



City Research Online

City, University of London Institutional Repository

Citation: Pouliasis, Panagiotis (2011). Essays on the empirical analysis of energy risk. (Unpublished Doctoral thesis, City University London)

This is the unspecified version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/1165/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Essays on the Empirical Analysis of Energy Risk

by

Panagiotis K. Pouliasis

A thesis submitted in fulfilment of the requirements for the Degree of Doctor of
Philosophy in the subject of Finance

City University London
Sir John Cass Business School
The Costas Grammenos International Centre for Shipping, Trade and Finance
London, UK
April, 2011

*In Memory of my Grandfathers,
Ioannis Vitas and Panagiotis Poulasis
And my Godmother Vicky Galani.*

Table of Contents

Table of Contents.....	iii
List of Figures.....	v
List of Tables.....	vi
List of Abbreviations and Mathematical Symbols.....	vii
Acknowledgements.....	x
Declaration.....	xi
Abstract.....	xii

Chapter 1: Introduction and Summary of the Thesis

1.1 Motivation and Aim of the Thesis.....	1
1.2 Thesis Objectives and Contribution.....	4
1.3 Summary of Thesis Structure.....	7

Chapter 2: Introduction to Oil Markets and Energy Risk

2.1 Introduction.....	8
2.2 Main Features of the Oil Market.....	9
2.3 Demand and Supply Framework.....	13
2.3.1 The OPEC effect.....	15
2.4 Historical Overview of Fluctuations in Oil Prices.....	17
2.5 Oil Price Volatility.....	21
2.6 The Petroleum Futures Market.....	24
2.6.1 Speculation and Investor Behaviour.....	27
2.7 Petroleum Price Risk Management.....	28
2.7.1 Quantifying Market Risk.....	31
2.7.2 Minimum Variance Futures Hedging.....	34
2.7.3 Metallgesellschaft Hedging Debacle.....	37
2.8 Term Structure of Futures Prices.....	40
2.9 Conclusion.....	45

Chapter 3: Regime Switching Models and Applications in Finance

3.1 Introduction.....	46
3.1.1 A Primer on Models of Changing Regime.....	48
3.2 Literature Review.....	49
3.2.1 General Review.....	49
3.2.2 Evidence from the Oil Markets.....	51
3.3 Fundamental Concepts of Markov Processes.....	54
3.3.1 Mixture of Distributions.....	54
3.3.2 Markov Chains.....	59
3.4 The Baseline Markov Regime Switching Model.....	62
3.4.1 Regime Inference and Maximum Likelihood.....	64
3.4.2 Path Dependency in Volatility.....	71
3.5 Conclusion.....	72

Chapter 4: Forecasting Petroleum Markets Volatility: The Role of Regimes and Market Conditions	
4.1	Introduction..... 74
4.2	Methodology 77
4.3	Description of the Data and Preliminary Analysis..... 80
4.4	Empirical results 84
4.4.1	Out-of-Sample Performance of Volatility Forecasts..... 88
4.4.2	Evaluating the Predictive Performance of Value-at-Risk Forecasts 95
4.4.2.1	Measuring Forecasting Performance with Risk Management Loss Functions.....100
4.5	Conclusions..... 104
APPENDIX 4.A: Alternative Distribution Assumptions 107	
APPENDIX 4.B: Tests of Two versus Three Regimes 109	
APPENDIX 4.C: The Stationary Bootstrap 112	
Chapter 5: A Markov Regime Switching Approach for Hedging Petroleum Commodities	
5.1	Introduction..... 113
5.2	Markov Regime Switching GARCH Models & Hedging..... 116
5.3	Description of the Data & Preliminary Analysis 120
5.4	Empirical Results 122
5.5	Time Varying Hedge Ratios & Hedging Effectiveness 126
5.6	Data Snooping Bias..... 131
5.7	Downside Risk Measures..... 133
5.8	Conclusions..... 136
APPENDIX 5.A : Time Varying Transition Probabilities 138	
APPENDIX 5.B: A Note on Seasonality and Hedging 141	
Chapter 6: Petroleum Term Structure Dynamics, Inter-Commodity Dependencies and the Role of Regimes	
6.1	Introduction..... 145
6.2	Methodology 148
6.2.1	Factor Decomposition..... 148
6.2.2	Modelling the Information in the Term Structure 151
6.3	Data Description and Preliminary Analysis..... 155
6.3.1	Unit Root and Co-integration Results 158
6.4	Empirical Results 161
6.5	Forecasting the Futures Curve Dynamics 172
6.5.1	Forecasting Petroleum Spreads..... 174
6.5.2	Forecasting the Variance Covariance Matrix..... 177
6.5.3	An Application to Value-at-Risk 181
6.6	Conclusions..... 185
APPENDIX 6.A: Factor Seasonality and Auto-correlogram 187	
Chapter 7: Concluding Remarks and Future Research	
7.1	Summary and Conclusions 188
7.1.1	Risk Measurement 189
7.1.2	Risk Management 190
7.1.3	Term Structure of Correlated Curves..... 191
7.2	Directions for Further Research..... 192
References 194	

List of Figures

Figure 2.1: Imbalances in the World Oil Market Structure	12
Figure 2.2: Middle East as a Swing Producer	13
Figure 2.3: Crude Oil Prices, 1869 - 1969.....	18
Figure 2.4: Crude Oil Prices, 1870 - 2000.....	19
Figure 2.5: Crude Oil Prices, 2001 - 2009.....	20
Figure 2.6: Annualised Spot Volatilities & Correlations of Main US Petroleum Commodities and Crack Spreads.....	23
Figure 2.7: Trading Activity of NYMEX & Brent Crude Oil Futures	26
Figure 2.8: Value-at-Risk approx. using WTI vs. a Diversified Portfolio	32
Figure 2.9: Dependence of Variance on the Hedge Ratio	36
Figure 2.10: Evolution of Heating Oil Term Structure vs. Spot (New York)	42
Figure 2.11: Term Structure of Prices and Volatilities.....	43
Figure 2.12: Premium/Discount of 1- over 6- & 18-Month WTI Futures.....	43
Figure 3.1: Fitted Mixture of Two Normals for WTI Crude Oil.....	58
Figure 3.2: Illustration of the Structure of a Three-State Markov Process.....	60
Figure 3.3: Ex-ante & Smooth Regime Probabilities for NYMEX Heating Crack.....	70
Figure 3.4: MRS Model Volatility vs. Actual Returns for NYMEX Heating Crack	70
Figure 3.5: Volatility Path-Dependency in the GARCH Model	71
Figure 4.1: Log-Prices of NYMEX WTI and Heating Oil Futures	81
Figure 4.2: NYMEX WTI Crude Oil Futures Log – Returns (standardised) Historical PDF vs. Standard Normal PDF	83
Figure 4.3: Heating Oil # 2 Futures Log – Returns (standardised) Historical PDF vs. Standard Normal PDF	84
Figure 4.4: WTI Crude Oil Futures Volatility	87
Figure 4.5: WTI Crude Oil Regime Probabilities of being in the Stable Regime (MRS-GARCH-X Model).....	87
Figure 4.6: WTI Crude Oil 5% VaR Estimates for Long & Short Positions	100
Figure 5.1: Smooth Regime Probabilities for WTI Crude Oil –Probability of being in the Low Variance State.....	125
Figure 5.2: Constant OLS, VECM-GARCH and MRS-BEKK Hedge Ratios for WTI Crude Oil	126
Figure 5.3: Basis for WTI Crude Oil.....	127
Figure 5.A.1: Time Varying Transition Probabilities of WTI Crude Oil.....	140
Figure 5.B.1: Monthly Seasonal Components of Spot-Futures Weekly Volatilities.....	142
Figure 5.B.2: Monthly Seasonal Components of Spot-Futures Correlations.....	143
Figure 6.1 Seasonality Adjusted Weights of Principal Components for Heating Oil & WTI Crude Oil.....	157
Figure 6.2 Estimated Level, Slope and Curvature Factor Prices for WTI Crude Oil and Heating Oil	159
Figure 6.3 NYMEX Heating Crack Regime Smoothed Probabilities for Level, Slope & Curvature Factors	169
Figure 6.4 WTI-Brent Regime Smoothed Probabilities for Level, Slope & Curvature Factors	169
Figure 6.5: NYMEX Heating Crack 1, 3 & 9 Month Futures Spread Volatilities	170
Figure 6.6: NYMEX Heating Crack Correlations of 1, 3 & 9 Month Futures	170
Figure 6.7: WTI- Brent 1, 3 & 9 Month Futures Spread Volatilities	171
Figure 6.8: WTI- Brent Correlations of 1, 3 & 9 Month Futures	171
Figure 6.9: NYMEX Crack 5% VaR Estimates for the Equally Weighted Portfolio	184
Figure 6.10: NYMEX Crack 5% VaR Estimates for the Calendar Portfolio	185
Figure 6.A.1: ACF of the 3 rd Factor (Curvature) for Heating & WTI Crude Oil, before and after the Adjustment for Seasonality	187

List of Tables

Table 2.1: Volatility Across Different Assets	30
Table 3.1: Fitted Finite Mixtures.....	57
Table 3.2: Simple MRS Models of Petroleum Spreads.....	69
Table 4.1: Summary Statistics & Unit Root Tests for NYMEX & ICE Petroleum Futures	82
Table 4.2: Estimates of Switching GARCH-X Models for NYMEX & ICE Petroleum Futures.....	85
Table 4.3: Comparisons of Out-Of-Sample Forecasting Performance of Volatility Models	90
Table 4.4: Comparisons of Out-Of-Sample Forecasting Performance of Volatility Models under different periods	93
Table 4.5: Value-at-Risk & Risk Management Loss Functions for Long & Short Positions	98
Table 4.6: Quantile Loss Across Different Market Conditions and Periods	102
Table 4.A.1: Estimates of Switching GARCH-X Models for NYMEX & ICE Petroleum Futures Under the Assumption of Generalised Error Distribution	108
Table 4.B.1: Model Selection Criteria.....	109
Table 4.B.2: Likelihood Ratio Tests – 2 vs. 3 Regimes	111
Table 5.1: Summary Statistics, Unit Root & Cointegration Tests for Spot and Futures Prices of WTI Crude Oil, Unleaded Gasoline and Heating Oil # 2	121
Table 5.2: Estimates of Markov Regime Switching BEKK Hedge Ratios for NYMEX Energy Commodities	123
Table 5.3: Hedging Effectiveness of Markov Regime Switching Against the Constant and Alternative Time-Varying Hedge Ratio Models	130
Table 5.4: Effectiveness Long/Short Hedging Positions of Markov Regime Switching Against the Constant and Alternative Time-Varying Hedge Ratio Models	135
Table 5.A.1: MRS-BEKK models with Transition Probabilities Conditioned on Inventories	139
Table 5.B.1: LR Tests on the Residuals of the MRS-BEKK Model	142
Table 6.1: Preliminary Data Analysis & PCA Results.....	156
Table 6.2: Unit Root & Johansen Cointegration Tests for Petroleum Futures Factors	160
Table 6.3: Estimates of Markov Regime Switching Models (Unrestricted Models)	163
Table 6.4: Unconditional Probabilities & Expected Duration.....	167
Table 6.5: Model Diagnostics	172
Table 6.6: Root Mean Squared Errors, Forecasting the Term Structure of Contemporaneous Spreads	175
Table 6.7: Forecasting the Variance Covariance Matrix of Correlated Futures Curves	179
Table 6.8: Forecasting Portfolio Value-at-Risk.....	183

List of Abbreviations and Mathematical Symbols

ADF	Augmented Dickey and Fuller (1979, 1981) unit root test.
AQLF	Average Quadratic Loss Function.
AR	Autoregressive.
BEKK	Baba, Engle, Kraft and Kroner (1987) formulation for multivariate GARCH.
BIC	Bayesian Information Criterion.
bpd	Barrels per day.
CB	ICE Brent Crude Oil.
CL	NYMEX West Texas Intermediate Light Sweet Crude oil.
C, Ct	Curvature factor.
CV	Cointegrating Vector.
DCC	Dynamic Conditional Correlation.
EVT	Extreme Value Theory.
Eq.	Equation.
GARCH	Generalised Autoregressive Conditional Heteroscedasticity.
GARCH-X	Augmented Generalised Autoregressive Conditional Heteroscedasticity.
GED	Generalised Error Distribution.
GO	ICE Gas oil.
GPD	Generalised Pareto Distribution.
HO	NYMEX Heating Oil # 2.
HS	Historical Simulation.
ICE	Intercontinental Exchange.
ICSS	Iterative Cumulative Sums-of-Squares.
I-MRS	Independent MRS i.e. system of equations, each following an independent Markov process.
J-B	Bera and Jarque (1980) test for normality.
KPSS	Kwiatkowski, Phillips, Schmidt and Shin (1992) unit root test.
LF	Loss function.
LogLik	Log Likelihood Function.
LogLik _{CON}	Constrained LogLik.
LogLik _{UNCON}	Unconstrained LogLik.
LR	Likelihood Ratio.
LR stat	Likelihood Ratio Statistic (equivalent to M).
LR _{CC}	Likelihood Ratio Test of Conditional Coverage.
LR _{IND}	Likelihood Ratio Test of Independence.
LR _{UC}	Likelihood Ratio Test of Unconditional Coverage.
L, Lt	Level factor.
MRS	Markov Regime Switching.
MAE	Mean Absolute Error.
Mix	Mixture of distributions (for GARCH models i.e. Mix-GARCH).
ML	Maximum Likelihood.
MME	Mixed Mean Error Statistic.
MME(O)	Mixed Mean Error Statistic of Over-prediction.
MME(U)	Mixed Mean Error Statistic of Under-prediction.
MVHR	Minimum Variance Hedge Ratio.
NYMEX	New York Mercantile Exchange.
OLS	Ordinary Least Squares.
OECD	Organisation for Economic Co-operation and Development.
OPEC	Organisation of Petroleum Exporting Countries.
PCA	Principal Components Analysis.
PF	Percentage of Failures.
PP	Philips and Perron (1988) unit root test.
QL	Quantile Loss.
RC	Reality Check.
RW	Random Walk.
RMSE	Root Mean Squared Error.
SBIC	Schwarz (1978) Bayesian Information Criterion (same as BIC).
S, St	Slope factor.
UL	Unexpected Loss.

