let us suppose that we need to take a control decision at time t_k . Then we can either decide to stay in the current regime at t_k and postpone our decision to switch regimes until time t_{k+1} ; or we can change immediately at t_k to some other regime j . The optimal decision maximises the expected total reward at time t_k in a risk-sensitive sense, i.e. with the exponential-type criterion, upon paying the switching cost.

Note that, when the conditional expectations appearing in (4.6) are known exactly or can be approximated for every k , the value of $V(t_0, x, i)$ can be easily computed using the backward dynamic programming equations (4.5). Then an optimal switching strategy is obtained as byproduct. Assume for now that for each k the corresponding conditional expectation can be approximated by

$$\mathbb{E} \left[V(t_{k+1}, X_{t_{k+1}}^{t,x}, i) \middle| \mathcal{F}_{t_k} \right] \approx \hat{\mathbb{E}} \left[V(t_{k+1}, X_{t_{k+1}}^{t,x}, i) \middle| \mathcal{F}_{t_k} \right].$$

Then the following algorithms output $V(t_0, x, i)$ and an optimal switching strategy.

Algorithm 1

1. Initialise $V(t_N, X_{t_N}^{t,x}, i) = -1$ for each regime $i \in \mathcal{R}$.

2. For each $k = N - 1, \dots, 1$ and for each i , set

$$V(t_k, X_{t_k}^{t,x}, i) = e^{\rho \int_{t_k}^{t_{k+1}} \psi_{j^*}(X_s^{t,x}) ds} \cdot \hat{\mathbb{E}} \left[V(t_{k+1}, X_{t_{k+1}}^{t,x}, j^*) \middle| \mathcal{F}_{t_k} \right] e^{-\rho K_{i,j^*}}, \quad (4.7)$$

$$\text{where } j^* := j_k^*(i) = \arg \max_j \left\{ e^{\rho \int_{t_k}^{t_{k+1}} \psi_j(X_s^{t,x}) ds} \cdot \hat{\mathbb{E}} \left[V(t_{k+1}, X_{t_{k+1}}^{t,x}, j) \middle| \mathcal{F}_{t_k} \right] e^{-\rho K_{i,j}} \right\}.$$

3. Output $V(t_0, x, i) = e^{\rho \int_{t_0}^{t_1} \psi_i(X_s^{t,x}) ds} \cdot \hat{\mathbb{E}} \left[V(t_1, X_{t_1}^{t,x}, i) \middle| \mathcal{F}_{t_1} \right]$.

Note that in this algorithm the estimate $\hat{\mathbb{E}}$ for a given step is used to approximate the value function at that step, which in turn is used in the next one to compute the conditional expectation. Hence the approximation error propagates through every step of Algorithm 1.

An alternative algorithm can be designed noting that at each step of Algorithm 1, rather than saving the value function at every step, we can store only the optimal switching control from that step onwards. This gives us a scheme to reconstruct the risk-sensitive value function and completely avoid the error propagation associated with Algorithm 1. Indeed, note that, if it is optimal to use regime j , say, from t_k up to some future time $\tau_k^*(j)$, we can use this knowledge explicitly to obtain V by first computing the risk-sensitive profit on $[t_k, \tau_k^*(j)]$ using regime j and then switching to some other regime at time $\tau_k^*(j)$. Then the value function is given by (cf. equation (4.3))

$$V(t_k, X_{t_k}^{t,x}, j) = \hat{\mathbb{E}} \left[e^{\rho \int_{t_k}^{\tau_k^*} \psi_j(X_s^{t,x}) ds} \cdot \mathcal{M}V(\tau_k^*, X_{\tau_k^*}^{t,x}, j) \middle| \mathcal{F}_{t_k} \right].$$

From (4.6), and recalling that at each time t_k we either switch or wait until t_{k+1} to take a switching decision, it follows that for each k ,

$$\tau_k^*(i) = \mathbb{I}_{[i=j^*]} \tau_{k+1}^*(i) + \mathbb{I}_{[i \neq j^*]} t_k,$$

where $j^* = j_k^*(i)$ is the argument maximising (4.6). The time $\tau_k^*(i)$ can be interpreted as the optimal switching time after t_k from regime i , whilst $j_k^*(i)$ is the optimal regime at t_k . If the optimal switching control from t_k onwards is known and given by $u_k^*(i)$ for each i , the future profit that can be obtained at t_k is given explicitly by

$$P_{t_k}(X_{t_k}^{t,x}, i) = -e^{\rho \int_{t_k}^{t_{N-1}} \psi_{u_k^*}(X_s^{t,x}) ds} e^{-\rho \sum_{n \geq k} K_{j_n^*, j_{n+1}^*}},$$

and recursively by the equations $P_{t_N}(X_{t_N}^{t,x}, i) = -1$ and

$$P_{t_k}(X_{t_k}^{t,x}, i) = e^{\rho \int_{t_k}^{t_{k+1}} \psi_{j^*}(X_s^{t,x}) ds} \cdot P_{t_{k+1}}(X_{t_{k+1}}^{t,x}, j^*) e^{-\rho K_{i,j^*}}.$$

Therefore $V(t_k, X_{t_k}^{t,x}, i)$ is given by the estimate $\hat{\mathbb{E}} \left[P_{t_k}(X_{t_k}^{t,x}, i) \middle| \mathcal{F}_{t_k} \right]$. Algorithm 2 below exploits the recursion associated with τ_k^* and the process P_{t_k} .

Algorithm 2

1. Initialise $P_{t_N}(X_{t_N}^{t,x}, i) = -1$ and $\tau_N^*(i) = t_N$ for each $i \in \mathcal{R}$.
2. For each $k = N - 1, \dots, 1$ and for each i ,

(a) Set $\tau_k^*(i) = \mathbb{I}_{[i=j^*]} \tau_{k+1}^*(i) + \mathbb{I}_{[i \neq j^*]} t_k$, where

$$j^* = j_k^*(i) := \arg \max_j \left\{ \hat{\mathbb{E}} \left[e^{\rho \int_{t_k}^{t_{k+1}} \psi_j(X_s^{t,x}) ds} \cdot P_{t_{k+1}}(X_{t_{k+1}}^{t,x}, j) \middle| \mathcal{F}_{t_k} \right] e^{-\rho K_{i,j}} \right\} \quad (4.8)$$

(b) Set $P_{t_k}(X_{t_k}^{t,x}, i) = e^{\rho \int_{t_k}^{t_{k+1}} \psi_{j^*}(X_s^{t,x}) ds} \cdot P_{t_{k+1}}(X_{t_{k+1}}^{t,x}, j^*) e^{-\rho K_{i,j^*}}$.

3. Output $V(t_0, x, i) = e^{\rho \int_{t_0}^{t_1} \psi_i(X_s^{t,x}) ds} \cdot \hat{\mathbb{E}} \left[P_{t_1}(X_{t_1}^{t,x}, i) \middle| \mathcal{F}_{t_1} \right]$.

We note that both Algorithm 1 and Algorithm 2 output as byproduct an optimal switching strategy j_k^* , which is computed for each t_k and regime i at step 2.

The algorithms presented above draw from numerical schemes to solving optimal stopping time problems through their dynamic programming formulations. In particular, for the case of valuing American-style options, Algorithm 1 was first suggested by Tsitsiklis and Van Roy (2001), whereas Algorithm 2 was developed by Longstaff and Schwartz (2001) as an efficient improvement to the work of Tsitsiklis and Van Roy. The difference here is that we use the multiplicative dynamic programming equations derived in the previous chapter to give a general approach to solving iterated optimal stopping problems. Furthermore, the numerical solutions arising in the risk-neutral framework utilised by Tsitsiklis and Van Roy (2001) and Longstaff and Schwartz (2001) result as a limiting case of our approach, which allows to specify different degrees of risk-sensitivity ρ .

As indicated above, the only difficulty in implementing either Algorithm 1 or Algorithm 2 is in finding a method for approximating the conditional expectation operators at each time step. This is a well-known statistical problem (Zanger, 2013) which is addressed in the next section. Before presenting methods to tackle this problem we briefly provide justifications for the time discretisation introduced in this section. The rest of the chapter focuses on a simulation-based algorithm to estimate the conditional expectations and how to improve the estimates or approximations implied by the algorithm.

4.1.1 Time discretisation error

It is worth mentioning that the optimal switching controls obtained with the algorithms above are by definition sub-optimal for the continuous-time problem. However, as we have already said, in practice it is necessary to consider a time discretisation of the original control problem. This approximation can have a significant effect on the numerical solution of the original problem which depends on the size of the time step Δt (in an non-regular partition this effect depends on the maximum distance between any two consecutive time points in the partition). As one would expect, the error induced by the approximation improves as $\Delta t \rightarrow 0$.

More specifically, when the state process $X^{t,x}$ can be perfectly simulated on \mathcal{P} , the discretisation error arising from the discrete time approximation to V is of the order $O(N^{-\frac{1}{4}})$ (Bouchard and Warin, 2012). That is to say, to double the accuracy we must in theory increase the number of time grid points by a factor of 16. With the additional assumptions of (weak) semi-convexity (Bally and Pagès, 2003a) of the payoff functions, the error bounds can be improved to $O(N^{-\frac{1}{2}})$, i.e. we would need only to increase the grid points by a factor of four. For details about this error bounds see for example Bouchard and Warin (2012), Bouchard and Chassagneux (2008) and Bally and Pagès (2003a). For the definition and sufficient conditions of semi-convex functions, and when the process X can be only approximated with the Euler scheme, we refer the interested reader to Bally and Pagès (2003a).

We would like to mention that the semi-convexity assumptions are non-restrictive and are satisfied by all the payoff functions used in this thesis. Additionally, all the stochastic processes used in this thesis can be simulated perfectly on any partition of

$[0, T]$. When this is not possible, we suggest to use strong Taylor approximations of stochastic differential equations (Kloeden and Platen, 1992, Platen and Bruti-Liberati, 2010).

4.2 A least-squares Monte Carlo approach

As shown above, in order to compute the optimal switching strategy and the risk-sensitive value function it is sufficient to develop a numerical method to approximate the conditional expectation functions at each t_k . In this section we focus on a regression-based Monte Carlo technique, which combines simulations of the state process along with statistical regression in order to approximate the conditional expectations. The methodology draws on simulation-based methods for valuing American-style financial derivatives. In particular we focus on the approaches proposed by Tsitsiklis and Van Roy (1999) and Longstaff and Schwartz (2001). See also the work of Egloff (2005), whose solution method gives the algorithms of Tsitsiklis and Van Roy and Longstaff and Schwartz as special cases.

The fundamental idea underlying regression-based Monte Carlo techniques consists of simulating N_p sample paths of the process $X^{t,x}, (x_{t_k}^{(n)}), k = 0, \dots, N, n = 1, \dots, N_p$ on the partition \mathcal{P} and expressing the conditional expectations at each t_k as a countable linear combination of \mathcal{F}_{t_k} -measurable basis functions. This relies on the assumption that the conditional expectation operators are square-integrable functions. Consequently, these operators are elements of a Hilbert space, for which one can choose a countable set of (orthonormal) basis functions which spans it (Royden et al., 1988).

Let us fix a set of F_{t_k} -measurable basis functions $\{B_m\}_{m=1,2,\dots}, B_m : \mathbb{R}^d \rightarrow \mathbb{R}$, and a finite positive integer N_b . Note that since $X^{t,x}$ is Markov the basis functions need only be functions of $x_{t_k}^{(n)}$ for each k . Let us also assume that at each t_k the future values $V(t_{k+1}, \cdot, \cdot)$ are known through the backward induction procedure of either Algorithm 1 or Algorithm 2. Accordingly, let us denote these known values generically by $Z(t_{k+1}, x_{t_{k+1}}^{(n)}, j)$, where for each path $x_{t_{k+1}}^n$ and regime j ,

$$Z(t_{k+1}, x_{t_{k+1}}^{(n)}, j) := V(t_{k+1}, x_{t_{k+1}}^{(n)}, j).$$

in the case of the Tsitsiklis–Van Roy approach (Algorithm 1), and

$$Z(t_{k+1}, x_{t_{k+1}}^{(n)}, j) := P_{t_{k+1}}(x_{t_{k+1}}^{(n)}, j),$$

in the case of the Longstaff–Schwartz approach (Algorithm 2). Recall equations (4.7) and (4.8). Then, for each j the estimate of the conditional expectation function at t_k is given by a finite linear combination of N_b basis functions,

$$\begin{aligned} \mathbb{E}\left[V(t_{k+1}, X_{t_{k+1}}^{t,x}, j) \middle| X_{t_k}^{t,x} = x_{t_k}^{(n)}\right] &\approx \hat{\mathbb{E}}\left[V(t_{k+1}, X_{t_{k+1}}^{t,x}, j) \middle| X_{t_k}^{t,x} = x_{t_k}^{(n)}\right] \\ &:= \sum_{m=1}^{N_b} \hat{a}_m B_m(x_{t_k}^{(n)}), \end{aligned}$$

where the coefficients $\hat{a}_m, m = 1, \dots, N_b$, are computed by regressing the known values $\left\{Z(t_{k+1}, x_{t_{k+1}}^{(n)}, j)\right\}_{n=1, \dots, N_p}$ against the basis functions $\left\{B_1(x_{t_k}^{(n)}), \dots, B_{N_b}(x_{t_k}^{(n)})\right\}_{n=1, \dots, N_p}$. That is, $\hat{a}_m, m = 1, \dots, N_b$ are such that the sum of squares

$$\sum_{n=1}^{N_p} \left| Z(t_{k+1}, x_{t_{k+1}}^{(n)}, j) - \sum_{m=1}^{N_b} a_m B_m(x_{t_k}^{(n)}) \right|^2, \quad (4.9)$$

is minimised with respect to $a_m, m = 1, \dots, N_b$.

The regression step

Solving the least squares problem (4.9) is equivalent to solving an associated system of linear equations known as the *normal equations*. For each t_k , the normal equations are given by the system of linear equations

$$\mathbf{B}'\mathbf{B}\mathbf{a} = \mathbf{B}'\mathbf{Z}, \quad (4.10)$$

where $\mathbf{B} = (b_{m,n})_{mn}$ is a $N_p \times N_b$ matrix whose (m, n) entry is given by $(B_m(x_{t_k}^{(n)}))$, $\mathbf{Z} = (Z(t_{k+1}, x_{t_{k+1}}^{(1)}, j), \dots, Z(t_{k+1}, x_{t_{k+1}}^{(N_p)}, j))'$ is a $N_p \times 1$ vector and $\mathbf{a} = (a_1, \dots, a_{N_b})'$ is a $N_b \times 1$ vector.

Many numerical routines are available for solving (4.10). The choice of a particular method will depend on the specific structure of the matrix \mathbf{B} (e.g. whether it is sparse or dense), and the desired compromise between accuracy and computational time.

For example, using a Cholesky decomposition to solving the normal equations is more efficient than using the QR decomposition, although the latter gives more accurate results than the former. Additionally, we can also use matrix factorisations in terms of singular value decompositions (SVD). In the case where the matrix \mathbf{B} is sparse, for example when the basis functions have local support, it is convenient to use specialised numerical packages to solve this kind of sparse matrix equations. This is particularly important for our choice of basis functions, which we introduce in Section 4.4.

It is worthwhile to mention that Algorithm 1 and Algorithm 2 are general algorithms that can be implemented using any alternative techniques to estimating conditional expectations. In addition to the linear regression model presented above, estimates based on non-parametric least squares regression models can be used as well. In the context of Bermudan options, regression based on non-parametric models has been studied and applied by Egloff (2005) and Egloff et al. (2007). Least squares neural networks estimates have been proposed by Kohler et al. (2010), while smoothing splines and kernel regression have been considered by Kohler (2008) and Barty et al. (2008), respectively. Other Monte Carlo methods for estimating conditional expectations are based on Malliavin calculus (which is used to find an alternative representation of conditional expectations), see e.g. Bouchard and Warin (2012) and references therein, and the optimal quantisation approach of Bally and Pagès (2003b), which has been recently used for the numerical solutions of (risk-neutral) optimal switching problems by Gassiat et al. (2012).

In this thesis we employ estimates of conditional expectations based on least squares parametric regression models. This is so because such estimates can be computed efficiently in highly dimensional problems, while maintaining simplicity and ease of implementation. Additionally, abstracting from these techniques allows us to focus on the methodological contributions of our risk-sensitive approach to problems of iterative optimal stopping.

An interesting aspect related with regression-based Monte Carlo methods is the study and analysis of their theoretical convergence rate. This has been a very active area of research since Tsitsiklis and Van Roy (1999) first suggested this class of numerical algorithms. The importance of these studies relies on the fact that they help to understand the relationship between the number of sample paths and the number of

basis functions used for the empirical regression. In addition, these studies give theoretical evidence on whether methods based on parametric or non-parametric techniques are preferable.

Although a rigorous treatment of the theoretical convergence of the algorithms is out of the scope of this thesis, we kindly refer the interested reader to the excellent review on the subject by Zanger (2013) for simpler case of optimal stopping problems, and to the recent publication by Aïd et al. (2014) for optimal switching problems. Our aim here is to provide practical suggestions that can be used to improve the Monte Carlo error associated with the algorithms presented in this chapter. In particular, we suggest in Section 4.4 to use (i) basis functions with local support to ensure that negativity assumption of the value functions is respected, and (ii) quasi-Monte Carlo techniques in Section 4.5 to improve the Monte Carlo error associated with simulation-based methods.

4.2.1 Extension of the algorithm for control-dependent drift coefficient

Next, we extend Algorithm 1 to the case where only the drift coefficient of X (but not the diffusion coefficient) depends on the control. This is the second special case we presented in Section 3.7. Specifically, let us suppose that the process $X^{t,x,u}$ can be written in terms of two (possibly multi-dimensional) components $X_s = (P_s, C_s)$ such that

$$dX_s = \begin{bmatrix} dP_s \\ dC_s \end{bmatrix} = \begin{bmatrix} \tilde{b}(P_s) \\ a(C_s, u_s) \end{bmatrix} ds + \begin{bmatrix} \tilde{\sigma}(P_s) & 0 \\ 0 & 0 \end{bmatrix} dW_s.$$

Our aim now is to compute the risk-sensitive value function $V(t, x, i) = V(t, p, c, i)$ for each initial data (t, p, x, i) . For clarity, we assume that C_s lives in \mathbb{R} , since the multi-dimensional case can be easily extended from this simpler case.

Let us suppose that for every s , $C_{\min} \leq C_s \leq C_{\max}$, and for some N_c let us consider a uniform partition of $[C_{\min}, C_{\max}]$,

$$\mathcal{C} = \{C_{\min} = c_0 < c_1 < \dots, c_{N_c} = C_{\max}\},$$

where $\delta c := c_q - c_{q-1} = (C_{\max} - C_{\min})/N_c$ for every $q = 1, \dots, N_c$. Note that because

of the path-dependency of C on u_s , it cannot be directly simulated, in contrast with the uncontrolled diffusion P_s . Hence we do not directly apply an empirical regression based on Monte Carlo simulations of (P_s, C_s) . Instead, for each $c_q \in \mathcal{C}$ we approximate the risk-sensitive value function $V(\cdot, \cdot, c_q, \cdot)$ with Algorithm 1.

Recall that at each step of the algorithm the values $V(t_{k+1}, P_{t_{k+1}}, C_{t_{k+1}}, j)$ are supposed known. However, given any $c_q \in \mathcal{C}$ and any switching control u_{t_k} , the future value

$$C_{t_{k+1}} = c_q + a(C_{t_k}, u_{t_k})\Delta t,$$

does not necessarily belong to \mathcal{C} , except when $a(C_{t_k}, u_{t_k})\Delta t = M\delta c$ for some positive integer M . Hence, in general we will need to interpolate $V(t_{k+1}, P_{t_{k+1}}, C_{t_{k+1}}, j)$ from the N_c future known values $V(t_{k+1}, P_{t_{k+1}}, \cdot, j)$, for example, by using polynomials, splines, trigonometric functions, and so forth. Algorithm 3 below uses the backward induction procedure to compute $V(t, p, x, i)$.

Algorithm 3

1. Initialise $V(t_N, P_{t_N}^{t,p}, c, i) = -1$ for each regime $i \in \mathcal{R}$ and $c \in \mathcal{C}$.
2. For each $k = N - 1, \dots, 1$, $c \in \mathcal{C}$ and regime i ,
if necessary, interpolate $V(t_{k+1}, P_{t_{k+1}}^{t,p}, C_{t_{k+1}}, j^*)$ and set

$$V(t_k, P_{t_k}^{t,p}, c, i) = e^{\rho \int_{t_k}^{t_{k+1}} \psi_{j^*}(P_s^{t,p}) ds} \cdot \hat{\mathbb{E}} \left[V(t_{k+1}, P_{t_{k+1}}^{t,p}, C_{t_{k+1}}, j^*) \middle| \mathcal{F}_{t_k} \right] e^{-\rho K_{i,j^*}}, \quad (4.11)$$

where

$$j^* := j_k^*(c, i) = \arg \max_j \left\{ e^{\rho \int_{t_k}^{t_{k+1}} \psi_j(P_s^{t,p}) ds} \cdot \hat{\mathbb{E}} \left[V(t_{k+1}, P_{t_{k+1}}^{t,p}, C_{t_{k+1}}, j) \middle| \mathcal{F}_{t_k} \right] e^{-\rho K_{i,j}} \right\}.$$

3. Output $V(t_0, p, c, i) = e^{\rho \int_{t_0}^{t_1} \psi_i(P_s^{t,p}) ds} \cdot \hat{\mathbb{E}} \left[V(t_1, P_{t_1}^{t,p}, C_{t_1}, i) \middle| \mathcal{F}_{t_1} \right]$, possibly by interpolating V for the required value of c .

Estimates of the conditional expectations in Algorithm 3 are obtained by an empirical least squares Monte Carlo regression based on simulations of the Markov process P .

Note that in the context of (risk-neutral) optimal operation of gas storage, similar procedures to Algorithm 3 have been implemented by Carmona and Ludkovski (2010)

and Warin (2012). Here we take advantage of the multiplicative dynamic programming equations derived for the general case of controlled processes, which allows us to incorporate risk-sensitivity to the above special class of control problems. Also, we note that hybrid algorithms using both the Tsitsiklis and Van Roy approach and the Longstaff and Schwartz approach can be designed (Carmona and Ludkovski, 2010). In our experience, although such algorithms tend to run faster than Algorithm 3, they produce a much larger variance and necessitate more sample paths and basis functions to give the same level of accuracy. We believe that Algorithm 3 offers a good compromise between ease of implementation, accuracy and computer time and memory usage. We shall apply this algorithm in Example 2, 3 and 4 in Chapter 6.

4.3 The RS-LSM algorithm

We now summarise our algorithm to find numerical solutions to risk-sensitive optimal switching problems. We call this algorithm the RS-LSM (risk-sensitive least squares Monte Carlo) algorithm.

RS-LSM algorithm

1. Initialisation: Choose $\rho \neq 0, N_p, N, T, N_c$, and a finite set of basis functions.
2. Set $\Delta = T/N$ and simulate N_p sample paths of the uncontrolled process $X^{t,x}$ on $\mathcal{P}, (x_{t_k}^{(n)}), k = 0, \dots, N, n = 1, \dots, N_p$ with fixed initial condition $x_{t_0}^{(n)} = x_0$.
3. For each regime i and sample path n , set the terminal values and optimal switching times after N along each path

$$P_{t_N}(x_{t_N}^{(n)}, i) = -1, \quad (4.12)$$

$$\tau^*(t_N, x_{t_N}^{(n)}, i) = N.$$

4. Backward induction: For each $k = N - 1$ to 1, repeat

- (a) For each regime i regress the known values $\left\{ P_{t_{k+1}}(x_{t_{k+1}}^{(n)}, i) \right\}_n$ onto the basis $\left\{ B_1(x_{t_k}^{(n)}), \dots, B_{N_b}(x_{t_k}^{(n)}) \right\}_n$ to obtain the estimate of the conditional expectation $\hat{\mathbb{E}}[V(t_{k+1}, X_{t_{k+1}}^{t,x}, i) | \mathcal{F}_{t_k}]$.

(b) Compute the continuation values

$$e^{\rho \int_{t_k}^{t_{k+1}} \psi_i(X_s^{t,x}) ds} \hat{\mathbb{E}} \left[V(t_{k+1}, x_{t_{k+1}}^{(n)}, i) | \mathcal{F}_{t_k} \right] e^{-\rho K_{i,j}}. \quad (4.13)$$

(c) Choose optimal switching along each path according to

$$j_k^*(x_{t_k}^{(n)}, i) = \arg \max_j \left\{ e^{\rho \int_{t_k}^{t_{k+1}} \psi_i(X_s^{t,x}) ds} \hat{\mathbb{E}} \left[P_{t_{k+1}}(x_{t_{k+1}}^{(n)}, i) | \mathcal{F}_{t_k} \right] e^{-\rho K_{i,j}} \right\}. \quad (4.14)$$

(d) Update the optimal switching times after k and the RS future rewards along each path,

$$\begin{aligned} \tau_k^*(k, x_k^{(n)}, i) &= \mathbb{I}_{[i=j^*]} \tau_{k+1}^* + \mathbb{I}_{[i \neq j^*]} t_k. \\ P_{t_k}(x_{t_k}^{(n)}, i) &= e^{\rho \int_{t_k}^{t_{k+1}} \psi_{j^*}(x_s^{(n)}) ds} P_{t_{k+1}}(x_{t_{k+1}}^{(n)}, j^*) e^{-\rho K_{i,j^*}}. \end{aligned} \quad (4.15)$$

5. Set $V(t_0, x_0, i) = e^{\rho \int_{t_0}^{t_1} \psi_i(x_0) ds} \cdot \sum_{n=1}^{N_p} P_{t_1}(x_{t_1}^{(n)}, i) / N_p$.

6. Re-scale the value function to compare it with the risk-neutral case,

$$\rho^{-1} \log(-V(t_0, x_0, i)).$$

For completeness, we also outline an algorithm for the solution of risk-neutral optimal switching problems case. We call it the RN-LSM algorithm to distinguish it from the risk-sensitive case. Since the algorithm is analogous to the RS-LSM algorithm we only provide the relevant steps (c.f. Carmona and Ludkovski (2008)).

RN-LSM algorithm

In the RS-LSM algorithm, replace equations (4.12), (4.13), (4.14), and (4.15) by equations (4.16), (4.17), (4.18), and (4.19), respectively,

$$P_{t_N}(x_{t_N}^{(n)}, i) = 0, \quad (4.16)$$

$$\hat{\mathbb{E}} \left[\int_{t_k}^{t_{k+1}} \psi_i(X_s^{t,x}) ds + V(t_{k+1}, x_{t_{k+1}}^{(n)}, i) | \mathcal{F}_{t_k} \right] - K_{i,j}, \quad (4.17)$$

$$j_k^*(x_{t_k}^{(n)}, i) = \arg \max_j \left\{ \hat{\mathbb{E}} \left[\rho \int_{t_k}^{t_{k+1}} \psi_i(X_s^{t,x}) ds + P_{t_{k+1}}(x_{t_{k+1}}^{(n)}, i) | \mathcal{F}_{t_k} \right] - K_{i,j} \right\}, \quad (4.18)$$

$$P_{t_k}(x_{t_k}^{(n)}, i) = \int_{t_k}^{t_{k+1}} \psi_{j^*}(x_s^{(n)}) ds + P_{t_{k+1}}(x_{t_{k+1}}^{(n)}, j^*) - K_{i,j^*}. \quad (4.19)$$

Finally, output $V(t_0, x_0, i) = \int_{t_0}^{t_1} \psi_i(x_0) ds + \sum_{n=1}^{N_p} P_{t_1}(x_{t_k}^{(n)}, i) / N_p$.

RS-LSM algorithm for control-dependent drift coefficient

The algorithm for the control-dependent drift coefficient case is completely analogous to the RS-LSM algorithm, except that now we run step 3. and 4. of the RS-LSM algorithm to each point of the partition \mathcal{C} . The optimal switching control and pathwise future rewards depend now on the variable C_s . The algorithm can be easily adapted from Algorithm 3, so we do not repeat it here. Similar comments to those for the RN-LSM algorithm follow in this case.

4.3.1 Switching regions

In order to describe the risk-sensitive optimal switching strategy, we approximate the *switching regions* and *switching boundaries* (Carmona and Ludkovski, 2008, Pham, 2009) using the output of the RS-LSM algorithm. We recall from equation (3.26) that the optimal switching strategy from regime i is defined by means of hitting times of $X^{t,x}$ to the set

$$\left\{ y \in \mathbb{R}^d : V(s, y, i) = \max_{j \neq i} (V(s, y, j) e^{-\rho K_{i,j}}) \right\}.$$

Accordingly, for each $i, j \in \mathcal{R}$ and time t' , we introduce the set

$$\mathcal{S}_{t'}(i, j) = \{ y \in \mathbb{R}^d : V(t', y, j) e^{-\rho K_{i,j}} > \max_{\ell \neq j} V(t', y, \ell) e^{-\rho K_{i,\ell}} \},$$

where it is optimal to immediately switch from regime i to j at time t' . We shall refer to $\mathcal{S}_{t'}(i, j)$ as the switching region from i to j . The switching boundary from i to j is the boundary of $\mathcal{S}_{t'}(i, j)$ regarded as a function of time. The switching boundary describes the threshold values at which it is optimal to change regimes. Recalling that

the value function $V(t, x, i)$ is a continuous function of t , it follows that the switching boundaries are also continuous in t . Also, at each time t and regime i there are $N_r - 1$ switching boundaries, where N_r is the total number of regimes.

Using the definitions above and our notation for the RS-LSM algorithm, we notice that for $i \neq j$ the set

$$\mathcal{S}_{t_k}(i, j) = \left\{ x_{t_k}^{(n)} : n \text{ such that } j_k^*(x_{t_k}^{(n)}, i) = j \right\},$$

defines the set of sample paths $x_{t_k}^{(n)}$ in which it is optimal to switch from regime i to j at time t_k . Thus, this set is an empirical approximation of the switching region $S_{t'}(i, j)$. By keeping track of the $j_k^*(x_{t_k}^{(n)}, i)$, we can then approximate each of the switching regions by summarising the graph of $(t_k, x_{t_k}^{(n)}, j_k^*(x_{t_k}^{(n)}, i))$ for each i as a final step in the algorithm. We notice that this approximation can be carried out in both special cases considered above.

4.4 Choice of basis functions

The accurate estimation of the conditional expectations is a critical step in any of the algorithms above. This implies that the choice of function bases is very important and must be considered carefully. The appropriate choice of function bases has been a subject of discussion in the literature. Originally, Tsitsiklis and Van Roy (1999) and Longstaff and Schwartz (2001) proposed to use bases of orthogonal (or orthonormal) functions, such as monomials, and Laguerre and Hermite polynomials. These orthogonal polynomial families have been studied by Stentoft (2004), who has also proposed to use Complete set of polynomials (Judd, 1998) for highly dimensional problems. Tensor products of the elements of these families were also proposed for multi-dimensional problems.

A common feature of these basis functions is that they all have global support, i.e. they are defined over the entire domain of the state. Although algorithms based on this choice of basis functions are easy to implement, they face issues such as oscillations of the value function, that the number of bases is increasing (in principle) with the problem dimension, and that increasing the number of functions does not necessarily

improve the convergence rate of the algorithm (Stentoft, 2004).

In our context, a major drawback of using this set of functions is that it is difficult to ensure that the projected values (i.e. the estimates of the continuation values given in equation (4.13)) are strictly negative (for maximisation problems) or strictly positive (for minimisation problems). Clearly, this violates the assumptions made in our mathematical framework and complicates the implementation of the algorithms. To illustrate this, the top panel of Figure 4.1 presents the regressed values $V(t_{k+1}, x_{t_{k+1}}^{(n)}, i)$ (blue markers) and the estimated continuation value (orange markers) obtained by fitting seven monomials for Example 1 in Chapter 6. We have used $\rho = 1.5$. Note the poor fitting of the regression at the far ends of the range of $x^{(n)}$. The reason for this is that the algorithm is trying to fit tail realisations (or rare events) arising from the joint distribution of $(X_{t_k}^{t,x}, X_{t_{k+1}}^{t,x})$. The relative small density of points in those areas gives little information for the regression procedure. As a consequence, when the regressed values are relatively close to zero the fitted values may be positive, as shown in the figure.

We have also experimented with different architectures for the set of functions bases, including orthonormal polynomial families and exponentials functions, all with global support. Moreover, we have tested function bases of different cardinality between three and 15 basis functions. In all these cases we have observed the same results, revealing that in general functions with global support are not robust for the RS-LSM algorithm. We point out however that the negativity assumption is violated in uncommon situations and in subsets where the data is sparse and close to zero. More specifically, our experience reveals that there exists a critical point $\check{\rho}$, depending on the operating regimes in \mathcal{R} , such that when $\rho \geq \check{\rho}$ the regression step produces positive estimates of the conditional expectations. This can be explained by noting that when ρ is large, $e^{-\rho\psi_i(x_{t_k}^{(n)})}$ can be very small depending on the relative values of $\psi_i(x_{t_k}^{(n)})$. Therefore, in such cases the regressed values will tend to accumulate in regions close to zero, causing possible unacceptable estimates of the RS value function.

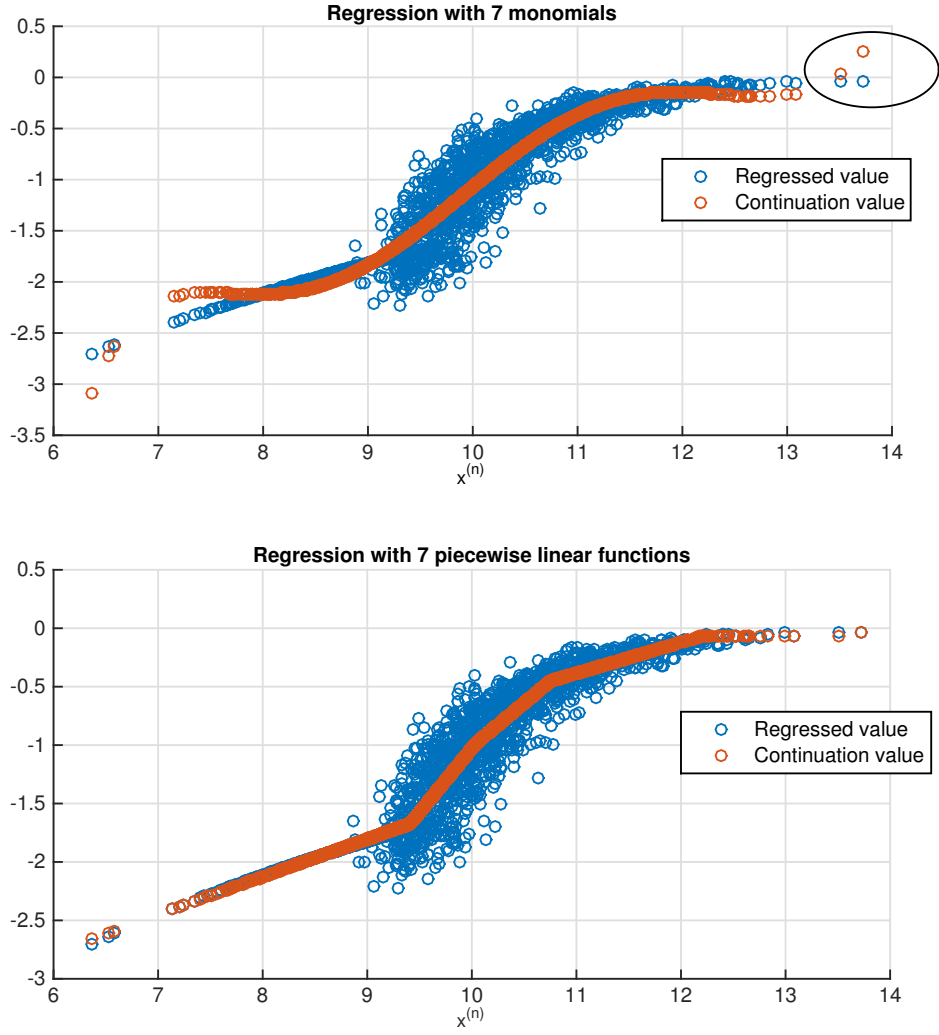


Figure 4.1: The estimated continuation values obtained with different sets of basis functions. The risk-sensitive value function must be strictly negative. The positive continuation values obtained by regressing against monomial functions (top panel) violate this assumption. The piecewise polynomials (bottom panel) provide a better fit and respect the negativity assumption.

One-dimensional local basis functions

To solve this issue, we suggest to use basis functions with local support rather than global support. More specifically, we take as basis functions a set of piecewise linear functions defined on a rectangular subdivision of a finite subdomain of \mathbb{R}^d . Further, we propose these bases to be functions of the step k , so that their support can be adapted to the distribution of $X^{t,x}$ at time t_k .

To fix ideas, let us consider the one-dimensional case and then generalise to multi-dimensional spaces. Accordingly, assume that $X^{t,x} \in \mathbb{R}$ and that at step k of the RS-LSM algorithm the simulated paths are such that $x_{t_k}^{(n)} \in [a, b], n = 1, \dots, N_p$. Let us consider a finite partition of $[a, b]$ into N_e subintervals $[x_e, x_{e+1}], e = 0, \dots, N_e - 1$. Let us define $N_e - 1$ functions $B_e : \mathbb{R} \rightarrow \mathbb{R}$ as follows:

$$B_e(x) = \begin{cases} \frac{x-x_e}{x_e-x_{e-1}}, & x \in [x_{e-1}, x_e] \\ \frac{x_{e+1}-x}{x_{e+1}-x_e} & x \in [x_e, x_{e+1}], \\ 0 & \text{otherwise} \end{cases} \quad e = 1, \dots, N_e - 1,$$

It can be easily shown that the functions B_e are linearly independent (Brenner and Scott, 2008), and clearly, $\text{supp} B_e = [x_{e-1}, x_{e+1}], e = 1, \dots, N_e - 1$. Therefore they form a basis for an $N_e - 1$ dimensional subspace of functions defined on $[a, b]$. We shall estimate the conditional expectations in (4.13) with a linear combination of the functions B_e . It remains to define how the local support of each basis function is defined.

We suggest to adapt the support of each function depending on the distribution $X^{t,x}$ at each t_k . Notice that the intervals $[x_{e-1}, x_e]$ need not be equally spaced. In fact, using this flexibility is what allows us to solve in general the issue at hand. As mentioned earlier, the core of the problem in the regression step lies in regions where the explanatory variables $x_{t_k}^{(n)}$ are sparse and close to zero, and the basis functions with global support do not provide a good fit of the response data.

Our idea is to employ relatively many basis functions with small support at the tails of the distribution of $X_{t_k}^{t,x}$, while using only a relative few number of basis functions but with large support at the central part of the distribution. By doing this, we aim at getting an improved fit where the data is sparse, using only as many basis functions as necessary.

Algorithmically, this is easy to implement. We start with a predetermined number of basis functions and equally spaced local support (or alternatively, having the same density of points). If the projected values in the regression step are all non-positive and the fit is satisfactory, we carry on with the backward induction step. Otherwise, if a projected value is positive, say, in the interval $[x_{e-1}, x_e]$, we split it into two intervals

$[x_{e-1}, \bar{x}_e]$ and $[\bar{x}_e, x_e]$ with $\bar{x}_e = \frac{x_e + x_{e-1}}{2}$ and replace the basis function B_e with two new basis functions on these new intervals. Notice that the number of basis functions increased only by one. We do this on all the problematic intervals and carry out the regression procedure with the new set of basis function. If any of the projected values is positive we keep splitting intervals until getting only non-positive projected values.

In the bottom panel of Figure 4.1, we try this technique using the same realisations of $x_{t_k}^{(n)}$ as in the top panel of the figure. We recall that in the top panel a set of basis functions based on monomials violated the negativity assumption of the risk-sensitive value function. As shown, using in total $N_e = 7$ piecewise linear functions for the regression step provides a much better fit not only in the region where the $x^{(n)}$ are close to zero, but also on the left end of the graph, as we wanted.

d-dimensional local basis functions

When the state has two or more dimensions we consider a bounded subspace D containing the points $x_{t_k}^{(n)} \in \mathbb{R}^d$, and a subdivision of D into disjoint subsets $D_e, e = 1, \dots, N_e$ such that $\cup_{e=1}^{N_e} D_e = D$. The subsets D_e can be chosen to be triangular shapes (in \mathbb{R}^2) or their analogous in higher dimensions, or rectangular shapes or hypercubes in higher dimensions. The support of each basis function is now restricted to the subsets D_e . The basis functions can be constructed by choosing tensor products of the functions given for the one-dimensional case. For example, when the state lives in \mathbb{R}^2 and the elements D_e are rectangular, we take the basis functions defined by

$$B_{e,f}(x, y) = \begin{cases} \frac{x-x_e}{x_e-x_{e-1}} \cdot \frac{y-y_f}{y_f-y_{f-1}}, & (x, y) \in [x_{e-1}, x_e] \times [y_{f-1}, y_f] \\ \frac{x_{e+1}-x}{x_{e+1}-x_e} \cdot \frac{y-y_f}{y_f-y_{f-1}}, & (x, y) \in [x_e, x_{e+1}] \times [y_{f-1}, y_f] \\ \frac{x-x_e}{x_e-x_{e-1}} \cdot \frac{y_{f+1}-y}{y_{f+1}-y_f}, & (x, y) \in [x_{e-1}, x_e] \times [y_f, y_{f+1}] \\ \frac{x_{e+1}-x}{x_{e+1}-x_e} \cdot \frac{y_{f+1}-y}{y_{f+1}-y_f}, & (x, y) \in [x_e, x_{e+1}] \times [y_f, y_{f+1}] \\ 0 & \text{otherwise} \end{cases}$$

Adapting the local support of each basis may be carried out in a similar fashion as in the one-dimensional case. For multi-dimensional problems, we suggest to change one dimension at a time to avoid an exponential increase in the complexity of the algorithm.