VaR	Value-at-Risk.
VAR	Vector Autoregressive process.
VAR-X	Augmented Vector Autoregressive process.
VECM	Vector Error Correction Model.
Vol	Volatility.
W-Sum	Weighted Sum of MME(U) and MME(O).
WTI	NYMEX West Texas Intermediate Light Sweet Crude oil.

A	Sensitivity to past error terms (GARCH model).
B	Sensitivity to past variance terms (GARCH model).
Covar(\cdot)	Covariance.
D	Diagonal matrix containing volatilities.
$e(\cdot)$, $\exp(\cdot)$	Exponential function.
$E[\cdot]$	Expectation operator.
$f(\cdot)$	Probability density function (pdf).
fn	Loss differential. Performance measure of the differential between two loss functions.
f	Vector of conditional distributions.
g	Degrees of freedom for GED.
F_t	Futures prices.
$F_{t,T}$	Futures prices at time t with maturity at time T.
$F(t,T)$	Futures prices at time t with maturity at time T.
h	Variance process, equivalent to σ^2 .
h_{ij}	Covariance process, equivalent to σ_{ij} .
H	Variance covariance matrix (equivalent to Σ and V).
i	Index variable.
$I(\cdot)$	Integrated of order (\cdot).
$I_{t,j}$	Indicator function.
j	Index variable.
K	Number of components – states.
k	Index for K. It also appears as an index variable, same as i and j .
$L(\cdot)$	Likelihood function.
LF	Loss Function.
M	Likelihood Ratio Statistic [for Davies (1987) test].
$N(\mu, \sigma^2)$	Normal distribution with mean μ and volatility σ .
Nu	Number of excesses above the threshold level u.
p_{ij}	Transition probability i.e. the probability that state i will be followed by state j .
p	Number of lags.
P	2x2 transition probability matrix.
$\Pr(\cdot)$	Probability operator.
q	Smoothing parameter.
Q	2x2 transition probability matrix.
$Q(\cdot)$	Ljung-Box Q statistics of autocorrelation.
r	Restrictions (number of – mainly for LR tests).
r_t	Log - returns.
rp_t	Return on the (producer's hedged) portfolio.
RV	Realised variance covariance matrix.
sv	Semi-Variance.
s_t	Discrete random state variable.
S_t	Spot prices.
t	Time index.
T	Set for t or maturity time.
T^{RC}	Observed statistic of the RC.
T^{RC*}	Simulated T^{RC} statistic.
T	Transpose operator (as a superscript).
u	Threshold value.
U	Eigenvector.
V	Variance covariance matrix (equivalent to Σ and H).
$\text{Var}(\cdot)$	Variance.
W_0	Wealth.
W	Vector containing the weights of a portfolio.

X_t	Random variable.
y	Excess negative shocks (losses) over a threshold.
Z	Term structure deviations (basis) - exogenous variable to GARCH model.
α	Speed of adjustment to long run mean.
$\mathbf{\alpha}$	Vector of coefficients measuring the speed of convergence to the long run mean.
β_0	Intercept/ long run mean (e.g. in the cointegration equation).
β_1	Slope (e.g. in the cointegrating equation).
$\mathbf{\beta}$	Vector of cointegrating parameters.
γ	Hedge ratio.
$\mathbf{\Gamma}$	Coefficient matrix measuring the short-run adjustment to changes in a system of equations.
$\Gamma(\cdot)$	Gamma function.
δ	Drift.
Δ	Difference operator.
ε_t	Residuals.
ζ	Scalar/constant parameter.
ϑ	Scale index for the GPD.
$\mathbf{\theta}$	Parameter vector.
Θ	this is for Davies (1987) bound test. Equivalent to $2LRstat^{0.5}$.
λ_{\max}	Maximum eigenvalue statistic [Johansen (1988) test].
λ_{trace}	Trace statistic [Johansen (1988) test].
$\mathbf{\Lambda}$	Vector of eigenvalues.
μ_t	Conditional mean equation.
v	Intercept.
ξ	Shape index for the GPD.
π	$\pi \approx 3.14159$.
π_k	Unconditional probability of state k.
$\pi_{kt T}$	Smoothed probability.
$\pi_{kt t-1}$	Ex-ante probability.
$\pi_{kt t}$	Filter probability.
$\mathbf{\pi}$	Vector of unconditional state probabilities.
$\mathbf{\Pi}$	Coefficient matrix measuring the long-run adjustment to changes in a system of equations
$\Pi(\cdot)$	Product operator.
ρ	Correlation.
σ	Volatility/ standard deviation.
$\sigma(\cdot)$	Volatility function.
$\Sigma(\cdot)$	Summation operator.
$\mathbf{\Sigma}$	Variance covariance matrix (equivalent to \mathbf{H} and \mathbf{V}).
φ	Sensitivity to exogenous GARCH term.
ϕ	Coefficients for logistic function.
$\Phi(c)$	Cumulative distribution function at confidence level 1-c.
$\chi^2(\cdot)$	chi squared distribution with degrees of freedom (\cdot).
$\mathbf{\Psi}$	Transition probability matrix.
ω	Intercept of variance equation.
Ω_t	Information set up to time t.
$\mathbf{1}$	Vector of ones.
1-c	Confidence level.
\odot	Element by element multiplication.
\otimes	Kronecker product.
\perp	Orthogonal complement.
* **, ***	Asterisks denote significance at 10%, 5% and 1% significance level unless otherwise stated.
\cdot	Transpose operator, same as \mathbf{T} .
$ \mathbf{H} $	Determinant of matrix \mathbf{H} .

Acknowledgements

I would like to express my deep appreciation to my supervisor *Prof. Nikos K. Nomikos* for his enthusiasm, assistance and patience throughout my time as a doctoral PhD candidate. Nikos has provided me with invaluable knowledge transfer to complete this thesis and constant encouragement. He has been an unconditional source of knowledge and I am truly indebted to his commitment.

My particular appreciation goes also to *Dr. Amir H. Alizadeh*. I thank Amir for all the creative discussions we had and for providing me valuable and critical comments during several meetings ever since my MSc dissertation. I am grateful for all his generous support.

Furthermore, my gratitude goes to *Prof. Costas Th. Grammenos*, for his sincere interest in my intellectual development, useful advice and for providing direction during my years in London and Cass Business School. I would also like to thank my cousin, *Dr. Nikos Papapostolou*, also a member of the Centre's Faculty, for his unending help and support. All the academic staff at the Centre for Shipping, Trade and Finance has fostered an ideal research environment providing tremendously helpful and constructive insights. Special thanks also go to *Marlene Stapleton* for her support and for being there when needed. Also I am thankful to *Dr. Ioannis Kyriakou* for the many discussions, academic and otherwise. I am also indebted to *Elina Malioti* and *Iliana Kristalli* who have greatly supported me during the difficult times of writing this PhD thesis.

Chapter 4 has benefited from the constructive suggestions of three anonymous referees and the Editor of the Energy Economics journal, *Richard Tol*. Chapter 5 has also benefited from the constructive suggestions of two anonymous referees, the participants at the 2007 Commodities and Finance Centre (CFC) Conference of Birkbeck University in London and the past Editor of the Journal of Banking and Finance, *Giorgio Szego*.

Finally, the unwavering support of family has been unparalleled. My parents, *Kostas and Maria*, have always believed in me and I would like to deeply thank them for all their unreserved dedication, infinite support and continuous encouragement. Without their love and inspiration, it would not have been possible for me to get this thesis together.

Panos K. Pouliasis
Cass Business School
February 2011

Declaration

I grant powers of discretion to the University Librarian to allow this thesis to be copied in whole or in part without further reference to me. This permission covers only single copies made for study purposes, subject to normal conditions of acknowledgement.

Abstract

Energy markets have become increasingly sophisticated, requiring modelling techniques of analogous calibre. This thesis deals with models of changing regime for the petroleum complex. Modelling the conditional distribution of energy prices as a regime switching process is motivated by the market-specific characteristics of oil: different market conditions, such as backwardation and contango, involve different dynamics. The first empirical part examines the very short-end of the futures curve volatility. To address in a realistic way the potential diverse response of oil volatility to fundamentals across high and low volatility regimes, augmented regime volatility models are employed. Results indicate that volatility can be decomposed to a highly persistent conditional volatility process and a relatively short-lived non-stationary process. Apart from evaluating the size of price risk, risk managers must also design a framework for mitigating their exposures. This is the focus of the second empirical part which estimates dynamic hedge ratios. Linking the concept of disequilibrium with that of uncertainty across high and low volatility regimes, a state-dependent error correction model with time-varying second moments is introduced. Finally, the third empirical part, examines the information content of the dependence structure between correlated petroleum futures curves. Term structure is decomposed into level, slope and curvature shocks. Introducing a multi-regime framework, these factors are utilised to study inter-commodity and inter-market spreads. Results suggest markedly different state-dependent speeds of mean reversion and volatility/correlation dynamics across regimes. Overall, the employed models provide superior forecasting performance and indicate that state-dependent dynamics may provide significant benefits to market participants. The findings of this thesis have important implications for energy market trading and risk management, as well as energy market operations, such as refining and budget planning, by providing valuable information on the oil price volatility dynamics and the ability to predict risk.

Chapter 1

Introduction and Summary of the Thesis

1.1 Motivation and Aim of the Thesis

After the two oil price shocks and the development of derivatives markets in the 1980's, oil consumption has increased by more than 20 million barrels per day whereas the total trading volume of futures contracts has far exceeded total world oil production. Since then, apart from "traditional" market players having exposure in the physical market (producers, refiners, marketers etc.), other participants such as commodity portfolio managers, hedge funds, index speculators and investment banks have progressively increased their share and exposure in the energy sector. As a result, energy commodity prices have experienced an unparalleled growth over the last decade with prices of crude oil showing an extremely persistent momentum going from \$20/bbl in the early 2000's to above \$80/bbl in the mid-2005 and over \$140/bbl in July 2008. Although the 2008 recession had a significant negative effect on commodity prices with crude oil falling, in less than six months, to below \$40/bbl, recovery was fast and at the end of 2009 oil fluctuated around \$80/bbl. What is more, the OPEC oil crisis and deregulation in the 1980's was followed by remarkable increases in energy price volatility. The main contributors of this phenomenon are the geographic concentration of oil supply at high political tension regions, weather sensitive demand, absence of readily available substitutes and the overall market structure, from the major determinants of prices of these commodities and the pricing mechanisms to cartel behaviour and the specific design of the supply chain network.

Therefore, given the policy implications at both macroeconomic and microeconomic level, modelling the dynamics of petroleum commodities has been a field where a vast amount of research has been conducted with the particular sector attracting considerable interest as a financial investment vehicle in recent years. This can also be attributed to the unique market forces driving exhaustible resources' price dynamics like oil and its products. Petroleum commodity markets have undergone fundamental changes, have become more sophisticated, and investors are continually confronted with new challenges.

In a naturally dynamic world that is characterised by continuously changing relationships, the energy industry has several reasons to promote applications in risk analysis. First, it is the capital intensive character of the industry. For example, oil field development and refinery capital investments call for reliable risk assessments and accurate decision making. Second, it is the diverse mix of participants involved in the physical markets. Households, corporations and governments are all involved in the industry, either as direct or indirect consumers, thus verifying, the significance of hydrocarbons which are indispensable for transportation, industrial and residential uses. For instance, crude oil represents a significant component of operating costs to large energy consumers such as refineries, shipping companies and airlines. Furthermore, petroleum importing countries are particularly susceptible to oil price increases as the price transmission to these economies is more consequential and governments are forced to adjust their revenue and expenditure policies accordingly. Third, it is the recent emergence of energy assets as financial investment vehicles. Significant amounts of funds have been and are being constantly allocated to energy commodities; they have become very popular among institutional investors of versatile risk attitudes either as a pure speculation instrument or as a diversification tool. Finally, although trading in petroleum commodities has existed for decades it is only around 30 years after deregulation and the organisation of exchanges around the world. It was only recently that a competitive market framework - where prices are determined freely under the fundamentals of supply and demand - was developed.

In such an exigent environment, price volatility has become an important feature of the market forming a market ripe with opportunities, but in turn increasing the need for risk measurement and management using derivative contracts such as futures. The main motivation of this thesis is to build on modern quantitative techniques with a view to address several issues of oil price modelling and risk management which are very relevant and fashionable topics in the industry. The driving force for the development of such models of petroleum markets is the need, by market agents, to ensure accurate estimation of risk measures, successful implementation of hedging strategies as well as thorough evaluation of investment policies. This thesis is a compilation of three closely related essays in petroleum risk modelling and management, dealing with several practically relevant issues in empirical energy economics. That said, three central aims are determined. The first is to quantify the risk of the more liquid and volatile near to maturity contracts where market activity is mainly concentrated. The second is to develop a methodology for futures hedging designed to support risk management programmes. The third is to understand and explore fundamental relationships, long-run

equilibria, and interdependencies between petroleum commodities and reveal the mechanics in the functioning of the term structure of futures and futures spreads in the energy complex.

Empirical stylised facts of petroleum return series suggest that risk is time-varying and depends on market conditions. This thesis addresses the explicit modelling of nonlinearities in the underlying data-generating process, as well as the conditional second moments in petroleum markets. Although significant quantitative advances in a Markov Regime Switching (MRS) framework have been made since the seminal paper of Hamilton (1989), moving these concepts into applied research in petroleum commodities is still underdeveloped. This thesis analyses the relative merits of regime switching models to describe change in the context of energy risk. The exploration of the dynamics of petroleum markets is aimed at improving the understanding and modelling of the real-world dynamics. We consider petroleum commodity cycles in the form of low and high volatility regimes and the switching between these cycles is assumed to be driven by Markov dynamics.