At this point we would like to mention that the idea of using functions with local support in regression problems is not new. Our work in this subsection has been inspired by the class of segmented regression models used in the statistics literature. See for example Vito M. R. Muggeo (2008)¹ and references therein. In these models, the relationship between the explanatory variable(s) and the response variable takes the form of piecewise linear functions connected at unknown vertices. This type of model allows to account for situations whereby the effect of the explanatory variable changes or is segmented, say, over time. We already saw an example of this in Figure 4.1.

A similar idea has been long used on the theory of finite element methods for the numerical resolution of partial differential equations (PDE). In this case the piecewise linear basis functions are defined on subdivisions of \mathbb{R}^d , resulting e.g. from a triangulation (or any unstructured or structure grid) of the space. Then the discretised PDE is approximated by a finite sum of basis functions and the coefficients are estimated by solving an associated linear system of equations. Recent references in this direction are Brenner and Scott (2008) and Lapidus and Pinder (2011).

In the context of regression-based algorithms in mathematical finance, Gobet et al. (2005) and Bouchard and Warin (2012) have recently used basis functions based on finite element methods. They choose hypercubes for the support of the basis functions. Our purpose here has been to adapt the support of the basis functions so that the sign of the value function is respected at each iteration of the RS-LSM algorithm.

4.5 Variance reduction techniques

One of the advantages of Monte Carlo schemes, such as the RS-LSM algorithm, is that they are more computationally efficient when solving highly dimensional problems, relative to alternative methods such as finite difference schemes and recombining trees (Glasserman, 2004). Additionally, Monte Carlo algorithms allow us to easily accommodate a wide range of model assumptions at implementation level, e.g. different models for X_t , complex payoff functions ψ and so forth. Nevertheless, the accuracy of the Monte Carlo results, and hence, of the RS-LSM algorithm, depends particularly on

¹I thank Dr Peter Foster for pointing out this reference to me

the number of samples paths N_p used in the procedure. Typically, a moderately large number of paths is necessitated to obtain a reasonable level of accuracy (Carmona and Ludkovski, 2008, 2010, Bouchard and Warin, 2012). In general, the error of the RS-LSM algorithm also depends on the discretisation step Δt , the number of basis functions N_b used in the regression step, the discretisation error on simulating $X^{t,x}$ (if any) and the discretisation on the C variable as in Section 4.2.1.

To improve the convergence rate of the RS-LSM algorithm, and hence the accuracy of numerical results, we employ various variance reduction techniques for Monte Carlo methods. In particular, we shall illustrate below the use of antithetic variables and the Sobol and Halton low-discrepancy sequences with permutations and non-zero leap values. For an overview of these techniques we refer the interested reader to Hull (2009). Quasi-Monte Carlo techniques based on low-discrepancy sequences are studied and analysed for example in Kocis and Whiten (1997) and Glasserman (2004).

One of the benefits of using variance reduction techniques in the RS-LSM algorithm is that they allow us to obtain an improved convergence rate with less computational effort. This is particularly relevant in highly dimensional problems with inventory variables (the second special case in Section 3.7). For accuracy purposes, the partition \mathcal{C} should be in principle sufficiently fine (in our examples taking $\delta c = 1/100$ provides a good approximation). Nonetheless, the use of this discretisation in conjunction with several Monte Carlo samples $x^{(n)}$ is computationally expensive (Kitapbayev et al., 2015), since at each time t_k we need to run N_c regression steps. We believe that variance reduction techniques provide improved efficiency and accuracy of the RS-LSM algorithm. Below, we illustrate this point with one of the examples in Chapter 6.

Comparison of variance reduction techniques

In this section, we study the potential that variance reduction techniques have on improving the Monte Carlo standard error associated with both the RN-LSM and the RS-LSM algorithm. For this, we apply these techniques to estimate the risk-neutral and risk-sensitive value function of Example 1 in Section 6.1.1. We recall that the risk-neutral value function is computed with the RN-LSM algorithm, whereas the risk-sensitive value function is computed with the RS-LSM algorithm. For the latter

we use $\rho = 0.5$ as an example, although similar results were obtained for other values of ρ .

We choose $T = 2, \Delta t = 1/200$ and run the algorithms with $N_p = 2000, 4000, 8000, 16000, 32000$ sample paths of the underlying process X . These paths are simulated using (i) the Mersenne Twister pseudorandom number generator (standard Monte Carlo) as implemented in MATLAB, (ii) antithetic variables ($N_p/2$ plus $N_p/2$ antithetic), (iii) Sobol sequences (scrambled, but not leaped) and (iv) Halton-Leaped sequences (leap value equal to 1949, note that this is a prime number). We refer to Kocis and Whiten (1997) for details on the Sobol and Halton sequences. After obtaining the optimal switching control with each algorithm, we implement it on a forward simulation procedure in order to obtain an estimate of $V(t_0, x_0, i)$, with $t_0 = 0, x_0 = 10$ and i is the ‘off’ regime. We repeat this exercise 60 times and compute the grand average of the estimates and the corresponding Monte Carlo standard error (SE) of the estimates.

Results

As the number of samples paths N_p increases, Table 4.1 provides the mean of the estimated values of $V(t_0, x_0, i)$ and its standard error (SE) across 60 separate runs of each algorithm, for $\rho = 0.5$ (top table) and $\rho = 0$ (bottom table). Inspection of the results reveals that the RS value function converges to a value around 5.78, whereas the RN value function converges to a value around 5.95. We note however that the convergence properties of the algorithms vary depending on the variance technique used. For a given technique, the SE decreases as more sample paths are used for the estimation, as expected. More interestingly, for the same value of N_p , the SE decreases as we proceed from the left column to the right column using the Mersenne Twister generator, antithetic variates, and Sobol and Halton-Leaped sequences. Thus, the SE is the smallest with the Halton-Leaped sequence and the largest with the standard random number generator. Furthermore, the variance obtained with the Halton-Leaped numbers is between 4 to 8 times smaller than that obtained with standard Monte Carlo. The results are consistent for both algorithms. Interestingly, the popular antithetic variates technique seems to not provide a significant reduction of variance in this example (cf. Stentoft (2004)). This is in line with to the discussion of Glasserman (2004).

RS-LSM algorithm, $\rho = 0.5$								
N_p	Standard		Antithetic		Sobol		Halton-Leaped	
	Mean	SE	Mean	SE	Mean	SE	Mean	SE
2000	5.8018	0.1461	5.7780	0.0850	5.7944	0.0711	5.7973	0.0660
4000	5.8036	0.1014	5.7823	0.0603	5.7971	0.0398	5.7929	0.0421
8000	5.7864	0.0610	5.7865	0.0422	5.7877	0.0340	5.7978	0.0296
16000	5.7914	0.0428	5.7803	0.0317	5.7813	0.0239	5.7856	0.0176
32000	5.7884	0.0318	5.7816	0.0227	5.7842	0.0194	5.7833	0.0120

RN-LSM algorithm								
N_p	Standard		Antithetic		Sobol		Halton-Leaped	
	Mean	SE	Mean	SE	Mean	SE	Mean	SE
2000	5.9668	0.1289	5.9866	0.0879	5.9616	0.0676	5.9639	0.0683
4000	5.9657	0.0984	5.9720	0.0583	5.9690	0.0524	5.9597	0.0429
8000	5.9437	0.0690	5.9452	0.0455	5.9549	0.0320	5.9656	0.0288
16000	5.9383	0.0543	5.9549	0.0330	5.9502	0.0298	5.9539	0.0177
32000	5.9476	0.0329	5.9483	0.0178	5.9539	0.0182	5.9520	0.0116

Table 4.1: Convergence of the Monte Carlo standard error (SE) for the value function of Example 1 in Section 6.1.1 using the RS-LSM and the RN-LSM algorithms. The results show the sample mean of $V(t_0, X_0, i_0)$ across 60 separate runs of each algorithm as the number of sample paths N_p increases (top to bottom) and when using different random number generation techniques (left to right).

Figure 4.2 compares the SE obtained with the standard Monte Carlo (orange line) and the Halton-Leaped sequence (blue line) both given with respect to the left y -axis, with the computational time (green line), given with respect to the right y -axis. The computational times are virtually the same for different variance reduction techniques, provided all other parameters remain unchanged. As shown, the computational time is linear with respect to N_p , while the standard error decreases approximately as $O(N_p)^{-1/2}$. Note that this is the theoretical form of the SE for Monte Carlo algorithms (Glasserman, 2004). We also notice that in our example, the SE for the low-discrepancy sequence decreases slower than $O(1/N_p)$, which is the typical convergence rate for this Quasi-Monte Carlo technique. Nonetheless, the convergence constant in this rate, and hence the SE, is much smaller for the Halton-Leaped sequence than for the standard Monte Carlo algorithm.

Combining the computational results with the empirical convergence rates, we

conclude that using Halton-Leaped sequences represents an important gain in computational time for this example. Indeed, instead of using the pure Monte Carlo simulations with 32000 sample paths, we could use only 8000 sample paths generated with Halton-Leaped sequences. By doing so, we obtain much more accurate numerical results using only 25% of the computational time.

As noted earlier, the overall error associated with RS-LSM algorithm depends on other algorithm parameters. Nevertheless, Quasi-Monte Carlo techniques offer a simple yet powerful way to improve the accuracy of the results. In Chapter 6, we will apply the algorithms developed in this chapter together with these simulation techniques on various examples arising from applications to energy systems.




Figure 4.2: The computational time versus the Monte Carlo standard error (SE) obtained with different random number generation techniques. The results are for the value function $V(t_0, X_0, i_0)$ of Example 1 in Section 6.1.1, using $\rho = 0.5$. The standard error (blue and orange lines) are given with respect to left y -axis, while the computational time (green line) is given with respect to the right y -axis.

Chapter 5

Bayesian inference of mixture of Gaussian and non-Gaussian OU processes

So far we have assumed given a process $(X_t)_{t \geq 0}$ representing the energy market prices driving the operation of, say, a power plant at any given time t . To appropriately account for risk factors, however, we need to calibrate a model of (X_t) so that its paths resemble the statistical properties and key features of the specified energy market. By doing this, the RS-LSM algorithm is able to learn a suitable risk-sensitive optimal switching policy for the system under study. The development, analysis and application of such calibration methods is exactly the purpose of this chapter.

Here we are concerned with Bayesian inference in a superposition of Ornstein-Uhlenbeck (OU) processes for energy spot price modelling. The superposition is composed of a diffusion and a number of jump processes, thus accounting for normal variations in prices, but also for different types of price spikes. We develop a missing data methodology for the inference problem based on marked Poisson processes, and propose efficient Markov Chain Monte Carlo algorithms for estimating the spot price models. The inference is exact, in the sense that the invariant distribution of the Markov chain is the posterior distribution of the parameters, free of any discretisation error. We perform extensive simulations to test the algorithms on both simulated and real data corresponding to the APXUK and EEX Phelix indices for electricity spot prices, and the National Grid system average price (SAP) for natural gas in the UK.

Additionally, we discuss the choice of the minimum number of jump components which is necessary to obtain a ‘reasonable’ model. Further, we incorporate non-positive jump processes in our model in order to capture negative price jumps in the EEX and SAP time series studied in this chapter.

The results of this chapter are loosely based on and extend our forthcoming journal article Gonzalez et al. (2015). To start off the chapter we provide an brief overview of the main features of liberalised energy markets.

5.1 Stylised facts of energy prices

Energy markets exhibit a number of key features or so-called stylised facts that make these markets quite different to other competitive markets. These stylised facts arise as a consequence of the mechanisms involving energy markets and the underlying characteristics of traded commodities. In the next paragraphs, we recall these characteristics and provide an overview of some of the most prominent stylised facts of energy prices. Extensive discussions about this subject, and in general, about deregulated energy markets, can be found for example in Weron (2007), Meyer-Brandis and Tankov (2008), Benth et al. (2008) and Benth (2013).

Characteristics of energy markets

Firstly, we recall that currently electricity cannot be directly and economically stored in large quantities. An obvious but significant consequence of this fact is that electricity must be consumed almost immediately after it is produced. It must be said that producers may store electricity indirectly in the form of hydropower, or via gas, oil or coal, from which power can be produced. Nevertheless, the end consumers of electricity cannot do this, and hence a constant balance between power supply and demand is necessitated (Weron, 2014). In contrast, natural gas can be stored in relatively large quantities in underground storage facilities (Felix, 2012), which helps smoothing gas supply levels in periods of high demand or low production. Capacity constraints, however, may lead gas markets to behave like power markets over short periods of time (Benth et al., 2008).

Secondly, in deregulated energy markets prices are set according to the rules of

demand and supply, i.e. at the temporary economic equilibrium for price and quantity resulting from the bidding strategies of market players. Both demand and supply can vary over time according to several factors. Demand is influenced (i) by weather; for instance, in Europe demand rises during the cold season due to an increase use of heating, while in the US demand increases in summer because of extensive use of cooling systems; (ii) by industrial activity, which tends to be high during workdays and low over weekends; and (iii) by patterns of consumption in households, with strong variations at the daily, weekly and seasonal time frames. At the same time, the supply side can be affected by congestion in transmission, outages (due to scheduled maintenance or breakdowns), extreme changes in temperature, which affect hydropower during extreme dry and cold seasons, and the increasing penetration of cheap renewable energy from wind farm and photovoltaics.

Finally, we recall that electricity demand is fairly inelastic, meaning that market players are willing to pay the necessary amount of money to obtain an uninterrupted supply of power (i.e. ‘to keep the lights on’). We remark here that the supply side usually consists of (i) relatively inexpensive but inflexible generation (nuclear and coal units), which can take up to a few hours to increase their output level, and (ii) relatively expensive but flexible generation (oil, gas-fired, hydro units) which can increase production in a matter of minutes. Hence the latter are typically used to cover demand during peak hours, while the former is used to satisfy the so-called base load. Because of substantial differences in production costs, inflexible generation, as opposed to flexible generation, is usually dispatched first.

These three characteristics of energy markets that we have recalled above largely influence the dynamics of energy prices, giving place to price spikes, seasonality in prices and mean reversion, among others. We turn now our attention to these stylised facts.

Price spikes

One of the most visible features of energy prices are the so-called price spikes. They can be described as extreme changes in the spot prices occurring over a short period of time of, say, a couple of hours up to a few days. Price spikes can be either upwards or downwards, so when a spike occurs, prices exhibit a rapid significant increase

(respectively, decrease) which is followed by a steep downwards (upwards) move to their previous level. Figure 5.1 illustrates some price spikes in the UK and German electricity spot markets. It is usual to refer to price spikes as price jumps and we use these terms interchangeably, although, as we shall see later on, jumps or price spikes come in different ‘shapes’.

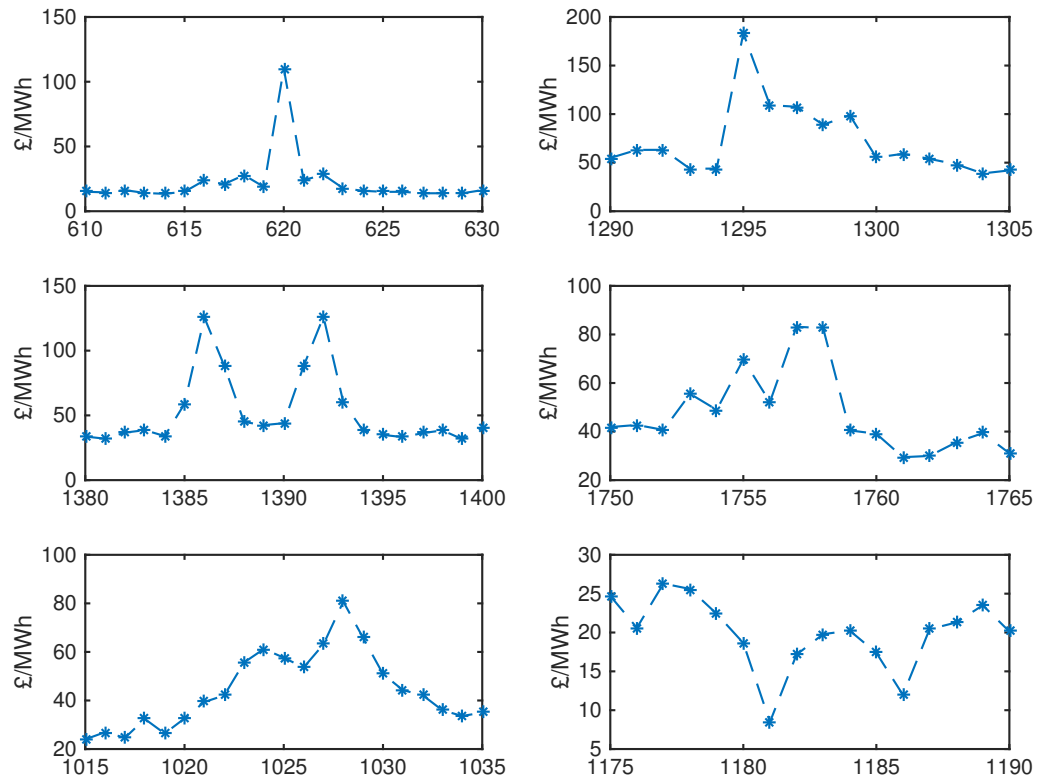


Figure 5.1: Typical patterns of price jumps in the APXUK and EEX indexes for daily average electricity spot prices. The x -axis is given in days.

Price spikes are mainly caused by imbalances between supply and demand due to large shocks on either side. For example, unexpected weather conditions, say, a rapid drop in temperature, can lead to a sudden rise in demand, and therefore to using flexible but expensive generation. In turn, this leads to considerable increases in price level. When the bad weather conditions come back to normal, demand and energy prices decrease accordingly, since the expensive units are needed no longer. A similar phenomenon occurs when a major power unit is closed down, and supply must be kept at the appropriate levels using peak, expensive generators.

Price spikes can be more severe in periods of high demand and when most of the

expensive units are already being used. In this case, small increments in demand level can produce large moves in prices, owing to the bidding strategies of market players (Weron, 2007). Under such market conditions, the price inelasticity of demand is heavily exploited by power generators, who want to maximise their profit. Also, they may be able to execute market power, given that buyers would be willing to accept the high power prices as indicated earlier.

Negative price spikes are mainly caused by the availability of cheap renewable energy and capacity constraints. For example, in the German power exchange EEX, negative spikes have been observed owing to the priority given to wind energy in that network (Benth, 2013). There, when wind-generated power suddenly increases, the electricity generated with other sources (say oil or coal) may be forced to be dumped. In many cases, it might be more economical to pay someone to consume that energy than decrease power output levels or shut down power units. In gas markets, it is the capacity constraints in gas pipeline systems that cause negative price spikes (Benth et al., 2008, p. 9).

It is worth mentioning that spikes tend to be more severe in some markets than in others, and as in the Polish Power Exchange PolPX, they may not occur at all (Weron, 2007). Further, the intensity of price spikes tends to be less in gas markets rather than in electricity markets. This is because gas can be stored in relatively large quantities, allowing to smooth output levels in times of high demand or low production. Additionally, positive spikes tend to be more severe and frequent than negative spikes. We shall see this in Section 5.8 when we study real data from European energy markets.

Typical patterns of price spikes

Price spikes or jumps exhibit different patterns depending on the underlying reason causing them. Figure 5.1 illustrates four price spike patterns that can be observed in the UK's and German's energy spot markets. On each panel, the x -axis is given in days, while the y -axis is given £/MWh.

The first jump pattern (top left panel) is a typical spike, whereby a positive jump in price is followed by a negative jump the next day. The second jump pattern (top right panel) is characterised again by a positive price jump, which however, decays quite slowly. The middle left panel shows that the two previous jump patterns can

occur consecutively. The third jump pattern (middle right panel) is a proper ‘jump’ or a persistent jump, whereby the price has effectively stayed shortly at the position it jumped to. We call the fourth jump pattern a ‘mushroom’ price jump (bottom left panel), since prices increase rapidly over a period of time (of more than one day) without any apparent price jump, and then decrease slowly after reaching a local maximum. The bottom right panel illustrates two negative price spikes. Differentiating between these types of spikes will help us to develop suitable models for energy spot prices, whereby different components account for different jump patterns. The benefits of doing so, in terms of spot price modelling, are exemplified in Section 5.8.

Seasonality

Another typical characteristic of energy markets is seasonality in both prices and demand levels. These seasonal dynamics are explained by the changing weather conditions throughout the year, which translates into differentiated patterns of consumption and/or production over time. Seasonal cycles on the demand and supply sides amount to seasonal cycles in energy spot prices, with these prices being relatively high during periods of high consumption and low during periods of low consumption.

Seasonality in power prices can be observed at hourly, daily and monthly time resolutions (Crespo Cuaresma et al., 2004, Weron, 2007). Hourly prices are driven by the aggregation of everyday people’s behaviour, and can be divided into off peak and on peak prices. Electricity prices tend to be high early in the morning (7am - 9am), when most people get up and start their day, and late in the afternoon (5pm - 7pm), when people arrive home and switch on different home appliances (TV, stove, oven, and so forth). Overnight prices tend to be below the average daily price. This intra-day consumption pattern gives place to a morning peak and an evening peak, with relative sizes varying across different energy markets.

When observed at the daily time resolution, electricity spot prices tend to be higher during weekdays than over weekends due to intense industry activity on workdays. Holidays and major events such as football matches may also have an impact at this time scale. The above intra-day and intra-week patterns remain all-year round, up to increasing and decreasing daily and weekly averages due to calendar seasons. During the cold season, say, in Europe, demand for energy increases due to massive utilisation

of energy for heating and street lightning. This situation is reversed in summer time, when demand drops due to an increase in temperature and long daytimes. At the generation level, hydropower levels are also affected by seasons, since production depends on reservoir levels, which may decrease during summer and freeze during winter.

Seasonality plays a key role in the risk mitigation strategies that we illustrate in the numerical examples of Chapter 6. Our main goal there will be to study the effect that these seasonal price dynamics have on risk-sensitive optimal strategies for different energy systems. In Section 5.2.1 we specify the models used in this thesis to account for intra-day and yearly seasonality in energy prices.

Mean reversion

Another interesting and quite well-known feature of energy prices is mean reversion, which can be described as the tendency of prices to move or revert back to their long term level (Clewlow and Strickland, 2000, Geman and Roncoroni, 2006). Thus, when prices are above their mean reversion level, they will tend to move down. And conversely, when prices are below their mean reversion level, they will tend to move up. We remark that this is only a tendency and long price excursions far from the long term average are possible due to different external factors affecting prices.

The mean reversion level emerges from the constant balance between supply and demand that is required over time, and hence it is linked to the (possible, time-dependent) marginal cost of production (Geman and Roncoroni, 2006). Recalling that demand presents seasonal fluctuations, it is clear that the mean reversion level is also seasonally varying, and can be deterministic or stochastic depending on the market under study.

A related concept to mean reversion and prices spikes is the so-called speed of or time to mean reversion, i.e. the rate at which prices are pulled back to the mean level. Further inspection of Figure 5.1 reveals that some spikes can have a relatively fast speed of mean reversion of up to one day, while others take longer (from two to five days) to revert back to normal levels. We note that this mixture of fast and slow decaying spikes can appear in the same market, making price modelling complex. The time to mean reversion changes for different commodities. For instance, electricity prices tend to have a fast speed of mean reversion, whereas gas prices has a strong

dependence on previous values which is translated into a relatively slow speed of mean reversion.

The empirical evidence regarding the spiky behaviour of energy prices combined with the existence of different rates of mean reversion makes price modelling challenging. Our approach to tackle the problem of price modelling is to incorporate different components in our spot price models to account for most of these different dynamics. In principle, having too many components, and hence too many parameters, makes it difficult to estimate the models. On the other hand, a very simplistic model can fail to capture the key features of energy prices. Therefore, we shall take a parsimonious approach, only increasing the number of components and parameters as necessary. This will be clear in Section 5.8 when studying real data from European energy markets.

5.2 Stochastic modelling of energy markets

Next, we present the class of spot price models that we use to capture the stylised-facts of energy prices, and for which we will propose a Bayesian inference methodology. The building block of this class of models is the Ornstein-Uhlenbeck (OU) process. A real process $Y(t)$, $0 \leq t \leq T$, is called a (non-Gaussian) OU process if it is the unique strong solution to the stochastic differential equation (SDE)

$$dY(t) = \lambda^{-1}(\mu - Y(t))dt + \sigma dL(t), \quad Y(0) = y_0, \quad (5.1)$$

where $L(t)$ is a driving noise process with independent increments. In equation (5.1), $\mu \in \mathbb{R}$ is the mean level to which the process tends to revert, $\lambda^{-1} > 0$ denotes the speed of mean reversion and $\sigma > 0$ is the volatility of the OU process. It can be easily shown (Benth et al., 2008) that the unique strong solution to the SDE (5.1) is given by

$$Y(t) = y_0 e^{-\lambda^{-1}t} + \int_0^t \sigma e^{-\lambda^{-1}(t-s)} dL(s). \quad (5.2)$$

In this thesis, we assume that the process $L(t)$ is a time-homogeneous Lévy process, and refer to it as the background driving Lévy process (BDLP). A common choice for L in the price modelling literature is a standard Wiener process, $L(t) = W(t)$. In this case, it follows that at each time $t \in [0, T]$, $Y(t)$ has a conditional Normal distribution

conditional on $Y(s) = y$, with mean

$$\mathbb{E}[Y(t)|Y(s) = y] = \mu + (y - \mu) e^{\lambda^{-1}s},$$

and variance

$$\text{Var}[Y(t)|Y(s) = y] = \lambda\sigma^2(1 - e^{-2\lambda^{-1}s})/2.$$

Under this specification, the process $Y(t)$ has a Gaussian stationary distribution, and so we will call it a Gaussian OU process.

Another common choice is to take $L(t)$ as a compound Poisson process, which has the so-called *interval representation*

$$L(t) = \sum_{j=1}^{N_t} \xi_j 1_{\{t \geq \tau_j\}}(t), \quad (5.3)$$

where N_t is the number of jumps until time t (inclusive), the τ_j are the arrival times of a Poisson process with intensity rate η and the ξ_j are i.i.d. random variables representing the jump size at time τ_j . In the particular case when $L(t)$ has Exponentially distributed jump sizes, $Y(t)$ has a Gamma stationary distribution (Barndorff-Nielsen and Shephard, 2001). Graphically, the Gamma-OU process $Y(t)$ moves with sharp upward jumps which decay exponentially fast according to the mean reversion parameter λ , making it particularly useful for modelling price spikes. Below we shall model the stochastic part of energy spot prices by superimposing a sum of OU processes.

5.2.1 A multi-factor model for energy spot prices

Let us denote by $S(t)$ the spot price at time $t \geq 0$. We assume that the spot price follows the stochastic process

$$S(t) = e^{f(t)} X(t), \quad (5.4)$$

where $f(t)$ is a deterministic function modelling price seasonality and the deseasonalised price $X(t)$ is a sum of OU processes

$$X(t) = \sum_{i=0}^n w_i Y_i(t), \quad (5.5)$$

where Y_0 is a Gaussian OU process

$$dY_0(t) = \lambda_0^{-1}(\mu - Y_0(t))dt + \sigma dW(t), \quad Y_0(0) = y_0, \quad (5.6)$$

and the $Y_i, i \geq 1$ are non-Gaussian OU process

$$dY_i(t) = -\lambda_i^{-1}Y_i(t)dt + dL_i(t), \quad Y_i(0) = y_i, i = 1, \dots, n, \quad (5.7)$$

with each L_i being a time-homogeneous compound Poisson process with intensity rate $\eta_i > 0$ and exponentially distributed jump sizes with mean β_i . The constants $w_i \in \{1, -1\}$ are used to indicate whether positive and negative jumps are being modelled. Note that since the L_i are increasing processes (in fact, they are subordinators) each of the processes $Y_i(t), i \geq 1$, is non-negative. Thus when $w_i = 1$, the process $Y_i(t)$ is used to capture positive price spikes, whereas when $w_i = -1$, the process $Y_i(t)$ models negative price spikes. We suppose that the following initial condition at $t = 0$ holds

$$S(0) = e^{f(0)} \left(\sum_{i=0}^n w_i Y_i(0) \right).$$

Using a sum of a number of different OU process specifications gives us great flexibility for modelling energy prices. On the one hand, with a Gaussian component we aim to model ‘normal’ market conditions characterised by small variations in energy price during normal energy trading. The jump components aim to model the arrival of information causing imbalances between supply and demand, as discussed earlier. On the other hand, by specifying two jump components, say, such that $\lambda_1 > \lambda_2$, we can capture slowly and quickly decaying prices spikes, which different shock sizes. In addition by taking $w_j = -1$, we can also incorporate negative price spikes, which as was said earlier, are typically found in the German EEX energy market. Although other specifications for the jump size distribution are possible, we take a parsimonious approach by specifying only one parameter for the jump size distribution. We will refer to the superposition model in (5.5) as the n -OU model, where n is the number of OU processes used in the model. In Section 5.7.3, we shall propose a diagnostic tool that will help us to choose a suitable minimum number n of OU components in our mixture model (see also sections 5.9.4, 5.10.3 and 5.11.3).

Seasonality

We assume that the seasonal function $f : [0, \infty) \rightarrow \mathbb{R}$ is bounded. The seasonal function can take different forms depending on the particular features of the data under study. As an example, Benth et al. (2007) apply a sine function with yearly period. To account for increasing prices due to inflation, a linear trend can also be included. More generally, Meyer-Brandis and Tankov (2008) suggest a combination of linear trend and sinusoidal functions with yearly and half yearly period. More sophisticated seasonal functions can also account for holidays, weekends and specific events during the year (Lucia and Schwartz, 2002, Seifert and Uhrig-Homburg, 2007), or they can be based on non-parametric models (Weron, 2007).

In this thesis we use a seasonality function based on a long term trend and superposition of sinusoids,

$$f(t) = a_1 + a_2 t + \sum_{k=1}^m b_k \sin(2\pi t/p_k) + \sum_{k=1}^m c_k \cos(2\pi t/p_k),$$

where p_k are appropriate periodicity constants, m is fixed and the a_k, b_k, c_k are constants to be estimated. Using $m = 1$, $p_1 = 365$ days and $c_1 = 0$, for example, $f(t)$ is used to model a seasonal trend with periodicity of one year. With $m = 2$ and $p_1 = 260, p_2 = 130$ we model yearly and half-yearly seasonality assuming a calendar year of 260 (trading) days.

At this point we would like to mention that other spot price models can be obtained from the deseasonalised price $X(t)$ and seasonal function $f(t)$. For example, geometric spot price models are obtained by taking

$$S(t) = f(t)e^{X(t)},$$

while the arithmetic models can be written as

$$S(t) = f(t) + X(t).$$

See Benth et al. (2008) for more details. The model in (5.4) can be seen as an arithmetic model with multiplicative (rather than additive) seasonality. Our empirical work suggests that geometric models better fit the gas spot prices we study below, whereas

arithmetic models are preferred for the electricity spot price datasets.

Apart from estimating the parameters of the seasonal function, our aim is to develop efficient MCMC algorithms for estimating the set of parameters $\Theta = \{\mu, \lambda_i, \sigma, \eta_i, \beta_i\}$. We shall do this after a literature review on the superposition models in (5.5) and the existing methods to make inference on them.

5.3 Literature review

Superposition of OU processes was introduced by Barndorff-Nielsen and Shephard (2001) as a building component for general and flexible stochastic volatility models. Motivated by this work, Benth et al. (2007) has recently suggested to employ this class of models for stochastic energy spot prices, including new model features relevant to energy markets. Since then, superposition of OU models has become fundamental for capturing the spot dynamics of energy markets (Benth et al., 2008). Apart from offering sufficient analytic tractability for pricing energy derivatives, superpositions of OU processes are attractive for price modelling because of their ability to capture the stylised facts of energy prices, including mean reversion and the short- and long-term price fluctuations of spot prices, by taking different specifications of the BDLP. This flexibility in price modelling is important since, as discussed above, energy markets are in general highly structured and have different fundamental characteristics depending on their geographical location and traded commodities.

Inference methods for superposition of OU processes have been already proposed in the literature. Initially, Benth et al. (2007) suggested two inference methods. The first is a filtering method based on a particle filter (see for example Chib et al. (2002), Johannes et al. (2009) and references therein), whilst the second method is based on prediction based-estimating functions, which are a generalisation of the martingale estimating functions that were suggested to estimate financial models based on diffusions (Sørensen, 2000). These methods were employed and analysed by Benth et al. (2012). A recursive filtering algorithm (Clewlow and Strickland, 2000) is used by Benth et al. (2008) to model UK gas markets, and is shown to perform well on a very restricted 2-OU model.

Drawing from the literature of non-parametric statistics, Meyer-Brandis and Tankov

(2008) suggest two alternative filtering methods for 2-OU models. The aim of these filters is to identify the spike path described by a non-negative process $Y_1(t)$, after which one can use the filtered series to estimate the Gaussian process $Y_0(t)$. The filtering methods rely on setting *a priori* a target volatility for the diffusion, which informs the stopping rule for the filters. In addition, the methods depend on the values of λ_0 and λ_1 , which are obtained by fitting the theoretical autocorrelation function of the model to the observed data. In an interesting application to the EEX electricity, Meyer-Brandis and Tankov provide reasonable estimates of model parameters and good flexibility in terms of modelling the BDLP $L(t)$.

A unified framework for the inference problem of superposition of (non-Gaussian) OU processes is based on a Bayesian paradigm. Bayesian inference for these models was suggested independently by many authors in the discussion section of Barndorff-Nielsen and Shephard (2001). See Roberts et al. (2004a), Griffin and Steel (2006) and Frühwirth-Schnatter and Sögner (2009) for a full report of the methods and simulation results in this context. All these authors, who deal with stochastic volatility models, noticed that the marginal likelihood of the observed data is not analytically tractable, and therefore suggested a missing data methodology whereby the state space is augmented with a latent process.

Inspired by these ideas, we adopt here a Bayesian approach and develop a missing data methodology for making Bayesian inference in the superposition of OU processes in (5.5). Furthermore, we develop efficient MCMC algorithms to simulate from the model parameters and the missing data, and apply them in models consisting of two and three OU processes. We show that the algorithms are easily extended to a general number of OU processes, including superposition of non-negative and non-positive jump OU processes.

5.4 Bayesian inference for a sum of two OU processes

We turn now to presenting techniques for Bayesian inference in the superposition model (5.5) using a missing data methodology. We shall first focus on the case of one jump OU component ($n = 1$), and then show in Section 5.6 how to extend our

proposed methodology to the case of multiple jump components. For simplicity, we drop the sub-script in the jump process L_1 and its parameters, so that $L = L_1$, $\beta = \beta_1$ and $\eta = \eta_1$.

Likelihood-based inference, and in particular Bayesian inference, require the use of the likelihood function

$$\ell(X \mid \mu, \lambda_0, \sigma, \eta, \beta) = \int \ell(X \mid \mu, \lambda_0, \sigma, Y_0, Y_1) \pi(dY_0 \mid \mu, \lambda_0, \sigma) \pi(dY_1 \mid \lambda_1, \eta, \beta) \quad (5.8)$$

$$= \int \ell(X \mid \mu, \lambda_0, \sigma, Y_0) \pi(dY_1 \mid \lambda_1, \eta, \beta), \quad (5.9)$$

where $\pi(dY_i \mid \cdot)$ denotes the prior measure of the process Y_i and is defined by (5.6), (5.7) and the specification of the process L . However, none of the integrals in (5.8) or (5.9) are either analytically tractable or amenable to numerical integration since they are over a highly dimensional space. This makes particularly difficult to use likelihood-based inference. Nevertheless, an alternative approach is to develop a missing data methodology, as discussed previously, whereby the state space is augmented by a suitable process in such way that the likelihood of the augmented space is tractable. We outline next the missing data methodology we shall employ in this thesis.

An advantage of the hybrid Gaussian/non-Gaussian model specified in equation (5.5) is that by means of a change of variables formula the conditional density of the observed data X given the jump process and the parameters can be easily derived. Indeed, since Y_0 is a Gaussian OU process (recall Section 5.2.1) its transition density is given by

$$Y_0(t+s) \mid (Y_0(t) = y) \sim N\left(\mu + (y - \mu)e^{\lambda_0^{-1}s}, \lambda_0\sigma^2(1 - e^{-2\lambda_0^{-1}s})/2\right).$$

Let $X = \{x_0, \dots, x_N\}$ denote observations of the process (5.5) at times $0 = t_0, \dots, t_N = T$, and $\Delta_i = t_i - t_{i-1} > 0$, $i = 1, \dots, N$, the time increment between consecutive observations. Let us suppose that Y_1 is also given at times t_i , $Y_1 = \{y_{1,0}, \dots, y_{1,N}\}$, and write

$$z_i = x_i - y_{1,i}, \quad i = 0, \dots, N. \quad (5.10)$$

Then $Z = \{z_0, \dots, z_N\}$ corresponds to observations of the process Y_0 and is independent of the parameters λ_1, η, β given Y_1 . Thanks to a simple change of variables we have that the augmented likelihood is given by

$$\ell(X \mid \mu, \lambda_0, \sigma, Y_1) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi}\Sigma_i} \exp \left\{ -\frac{1}{2\Sigma_i^2} \left(z_i - \mu - (z_{i-1} - \mu) e^{-\lambda_0^{-1}\Delta_i} \right)^2 \right\}, \quad (5.11)$$

where $\Sigma_i^2 = \lambda_0 \sigma^2 (1 - e^{-2\lambda_0^{-1}\Delta_i})/2$. In principle, the above discussion suggests to treat the process Y_1 as missing data and to develop MCMC algorithms to update the parameters of interest along with Y_1 . This kind of state space augmentation was successfully applied in the context of discrete-time log-normal stochastic volatility models (Jacquier et al., 1994, Kim et al., 1998). However, Barndorff-Nielsen and Shephard (2001, p. 188) argue that this type of augmentation scheme suffers from high posterior correlation because the parameter λ_1 is uniquely determined by knowledge of the trajectory of the process Y_1 , and therefore MCMC samplers based on this parameterisation tend to perform poorly. Further discussion about this issue can also be found in Kastner and Frühwirth-Schnatter (2014).

Instead, Barndorff-Nielsen and Shephard propose an alternative data augmentation scheme based on series representations of integrals with respect to Lévy jump processes $L(t)$ (as in equation (5.2)). With this series representation, also known as the Rosiński representation or Ferguson-Klass representation (see Barndorff-Nielsen and Shephard (2001) and references therein), the process Y_1 can be written explicitly in terms of a sequence of arrival times of a Poisson process and a sequence of uniform random variates in the unit interval. Then the idea is to treat as missing data these two sequences and recover the process Y_1 via the series representation. Bayesian inference for a stochastic volatility model under this parameterisation was first explored in the discussion section of Barndorff-Nielsen and Shephard (2001) and further developed in Griffin and Steel (2006) and Frühwirth-Schnatter and Sögner (2009).

An alternative parameterisation in terms of the arrival times and jump sizes of the compound Poisson process L was independently suggested by different researchers in the discussion of Barndorff-Nielsen and Shephard (2001) and results have been reported by Roberts et al. (2004a). The latter article uses marked Poisson processes to

represent the compound Poisson process L . Further, it proposes two different parameterisations of the augmented model (so called centred and non-centred parameterisations) and provides some guidelines about the context in which they will perform best. In this thesis we follow this approach due to its simplicity and wide applicability to different specifications of the process Y_1 . Nevertheless, from a theoretical perspective, the parameterisation employed here differs from the ones proposed by Roberts et al. (2004a).

5.4.1 Data augmentation based on a marked Poisson process

Next we present a data augmentation methodology for the class of models in (5.5), which takes advantage of the unique correspondence existing between marked Poisson processes and compound Poisson processes.

In Section 5.2.1 we suppose that the BDLP process $L(t)$ of $Y_1(t)$ is a compound Poisson process. Let us denote by $\{(\tau_j, \xi_j)\}$ the set of arrival times and corresponding jump sizes of $L(t)$. Recalling the representation (5.3) and equation (5.2), we have that the dynamics of $Y_1(t)$ are explicitly given by

$$Y_1(t) = Y_1(s)e^{-\lambda_1^{-1}(t-s)} + \sum_{s < \tau_j \leq t} e^{-\lambda_1^{-1}(t-\tau_j)} \xi_j. \quad (5.12)$$

This equation readily shows that the trajectory of Y_1 on $[0, T]$ is completely known given the initial data $Y_1(0)$, the time to mean reversion λ_1 and the pairs of arrival times and jump sizes (τ_j, ξ_j) of the compound Poisson process L . This motivates a data augmentation methodology where the process L , or equivalently, the set of pairs $\{(\tau_i, \xi_i)\}$ is treated as missing data instead of the process Y_1 . The main benefit of doing this is the introduction of statistical independence between λ_1 and the latent process $Y_1(t)$, thus improving the mixing properties of our proposed MCMC algorithms. For the sake of simplicity and without substantial practical impact, we will assume hereafter that $Y_i(0) = 0, i = 1, \dots, n$.

Let Φ be a marked Poisson process living on $S = [0, T] \times (0, \infty)$ with points τ_j on $[0, T]$ and random marks ξ_j on the mark space $(0, \infty)$. Let us denote by $\Phi(S)$ the number of points of Φ on S , with the special case $\Phi(S) = 0$ if Φ has no jumps on S . Note that by construction, $\Phi(S) < \infty$, because L is a compound Poisson process

restricted to the set S .

To be able to make Bayesian inferences in the process Φ we need available its likelihood function. However, the likelihood is not available with respect to the Lebesgue measure since the outcome space for Φ is infinite-dimensional (Geyer and Møller, 1994). Instead, it must be defined with respect to a dominating measure. Accordingly, we choose as dominating measure the measure on S of a marked Poisson process with unit intensity rate and mean jump size equal to one. The density of Φ with respect to this measure is given by

$$\pi(\Phi \mid \eta, \beta) = e^{-(\eta-1)T} \left(\frac{\eta}{\beta} \right)^{\Phi(S)} \exp \left\{ -(\beta^{-1} - 1) \sum_{i=1}^{\Phi(S)} \xi_i \right\}. \quad (5.13)$$

See for example Roberts et al. (2004a), and specially Geyer and Møller (1994) and Møller and Waagepetersen (2004) for more details on marked Poisson process. In Section 5.5.1 we present MH algorithms for posterior inference of Φ based on this density.

5.4.2 A partially non-centred parameterisation

Having incorporated the latent process Φ in our inference problem, our aim now is to simulate from the posterior distribution of the parameters and the missing data given the observed data,

$$\pi(\mu, \lambda_0, \sigma, \lambda_1, \eta, \beta, \Phi \mid X) \propto \ell(X \mid \mu, \lambda_0, \sigma, \lambda_1, \Phi) \pi(\Phi \mid \eta, \beta) \pi(\mu, \lambda_0, \sigma, \lambda_1, \eta, \beta), \quad (5.14)$$

where $\pi(\mu, \lambda_0, \sigma, \lambda_1, \eta, \beta)$ is the joint prior density of the parameters and the likelihood of the data is computed as in (5.11), $\ell(X \mid \mu, \lambda_0, \sigma, \lambda_1, \Phi) = \ell(X \mid \mu, \lambda_0, \sigma, Y_1)$. To do this we will develop a Hastings-within-Gibbs sampler using the data augmentation algorithm reviewed in Section 2.8.

The prior distribution

To complete our Bayesian model we need to specify our choice of prior distributions for the parameters. Throughout this chapter we assume that the parameters of the superposition model with one jump component ($n = 1$) are *a priori* independent of

each other. When $n \geq 2$ we introduce statistical dependence between the mean reversion parameters $\lambda_j, j = 1, \dots, n$ to be able to identify these parameters (see Section 5.6). Also, we choose conjugate priors when possible, so that it is computationally easy to obtain samples from the corresponding posterior distributions. In all cases we use weakly uninformative but proper priors to ensure that a posterior distribution exists.

Our choice of prior distributions is as follows. We assume a $N(a_\mu, b_\mu^2)$ prior distribution for μ , a $IG(a_\sigma, b_\sigma)$ for σ^2 , a $Ga(a_\eta, b_\eta)$ for η and a $IG(a_\beta, b_\beta)$ for β . For the mean reversion parameters $\lambda_j, j = 0, \dots, n$ we formulate $Ga(a_{\lambda_i}, b_{\lambda_i})$ prior distributions. The prior of the latent process Φ is the density of a Poisson process with intensity η and exponentially distributed jumps with mean parameter β .

Prior elicitation for the parameters is facilitated by exploratory analysis of the data under study. Since μ represents the level of mean reversion for the diffusion, we can obtain prior information about this parameter by computing the mean of the deseasonalised data. Additionally, the parameter σ indicates the volatility of the data without jumps, while η is the jump intensity rate. To obtain prior information about these parameters we can remove jumps from the data by using filtering methods such as the recursive filtering algorithm of Clewlow and Strickland (2000, p. 30) or the non-parametric methods of Meyer-Brandis and Tankov (2008). As byproduct we obtain a point estimate of the mean jump size β . Prior information for λ_j may be obtained by fitting the theoretical autocorrelation function (ACF) of the model (5.5) to the empirical ACF of the data. Indeed, assuming that all the Y_i are stationary and square integrable, it follows that the ACF of the process $X(t)$ at lag τ is (Benth et al., 2008)

$$\text{Corr}[X(t), X(t + \tau)] = \sum_{i=0}^n \omega_i e^{-\lambda_i^{-1} \tau}, \quad (5.15)$$

with weights ω_i summing up to 1. We notice, however, that the fitted ACF can give point estimates of the λ_i which are far from estimates obtained directly from the data (see for instance Meyer-Brandis and Tankov (2008), Benth et al. (2012) and our MCMC applications to real data below). Hence, it is not advisable to elicit these prior distributions from the fitted ACF. Nonetheless, expression (5.15) suggests to use the transformation $\rho_i := e^{-\lambda_i^{-1}}, i = 0, \dots, n$, for which we can then use a Uniform prior

distribution on $(0, 1)$ and thus avoid imposing any prior information of the λ_i . We note that choosing a $\text{Ga}(1, 1)$ prior for λ_i implies a $\text{Un}(0, 1)$ prior for ρ_i . Our simulations below are based on the parameterisation $(\mu, \sigma, \rho_i, \eta, \beta)$ and a Uniform prior for ρ_i .

Comments on the model parameterisation

Parameterisations of hierarchical models are often classified as centred and non-centred, according to the structure dependence among the parameters, the observed data and the missing data (see e.g. Roberts et al. (2004a), Sermaidis et al. (2013) and Kastner and Frühwirth-Schnatter (2014), among others). A model parameterisation is said to be *centred* if, given the missing data, the parameters are conditionally independent of the observed data. Alternatively, the parameterisation is said to be *non-centred* if the parameters and the missing data are *a priori* independent. We refer to Papaspiliopoulos et al. (2003) for more details about this terminology. An important feature of the model parameterisation in (5.14) is that the Poisson intensity rate η and the mean jump size parameter β are conditionally independent of the observed data X given the missing data Φ . Thus it is a centred parameterisation with respect to Φ , η and β . Nonetheless, since the missing data is *a priori* independent of the remaining parameters, (5.14) can also be seen as a non-centred parameterisation in terms of Φ and $\mu, \sigma, \lambda_0, \lambda_1$. Therefore the model parameterisation (5.14) may be classified as a partially non-centred parameterisation, since it lies intermediate between fully centred and non-centred parameterisations.

A specification of the missing data that breaks the dependence structure of η, β and Φ (a non-centred parametrisation) involves a marked Poisson point process with unit intensity $\tilde{\Phi} = \{(\tau_i, m_i, \xi_i)\}$ on $[0, T] \times (0, \infty) \times [0, 1]$ and marks $\{\xi_i\}$ Uniformly distributed on $[0, 1]$. A Poisson thinning argument shows that $\{\tau_i\}$ such that $m_i \leq \eta$ form a Poisson process with intensity η on $[0, T]$ and corresponding marks $-\beta^{-1} \log(\xi_i)$ are exponentially distributed with parameter β (Roberts et al., 2004a). Although preferred for certain models, our practical experience has showed that this parametrisation does not outperform the partially non-centred one when working with both simulated and real data. We think the reason for this is that the data tends to be quite informative about the occurrence of jumps in our model, given that sizeable jumps in the observed data can hardly be explained by the Gaussian OU process, and thus they tend to

appear clearly in the augmented likelihood of Z , or equivalently, of X .

Parameterisations of hierarchical models play a crucial role in the performance of MCMC methods, in particular, when Gibbs steps are employed (Papaspiliopoulos et al., 2003, Sermaidis et al., 2013). On the one hand, high posterior correlation between the parameters induces slow mixing and hence slow exploration of the parameter space. On the other hand, when the missing data is well identified, the Gibbs sampler becomes very efficient. Fully centred and non-centred parameterisations have their advantages, and there is no definite word on when one will perform better than the other, since this depends on the problem at hand and the value of the parameters to be estimated (Kastner and Frühwirth-Schnatter, 2014). It is worthwhile to note, however, that partially non-centred parameterisations, as the one given in (5.14), can have superior performance properties than both centred and non-centred parameterisations (Papaspiliopoulos et al., 2003).

An alternative approach would have been to develop firstly a fully non-centred parameterisation, as explained above, and then *interweave* it with (5.14) using an ancillarity-sufficiency interweaving strategy (ASIS) recently introduced by Yu and Meng (2011). The purpose of this is to develop robust algorithms which exploit the advantages of MCMC schemes based on these separate parameterisations. Practical applications of ASIS can be found in Sermaidis et al. (2013) and Kastner and Frühwirth-Schnatter (2014). Nevertheless, we have decided to focus on the parameterisation in (5.14), since in practical applications it demonstrated good robustness properties when working with simulated and real data. We illustrate this point in Section 5.7 and Section 5.8.

5.5 MCMC implementation

We use a Hastings-within-Gibbs algorithm to simulate from the posterior distribution of the parameters given in (5.14). The choice of prior distributions that was suggested in the previous section will prove computationally convenient since it will allow us to design Gibbs samplers for some of the parameters. Our proposed MCMC algorithm for the 2-OU model cycles through the following steps:

MCMC algorithm for the 2-OU model

- Step 1: update $\mu \sim \pi(\mu \mid \sigma, \rho_0, \rho_1, X, \Phi)$.
 Step 2: update $\sigma^2 \sim \pi(\sigma^2 \mid \rho_0, \rho_1, X, \Phi)$.
 Step 3: update $(\rho_0, \rho_1) \sim \pi(\rho_0, \rho_1 \mid \mu, \sigma, X, \Phi)$.
 Step 4: update $\eta \sim \pi(\eta \mid \Phi)$.
 Step 5: update $\beta \sim \pi(\beta \mid \Phi)$.
 Step 6: update $\Phi \sim \pi(\Phi \mid \mu, \rho_0, \sigma, \rho_1, \eta, \beta, X)$.

Below we give more details about each of these updating steps. We recall first the expression for the augmented likelihood (see equation (5.11) and also Section 5.4.1)

$$\ell(X \mid \mu, \lambda_0, \sigma, \rho_1, \Phi) \propto \frac{1}{\prod_{i=1}^N \Sigma_i} \exp \left\{ -\frac{1}{2} \sum_{i=1}^N \frac{1}{\Sigma_i^2} \left(z_i - z_{i-1} e^{-\lambda_0^{-1} \Delta_i} + \mu \left(e^{-\lambda_0^{-1} \Delta_i} - 1 \right) \right)^2 \right\},$$

where $\Sigma_i^2 = \lambda_0 \sigma^2 (1 - e^{-2\lambda_0^{-1} \Delta_i})/2$, and the z_i 's are computed as the difference between observations of X and the trajectory of Y_1 implied by a realisation of the marked Poisson process Φ . We recall also the transformation $\rho_j = e^{-\lambda_j^{-1}}$.

Update μ

Using the conjugate prior for μ specified in the previous section, it can be easily shown that its full conditional distribution is

$$\mu \mid \cdot \sim N \left(\frac{\sum_{i=1}^N \left(1 - e^{-\lambda_0^{-1} \Delta_i} \right) \Sigma_i^{-2} \left(z_i - z_{i-1} e^{-\lambda_0^{-1} \Delta_i} \right) + \frac{a_\mu}{\sigma_0^2}}{\sum_{i=1}^N \left(1 - e^{-\lambda_0^{-1} \Delta_i} \right)^2 \Sigma_i^{-2} + \frac{1}{b_\mu^2}}, \frac{1}{\sum_{i=1}^N \left(1 - e^{-\lambda_0^{-1} \Delta_i} \right)^2 \Sigma_i^{-2} + \frac{1}{b_\mu^2}} \right).$$

Hence we use a Gibbs sampler to update μ .

Update σ^2

A computational advantage of choosing an IG prior for σ^2 is that its full conditional is also available and easy to sample from,

$$\sigma^2 \mid \cdot \sim \text{IG} \left(\frac{N}{2} + a_\sigma, \frac{1}{\lambda_0} \sum_{i=1}^N \frac{s_i}{(1 - e^{-2\lambda_0^{-1}\Delta_i})} + b_\sigma \right),$$

where

$$s_i = \left(z_i - z_{i-1}e^{-\lambda_0^{-1}\Delta_i} + \mu \left(e^{-\lambda_0^{-1}\Delta_i} - 1 \right) \right)^2.$$

Update ρ_0 and ρ_1

Full conditionals are not available for $\rho_i, i = 1, 2$. Hence we use a multiplicative random-walk Metropolis-Hastings algorithm to simulate from the posterior density of ρ_0

$$\pi(\rho_0 \mid \mu, \sigma, \rho_1, X, \Phi) \propto \ell(X \mid \mu, \rho_0, \sigma, \rho_1, \Phi) \pi(\rho_0),$$

and from the posterior density of ρ_1

$$\pi(\rho_1 \mid \mu, \sigma, \rho_0, X, \Phi) \propto \ell(X \mid \mu, \rho_0, \sigma, \rho_1, \Phi) \pi(\rho_1).$$

The variance of the proposal distribution is tuned after pilot runs, so that the acceptance rate is between 20% and 50%.

Update η and β

Since we specified conjugate priors for both η and β , their conditional distributions are given by

$$\begin{aligned} \eta \mid \Phi &\sim \text{Ga}(a_\eta + \Phi(S), T + b_\eta), \\ \beta \mid \Phi &\sim \text{IG} \left(a_\beta + \Phi(S), \sum_{i=1}^{\Phi(S)} \xi_i + b_\beta \right). \end{aligned}$$

Thus we can use Gibbs samplers to simulate η and β .

5.5.1 Update step for the latent process Φ

To sample from the conditional posterior distribution of the marked Poisson process Φ we use a Metropolis-Hastings algorithm. The algorithm was first suggested by Geyer and Møller (1994) in the context of general point processes. Below we describe three types of updates used in our algorithm: birth-and-death (BD) algorithm, local displacement (LD) move and a block update of jump sizes. For references about these algorithms we refer to Geyer and Møller (1994), Roberts et al. (2004a) and Frühwirth-Schnatter and Sögner (2009). See also Moller and Waagepetersen (2004).

Let us assume that the current state of the Markov chain is $\Phi = \{(\tau_1, \xi_1), \dots, (\tau_k, \xi_k)\}$, that is there are k points on the set S with jump times given by τ_j and corresponding jump sizes by ξ_j , $j = 1, \dots, k$.

Birth-and-death algorithm

In order to update Φ , at each MCMC iteration we choose one of two kind of moves. With a probability $p \in (0, 1)$, we choose a birth move, whereby a new-born point (τ, ξ) is added to the current configuration of Poisson points. The new proposed state is then $\Phi \cup \{(\tau, \xi)\}$. The point τ is drawn uniformly from $[0, T]$, whilst ξ is drawn from the jump size distribution $\text{Ex}(\beta)$. For this move the proposal transition kernel has the following density with respect to the reference measure suggested in section 5.4.1

$$b(\Phi, \Phi \cup \{(\tau, \xi)\}) = \frac{1}{T} \exp(-(\beta^{-1} - 1)\xi).$$

The second type of move is a death move whereby with probability $1 - p$ a point $\{(\tau, \xi)\}$ is removed from Φ . Nothing changes if the configuration Φ is empty. The proposal transition kernel for this move is

$$d(\Phi, \Phi \setminus \{(\tau, \xi)\}) = \frac{1}{k}.$$

Let

$$r(\Phi, (\tau, \xi)) = \frac{\ell(X \mid \mu, \rho_0, \sigma, \rho_1, \Phi \cup \{(\tau, \xi)\})}{\ell(X \mid \mu, \rho_0, \sigma, \rho_1, \Phi)} \frac{1-p}{p} \frac{\eta T}{k+1}.$$

Hence the acceptance probability for a birth move is

$$\alpha(\Phi, \Phi \cup \{(\tau, \xi)\}) = \min \{1, r(\Phi, (\tau, \xi))\},$$

and for a death move

$$\alpha(\Phi, \Phi \setminus \{(\tau, \xi)\}) = \min \{1, 1/r(\Phi, (\tau, \xi))\}.$$

We notice that p may be dependent on the current number of Poisson points. This is particularly useful when the initial state of Φ is the empty configuration.

Local displacement move

We also consider a scheme which produces a local change to the latent process Φ . The basic idea is to choose one of the current points of Φ and displace it over a local neighbourhood, so the likelihood outside this subset remains unchanged. The MH transition kernel of this move results as the mixture of k transition kernels, each of them reversible with respect to the posterior of the chosen point (τ_j, ξ_j) conditioned on $\Phi \setminus \{(\tau_j, \xi_j)\}$ (Roberts et al., 2004a). At each iteration of the algorithm we choose one of the k transition kernels randomly.

The local displacement move is as follows. Let us assume that the jump times of the Poisson process are such that $\tau_1 < \dots < \tau_k$ (note that this can be done without loss of generality, since we can always consider a re-ordering of the jump times). Then we choose randomly one of the jump times, say, τ_j , and generate a new jump time τ_{new} uniformly on $[\tau_{j-1}, \tau_{j+1}]$, where we assume that $\tau_0 = 0$ and $\tau_{k+1} = T$, for consistency. The point (τ_j, ξ_j) is then displaced and re-sized applying a deterministic invertible transformation

$$h(\xi_j, \tau_j, \tau_{new}) = (e^{-\lambda_1^{-1}(\tau_{new} - \tau_j)} \xi_j, \tau_{new}, \tau_j) = (\xi_{new}, \tau_{new}, \tau_j),$$

whose Jacobian is given by $J(\xi_j, \tau_j, \tau_{new}) = -e^{-\lambda_1^{-1}(\tau_{new} - \tau_j)}$.

Let us set $\Phi_{new} := \Phi \cup \{(\tau_{new}, \xi_{new})\} \setminus \{(\tau_j, \xi_j)\}$, i.e. Φ_{new} is the resulting configuration of the marked Poisson process after applying the displacement move to Φ . Hence

the acceptance ratio of the resulting MH algorithm is (Tierney, 1998)

$$\alpha(\Phi, \Phi_{new}) = \min \{1, r(\Phi, (\tau, \xi))\},$$

where

$$\begin{aligned} r(\Phi, (\tau_{new}, \xi_{new})) &= \frac{\ell(X \mid \mu, \rho_0, \sigma, \rho_1, \Phi_{new})}{\ell(X \mid \mu, \rho_0, \sigma, \rho_1, \Phi)} \frac{\pi(\Phi_{new} \mid \eta, \beta)}{\pi(\Phi \mid \eta, \beta)} |J(\xi_j, \tau_j, \tau_{new})| \\ &= \frac{\ell(X \mid \mu, \rho_0, \sigma, \rho_1, \Phi_{new})}{\ell(X \mid \mu, \rho_0, \sigma, \rho_1, \Phi)} e^{-\beta^{-1}(\xi_{new} - \xi_j) - \lambda_1^{-1}(\tau_{new} - \tau_j)}. \end{aligned}$$

For more details on MCMC algorithms based on deterministic proposals we refer to Green (1995) and Section 2 in Tierney (1998). An introduction to these class of algorithms can also be found in Chapter 11 in Robert and Casella (2004).

Blocked multiplicative random-walk MH for jump sizes

To improve mixing, at each iteration of the algorithm the jump sizes of Φ are updated conditionally on the current jump times, using a blocked multiplicative random-walk Metropolis-Hastings algorithm. This updating scheme was first suggested by Roberts et al. (2004a) and has been also used by Frühwirth-Schnatter and Sögner (2009). Specifically, for each jump size ξ_j we propose a new jump size ξ'_j such that $\xi'_j = \xi_j \phi_j$, where the ϕ_j are i.i.d. with distribution $\text{LN}(0, c^2)$. The variance c^2 is chosen inversely proportional to the current number of jumps. Denoting by Φ' the proposed Poisson point process, the acceptance probability for this move can be easily shown to be $\min \{1, r(\Phi, \Phi')\}$, where

$$\begin{aligned} r(\Phi, \Phi') &= \frac{\pi(\Phi' \mid \mu, \rho_1, \sigma, \lambda_2, \eta, \beta, X)}{\pi(\Phi \mid \mu, \rho_1, \sigma, \lambda_2, \eta, \beta, X)} \prod_{i=1}^k \frac{\xi'_i}{\xi_i} \\ &= \frac{\ell(X \mid \mu, \rho_1, \sigma, \lambda_2, \eta, \beta, \Phi')}{\ell(X \mid \mu, \rho_1, \sigma, \lambda_2, \eta, \beta, \Phi)} \exp \left(-(\beta^{-1} - 1) \sum_{i=1}^k (\xi'_i - \xi_i) \right) \prod_{i=1}^k \frac{\xi'_i}{\xi_i}. \end{aligned}$$

Note that the product $\prod_{j=1}^k \frac{\xi'_j}{\xi_j}$ is equivalent to $\prod_{j=1}^k \phi_j$.

In the simulations reported in this thesis we used $c^2 = 2 \log(1 - 4 \log(a)/\Phi(S))$, with $a = 0.8$, as suggested by Frühwirth-Schnatter and Sögner (2009). Experiments with other variances, $c^2 = 1/(1 + \Phi(S))^b$, with $b = 0.5, 1, 2$, lead to similar outcomes,

which suggests that there is a great deal of robustness in the choice of the relation between c^2 and the number of jumps $\Phi(S)$.

5.6 Bayesian inference for a sum of three OU processes

Next we turn our attention to a mixture of a Gaussian OU process and two Gamma OU processes

$$X(t) = Y_0(t) + Y_1(t) + Y_2(t), \quad (5.16)$$

where

$$dY_0(t) = \lambda_0^{-1}(\mu - Y_0(t))dt + \sigma dW(t), \quad (5.17)$$

with $W(t)$ a standard Wiener process, and for each $i = 1, 2$,

$$dY_i(t) = -\lambda_i^{-1}Y_i(t)dt + dL_i(t), \quad Y_i(0) = 0, \quad (5.18)$$

where $L_i(t)$ is a time-homogeneous compound Poisson processes (CPP) with intensity rate $\eta_i > 0$ and Exponentially distributed jump sizes with mean β_i . We assume that $W(t)$, $L_1(t)$ and $L_2(t)$ are independent of each other, and for identifiability we suppose that $\lambda_1 > \lambda_2$, so that the jumps of $Y_1(t)$ have a slower decay rate than those of $Y_2(t)$. When modelling negative jumps with a 3-OU model, we do not impose any restrictions on the parameters, since in this case the identifiability problem vanishes.

We are now interested in estimating the parameters $\mu, \lambda_0, \sigma, \lambda_1, \eta_1, \lambda_2, \eta_2, \beta_1$ and β_2 . To do this, we shall employ a missing data methodology largely based on the case of 2-OU model developed in the previous sections. In the present case the latent variable is given by two marked Poisson processes corresponding to $L_1(t)$ and $L_2(t)$. The augmented likelihood $\ell(X \mid \mu, \lambda_1, \sigma, Y_2, Y_3)$ is given by Eq. (5.11) with $z_j = x_j - Y_1(t_j) - Y_2(t_j)$. Denoting by Φ_1 and Φ_2 the marked Poisson process corresponding to $L_1(t)$ and $L_2(t)$, the likelihood of $\Phi = (\Phi_1, \Phi_2)$ with respect to the dominating measure generated by the marked Poisson process with intensity 1 and jumps $\text{Ex}(1)$ is given by

$$\pi(\Phi \mid \eta_1, \eta_2, \beta_1, \beta_2) = \pi(\Phi_1 \mid \eta_1, \beta_1)\pi(\Phi_2 \mid \eta_2, \beta_2),$$

where for $i = 1, 2$,

$$\pi(\Phi_i \mid \eta_i, \beta_i) = e^{-(\eta_i-1)T} \left(\frac{\eta_i}{\beta_i} \right)^{\Phi_i(S)} \exp \left\{ -(\beta_i^{-1} - 1) \sum_{j=1}^{\Phi_i(S)} \xi_{i,j} \right\}.$$

Here $\Phi_i(S)$ denotes the number of points of Φ_i in the set S and

$$\Phi_i = \{(\tau_{i,1}, \xi_{i,1}), \dots, (\tau_{i,K_i}, \xi_{i,K_i})\}.$$

As before, we parameterise the model in terms of $\rho_1 = e^{-1/\lambda_1}$ and $\rho_2 = e^{-1/\lambda_2}$ and assume that

$$\rho_1 \sim \text{Beta}(a_{\rho_1}, b_{\rho_1}), \quad \rho_2 \mid \rho_1 \sim \rho_1 \text{Beta}(a_{\rho_2}, b_{\rho_2}),$$

so that λ_2 is *a priori* dependent on λ_1 . This choice of prior distributions induces the ordering $\lambda_1 > \lambda_2$ (Frühwirth-Schnatter and Sögner, 2009). The prior distributions for the remaining parameters are the same as those given in Section 5.4.2 for the case of one jump component.

The parameters of the diffusion component, $\mu, \rho_0 = e^{-\lambda_0^{-1}}$ and σ^2 , will be updated in a similar way to that proposed in Section 5.5. The parameters η_1 and η_2 are independent of the observed data and other parameters given Φ_1 and Φ_2 , respectively, so their conditional distributions are the same as in the case of one jump OU process, with the obvious modifications. The same applies to the mean jump sizes β_1 and β_2 . Since the posterior conditional densities of ρ_1 and ρ_2 are non-standard, they are simulated using a multiplicative random-walk Metropolis-Hastings algorithm as in Section 5.4.2. The three updating steps of section 5.5 shall be used to update first the point process Φ_1 and then the process Φ_2 given the updated process Φ_1 . In our applications to real data, we also employ a restricted 3-OU model where we assume that $\beta = \beta_1 = \beta_2$. In this case we sample β from

$$\begin{aligned} \pi(\beta \mid \Phi_1, \Phi_2) &= \pi(\Phi_1 \mid \eta_1, \beta) \pi(\Phi_2 \mid \eta_2, \beta) \pi(\beta) \\ &\sim \text{IG} \left(a_\beta + \Phi_1(S) + \Phi_2(S), \sum_{j=1}^{\Phi_1(S)} \xi_{1,j} + \sum_{j=1}^{\Phi_2(S)} \xi_{2,j} + b_\beta \right). \end{aligned}$$

Obvious modifications follow to other updating steps. When imposing this restriction

on the mean jump size parameters, we will refer to the 3-OU as the restricted 3-OU model. If no restrictions are imposed (except that of the $\lambda_i, i = 1, 2, \dots$), we will call it the unrestricted 3-OU model.

5.7 Analysis of algorithms on simulated data

5.7.1 2-OU model

To assess the performance of the MCMC algorithm for the 2-OU model, we apply it to a collection of simulated datasets from the model in equation (5.5) with $n = 1$. It is of particular interest to explore the simulation efficiency of the algorithm for different values of the intensity rate η . This is because this parameter can affect significantly the mixing properties of the algorithm, owing to the existence dependence between η and the latent process Φ . The proposed exercise is then to keep fixed the parameters $\mu, \rho_0, \sigma, \rho_1, \beta$ and vary η in the set $\{0.05, 0.1, 0.2, 0.3\}$. Thus we consider datasets with low and high frequency of jumps, covering the range of jump rates typically reported in the literature (cf. Seifert and Uhrig-Homburg (2007), Meyer-Brandis and Tankov (2008), Benth et al. (2008, 2012)). The true values of the parameters used in this simulation study are summarised in Table 5.1. The situations presented in the table intend to cover ‘typical’ values for energy spot prices. Under normal market conditions, we are assuming a time to mean reversion equal to one week with a daily volatility of 10%. Also we suppose that price spikes have a half-life of about $2 \log 2 \approx 1.3$ days. We shall consider other values for the speed of mean reversion λ^{-1} when studying the 3-OU model. Finally, the different values of η imply that the frequency of price jumps goes from 18 to 108 jumps a year on average.

Our choice of prior distributions was discussed in detail in Section 5.4.2. In all our MCMC simulations, hyperparameters are chosen so that prior distributions are relatively flat around the true value of the parameters. To design a fair exercise, however, for each value of η its prior is chosen so that its prior mean is equal to its true value. Our specific choice of prior distributions and hyperparameters are provided in Table 5.1. Starting values for the parameters are $(\mu, \rho_0, \sigma, \rho_1, \beta) = (1, 5, 0.15, 2, 0.5)$ with η starting at 0.001. The initial state of the process Φ is the empty configuration, i.e. we assume that initially there are no jumps. A different initial state for Φ could be,

for example, one simulated from its prior distribution with $\eta = \eta_{true}$. When doing this, we noticed that in most simulations the MCMC algorithm spent the first iterations killing most Poisson points, after which it placed new points at ‘corrected’ locations. Thus we prefer not to impose any initial configuration on Φ to avoid having a possibly long burn-in period.

For each simulated dataset, we run the MCMC algorithm for two million iterations, remove the first half million as burn-in and thin the Markov chain retaining one every 100 iterations. Since the results depend on the sample paths used as input to the MCMC algorithm, we have repeated this exercise for 60 different realisations of the model. This is particularly useful when obtaining point estimates of the jump parameters, due to the possibility of having distinct realisations of the Poisson process. For the birth-and-death algorithm the probability of placing a new-born point is set to $p = 0.5$. In all cases we assumed that the simulated data was daily, i.e. $\Delta_j = 1, j = 1, \dots, N$, and set the length of each time series to $T = 1000$.

The algorithm was implemented by combining Matlab and C++ MEX code, and was run on a Intel Xeon E5 processor running at 2.5GHz (with Intel Turbo Boost technology giving a maximum clock speed of 3.0GHz). By employing a parallel computing architecture with 12 cores, the algorithm completed 60×4 runs in about 20 hours. The mean computation time for completing 1000 MCMC iterations on the same machine using only a single core was about 1.9 seconds.

Parameter	True value	Prior properties		
		Prior	Mean	SD
μ	1	$N(1, 20^2)$	1	20
ρ_0	$e^{-1/8} \approx 0.8825$	$U(0, 1)$	0.5	0.29
λ_0	8	- -	- -	- -
σ^2	0.01	$IG(1.5, 0.005)$	0.01	- -
ρ_1	$e^{-1/2} \approx 0.6065$	$U(0, 1)$	0.5	0.29
η	$\{0.05, 0.1, 0.2, 0.3\}$	$Ga(1, \eta_{true}^{-1})$	η_{true}	η_{true}
β	0.7	$IG(1, 1)$	- -	- -

Table 5.1: True value of parameters and the prior distributions used to asses the performance of the MCMC algorithm on the 2-OU model.

Point estimates of the parameters

For each set of parameters, Table 5.2 provides point estimates of all parameters alongside their standard deviations (SD) based on 60 realisations of the model. The mean of each parameter was obtained by computing the grand average of the posterior means resulting from fitting each of the 60 sample paths. As reported in the table, the algorithm has estimated with good accuracy all the parameters, giving a relative error of less than 12% in all cases. Also, we notice that the MCMC algorithm is robust to the different sets of parameters used in this study. The only discrepancies are seen for the estimates of β , which were overestimated on average in the four cases considered here. Note, however, that the estimates of β improve as η becomes larger. This is due to the increasing number of jumps on average on the simulated paths, which provide more and more information about the true jump size. This is corroborated by looking at the SD of the estimates of β , which are seen to decrease with the true jump rate η_{true} . Overall, we consider that the results are quite satisfactory given the complexity of the model, and thus we expect a good performance of the algorithm when applying it to real data.

Convergence of Markov chain for different values of η

Next we analyse the convergence of the Markov chain for the jump intensity rate η . Figure 5.2 shows the first million MCMC draws for η obtained by running the MCMC algorithm for a particular realisation of the 2-OU model. We recall that at each MCMC iteration the birth-and-death algorithm can place only one new-born point or kill one of the existing points. This of course implies that the algorithm will take longer to converge for high jump frequency datasets rather than for low jump frequency datasets. This is exactly what we observe in Figure 5.2. As shown, when $\eta = 0.05, 0.1$ the chain converges to the true value within the first 50000 iterations, while for $\eta = 0.2$ and $\eta = 0.3$ it takes about 100000 and 500000 iterations, respectively. In all cases it seems that the chains have reached stationarity within half million iterations. This result is particularly interesting in the case $\eta = 0.3$. Although η was started comparatively far from its prior mean and true value, its chain moved to the modal area of the posterior distribution rather quickly (in less than 10000 iterations), illustrating the robustness of the proposed algorithm to initial values.

True value		Posterior properties							
		$\eta = 0.05$		$\eta = 0.1$		$\eta = 0.2$		$\eta = 0.3$	
		Mean	SD	Mean	SD	Mean	SD	Mean	SD
μ	1	0.99989	0.02408	0.99593	0.02161	0.99542	0.02809	0.99876	0.02735
λ_0	8	8.04637	1.28373	8.08860	1.05722	7.75620	1.16360	8.10754	1.13850
σ^2	0.01	0.01012	0.00055	0.01006	0.00055	0.01001	0.00064	0.01014	0.00067
λ_1	2	2.01580	0.11290	2.00861	0.09297	2.01366	0.07173	2.00885	0.05731
η	--	0.04742	0.00992	0.09815	0.01448	0.19541	0.02214	0.29973	0.02274
β	0.7	0.79122	0.12770	0.74580	0.08512	0.73016	0.06973	0.71427	0.06399

Table 5.2: For different values of the true intensity rate η , the sample mean and standard deviation of the MCMC posterior means across 60 simulated paths. The posterior means were obtained by running the MCMC algorithm two million iterations, removing the first half-million as burn-in and thinning the chain one every 100.

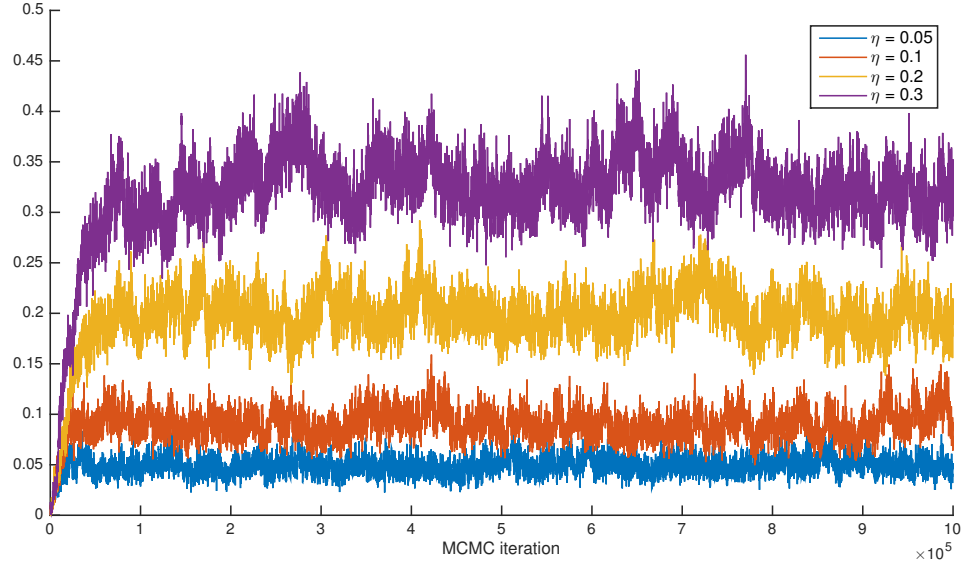


Figure 5.2: Trace plots of η corresponding to simulated datasets from the 2-OU model with different true values of η , $\eta_{true} = \{0.05, 0.1, 0.2, 0.3\}$.

Simulation efficiency

To assess the relative performance of the algorithm for different datasets, the ACF of the Markov chain is used as a measure of simulation efficiency. We notice that it would be impossible to derive explicit convergence rates in the present non-Gaussian context. Figure 5.3 presents the ACFs of the jump intensity rate (after removing the first 500000 MCMC iterations and thinning the chain one every hundred) for different datasets simulated with the parameters in Table 5.1. Inspection of the figure reveals that the shortest autocorrelation time is obtained for $\eta_{true} = 0.05$, while the ACF takes comparatively longer to decay as η_{true} increases. Nevertheless, in all cases the ACF for η dies away within the first 500 lags. The results show that simulation performance tends to decrease with the value of the true jump rate η_{true} , indicating that the algorithm has slow mixing for high jump frequency datasets. We note however that these long autocorrelation times are typical for mixture of OU process in missing data contexts (c.f Griffin and Steel (2006) and Frühwirth-Schnatter and Sögner (2009)).

We would like to mention that the ACFs for the other jump process parameters ρ_1 and β are similar to those obtained for η , whereas the ACFs for μ , ρ_0 and σ decay much faster and within the first 50 lags (not shown here). The better mixing properties of the diffusion parameters are due to the fact that the data is quite informative about

these parameters. Indeed, during periods when no jumps occur, or similarly, when the jump process is well identified, the 2-OU model reduces to a diffusion process with no missing data, so all the observed data corresponds to the Gaussian OU process Y_0 .

To improve mixing for high jump frequency data, this and the previous subsection suggest that the probability of birth moves p can be made dependent on the number of jump times as in Geyer and Møller (1994). For example, p can be relatively large over the first MCMC iterations and decrease up to some point over time. This has the effect of proposing and placing many jumps over the initial iterations and then balancing the birth and death of Poisson points. We experimented with different values of p using this idea, and obtained only slightly faster convergence rates than those presented in Figure 5.2. The greatest benefit of doing this was a rapid convergence of the chain of η for the high jump frequency dataset. However we note that tuning p is dataset-dependent, and thus pilot runs are often necessary for doing this. In all simulations we assume p is fixed and equal to 0.5.

An alternative is to employ multiple updates of Φ per update of the remaining model parameters. In our practical experience these multiple updates allow to explore the space faster and thus improve the convergence of the Markov chain. Although we do not employ this technique with the 2-OU model, it will be particularly useful when fitting the more complex 3-OU model to simulated and real data.




Figure 5.3: Autocorrelation function of the Markov chain of η for sample paths simulated with $\eta_{true} = \{0.05, 0.1, 0.2, 0.3\}$. The results were obtained by running the MCMC algorithm two million iterations, removing 500000 as burn-in and thinning the chain one every 100.

5.7.2 3-OU model

Next we apply our MCMC algorithm to simulated data generated with the 3-component mixture of OU process in (5.16). We shall focus only on the case of positive jumps, since similar results are obtained for the 3-OU model with one negative component.

Table 5.3 provides the true value of the parameters used for the simulation study of the 3-OU model, alongside our specific choice of prior distributions and hyperparameters. The situation presented in Table 5.3 corresponds to a relatively high volatile market with two different jump components: one with long-decaying and frequent medium-size jumps, and another one with large but rare relatively quick-decaying jumps. The initial state of the parameters is $(\mu, \lambda_0, \sigma, \lambda_1, \eta_1, \beta_1, \lambda_2, \eta_2, \beta_2) = (1, 5, 0.2, 5, 0.001, .5, 1, 0.001, 0.5)$, and we assume that initially there are no jumps on $\Phi = (\Phi_1, \Phi_2)$ and that $p = 0.5$, $\Delta_j = 1$ and $T = 1000$. For computing posterior properties we have run the algorithm for two million iterations, removed the first 500000 as burn-in and thinned the chain one every 100th iteration. The average time to complete 1000 MCMC iterations was about 3.6 seconds.

MCMC posterior results

The last two columns of Table 5.3 present the grand average of the posterior means and the SD of each of the model parameters across 60 sample paths on which we have apply the MCMC algorithm. As displayed in the table, all the parameters in the 3-OU model have been estimated with good accuracy, similar to the case of the 2-OU model. Further, the last column, which reports the SD of the posterior estimates, shows that relatively small variations are obtained in the estimates across all sample paths. We would expect to get even smaller values of SD with increasing length of the sample paths.