We argue that traditional single state models are not sufficiently flexible to explain real world dynamics. Price, volatility and correlation change as new information arrives in the market, causing market dynamics to switch back and forth among different processes. The focus of this thesis will be on explaining this behaviour in oil markets and further demonstrate whether the existence of such states prompts for the need to assess risk differently. In doing so, we benefit from the flexible family of MRS models that permit us to accommodate many of the stylised facts that these markets exhibit such as non-normality, asymmetries and time-varying dependence. The information content derived from MRS models will be thoroughly discussed with the aim to assess their role and effectiveness in quantifying risk under different market conditions, evaluate the extent to which regimes convey relevant information on risk management objectives and finally uncover fundamental interactions in a multi-regime framework.

The topics studied range from risk quantification, volatility/ correlation forecasting, futures hedging as well as identification of risk factors, term structure dynamics and co-integration. All essays have many things in common; first, they all focus on time series properties of petroleum prices; second, they all explicitly model the return volatilities and/or correlations of these assets as time varying; third, they all deal with nonlinear models; and forth they all aim on accurate risk assessment and enhanced forecasting ability. Policy makers, investors and, in general, all market players (crude oil producers and consumers, refiners, portfolio managers, commodity traders etc.) need to address these issues by delicately measuring the degree of energy risk exposure and the impact of price and volatility variability

on their cash flows in order to devise sound risk management strategies and reduce income uncertainty.

1.2 Thesis Objectives and Contribution

This thesis consists of three essays that discuss both theory and applications of regime switching models to energy futures markets. The thesis contributes to the existing literature by addressing three main issues: the application of regime switching processes to the volatility of short-term energy futures, the regime switching behaviour of the futures-spot relationship with application to minimum variance hedging and empirical evidence of regime shifts in the petroleum commodities market with specific interest in the interdependence between different commodities comprising economically meaningful spreads.

In the second chapter, *Introduction to oil markets and energy risk*, we review fundamental concepts of the petroleum market structure and dynamics. The chapter begins with an introduction on how petroleum markets have evolved. This section is followed by an overview of the particular market-specific characteristics and illustrates the environment that the industry operates in as well as the risks inherent in the energy sector. After an outline of the fundamentals (supply-demand) and a brief reference to market structure and the role of OPEC (Organisation of Petroleum Exporting Countries), we provide a synopsis of historical developments in oil prices; this serves as a bridge to discuss the implications of oil price volatility and the emergence of organised exchanges for petroleum commodities. The term structure of future prices and the incentives for risk management, including risk quantification and minimum variance hedge ratios, are also discussed. Next, the chapter discusses the importance of flexible risk management programmes by means of a case study (Metallgesellschaft Refining and Marketing, MGRM; 1993) highlighting the lessons that can be learnt from the past.

In the third chapter, *Regime switching models and applications in finance*, we provide a literature review with the objective to present several applications of regime switching models in finance and furthermore, showcase some important findings from the energy markets in general. A conformable introduction of the basic concepts behind MRS models is also presented; for this reason we briefly review mixture distribution models and Markov Chains. This is followed by the basic set up of the MRS model and estimation techniques.

The fourth chapter, *Forecasting petroleum futures market volatility: the role of regimes and market conditions*, is the first empirical chapter of the thesis and proposes the use of various volatility regime models (mixture distribution and MRS GARCH) in the petroleum futures

markets. Only few studies have analysed in depth the nature of the volatility regimes of oil futures prices and their forecasting ability. We extend previous research by accounting for the effect of deviations of the term structure (as measured by the squared lagged basis of futures prices) in the conditional volatility processes. Investigating volatility components under different regimes will enable us to investigate for the first time the asymmetric dependence of volatility to the basis and draw some new interesting insights regarding the effect of disequilibrium and the persistence of volatility under different market conditions. State dependent models are found superior in representing volatility persistence than the traditional GARCH models, and also tend to perform better in an out-of-sample basis. The conditional regime volatility process can be described by long memory and low sensitivity to market shocks, when the market is in the low variance state, and a relatively short-lived nonstationary process with higher sensitivity to shocks, when the market is in the high variance state. In addition, we link the regime volatility framework with tail estimation by examining the tails of the conditional distributions of the models and extending the above framework to a conditional extreme value theory setting. Volatility and Value-at-Risk forecasts are tested across periods of backwardation and contango, since the risk-return profile of energy prices is known to change fundamentally between the two different states. Overall, by identifying different volatility components for normal and highly volatile periods, market participants may benefit in terms of accurate risk quantification.

The fifth chapter, *A Markov regime switching approach for hedging petroleum commodities*, proposes a new way to estimate time-varying hedge ratios and compares it with several other benchmark methods to establish its accuracy. The innovation is in generalising the computation of hedge ratios to allow for both discrete shifts in the distribution (MRS) and GARCH effects. Moreover, the inclusion of the error correction mechanism in the regime switching framework enables us to examine whether the speed of adjustment of spot and futures petroleum prices to the long-run relationship changes across different regimes introducing an informative link between volatility and cointegration allowing for both time dependency and asymmetric behaviour across different states in the market. The suggested two state MRS vector error correction GARCH model shows improved in-sample fit and superior forecasting performance for both long and short hedges. Overall, by identifying time varying state dependent hedge ratios for normal and highly volatile periods, market agents may be able to obtain significantly superior gains, measured in terms of variance reduction and increase in utility.

The sixth chapter, *Petroleum Term Structure Dynamics, Inter-Commodity Dependencies and the Role of Regimes*, proposes non-linear multivariate equilibrium models of the term structure of correlated petroleum forward curves. We decompose the term structure into level, slope and curvature shocks and examine their mean-reverting and co-integrating properties in a futures spreads setting. Then, for the first time in the literature, a multi-regime multivariate MRS model is fitted to describe the risk factors, motivated by the fact that factor-specific features are typically inherited by the asset returns. In addition, we extend this model by allowing for independent switching among factors and commodities, in an attempt to capture the complex interaction of petroleum market mechanisms and accommodate several stylised features of forward curves which are observed real life phenomena¹. We find evidence in favour of the existence of a long-run relationship between level and slope factors, however, curvatures are found to be mean-reverting to commodity-specific equilibria. Results indicate that each regime clearly differentiates two distinct market dynamics for both the conditional mean and the volatility of the underlying process. Moreover, it seems that when one market is in the low and the other in the high variance state, it is more likely to observe lower correlations. Although the evolution of the oil term structure in the market has important implications in the fields of energy risk management and derivatives pricing, the issue of predictability of oil price curves has surprisingly received little attention. While the model can in principle be employed to analyse interrelationships of correlated petroleum futures curve dynamics, we also aim to fill in this gap in the literature by providing a new unified approach to obtain forecasts of the term structure of futures spreads, the variance-covariance matrices and risk management downside risk measures. Results from these exercises indicate that the multi-regime factor model can sometimes achieve significant gains compared to competing models.

In the seventh Chapter *Concluding remarks and further research*, we conclude by summarising the main empirical findings of this study. We also examine some common themes that appear throughout the thesis and outline potential interesting and challenging paths of future research as directed by the findings of this thesis.

All empirical applications serve the purpose to analyse which modelling technique is superior by employing appropriate benchmarks. The benefits of modelling and forecasting time-varying risk are evaluated appropriately using relevant loss functions. Robustness with respect to data snooping bias is also addressed by employing contemporary methodologies based on

¹ For instance, it is quite common to observe simultaneously high volatility in the product market and low in the crude oil market, due to the presence of backwardation and contango in the two curves. The high price volatility in the product market may be due to refining capacity constraints no matter whether crude oil production flows smoothly.

bootstrap simulations; this way, we also assess the practical relevance of taking time-variation of model parameters in a Markov framework.

To conclude, all the above topics have never previously been examined in the energy economics literature in a similar approach as offered by this thesis, thus making its contribution an original source of reference for academics and a practical tool for practitioners. The findings of this thesis have important implications for energy market participants that deal with trading and risk management as well as energy market operations, such as refining and budget planning, by providing valuable information on oil price differentials, volatility behaviour and co-dependence as well as their predictability. Overall, market agents may be also able to improve the forecasting accuracy and enhance the performance of their hedges.

1.3 Summary of Thesis Structure

The original contribution of this work commences in Chapter 4. The empirical body of the thesis involves Chapters 4 to 6. Note that each chapter covers a topic on its own, so that they can be read independently of previous and subsequent chapters. Part of Chapter 4 has been published in *Energy Economics* (Nomikos and Pouliasis, 2011). Part of Chapter 5 has been published in the *Journal of Banking and Finance* (Alizadeh, Nomikos, and Pouliasis, 2008) and an earlier version was presented at the Commodities and Finance Centre (CFC) Conference of Birkbeck University in London. The specific organisation of the thesis follows the objectives mentioned above in section 1.2 and the remainder of this study is organised as follows:

Chapter 2 offers an outlook of energy markets and the market structure and also provides some basic background on energy risk, risk measurement and risk management. Chapter 3 provides the necessary literature review on the employed models and the mathematical foundation that will be used throughout the thesis; it is also meant to fix notation. Chapter 4 is the first empirical study of this thesis. It deals with volatility forecasting and risk quantification of petroleum futures suggesting the use of Markov regime switching augmented models of the conditional moments. Chapter 5 proposes a new method of futures hedging using multivariate Markov regime switching vector error correction models with conditionally heteroscedastic error structure. Chapter 6 deals with the information content of the dependence structure between correlated petroleum futures curves. Using a factor decomposition of the term structure of futures prices, factors of petroleum spreads are modelled as multivariate MRS models and are used to replicate volatility and correlation dynamics. Finally, Chapter 7 concludes this thesis and gives directions for future research.

Chapter 2

Introduction to Oil Markets and Energy Risk

2.1 Introduction

Petroleum commodities constitute a relatively young market that has emerged over time to be a vital resource of our modern civilisation. In ancient times, oil was collected from oil seepages and it was not until 1859 when the first successful commercial oil well was drilled in Titusville, Pennsylvania, signalling the beginning of a new era for mankind. Early drillers began large scale oil production (mainly of kerosene) in the 1860's and, as an effect energy prices for illumination were significantly reduced. With more quantities being released into the market, oil prices in the US experienced a sharp decrease falling from \$10/bbl to \$0.5/bbl by the end of 1861. Although crude oil was still principally used for lighting, the new industry started to evolve with the development of sophisticated and more efficient technologies in the refining process. As a result, the systematic exploration, extraction, production and refinement of crude oil for commercial use started growing at a very fast pace. By 1890, kerosene was the only major by-product of oil until gasoline and fuel oil came on scene. Since then, rapid technological changes eventually led to the dominance of petroleum industry as a source of energy. Gasoline as a motor fuel became commercially viable after the invention of thermal cracking in 1913 and, with the boom of the automobile industry in the 1920's, a new market was created for oil, not only as a fuel but also as asphalt to construct roadways. Furthermore, heating oil was first produced on a large scale basis for heating purposes in the 1930's but was widely used only after 1945. The switch of railroad locomotives and ships from coal to oil, as well as the growth of the aviation industry critically raised petroleum demand for transportation. Soon after World War II (WWII), oil replaced coal and steam power and essentially became the main contributor of energy for transportation and commercial purposes.

More than half a century later, crude oil has steadily increased its share of providing energy for human activity in every sense. Since 1965, both oil production and consumption have almost tripled. Households, corporations and governments are all involved in the industry,

either as direct or indirect consumers, thus verifying, the significance of hydrocarbons which are indispensable for transportation, industrial and residential uses to our urbanised society. As a result, it is not surprising that oil is the world's most actively traded commodity with recorded trade movements in excess of 50 million barrels per day (bpd¹) over the last 5 years; this figure represents approx. 65% of the world's oil production. Additionally, in 2004 world consumption exceeded the threshold of 80 million bpd and has remained above that ever since, reaching its maximum in 2007 at 85.6 million bpd. Even though recent technological advances enhanced the development of substitute energy sources, oil delivers superior efficiency of use; therefore, industries are still vastly dependent on oil. The dominant position of oil in the energy sector is also apparent from the fact that oil corresponded to more than a third of the worldwide energy consumption for the years 2008 and 2009; second and third place was occupied by coal and natural gas with shares of 29% and 24%, respectively.

In this chapter we describe the structure of the oil markets. The next section presents the stylised facts of the world oil market. This is followed by a presentation of the demand and supply framework and the role of OPEC (Organization of Petroleum Exporting Countries). Section 2.4 offers a historical overview of fluctuations in oil prices as far as the 1860's. Section 2.5 is devoted on oil price volatility. Section 2.6 provides an introduction in petroleum futures markets and gives a brief note on the role of market speculation in price formation. The next section deals with risk management with separate sub-sections on quantifying oil price risk and minimum variance hedging whereas we also highlight the need of sound risk measurement and risk management strategies in the modern energy markets with a case study: the hedging debacle of Metallgesellschaft. Section 2.8 offers an introduction to oil term structure of futures and, finally, the last section concludes.

2.2 Main Features of the Oil Market

This section provides an overview of facets that make the oil market special. The analysis describes the environment that the industry operates in and the inherent energy sector risks. Throughout the last century, the specific market characteristics and the rapidly evolving setting in the hydrocarbon market have added several complexities to the price determination process. As for any other commodity, the price of oil and its products is determined by the interaction of demand and supply. Yet, understanding the oil markets goes far beyond basic economics. Key elements such as the response of the market to several events (such as wars),

¹ The statistics used in this chapter are from the British Petroleum Review 2010 report unless otherwise stated.

the nature of the commodity, the structure of the industry, the regulatory framework, and the behaviour and interaction of market participants (such as OPEC) define the mechanics of oil fundamentals and bear immense risks for the participants of the sector. The process of searching for economically viable oil wells until usage by the ultimate consumer divides the supply chain into three segments: upstream, midstream and downstream². It is widely accepted that the oil business has an oligopolistic structure with high-cost producers operating at maximum capacity and low-cost producers controlling the excess supply to satisfy demand surges or even cut production to balance market shares and protect their interests.