It is worth mentioning that the different types of price jumps in the sample paths are captured by the algorithm quite precisely. This is evidenced by the accurate estimates obtained for the parameters of Y_1 and Y_2 , which means that the algorithm is able to differentiate between different dynamics of jumps in the data, and different shock sizes. Clearly, this ability is particularly useful recalling our discussion of the key features of energy prices at the beginning of this chapter. A detailed analysis of the Markov chain for Φ_1 and Φ_2 shows that jump locations and jump sizes have been

Parameter	True value	Prior properties			Properties of posterior means	
		Prior	Mean	SD	Mean	SD
μ	1	$N(1, 20^2)$	1	20	0.9986	0.0446
e^{-1/λ_0}	0.8825	$U(0, 1)$	0.5	$1/\sqrt{12}$	0.8815	0.0148
λ_0	8	--	--	--	8.0604	1.0776
σ^2	0.0225	$IG(1.5, 0.005)$	0.01	--	0.0230	0.0016
e^{-1/λ_1}	0.7165	$U(0, 1)$	0.5	$1/\sqrt{12}$	0.7171	0.0336
λ_1	3	--	--	--	3.0584	0.4458
η_1	0.1	$Ga(1, 10)$	0.1	0.1	0.0964	0.0253
e^{-1/λ_2}	0.1353	$U(0, 1)$	0.5	$1/\sqrt{12}$	0.1522	0.0511
λ_2	0.5	--	--	--	0.5323	0.0981
η_2	0.05	$Ga(1, 50)$	0.05	0.05	0.0438	0.0147
β_1	0.5	$IG(1, 1)$	--	--	0.4585	0.0976
β_2	1	$IG(1, 1)$	--	--	1.0549	0.1480

Table 5.3: The true value and prior distribution of each parameter of the 3-OU model, alongside the properties of the posterior means (sample mean and standard deviation) across 60 samples paths. For each path the MCMC algorithm was run two million iterations, the first half-million were removed as burn-in and the resulting chain was thinned one every 100.

well estimated. This is in good agreement with the fact that the intensity rates and mean jump sizes of each marked Poisson process have been accurately estimated by the algorithm. We will present an analysis of the latent process when studying real data. See Section 5.8.

Figure 5.5 provides the trace, the estimated posterior distribution and the empirical ACF for each parameter of the 3-OU model. The ACFs, in particular, show that the algorithm is mixing relatively well. Nonetheless, we notice that the autocorrelation times for the 3-OU model tend to be larger than those obtained for the 2-OU model when comparing similar model parameters. This is obvious in the case of the intensity rates η_i and the mean reversion parameters λ_i . Owing to the complexity of the model, we consider that the results are quite satisfactory since in all cases the ACFs die away within the first 800 lags.

As said earlier, we can improve the mixing properties of the algorithm by updating Φ multiple times per update of the model parameters. This allows the algorithm to explore the space faster and better identify the missing data Φ . Accordingly, we have

modified the MCMC algorithm for the 3-OU model so that both Φ_1 and Φ_2 are updated five times per MCMC iteration using the MH algorithms of Section 5.5.1. The mean computation time for completing 1000 iterations is now 8.8 seconds on average, i.e. 2.4 times slower than with the single updating scheme for Φ . Nevertheless the benefits in term of the mixing properties of the algorithm are significant. To illustrate this point, Figure 5.4 provides the ACF of η_1 when using the original and the modified schemes for updating Φ . As shown, with the original algorithm the ACF takes about 700 lags to die away, whilst with the modified algorithm the ACF takes almost half the lags to do the same. Similar results are obtained for all parameters of the jump processes Y_1 and Y_2 .




Figure 5.4: The ACF of η_1 when fitting the 3-OU model to simulated data, using one (blue line) and five (orange line) update(s) of the latent process Φ per update of the remaining parameters.

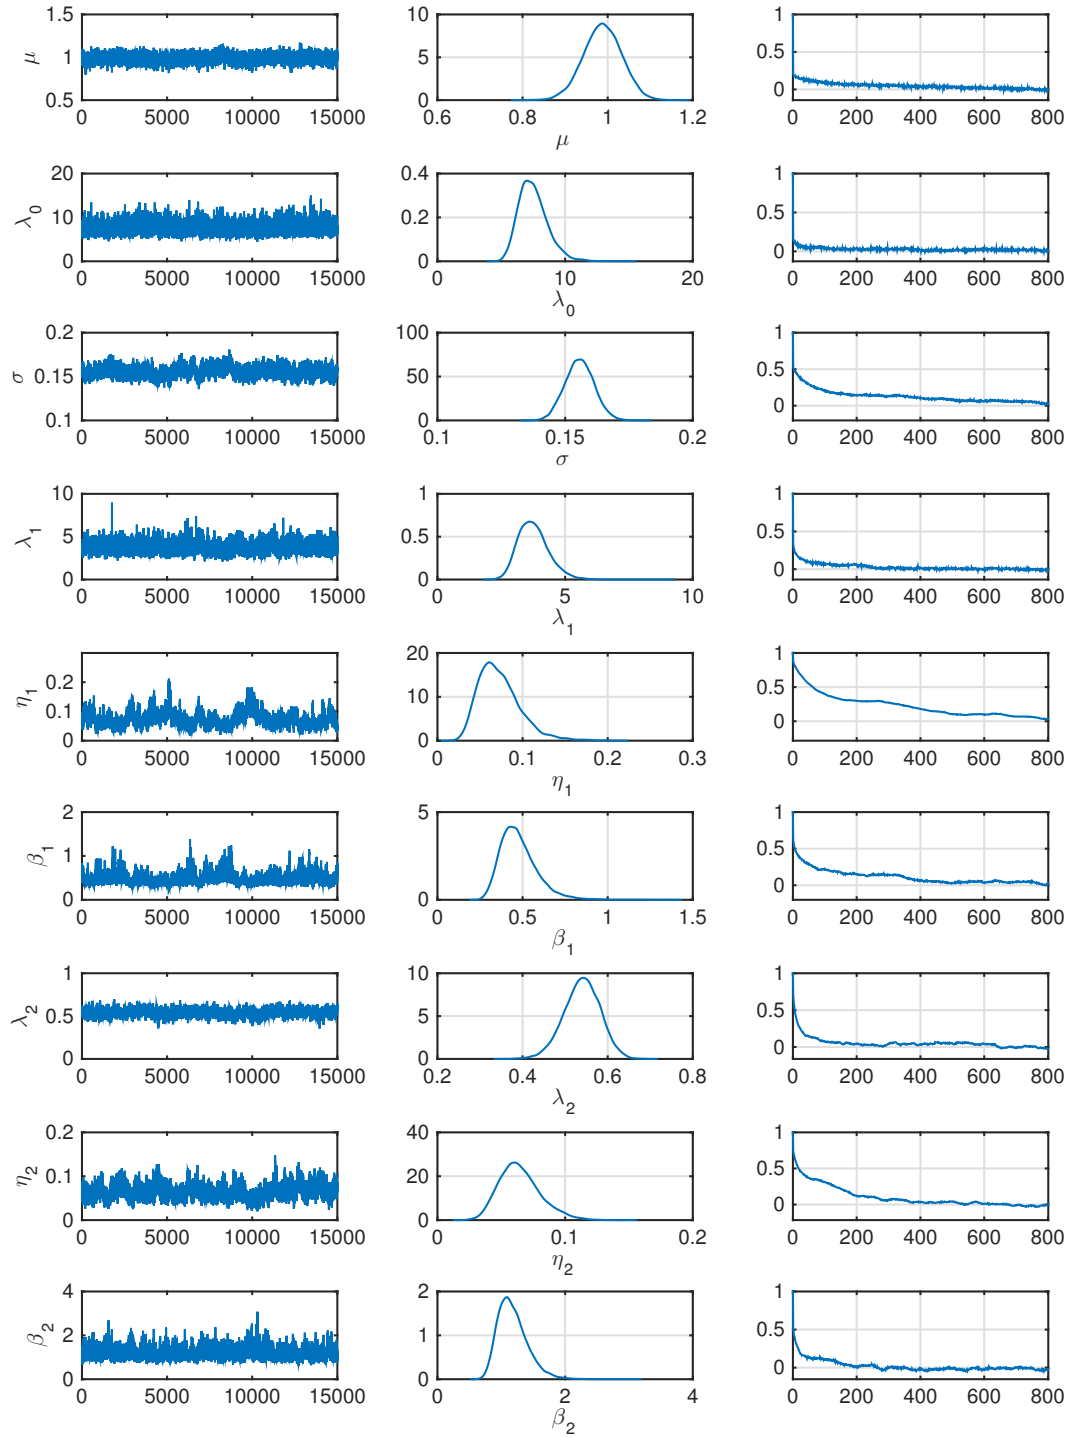


Figure 5.5: The figure shows the MCMC chain (left), the posterior distribution (middle) and the ACF (right) of the model parameters (top to bottom) when fitting the 3-OU model to simulated data. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100.

5.7.3 Diagnostics

Next, we propose a diagnostic tool to asses in a Bayesian framework the goodness-of-fit of the estimated mixture model. The basic idea relies on checking the theoretical Gaussian assumption on the process Y_0 using the MCMC output of the parameters and the latent process Φ .

In Section 5.2.1, we have assumed that Y_0 is a Gaussian OU process. Hence, its sample paths can be simulated exactly from

$$Y_0(t) = \mu + (Y_0(t + \Delta t) - \mu)e^{-\Delta t/\lambda_0} + \left(\frac{\sigma^2 \lambda_0}{2} (1 - e^{-2\Delta t/\lambda_0}) \right)^{1/2} \varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad (5.19)$$

See for example Gillespie (1996) or Glasserman (2004). Let us assume that we have available a path of $Y_0(t)$ at times $t = 0, 1, \dots, N$. Hence, thanks to equation (5.19), the noise process (ε_t) is uniquely determined by knowledge of μ, σ, λ_0 and $Y_0(t)$ at each $t = 0, 1, \dots, N$. By hypothesis, the noise data $\{\varepsilon_1, \dots, \varepsilon_N\}$ should have a Normal distribution, so we can conduct a statistical hypothesis testing procedure to check whether this assumption is well supported across the MCMC iterations.

To do this, at each iteration k of the MCMC algorithm we use the Markov chain of the parameters $\Theta^{(k)} = \{\mu^{(k)}, \lambda_i^{(k)}, \sigma^{(k)}, \eta_i^{(k)}, \beta_i^{(k)}\}$ and the missing data $\Phi^{(k)}$ to recover the path of each jump process $Y_i^{(k)}(t), i = 1, \dots, n$ (recall equation (5.12)). From this, we can then compute the path of $Y_0^{(k)}$ by taking

$$Y_0^{(k)}(t_i) = x_i - \sum_{i=1}^n Y_i^{(k)}(t_i), \quad (5.20)$$

where x_i is the observation of the process in (5.5) at time t_i . For simplicity, we have assumed that $w_i = 1, i = 1, \dots, n$. Consequently, for each MCMC iteration we have the noise data $\{\varepsilon_{t_i}^{(k)}\}$ on which we can test the Normality assumption. In this thesis, we use the one-sample Kolmogorov-Smirnov (KS) test of the null hypothesis that the cumulative distribution function (cdf) of the sample $\{\varepsilon_{t_i}^{(k)}\}$ is equal to the standard Normal cdf. The alternative is that the cdfs are not equal. By carrying out this test, we obtain a p -value $p^{(k)}$ for every MCMC iteration k .

In a Bayesian framework, these p -values should not be considered individually since there is uncertainty in the true value of the parameters. Nevertheless, following

Gelman et al. (1996) and Gelman (2003) we can get a Bayesian p -value, i.e. a posterior probability under our modelling assumptions, by averaging the $p^{(k)}$ -values over the posterior distribution of Θ . We notice that since we are integrating Θ out, the Bayesian p -value is a function of the observed data X only. This Bayesian p -value is interpreted in the classical sense: if it is smaller than the significance level of the test, then the null hypothesis is rejected in favour of the alternative. Our diagnostic tool of goodness-of-fit is based on a Bayesian p -value computed from a MCMC run on the superposition of OU processes. If the model fits the data well, we expect the Bayesian p -value to be higher than the significance level of the test, otherwise it will be smaller.

Notice that when each process $Y_i, i = 1, \dots, n$ (or equivalently the missing data Φ) has been well identified from the observed data, the resulting process Y_0 should have in principle a Gaussian distribution. So, although we are not directly checking the prior assumptions on the jump processes, the misspecification of Φ is readily seen when the noise process in Y_0 is tested for Normality.

Let us try the proposed diagnostic tool on the MCMC results for the 3-OU model of the previous section. Figure 5.6 provides the estimated posterior distribution of p -values from the KS test when applied to the residuals $\{\varepsilon_t\}$ obtained from fitting the 3-OU model. The results are based on 1000 samples, taken after a burn-in period. For each sample, the test was carried out at the 5% significance level. The results show that the posterior p -values have an approximate Uniform distribution on $(0, 1)$. Moreover, the Bayesian p -value, i.e. the sample mean of the distribution, is 0.54, indicating an excellent fit of the estimated model to the simulated data.

Next we check how the diagnostic tool performs with a misspecified model. To do this, we have carried out again the KS tests using the same MCMC results for the 3-OU model above. However, this time we use a different simulated path generated using the true values given in Table 5.3. In this case, the Bayesian p -value is 0.32×10^{-8} , indicating a rejection of the null hypothesis of the KS test. This result is explained by noting that, in this case, the estimated latent process Φ does not correspond to the jump processes of the new observed data. Therefore the processes $Y_0^{(k)}$ computed with (5.20) are not Gaussian, which is revealed by testing the Normality assumption of the residuals $\{\varepsilon_t^{(k)}\}$.




Figure 5.6: Posterior of p -values obtained when testing the estimated noise process of Y_0 for Normality using the Kolmogorov-Smirnov test at the 5% significance level. The results were obtained when fitting the 3-OU model to simulated data.

5.8 Application to real data

In the next three sections we make inferences on the mixture of Gaussian and non-Gaussian OU processes in (5.4) using real data corresponding to the APX Power UK spot base index (APXUK hereafter)¹, the European Energy Exchange Phelix day base index (EEX)² and the National Grid system average price (SAP)³. The APXUK and the EEX Phelix indices represent the daily average electricity spot price in the UK and Germany, respectively, whilst the SAP represents the daily average gas spot price in the UK. The APXUK and EEX datasets used in this chapter have been retrieved from Thompson Reuters Datastream, whereas the SAP dataset was retrieved from Bloomberg, all quoted in £/MWh.

We shall analyse 2-OU and 3-OU models all with positive jumps. In addition, the 3-OU model is studied with one positive and one negative jump process. This particular mixture model will be proved successful in capturing mixture of positive and negative price spikes in the EEX and SAP datasets.

For making inferences on the mixture in (5.4), we first fit a deterministic seasonal function to the given dataset, and then proceed to fit the deseasonalised data to the superposition process in (5.5) using our proposed MCMC algorithms. Our aim is

¹<https://www.apxgroup.com/market-results/apx-power-uk/ukpx-rpd-index-methodology/>

²<http://www.epexspot.com/en/>

³<http://marketinformation.natgrid.co.uk/gas/>

to study the potential that mixtures of OU processes have to explain the dynamics of different energy spot price time series. The parameter estimates of these models will be used in Chapter 6. There, we apply the theory and algorithms developed in the two previous chapters, alongside the models studied in this chapter, to the risk management of energy assets driven by energy spot prices based on mixtures of OU processes.

5.9 APXUK index for electricity spot prices

We start by fitting the superposition of OU processes in (5.4) to the APXUK time series covering the period from March 27, 2001 to November 21, 2006, weekends are not included. For this dataset we have in total $T = 1476$ observations (about 5.6 years assuming 260 trading days per calendar year). We take $\Delta_j = 1$.

5.9.1 Deterministic part

The first step is to remove seasonality from the data. For this we fit the seasonal trend

$$f(t) = a_1 + a_2 t + a_3 \sin(2\pi t/260) + a_4 \cos(2\pi t/260) + a_5 \sin(4\pi t/260) + a_6 \cos(4\pi t/260), \quad (5.21)$$

by solving the non-linear weighted least squares problem ⁴

$$\min_{a_i, i=1, \dots, 6} \sum_{t=1}^T (\log X_{obs}(t) - f(t))^2,$$

where $X_{obs}(t) = S(t)$ denotes the observed spot price at time t . The fitting procedure has been carried out using weighted least squares in order to reduce the influence of outliers and price spikes in the estimation results. Similar estimates are obtained by first removing any outliers from the data and then employing a standard least squares procedure. The estimates of the parameters are given in Table 5.4 and are all statistically significant at the 5% level. The coefficient of determination R^2 is 0.5910, indicating that the proportion of total variation of the data explained by the seasonal function $f(t)$ is relatively small.

⁴We have used the `nlinfit` built-in function in MATLAB.

a_1	a_2	a_3	a_4	a_5	a_6
2.5770	0.0008	-0.0817	0.0443	-0.0097	-0.0395

Table 5.4: Fit of the seasonal trend parameters in equation (5.21) to the APXUK time series.

Figure 5.7 shows the log-prices along with the estimated seasonal trend (top graph) and the deseasonalised time series (bottom graph), which in this case is given by

$$X(t) = \frac{X_{obs}(t)}{e^{f(t)}}.$$

We notice the strong seasonality present in the data, showing comparatively low average prices during the warm season (April-September) and high average prices in the cold season (October- March). The deseasonalised data shows a strong mean reversion to values around 1 (the sample price mean is about 1.045), and various price spikes smaller than 6 in the corresponding scale. We note that in this dataset price spikes tend to occur not only over the cold season but also during the warm season.

To have an idea of the short- and long-term dynamics of the APXUK dataset, we fit the theoretical ACF of our model (equation (5.15)) to the deseasonalised APXUK time series. Table 5.5 shows the weights ω_i and the estimates of λ_i in equation (5.15), alongside the sum of squared residuals (SS) resulting from the least squares routine, when fitting the sum of exponential functions in (5.15) with $n = 0, 1, 2$ to the deseasonalised APXUK series. It is evident from the results that the data exhibits short- and long-term dynamics, which becomes more and more clear as we increase the number of OU components. Note also that the SS decreases by fitting three exponential functions rather than one, although the cases $n = 1$ and $n = 2$ provide similar fit to the data. The short-term dynamics range between 12 hours and one week, whilst the long-term dynamics are of a few months. Thus we expect a similar situation when fitting 2-OU and 3-OU models to the APXUK dataset.




Figure 5.7: The daily average APXUK electricity spot log-prices along with fitted seasonal trend (top graph), and the resulting deseasonalised time series (bottom graph).

$n + 1$	ω_i			λ_i			SS
1	1	-	-	17.2498	-	-	3.2792
2	0.6826	0.3174	-	2.9055	117.983	-	0.3081
3	0.2855	0.4205	0.2940	0.5272	5.6868	131.4472	0.2767

Table 5.5: Fit of the autocorrelation function of the APXUK time series using superpositions of OU processes. $n + 1$ denotes the number of OU components, the ω_i are the estimated weights of the exponential functions, the λ_i are the estimated times to mean reversion and SS is the sum of squares resulting from the procedure.

5.9.2 Stochastic part, 2-OU model

Next we proceed to apply the MCMC algorithm for the 2-OU model to the deseasonalised APXUK time series. For this we run two million MCMC iterations, remove the first 500000 as burn-in and thin the chain taking one every 100 iterations. The parameters were started at $(\mu, \lambda_0, \sigma, \lambda_1, \eta, \beta) = (1, 5, 0.1, 2, 0.1, 0.5)$, and we assume again that the initial state of the process Φ was the empty configuration and that $p = 0.5$. To improve the mixing properties of the algorithm, we update the latent process Φ five times per update of the model parameters. In fact, from now on we used this

modified algorithm for sampling Φ .

Table 5.6 provides the prior distributions used in the algorithm, along with the posterior mean and standard deviation of each model parameter. The estimated level of mean reversion is in good agreement with the mean price of the deseasonalised APXUK time series which is equal to 1.045. The estimated daily volatility is 9.8%, while the time to mean reversion in normal market conditions is around 12 trading days. Interestingly, the results show a relatively high jump frequency of around 62 jumps per year on average with a quick decay rate of about 13 hours. Note that the MCMC estimates of λ_i resemble the short- and long-term dynamics observed when fitting the theoretical ACF of the model. We notice that the posterior SD is much smaller than the prior SD, from which we can conclude that the data is quite informative about the model parameters.

Parameter	Prior properties			Posterior properties	
	Prior	Mean	SD	Mean	SD
μ	N(1, 20 ²)	1	20	0.9592	0.0308
e^{-1/λ_0}	U(0, 1)	0.5	0.2887	0.9170	0.0116
λ_0	- -	- -	- -	11.7936	1.8298
σ^2	IG(1.5, 0.005)	0.01	- -	0.0096	0.0008
e^{-1/λ_1}	U(0, 1)	0.5	0.2887	0.1570	0.0189
λ_1	- -	- -	- -	0.5403	0.0352
η	Ga(1, 10)	0.1	0.1	0.2499	0.0297
$1/\beta$	Ga(1, 1)	1	1	1.4116	0.1440

Table 5.6: Prior distributions and posterior properties obtained for the APXUK dataset when fitting the 2-OU model. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the chain one every 100. The latent process Φ was updated five times per update of the model parameters.

Figure 5.8 provides the trace, the estimated posterior distribution and the empirical ACF for each model parameter. A close inspection of the graphs shows that the Markov chain has reached stationarity and that the MCMC algorithm has a good mixing behaviour. Indeed, the ACFs of μ , λ_0 , λ_1 and β have a rapid decay rate (notice in particular the very fast decay of the ACF of μ and λ_0), whereas the ACFs of σ^2 and η take longer to die away. At this point we would like to mention that for our model there exists a relatively high statistical correlation between σ^2 and η . This is because

both parameters help to explain the variations in the data, σ^2 for the Gaussian-OU component and η for the Gamma-OU component (the parameters λ_1 and β also explain the variance of the latter component, although their influence appears mainly due to the average number of jumps η). Hence one can have situations where a high volatile time series can be modelled with a high σ and a low frequency of large jumps, or small σ and high frequency of small jumps (see for example Seifert and Uhrig-Homburg (2007)). The interdependence between these parameters can cause acute problems in the Gibbs sampler when the jump component is not well identified. In our case, the multiple updates of Φ help the algorithm to explore the space efficiently and thus improve the autocorrelation times of σ^2 and η .

5.9.3 Stochastic part, 3-OU model

We have also fitted a 3-OU model with positive jumps to the deseasonalised APXUK time series. Although the results in the previous subsection seem promising, exploratory analysis of the APXUK dataset at hand shows differences between prices spikes relating their decay and average size properties. This motivates to employ a second jump component where we imposed $\lambda_1 > \lambda_2$, so that the OU process Y_2 has jumps decaying faster than those of Y_1 . We explore two situations, namely, one where the mean jump sizes β_i are allowed to have different values and another one where they are assumed to be equal. The results in this section have been obtained by running two million MCMC iterations, removing a burn-in of 500000 and retaining every 100th iteration. We have use the prior distributions given in Table 5.3, with the exception that η_2 has a $\text{Ga}(1, 10)$ prior distribution.

Figure 5.9 presents the trace, the estimated posterior distribution and the ACF of each model parameter for the unrestricted 3-OU model. Similar results are obtained for the case when $\beta_1 = \beta_2$ so we do not report them here. Having a look at the trace of each parameter we can conclude that the Markov chain have reached stationarity within the specified burn-in period. Furthermore, inspection of the ACFs reveal that the chain has relatively good mixing properties, all dying away within the first 1000 lags. These results are quite promising taking into account the complexity of the 9-parameter model studied here, and the reports of Roberts et al. (2004a), Griffin and

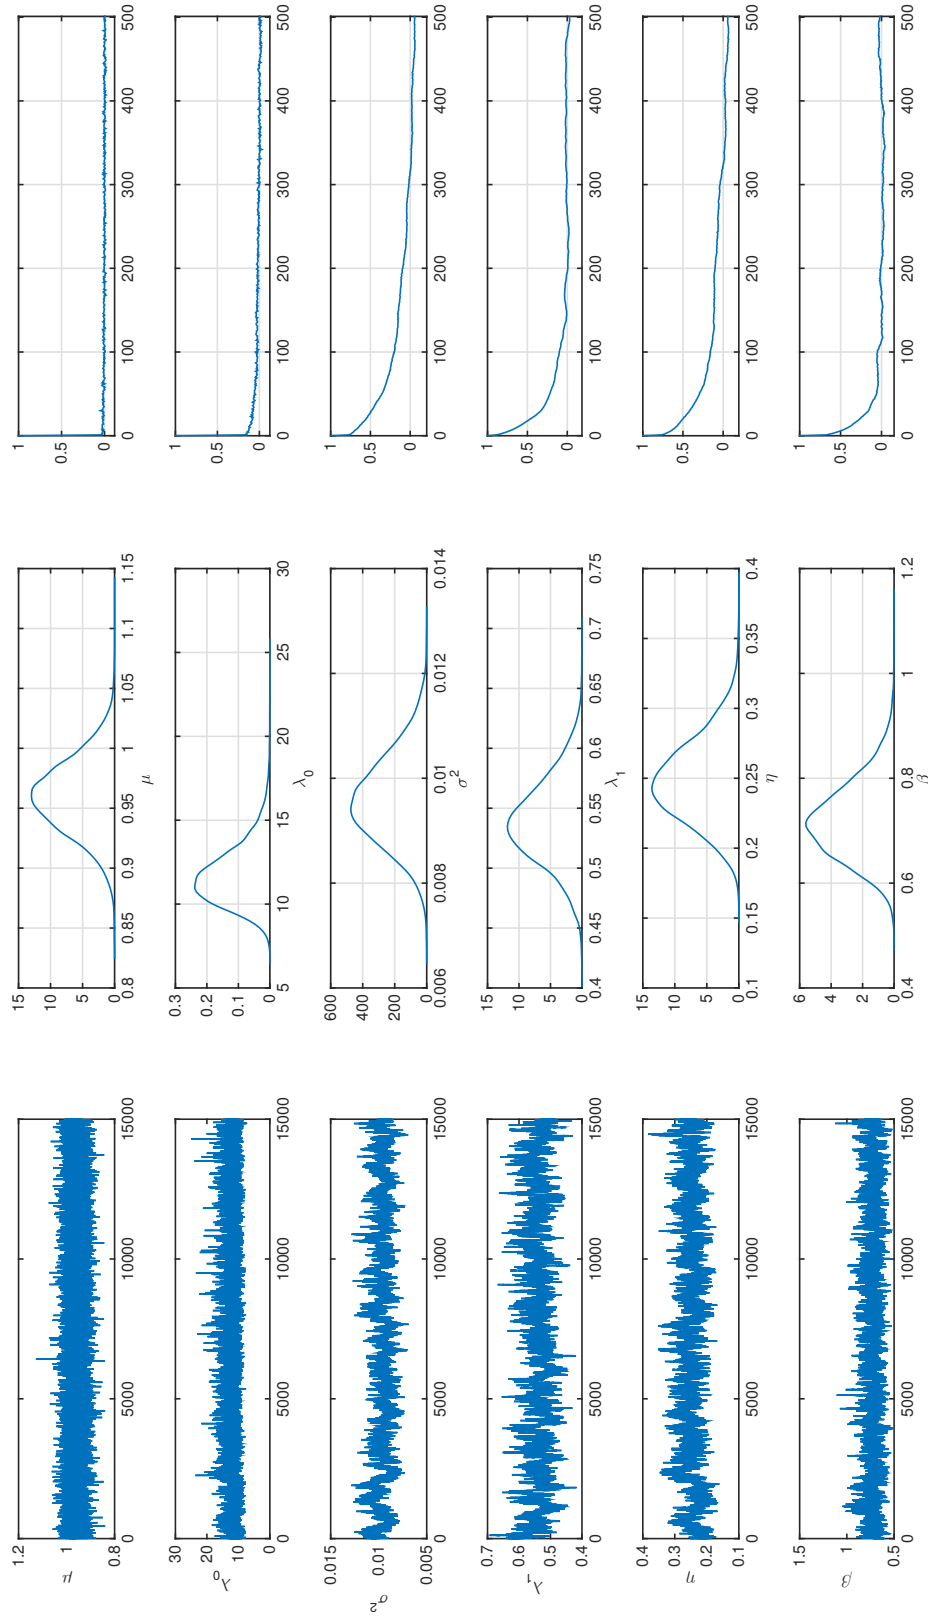


Figure 5.8: The figure shows the MCMC chain (left), the posterior distribution (middle) and the ACF (right) of the model parameters (top to bottom) when fitting the 2-OU model to the APXUK dataset. The results were obtained by running two million MCMC iterations, removing a burn-in period of 500000 and thinning the resulting chain one every 100. The latent process Φ was updated five times per update of the model parameters.

Steel (2006) and Frühwirth-Schnatter and Sögner (2009). In these articles, the proposed MCMC algorithm tends to mix comparatively slowly when applied to real data and long autocorrelation times are obtained for η , specially when its value is relatively large. Additionally, our simple multiple-updating scheme for Φ is competitive with the algorithms found in Frühwirth-Schnatter and Sögner (2009), which also require careful tuning of acceptance rates.

Table 5.7 provides the posterior properties of all the parameters for the 3-OU model when $\beta_1 \neq \beta_2$ and when $\beta_1 = \beta_2$. We notice that the estimates of μ and λ_0 are very similar to those estimated for the 2-OU, c.f. Table 5.6. However, there are noticeable differences in the estimated volatility level and the properties of jumps. Our results show that the volatility level of Y_0 is the smallest when $\beta_1 \neq \beta_2$, and it is followed in size when $\beta_1 = \beta_2$ and then when fitting the simple 2-OU model of the previous subsection. Thus the flexibility given to the 3-OU model by specifying different mean jump sizes seems to better capture the large daily variations in price level of the dataset at hand.

We notice that a fast decaying jump component Y_2 with a decay rate similar to that found for the 2-OU model has been estimated. This fast decay rate ranges between 9 and 12 hours. However, the jump frequency and jump size of the fast decaying component varies depending on the model. The largest frequency of fast-decaying jumps is obtained for the 2-OU model, followed by the 3-OU model when $\beta_1 = \beta_2$ and lastly by the case $\beta_1 \neq \beta_2$, while the mean jump size is largest for the unrestricted 3-OU model, followed by the 2-OU model and finally the restricted 3-OU model.

The most interesting result is the appearance of a clear, quite different jump process specification Y_1 with a half life of about 2 days. When $\beta_1 \neq \beta_2$, the results indicate that this is high activity jump component with an average of 62 jumps per year but relatively small jump size, $\beta_1 = 0.22$. Notice that in this case Y_2 has a jump frequency of 42 jumps per year, a mean life of 7 hours and a relatively large jump size of 0.8. Thus the process Y_2 can be thought of describing the arrival of information causing short-term but large shocks on electricity supply and demand, whereas Y_0 and Y_1 capture the systematic variations of power prices, with daily energy trading activities having mild and relatively strong effects on the settlement process of prices.

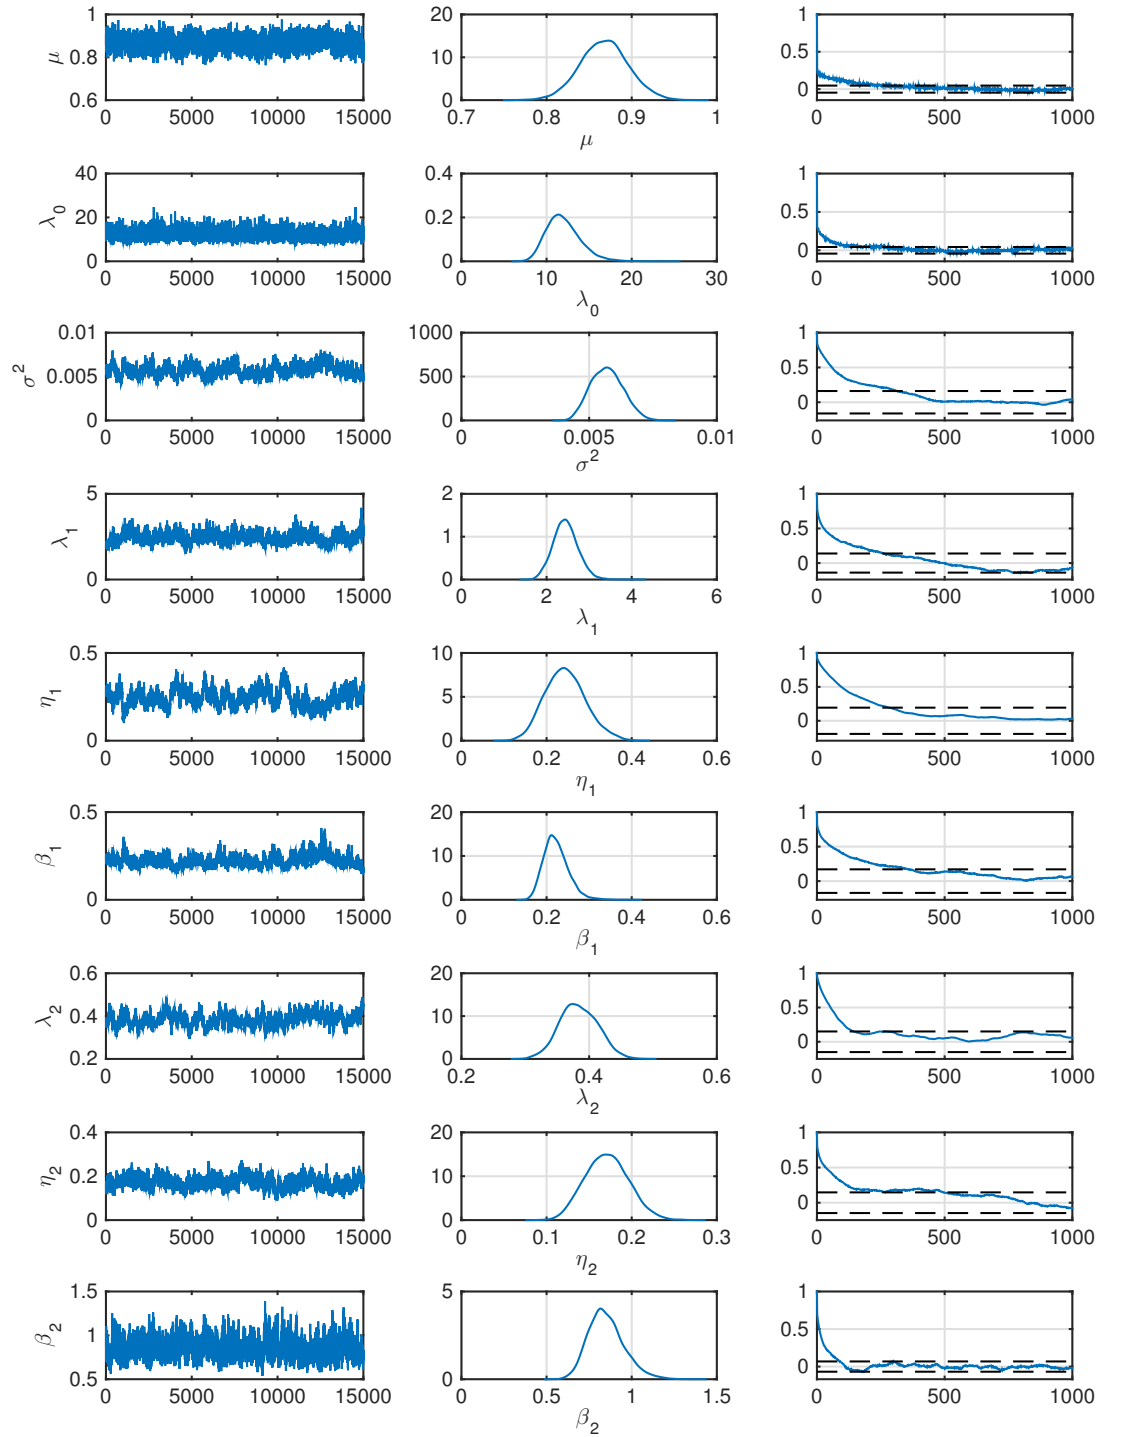


Figure 5.9: The figure shows the MCMC chain (left), the posterior distribution (middle) and the ACF (right) of the model parameters (top to bottom) when fitting the 3-OU model to the APXUK dataset. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100.

			Posterior properties			
2-OU model			3-OU model			
Parameter	Mean	SD	$\beta_1 \neq \beta_2$		$\beta_1 = \beta_2$	
			Mean	SD	Mean	SD
μ	0.9592	0.0308	0.8693	0.0284	0.9168	0.0253
e^{-1/λ_0}	0.9170	0.0116	0.9176	0.0128	0.9044	0.0139
λ_0	11.7936	1.8298	11.9352	1.9892	10.1819	1.6152
σ^2	0.0096	0.0008	0.0057	0.0006	0.0076	0.0007
e^{-1/λ_1}	0.1570	0.0189	0.6605	0.0342	0.6253	0.0610
λ_1	0.5403	0.0352	2.4408	0.3075	2.1960	0.4253
η_1	0.2499	0.0297	0.2412	0.0487	0.0777	0.0197
e^{-1/λ_2}	--	--	0.0742	0.0148	0.1239	0.0145
λ_2	--	--	0.3839	0.0297	0.4788	0.0270
η_2	--	--	0.1698	0.0250	0.2070	0.0313
β_1	0.7159	0.0738	0.2243	0.0291	0.5991	0.0612
β_2	--	--	0.8544	0.1047	0.5991	0.0612

Table 5.7: Summary of posterior properties for each parameter when fitting the 2- and 3-OU models to the APXUK series. The results were obtained by running two million MCMC iterations, removing 500000 as burn-in and thinning the resulting chain one every 100. The latent process Φ was updated five times per update of the remaining parameters.

When $\beta_1 = \beta_2 = 0.6$, Y_1 is a low activity jump process with 18 jumps per year on average and comparatively large jump sizes. In this case, Y_1 is modelling rare but long-decaying big shocks on power prices, whilst Y_2 is modelling fast-decaying jumps of small and medium size and Y_0 accounts for normal market conditions.

To get a picture of how the 2- and 3-OU models capture price jumps, it is helpful to plot estimated paths of Y_0 and the latent process Φ for both models. For a period of two hundred days, Figure 5.10 presents the deseasonalised APXUK time series superimposed to the estimated Y_0 on a particular MCMC iteration and the corresponding estimated latent process(es) for each model. We show only the case $\beta_1 \neq \beta_2$ and recall that Y_0 can be easily obtained by knowledge of the observed data and missing data (equation (5.10)). As shown, the MCMC algorithm for the 2-OU model tends to place many consecutive fast-decaying jumps in order to estimate what seems to be a slowly-decaying jump. This is noticeable for example at time $t = 460$ and in particular at time $t = 600$. In contrast, with the 3-OU model, the algorithm places a

single medium-size jump in the slowly-decaying jump component Y_1 at time $t = 600$, avoiding a ‘cluster’ of jumps altogether.

This analysis allows us to get some insight into the results for the 2- and 3-OU models of Table 5.6 and Table 5.7, respectively. Since the 2-OU has only one jump process, this jump component is used to explain large variations in the data which cannot be explained as tail realisations of the Gaussian OU process. The algorithm forms cluster of jumps when necessary (which may be theoretically invalid, owing to the Exponential distribution of jump arrivals) and the mean jump size results as an average of all large and small jumps. However, the 3-OU model allows more flexibility in differentiating between types of price spikes in the data which may also have different mean jump sizes as illustrated in Figure 5.10. When only one mean jump size parameter is used in the 3-OU model, the fast-decaying component has a strong effect due to the relatively large mean jump size induced by the combination of jumps with fast and slow decay rate.

5.9.4 Diagnostics

Next, we apply the diagnostic tool developed at the end of Section 5.7.2 to the MCMC results obtained when fitting a superposition of two and three OU processes to the APXUK dataset. Figure 5.11 presents the posterior p -values when testing the Normality assumption of the residuals of Y_0 . We show results for the 2-OU model (top panel) and the 3-OU models assuming $\beta_1 \neq \beta_2$ (middle panel) and $\beta_1 = \beta_2$ (bottom panel). We have used the KS test with a 5% significance level and 1000 posterior samples of the corresponding Markov chain.

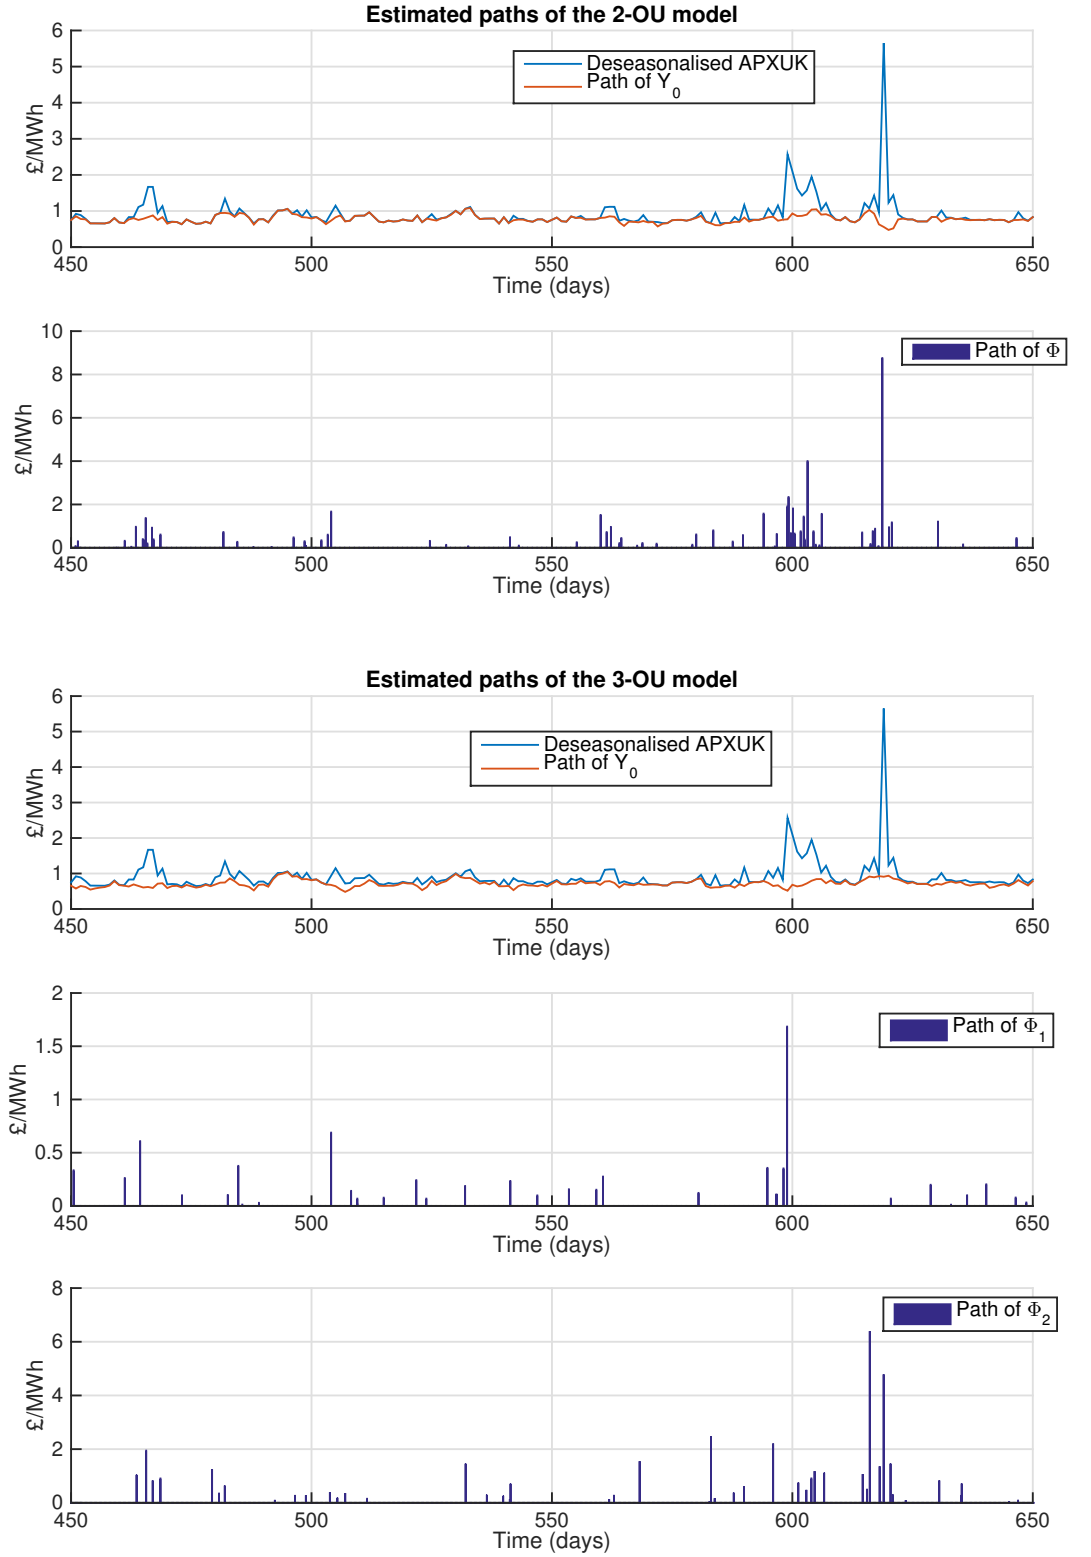


Figure 5.10: The deseasonalised APXUK time series superimposed to a MCMC estimate of the path of the Gaussian OU process, and the corresponding estimate of the latent process Φ for the 2-OU model and the 3-OU model.

As shown, the posterior distribution of p -values is skewed to the right in all cases, in contrast with the results for simulated data (recall Section 5.7.2). This illustrates the expected situation that when fitting statistical models to real data the goodness-of-fit of the model tends to decrease considerably. A remarkable result is that, while the mean p -value for the 2-component case is only slightly larger than 5%, the mean p -value for the restricted 3-OU model is 16% and much larger, 33%, for the unrestricted 3-OU model. Furthermore, the median of p -values have a similar behaviour increasing from 2.3% to 11% and 27.6% when increasing the flexibility of the model. Hence the results support at a greater extent the use of a 3-OU model with different jump sizes rather than 2- or a restricted 3-OU model for the APXUK dataset under study. We note nevertheless that all these models seem to be adequate for the data judging from the mean p -value alone, which in all cases is larger than the significance level specified for the KS test.

Ideally, the data generated by a proposed statistical model should ‘look’ like the observed data used to estimate it. This allows us to get an idea of whether a specified model describes the statistical features of the data under study. This observation motivates a second diagnostic tool for the mixture of OU processes, whereby summary statistics of simulated and observed data are compared against each other. In the following paragraphs, we carry out this analysis for the APXUK dataset when modelled with a superposition of two and three OU processes.




Figure 5.11: Posterior of p -values obtained when testing the estimated noise process of Y_0 for Normality using the Kolmogorov-Smirnov test at the 5% significance level.

To do this, we simulate 500 sample paths generated with each estimated model according to the point estimates given in Table 5.6 and Table 5.7, and compute summary statistics for each of these sample paths. Specifically, we employ the sample mean and SD of each path, and sample skewness, kurtosis and the first, second and third quartile ($Q_i, i = 1, 2, 3$) of the sample distribution of each sample path. Table 5.8 shows

the grand average across the 500 sample paths of these summary statistics and the corresponding standard deviation in parentheses. Having a look at the sample mean, SD and quartiles, we notice that all models produce sample paths with similar sample statistics to those of the observed data. Nevertheless, discrepancies are seen in the third and fourth moments of the distribution of simulated data. Indeed, all estimated models underestimate the sample skewness and kurtosis of the APXUK dataset. The model that better resembles the real data is the flexible 3-OU model where the mean jump sizes of Y_1 and Y_2 are unrestricted.

	Mean	SD	Skewness	Kurtosis	Q_1	Q_2	Q_3
APXUK dataset	1.0450	0.3716	3.8631	33.3966	0.8154	0.9492	1.1696
2-OU model	1.0559 (0.0308)	0.3522 (0.0205)	2.1673 (0.5383)	15.3691 (6.5621)	0.8438 (0.0326)	1.0175 (0.0314)	1.2067 (0.0336)
3-OU $\beta_1 \neq \beta_2$	1.0579 (0.0261)	0.3326 (0.0217)	2.4863 (0.8337)	17.6010 (10.6037)	0.8561 (0.0263)	1.0103 (0.0261)	1.1908 (0.0291)
3-OU $\beta_1 = \beta_2$	1.0748 (0.0276)	0.3594 (0.0234)	2.2001 (0.4596)	13.2047 (5.7388)	0.8581 (0.0262)	1.0162 (0.0259)	1.2035 (0.0307)

Table 5.8: Summary statistics of the APXUK time series and simulated data generated with the estimated 2- and 3-OU models. The results for simulated data are based on the grand average across 500 simulated sample paths. The standard deviation is provided in parenthesis and Q_i denotes the i -th quartile.

These discrepancies in skewness and kurtosis may be due to the extremely rare price spikes present in the APXUK data, see for example the large price spikes around time $t = 625$ and $t = 1300$ in Figure 5.7. By removing these two large price jumps from the dataset, its sample skewness and kurtosis decrease to 2.38 and 12.46, respectively, which confirms our expectations. Notice that the unrestricted 3-OU model produces the largest kurtosis amongst all models studied here, albeit we should take into account that the standard deviation of this figure is relatively high. This large sample kurtosis is explained by recalling that this model has the smallest volatility σ among the three models considered (see Table 5.7), which implies a greater occurrence of small variations and hence more kurtosis (peakedness) on the distribution of simulated data.

5.10 EEX Phelix index for electricity spot prices

We now turn our attention to the EEX Phelix time series covering the period from June 16, 2000 to November 11, 2006, weekends removed. Statistical procedures for fitting a sum of two OU processes on this dataset have been previously suggested e.g. by Meyer-Brandis and Tankov (2008), and thus this dataset provides an interesting test data to apply our MCMC algorithm. As before, we make inference in the spot price model of equation (5.4) by first removing seasonality from the data and then proceeding to fit a superposition of OU processes to the deseasonalised data. An interesting aspect of the EEX series is the presence of negative jump prices. Thus we will see that there is a great advantage on using a 3-OU model with both positive and negative jump components rather than a simple 2-OU model. Conceptually, this is similar to the results of the previous section, where adding a second jump component provided some benefits in terms of the statistical properties of the estimated OU components.

5.10.1 Deterministic part

To remove seasonality from the data, we use the same seasonal function (equation (5.21)) as in the case of the APXUK dataset. Table 5.9 provides our estimates of the parameters when fitting the EEX log-prices using a non-linear weighted least squares routine. All parameters are statistically significant at the 5% level, and the R^2 is 0.5750, only slightly larger than when fitting the APXUK data. Thus $f(t)$ accounts only for a small fraction of the total variability of the EEX time series.

The fitted seasonal function has been superimposed to the log-prices in the top panel of Figure 5.12, whilst the resulting deseasonalised series is shown in the bottom panel of the figure. Inspection of the figure reveals some interesting dynamics of the EEX series. Firstly, the seasonality effect is strong in the EEX series with price spikes tending to occur on both the cold season and the warm season. Secondly, we notice the presence of negative prices spikes in the deseasonalised data, which are however more noticeable on the logarithmic scale (top panel). Finally, the EEX series presents larger positive price jumps when compared to the APXUK series. Nonetheless, most price jumps are seen to decay quite quickly.

Table 5.10 shows the weights ω_i and the estimates of λ_i in equation (5.15), alongside

the sum of squared residuals (SS) resulting from the least squares routine, when fitting the sum of exponential functions in (5.15) with $n = 0, 1, 2$ to the deseasonalised EEX series. Qualitatively, the results are similar to those obtained for the APXUK series. Indeed, we observe that the EEX data exhibits a short- and long-term price dynamics. However, the long-term variations are considerably shorter than in APXUK data, ranging between half- and three weeks. Note also that the ACFs have been closely estimated, judging from the low SS values in Table 5.10.

a_1	a_2	a_3	a_4	a_5	a_6
2.4369	0.0007	-0.0005	-0.0797	0.0457	-0.0122

Table 5.9: Fit of the seasonal trend parameters in equation (5.21) to the EEX time series.




Figure 5.12: The daily average APXUK electricity spot log-prices along with a fitted seasonal trend (top graph), and the resulting deseasonalised time series (bottom graph).

n	ω_i			λ_i			SS
1	1	-	-	3.77	-	-	0.2433
2	0.1814	0.8186	-	2.295	23.32	-	0.0498
3	0.1850	0.6252	0.1898	0.02	2.52	22.05	0.0281

Table 5.10: Fit of the autocorrelation function of the EEX time series using superpositions of OU processes. n denotes the number of OU components, the ω_i are the estimated weights of the exponential functions, the λ_i are the estimated times to mean reversion and SS is the sum of squares resulting from the procedure.

5.10.2 Stochastic part, 2- and 3-OU models

In this section we fit a mixture of two and three OU processes to the deseasonalised EEX series. In the case of the 3-OU model, we employ a non-negative and a non-positive jump component with and without restrictions on the mean jump size parameters. Although a 3-component model seems to be necessary for the EEX dataset, fitting a 2-component model allows us to compare these different models. All the results in this section have been obtained by running two million MCMC iterations, removing the first 500000 as burn-in and retaining one every 100 iteration. Our choice of prior distributions is the same used for the APXUK dataset. In addition, for the 3-OU model given in (5.5) with $n = 3$, we have set $w_0 = w_1 = 1, w_2 = -1$, so that the process Y_2 accounts for negative price jumps.

Table 5.11 presents the posterior means and standard deviation of all parameters when fitting a 2-OU model and a 3-OU model (both restricted and unrestricted models) to the EEX series. The set of parameter estimates given in the table show some fundamental differences between the EEX and the APXUK datasets. In comparison with the APXUK series, the algorithm has estimated a larger volatility σ and a shorter time to mean reversion λ_0 (between four and eight days) for the EEX series. Additionally, the EEX Phelix index exhibits positive jumps with a comparatively high mean jump size but smaller jump frequency of about 30 jumps per year. Negative price jumps also have a similar intensity rate as positive jumps, albeit their mean jump sizes differ depending on the model. With the 2-OU model the mean size of positive jumps is about 1.11, while the restricted 3-OU has an estimated mean jump size of 0.77 (for both positive and negative jumps). In contrast, in the unrestricted 3-OU

model the mean size of positive jumps and negative jumps is 0.9 and 0.42, respectively, indicating large positive price spikes and relatively small negative price spikes. We notice however that the times to mean reversion $\lambda_i, i = 1, 2$, are almost the same for both positive spikes and negative spikes, of about 12 hours, a situation which is clearly different to the case of the APXUK dataset.

Parameter	Posterior properties					
	2-OU model			3-OU model		
	Mean	SD	$\beta_1 \neq \beta_2$		$\beta_1 = \beta_2$	
			Mean	SD	Mean	SD
μ	0.9979	0.0182	1.0163	0.0223	1.0037	0.0210
e^{-1/λ_0}	0.7954	0.0168	0.8791	0.0150	0.8695	0.0149
λ_0	4.4034	0.4151	7.8914	1.0802	7.2518	0.9165
σ^2	0.0273	0.0015	0.0127	0.0013	0.0133	0.0012
e^{-1/λ_1}	0.1006	0.0297	0.1907	0.0397	0.2419	0.0288
λ_1	0.4336	0.0589	0.6049	0.0762	0.7061	0.0596
η_1	0.1045	0.0205	0.1227	0.0187	0.1276	0.0191
e^{-1/λ_2}	--	--	0.2275	0.0377	0.1408	0.0331
λ_2	--	--	0.6777	0.0769	0.5102	0.0620
η_2	--	--	0.1340	0.0314	0.0912	0.0175
β_1	1.1207	0.1545	0.9155	0.1216	0.7730	0.0862
β_2	--	--	0.4343	0.0830	0.7729	0.0858

Table 5.11: Summary of posterior properties for each parameters when fitting the 2- and 3-OU models to the EEX series. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100. The latent process Φ was updated 5 times per update of the remaining parameters. The process Y_2 is used to model negative jumps.

Our MCMC results for the 2-OU model can be compared with those in Meyer-Brandis and Tankov (2008). Their estimates of μ, λ_0 and σ are 1.025, 6.2 and 0.148, respectively, which are quite close to our MCMC estimates. Using the ACF as we did above, they indirectly find an estimate of λ_1 equal to 2.9 (recall Table 5.10), which clearly is quite different to our MCMC estimate of 0.43. We note, however, that these estimates for λ_1 need not be equal, since they have been calculated with two different approaches. The first one was calculated from the empirical ACF of the data, whilst the MCMC estimate was obtained directly from the data. An interesting result is that Meyer-Brandis and Tankov needed to filter less jumps than us to get similar volatility

σ for the diffusion. Indeed, assuming constant jump intensity rate their estimate of η_1 is 0.037, around 10 jumps per year. Our MCMC estimate of η_1 is 0.1045, implying 26 jumps per year on average. Also, they estimate a mean jump size equal to 1.47, which is larger than our corresponding estimate of 1.02. The discrepancies in the jump process parameters are due to the fact that the MCMC algorithm does not focus only on the relative large price spikes, but also on those with a relatively small size. This implies obtaining many more jumps on average (larger η_1) and a possibly smaller mean jump size β_1 .

The mixing properties of the algorithm for the EEX dataset are very good and comparable to those obtained for the APXUK dataset. To illustrate the results Figure 5.13 presents the trace, the estimated posterior distribution and the ACFs corresponding to each parameter of the restricted 3-OU model. In all cases the ACFs decay rapidly within the first 500 lags, indicating relatively good mixing properties of the MCMC algorithm for the 3-OU model.

Figure 5.14 shows the deseasonalised EEX series and the estimated processes Y_0 , Y_1 and Y_2 for a particular iteration of the MCMC algorithm. We present results for the 2-OU model and for the unrestricted 3-OU model. As shown, the 2-OU model captures quite precisely the positive price spikes in the period considered in the figure. Negative price jumps, however, are accommodated using the diffusion process, which may result in a poor performance of the 2-OU model when using our diagnostics tools for Y_0 . In contrast, the 3-OU model captures large positive and negative price variations in the EEX series, and overall it seems to explain the dynamics of the data quite well.

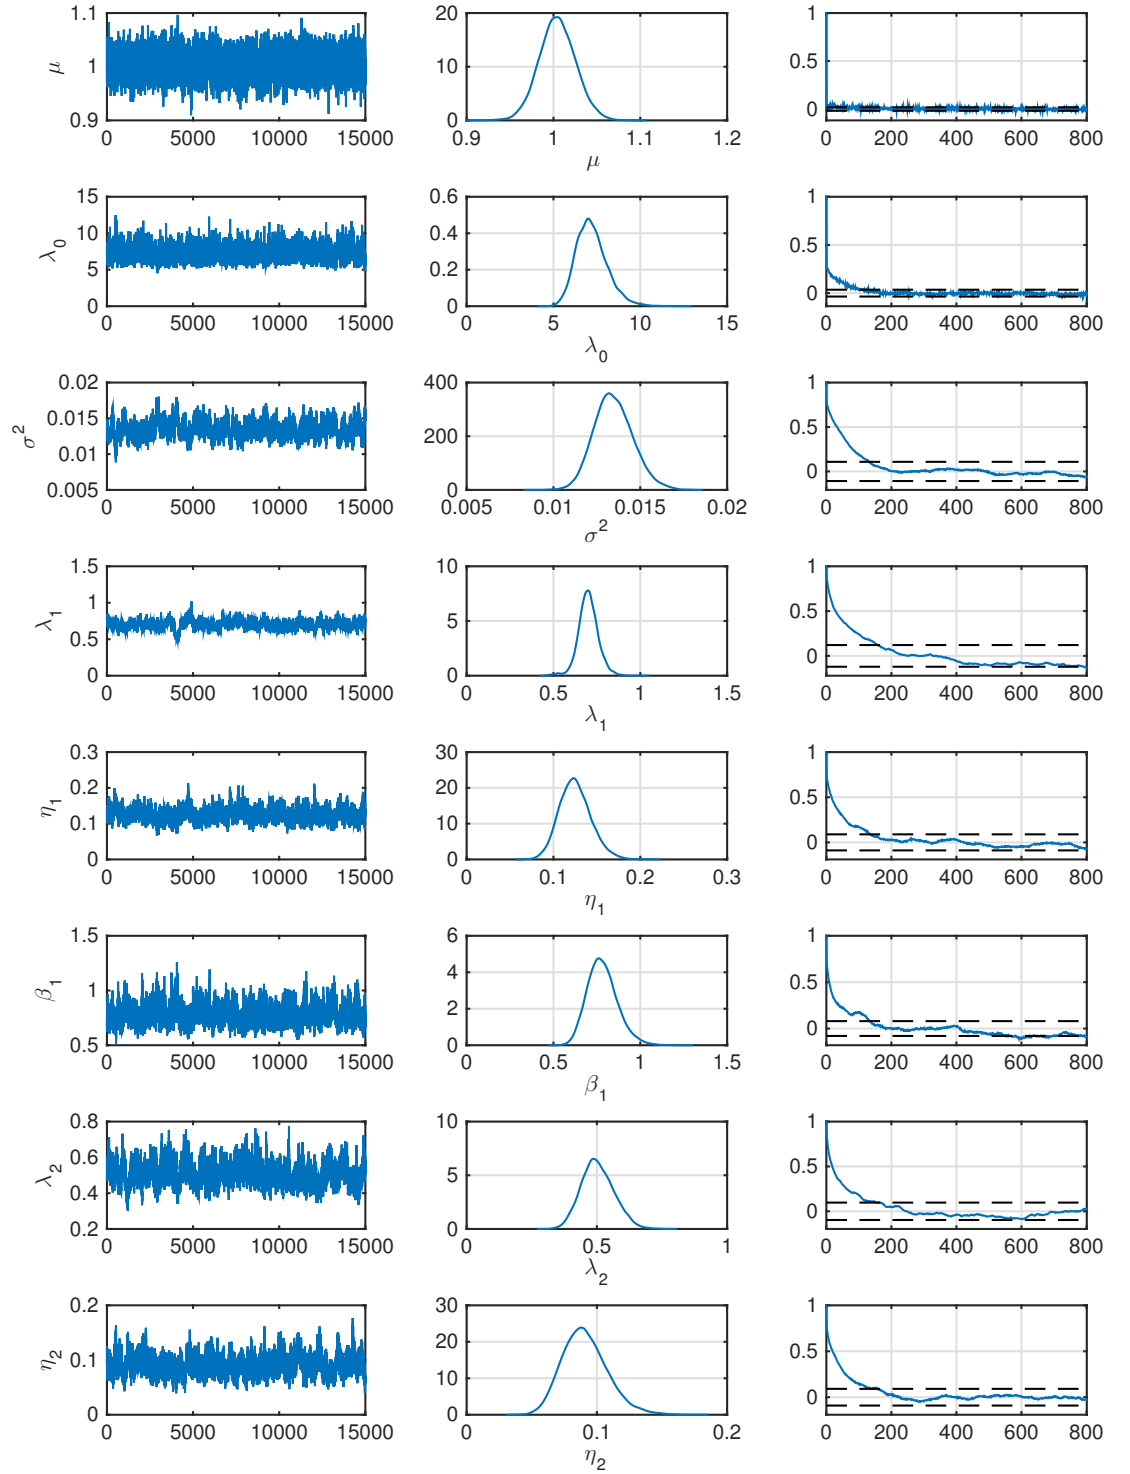


Figure 5.13: The figure shows the MCMC chain (left), the posterior distribution (middle) and the ACF (right) of the model parameters (top to bottom) when fitting the 3-OU model to the EEX dataset. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the chain one every 100.

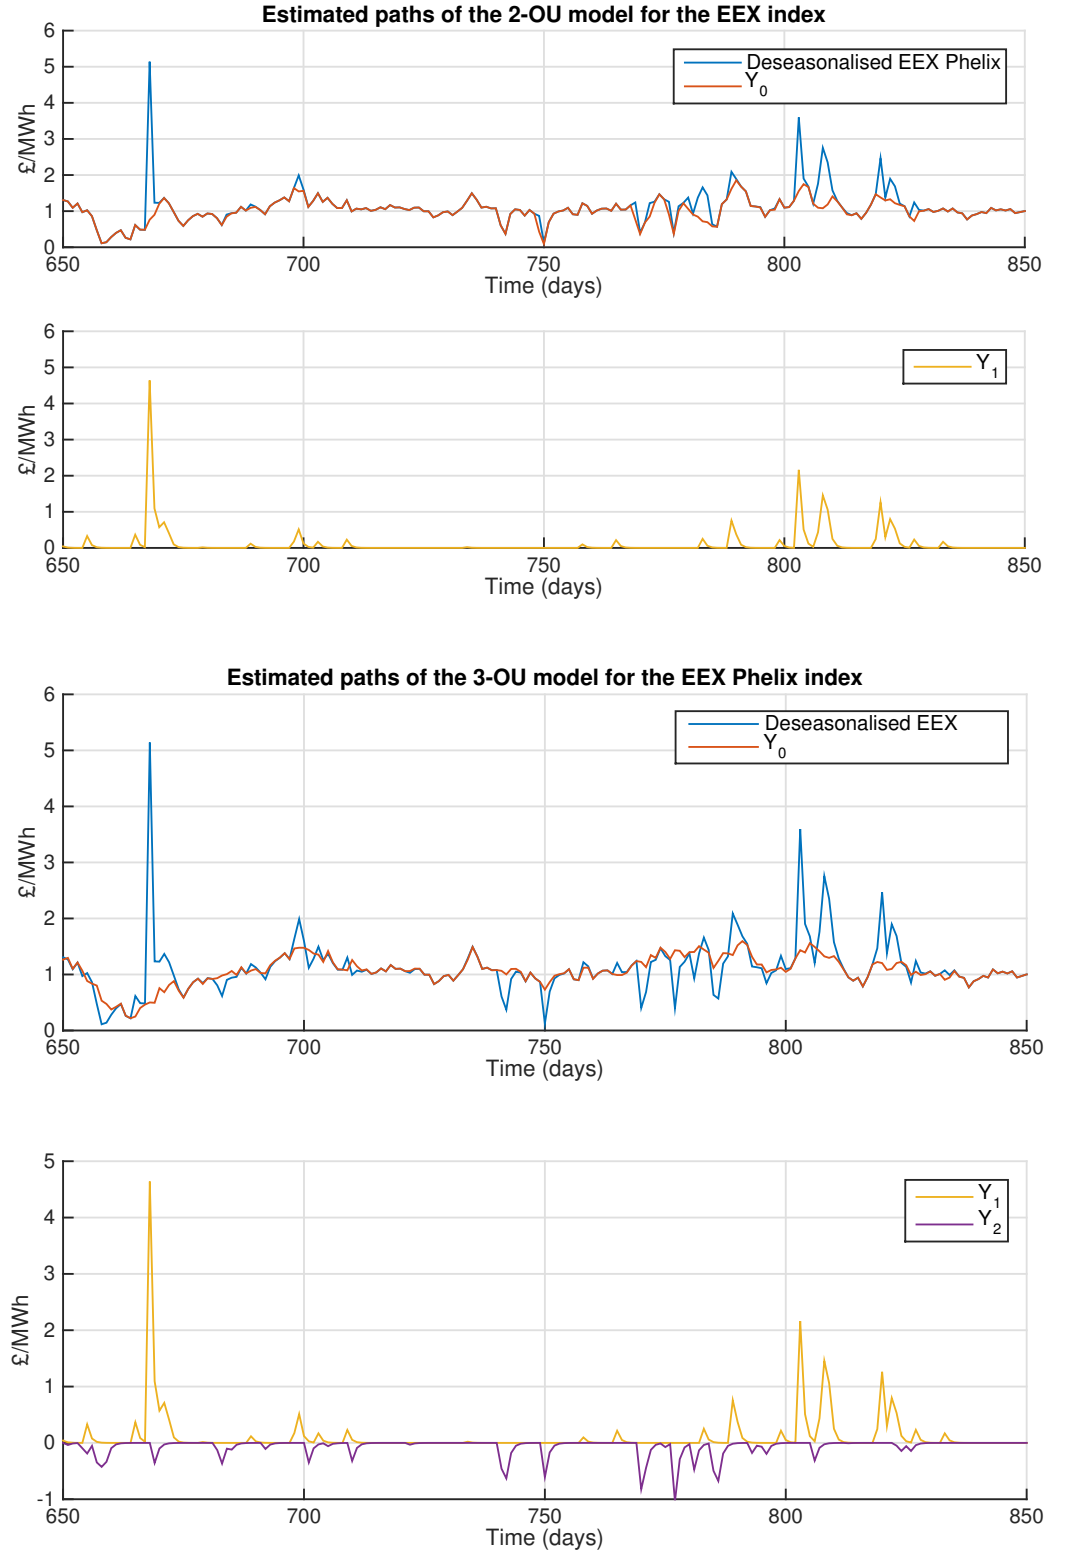


Figure 5.14: The deseasonalised EEX time series superimposed to a MCMC estimate of the path of the Gaussian OU process, and the corresponding estimate of the latent process Φ for the 2-OU model and the unrestricted 3-OU model.

5.10.3 Diagnostics

Next we apply the proposed diagnostics tools to the MCMC results for the 2- and 3-OU models. Figure 5.15 provides the posterior p -values when testing the Normality assumption of the residuals of the estimated Y_0 with the KS test at the 5% level. Results are provided for the 2-OU model (top panel) and the 3-OU model with unrestricted (middle panel) and restricted (bottom panel) mean jump size parameters β_i . We have used 1000 posterior samples for each model.

As expected, the simple 2-OU model is rejected as a suitable model for the EEX time series, in the sense that removal of the estimated jumps does not give a Gaussian OU model for Y_0 . This is basically due to the fact that negative price jumps have not been captured, and so the residuals of Y_0 present large negative returns. On the other hand, the 3-OU models with negative jumps are well supported by the test, which gives Bayesian p -values well above 5%. We notice that the mean p -value for the restricted 3-OU model is only slightly larger than that for the unrestricted model, and thus both models are equally significant to the KS test.

Next we compare the statistical properties of the data at hand and of the estimated models. Table 5.12 presents summary statistics of the deseasonalised EEX series and of simulated data generated with the estimated parameters provided in Table 5.11. The results are based on the grand average of 500 realisations of each model.

As shown, the mean, SD and quartiles of the deseasonalised data are resembled by the estimated 2-OU and 3-OU models. Nonetheless, the skewness and kurtosis of the distribution of the EEX data seem to be underestimated by the fitted models. We already know from the analysis of the APXUK dataset that the occurrence of jumps can affect significantly these two statistics. Indeed, removing the largest jumps from the observed data gives a SD of 0.31, a skewness of 1.879 and a kurtosis equal to 11.1, which are much closer to the statistics of the 2-OU and the 3-OU models. However, we notice the restricted 3-OU model has an extremely low skewness, far away from the observed sample skewness. The unrestricted 3-OU models seem to be doing really well judging from its statistical properties and the hypothesis testing results with the KS test which we showed above.

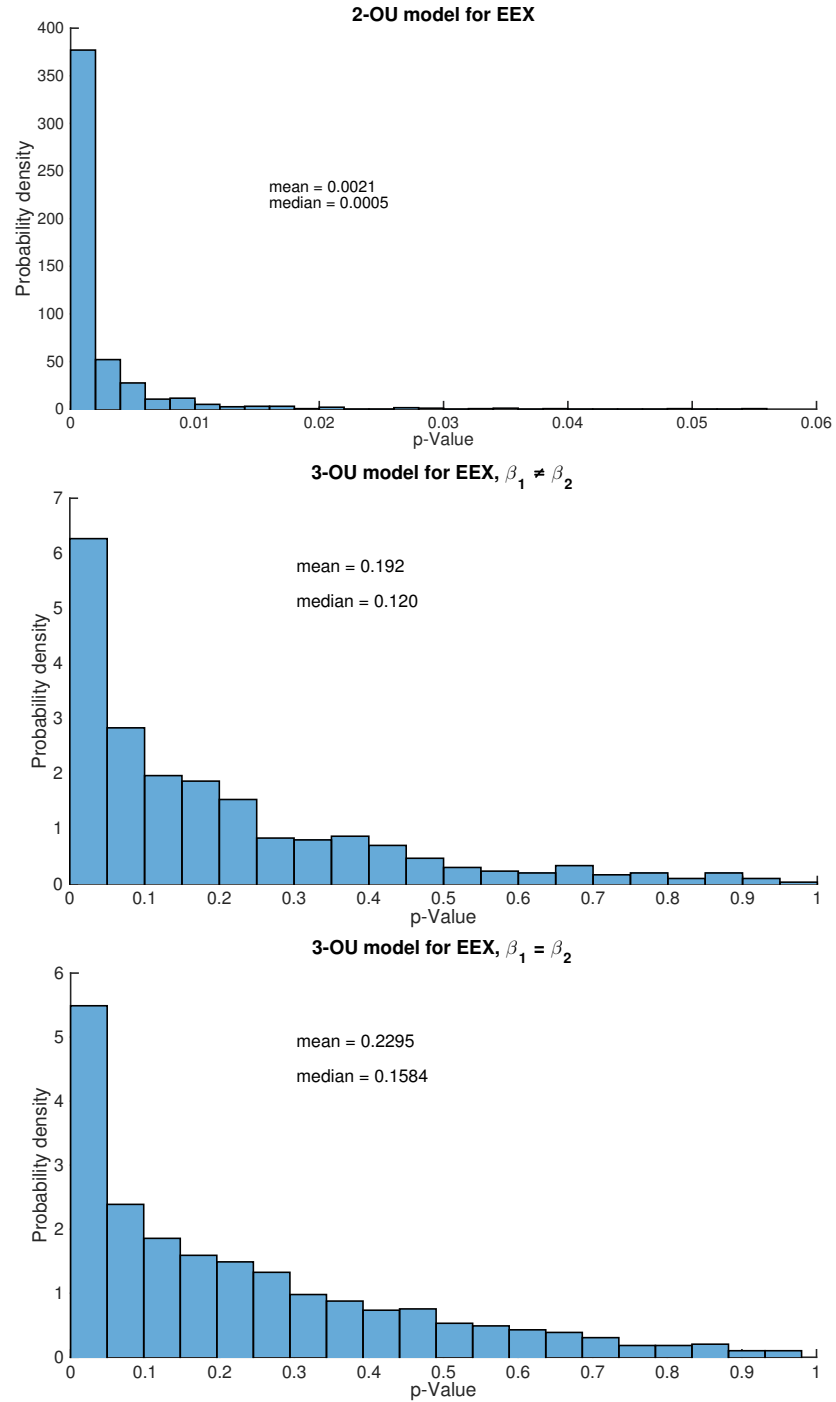


Figure 5.15: Posterior of p -values obtained when testing the estimated noise process of Y_0 for Normality using the Kolmogorov-Smirnov test at the 5% significance level.

5.10.4 Prior sensitivity analysis

In this section, we conduct a prior sensitivity analysis for the parameters of the unrestricted 3-OU model. Our aim is to explore the implications of alternative priors on the

	Mean	SD	Skewness	Kurtosis	Q_1	Q_2	Q_3
EEX dataset	1.0523	0.4283	7.3419	107.4844	0.8489	0.9991	1.1546
2-OU model	1.0496 (0.0186)	0.3418 (0.0244)	2.9987 (1.1019)	29.6479 (19.3840)	0.8517 (0.0192)	1.0229 (0.0184)	1.2016 (0.0198)
3-OU $\beta_1 \neq \beta_2$	1.0447 (0.0237)	0.3576 (0.0220)	2.0153 (0.8001)	20.6779 (10.1044)	0.8539 (0.0242)	1.0237 (0.0234)	1.1984 (0.0254)
3-OU $\beta_1 = \beta_2$	1.0427 (0.0236)	0.3737 (0.0225)	0.8305 (0.9599)	19.1751 (10.5215)	0.8564 (0.0240)	1.0268 (0.0228)	1.2027 (0.0247)

Table 5.12: Summary statistics of the EEX time series and simulated data generated with the estimated 2- and 3-OU models. The results for simulated data are based on the grand average across 500 simulated sample paths. The standard deviation is provided in parenthesis and Q_i denotes the i -th quartile.

posterior estimates and the performance of the algorithm. As we mentioned earlier, the chains for the parameters of the diffusion have excellent mixing properties. Indeed, after conducting prior sensitivity checks on μ and λ_0 we found no significant change on the posterior estimates and ACFs of the parameters. Also, since we have been using a Uniform distribution on $\rho_i = e^{-\lambda_i}, i = 1, 2$, it is not necessary to check other prior specifications for these parameters. However, statistical correlation between the parameters of the jump processes, in particular of the jump intensity rate η_i , and the volatility σ can cause some problems. Thus it is interesting to check the robustness of the algorithm to the specification of the priors for σ, η_1 and η_2 .

First, let us check the Inverse Gamma prior with mean 0.01 that we have used for σ^2 . We fit again the 3-OU model to the deseasonalised EEX series using a flat prior for σ^2 :

- Prior 2: $\sigma^2 \sim U(0, 0.25^2)$.

This distribution covers all the values of σ^2 that have estimated for the datasets studied in this chapter. The priors of all other parameters remain unchanged. Table 5.13 reports the posterior estimates of the 3-OU model parameters when using Prior 2. A comparison with the results in Table 5.11 reveals that our new choice of prior has no effect on the posterior estimates, which have remained virtually unchanged. The prior estimates of σ^2, η_1 and η_2 are still around 0.012, 0.12 and 0.13, respectively. Concerning the mixing of the algorithm, we notice that the ACFs of σ^2, λ_1 and η_2 are

affected, taking slightly longer to die away with Prior 2 than with our initial choice of prior distribution. This is because we have replaced Gibbs steps for less efficient MH steps to update σ^2 , since in this case the full conditional is not easy to sample from directly. In addition, correlation between the volatility and the jump parameters also plays a role in the loss of performance observed in this case.

Next we check the prior distributions for the jump rates η_1 and η_2 . In this case we substitute the Gamma prior distributions with mean and SD equal to 0.1 for flat distributions:

- Prior 3: $\eta_i \sim U(0, 0.5), i = 1, 2$.

Once more we replace Gibbs steps with MH steps and update η_1 and η_2 jointly using a block random-walk MH algorithm. Table 5.11 provides posterior summaries of all parameters when using Prior 3. The results provide further evidence of the robustness of the MCMC algorithm to alternative prior specifications. As shown, Prior 3 has no effect on the posterior estimates of the parameters, which are unchanged when compared to those in Table 5.11. Nevertheless, we notice that the mixing properties of the algorithm tend to deteriorate in a similar fashion when using Prior 2. Nonetheless, all ACFs die away within the first 800 lags, showing a relatively good performance of our algorithm for this 9-parameter model with two latent processes.

5.11 National grid SAP for gas spot prices

We finish off this chapter with an application to the SAP dataset covering the period from April 6, 2000 until July 31, 2007. This dataset is much larger than the APXUK or EEX datasets, with a total of 2680 observations including weekends.

Gas prices are known to have a relatively long-term dynamics due to the storability of natural gas. Additionally, although seasonality is present in gas prices, the intra-day and intra-week behaviour of these prices differs from that of electricity prices. Nonetheless, as we noted earlier, capacity constraints in the gas pipeline system make gas markets to behave as electricity markets. Thus we expect to find not only pronounced positive price spikes but also mild negative price spikes. In terms of modelling, we shall see from the results below that the gas index we study here is similar to the

Parameter	Posterior properties			
	Prior 2		Prior 3	
	Mean	SD	Mean	SD
μ	1.0152	0.0220	1.0155	0.0225
e^{-1/λ_0}	0.8776	0.0157	0.8814	0.0154
λ_0	7.7998	1.1199	8.0708	1.1680
σ^2	0.0126	0.0013	0.0121	0.0014
e^{-1/λ_1}	0.2029	0.0331	0.2028	0.0327
λ_1	0.6281	0.0642	0.6279	0.0635
η_1	0.1227	0.0205	0.1288	0.0216
e^{-1/λ_2}	0.2248	0.0370	0.2289	0.0378
λ_2	0.6720	0.0752	0.6804	0.0771
η_2	0.1336	0.0303	0.1451	0.0361
β_1	0.9038	0.1119	0.8834	0.1121
β_2	0.4398	0.0875	0.4242	0.0837

Table 5.13: Summary of posterior properties for each parameters when fitting the 2- and 3-OU models to the EEX series. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100. The latent process Φ was updated five times per update of the remaining parameters. The process Y_2 is used to model negative jumps.

EEX Phelix we studied in the previous section. We start our inference methodology by fitting a suitable seasonal trend for gas prices.

5.11.1 Deterministic part

Following Benth et al. (2008), we decide to fit a seasonal trend with only yearly periodicity to the gas log-prices,

$$f(t) = a_1 + a_2 s + \cos(2\pi(t - a_3)/365).$$

Using a weighted non-linear least squares algorithm we find $a_1 = 1.66$, $a_2 = 0.0002776$ and $a_3 = 95.836$, with all parameters being significant at the 5% level. The coefficient of determination is $R^2 = 0.362$, which is much smaller than for the APXUK and EEX datasets. Our model for the gas spot prices differs slightly from the model for

electricity prices. Now we assume that the deseasonalised data is modelled by

$$X(t) = \log X_{obs}(t) - f(t).$$

Our practical work suggests that this model better fits the data than the model with multiplicative seasonality in equation (5.4). Figure 5.16 provides the gas log-spot prices superimposed to the fitted seasonal trend (top panel) and the deseasonalised data (bottom panel). Similar to the EEX series, negative price jumps are evident in the SAP series, and the seasonality effects produce relatively high gas prices during the cold season (October-March) and low prices over the warm season (April-September).

The long-term behaviour of the deseasonalised SAP data is evident from the results in Table 5.14, which provides the fitted theoretical ACF of our spot price models to the empirical ACF of the data. As shown, for all models considered there exists a long decay rate in the ACF ranging between 100 and 180 days, whose weight is always the largest among all possible options. This is a significant difference with the results for the APXUK and EEX datasets reported in Table 5.5 and Table 5.10, where the greatest weight was given to comparatively short decay rates.