Market participants usually price crude oil at a discount or premium with respect to particular benchmarks and subject to differences in quality and location-specific characteristics. There are many specifications of crude oil and these are based on the chemical composition and physical properties. Oil is a non-standard commodity and in particular, its sulphur content (sweet or sour), gravity (light or heavy), viscosity and acidity define how easy it is to be refined, affecting, as a result, the operating costs and the refining yields. In the US, the most prominent crude oil grade and the primary pricing marker for North American crude is West Texas Intermediate (WTI) which is a high quality crude with small amounts of sulphur. In Europe, the Brent blend crude oil, sourced from the North Sea oil fields, constitutes the benchmark for pricing European, Middle Eastern and African crude oil. This is also a high grade crude but of relatively lower quality than WTI. Another benchmark is the heavy sour Dubai/Oman crude, a marker for Middle Eastern crude sold in the Far East markets. Heavy crude oil achieves lower yields of the more valuable light distillates and higher yields of the less valuable heavy distillates, thus reducing the refining industry's profit margins. Moreover, OPEC collects pricing data for the different types of export crude produced by its members so as to provide systematic crude oil pricing information. This is known as the OPEC Basket, which is heavier than WTI and Brent.

In the literature, there is evidence of strong regionalisation across oil markets (Weiner, 1991). However, more recent studies such as Gülen (1996) and Kleit (2001) support the

² *Upstream* refers to exploration and production. The capital intensive nature of this industry and the varying costs (conditional on various geopolitical factors, climate, land or offshore wells etc.) - call for reliable and precise valuation techniques. Sophisticated methods are employed including gravity, magnetic and seismic surveys before the phase of exploratory oil drilling (wildcat) to further reveal the geological formation of potential reservoirs (well logging). On completion of successful projects the well is equipped with a drilling rig and necessary apparatus to facilitate extraction. *Midstream* refers to the collection, storage and transportation of crude oil to refiners. *Downstream* refers to the refining and distribution of the processed products. Refineries are chemical plants that convert crude oil to light, middle and heavy distillates. Crude is heated and separated into its component hydrocarbons (fractional distillation) which are then purified - further chemically altered - to fit commercial purposes (desulphurisation). Heavy cuts are usually reprocessed using various forms of cracking (thermal, catalytic, hydro) to manufacture more valuable lighter distillates, resulting improved yields.

hypothesis that crude oil markets have become more unified and move closely together. Milonas and Henker (2001) show that the price spreads of different crude grades represent variations in quality and regional supply/demand as well as transportation costs, seasonal factors and volatility. For instance, due to the location factor, squeezes or seasonality in one market can switch the price differential from positive to negative and vice versa, irrespective of quality. Furthermore, they find that the Brent and WTI markets are not fully integrated and there may be periods that prices evolve independently in the short term. Price differentials increase or decrease with certain bounds, linked through the cost of carry relationship, and any divergence can be restored by arbitrage. Therefore, although the oil market is global, wide variations in price differentials are common and a single unified price cannot serve the industry accurately; for instance, in the case of forecasts.

As a mineral commodity, every owner of crude oil has the option to either extract the resource now or hold on to it to extract it in the future. Based on the seminal paper of Hotelling (1931), the per unit price of a non-renewable resource over its extraction cost grows at the rate of interest (at least in the long run) and the production trajectory depletes monotonically until exhaustion. If oil prices rise at a slower rate, all producers would keep their stocks in the ground, decreasing the current supply in the market and thus increasing the current price. Oil exhaustibility implies that the production and consumption levels of upcoming periods depend on the production and consumption of past periods. Unlike standard commodities, producers receive a scarcity rent, a premium representing the compensation for holding stocks in reserve for use tomorrow. Petroleum markets exhibit incomparable marketplace dynamics because, on the one hand, both developed and developing economies are highly reliant on oil-driven technologies for sustaining their regular energy requisites or fuelling their growth and industrialisation rates and, on the other hand, there are marked geopolitical asymmetries regarding the location of the supply and demand centres. This in turn, has triggered political intervention originating from the need to guarantee energy security in the long-term, buffer against physical disruptions in the short term and mitigate market power risks.

Figure 2.1 demonstrates the fundamental imbalances in the setting of the world oil industry for 2009. For instance, production in North America, Europe and Asia Pacific regions' is much less than the levels of consumption as opposed to Middle East, Former Soviet Union and Africa. In addition, proven reserves are concentrated in the Middle East region. As of the end of 2009, 20% of the proven reserves are located in Saudi Arabia whereas, in total, 57% of the reserves lie in the Middle East. Total proven reserves are currently estimated to be 1.3 trillion barrels. Regarding oil consumption, Middle East is responsible for more than one third of the world's production but consumes less than 10%. United States and Europe (excl. Former

Soviet Union) are the largest importers of crude and consume jointly around half of the global production but produce less than 20% of it. In addition, China and Japan are also large consumers with shares of 10% and 5%, respectively and they have less than 5% stake in world production together. In South and Central America, Venezuela holds 87% of the proven reserves in the region (172 billion barrels), with a share of less than 1% of the total consumption and more than 3% of the total production. Nigeria, Angola, Algeria and Libya produce jointly less than 10% of the world production but African countries, in total, consume less than 4%.

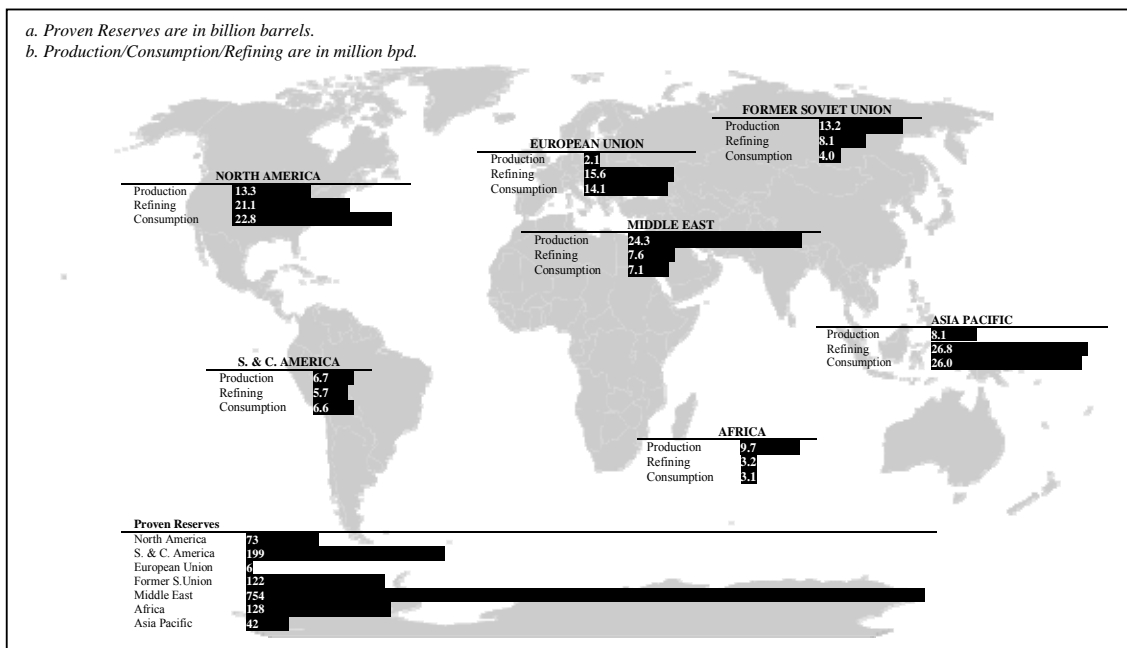


Figure 2.1: Imbalances in the World Oil Market Structure

The outlook has not been constant throughout the years. In America, for instance, since 1980 the Northern part's proven reserves have depleted by 20% to 73 billion barrels, whereas in South and Central America, reserves have experienced a sevenfold increase to 199 billion barrels. Overall, production and consumption have increased by 27% and 37%, respectively, and proven reserves have doubled with OPEC members controlling more than 77% of the reserves compared to 67% in the 1980's. Additionally, Middle East has expanded its domineering supply-side role throughout the years.

Figure 2.2 confirms the role of Middle East as a swing producer since most of the spare capacity is concentrated in this area. The correlation of the production deficit of OECD (Organisation for Economic Co-operation and Development) with the Middle East surplus since 1965 is 95%. The deficit is greater than the surplus which is balanced from Former Soviet Union, African and South & Central American imported crudes.




Figure 2.2: Middle East as a Swing Producer

Finally, refining capacity has developed steadily following the consumption requirements of each region. In 1980 refining throughput i.e. the quantity of crude being processed per day was around 60 million barrels (almost twice over that of 1965) with an estimated daily capacity (upper limit) of 80 million, as opposed to 73 million bpd throughput (91 million bpd capacity) in 2009. Before WWI, international trade of crude oil was too costly due to its low value and refineries were located near production oil sites. However, the market's nexus until then was the biggest producer and exporter of oil at the time; that is the US. With the demand boom after 1920, the gradually strengthening position of Middle East's production and the developments in the shipping industry, the market was essentially restructured with refineries being relocated next to demand poles since the transportation of low value crude could benefit from economies of scale.

2.3 Demand and Supply Framework

Traditionally, energy markets function with a unique structure of supply and demand mechanisms, which introduce a degree of complexity along with elevated levels of volatility. Economic theory asserts that excess demand (supply) results in an upward (downward) pressure on prices. Those forces that affect the price levels merit discussion at the outset because should any potential imbalances arise this will have a direct impact on volatility.

Crude oil is characterised by a *global demand curve* which is derived from the energy consumption rates of the finished and intermediate products, given the transportation, industrial and residential needs. Obviously this is directly linked to the global economic activity; hence population growth and the degree of industrialisation are of paramount importance. In general, demand for oil is a composite of global economic activity, demographics, competing energies

and consumer preferences, refining capacity, storage and advances in demand-side technologies. Price controls and government/organisation policies are also key elements of the demand side. For instance, importing countries influence the market via environmental regulations (e.g. Kyoto protocol, 1997), research and development of alternative energy resources or even political intervention to ensure sufficient flow of the commodity from secure supply sources and national investments in oil ventures. On the other hand, the product markets in the world differ from one place to another and are characterised by a *regional demand curve*. Different economic activities, climate, level of technological development and diverse lifestyles around the globe lead to different consumption patterns; for instance, 50% of the total oil product consumption of the US concerns light distillates as opposed to Europe where the prominent role is held by middle distillates.

In the short-term, global demand for crude oil may be mismatched with the underlying regional demand for petroleum products as a result of regional inventory building for products to meet seasonal demand and the timing effect of production. Regional consumption is susceptible to the refineries' flexibility in adjusting the yields, storage policies and capacity constraints, implying that when imbalances occur, international trade will accommodate the need for oil. Among the importing countries, as already mentioned, the largest importer and consumer of oil is the US. In 2009 over 11 million bpd were imported in the United States, with 20% representing petroleum products. Canada, Mexico and South and Central America feed more than half of the US oil needs, whereas imports from the Middle East and Africa account for more than 15% and 19% of the total figure, respectively.

Turning now to the supply side, this aspect concerns the amount of oil offered by producers based on the optimisation of their revenue. The capital intensive nature of the industry, the varying extraction costs over time, heterogeneity of the commodity and the depleting reserves that are occasionally augmented through exploration and development projects, are some of the stylised facts of the supply function. This side is a composite of proven reserves, estimates of undiscovered reserves, stocks, supply-side technologies to improve production process and rates of extraction, geopolitical uncertainty arising from the imbalances in production, as well as political events. Oil supply chain disruptions might occur at every stage of production i.e. upstream (extraction capacity, cost of drilling, environmental policies, etc.), midstream (transportation infrastructure, extreme weather conditions etc.) and downstream (availability and location of refineries, capacity constraints, taxes and legal systems etc.), affecting the price either by causing bottlenecks in the production process or by changing the overall costs.

In addition, the behaviour, role and interests of OPEC and non-OPEC suppliers add a certain complexity to the supply side, introducing the necessity to distinguish between two supplier profiles (see also next section). OPEC, founded in 1960 in Baghdad, consists of twelve countries controlling more than 70% of the world's proven oil reserves and is responsible for less than half of the world oil production (40% on average since 1965). After the first oil crisis of the 1970's non-OPEC suppliers increased their share and the resulting geographical dispersion of the oil fields served to smooth the supply process. However, non-OPEC production is associated with more technological difficulties and higher costs; for instance in North Sea and Alaska. Moreover, the industry involves six large multinational vertically integrated corporations (ExxonMobil, ChevronTexaco, ConocoPhillips, Total, British Petroleum and Royal Dutch Shell) and several smaller independent firms. The world's largest oil corporation in terms of both reserves and production is state-owned, the Saudi Arabian Aramco which was fully nationalised in 1980.

In summary, oil prices at any point in time should reflect the balance between supply and demand for crude oil. First, in the short run, both demand and supply curves are known to be very inelastic implying that supply shortages or severe positive demand shocks are translated to large price increases (Krichene, 2006). Inelastic demand is due to the fact that substitution and energy conservation requires huge investment and certain time to set up the proper infrastructure. Inelastic supply is due to the fact that releasing additional quantities in the market given low inventory levels is impracticable: neither non-OPEC members (they already operate at full capacity) nor OPEC members that control spare capacity (due to the timing effect of production) can respond instantaneously. Second, a more elastic, though still relatively inelastic, demand and supply is observed in the long run. This is due to the fact that more supplies can be brought in the market by increasing OPEC production, growing activity in exploration-development and utilisation of unexploited wells, technological advances that lead to more efficient use or even substituting oil with alternative energies. In such a setting of fundamentals, demand and supply shocks under tight market conditions are translated to large price movements, which in turn introduce increased volatility and have direct implications on the energy policies for both governments and companies susceptible to energy risk.

2.3.1 The OPEC effect

The strategic importance of oil was first depicted in the years after John D. Rockefeller and the Standard Oil Company came on the scene in 1870. From the beginning of the 20th century and after the break up of Standard Oil, it was obvious that *laissez faire et laissez passer*

did not fit in the nature of the industry and limited competition was always a distinctive feature of the market. After WWI and the dissolution of the Ottoman Empire, an early attempt for the creation of an international cartel took place in Achnacarry involving Royal Dutch Shell, Standard Oil and Anglo-Persian companies to stabilise the market, prevent price wars and freeze market shares - known also as the "As is" meeting, in 1928. The products of this association were the Red Line Agreement and the creation of the Iraq Petroleum Company (IPC) which controlled most oil of the Iraq region up to 1961: each partner agreed to jointly exploit any new reserves. Since the 1960's, a key factor in the price determination process is the behaviour of OPEC countries.

The bargaining power of the major producing countries is rooted in the strategic importance of oil due the degree of concentration of the reserves and the need for energy security. Outside the cartel several suppliers exist that do behave competitively, holding 23% of the proven reserves and possessing a 60% production share, on average, for the last 45 years. OPEC's mission, being the low cost producer, is to secure a regular supply of oil and smooth the market whilst at the same time ensuring fair prices for its members by coordinating its production output and allocating quotas. On the other hand, the response of non-OPEC countries to oil prices is relatively lower (Krichene, 2006). Oil prices do fluctuate depending on how OPEC calculates quotas and how its members comply with these decisions. For example, although producers do not incur any storage costs by altering production rates -oil is basically left in the ground- most exporting countries rely heavily on oil revenues and might be unwilling to lessen output.

The market power of the organisation has been challenged in a plethora of studies. The branch of empirical research studying the behaviour of OPEC starts with the seminal paper of Griffin (1985). Studying the supply functions for individual countries, the author concludes that most OPEC members behave as if they were part of a collusive cartel. Later, several studies attempted to analyse OPEC influence and the determinants of OPEC supply. For instance, Kaufmann et al. (2004) find that OPEC can manipulate real oil prices via altering production quotas and operable capacity whereas OECD stocks and the amount of cheating those quotas are critical aspects. Generally, literature is far from consistent on the issue but, overall, confirms the oligopolistic market character of oil; for a related review see Smith (2005). As Kaufmann et al. (2008) note "*There is no reason to expect a simple model to describe the production behaviour by members of an international organisation that consists of sovereign nations, which have vastly different geological endowments, economic structures, and political/social aspirations*". Another interesting observation is the fact that the reaction of OPEC members and the corresponding influence in the market is not clear-cut and has proven to be asymmetric and

complex with uncertain success rate. For instance, De Santis (2003) confirms that Saudi Arabia, the most powerful member of OPEC, will cut production to protect the revenues of her producers and the organisation members but at the same time has little incentive to adjust output accordingly in response to positive demand shocks. Of course, it is not a part of OPEC's objective to set unjustified price ceilings. However, maintaining healthy growth and preventing strong trends in energy substitution rates are natural concerns since persistence of high prices can be proven damaging for the whole market in the long-run.

Real practice shows that OPEC's political economic response depends on the market sentiment and expectations regarding the ability of the cartel to operate effectively. Some OPEC meetings are overlooked whereas others create excessive notice, prompting speculative activity. Lin and Tamvakis (2010) carry out an event study examining how OPEC announcements influence major international crude prices and find that the magnitude of these effects depends mainly on the prevailing price zone. In particular, they identify low, normal and high price regimes and report that quota cuts (increases) lead to price increases (decreases), except in the low (high) price band. In theory, the pricing function does rely on OPEC and the organisation is expected to gain power, given the current and potential tight fundamentals. The ratio of OPEC to non-OPEC proven reserves has grown steadily over the years from 2.7 in 1980 to 5.7 in 2009, while the corresponding production ratio fluctuates around unity for the last 5 years, with an increasing trend since the 1980's. Shrinking reserves in the North Sea, US and Mexico and the concentration of supply capacity at high political tension regions define new risk perceptions for demand, supply and volatility effects. Finally, although the process of adjusting OPEC supply can only hope to put pressure on prices, given the high response of the oil market to news, OPEC meetings and decisions, the degree of timely response, flexibility to follow quotas and expectations regarding compliance and credibility of its members as well as speculation of the overall cartel's behaviour and interests, do contribute to raise volatility in the short term. This is especially true in turbulent periods where recovery takes time and planning.

2.4 Historical Overview of Fluctuations in Oil Prices

In the 1860's the industry was characterised by considerable levels of volatility mainly due to market disorder, poor rules of ownership and the particular market structure comprising small independent firms or individuals. As Dvir and Rogoff (2010) note, during these early years, the persistent growth shocks due to intense industrialisation, on the demand side, and the uncertainty surrounding easy consumer access, on the supply side, had a significant effect on price fluctuations. Moreover, oil production was concentrated mainly to north-western

Pennsylvania and railroad monopoly over the transportation of oil was a constraining factor against producers' margins. Consequently, temporal variations and cycles were common. After 1870, Standard Oil initiated a tactic to exploit that disorder and by gradually acquiring most of the existing businesses it managed to transform the market into a monopoly.




Figure 2.3: Crude Oil Prices, 1869 - 1969

By 1880, Standard Oil controlled 90% of the oil production and trade in the US at all stages of production, which had a stabilising effect on prices: from \$8/bbl in 1864, prices dropped and remained in the \$1/bbl region up and until the dissolution of Standard Oil in 1911 (Figure 2.3). In fact, prices remained relatively low for many years forward regardless of the political turmoil in the international arena such as the Mexican nationalisation, the Iranian nationalisation, the Suez Canal closure (which increased considerably transportation requirements for oil) and the apparent demand surge. This is attributed to the end of the railroad monopoly over the transportation of oil in the US with the construction of Tidewater, the continuous discovery of new reserves, such as the East Texas oilfield in the early 1930's, the rise of Middle East as a major producer and the expanding production so that depleting reservoirs were not yet an issue. As a result, crude oil was in abundance and producing countries would produce as much it was needed and importing countries could accumulate this way a safe amount of inventories, sufficient to absorb any potential demand shocks. However, with the strong demand growth and given that oil was now cheaper than coal, dependence on oil mounted in the 1960's.

Due to the supply abundance, prices remained depressed and stable up to 1973 (Figure 2.4). With OPEC now in the scene and the excessive dependence of the industrialised economies on oil imports, the market was about to change fundamentally: the industry was evolving rapidly and becoming more and more sophisticated. To add, US production had already reached its peak, implying that excess capacity was concentrated exclusively in the Middle East. A falling dollar, in the beginning of the decade, meant that producers received decreasing streams of income for their production. Therefore, OPEC asserted its power and started negotiations.




Figure 2.4: Crude Oil Prices, 1970 - 2000

In 1971, after a process of nationalisation of western oil companies' concessions in the regions of oil exporting countries, OPEC raised tax rates, signed the Tehran and the Tripoli Agreement and decided on an increase in the posted prices to be followed by further increases in order to counterbalance the dollar depreciation. In October 1973, Saudi Arabia, Libya, and other Arab states proclaimed an oil-exports embargo to retaliate for the US decision to assist Israel in the Yom Kippur War. Although short-lived (until March, 1974) the embargo led to what is known as the *first oil price shock* with the cost of crude more than quadrupling to nearly 12\$/bbl. Yet, the oil craving of the economies was so intense that demand, production and exports remained strong throughout the 1970's. Consequently, the *second oil price shock* was not far ahead when another political event agitated the status quo: the Iranian Revolution in 1979 which resulted in the dissolution of the western oil companies in the region.

With the price of oil reaching \$40/bbl, a new phase begun and the entire humankind was forced to reconsider many issues. First, demand starts to drop in view of the high prices

since energy conservation and fuel switching was encouraged. Second, non-OPEC production was increasing since the two crises revealed that the value of politically safe reserves and energy security became once again a first priority issue for many governments (providing incentives for domestic exploration and development or alternative projects). Hence, new supply patterns emerged with high-cost suppliers such as the UK, Alaska and Canada increasing their share and bringing forward their own oil industries. Obviously, OPEC's control deteriorated and any decisions for production cuts did not have the desirable effect because they were offset from the non-OPEC supply.

Saudi Arabia's crude oil output, from over 10 million barrels bpd in 1980, fell to just 3.6 million bpd in 1985. By the mid 1980 prices had fallen to lower levels than the aftermath of the 1973 embargo. Saudi Arabia's decision to aggressively increase its market share and abandon the role of the swing producer also contributed to this end. In 1986 Saudis increased production by 45% to 5.2 million bpd and applied the netback pricing system, thereby bearing the price risk of their customers by guaranteeing refiners' profit margins. This led to a record low price of around \$10/bbl in 1986 causing struggle to all producers, OPEC and non-OPEC. Since then, several events contributed to a completely different price behaviour compared to the preceding two decades. Markets became more competitive and efficient and, especially after the introduction of derivatives in 1983, more transparent, liquid and open. The oil business was now responding much more frequently to the arrival of new information and political events (e.g. Persian sGulf Crisis, Nigeria workers' strike), weather conditions and supply/demand dynamics; once the balance was restored prices reverted back i.e. shocks became less persistent. We can observe (Figure 2.4) that after the Asian Crisis of 1998 a barrel of crude was traded for even less than \$10, a historical low of more than two decades.




Figure 2.5: Crude Oil Prices, 2001 – 2009

In July 2008, the refiner acquisition cost of imported crude in the US reached the \$130/bbl range and the corresponding WTI and Brent spot markets reached \$145/bbl. Apart from the rapid growth of demand for commodities in emerging countries (mainly China) and the relatively low US stocks, several events contributed to the apparent upward trend (see Figure 2.5), including the 9/11 attacks, the Venezuelan strike in late 2002, the US military action in Iraq after 2003, North Korea's missile launches, the Hurricanes in the Gulf of Mexico, the conflict between Israel and Lebanon in 2006 and the Iranian nuclear brinkmanship. Changes in federal oil policies, a falling dollar and of course the sizeable entry of index speculators into the futures markets also contributed to the July peak and what some call the *third oil price shock*. The outcome was for prices to collapse below \$40/bbl due to a drop in demand for oil in combination with oversupply, and the financial crisis which led to the subsequent deleveraging of commodity portfolios from risky assets. To conclude, the third oil price shock was realised in the course of 5 years (2003-2008) rather than being caused essentially by OPEC members (1st and 2nd oil price shocks). Nevertheless, all three shocks were followed by economic recessions.

2.5 Oil Price Volatility

Oil prices were mainly controlled by the Seven Sisters³ during the 1950's and 1960's and by OPEC during the mid-1970's to the mid-1980's. After the two oil price crises and the introduction of the market-based pricing system in the 1980's individual investors and energy market participants have always been faced with high levels of uncertainty. As previously demonstrated, the price of oil is determined by distinctive supply and demand interactions augmented by a complex game of interdependencies among market participants. As a result, price volatility has become an important feature of the industry due to the detrimental effects that can occur from under- or over- estimating its impact on the revenue and cost sides (in all lines of oil-related business) and in general, the cash flows and earnings from relevant investment strategies. It is vital for oil price related decision makers such as governments, firms, individuals and multinational organisations (such as OPEC) to understand, quantify, monitor and control the risk matrix associated with the petroleum industry.

Oil price volatility has been studied on several aspects since it has economic consequences of general interest. First, it plays an important role to regional and global economic activity because sharp price fluctuations lead to economic instability for both oil-

³ This is a common term referring to the seven major Anglo-American oil companies that structured the *Consortium for Iran* after WWII. These were Royal Dutch Shell, Anglo-Persian Oil Company, Gulf Oil, Texaco, Standard Oil of New Jersey, Standard Oil New York and Standard Oil of California. Operating as a cartel, they essentially dominated the global oil markets until OPEC raised its prominence.

exporting and oil-importing countries affecting government and companies corporate policies as well as individuals. Oil volatility has a certain impact on the macroeconomy (see Lee et al., 1995; Federer, 1996; Hamilton, 2003 and Chen and Chen, 2007) and on the financial stock markets (see Sadorsky, 1999, 2003 and Driesprong et al., 2008). Second, oil price volatility has an unfavourable effect on investments by increasing the uncertainty regarding future cash flows (especially on high capital intensive projects such as exploration and production) and causing project delays (see for instance Pindyck, 1991) whereas persistent uncertainty induces the longer term effect of reallocating the available resources to less volatile sectors and competing energies. Third, Pindyck (2004a) argues that volatility affects the demand for storage, thus, the firms' operating options and opportunity cost of the current production. Finally, being a pivotal input to the value of contingent claims, volatility behaviour is indispensable for pricing, hedging and evaluating strategic alternatives.

Economic theory manifests that asset returns tend to exhibit volatility clustering; in other words, large (small) price changes tend to be followed by large (small) price changes (Mandelbrot, 1963). This indicates that oil price changes might follow time-varying distributions and, therefore, if this is the case, risk should also be time-varying; this has been a central issue in various studies. To accommodate autocorrelation in the squared return process, Engle (1982) provided the first insight into modelling the time dependency of volatility in the financial markets with the development of the Autoregressive Conditional Heteroscedasticity (ARCH) models; later generalised by Bollerslev (1986) (GARCH).

To provide an overview of volatility movements, Figure 2.6 displays the volatilities and correlations of WTI crude oil spot prices at Cushing Oklahoma and conventional gasoline and heating oil at the New York Harbour as estimated by a simple tri-variate GARCH(1,1) model. Annualised volatility lies within 22%-90% for WTI, 27%-105% for gasoline and 22%-120% for heating oil. Largely, the time varying nature of risk in the oil markets has been the norm in the literature for modelling either individual petroleum commodities (Kang et. al, 2009 and Agnolucci, 2009) or portfolios of such commodities (Haigh and Holt, 2002). As an example, Pindyck (2004b) finds significant fluctuation in crude oil volatility with short-lived shocks (reporting a half life 5 to 10 weeks) and a small upward trend which is however of no economic significance; this implies that the volatility path is more relevant to shorter duration oil-based derivatives rather than long term real options. Narayan and Narayan (2007) show that oil price shocks have asymmetric and persistent effects on volatility. However, this relationship weakens or even disappears across sub-periods implying that oil price behaviour experiences sudden changes.