$n + 1$	ω_i			λ_i			SS
1	1	-	-	107.3	-	-	0.5342
2	0.4482	0.5518	-	6.24	174.88	-	0.4288
3	0.0879	0.3844	0.5177	0.79	8.76	179.7	0.2966

Table 5.14: Fit of the autocorrelation function of the SAP time series using superpositions of OU processes. $n + 1$ denotes the number of OU components, the ω_i are the estimated weights of the exponential functions, the λ_i are the estimated times to mean reversion and SS is the sum of squares resulting from the procedure.




Figure 5.16: The daily average SAP series of gas spot log-prices along with fitted seasonal trend (top graph), and the resulting deseasonalised time series (bottom graph).

5.11.2 Stochastic part, 2- and 3-OU models

To fit a sum of 2-OU and 3-OU processes to the SAP, we have run the MCMC algorithms two million iterations, removed a burn-in period of 500000 and thinned the chain taking one every 100th iteration. We have employed the same prior distributions used for the APXUK dataset in Section 5.9. For the 3-OU model we take the processes Y_1 and Y_2 to model possible positive and negative jumps in the data under study, respectively.

Table 5.15 provides the posterior mean and standard deviation of each model parameter obtained when fitting the 2-OU model and the restricted and unrestricted 3-OU models to the SAP time series. In line with our previous discussion, an interesting result is the relatively long time to mean reversion λ_0 of the diffusion process implied by the 2-OU model and the 3-OU model. This indicates a long-term memory of gas prices under normal market conditions, a consequence of the storability and large reserves available of natural gas. We notice that over the period of time under study the volatility of gas prices is similar to the volatility of APXUK electricity prices, but smaller than the EEX prices. It is also worthwhile noting that the restricted and

unrestricted 3-OU models imply similar dynamics for the spike prices.

Regarding jump frequency, we observe that the SAP series has a comparatively higher frequency of both positive price jumps and negative price jumps than the EEX series. Nonetheless, the SAP price jumps are much smaller than the EEX price jumps on average.

Parameter	Posterior properties					
	2-OU model		3-OU model			
	Mean	SD	$\beta_1 \neq \beta_2$		$\beta_1 = \beta_2$	
			Mean	SD	Mean	SD
μ	-0.0169	0.0430	-0.0002	0.0534	0.0127	0.0534
e^{-1/λ_0}	0.9418	0.0071	0.9710	0.0049	0.9719	0.0047
λ_0	16.9321	2.2150	34.9776	6.3094	36.0763	6.4657
σ^2	0.0171	0.0007	0.0059	0.0006	0.0057	0.0005
e^{-1/λ_1}	0.0213	0.0133	0.5791	0.0395	0.5725	0.0368
λ_1	0.2535	0.0420	1.8513	0.2304	1.8103	0.2095
η_1	0.0796	0.0165	0.1785	0.0339	0.1520	0.0237
e^{-1/λ_2}	--	--	0.4913	0.0314	0.5078	0.0301
λ_2	--	--	1.4145	0.1287	1.4833	0.1323
η_2	--	--	0.1608	0.0301	0.2095	0.0324
β_1	1.0024	0.1769	0.1751	0.0210	0.2011	0.0176
β_2	--	--	0.2493	0.0310	0.2011	0.0176

Table 5.15: Summary of posterior properties for each parameter when fitting the 2- and 3-OU model to the gas SAP series. The results were obtained by running two million MCMC iterations, removing the first 500000 as burn-in and thinning the resulting chain one every 100. The latent process Φ was updated 5 times per update of the remaining parameters. The process Y_2 is used to model negative jumps.

5.11.3 Diagnostics

Figure 5.17 presents the usual plots of the posterior p -values of the KS test when testing the residuals of Y_0 for Normality at the 5% level. The results show a massive improvement on the test when incorporating an additional jump component to account for negative prices. The present results are conceptually similar to those for the EEX, where we found that the restricted 3-OU model and the unrestricted one give similar mean p -values.

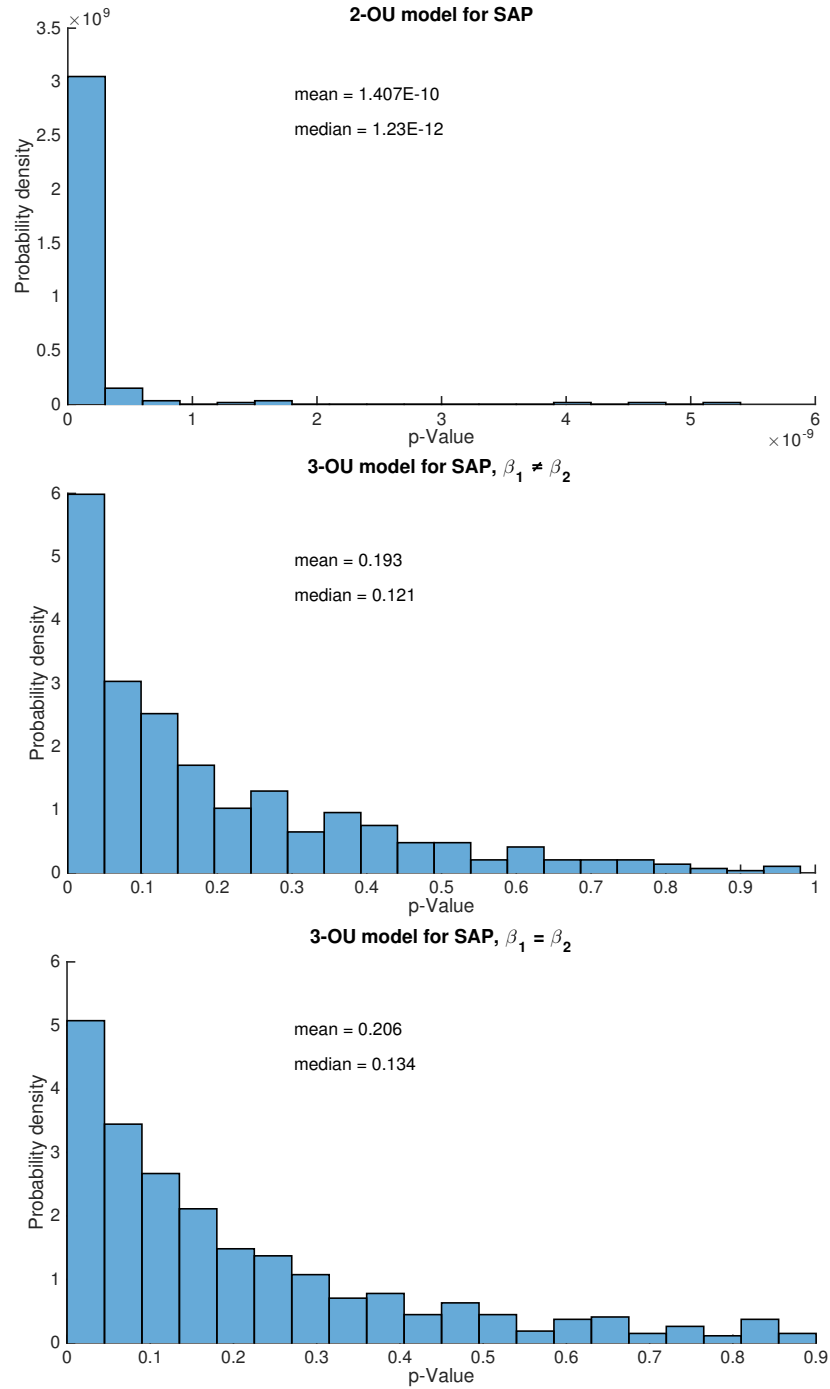


Figure 5.17: Posterior of p -values obtained when testing the estimated noise process of Y_0 for Normality using the Kolmogorov-Smirnov test at the 5% significance level.

Finally, Table 5.16 provides summary statistics of the SAP dataset and of data simulated with each of the estimated models in Table 5.15. The summary statistics for each superposition of OU processes is based on the grand average across 500 realisations. As shown, while the mean, SD and quartiles of the simulated data are

quite close to those of the deseasonalised data, there are small discrepancies in skewness and kurtosis. After removing the largest jumps from the SAP, we find that the empirical skewness and kurtosis of the SAP are 0.1085 and 4.0846, respectively, and thus in better agreement with the MCMC results. This shows again that the sample statistics and estimation results tend to be influenced by large shocks in the data. In any case, as we have already observed in the case of the APXUK and EEX dataset, the unrestricted 3-OU model resembles closely the empirical distribution of the SAP dataset, while being well supported by our diagnostic tool.

	Mean	SD	Skewness	Kurtosis	Q_1	Q_2	Q_3
SAP dataset	0.0	0.3600	-0.2623	7.6402	-0.2195	0.0078	0.2222
2-OU model	0.0125 (0.0409)	0.4078 (0.0214)	0.5874 (0.4077)	6.9756 (5.6377)	-0.2567 (0.0440)	0.0029 (0.0425)	0.2671 (0.0443)
3-OU model $\beta_1 \neq \beta_2$	0.0166 (0.0524)	0.3598 (0.0230)	-0.0103 (0.1688)	3.4205 (0.3618)	-0.2207 (0.0560)	0.0154 (0.0548)	0.2519 (0.0575)
3-OU model $\beta_1 = \beta_2$	0.0141 (0.0512)	0.3562 (0.0237)	0.0395 (0.1810)	3.3653 (0.3433)	-0.2213 (0.0561)	0.0117 (0.0539)	0.2465 (0.0554)

Table 5.16: Summary statistics of the SAP time series and simulated data generated with the estimated 2- and 3-OU models. The results for simulated data are based on the grand average across 500 simulated sample paths. The standard deviation is provided in parenthesis and Q_i denotes the i -th quartile.

Chapter 6

Numerical examples

The purpose of this chapter is to illustrate our risk-sensitive approach to optimal switching in different problems arising from applications of stochastic optimal control to energy systems. Throughout this chapter we take the approach of comparing our methodology to the risk-neutral case, which provides a baseline with which we may assess the potential benefits of risk-sensitive control. This also allows to draw conclusions regarding differences in risk attitudes of system controllers, and to demonstrate the fundamental differences between risk-neutral and risk-sensitive control policies. As indicators of risk, we use the risk measures VaR, CVaR and EVaR, which were briefly reviewed in Section 2.5. In addition, we suggest an indicator tool of the balance between average case performance and risk offered by our methodology. The indicator helps to choose an optimal degree of risk-sensitivity, in the sense that it will provide the minimum sacrifice in performance (e.g. total expected revenue/cost) and maximum decrease in risk (as measure by CVaR, for example).

The series of examples considered here are of increasing complexity. For the reader's convenience, Table 6.1 summarises their main characteristics, such as number of switching regimes, dimension of the problem and whether the underlying process has been calibrated to real data.

Throughout this chapter, we use the convention that a control strategy for $\rho = 0$ means that the optimal control was calculated with the RN-LSM algorithm. Naturally, when $\rho \neq 0$ we use the RS-LSM algorithm. Also, we always report the absolute value of ρ . However, we clarify here that we are interested only in risk-averse situations. Thus, in equation (3.24), we formally use $\rho < 0$ when maximising a total reward

function (Example 1, 2 and 3), whereas we use $\rho > 0$ when minimising a total cost function (Example 4).

Example	# of regimes	Dimension	Prices	Time resolution
Gas-fired power plant	2	2	Simulated	Daily
Energy storage	3	3	Simulated	Daily
Energy storage	3	5	Fitted UK	Daily
District energy systems	15	7	Fitted UK	Hourly

Table 6.1: Main characteristics of the examples considered in this chapter.

Comments on numerical implementation

We would like to mention that for each numerical application considered in this chapter, we study the effect of ρ on the risk-sensitive value function, switching regions, risk measures, and so forth. Therefore it is necessary to choose a ‘suitable’ range of ρ which is of special interest. Here, we are interested not only in optimising the risk-sensitive criterion, but also in finding values of ρ for which the risk measures VaR, CVaR and EVaR decrease. Recall Section 2.5 and our discussion in Section 2.6. Moreover, when possible, we will attempt to find a value of ρ offering the ‘best’ tradeoff between a reduction in average-case performance and a decrease in, say, the CVaR risk measure.

Although these suitable values of ρ are problem-specific and some trial and error may be necessary, we suggest to initially use relative small values of ρ (~ 0.1) and compute the risk measures to guide the search. Also, it is worth mentioning that this range may depend on the relative scale of the value function, so when the (risk-neutral) value function is large, interesting values of ρ will tend to be small, and vice versa.

In addition, since the exponential function is calculated at every step of the algorithm it is advisable to maintain its argument relatively small (e.g. in the interval ranging from -10 to 10). We recommend normalising variables (prices, capacity rates, etc.) appropriately to avoid any computational overflows. Note that this is only for numerical convenience, and that the value function is unaffected provided we renormalise back all variables.

Finally, we note that computing a suitable range of ρ does not necessarily imply a long computational time, since it can be easily implemented using a parallel computing

architecture. For instance, we completed 12 runs of the algorithm for Example 1 below in under one minute using a computer with a 12-core Intel Xeon E5 processor running at 2.5GHz. As general purpose GPUs (graphics processing units) with several hundreds of cores become widely available in modern personal computers, we imagine that many more runs of the algorithm can be completed in a relatively short time span, making this step much more computationally efficient.

6.1 Energy tolling agreements

Our first example concerns the optimal operation of a gas-fired power plant in a liberalised energy market. This problem arises in the context of pricing energy tolling agreements, a useful mechanism introduced by the power industry which allows market participants to share and mitigate the risk associated with their business. In the following we briefly introduce energy tolling agreements and explain the underlying optimal control problem of a gas-fired plant in this context. Important references here are Deng and Xia (2006) and Carmona and Ludkovski (2008).

Background

A tolling agreement involves two parties: the owner of a fossil-fuel power plant, who wants to rent or temporarily transfer the ownership of their power plant, and the buyer or holder of the agreement, who wants to take control of such power plant. A tolling agreement gives the holder the right to either operate the power plant over a specified horizon or take the produced output electricity over pre-specified periods of time. In return, the plant owner receives a fixed payment, which represents the value that can be obtained by fully exploiting the operational flexibility of the power plant.

As the name suggests, the operation of a gas-fired plant implies buying natural gas from the market, converting it into electricity and selling the output back on the market. This conversion process, however, is not perfect and occurs at a known rate called the heat rate, which depends on the particular power plant under consideration. The heat rate \overline{Hr} measures the amount of fuel units (in our case, of gas) that are needed to produce one unit of electricity. In Europe, the heat rate is quoted in GJ/MWh (giga-joule per megawatt-hour), whilst in the US it is given in units of MMBtu/MWh

(millions of British thermal units per megawatt-hour). A power plant with a lower heat rate is more efficient than one with a higher heat rate, since in the former case more units of fuel can be converted into electricity.

Since the power plant produces electricity by converting natural gas, the revenue accrued by the plant operator is a function of the difference in price between an electricity unit and the amount of gas used to generate it. This price difference is known as the spark spread (between electricity and gas). It measures the ‘spread’ between the commodity price (electricity) and its production cost (from converting gas). More specifically, let E_t and G_t denote respectively the market price of electricity and of gas at time t , Cap the capacity of the plant (i.e. the amount of electricity which can be produced, say, annually) and K the operating costs. Then the payoff of the operator takes the form

$$\psi(E_t, G_t) = Cap(E_t - \overline{Hr}G_t) - K, \quad (6.1)$$

where $E_t - \overline{Hr}G_t$ represents the spark spread at time t . In an energy tolling agreement, the power plant operator has the right but not the obligation to generate electricity. That is, she can decide when to turn on the plant and when to turn it off. Hence, a rational plant operator will exploit the relative values of the spark spread (net the operational costs) as time evolves. Ignoring any operational constraints, the operator will bring the power plant on-line if the spark spread is positive, since doing so generates profits, and conversely, she will bring the plant off-line when the spark spread is negative.

We note however that control decisions depend not only on the relative values of the spark spread but also on the operational constraints of the power plant and the contractual constraints of tolling agreements. On the operational side, typical constraints are startup and shutdown costs, ramp-up/down constraints, and different heat rates associated to different electricity output levels. On the other hand, since often restarts can deteriorate a power plant, tolling contracts may specify a maximum number of restarts of the power plant (e.g. a maximum of one restart every month) and possibly a financial penalty if this number is exceeded. Then this type of constraints are imposed to increase the expected lifetime of a power plant and decrease the likelihood

of a forced outage. In general, power plant operators use the spark spread to cover the plant operational costs. In this context, for instance, a power plant may be brought on-line if the spark spread is positive and the expected profit gained from this decision also covers the associated switching costs (the ramp-up costs).

The control problem associated with running a gas-fired plant refers to exploiting over time the available operational options of the plant, i.e. to start-up or to shut-down the plant, in order to maximise the total profit of the plant operator. The control policy takes into account a vector of stochastic energy prices (both gas and electricity, or directly the spark spread) and the constraints associated with running the plant. Furthermore, due to the high price volatility found in energy markets, in particular for electricity prices (recall Section 5.1), it is important to consider and control the price risk faced by gas-fired power plant operators.

The next example, which is borrowed from Carmona and Ludkovski (2008), aims to study this control problem in both risk-averse and risk-neutral situations. In order to focus on the conceptual contribution of our risk-sensitive approach to optimal switching, the spark spread process is modelled directly as a one-dimensional process. By doing this, we are able to plot the switching regions as functions of the whole system state vector (in this case, time and price). This is not possible in the high multi-dimensional problems considered in subsequent sections.

6.1.1 Example 1. An on/off optimal switching problem

Consider a gas-fired power plant with a maximum annual capacity of 10 GWh, and which is operated over a time horizon of $T = 2$ years. The power plant has two operating regimes available, namely, the ‘off’ regime, in which the plant is shut-down, and the ‘on’ regime, in which the plant is running at full capacity. We assume that the spark spread between electricity and gas (in units of £/MWh) is given by

$$dX_t = 2(10 - X_t)dt + 2dW_t, \quad X_0 = 10.$$

Further, we suppose that there are no operating costs when the power plant is off-line, whereas the annualised operating costs in the ‘on’ regime are equal to k£100

(thousands of pounds). Accordingly, the payoff rates (in k£ per year) associated with each operating regime are $\psi_1(X_t) = 0$ for the ‘off’ regime, and $\psi_2(X_t) = 10X_t - 100$ for the ‘on’ regime. The switching cost between operating regimes is constant and equal to $K_{12} = K_{21} = 0.3$ (k£).

Numerical implementation

We implement the RN-LSM and the RS-LSM algorithms of Section 4.3 using $N = 200$ time steps, $N_p = 16000$ sample paths of X_t and at least seven linear piecewise basis functions for the regression step. To study the effect of the risk-sensitivity parameter ρ , we run the RS-LSM algorithm with $\rho = 0.1k, k = 1, \dots, 20$.

Switching Boundaries

Let us first analyse the optimal switching boundaries for the present example. We recall that the switching boundaries describe the optimal switching strategy at each time t , regime i and state x_t (see Section 4.3.1). Since in this example there are available two switching regimes (with continuous payoff functions), we have that for each regime, there exist (Carmona and Ludkovski, 2008) two connected switching regions delimited by one switching boundary: for example, when the current regime is ‘off’ there will be a (two-dimensional) region where it is optimal to continue in this regime and another one where it will be optimal to switch to the ‘on’ regime. Clearly, switching from one regime to the other depends on the relative values of the spark spread. In this example, the switching boundary consists of threshold values that indicate when it is optimal to change operating regimes.

To illustrate this, the top panel of Figure 6.1 provides the switching boundaries of $V(t, x, i)$ for $\rho = 0, 0.5, 2$ and for each operating regime over the entire operating period $[0, T]$. Superimposed to the switching boundaries it is one possible realisation of the spark spread process over the same period (black line).

The switching boundaries for $\rho = 0, 0.5, 2$ are given by the blue, red and yellow lines, respectively. Let us take for example the switching boundaries of the risk-neutral strategy, $\rho = 0$. The upper blue line represents the switching boundary from the ‘off’ to the ‘on’ regime, whereas the lower blue line indicates the switching boundary from

the ‘on’ to the ‘off’ regime. So, for instance, when the power plant is off-line, it will be optimal to bring it on-line when the spark spread is above the blue upper boundary (i.e. when the spark spread is larger than 10.7). If the plant is on-line and the spark spread is below the lower blue boundary, it will be optimal to switch the plant to the off mode. Similar conclusions apply to the cases $\rho = 0.5$ and $\rho = 2$.




Figure 6.1: Top panel: The switching boundaries of $V(t, X_t, i)$ for Example 1 obtained with different degrees of risk-sensitivity ρ , superimposed to a sample path of the spark spread process X_t . Bottom panel: The corresponding cumulative revenue of the power plant operator for the given path of the spark spread and each value of ρ .

The region between the upper and lower boundary (say, between the upper blue line and the lower blue line) is the so-called hysteresis band (Dixit, 1989). In this area no switching occurs, owing to the fixed costs associated with switching regimes. When the spark spread is within the hysteresis band, the expected revenue obtained by keeping the current operating regime is larger than the expected revenue obtained by switching regimes and paying the corresponding fixed switching cost. Hence the plant operator finds sub-optimal to change regimes and is willing to wait until the spark spread is favourable. The hysteresis band widens as the switching costs increase, and conversely it narrows as the switching costs decrease. In the absence of switching costs the hysteresis band disappears.

The top panel of Figure 6.1 also shows that the switching boundaries in this example are basically constant when the time to maturity $T - t$ is relatively large. This is because the spark spread process, the payoff functions and the switching costs are all stationary, and hence, time has little effect over the optimal switching policy. The effect of non-stationarity on energy prices will be illustrated in Example 3 and Example 4 below. Nevertheless, close to maturity T , the switching boundaries get away from the mean level of the spark spread process. This is explained by noting that, there, the switching cost is much greater than the potential revenue that can be obtained by switching regimes. This holds true for both the risk-neutral switching strategy and the risk-sensitive one.

Difference between risk-neutral and risk-sensitive boundaries

We turn now our attention to the effect that the risk-sensitivity parameter has on the switching boundaries. Although close to T the risk-neutral and the risk-sensitive switching boundaries coincide, there are noticeable differences between them far from maturity. Inspection of the top panel of Figure 6.1, reveals that the switching boundaries are shifted upwards as ρ increases. This means that a risk-sensitive plant operator waits longer to bring the unit on-line, and turns it off sooner, than a risk-neutral operator. This is related to the risk attitudes of the plant operators. We recall that in the risk-neutral case a (rational) power plant operator aims only to maximise the expected total revenue over the finite horizon of the tolling agreement. Then the risk-neutral optimal scheduling of the power plant exploits the spark spread (up to switching costs) and the flexibility of the plant to achieve this goal, regardless of the financial risk associated with the control strategy. A risk-averse power plant operator, however, is partly motivated by the possible decreases in risk (as measured by the entropic risk measure and, if ρ is small, by the variance) with respect to the total revenue that could be obtained over the given horizon.

Hence, in contrast with a risk-neutral operator, a risk-sensitive one will not take advantage of the available market opportunities to maximise her total expected revenue. This is because the relative (upward and downward) movements of the market may produce large variations of her total revenue, and consequently a large amount of price risk. Instead, she will wait until the spark spread is high enough above its

mean level that bringing the unit on-line, and hence, ‘playing’ in the market pays a premium for the risk associated with this action. Conversely, she will bring the unit off-line soon enough to prevent large variations in revenue caused by a downward trend in the market. Thus, in this example a risk-sensitive controller prefers to wait longer to enter in a risk situation and to exit sooner from it, relative to a risk-neutral controller.




Figure 6.2: Top panel: The switching boundaries of $V(t, X_t, i)$ for Example 1 obtained with different degrees of risk-sensitivity ρ , superimposed to a sample path of the spark spread process X_t . Bottom panel: The corresponding cumulative revenue of the power plant operator for the given path of the spark spread and each value of ρ .

The effect that these different risk attitudes have on the total expected revenue of the operator are significant. As an illustration, the bottom panel of Figure 6.1 shows the accumulated revenue over time when $\rho = 0, 0.5, 2$, that the operator would obtain with the particular realisation of the spark spread process provided on the top panel of the figure. Although the risk-neutral strategy seems to be doing really well over the first year, the subsequent downward market trend makes that some part of the accumulated wealth is lost, since the power plant is turned off too late. In contrast, the risk-sensitive optimal scheduling limits this loss by bringing the unit off-line a few months sooner. This situation is exacerbated in the illustration of Figure 6.2, whereby the market conditions work against the risk-neutral operator, while the risk-sensitive

one successfully limits any potential financial losses. Nonetheless, the risk-neutral scheduling policy for the power plant gives a higher expected revenue than the risk-sensitive one, as will see next.

Effect of ρ on risk indices

Figure 6.3 shows the effect of ρ on the mean total revenue, its standard deviation (SD) and risk measures VaR, CVaR and EVaR, all at the 95% confidence level (similar results are obtained for confidence levels in the range 90% - 99%, but for clarity we focus only on a fixed confidence level). The results are obtained by implementing the resulting optimal switching controls to each of the $N_p = 16000$ sample paths in a forward Monte Carlo simulation. In this example we are dealing with a profit function, and so we take the convention that the larger the value given by the risk measure the smaller the risk. Additionally, we recall here that the entropic risk measure is a monotone function of ρ (see Section 2.5; our numerical results also verify this assumption). Hence, plotting it against ρ provides little information regarding which values of ρ should be used in practice. Thus we have decided to focus on the VaR, CVaR and EVaR, and study whether our risk-sensitive approach can also minimise these risk measures. Recall our discussion in Section 2.6.

The top panel of Figure 6.3 shows that the risk-neutral switching policy generates the maximum expected revenue across all other scheduling policies for the power plant, albeit it also produces the largest variations of the total revenue and the largest price risk, as measured by VaR, CVaR and EVaR. On the other hand, as ρ increases, the expected total revenue and its variance decrease, while all the risk measures increase (i.e. the indicated risk decreases). This illustrates clearly the tradeoff between expected revenue and risk offered by our risk-sensitive optimisation, in the sense that for a sacrifice in the total expected revenue we receive in return a decrease in price risk. We note however that this behaviour occurs only locally in a neighbourhood of $\rho = 0$ (as we would expect from our discussion in Section 2.6). Our numerical results (which for clarity we have omitted) show that the variance and risk indices increase again as ρ gets large.




Figure 6.3: For the total revenue function of Example 1, its mean, standard deviation (SD) and risk measures VaR, CVaR and EVaR all at the 95% confidence level, as a function of the degree of risk-sensitivity ρ . In this example the greater the risk measures the smaller the risk faced by the plant operator. The results are based on 16000 sample paths of the spark spread process (X_t) .

An interesting point arising from the numerical results is that the benefits with respect to risk minimisation seem to outweigh the decrease in expected total revenue. That is, we may gain more by taking a tradeoff between revenue and risk than doing otherwise. To better illustrate this point, we have explicitly computed the change in mean and risk measures of each switching strategy with respect to the risk-neutral one. More specifically, for each value of $\rho_k = 0.1k, k = 1, \dots, 20$ we calculate the difference between the mean total revenue obtained with ρ_0 and with ρ_k ,

$$V_{\rho_0}(0, X_0, i) - V_{\rho_k}(0, X_0, i). \quad (6.2)$$

Hence a positive value in (6.2) indicates that the expected total revenue obtained with ρ_k increased with respect to the risk-neutral strategy and conversely, if the difference is negative it means that the expected total revenue decreased. For a fixed confidence level, the same calculations are carried out for the risk measures VaR, CVaR and EVaR. In this case, we indicate whether the risk index decreases or increases with respect to

the risk-neutral case. The results are given in Figure 6.4 and can be compared directly with the results in Figure 6.3. It is now obvious from Figure 6.4 that, when measured with respect to the baseline case $\rho = 0$, the risk indices decrease faster than the mean total revenue when $\rho \in [0, 1.5]$. Therefore, from a risk-averse perspective, it is advantageous to use a risk-sensitive optimisation since for a small sacrifice in average revenue one receives a larger decrease in the amount of price risk.




Figure 6.4: For Example 1, the net change with respect to the risk-neutral case $\rho = 0$ of the mean and risk measures of the total revenue, as the risk-sensitivity ρ increases. A negative value indicates a decrease with respect to the risk-neutral case. The results are based on 16000 sample paths of the spark spread process (X_t) .

A risk-averse controller would then choose the strategy giving the minimum decrease in expected revenue and the maximum reductions in price risk. In this sense, she would choose some optimal degree of risk-sensitivity ρ^* to achieve this, if such value exists. Accordingly, we propose to use the following indicator of the optimal degree of risk-sensitivity ρ^* . For each $\rho \geq 0$, let $P_\rho := P_\rho(t, x, i; u)$ denote the total profit function using the admissible switching control u when the initial data is (t, x, i) , and let $V_\rho := V_\rho(t, x, i)$ be the expected value of P_ρ . Then for a finite set of possible

risk-sensitivity degrees $\{\rho_k\}_{k=1,\dots,N_\rho}$, $\rho_0 = 0$, we define ρ^* as

$$\rho^* = \arg \max_{\rho_k} \{V_{\rho_0} - V_{\rho_k} - (\text{CVaR}_{1-\alpha}(P_{\rho_0}) - \text{CVaR}_{1-\alpha}(P_{\rho_k}))\}. \quad (6.3)$$

For clarity we have dropped the arguments in the value functions. The indicator in equation (6.3) is simply the difference between the blue line and the yellow line in Figure 6.4. From the results in this figure, we find that $\rho^* = 0.9$ when $1 - \alpha$ is 0.95. If we chose, say, $\rho = 0.9$, the reduction in risk is almost the same as the reduction in revenue, so this strategy is not attractive from a risk-averse perspective. We note that instead of using CVaR in equation (6.3) one could also use any of the alternative risk measures. We have decided to use CVaR because it is a coherent risk measure with an intuitive financial interpretation of risk. Finally, we note that the indicator is easily adapted to the case when the random variable under study is a loss/cost function (rather than a profit function, as in this example).

6.2 Energy storage

In our next example we study the optimal operation and risk management of energy storage. In particular, we focus on the stochastic optimal control problem associated with the operation of a gas storage facility on a liberalised energy market. By applying the RS-LSM algorithm to this problem, we aim to strike a balance between the expected revenue of the storage operator and the financial risk arising from dealing with stochastic gas prices. Our simulation results and real-world applications demonstrate that it is possible to do this in such way that the loss in average total revenue is smaller than the reductions in risk. Throughout we will assume gas markets are liquid and efficient.

Background

Natural gas has come to play an important role in energy markets. In addition to being used for heating purposes during the cold season, natural gas has increasingly become the preferred fuel for power generation (Kjärstad and Johnsson, 2007). In fact, as reported by Kjärstad and Johnsson, gas has been progressively substituting

traditional fuels such as coal, lignite and also nuclear generation, owing to restrictions on carbon dioxide emissions and possible permanent shut-downs of nuclear power plants.

Thus, gas prices not only exhibit seasonal patterns resembling the yearly variations in gas and electricity demand, but also extreme short-term fluctuations due to weather conditions and external factors affecting gas and electricity supply. As illustrated in Section 5.11, UK gas spot prices tend to be relatively high during winter and relatively low during the typically mild British summer. Moreover, during the period 2001-2007, which we studied in the previous chapter, UK gas spot prices had a relatively high volatility similar to that of UK electricity spot prices.

In this context, gas storage has an intrinsic economic value since it offers the possibility to balance supply and demand in gas markets (Cyriel de Jong and Kasper Walet, 2003). Taking advantage of seasonality, a gas storage manager can inject gas into the storage which she buys when prices are relatively low, and withdraw the stored gas selling it at the market when prices are relatively high. This long-term strategy however is not the only one bringing value to the storage. The high daily volatility of gas prices, along with their mean reversion properties, also make possible to realise profits in the short-term. For example, if prices are expected to temporarily increase over the next few days, the manager can store some extra gas now and sell it appropriately as new market information arrives.

Managing a gas storage facility amounts then to find the optimal time to take any of three decisions: either buy gas and inject it into the storage, withdraw it and sell it to the market or simply hold it in the storage. Thus the problem of finding a managerial strategy for the storage facility fits naturally into the class of optimal switching problems considered in this thesis, in particular to the second special case of problem in Section 3.7. Clearly, these managerial decisions depend on the current and expected gas prices, and the current level of gas in store. Furthermore, the constraints on the storage must also be taken into account, including, among others, minimum and maximum storage capacity, constrained injection and withdrawal rates which may be volume-dependent and gas losses.

Over the past 15 years, the optimal operation and valuation of gas storage facilities

have been an active area of research. Many different solution approaches and mathematical frameworks have been employed to tackle this problem (see for example Felix (2012), Warin (2012) and references therein). Notably, methods based on PDEs (Chen and Forsyth, 2007, Thompson et al., 2009), recombining trees (Felix and Weber, 2012, Felix, 2012) and least squares Monte Carlo simulations (Cyriel de Jong and Kasper Walet, 2003, Carmona and Ludkovski, 2010) have been shown to be well suited to account for the complexities of the problem. In an interesting simulation study, Felix (2012) compares the computational performance of these three classes of numerical methods. He shows that PDE methods based on implicit finite difference schemes needed the longest computational time, although they were also the most accurate. In contrast, the recombining tree approach required the lowest computational time but converged slower than the other options. In addition, it is shown that Monte Carlo methods are quite demanding in terms of computational time and computer storage space, albeit they are ranked intermediate in terms of accuracy and computational performance. We notice, however, that by using the Quasi-Monte Carlo approach suggested in Section 4.5, we need a comparatively small number of paths to get excellent accuracy. As shown there, this considerably improves the computational time and memory requirements of the algorithm.

Surprisingly, all the above papers focus on the risk-neutral case, with the only exception being Felix and Weber (2012). In fact, this is the only paper we know where non-linear risk objectives have been introduced in the optimisation problem of gas storage facilities. The approach of Felix and Weber consists of introducing the CVaR risk measure (recall Section 2.5) in the optimisation objective. More precisely, they aim at finding a tradeoff between expected reward and price risk using the optimisation criterion

$$\mathbb{E}(P(u; [0, T])) + \beta \text{CVaR}_{1-\alpha}(P(u; [0, T])),$$

where $P(u; [0, T])$ indicates the total reward obtained using the control u over the horizon $[0, T]$, $\beta \in [0, 1]$ and $1 - \alpha$ is a specified confidence level. Thus the cases $\beta = 0$ and $\beta = 1$ can be seen as two extreme cases, where we either maximise the expected total reward or we minimise price risk. In a realistic application to German storage facilities, Felix and Weber (2012) show that when increasing β from 0 to 1, the loss in average total reward increases faster than the reductions in risk. We shall show below

that with our risk-sensitive approach it is possible to strike a tradeoff between average reward and risk in a more favourable way.

Gas storage model

As mentioned earlier, the operation of a gas storage can be seen as a (risk-sensitive) optimal switching problem with inventory. Here, we assume that C_s indicates the level of gas in store at time s , and that $C_s \in [C_{min}, C_{max}]$. Further, we assume known the storage withdrawal and injection rate functions, $a_{in}(C_s)$ and $a_{out}(C_s)$, along with the gas loss rates $\ell_{in}(C_s)$ and $\ell_{out}(C_s)$. We assume that these functions are measurable Lipschitz functions.

For simplicity, we suppose that when optimal, injection and withdrawal are carried out at the maximum possible rate. Thus, our control strategy is of bang-bang type. As shown in Chen and Forsyth (2007), this assumption has only mild implications on the numerical results. Thus the gas inventory level satisfies

$$dC_s = a_{u_t}(C_s)ds, \quad C_t = c,$$

where $u_s \in \{1, 2, 3\}$ indicates operating regime at time s . We take the convention that the regime ‘1’ is for injection, ‘2’ for hold/store and ‘3’ for withdraw, so $a_1 \equiv a_{in}$, $a_3 \equiv -a_{out}$ and $a_2 \equiv 0$. Similar rules apply to ℓ_{in} and ℓ_{out} . Following Carmona and Ludkovski (2010), we assume the payoff functions associated with each of these regimes take the form

$$\psi_{u_s}(G_s, C_s) = G_s(-a_{u_s}(C_s) - \ell_{u_s}(C_t)) - K_{u_s}(C_s),$$

where G_s is the gas price at time s and K_{u_s} represent various capacity charges proper to each regime. Having these ingredients, our aim is to find the value function and risk-sensitive optimal switching strategy in (3.6).

For comparison, we first apply our risk-sensitive approach to the risk-neutral situation presented in Example 5.1 in Carmona and Ludkovski (2010). After this, we consider a realistic situation using the fitted UK gas prices of the previous chapter.

Below we use the units Bcf (billion cubic feet) and MMBtu (million of British thermal units), with 1 Bcf being equivalent to 10^6 MMBtu.

6.2.1 Example 2. Without seasonality

In this example we consider a gas storage facility with a maximum working capacity of $C_{\max} = 8$ Bcf. The payoff functions (in \$/year) associated to each operating regime are

$$\begin{aligned} \text{Inject:} \quad & \psi_1(G_t, C_t) = -G_t a_1 - 0.1 C_t \\ \text{Store:} \quad & \psi_2(G_t, C_t) = -0.1 C_t \\ \text{Withdraw:} \quad & \psi_3(G_t, C_t) = -G_t a_3 - 0.1 C_t \end{aligned}$$

where G_t and C_t are, respectively, the gas spot price and the gas inventory level at time t , and the constant injection and withdrawal rates are $a_1 = 0.06 \cdot 365$ and $a_3 = -0.24 \cdot 365$, respectively, both in Bcf/year (in the next example we shall consider volume-dependent injection/withdrawal rates). The inventory variable C_t satisfies the ODE

$$dC_t = a_{u_t} dt, \quad C_0 = 4, \quad C_t \in [0, 8],$$

where $u_t \in \{1, 2, 3\}$ and $a_2 = 0$ for the ‘store’ regime. Thus the storage can be filled completely in about $8/0.06 \approx 133.34$ days and emptied in $8/0.25 = 32$ days. To model gas prices (quoted here in \$/MMBtu) we use a mean reverting process,

$$d \log G_t = 17.1(\log 3 - \log G_t)dt + 1.33dW_t, \quad (6.4)$$

and assume that the initial gas price is $G_0 = 3$. The yearly volatility implied by this model is 133%, while the time to mean reversion is about $365/17.1 \approx 21.3$ days. We assume that the initial gas level in storage is $C_0 = 4$ Bcf and that the storage is managed over a period of one year, $T = 1$. The unit salvage value, i.e. the value of the gas left in storage at time T takes the form of a penalty function which depends on the terminal inventory. Accordingly, the storage operator is charged two times the market price of gas if the final gas inventory is less than the initial inventory C_0 and receives nothing otherwise. Thus, the terminal condition in the risk-neutral case reads

$$V(T, G_T, C_T, i) = -2G_T (C_0 - C_T)^+,$$

while in the risk-sensitive case it is given by

$$V(T, G_T, C_T, i) = -\exp(-2\rho G_T (C_0 - C_T)^+). \quad (6.5)$$

Naturally, this terminal condition motivates the manager of the gas facility to keep the gas inventory at T as close to 4 Bcf as possible to avoid any revenue loss. Finally, we suppose that changing between any of the regimes costs $K_{i,j} = \text{MM}\$0.25, i \neq j$, and that the interest rate is $r = 0.06$ per year.

Numerical implementation

We implement both the RN-LSM algorithm and the RS-LSM algorithm using 200 time steps, $N_p = 2000$ price sample paths and a discretisation of the inventory variable C with $N_c = 80$ grid points. Following the recommendation given on page 187, we run the RS-LSM algorithm with $\rho = 0.05k, k = 1, \dots, 5$, in order to study the effect of the risk-sensitivity parameter ρ on the value function and switching regions. For this example, we completed 12 runs of the algorithm within 30 minutes using a parallel computing architecture.

Tradeoff between expected total revenue and risk measures

First, we investigate the behaviour of the expected total revenue of the gas storage operator and the corresponding risk measures when using different degrees of risk-sensitivity ρ . For increasing values of ρ , Figure 6.5 provides the mean total revenue, its SD (both in the top panel), and its risk measures VaR, CVaR and EVaR (bottom panel), when $t = 0$ and initial regime is ‘store’. The mean total revenue is given with respect to the left (blue) y -axis, while the SD is given with respect to the right (orange) y -axis.

The values in Figure 6.5 were computed by implementing the resulting optimal switching strategies given by the RN-LSM and RS-LSM algorithms in a forward Monte Carlo simulation of 2000 sample paths of the gas price model. The risk indices are computed from the probability distribution of the expected total revenue, all at the 99% confidence level. Thus, in this case, the greater the value indicated by the risk measures the better from the perspective of a risk-averse controller. Similar results (not shown here) were obtained for confidence levels in the range 90% – 98%. We

would like to point out that our results for the risk-neutral case presented in Figure 6.5 are in good agreement with those reported in Carmona and Ludkovski (2010), see their Table 1 and Table 2. Furthermore, as ρ goes to zero the risk-sensitive value function approaches the risk-neutral value function, as is expected from the results in Section 2.2. This serves as a verification step of our numerical implementation and gives confidence in the results presented here.

The results in Figure 6.5 show clearly the existing tradeoff between the mean total revenue and the risk indices associated with the operation of the storage. As ρ goes from 0 to 0.25, the average total revenue realised by the operator falls from 9.57 to 9.08 price units, a percentage change of about 5.2%. However, this is compensated by a decrease in the SD of total revenue and risk measures. The SD decreases by about 9.2% from 4.1 to 3.72 price units, while the price risk, as measured by VaR, CVaR and EVaR, decreased by 0.85, 1.1 and 1.4 price units, respectively.

Note that the risk-neutral optimal control implies the largest risk among all the strategies considered here, while the risk indices CVaR and EVaR attain local minima at $\rho = 0.225$. Nonetheless, the optimal value of ρ , in the sense of the indicator proposed on page 198, is given by $\rho^* = 0.2$ (this value can be computed directly from Figure 6.6, as the ρ where the blue line and the yellow line are separated the most). This value represents the best compromise between loss in average total revenue and reduction in CVaR offered by the risk-sensitive optimisation algorithm. After $\rho = 0.225$, however, all risk measures appear to increase, albeit the SD continues to decrease over some range of ρ . It is worth mentioning that the CVaR, given in the bottom panel of Figure 6.5, is close to zero for ρ^* . Under the assumptions of this example, this means that although the storage operator makes a reduced total revenue when compared to the case $\rho = 0$, she can be 99% confident that the expected losses will not surpass a few dozen thousands USD (\$). In the risk-neutral case the expected financial losses are in the order of one million USD.




Figure 6.5: For the total revenue function of Example 2, its mean, standard deviation (SD) and risk measures VaR, CVaR and EVaR all at the 99% confidence level, as a function of the degree of risk-sensitivity ρ . In this example the greater the risk measures the smaller the risk faced by the plant operator. The results are based on 2000 sample paths.




Figure 6.6: For Example 2, the net change with respect to the risk-neutral case $\rho = 0$ of the mean and risk measures of the total revenue, as the risk-sensitivity ρ increases. A negative value indicates a decrease with respect to the risk-neutral case. The results are based on 2000 sample paths of the gas price model.

Optimal control maps

The previous results are a direct consequence of the optimal control policy obtained with the RN-LSM and RS-LSM algorithms. Figure 6.7 illustrates the optimal switching strategy for $\rho = 0, 0.1, 0.15, 0.2$ as a function of the gas price (x -axis) and the current inventory level (y -axis) at time $t = 0.75$ years, when the current regime is ‘store’. The marker colours in each graph indicate the optimal operating regime for each pair (G_t, C_t) of gas-inventory values: red for regime inject, black for regime store and blue for regime withdraw. For example, when $\rho = 0$ (top-left graph) and the current inventory level is 4 MMBtu, it is optimal to inject gas into the storage if the gas market price is less than 3\$/MMBtu, whereas it is optimal to withdraw gas from the storage if the gas price is more than 4\$/MMBtu.

As shown, in all cases of ρ the optimal switching control can be translated as a ‘buy low - sell high’ trading strategy, whereby one injects gas into the storage when the market gas price is relatively low and the storage is not full, and withdraws gas from the storage if prices are relatively high, provided the storage is not empty. We notice however that the expected cash flows of such strategy must cover the instantaneous operating costs and switching costs. The black region on each graph indicates that there it is not optimal to either inject or withdraw gas, owing to the fixed switching costs. As explained in the previous example, in this no-action region the potential expected revenue that can be obtained by staying in regime ‘store’ is larger than the expected revenue that may be obtained by switching regimes and paying the corresponding switching cost. Consequently the storage operator finds optimal to wait until gas prices change.

The reductions in risk of the previous section are clarified now by inspection of the graphs in Figure 6.7. Observe how the hysteresis band progressively widens, and the red and blue regions become less wide, as ρ increases from $\rho = 0$ to $\rho = 0.225$. The overall effect of this is a decreased frequency in responses to the opportunity costs signalled by the energy market when the system state is within the hysteresis band. Thus, in this case risk-sensitivity acts by taking less aggressive trading positions on the market, implying that a risk-sensitive storage operator will tend to avoid high price volatility earnings as their risk-aversion increases. Doing otherwise would increase the variance and, as we have demonstrated, the risk index of their expected total revenue.

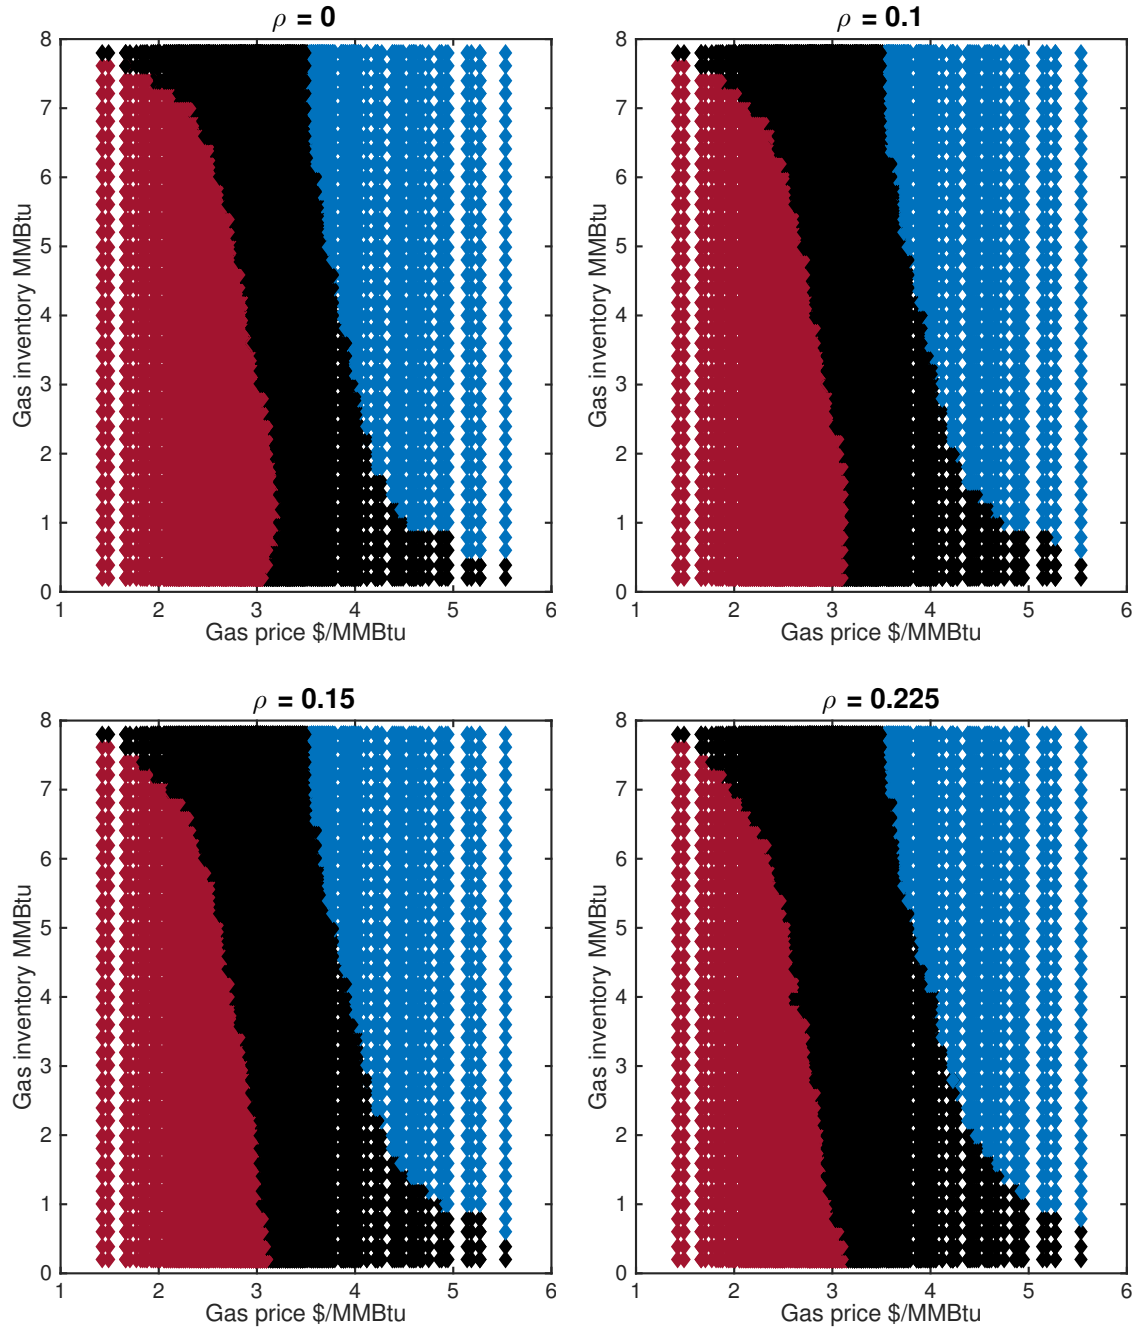


Figure 6.7: Switching regions from regime store obtained for different values of the risk-sensitivity parameter ρ . The marker colours in each graph indicate the optimal operating regime for each pair (G_t, C_t) of gas-inventory values: red for regime inject, black for regime store and blue for regime withdraw. The results are based on 2000 simulations of the gas price process.

We recall here that storage offers a mechanism to exploit large fluctuations in market prices. In particular, the role of a gas storage is to act as a financial straddle option (Carmona and Ludkovski, 2010), whereby the (risk-neutral) owner considers favourable either large positive or large negative price movements in high volatility

markets. In this case it is seen optimal to exploit the flexibility or options of the gas storage, and act quickly to market changes, selling and buying gas as prices cross given threshold values. From a risk-return point of view, however, we have seen that a risk-averse controller (in the words of Whittle (1990)) ‘fears’ that these large price fluctuations will move to his disadvantage. Hence, the controller decides to take a more ‘cautious’ trading position and avoid comparatively fast changes of the gas storage operating regimes.

The effect of risk-sensitivity is increased by considering the final inventory penalty specified by equation (6.5). As time approaches maturity, the optimal control policy is dominated by this terminal condition. Therefore with a risk-sensitive attitude one stops trading sooner than a risk-neutral controller, trying to lock the storage at $C_t = 4$ even when it may be optimal (with a risk-neutral attitude) to take positions on the market. Our simulations also demonstrate (omitted here for clarity) that an ‘extreme’ risk-averse controller would keep the same initial inventory throughout the whole period, as this would minimise their exposure to the high volatility gas market. Obviously this would not generate any profits, and so this is a pessimistic attitude from part of the controller (Whittle, 1990, p. 5).

6.2.2 Example 3. A depleted gas reservoir on the UK market

In the next example we consider a realistic natural gas storage operating on the UK gas market. The gas storage is based on a typical depleted gas reservoir with a large working capacity and volume-dependent withdrawal and injection rates. This type of storage is the most common form of underground storage, offering a large storage value. We consider seasonality in gas prices and the possible occurrence of price spikes. The gas storage data is based on the situation presented by Felix (2012). For consistency with the previous example, the storage capacities are quoted here in Bcf and in yearly units.

Gas storage characteristics

Now we consider a gas storage with a maximum working capacity of $C_{\max} = 44.11\text{Bcf}$, and maximum withdrawal and injection capacities given by $a_{\text{out}}^{\max} = 733.6$ and $a_{\text{in}}^{\max} =$

439.6. The withdrawal and injection rates are volume-dependent and given by

$$a_{out}(c) = a_{out}^{\max} \begin{cases} 0.2 + 1.6 \frac{c}{C_{\max}}, & c \leq 0.5C_{\max} \\ 1 & c \geq 0.5C_{\max} \end{cases},$$

$$a_{in}(c) = a_{in}^{\max} \begin{cases} 1 & c < 0.5C_{\max} \\ 1 - 1.6 \frac{(c - 0.5C_{\max})}{C_{\max}}, & c \geq 0.5C_{\max} \end{cases}.$$

Note that the withdrawal (respectively, injection) rate is proportional (inversely proportional) to the gas inventory level. Thus more gas can be withdrawn when the storage is full than when it is almost empty, and conversely more gas can be injected into the storage when it is almost empty than when it is full. We assume that the gas storage is subject to losses of 1% of the gas volume being extracted or injected.

Gas spot prices

To model gas spot prices $G(s)$ we use the unrestricted 3-OU price model studied in Section 5.11. For convenience, we recall that the model consists of one Gaussian OU process describing ‘normal’ market conditions and two Gamma OU processes capturing positive and negative price spikes,

$$\ln G(s) = f(s) + Y_0(s) + Y_1(s) - Y_2(s). \quad (6.6)$$

The parameter estimates of the seasonal component $f(s)$ and the processes $Y_i(s)$ can be found in Section 5.11.1 and Table 5.15. As mentioned in Chapter 5, this gas price model is able to describe the seasonality present in UK gas spot prices, where relatively high prices are observed in the cold season (October-March). Also, we found that UK gas prices are characterised by a slow speed of mean reversion and slowly-decaying but frequent small jumps. In addition, we demonstrated that the estimated price model resembles closely the statistical properties of the observed data. See Table 5.16.

Costs and payoff functions

We assume the payoff functions (£/year) are given by

$$\begin{aligned} \text{Inject:} \quad & \psi_1(G_t, C_t) = -G_t(a_{in}(C_t) - 0.01C_t) - 0.1C_t \\ \text{Store:} \quad & \psi_2(G_t, C_t) = -0.1C_t \\ \text{Withdraw:} \quad & \psi_3(G_t, C_t) = G_t(a_{out}(C_t) - 0.01C_t) - 0.1C_t, \end{aligned}$$

In these equations the term $0.01C_t$ represents the losses associated with injection and withdrawal of gas from the gas storage. Also, note that we are assuming direct storage costs of $0.1C_t$ £/year for each regime. We take the same terminal condition of Felix (2012), namely,

$$V(T, G_T, C_T, i) = \bar{G}_T C_T,$$

where \bar{G}_T is the expected gas price at maturity T . Hence the storage operator receives the average market price of gas for any gas unit left in storage at time T . Additionally, we assume that the storage is operated over one year $T = 1$, that switching decisions are taken daily, so $N_t = 365$, and that $K_{i,j} = \text{MM}\pounds 0.25, i \neq j$. The RN-LSM and RS-LSM algorithms are implemented with $M = 2000$, $N_c = 80$, and assuming an interest rate of $r = 0.06$ on an annual basis.

At this stage, we point out that our RS-LSM algorithm requires to work with a Markov process. Hence, we cannot use realisations of the price process $G(s)$ in equation (6.6) since it is not Markov. Instead we consider the Markov process $(Y_0(s), Y_1(s), Y_2(s))$, and define three-dimensional basis functions $B_m(x, y, z)$ for the regression step. The rest of the algorithm remains unchanged.

Tradeoff between expected total revenue and financial risk

For increasing values of ρ , Figure 6.8 provides custom box plots of the total revenue function over the entire operating period when the initial regime is ‘store’, the storage is working at half capacity and the initial gas price is 1.57 £/Bcf. The two lower whiskers below each box indicate the risk measures $\text{CVaR}_{95\%}$ and $\text{CVaR}_{99\%}$. The red line is the median, the edges (top and bottom) of the box are the 25th and 75th percentiles, the upper whisker is the highest datum that is not an outlier, i.e. it is still within 1.5IQR (interquartile) of the upper quartile. The results are obtained by implementing the

resulting optimal switching controls in a forward Monte Carlo simulation with 2000 simulated gas price paths.

Figure 6.8 illustrates how the distribution of the total revenue changes with increasing risk-sensitivity from part of the storage controller. As shown, the average total expected revenue and its median decreases with ρ . Nonetheless, we also observe that the risk measures CVaR (at the 95% and 99% confidence levels) have an increasing relationship with ρ , which means that the price risk is decreasing as we increase the degree of risk-sensitivity. This tradeoff however is not only between expected revenue and risk. As shown, the standard deviation (note the decreasing spread of the box plots), skewness and kurtosis of the total revenue distribution also decrease with ρ . On the one hand, this means that the probability of receiving extreme gains and/or extreme losses associated with the storage operation is reduced on average by taking higher values of ρ . From a risk-averse perspective, this is advantageous since the risk-sensitive strategies are limiting the probability of extreme financial losses. On the other hand, we notice that a risk-sensitive controller tends to prefer a sure although smaller average return from their control strategy, than a higher but more variable return. This specific tradeoff between risk and total revenue is made precise by the results in Figure 6.8. For a given level of risk, the gas storage manager receives on average the total revenue specified there.

It is worth mentioning that the balance between average revenue and price risk may be taken in such way that the reductions in the latter are greater than those in the former. Indeed, as shown in Figure 6.9, the expected total revenue decreases more quickly than both $\text{VaR}_{99\%}$ and $\text{CVaR}_{99\%}$, with the best tradeoff achieved when $\rho^* = 0.072$. With this degree of risk-sensitivity, the reductions in average total revenue is about 4.9MM£, while the expected losses at the 99% confidence level (if there are any) decrease by 15.8MM£. This results can be compared with the risk optimisation approach in Felix (2012). There, the expected return of the storage manager decreases faster than both the standard deviation and the $\text{VaR}_{99\%}$, and consequently the $\text{CVaR}_{99\%}$. Our results show that it is possible to achieve considerable reductions in risk with little loss in total expected revenue.




Figure 6.8: Custom box plots of the total revenue function of Example 3. The two lower whiskers below each box indicate the risk measures $\text{CVaR}_{95\%}$ and $\text{CVaR}_{99\%}$. The red line is the median, the edges (top and bottom) of the box are the 25th and 75th percentiles, the upper whisker is the highest datum that is not an outlier, i.e. it is still within 1.5IQR (interquartile) of the upper quartile.




Figure 6.9: Change with respect to the risk-neutral case $\rho = 0$ in the expected total revenue (mean), and risk measures $\text{VaR}_{99\%}$ and $\text{CVaR}_{99\%}$ as ρ increases. A negative value means that the corresponding index has decreased.

Optimal switching regions

To illustrate the effect of seasonal gas prices on the optimal switching boundaries, we fix the inventory level dimension and plot the control policy in the time-price space. Accordingly, when the inventory level is fixed and equal to 30 Bcf, Figure 6.10 provides the optimal switching control from regime ‘store’ as a function of gas prices over the period of one year, starting on 1st April. The top panel presents the risk-neutral case, whereas the bottom panel presents the risk-sensitive case with $\rho = 0.072$. The convention for the colour markers is the same as before: red for regime inject, black for store and blue for regime withdraw. Our empirical analysis of UK gas prices shows that prices tend to be relatively high during the cold season (October-March) and relatively low in the warm season (April-September). This is quantitatively revealed in the fitted seasonal component of $f(s)$, which attains local maxima during the cold season and local minima during the warm season. This seasonal cycle is clearly reflected in the optimal switching regions for both the risk-neutral case and the risk-sensitive case. Furthermore, in the risk-neutral the hysteresis band (black region) seems to be centred on the time-dependent price mean level $f(s)$ over the entire year, showing the strong dependency of the control policy on time. Thus, in contrast with the previous examples the optimal switching control in this example is non-stationary.

The optimal control strategies are also exploiting the high volatility of gas prices and the fact that they mean-revert to $f(s)$ over time, allowing to buy-low and sell-high. Note, however, that seasonality in the prices is not the only driving element for the optimal control of the storage. That is to say, only buying in the warm period and selling in the cold season is a sub-optimal strategy. The relative fast injection and withdrawal capacity rates of the storage unit allows to exploit large changes in prices. Hence, as time progresses, the control policy is adjusted in such way that the maximum buy price as well as the minimum sell price increase or decrease, taking into account the seasonal market opportunities.

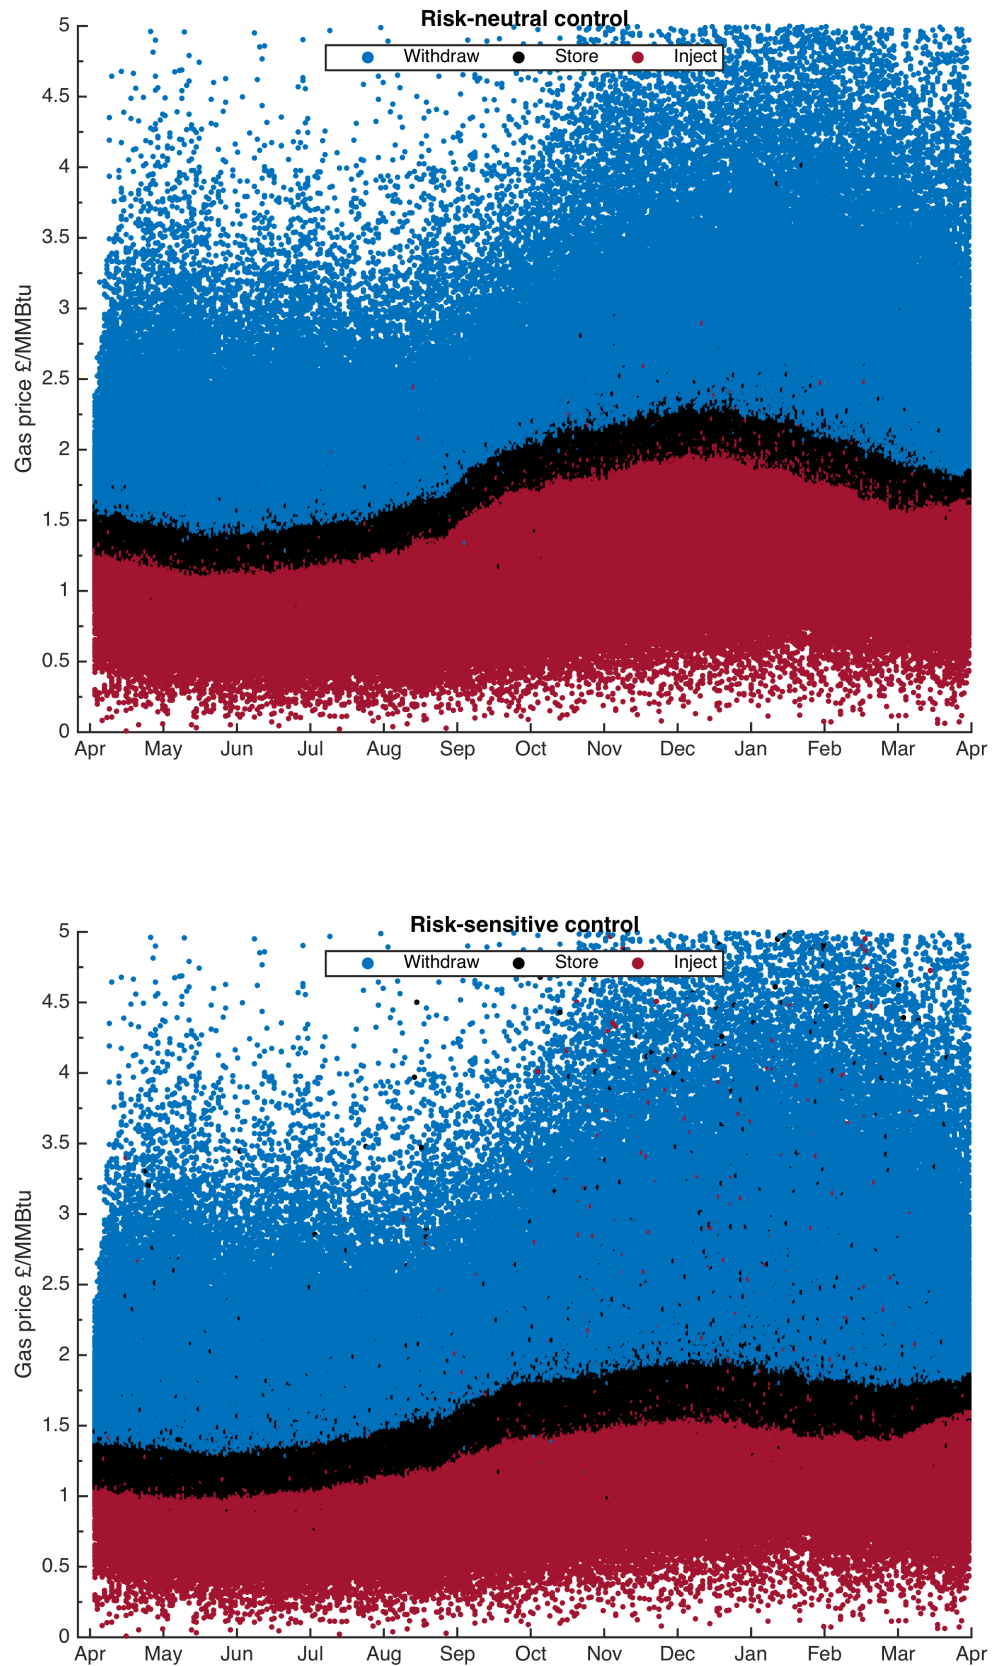


Figure 6.10: For Example 3, the switching regions from regime 'store' as function of time and gas price when the gas inventory level is fixed and equal to 30 Bcf. The top panel presents the risk-neutral case, whereas the bottom panel presents the risk-sensitive case with $\rho = 0.072$.

The significant reductions in risk shown above are only slightly perceived in the optimal switching maps in Figure 6.10. Although superficially, they are quite similar, the risk-sensitive optimal boundaries are actually shifted downwards relative to the risk-neutral ones. This is revealed by carefully observing the black region in both cases relative to the reference gas price 1.5£/MMBtu. It can also be observed from Figure 6.10 that, when compared to the risk-neutral case, the hysteresis band in the risk-sensitive case is thicker and the curves of the switching boundaries are less pronounced (i.e. they are flatter).

Thus when $\rho = 0.072$ the optimal switching strategy from regime store injects and withdraws gas at lower prices, and also tends to stay longer in the non-action region until changing market prices, relative to the case $\rho = 0$. This risk-attitude is a combination of that found in Example 1 and in Example 2, except that here the switching boundaries are shifted downwards and not upwards in contrast with Example 1.

6.3 Flexible district energy systems

Our final example concerns the optimal operation of a multi-generation plant in a district energy system. The plant consists of a Combined Heat and Power (CHP) plant and a gas boiler acting as back-up. Using gas as fuel, the CHP produces heat and electricity through a conversion mechanism, while the gas boiler only produces heat. A new feature appearing in this example is the constraint that a given heat and electricity demand profile must be met partly with the output of both the CHP and the boiler. Any electricity deficit must be covered by buying the required power at the electricity market. This implies that the system controller may need to temporarily rely on the (liberalised) electricity market, which typically is much more volatile than the gas market.

In the context of risk-neutral optimal switching, this problem has been recently studied by Kitapbayev et al. (2015). See also Kitapbayev et al. (2013). In Gonzalez and Moriarty (2014), we have already applied our risk-sensitive approach on this example, using sophisticated electricity and gas price models suitable for short-term operation. However, the above papers do not consider the spiky behaviour of energy prices and the

implications it has on the optimisation problem. Our aim here is to extend these results and make a more detailed analysis of our risk-sensitive approach to operating multi-generation plants. Furthermore, since the district energy system relies on electricity and gas markets, this example creates an opportunity to demonstrate the potential benefits of our risk-sensitive optimisation algorithm in combination with the multi-factor jump diffusion models for electricity and gas studied in Chapter 5. Thus, we aim at bringing together the novelties presented in the thesis, showing that our risk-sensitive control methodology is able to learn the jump dynamics of energy prices and thus output appropriate control strategies for energy district systems.

In this example, risk is related to the total operational costs that can be incurred by running the system under stochastic energy prices. Thus, our optimisation objective is to strike a balance between expected operational costs and the associated price/cost risk, while capturing the interplay between uncertainty and operational flexibility of the system. The example below is taken directly from Kitapbayev et al. (2015).

6.3.1 Example 4. District energy systems

The district energy system, illustrated schematically in Figure 6.11, consists of a CHP unit, gas boiler and thermal storage. Local heat demand must be satisfied by the CHP and gas boiler, while local electricity demand must be satisfied by the CHP and the electricity market. As in Kitapbayev et al. (2015) the electricity and heat demand (Figure 6.12) are modelled as known deterministic functions of time; we consider the winter period for the purposes of illustration. The thermal storage capacity is 15 MWh with standing heat losses of 0.5% per hour and the conversion efficiency of the CHP unit is $e = 0.4$ to electricity and $h_1 = 0.4$ to heat, whereas the efficiency for the gas boiler is $h_2 = 0.91$.




Figure 6.11: The smart district energy system studied in Example 4.




Figure 6.12: The winter heat and electricity demand profiles of Example 4.

Operational flexibility is modelled by allowing the input gas rates for both the CHP unit and the gas boiler to be chosen from a given finite set of states. In particular, we assume three possible gas input rates for the CHP unit, given by

$$C_j = 3.75j \text{ MW}, j = 0, 1, 2,$$

and five for the gas boiler, given by

$$B_k = 1.65k \text{ MW}, k = 0, 1, \dots, 4.$$

An operating regime is then composed of two gas input rates, one for the CHP unit and one for the gas boiler, making in total 15 operating regimes. For example, the regime C_2B_4 indicates that both the CHP unit and the boiler operate at maximum capacity. Although using more regimes is in principle straightforward to implement,

having this relatively limited number of regimes nevertheless allows us to illustrate numerical results for risk-sensitive control and to compare these with results under risk neutrality. For consistency with the notation of Chapter 3, we note that the (apparent) 2-dimensional control space $\{0, 1, 2\} \times \{0, 1, 2, 3, 4\}$ can be mapped onto the one-dimensional space $\{1, 2, \dots, 15\}$ in an obvious way. Let $u_s^{(1)} \in \{0, 1, 2\}$ and $u_s^{(2)} \in \{0, 1, 2, 3, 4\}$ represent the input gas rates for the CHP and the gas boiler, respectively. Then the operating regime at time s can be written as

$$u_s = 1 + u_s^{(1)} + 3 \cdot u_s^{(2)},$$

taking the convention that the regime ‘1’ correspond to C_0B_0 , the regime ‘2’ to C_1B_0 , and so forth. Note that u_s is uniquely determined by $u_s^{(1)}$ and $u_s^{(2)}$ and that the map is bijective. We suppose that the thermal storage unit is initially empty and that the change in storage level over any period is equal to the total heat output from the boiler and the CHP unit minus heat demand and standing losses. Thus the heat inventory level H_s satisfies the ODE

$$dH_s = (h_1 C_{u_s^{(1)}} + h_2 B_{u_s^{(2)}} - 0.05 H_s - D_s^h) ds,$$

with D_s^h denoting the heat demand level at time s , $H_s \in [0, 15]$ MW and $H_0 = 0$. Access to the electricity and gas spot markets is assumed to be through a network connection, with a fixed network charge of 50£/MWh for electricity and 15£/MWh for gas, so that the buy-price for each vector equals the spot wholesale price (whose stochastic dynamics are described below) plus the corresponding network charge. Any surplus electricity generated by CHP is sold back to the market at the wholesale price. The total operational cost for this system is equal to the cost of buying gas minus the proceeds from selling back surplus electricity over the specified planning period. Accordingly, the payoff functions are given by

$$\begin{aligned} \psi_{u_s}(E_s, G_s) &= C_{u_s^{(1)}} (G_s + 15) + \left(D_s^e - e_1 C_{u_s^{(1)}} \right)^+ (E_s + 50) \\ &\quad - \left(e_1 C_{u_s^{(1)}} + D_s^e \right)^+ E_s + B_{u_s^{(2)}} (G_s + 15), \end{aligned}$$

where E_s and G_s are the electricity and gas prices, respectively, and D_s^e is the electricity

demand level, all at time s . Here $(x)^+ := \max(0, x)$.

Electricity and gas spot price models

The spot prices for electricity and gas (in hourly units) are modelled by (the superscripts e and g refers to electricity and gas, respectively)

$$\begin{aligned} E(s) &= e^{f^e(s)}(Y_0^e(s) + Y_1^e(s)) \\ \log G(s) &= f^g(s) + Y_0^g(s) + Y_1^g(s), \end{aligned} \tag{6.7}$$

where Y_0^e and Y_0^g satisfy the SDE

$$dY_0(s) = \lambda_0^{-1}(\mu - Y_0(s))ds + \sigma dW(s),$$

and Y_1^e and Y_1^g satisfy

$$dY_1(s) = -\lambda_1^{-1}Y_1(s)ds + dL(s),$$

where W^g and W^e are two correlated Wiener processes, and L^g and L^e are independent compound Poisson processes each with Exponentially distributed jump sizes with mean β^e and β^g . The functions f^e and f^g capture the linear trend and seasonal component of energy prices, and are given by

$$\begin{aligned} f^e(s) &= a_1 + a_2s + a_3 \sin(2\pi s/p_1) + a_4 \cos(2\pi s/p_1) + a_5 \sin(4\pi s/p_2) + a_6 \cos(4\pi s/p_2), \\ f^g(s) &= a_1 + a_2s + \cos(2\pi(s - a_3)/p_2). \end{aligned}$$

The district energy system is operated on an hourly basis. Thus, we let $p_1 = 24$ h (hours) to model the intra-day pattern of electricity prices, and $p_2 = 24 \times 365 = 8760$ h to model the seasonal yearly cycle. Due to the storability of gas, gas prices exhibit smoother fluctuations than electricity prices, and hence we do not model any intra-day effects on gas prices but only a possible seasonal term over the period of one year.

We use the MCMC estimates of the model parameters given in Table 6.2 and the estimated coefficients a_j provided in Table 6.3. These values represent the UK electricity and gas hourly spot prices over the period from January 1, 2009 until July 31,

2012. The estimated correlation between W^g and W^e is equal to 0.43. Notice the relatively low volatility of gas prices with respect to electricity prices, and the long-term dynamics of gas prices. It is also interesting that prices spikes are much more pronounced and frequent in the electricity market than in the gas market. Nevertheless, electricity spikes occur one every five weeks on average, while gas price spikes occur one every four months. Thus, price spikes represent the high-impact low probability events we referred to in Chapter 1 and in Section 2.1. As shown below, all these price characteristics shape the optimal switching strategy for the multi-generation plant of the present example.

Parameter	Electricity	Gas
μ	0.9822	-0.01
λ_0	11.789	43.037
σ	0.0978	0.03
λ_1	0.5403	2.4389
η	0.0083	0.0025
β	0.9084	0.2115

Table 6.2: Parameter estimates in hourly units of the spot price model parameters in Example 4.

	a_1	a_2	a_3	a_4	a_5	a_6
Electricity	3.53563	1.10E-05	-0.11409	-0.16882	0.00688	0.03023
Gas	2.903	4.88E-6	-.059505	-384.14	- -	- -

Table 6.3: Estimated coefficients of the mean price level functions f^e and f^g in Example 4.

Numerical implementation

We implement the RN-LSM and RS-LSM algorithms with $M = 5000$ sample paths, $N_c = 80$, and using a horizon of $N_t = 24 \times 3$ hours starting at 00:00 on January 1, 2009. We remark that the process $(E(s), G(s))$ does not have the Markov property, as required by the RS-LSM algorithm. Hence, we consider the process $(Y_0^e(s), Y_1^e(s), Y_0^g(s), Y_1^g(s))$, and use 4-dimensional basis functions in order to account for each factor $Y_i^e, Y_i^g, i = 0, 1$. To explore the effect of ρ on the optimal switching policy, we take $\rho = 0.02k, k = 0, \dots, 15$. Our numerical implementation completed 12

runs of the algorithm in 81.6 seconds on average.

Expected total operating costs versus risk measures

For the total cost function associated with running the district energy system, Figure 6.13 shows its mean, SD and risk measures $\text{VaR}_{99\%}$, $\text{CVaR}_{99\%}$ and $\text{EVaR}_{99\%}$, as a function of ρ . In this example the smaller the value of the risk measures the smaller the risk with respect to the operational costs. In the top panel of Figure 6.13, the mean is given with respect to the left y -axis, whereas the SD is given with respect to the right y -axis. The results illustrate the usual tradeoff between expected performance and risk that we have already seen in the previous examples. Since in this example we are dealing with a cost function, we see that the risk-sensitive criterion produces control policies giving a larger expected total cost relative to the risk-neutral case. Nonetheless, this is compensated by reductions in the risk indices, as shown in the bottom panel of Figure 6.13. The relative change in mean and risk indices with respect to the risk-neutral strategy is shown in Figure 6.14. Under the given assumptions, we find that a good compromise between expected operating costs and price risk is given with the control strategy for $\rho^* = 0.15$ (recall equation (6.3)). We also notice from the figure that substantial benefits are obtained by large reductions in the SD with increasing ρ .

Optimal switching boundaries and control maps

The optimal switching control for the multi-generation plant is a function of time, demand levels, energy prices (given by four price factors) and level of stored heat. For illustration purposes, we fix time and storage level. Figure 6.15 provides the optimal operating states for the gas boiler (top panels) and the CHP (bottom panels) obtained with the risk-neutral control (left) and risk-sensitive control for $\rho^* = 0.15$ (right), at 9am on 2nd January assuming that the current level of stored heat is 4.5MWh. The results are given as a function of the total electricity price E_s (x -axis) and the total gas price G_s (y -axis). The marker colours on each plot indicate the optimal input gas rate for each component, i.e. each colour represents the operating states C_i and B_j separately, with a more intense colouring indicating higher gas input rates. We recall that an operating regime for the district energy system is made of the pairs $C_i B_j$.




Figure 6.13: For the total cost function of Example 4, its mean, standard deviation (SD) and risk measures VaR, CVaR and EVaR all at the 99% confidence level, as a function of the degree of risk-sensitivity ρ . The smaller the value of the risk measures the smaller the associated risk. The results are based on 5000 sample paths.




Figure 6.14: For Example 4, the net change with respect to the risk-neutral case $\rho = 0$ of the mean and risk measures of the total revenue, as the risk-sensitivity ρ increases. A negative value indicates a decrease with respect to the risk-neutral case. The results are based on 5000 sample paths.

The results displayed in Figure 6.15 show that the risk-neutral strategy (graphs on the left) for the gas boiler consists of shifting heat demand through time, whereas the CHP unit exploits the variations in the spark spread between electricity and gas (cf. Kitapbayev et al. (2015)). Indeed, the top-left graph shows that the gas boiler is mainly used when the gas price is below its average. This amounts to an increased level of stored heat when gas prices are relatively low, which is then used to cover heat demand in times of relatively high gas prices. Also, by fixing the gas price on the bottom-left graph and increasing the electricity price, we observe that a higher gas input rate for the CHP unit is used. This means that the optimal strategy for the CHP is to convert the relatively cheap gas into expensive electricity, effectively exploiting the spark spread through the energy conversion process. The linear switching boundaries in this graph are upward-sloped due to the positive correlation between the stochastic energy vectors (gas and electricity).

Inspection of the risk-sensitive strategy (graphs on the right of Figure 6.15) reveals noticeable differences with the risk-neutral strategy. Looking at top-right graph it can be seen that the gas boiler is now used more heavily when the gas price is above its average, indicating a preference to rely on the less volatile gas market than on the highly volatile electricity market (recall our estimates of price volatility in Table 6.2). In addition, the linear switching boundaries of the risk-neutral strategy for the CHP are now replaced by curved boundaries. Surprisingly, the region covering the half-capacity gas input rate (3.75 MW) has a ‘swoosh’ shape: as shown, when electricity prices are comparatively high it is not optimal to inject gas into the CHP unit. This is explained by noting that the boiler and the CHP, whose control strategies are plotted here separately for clarity, work together towards meeting the specified heat and electricity demand profile. If there is any surplus electricity, it must be sold into the electricity market, whose high volatility induces high risk and variance in the total operating cost. Thus the boiler is used more heavily than the CHP unit in order to avoid generating too much surplus electricity. This control strategy also compensates for the corresponding reduced heat output level from the CHP.