Figure 2.6: Annualised Spot Volatilities & Correlations of Main US Petroleum Commodities and Crack Spreads

Volatility behaviour has also implications for refiners or portfolio managers since dynamic interdependence of risks implies correlation risk and diversification effects for market participants. For example, the annualised volatility of the 3:2:1 crack spread varies from 13% to 76% with correlations of WTI vs. heating, WTI vs. gasoline and heating vs. gasoline ranging from 10% to 90%, -1% to 80% and 6% to 81%, respectively. When producers, refiners, consumers, portfolio managers and, in general, investors, are risk averse there is reason to mitigate the risks arising from oil price volatility. Theoretically, to eliminate the effects of severe price variation there are several choices available such as integration, diversification, and inventory management. However, for agents involved in the physical market this entails huge investment whereas the efficiency of such procedures to limit risk is questionable. Usually the most efficient way to deal with market risk is by using derivatives (paper contracts) such as forwards, swaps, futures and options. These are discussed next.

2.6 The Petroleum Futures Market

Derivative securities are financial contracts whose value is derived from some underlying asset. Over-the-Counter (OTC) trading of such contracts has existed since ancient times, even for options. A form of bilateral tailor-made forward agreements was always needed to facilitate trades, improve communication between the buy and sell side and deal with the typical arrangements of transportation, delivery, regulation, etc. Modern derivative instruments date back to the mid 1860's with the introduction of standardised commodity future contracts at the Chicago Board of Trade. The birth of exchange-regulated contracts was mainly due to the concerns of U.S. merchants to improve the effectiveness of the commercial marketplace by ensuring liquidity (bringing together potential buyers and sellers), creating the opportunity to hedge against adverse price changes and most importantly mitigating credit risk which was a severe hazard of the financial system at that time. A petroleum futures contract is a legally binding standardised agreement between two parties to buy or sell a given amount of the underlying commodity at an agreed forthcoming date with pre-specified arrangements regarding the quality, location and method of delivery. A party who holds contracts at the expiration date is obliged to make or take physical delivery unless otherwise specified (some contracts are only cash-settled). Throughout the life of a contract the buyers'/sellers' gains or losses are daily settled (marked to market). The security and performance of the contract is guaranteed by the exchange (no credit risk). To initiate a trade two types of margins are maintained: an initial deposit to the exchange's clearing house and a daily variation margin which covers deductions that arise after the daily revaluation of the futures portfolio.

In 1872, the New York Mercantile Exchange (NYMEX) was established by a group of dairy merchants and was later developed to the largest energy commodities exchange. Sixteen years later the National Petroleum Exchange in Manhattan facilitated the first oil futures-like derivative instrument when John D. Rockefeller issued certificates against oil stored in pipelines. The first energy exchange-regulated market was launched in 1978 with the introduction of the NYMEX heating oil futures contract in view of the first oil price shock and the subsequent upward trend in oil prices. Later, in 1983 and 1984 crude oil and gasoline futures were also introduced, respectively. Since August 2008, NYMEX has been integrated with Chicago Mercantile Exchange (CME Group) and now constitutes the most mature futures market trading petroleum futures on WTI, heating oil, unleaded gasoline, Brent, 3:2:1 and 1:1 cracks as well as New York Jet Fuel, Gulf Coast Heating Oil, European Jet Kerosene, among others. Another major exchange providing oil derivative contracts was the International Petroleum Exchange (IPE) launched in 1980 in London which is now known as the

Intercontinental Exchange (ICE) after its acquisition in 2001 by the homonym company. The first contract for gas oil futures was introduced in 1981 followed by the Brent crude oil contract in 1988. The exchange trades a variety of energy futures including gas oil crack, Brent-WTI spread and heating oil - gas oil spread and US oil futures, among others. Other exchanges that trade oil-related contracts are the Tokyo Commodity Exchange (TOCOM) since 1999 (crude oil, gas oil, gasoline and kerosene futures) and the Dubai Mercantile Exchange (DME) since 2007 (crude oil futures). Finally, Singapore Mercantile Exchange (SMX) opened the third quarter of 2010 and launched petroleum futures (WTI and euro priced Brent).

NYMEX WTI contracts are traded for all consecutive months within the current and the next 5 years. Contracts are also listed for every June and December delivery up to 9 years forward. New contracts are listed on an annual basis, after expiration of the December contract. Each contract is traded until the close of business on the 3rd business day prior to the 25th calendar day of the month preceding the delivery month. If the 25th calendar day of the month preceding the delivery month is a non-business day, trading shall cease on the 3rd business day prior to the business day preceding the 25th calendar day. Delivery shall be made free-on-board (FOB) at any pipeline or storage facility in Cushing, Oklahoma, whereas delivery shall come to effect between the 1st and last calendar day of the contract month. The size of the contract is 1,000 barrels and is quoted in US dollars per barrel (US\$/bbl). NYMEX gasoline contracts are traded for all 36 consecutive months and the underlying is Reformulated Gasoline Blendstock (RBOB) for delivery at the New York Harbor. Each contract is terminated the last business day of the month preceding the delivery month. Contract size is 42,000 US gallons is quoted in US\$ cents per gallon. Similar are the specifications for NYMEX heating oil contracts based on No. 2 fuel oil, deliverable at New York Harbor⁴.

Turning to ICE, Brent crude oil contracts are traded for all deliveries within the next 72 consecutive months. Six additional contracts are listed for June and December deliveries for up to 9 years forward. Each contract is traded until the close of business of the 15th day before the 1st day of the delivery month. If such day is a non-business day, trading shall cease on the next business day. The ICE Brent crude contract is a deliverable contract containing 1,000 barrels of crude, based on Exchange Futures for Physical (EFP) with an option for cash settlement and

⁴ Note that the two refined products, accounting for more than 70% of the refining yield, have undergone changes in their respective contract specifications throughout the years. First, heating oil was regularly traded for all deliveries within the next 18 months up to 2007 when delivery months increased to 36; the specifications of heating oil contract is also expected to change due to regulatory changes that intend to reduce the sulphur content of the commodity in the New York Harbour area - the last listed contract expires in January 2013. Second, in 2006 the unleaded gasoline contract with 12 forward delivery months was replaced by the RBOB for Blending with 10% Denatured Fuel Ethanol (92% purity), a change imposed to meet government emissions regulations.

reference price the ICE Brent Index for the day following the last trading day of the contract. Settlement price is the weighted average price of all trades during a 3 minute settlement period commencing at 19:27, London time. Size and quotation are the same with WTI. ICE gas oil contracts are traded for all deliveries for 36 consecutive months forward, then quarterly out to 48 months and then half-yearly out to 60 months. Contracts expire at 12.00 hours, 2 business days prior to the 14th calendar day of the delivery month. Its underlying physical market is heating oil barges (or coasters up to 10,000 deadweight) or in-tank or inter-tank transfer from an Exchange Recognised Customs and Excise bonded storage installation or refinery delivered in ARA (Antwerp, Rotterdam, Amsterdam). Contract size is 100 tonnes of gasoil at a density of 0.845 kg/litre and is quoted in US\$ and cents per tonne.

Figure 2.7 (left panel) displays the annual aggregated volume of crude oil futures traded in NYMEX and ICE. The dashed line is the annual world production. Notably, as early as 1990 the traded volume of the recently launched NYMEX contract matched global output. Apart from hedgers, the success of exchange traded oil-related contracts stimulated trading activity, attracting a broad range of new participants such as portfolio managers and index speculators. The same figure (on the right) depicts the evolution and variability of daily traded volume of the nearest to expiration NYMEX WTI contract since 1988. Again, the upward trend after 2002 is noticeable.

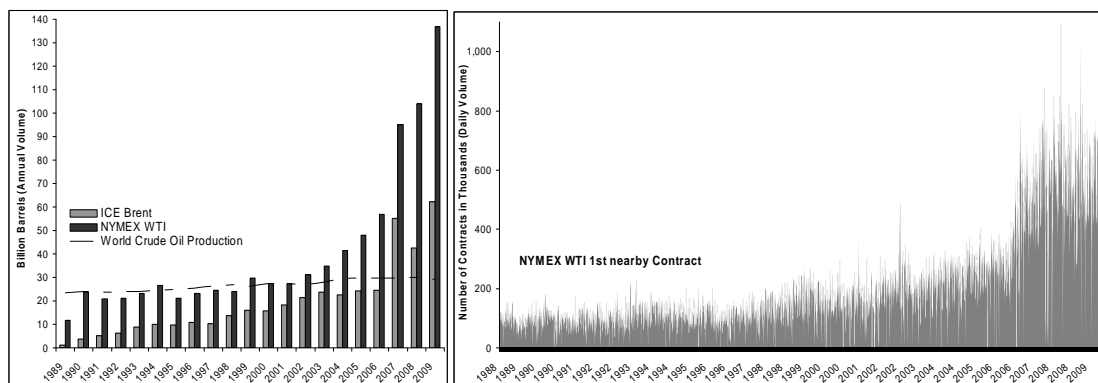


Figure 2.7: Trading Activity of NYMEX & Brent Crude Oil Futures

Fleming and Ostdiek (1999) examined the impact of energy derivatives trading on the crude oil market and found evidence of a sudden increase in volatility after the introduction of crude oil futures (for a period of 3-4 weeks). In the longer term, despite a rise in volatility estimates for the following year, this effect cannot be disentangled from the impact of several coincided exogenous factors such as deregulation and the rapid growth of the industry. Overall, the relation between futures trading activity and spot market volatility showed that futures

trading improved both depth and liquidity, having a mitigating impact on volatility rather than destabilising the underlying market. Furthermore, focusing on the subsequent introduction of new derivative contracts - including options-, volatility effects seem to disappear as the market gradually becomes more complete.

The rationale for the existence of derivative markets is to facilitate price discovery and offer the means to price and hedge risk. There is a plethora of studies in the literature that have investigated the extent to which this dual role of the futures market is indeed performed. For instance, Moosa (2002) showed that futures account for a rather high portion of the price discovery function (60%) and are also successful in transferring the risk from participants who want to reduce the variance of their portfolio to participants that are willing to bear those risks; confirming earlier studies supporting futures as the leader in the price discovery process (Silvapulle and Moosa, 1999). This is not a surprising fact though since the physical oil market is relative illiquid characterised by a declining physical volume of the main benchmarks (WTI and Brent) and is also much less transparent with fewer participants compared to futures. However, spot prices endow supply and demand forces with economic substance and it is not surprising that they also play a key role rather than just being satellites of derivatives prices. Moreover, futures have been found to be unbiased predictors of future spot prices (Crowder and Hamed, 1993; Schwarz and Szakmary, 1994 etc.) On the other hand, hedging effectiveness has also been a fashionable topic and several studies support the ability of oil contracts as a risk management tool. Due to liquidity limitations it is accepted that the “*most effective hedge is the nearby contract*” (Chen et al., 1987) where trading volume is mainly concentrated. That said, longer term hedges are expected to be less effective mainly due to varying convenience yields that create basis risk (lower futures-spot correlation) in the process of rolling futures positions forward. Haigh and Holt (2002) analyse the problem of refiner who is exposed to crack spread fluctuations and his ability to trim down efficiently the price risks involved, using NYMEX futures. Results illustrate substantial rewards in terms of risk reduction and certainty equivalent income.

2.6.1 Speculation and Investor Behaviour

After the development of organised exchanges, derivatives products expanded giving easy access to the industry. They increasingly gained importance, motivating a large entry of new financial players. According to the purpose of involvement in the market, active investors can be classified to hedgers, arbitrageurs and speculators. Hedgers provide the founding economic substance of the market linked to the physical underlying market. Their aim is to

reduce the risks to which they are exposed. Arbitrageurs play a correction role in the market; engaging simultaneously in two or more related markets, they restore the balance by exploiting economically meaningful counterfactual relationships e.g. between two different quality crudes, two different location crudes, spot vs. futures, deferred futures, crude vs. products etc. Their aim is to profit from deviations of fundamental relationships. Lastly, speculators are investors who willingly bear price risks in view of the profit potential; in effect, they constitute the polar opposite of hedging. Especially nowadays, there is always a speculative part in the demand for oil, either by hedge funds, commodity traders, institutional and individual investors with primary reasons to exploit the tight fundamentals (i.e. stagnant supply and demand surge) and/or use petroleum commodities as a diversification tool.

In 2006, the US Senate Subcommittee on Investigations (*“The role of market speculation in rising oil and gas prices”*) reported that increased speculation activity in turn swelled paper demand and prolonged the bullish markets. The speculative money released was believed to have changed the fundamentals i.e. crude oil market was characterised by both high prices and large inventories. In 2008, Michael Masters with a written testimony to the Committee on Homeland Security and Governmental Affairs of the US Senate found that assets allocated to commodity index trading strategies had increased from \$13 to \$260 billion between 2003 and 2008. Over the same five-year period, index speculators demand for petroleum futures has increased by 848 million barrels (equivalent to the increase in China’s demand over the last five years - 920 million barrels) and have stockpiled futures with an underlying quantity of 1.1 billion barrels (eight times higher than the oil added to the Strategic Petroleum Reserves of the US).

Although speculators serve an important role regarding market efficiency, transparency and enhancing liquidity, some side effects cause deviations from the equilibrium prices and increased volatility, at least temporarily. Kaufmann and Ullman (2009) confirm the significance of speculation in the oil market, concluding that high prices and the upward trend up to March 2008 has been indeed triggered by a change in fundamentals with increasing demand and sluggish non-OPEC supply. This setting, being identified by speculators, caused oil prices to overshoot their fundamental equilibrium, slowing in effect demand growth and economic activity.