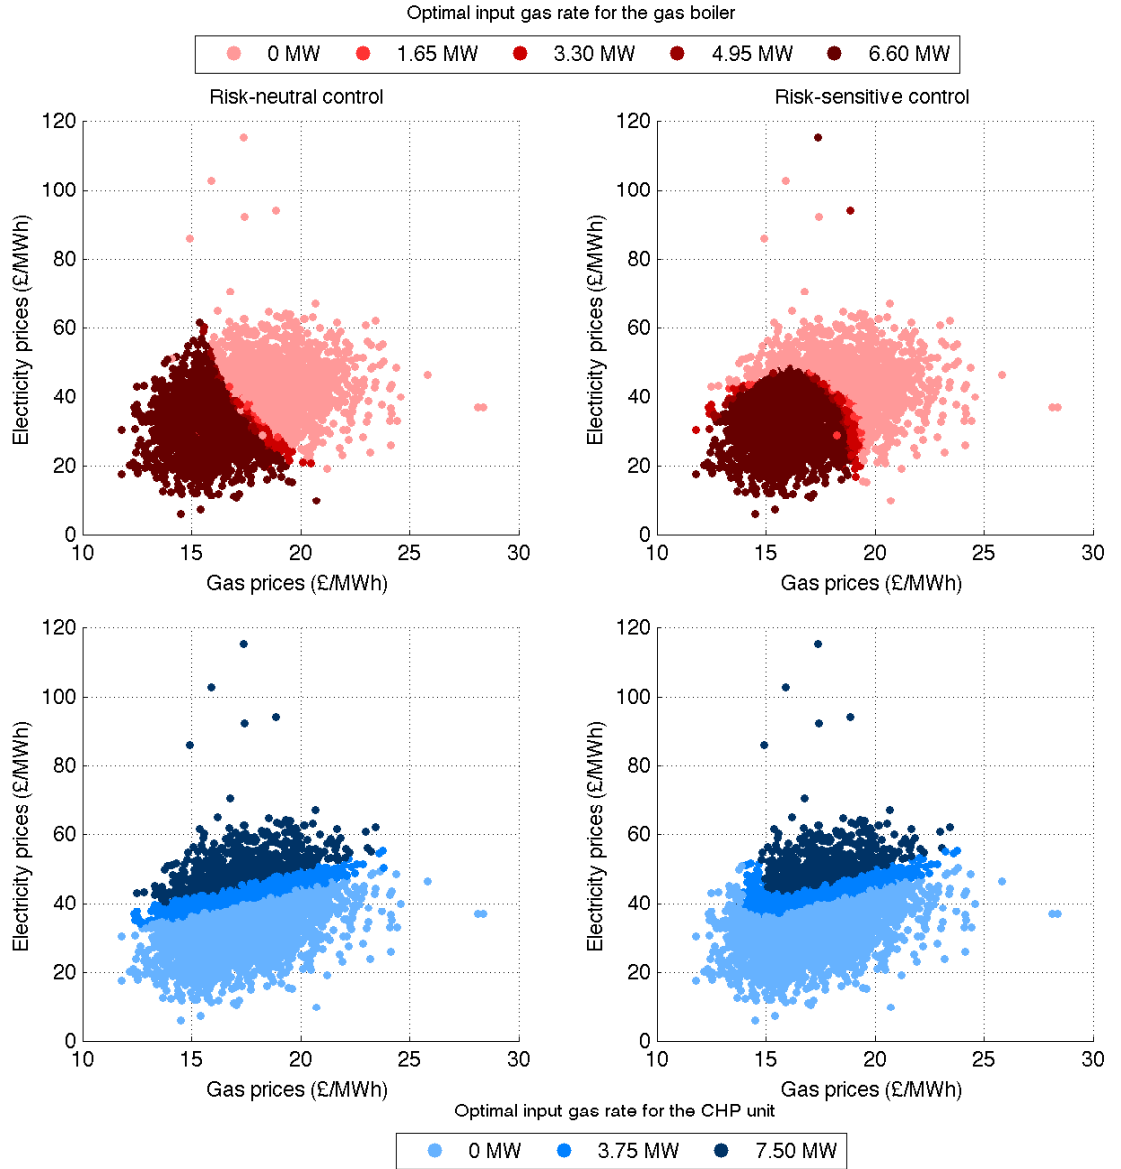


Figure 6.15: Optimal operating regimes for the gas boiler (top) and the CHP unit (bottom) at 9am on 2nd January, using risk-neutral control (left) and risk-sensitive control with $\rho = 0.15$ (right). The current level of stored heat is 4.5 MWh.

Dependence of the value function on the price factors

Price spikes can have a significant effect on the optimal control strategies, and hence on the value function for the district energy system problem. To understand this effect, we visualise the value function $V(t, (Y_0^e, Y_1^e, Y_0^g, Y_1^g), C_t, i; \rho)$ for the district energy system problem by fixing the time variable, the level of stored heat and either the total gas price $G(s)$ or the total electricity price $E(s)$. Accordingly, in the results below, we suppose the time to maturity is 24 hours, the inventory level is $C_t = 0$ and

the current operating regime is $i = C_0 B_0$. Additionally, we notice that it is enough to consider the case $\rho = 0$. This is because the value function reports only the expected total costs under different strategies. Hence, knowledge of one case allows us to draw conclusions for all other cases, noting that the value function increases monotonically with ρ .

Let us focus first on the effect of the gas price factors on the value function. The top panel of Figure 6.16 presents a contour plot of V as function of the gas price factors Y_0^g and Y_1^g when $Y_0^e = 1, Y_1^e = 0$. Note that in this case the electricity price level is about its long-term average. Unsurprisingly, the value function has an increasing relationship with both Y_0^g and Y_1^g when the total electricity price is fixed. This is due to the fact that gas is the main fuel used to operate the district energy system. Thus, when the total gas price $G(s)$ is high the expected total costs of satisfying the demand profile is high, and conversely, when gas prices are low the expected total costs are low.

We notice, however, that the value function has more sensitivity to values of Y_0^g away from its long term mean (which is equal to zero) than to the spike component Y_1^g . Indeed, keeping the other variable fixed, an increase in Y_0^g of 0.5, say, increases the value function much more than an equal increase in Y_1^g . This is a consequence of the relative slow speed of mean reversion of ‘normal’ gas market prices ($\lambda_0^g = 43$), accompanied with a fast speed of mean reversion of gas price spikes ($\lambda_1^g = 2.5$). Recall Table 6.2. Thus, when Y_0^g is large, gas prices will tend to remain large for a long time, implying a high expected total cost. In contrast, when the current gas price is dominated by a large spike in Y_1^g , prices will tend to revert back to the mean level more quickly, which results in a comparatively small expected total cost.

We would then expect that if normal gas prices reverted back in a shorter time, then the value function had less sensitivity to the Y_0^g component. To confirm this, in the bottom panel of Figure 6.16, we have modified the speed of mean reversion of Y_0^g by setting $\lambda_0^g = 5$, so that ‘normal’ gas prices revert back much faster and at double the rate of price spikes. As shown, the value function still has an increasing relationship with each gas price factor. This time, however, the sensitivity of the value function has decreased in the Y_0^g component, confirming our intuition about the effect of the speed of mean reversion parameters. A similar effect is obtained by increasing λ_1^g so

that the ratio of mean reversion parameters is also about 2. Notice that this analysis carries on when the total gas price $G(s) = \exp(f^g(s) + Y_0^g(s) + Y_1^g(s))$ is kept fixed. Indeed, along the planes $Y_0^g + Y_1^g = c \in \mathbb{R}$, the weight that either Y_0^g or Y_1^g has on the value function is significantly affected by the ratio of the mean reversion parameters.

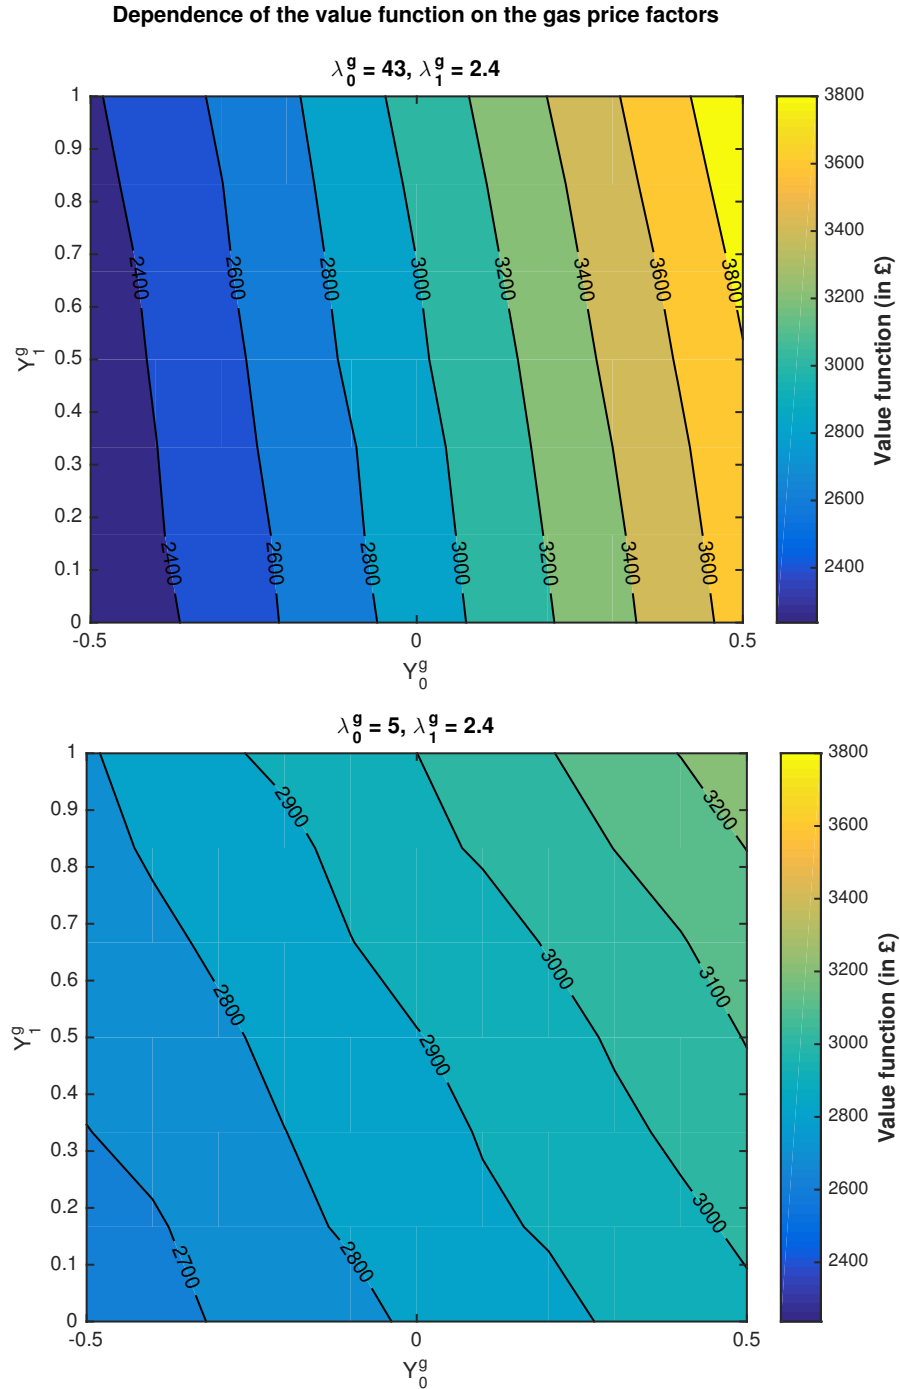


Figure 6.16: Contours of the value function $V(t, (Y_0^e, Y_1^e, Y_0^g, Y_1^g), C_t, i; \rho)$ for Example 4 as function of the gas price factors Y_0^g and Y_1^g when $Y_0^e = 1, Y_1^e = 0$. The time to maturity is 24 hours, $\rho = 0$, the inventory level is $C_t = 0$ and the current operating regime is $i = C_0 B_0$.

The dependence of the value function on the electricity price factors, Y_0^e and Y_1^e , is shown in the top panel of Figure 6.17. We recall that the optimal control strategies use the operational flexibility of the CHP to exploit the spark spread between electricity and gas. As a consequence, the figure shows that the value function has a decreasing relationship with each electricity price factor Y_0^e and Y_1^e for fixed total gas price. Formally, any of these price factors increase the total electricity price $E(s)$. Hence, as any of them increases the proceeds obtained from selling any excess electricity output from the CHP also increase. The overall effect of this is that the expected total costs of operating the district energy system decreases.

We have modified the electricity spike component by increasing λ_1^e from 0.5 to 6 hours, so that now spikes decay much slower than before. The results are shown in the bottom panel of Figure 6.17. Again, we notice that the speed of mean reversion parameters have an important effect on the value function. Since in this situation electricity prices will tend to stay high for longer whenever a spike occurs, we observe that the value function has considerably decreased with respect to the previous case, especially when the prices are dominated by a price spike. Additionally, we note that the value function has an increased sensitivity in the Y_1^e component. Under the given assumptions, we can then conclude that the presence of electricity spikes is advantageous when gas prices remain relatively low, since this presents economic opportunities that help to minimise the total expected cost. Importantly, both the RN-LSM and RS-LSM algorithms are able to learn the dynamics of the electricity and gas price spike components in order to output appropriate risk-neutral and risk-sensitive strategies for the district energy system.




Figure 6.17: Contours of the value function $V(t, (Y_0^e, Y_1^e, Y_0^g, Y_1^g), C_t, i; \rho)$ for Example 4 as function of the electricity price factors Y_0^e and Y_1^e when $Y_0^g = 0, Y_1^g = 0$. The time to maturity is 24 hours, $\rho = 0$, the inventory level is $C_t = 0$ and the current operating regime is $i = C_0 B_0$.

Chapter 7

Conclusions and future work

In this thesis we have investigated two distinct but related problems concerning the risk management of energy assets operating in liberalised energy markets. In the first part of the thesis, we have proposed a novel methodology to produce managerial strategies for energy assets, taking into account the energy price risk associated with uncertain and highly volatile markets. From a theoretical perspective, our optimal control methodology generalises the class of optimal switching problems to the case of risk-sensitive criteria. The performance index of the optimisation problem involves a controlled diffusion and an exponential function of the total revenue and the switching costs. We recast the risk-sensitive optimal switching problem as an iterative optimal stopping problem for controlled diffusions. We then construct and characterise optimal risk-sensitive control strategies employing the theory of Snell envelopes. Our main result demonstrates that the value function of the risk-sensitive control problem satisfies a multiplicative version of the dynamic programming equation for optimal switching problems. Based on this equation, we develop numerical methods for solving risk-sensitive optimal switching problems. The methods combine Monte Carlo simulations, alongside regression-based approximations of conditional expectations computed with piecewise linear basis functions. Additionally, we show how to improve the convergence rate of our algorithms by employing Quasi-Monte Carlo techniques based on low-discrepancy sequences.

In the second part of the thesis, we propose a unified framework for the parametric estimation of energy spot price models based on mixtures of Gaussian and

non-Gaussian OU processes. This class of models have become established in modelling energy spot prices, offering sufficient analytical tractability for derivatives pricing and great flexibility to capture price dynamics. As opposed to current estimating techniques for this class of models, our estimating technique is based on a Bayesian paradigm, aided with the use of MCMC algorithms and data augmentation schemes. The robustness, accuracy and efficiency of our proposed MCMC algorithms have been extensively demonstrated on simulated data. Furthermore, our analysis of the UK and German energy markets demonstrates that the MCMC-estimated models are able to capture not only long- and short-lived positive price spikes, but also short-lived negative price spikes. An ability of our methodology is that the proposed diagnostic tool for goodness-of-fit, which is based on Bayesian p -values, allows to determine a suitable minimum number of components in the mixture model. With this tool, we have shown that models based on two OU components may be inadequate to capture relevant features of energy prices. Clearly, understanding and suitably modelling energy price dynamics have a direct impact on risk management systems. Therefore we believe that the methods developed in the thesis can be used widely outside the context in which we work here.

The final part of the thesis brings together our risk-sensitive approach to managing energy assets and the MCMC-estimated price models. Here, we have strived to capture the interplay between risk, uncertainty, flexibility and performance in various applications to energy systems, including power plants, energy storage and district energy systems. In all the examples, we have compared the classical risk-neutral approach to our novel risk-sensitive approach. Also, we have shown how risk attitudes change with different degrees of risk-sensitivity. Our findings can be summarised as follows. On the one hand, our risk management methodology offers a tradeoff between average performance (e.g. total revenue or total operating costs) and price risk. Indeed, we have seen that by setting a suitable degree of risk-sensitivity, although the average performance decreases, modern measures of risk such as VaR, CVaR and EVaR all decrease. Additionally, it is also possible to find an appropriate degree of risk-sensitivity which provides a best tradeoff between performance and risk. On the other hand, all our examples illustrate that this tradeoff can be achieved in such a way that the reductions in risk outweigh the loss in average performance. In other words, by paying

a relatively small ‘risk insurance’, we are able to obtain relatively large reductions in risk, making us to consider more carefully the classical risk-neutral approach.

An interesting result obtained in the district energy systems example is that price spikes are not always a source of risk. Instead, they might be beneficial from an economical standpoint to some market players. As demonstrated in this example, one can use the spark spread to arbitrage the market by buying cheap gas (if available) and selling very expensive electricity. However, we need to bear in mind that correlation between fuel prices, particular dynamics of energy markets and constraints of the power system need to be taken into account when exploiting these arbitrage opportunities.

Before concluding, we would like to mention some possible avenues of future research. One of them is the extension of the RS-LSM algorithm for the general case of controlled processes, i.e. when both the drift and the diffusion matrix depend explicitly on the switching control. This might prove useful, for instance, when dealing with markets with limited liquidity. In this case, the managerial strategies of some market players might influence energy prices, and thus we cannot longer assume that energy asset managers are price takers.

Another aspect relates to using alternative risk-criteria in the optimisation problem. Our risk-sensitive control approach minimises directly the entropic risk measure, and indirectly, the risk measures VaR, CVaR and EVaR (under suitable assumptions). It would be interesting to see how the methodology and algorithms can be modified to minimise directly multi-objective optimisation criteria based on expectations and coherent risk measures such as the CVaR and EVaR.

Regarding price models and MCMC algorithms, we believe that a number of features can be incorporated depending on the specific market under study. For example, one could be interested in stochastic volatility and/or stochastic mean reversion levels. Additionally, we might consider seasonal dynamics in the underlying Poisson processes and alternative specifications to jump size distributions. This might provide additional flexibility to price modelling, albeit estimating routines for these more complex models would need to be carefully designed for efficiency and computational performance. In this regard, we hope we have achieved a good balance between model adequacy and parsimony, providing price models that capture the key statistical features of energy prices, yet with an estimating technique appropriate for practical applications.

Bibliography

- Ahmadi-Javid, A. (2012a). Addendum to: Entropic Value-at-Risk: A New Coherent Risk Measure. *J Optim Theory Appl*, 155(3):1124–1128.
- Ahmadi-Javid, A. (2012b). Entropic Value-at-Risk: A New Coherent Risk Measure. *J Optim Theory Appl*, 155(3):1105–1123.
- Artzner, P. (1999). Application of Coherent Risk Measures to Capital Requirements in Insurance. *North American Actuarial Journal*, 3(2):11–25.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228.
- Asri, B. E. and Hamadene, S. (2009). The Finite Horizon Optimal Multi-Modes Switching Problem: The Viscosity Solution Approach. *Appl Math Optim*, 60(2):213–235.
- Aïd, R., Campi, L., Langrené, N., and Pham, H. (2014). A Probabilistic Numerical Method for Optimal Multiple Switching Problems in High Dimension. *SIAM J. Finan. Math.*, pages 191–231.
- Bally, V. and Pagès, G. (2003a). Error analysis of the optimal quantization algorithm for obstacle problems. *Stochastic Processes and their Applications*, 106(1):1–40.
- Bally, V. and Pagès, G. (2003b). A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. *Bernoulli*, 9(6):1003–1049.
- Barndorff-Nielsen, O. E. and Shephard, N. (2001). Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):167–241.

- Barrieu, P. and El Karoui, N. (2008). Pricing, hedging and optimally designing derivatives via minimization of risk measures. In Carmona, R., editor, *Indifference Pricing: Theory and Applications*. Princeton University Press, Princeton, USA.
- Barty, K., Girardeau, P., Strugarek, C., and Roy, J.-S. (2008). Application of kernel-based stochastic gradient algorithms to option pricing. *mcsa*, 14(2):99–127.
- Bensoussan, A., Frehse, J., and Nagai, H. (1998). Some Results on Risk-Sensitive Control with Full Observation. *Applied Mathematics & Optimization*, 37(1):1–41.
- Bensoussan, A. and Nagai, H. (2000). Conditions for No Breakdown and Bellman Equations of Risk-Sensitive Control. *Applied Mathematics & Optimization*, 42(2):91.
- Bensoussan, A. and van Schuppen, J. H. (1985). Optimal Control of Partially Observable Stochastic Systems with an Exponential-of-Integral Performance Index. *SIAM Journal on Control and Optimization*, 23(4):599–613.
- Benth, F. E. (2013). Stochastic Volatility and Dependency in Energy Markets: Multi-Factor Modelling. In Henderson, V. and Sircar, R., editors, *Paris-princeton Lectures on Mathematical Finance 2013*. Springer.
- Benth, F. E., Benth, J. S., and Koekebakker, S. (2008). *Stochastic Modelling of Electricity and Related Markets*. World Scientific.
- Benth, F. E., Kallsen, J., and Meyer-Brandis, T. (2007). A Non-Gaussian Ornstein-Uhlenbeck Process for Electricity Spot Price Modeling and Derivatives Pricing. *Applied Mathematical Finance*, 14(2):153–169.
- Benth, F. E., Kiesel, R., and Nazarova, A. (2012). A critical empirical study of three electricity spot price models. *Energy Economics*, 34(5):1589–1616.
- Bernardo, J. M. and Smith, A. F. M. (2009). *Bayesian Theory*. John Wiley & Sons.
- Bernoulli, D. (1954). Exposition of a New Theory on the Measurement of Risk. *Econometrica*, 22(1):23–36.
- Bertsekas (1976). *Dynamic programming and stochastic control*. Academic Press.

- Bielecki, T. R. and Pliska, S. R. (1999). Risk-Sensitive Dynamic Asset Management. *Appl Math Optim*, 39(3):337–360.
- Bouchard, B. and Chassagneux, J.-F. (2008). Discrete-time approximation for continuously and discretely reflected BSDEs. *Stochastic Processes and their Applications*, 118(12):2269–2293.
- Bouchard, B. and Warin, X. (2012). Monte-Carlo Valuation of American Options: Facts and New Algorithms to Improve Existing Methods. In Carmona, R. A., Moral, P. D., Hu, P., and Oudjane, N., editors, *Numerical Methods in Finance*, number 12 in Springer Proceedings in Mathematics, pages 215–255. Springer Berlin Heidelberg.
- Brekke, K. and Øksendal, B. (1994). Optimal Switching in an Economic Activity under Uncertainty. *SIAM Journal on Control and Optimization*, 32(4):1021–1036.
- Brennan, M. J. and Schwartz, E. S. (1985). Evaluating Natural Resource Investments. *The Journal of Business*, 58(2):135–157.
- Brenner, S. C. and Scott, R. (2008). *The Mathematical Theory of Finite Element Methods*. Springer Science & Business Media.
- Carmona, R. and Ludkovski, M. (2008). Pricing Asset Scheduling Flexibility using Optimal Switching. *Applied Mathematical Finance*, 15(5-6):405–447.
- Carmona, R. and Ludkovski, M. (2010). Valuation of energy storage: an optimal switching approach. *Quantitative Finance*, 10(4):359–374.
- Casella, G. and George, E. I. (1992). Explaining the Gibbs Sampler. *The American Statistician*, 46(3):167–174.
- Chen, Z. and Forsyth, P. (2007). A Semi-Lagrangian Approach for Natural Gas Storage Valuation and Optimal Operation. *SIAM J. Sci. Comput.*, 30(1):339–368.
- Chernoff, H. (1952). A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the sum of Observations. *The Annals of Mathematical Statistics*, 23(4):493–507.
- Chib, S. and Greenberg, E. (1995). Understanding the Metropolis-Hastings Algorithm. *The American Statistician*, 49(4):327–335.

- Chib, S., Nardari, F., and Shephard, N. (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, 108(2):281–316.
- Clewlow, L. and Strickland, C. (2000). *Energy Derivatives: Pricing and Risk Management*. Lacima Publications.
- Crespo Cuaresma, J., Hlouskova, J., Kossmeier, S., and Obersteiner, M. (2004). Forecasting electricity spot-prices using linear univariate time-series models. *Applied Energy*, 77(1):87–106.
- Cvitanic, J. and Karatzas, I. (1996). Backward Stochastic Differential Equations with Reflection and Dynkin Games. *The Annals of Probability*, 24(4):2024–2056.
- Cyriel de Jong and Kasper Walet (2003). To store or not to store. *Energy Power Risk Mgmt*, pages 8–11.
- Davis, M. and Lleo, S. (2011). Jump-Diffusion Risk-Sensitive Asset Management I: Diffusion Factor Model. *SIAM J. Finan. Math.*, 2(1):22–54.
- DeGroot, M. H. (2005). *Optimal Statistical Decisions*. John Wiley & Sons.
- Dellacherie, C. and Meyer, P.-A. (1982). Probabilities and potential. B. Theory of martingales. Translated from the French by JP Wilson. *North-Holland Math. Stud*, 72.
- Deng, S.-J. and Xia, Z. (2006). A real options approach for pricing electricity tolling agreements. *Int. J. Info. Tech. Dec. Mak.*, 05(03):421–436.
- Dixit, A. (1989). Entry and Exit Decisions under Uncertainty. *Journal of Political Economy*, 97(3):620–638.
- Djehiche, B., Hamadene, S., and Popier, A. (2009). A finite horizon optimal multiple switching problem. *SIAM Journal on Control and Optimization*, 48(4):2751–2770.
- Djehiche, B., Hamadène, S., and Hdhiri, I. (2010). Stochastic Impulse Control of Non-Markovian Processes. *Appl Math Optim*, 61(1):1–26.
- Dobson, I., Carreras, B. A., Lynch, V. E., and Newman, D. E. (2007). Complex systems analysis of series of blackouts: Cascading failure, critical points,

- and self-organization. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 17(2):026103.
- Dupuis, P., James, M. R., and Petersen, I. (2000). Robust Properties of Risk-Sensitive Control. *Math. Control Signals Systems*, 13(4):318–332.
- Egloff, D. (2005). Monte Carlo algorithms for optimal stopping and statistical learning. *Ann. Appl. Probab.*, 15(2):1396–1432.
- Egloff, D., Kohler, M., and Todorovic, N. (2007). A dynamic look-ahead Monte Carlo algorithm for pricing Bermudan options. *Ann. Appl. Probab.*, 17(4):1138–1171.
- El-Karoui, N. and Hamadène, S. (2003). BSDEs and risk-sensitive control, zero-sum and nonzero-sum game problems of stochastic functional differential equations. *Stochastic Processes and their Applications*, 107(1):145–169.
- Evatt, G. W., Johnson, P. V., Duck, P. W., Howell, S. D., and Moriarty, J. (2010). The expected lifetime of an extraction project. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, page rspa20100247.
- Felix, B. (2012). Gas Storage Valuation: A Comparative Simulation Study. SSRN Scholarly Paper ID 2089268, Social Science Research Network, Rochester, NY.
- Felix, B. J. and Weber, C. (2012). Gas storage valuation applying numerically constructed recombining trees. *European Journal of Operational Research*, 216(1):178–187.
- Fink, D. (1997). A compendium of conjugate priors. Technical report, Montana State Univeristy.
- Fleming, W. H. and McEneaney, W. M. (1992). Risk sensitive optimal control and differential games. In Duncan, T. E. and Pasik-Duncan, B., editors, *Stochastic Theory and Adaptive Control*, number 184 in Lecture Notes in Control and Information Sciences, pages 185–197. Springer Berlin Heidelberg.
- Fleming, W. H. and Soner, H. M. (2006). *Controlled Markov Processes and Viscosity Solutions*. Springer.

- Frühwirth-Schnatter, S. and Sögner, L. (2009). Bayesian estimation of stochastic volatility models based on OU processes with marginal Gamma law. *Ann Inst Stat Math*, 61(1):159–179.
- Föllmer, H. and Schied, A. (2002a). Convex measures of risk and trading constraints. *Finance Stochast*, 6(4):429–447.
- Föllmer, H. and Schied, A. (2002b). *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter.
- Föllmer, H. and Schied, A. (2010). Convex Risk Measures. In *Encyclopedia of Quantitative Finance*. John Wiley & Sons, Ltd.
- Gassiat, P., Kharroubi, I., and Pham, H. (2012). Time discretization and quantization methods for optimal multiple switching problem. *Stochastic Processes and their Applications*, 122(5):2019–2052.
- Gelfand, A. E. and Smith, A. F. M. (1990). Sampling-Based Approaches to Calculating Marginal Densities. *Journal of the American Statistical Association*, 85(410):398–409.
- Gelman, A. (2003). A Bayesian Formulation of Exploratory Data Analysis and Goodness-of-fit Testing*. *International Statistical Review*, 71(2):369–382.
- Gelman, A., Meng, X.-L., and Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica sinica*, 6(4):733–760.
- Geman, H. and Roncoroni, A. (2006). Understanding the Fine Structure of Electricity Prices. *The Journal of Business*, 79(3):1225–1261. ArticleType: research-article / Full publication date: May 2006 / Copyright © 2006 The University of Chicago Press.
- Geman, S. and Geman, D. (1984). Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741.
- Geyer, C. J. and Møller, J. (1994). Simulation Procedures and Likelihood Inference for Spatial Point Processes. *Scandinavian Journal of Statistics*, 21(4):359–373.

- Gilks, W. R., Richardson, S., and Spiegelhalter, D. (1995). *Markov Chain Monte Carlo in Practice*. CRC Press.
- Gillespie, D. T. (1996). Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral. *Phys. Rev. E*, 54(2):2084–2091.
- Glasserman, P. (2004). *Monte Carlo methods in financial engineering*, volume 53. Springer.
- Gobet, E., Lemor, J.-P., and Warin, X. (2005). A regression-based Monte Carlo method to solve backward stochastic differential equations. *Ann. Appl. Probab.*, 15(3):2172–2202.
- Gonzalez, J. and Moriarty, J. (2014). Risk-sensitive optimal switching and applications to district energy systems. In *2014 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, pages 1–6.
- Gonzalez, J., Moriarty, J., and Palczewski, J. (2015). Markov Chain Monte Carlo calibration of a jump-diffusion electricity spot price model. *In progress*.
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82(4):711–732.
- Griffin, J. E. and Steel, M. F. J. (2006). Inference with non-Gaussian Ornstein-Uhlenbeck processes for stochastic volatility. *Journal of Econometrics*, 134(2):605–644.
- Hamadène, S. and Jeanblanc, M. (2007). On the Starting and Stopping Problem: Application in Reversible Investments. *Mathematics of Operations Research*, 32(1):182–192.
- Hamadène, S. and Zhang, J. (2010). Switching problem and related system of reflected backward SDEs. *Stochastic Processes and their Applications*, 120(4):403–426.
- Hastings, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109.
- Hata, H. and Iida, Y. (2006). A risk-sensitive stochastic control approach to an optimal investment problem with partial information. *Finance Stoch.*, 10(3):395–426.

- Hata, H. and Sekine, J. (2010). Explicit Solution to a Certain Non-ELQG Risk-sensitive Stochastic Control Problem. *Appl Math Optim*, 62(3):341–380.
- Hull, J. (2009). *Options, Futures and Other Derivatives*. Pearson/Prentice Hall, Upper Saddle River, 7th edition.
- Jacobson, D. (1973). Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Transactions on Automatic Control*, 18(2):124–131.
- Jacquier, E., Polson, N. G., and Rossi, P. E. (1994). Bayesian Analysis of Stochastic Volatility Models. *Journal of Business & Economic Statistics*, 12(4):371–389. ArticleType: research-article / Full publication date: Oct., 1994 / Copyright © 1994 American Statistical Association.
- James, M. R. (1992). Asymptotic analysis of nonlinear stochastic risk-sensitive control and differential games. *Math. Control Signal Systems*, 5(4):401–417.
- Johannes, M. S., Polson, N. G., and Stroud, J. R. (2009). Optimal Filtering of Jump Diffusions: Extracting Latent States from Asset Prices. *Rev. Financ. Stud.*, 22(7):2759–2799.
- Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press.
- Karatzas, I. and Shreve, S. E. (1998a). *Methods of Mathematical Finance*. Springer Science & Business Media.
- Karatzas, I. A. and Shreve, S. E. (1998b). *Brownian Motion and Stochastic Calculus*. Springer, New York, 2nd edition.
- Kastner, G. and Frühwirth-Schnatter, S. (2014). Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Computational Statistics & Data Analysis*, 76:408–423.
- Kephart, J. and Chess, D. (2003). The vision of autonomic computing. *Computer*, 36(1):41–50.
- Kim, S., Shephard, N., and Chib, S. (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies*, 65(3):361–393.

- Kitapbayev, Y., Moriarty, J., and Mancarella, P. (2015). Stochastic control and real options valuation of thermal storage-enabled demand response from flexible district energy systems. *Applied Energy*, 137:823–831.
- Kitapbayev, Y., Moriarty, J., Mancarella, P., and Blochle, M. (2013). A real options assessment of operational flexibility in district energy systems. In *European Energy Market (EEM), 2013 10th International Conference on the*, pages 1–7.
- Kjärstad, J. and Johnsson, F. (2007). Prospects of the European gas market. *Energy Policy*, 35(2):869–888.
- Kloeden, P. E. and Platen, E. (1992). *Numerical Solution of Stochastic Differential Equations*. Springer Science & Business Media.
- Kocis, L. and Whiten, W. J. (1997). Computational Investigations of Low-discrepancy Sequences. *ACM Trans. Math. Softw.*, 23(2):266–294.
- Kohler, M. (2008). A regression-based smoothing spline Monte Carlo algorithm for pricing American options in discrete time. *AStA*, 92(2):153–178.
- Kohler, M., Krzyzak, A., and Todorovic, N. (2010). Pricing of High-Dimensional American Options by Neural Networks. *Mathematical Finance*, 20(3):383–410.
- Krylov, N. V. (2008). *Controlled Diffusion Processes*. Springer Science & Business Media.
- Kupper, M. and Schachermayer, W. (2009). Representation results for law invariant time consistent functions. *Math Finan Econ*, 2(3):189–210.
- Lapidus, L. and Pinder, G. F. (2011). *Numerical Solution of Partial Differential Equations in Science and Engineering*. John Wiley & Sons.
- Liu, J. S., Wong, W. H., and Kong, A. (1994). Covariance Structure of the Gibbs Sampler with Applications to the Comparisons of Estimators and Augmentation Schemes. *Biometrika*, 81(1):27–40.
- Liu, J. S., Wong, W. H., and Kong, A. (1995). Covariance Structure and Convergence Rate of the Gibbs Sampler with Various Scans. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):157–169.

- Longstaff, F. A. and Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial studies*, 14(1):113–147.
- Lucia, J. J. and Schwartz, E. S. (2002). Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, 5(1):5–50.
- Markowitz, H. (1952). Portfolio Selection*. *The Journal of Finance*, 7(1):77–91.
- McArthur, S., Taylor, P., Ault, G., King, J., Athanasiadis, D., Alimisis, V., and Czaplewski, M. (2012). The Autonomic Power System - Network operation and control beyond smart grids. In *2012 3rd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies (ISGT Europe)*, pages 1–7.
- Meng, X.-L. (2000). Missing Data: Dial M for ??? *Journal of the American Statistical Association*, 95(452):1325–1330.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092.
- Meyer-Brandis, T. and Tankov, P. (2008). Multi-factor jump-diffusion models of electricity prices. *International Journal of Theoretical and Applied Finance*, 11(05):503–528.
- Meyn, S. and Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability*. Springer-Verlag, London.
- Moller, J. and Waagepetersen, R. P. (2004). *Statistical Inference and Simulation for Spatial Point Processes*. CRC Press.
- Müller, P. (1991). A generic approach to posterior integration and Gibbs sampling. *Technical Report, Purdue Univ., West Lafayette, Indiana*.
- Müller, P. (1993). Alternatives to the Gibbs sampling scheme. *Technical report, Institute of Statistics and Decision Sciences, Duke Univ.*
- Nagai, H. (1996). Bellman Equations of Risk-Sensitive Control. *SIAM J. Control Optim.*, 34(1):74–101.

- Nagai, H. (2003). Optimal Strategies for Risk-Sensitive Portfolio Optimization Problems for General Factor Models. *SIAM J. Control Optim.*, 41(6):1779–1800.
- Nagai, H. (2014). Risk-Sensitive Stochastic Control. In Baillieul, John and Samad, Tariq, editors, *Encyclopedia of Systems and Control*. Springer London.
- Papaspiliopoulos, O., Robert, G., and Sköld, M. (2003). Non-centered parameterisations for hierarchical models and data augmentation. In Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., editors, *Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting*, page 307. Oxford University Press, USA.
- Peskir, G. and Shiryaev, A. (2006). *Optimal Stopping and Free-Boundary Problems*. Springer Science & Business Media.
- Pham, H. (2009). *Continuous-time Stochastic Control and Optimization with Financial Applications*. Springer-Verlag Heidelberg.
- Pham, H., Vath, V., and Zhou, X. (2009). Optimal Switching over Multiple Regimes. *SIAM J. Control Optim.*, 48(4):2217–2253.
- Pilipovic, D. (2007). *Energy Risk: Valuing and Managing Energy Derivatives: Valuing and Managing Energy Derivatives*. McGraw Hill Professional.
- Platen, E. and Bruti-Liberati, N. (2010). *Numerical Solution of Stochastic Differential Equations with Jumps in Finance*. Springer Science & Business Media.
- Ploeg, F. v. d. (1984). Economic Policy Rules for Risk-Sensitive Decision Making. *Zeitschrift für Nationalökonomie/Journal of Economics*, 44.
- Robert, C. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer Science & Business Media.
- Robert, C. and Casella, G. (2009). *Introducing Monte Carlo Methods with R*. Springer Science & Business Media.
- Robert, C. P. (2007). *The Bayesian choice from decision-theoretic foundations to computational implementation*. Springer, New York.

- Roberts, G. O., Papaspiliopoulos, O., and Dellaportas, P. (2004a). Bayesian Inference for Non-Gaussian Ornstein-Uhlenbeck Stochastic Volatility Processes. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 66(2):369–393.
- Roberts, G. O. and Rosenthal, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statistical science*, 16(4):351–367.
- Roberts, G. O., Rosenthal, J. S., and others (2004b). General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1:20–71.
- Roberts, G. O. and Sahu, S. K. (1997). Updating Schemes, Correlation Structure, Blocking and Parameterization for the Gibbs Sampler. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(2):291–317.
- Rockafellar, R. T. and Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42.
- Royden, H. L., Fitzpatrick, P., and Hall, P. (1988). *Real analysis*, volume 198. Macmillan New York.
- Seifert, J. and Uhrig-Homburg, M. (2007). Modelling jumps in electricity prices: theory and empirical evidence. *Review of Derivatives Research*, 10(1):59–85.
- Sermaidis, G., Papaspiliopoulos, O., Roberts, G. O., Beskos, A., and Fearnhead, P. (2013). Markov Chain Monte Carlo for Exact Inference for Diffusions. *Scandinavian Journal of Statistics*, 40(2):294–321.
- Sherlock, C., Fearnhead, P., and Roberts, G. O. (2010). The Random Walk Metropolis: Linking Theory and Practice Through a Case Study. *Statistical Science*, 25(2):172–190.
- Simkins, B. and Simkins, R. (2013). *Energy Finance and Economics: Analysis and Valuation, Risk Management, and the Future of Energy*. John Wiley & Sons.
- Stentoft, L. (2004). Assessing the least squares Monte-Carlo approach to American option valuation. *Review of Derivatives research*, 7(2):129–168.
- Sørensen, M. (2000). Prediction-based estimating functions. *Econometrics Journal*, 3(2):123–147.

- Tanner, M. A. and Wong, W. H. (1987). The calculation of posterior distributions by data augmentation. *Journal of the American statistical Association*, 82(398):528–540.
- Thompson, M., Davison, M., and Rasmussen, H. (2004). Valuation and Optimal Operation of Electric Power Plants in Competitive Markets. *Operations Research*, 52(4):546–562.
- Thompson, M., Davison, M., and Rasmussen, H. (2009). Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics*, 56(3):226–238.
- Tierney, L. (1994). Markov Chains for Exploring Posterior Distributions. *The Annals of Statistics*, 22(4):1701–1728.
- Tierney, L. (1998). A Note on Metropolis-Hastings Kernels for General State Spaces. *The Annals of Applied Probability*, 8(1):1–9.
- Tsitsiklis, J. and Van Roy, B. (1999). Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control*, 44(10):1840–1851.
- Tsitsiklis, J. and Van Roy, B. (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703.
- van de Ven, P., Hegde, N., Massoulié, L., and Salonidis, T. (2013). Optimal Control of End-User Energy Storage. *IEEE Transactions on Smart Grid*, 4(2):789–797.
- Vito M. R. Muggeo (2008). segmented: An R Package to Fit Regression Models with Broken-Line Relationships. *R News*, 8(1).
- Warin, X. (2012). Gas Storage Hedging. In Carmona, R. A., Del Moral, P., Hu, P., and Oudjane, N., editors, *Numerical Methods in Finance*, volume 12 of *Springer Proceedings in Mathematics*, pages 421–445. Springer Berlin Heidelberg.
- Weron, R. (2007). *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach*. John Wiley & Sons.

- Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30(4):1030–1081.
- Whittle, P. (1981). Risk-Sensitive Linear/Quadratic/Gaussian Control. *Advances in Applied Probability*, 13(4):764–777. ArticleType: research-article / Full publication date: Dec., 1981 / Copyright © 1981 Applied Probability Trust.
- Whittle, P. (1990). *Risk-sensitive optimal control*. Wiley New York.
- Yu, Y. and Meng, X.-L. (2011). To Center or Not to Center: That Is Not the Question—An Ancillarity–Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Efficiency. *Journal of Computational and Graphical Statistics*, 20(3):531–570.
- Zanger, D. Z. (2013). Quantitative error estimates for a least-squares Monte Carlo algorithm for American option pricing. *Finance Stoch*, 17(3):503–534.