2.7 Petroleum Price Risk Management

Since the first oil price shock, oil price volatility has clearly demonstrated the potential to significantly impact the ordinary conduct of business of many companies and consumers’

income. Numerous pressures arising from geo-political instability, unexpected or extreme weather conditions, political decisions, refining capacity constraints, limited and concentrated spare supply and sudden demand surges can and often do create price swings which affect revenues, financial performance and elevate budgetary requirements. Derivative markets provide an essential tool to transmit price exposure and reduce the portfolio uncertainty. Reducing cash flow volatility also entails indirect benefits for companies. First, it reduces the cost of financial distress by avoiding large firm value changes or even limiting the downside during bankruptcy proceedings. Second, it reduces the expected value of income tax payments by smoothing taxable earnings throughout time. Third, it can improve efficiency by avoiding under-investment. Fourth, it can increase debt capacity and lower the cost of funds by reducing the possibility of sudden cash shortages leading to costly financing (see Smith and Stulz, 1985 for the benefits of hedging to the firm's value).

In the oil market, producers act as natural sellers of futures contracts (short hedge) to protect themselves against a decline in crude oil prices. In contrast, refiners act as natural buyers of crude oil futures contracts (long hedge) to protect against a price increase which would in effect increase their production costs. Moreover, refiners are also short hedgers of their production, thus they will often sell futures contracts of refined products to protect their margins. Other natural hedgers in the industry (i.e. investors with commercial interest in the physical commodity) are governments, marketers, distributors and in general, everyone engaged in the supply chain of oil up and to the final consumer being either a household or a business. There is expected to be a wide variation in the value of each market participant derivative holdings for hedging purposes. For instance, vertical integration in the oil business can act substitutive to other means of risk management; petroleum firms have limited need to hedge in this case, since they are involved in all stages of production process. However, the mainstream of the business is active in a certain field of expertise e.g. either producing or refining or trading/shipping etc. and the need for sound hedging strategies is vital.

Hedging using financial derivatives is a challenging task because hedging strategies, if not appropriately utilised and fully understood, can be equally problematic to unhedged positions or even worse; by creating a deceptive sense of security. Improper control and supervision of risk management systems, inadequately defined rules and inaccurate valuation of the open positions as well as poorly defined strategies can lead to a debacle. For optimum risk management strategies, financial management needs to create accountability to prevent extreme unforeseen losses and understand the financial consequences of the hedged portfolio - in various market scenarios - through a well defined tested structure, since by eliminating price risk other risks might be introduced such as basis, liquidity and credit risk, among others. The main

hurdles of every hedging plan are hedging costs, absence of suitable products and the perception that shareholders use the firm as a vehicle to obtain oil price risk exposure. Most importantly, basis risk, arising from differences in the derivative contract written and the actual underlying asset could prove disastrous in hedging due to fragile correlation structure. The steeper the basis risk, the larger the disincentive to hedge. In particular, Haushalter (2000) reports evidence that oil and gas producers' fraction of production hedged decreases with basis risk. As an extreme example of basis risk consider marine bunker prices. Bunker fuel oil is a residual oil that is used as a fuel for vessels (IFO180, IFO380 and marine diesel oil) and represents a significant input in the cost function of shipowners and/or ship operators. Alizadeh et al. (2004) find crude oil and petroleum futures inadequate to hedge marine bunker prices in Rotterdam, Singapore and Houston. The reduction in the portfolio variance (hedged vs. unhedged) reaches at best 43%; this is attributed to basis risk since first, the underlying commodity is different and second, the balance of supply and demand is regional.

Table 2.1: Volatility Across Different Assets

Asset	Annualised % Volatility			Return Percentiles (%)			
	Overall	1989-1999	2000-2009	1% Tail	5% Tail	99%Tail	95% Tail
Energies							
Heating oil	39.80	38.15	41.47	-6.66	-3.65	6.29	3.68
Gasoline	41.51	36.32	46.29	-7.33	-4.11	6.46	3.88
WTI	39.99	38.21	41.76	-7.09	-3.75	6.38	3.67
Natural Gas	58.33	55.15	61.29	-9.33	-5.50	10.13	5.68
Electricity PJM	59.86		59.86	-8.64	-4.44	13.31	4.73
Metals:							
Gold	16.08	12.53	19.07	-3.05	-1.56	2.58	1.54
Silver	27.97	24.23	31.38	-5.15	-2.77	4.62	2.58
Copper	27.48	23.70	30.91	-4.93	-2.67	4.45	2.72
Palladium	32.11	27.55	36.25	-6.00	-2.99	5.56	3.13
Platinum	22.38	17.41	26.55	-3.96	-2.11	3.57	2.03
Agriculture							
Kansas Wheat	25.23	21.42	28.66	-4.45	-2.33	4.36	2.54
Minneapolis Wheat	25.81	21.10	29.93	-4.32	-2.27	4.40	2.43
CBOT Wheat	28.40	23.99	32.35	-4.54	-2.63	4.89	2.84
Corn	25.75	21.55	29.48	-4.21	-2.44	4.69	2.54
Oats	33.29	30.54	35.95	-5.47	-3.30	5.96	3.34
Cotton	28.14	24.17	31.74	-4.79	-2.65	4.62	2.64
Soybean Meal	26.46	22.38	30.12	-4.87	-2.47	4.43	2.59
Soybean Oil	23.76	20.73	26.53	-3.92	-2.29	4.14	2.53
Soybeans	23.84	20.09	27.17	-4.39	-2.33	4.06	2.28
Sugar	34.79	34.07	35.54	-6.26	-3.34	5.68	3.46
Cocoa	31.91	29.37	34.33	-5.25	-3.07	5.65	3.28
Coffee	39.46	42.79	35.64	-6.68	-3.73	6.68	3.87
Orange Juice	33.08	35.00	30.96	-5.72	-3.16	5.63	3.07
Lumber	30.62	28.37	32.81	-4.00	-3.05	4.79	3.09
Meats							
Feeder Cattle	13.70	12.69	14.69	-2.47	-1.42	2.18	1.37
Live Cattle	16.02	14.81	17.19	-2.52	-1.56	2.50	1.54
Pork Bellies	37.66	42.16	32.42	-5.07	-3.50	5.38	3.52
Lean Hogs	33.27	32.07	34.48	-4.53	-2.77	5.19	2.55
Financial & Others							
S&P 500	18.60	14.19	22.25	-3.15	-1.79	3.30	1.69
30 Year T- Bond	9.69	8.83	10.51	-1.70	-1.01	1.50	0.94
Com. Res. Bureau	6.68	5.25	7.88	-1.29	-0.61	1.13	0.64
Goldman Sachs CI	22.19	17.23	26.36	-3.86	-2.17	3.55	2.21
Baltic Dry Index	21.40	11.30	28.28	-4.04	-1.78	4.17	1.88

To illustrate the importance of risk management in the energy markets the annual historical price volatility for a number of commodities from 1989 to 2009 is shown in Table 2.1. The financial group has the lowest overall volatility, and the energy group has by far the highest. Generally, energy commodities have distinctly higher volatility than other types of commodities. The properties of the tails of the distribution are also shown. For the petroleum market, for instance, the 1% tail is around 6-7% representing the maximum expected loss on a daily basis. Comparable markets are Palladium, Coffee and Sugar with also 6% tail and volatilities above 32%.

2.7.1 Quantifying Market Risk

After a series of derivatives disasters in the 1990's such as Barings Bank, Metallgesellschaft and Orange County, the need to accurately measure market risk exposure became a demanding task. According to 1988 Basle's Committee Accord, banks are required to preserve a certain level of capital requirements to guarantee that potential losses will not lead to financial distress. One popular method widely adopted in the mid 1990's to facilitate risk management practices, is Value-at-Risk (VaR); the maximum expected loss with a specified probability over a given horizon. Therefore, based on Basle Accord and Capital Adequacy Rules (Basle II), VaR is indispensable for regulatory requirements to discourage irrational risk taking and justify risk with sufficient maintained funds. Risk managers have accepted VaR as a key measure to quantify market risk. In 1994⁵ *J.P. Morgan* publicised its internal set of assumptions and estimation procedure of VaR (RiskMetrics) triggering the attention of both academics and practitioners. The energy complex and their financial derivatives form a key element of the present financial system, and definition of unambiguous risk measurement policies is crucial. However, measuring the market risk of oil and petroleum products is a delicate issue. This is due to a combination of factors such as time-dependence (leading to changes volatility behaviour), non-linear dynamics, heavy tails in oil returns and the complications associated with multiple risk factors.

VaR acts as a practical decision making tool for risk management, indicating the potential downside risk of a portfolio in a single number, easily communicated to all interested parties such as shareholders, management and regulators. As such, VaR plays a manifold role in the modern energy markets. First, being aware of the amount at risk, oil-related businesses can

⁵ For banks, the first regulatory capital requirements was imposed after the *Great Depression* era that followed the stock market crash (October 29, 1929) with the establishment of the Securities and Exchange Commission (SEC) in 1934. For a historical review of the VaR theory and practice the reader is referred to Holton (2003).

employ hedging instruments to mitigate unwanted exposure. Second, lenders (banks) financing energy projects can assess their risk in providing funds. Third, traders can assess any changes in the value of their maintained portfolios under the probability of adverse scenarios. Fourth, when trading oil derivatives it is always helpful to know the potential loss, especially when contracts are cleared and margins should be maintained. For instance, consider the case of Amaranth Advisors LLC natural gas derivatives debacle in September 2006 which resulted in a loss of approx. \$6 billion of the \$9 billion assets under management: Chincarini (2007) finds that Amaranth portfolio strategy was aggressive and - assuming hypothesised positions- estimates that a simple VaR calculation would explain 65% of their losses, the rest being explained by liquidity risk. Finally, there is a growing concern for energy price risk quantification among market participants, given the complexity inherent in the fundamentals and the volatile nature of the industry.




Figure 2.8: Value at Risk approx. using WTI vs. a Diversified Portfolio

As a simple example on how VaR works Figure 2.8 represents the empirical distributions of nearby WTI crude oil futures daily returns for the period 1990-2009 together with a diversified portfolio where only 50% of the wealth is allocated to WTI crude oil futures and the rest 50% are divided into equal amounts (12.5%) to S&P 500 index, gold futures, 30-year US government bond futures and \$ US / GBP currency. We can observe diversity in shapes illustrating interesting risk differences. The bottom panel presents some statistics showing that the diversified portfolio has reduced risk by 47% ($= [36.3 - 19.1] / 36.3$) with a daily VaR figure of 10% at 1% confidence level, as opposed to 19% for WTI. In addition, the expected shortfall i.e. the expected value of the loss under the condition that maximum expected loss has been exceeded provides more information in the tail. In that case the 10-day loss for a \$5 million fund will be 0.23 and 0.43 million for the portfolio and the WTI contract, respectively. Bottom right panel of Figure 2.8 shows a closer look on the left tail of the distribution.

Oil price risk can have diverse effects on different market participants; for instance, when crude oil prices fall this has a negative impact on the producers' cash flows, thus reducing refiners' production costs. However, it is the price transmission mechanisms and crude-products spillover effects that determine whether this outcome will subsequently lead to improved profit margins. From Figure 2.8 we can also observe a small asymmetry in the distribution of WTI returns, implying that the downside risk of producers is higher than that of refiners (19.3% vs. 18.5% for 10-day VaR at 1% level, respectively which is equivalent to \$40K per \$5 million value of crude with an expected shortfall difference of \$90K= $[27.5 - 25.2] \times 5$).

By exploring the particular structure of the tails, VaR is commonly applied by practitioners to disclose market risk and avert higher than sought levels of uncertainty. While the application of VaR is not infallible, it provides a simple and moderately safe method for extrapolating information under difference tolerance levels and horizons, being used to allocate capital, measure diversification effects, compare riskiness of portfolios or projects, estimate the impact of price changes on cash flows, provide a measure of credibility for companies, evaluate the effectiveness of hedges etc. Nonetheless, given the implications of mispricing capital at risk in the oil markets, rather than relying on a single metric, risk measurement includes a *modus operandi* that should be purposefully adopt to new information and changing market conditions. Therefore, an extensive set of actions is necessary, such as educating risk analysts and risk managers, reviewing the models, back-testing, stress testing, scenario analysis as well as clearly defining budgets, position limits, guidelines and policies.

In Chapter 4 we will have a closer look on market price variation and evolution of VaR across time and across different market conditions for the main petroleum commodities. Chapter 4 is completely devoted to examining volatility dynamics with a view of obtaining efficient risk

metrics, as well as consistent VaR and volatility forecasts. Particular focus will be given to the benefits of accommodating within the GARCH framework, changes of fundamentals and changes of the overall volatility behaviour. For this reason, our models will be dependent on both changes in the term structure and changes of unobserved regimes in order to realistically represent some of the stylised facts of the examined markets. The robustness of such forecasting strategies will be compared to benchmarks, using both statistical tests and risk management loss functions. After the two oil price crises in the 1970's energy market participants have always been faced with high levels of uncertainty and, as mentioned above, it is vital not only to develop sound models for risk quantification but review and back-test their performance.

2.7.2 Minimum Variance Futures Hedging

Keynes (1930) was the first to assume that futures act as an insurance scheme for hedgers, who pay premiums to speculators to carry their risk. For example, if a refiner holds a barrel of crude oil and the price falls (rises), he realises a certain capital loss (gain). Thus, a risk-averse refiner would want to unwind such price risk by simultaneously taking an equivalent reverse position in the futures market - to be settled on cash upon delivery - in order to offset any capital loss, or in other words, lock the price today to avoid unfavourable surprises. The proportion of futures contracts that should be held to effectively reduce the risk of each unit of the assumed inventory is called hedge ratio. Although conventional wisdom suggests that this should be one-to-one (naïve hedge), this strategy fails to deliver because due to imperfect futures-spot correlation and the term structure of volatility, a residual capital gain or loss is expected; price movements are neither parallel nor synchronised. This has triggered the interest at an academic level by the works of Johnson (1960) and Stein (1961) who introduced the concept of portfolio theory in the hedging problem. The foundations of modern risk analysis date back to Markowitz (1952). According to portfolio theory, investors construct the optimal portfolio by combining risky assets in such a way that offers the highest reward for the minimum amount of risk. Ederington (1979) applied this concept in determining a minimum-variance hedge ratio and proposed a measure of hedging effectiveness. Physical inventories are viewed as fixed and the decision on how much to hedge is determined by the optimum hedge ratio which minimises the variance of the portfolio of futures and spot positions; this is essentially the ratio of the unconditional covariance between spot and futures price changes over the unconditional variance of futures price changes. Consider the case of an oil producer who wants to secure his income in the petroleum futures market. The return on the producer's hedged portfolio, rp_t , is:

$$rp_t = \Delta S_t - \gamma \Delta F_t \quad (2.1)$$

where Δ is the difference operator, S_t and F_t , spot and futures prices, respectively, such as that the change in the spot (futures) position between $t-1$ and t is $\Delta S_t = S_t - S_{t-1}$ ($\Delta F_t = F_t - F_{t-1}$); γ is the hedge ratio. Let $\text{Var}(\Delta S_t)$, $\text{Var}(\Delta F_t)$ and $\text{Cov}(\Delta S_t, \Delta F_t)$ be, respectively, the unconditional variance of the spot and futures returns and their unconditional covariance. The producer must choose the value of γ that minimises the variance of his portfolio returns which is found as the solution to:

$$\text{Var}(rp_t) = \text{Var}(\Delta S_t) - 2\gamma \text{Cov}(\Delta S_t, \Delta F_t) + \gamma^2 \text{Var}(\Delta F_t) \quad (2.2)$$

Taking the partial derivative of Equation (2.2) with respect to γ and setting it equal to zero yields the variance minimising hedge ratio (MVHR), γ^* :

$$\gamma^* = \frac{\text{Cov}(\Delta S_t, \Delta F_t)}{\text{Var}(\Delta F_t)} \quad (2.3)$$

The value of γ^* is equivalent to the slope coefficient of a simple regression of spot against futures returns. For example a value of $\gamma^* = 0.9$ would indicate that the hedger should sell 0.9 barrels in futures for every unit held in the underlying. The variance reduction of the hedged vs. the unhedged portfolio is equal to the coefficient of determination R^2 [$= 1 - \text{VaR}(rp_t) / \text{VaR}(\Delta S_t)$].

Figure 2.9 displays the variance of the WTI and heating oil producer's hedged position. The optimum hedge ratio is in both cases less than one (at 1% significance level). Compared to a naïve hedger, a minimum variance hedge strategy would have achieved an additional reduction in risk of 100 basis points. However, note that these hedge ratios are estimated historically for the period January 1995 to December 2009 using daily data. Investors are mainly concerned for forecasts of the hedge ratios to cover their current and future positions.

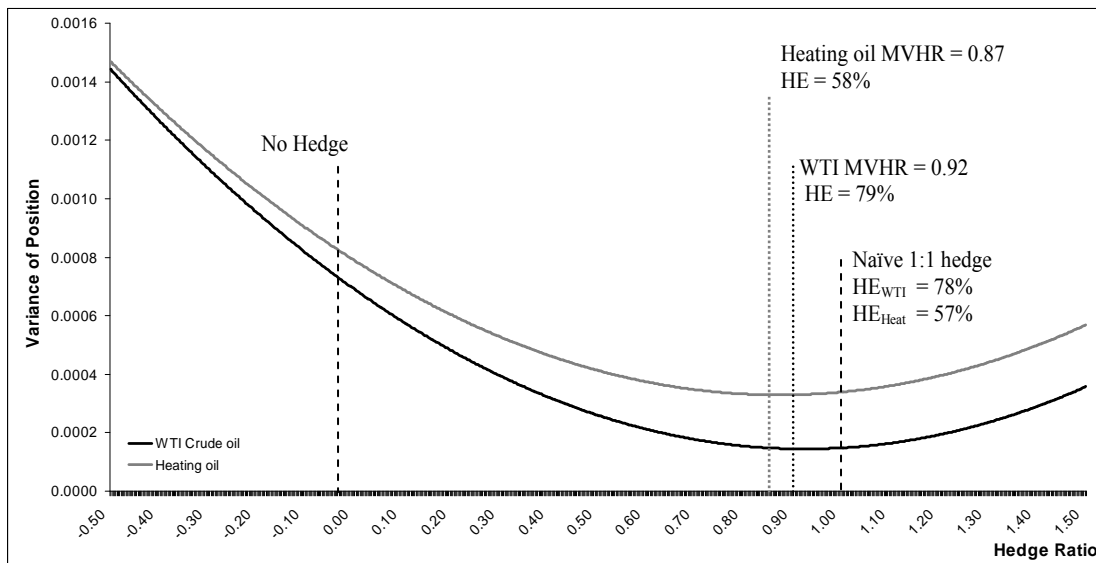


Figure 2.9: Dependence of Variance on the Hedge Ratio

Myers and Thompson (1989) generalised the estimation of optimal hedge ratios to account for conditioning information that is available up to the time of the hedging decision, demonstrating that conventional MVHR is too restrictive and it might be non minimum-variance efficient; the processes generating the covariance matrix of futures and cash prices are usually non-constant throughout time and the expected return to holding a futures contract might be nonzero. Consequently, it is more appropriate to establish a market model of equilibrium including lags of futures and spot prices, plus any other known key price drivers (such as stocks and storage costs) to make an informed decision. They suggest that as new information arrives in the market hedge ratio changes to reflect new market conditions and hedging models that manage to accommodate time-variation in risk time are expected to generate superior performance.

We will discuss in more detail the construction of optimum hedge ratios in Chapter 5. We will demonstrate how multivariate analysis may be employed for describing the joint return distribution of futures and spot prices. Following the suggestion of Myers and Thompson (1989), to devise our hedging strategy we will base our approach on an equilibrium model that permits both asymmetric and non-linear adjustment in the futures-spot relationship (see for instance Ng and Pirrong, 1996) and asymmetric persistence of volatility shocks while also allowing for sudden changes in the unconditional variance covariance matrix. This way, Chapter 5 presents a model that encompasses all the features discussed so far, used to generate time-

varying hedge ratios that produce economically significant results, having relevant implications for the locus of hedgers and portfolio managers.

2.7.3 Metallgesellschaft Hedging Debacle

The purpose of this section is to demonstrate what can go wrong in the risk management process through the Metallgesellschaft AG (MG) experience. MG was formerly one of Germany's largest industrial conglomerates being active in quite a few areas, from mining and engineering to commodity trading and financial services. In the early 1990's, Metallgesellschaft's U.S. oil subsidiary MGRM (Metallgesellschaft Refining and Marketing) reported huge derivatives losses, later estimated to \$1.4 billion. The architecture of the hedges that MGRM devised to protect against oil price adverse movements created a large controversy over their capability to facilitate risk reduction and lock merchandising profits. MGRM's case study deals with the execution of a failed oil price hedging strategy where a firm, ignoring the stochastic behaviour of the term structure of petroleum prices as well as cash flow requirements to support their hedging plan managed to escalate the risk matrix function of the corporation.

In 1992, MGRM set in motion an innovative marketing plan offering long-term fixed price guarantees on deliveries for gasoline, heating oil, and diesel fuel, for up to 10 years. These contracts were mainly of two types. The first type included contracts wherein delivery schedules were pre-specified (firm-fixed) with an attached option-like feature to terminate at customer's will and receive as a payment half the difference of the prevailing WTI spot price and the guaranteed fixed price, multiplied by the remaining contracted quantity. In the second type contracts, the price and total quantity was fixed (firm-flexible), allowing customers to decide on the timing and volume of deliveries. Of course, the contracted total quantity should have been exhausted by expiration; these contracts also gave customers the option to terminate and receive the full difference between the 2-month futures price and the contract price. By September 1993, MGRM had committed to deliver over 150 million barrels of petroleum products at fixed prices the bulk of which was negotiated during the summer of 1993, when the oil prices fluctuated around \$20/bbl.

The large underlying volumes (around 20% of the futures market total open interest) indicated that MGRM was facing a substantial amount of price risk. Due to the structural design of the underlying marketing program, hedging was not a plain-vanilla case and MGRM's implemented a stack and roll strategy⁶: long-term exposure was offset by buying short term

⁶ Although, the management had two other alternatives, they were associated with significant costs (see Edwards and Canter, 1995). First, MGRM could commit to physical storage by purchasing and storing

futures and OTC swaps (stack and roll) on a barrel-to-barrel basis. The main shortcoming of the stack-and-roll strategy was the systematic process of selling the maturing contracts and simultaneously buying new short-dated contracts to maintain the hedge, implying a rollover cost if energy prices were in contango. At first, based on historical data it appeared that since backwardation is the norm for the oil markets (Litzenberger and Rabinowitz, 1995) this strategy was indeed the most cost-effective even offering a potential for making profits from the trades. Nevertheless, market conditions changed and the expectations to receive rollover gains from inverted markets collapsed leading MGRM to the brink of bankruptcy. In fact, given the high traded volumes rollover costs became so acute that repetitively margin calls soon led to a liquidity crisis; while changes in the value of the actual commodity do not generate matching inflows or outflows until the realisation of delivery, futures positions do because of marking to market. In view of the large losses and funding requirements of MGRM, the supervisory board dictated liquidation of the hedges and started a process of negotiating the withdrawal of MGRM's long-term contracts. In the meantime, NYMEX revoked MGRM's "hedging exemption" demanding higher margins which further accelerated the liquidation process.

The Metallgesellschaft debacle including the original long-term strategy and the manner in which it was hedged as well as the decision to rapidly liquidate the hedges triggered long discussions in the academic community as well as among practitioners. Supporters of the strategy (such as Culp and Miller, 1994; 1995a; 1995b) claim that the firm's plan was economically sound and MGRM would not have suffered such losses if the hedge position had not been hastily terminated. In fact, they argue that forward delivery contracts increased in value by the same amount as the short term contracts decreased in value when energy prices fall. It was a plain liquidity crisis that should have been dealt with by the MGRM's bankers so as to realise the long-run profit potential of the strategy. Regarding the concept of the amount hedged (1:1) as opposed to a MVHR alternative, Culp and Miller argue first, the data are subject to considerable error that will inevitably produce imperfect hedges and second, MVHR does not maximise firm value; informed speculation is a regular part of risk management strategies. Finally, Culp and Miller (1995b) support the hedging program, pointing out that basis risk was

the necessary quantity of oil to meet supply obligations. This would imply prohibitive costs (including financing costs for the immediate purchase of oil, storage costs and a certain investment to facilitate such stock) and a non-negligible residual risk (uncertainty regarding future carrying costs such as interest rates and some remaining market risk arising from the possibility that customers would exercise their option). Second, the firm could engage in long-term forwards matching exactly the expiration dates of the supply obligations. Since there was not much of a market for oil up to 10 years (futures were traded for up to 3 years and illiquidity in the distant end of the futures curve was a serious constraint) OTC dealers would have requested a premium for accommodating the rollover cost and credit risk, creating also significant costs.

not a real threat due to high correlation of nearby and spot which would definitely reduce the risk of an outright position.

Conversely, opponents of the MGRM strategy (inter alios, Mello and Parsons, 1995) argue that the firm's hedging strategy had an inherent speculative component: provided that backwardation persists when the roll-over takes place, the nearest to maturity contract price will exceed deferred contracts' prices thus a profit will be realised (selling the nearby at $F_{t,T}$ with T days to expiry and buying a deferred contract with maturity $T+n$ at $F_{t,T+n}$); it essentially was designed to exploit the term structure. According to Edwards and Canter (1995), MGRM's hedging strategy rollover risk was 15% of its price risk, a risk that apparently the MGRM managers were willing to bear, justified on the grounds that crude oil is more often in backwardation than contango⁷. Moreover, Pirrong (1997) estimates that the MVHR for delivery obligations with maturities of 15 months was typically around 0.5 (similar to Edwards and Canter, 1995 and Mello and Parsons, 1995): less than 0.5 for the September 1992 - June 1993 period and between 0.5 and 0.6 for the period June - December, 1993; this implies that MGRM did not possess superior information and a barrel-for-barrel strategy actually increased oil price risk by overhedging, creating thus excessive basis risk due to the Samuelson effect (1965). Pirrong (1997) further shows that an implementation of a minimum variance hedging program would have saved roughly \$1 billion, generating 57% less losses than the ones realised. Furthermore, Edwards and Canter (1995) argue against Culp and Miller by emphasising that an economically sound hedging strategy should allow the hedger to unwind its positions at any time without sustaining extensive losses rather than locking him until the end of the original scheme.

In 1993, with OPEC overproduction, surging North Sea output and weak demand, oil prices plummeted and energy markets went into contango. The shifting of fundamentals was the main affair that caused the debacle. In retrospect, it seems that MGRM was indeed hedged against price risk but backwardation prevalence was a vital assumption to prevent other risks from exacerbating. Coming across significant *term structure risk* (or *calendar basis risk*) in combination with overhedging proved that was sufficient to create a domino effect in the existing risk matrix. *Rollover cash-flow risk* created excessive price risk of cash-settling the stack of short term futures contracts whereas the corresponding cash inflow from declining prices could not support the mark-to-market outflows since it was to be realised gradually over

⁷ For instance, an investor who wanted to roll the September '93 contract 1-week prior to expiry i.e. September 14, 1993 selling the nearby at \$16.96/bbl and buying the 3-month contract (December '93) at 17.67 with $F_{t,T} < F_{t,T+1}$ realises a loss of \$0.71/bbl, accumulating to more than \$0.1 billion for the equivalent position of MG (=150 million barrels x \$ 0.71/bbl).

the next 10 years. *Funding liquidity risk* for MGRM was gradually aggravated by margin calls which resulted in accumulated losses. Although MGRM was benefiting from a network of large financial institutions as shareholders, poor communication between the management and the parent MG added a liquidity burden on the parent who was not prepared to bear. There were a series of risks inbuilt in MGRM's hedging strategy apart from the aforementioned. For instance MG was also facing *credit risk* in that the counterparties might default on such long-dated physical obligations. Furthermore, operational risk was an important obstacle since the firm failed to accurately recognise, communicate and quantify the perils of the strategy designed.

Establishing a framework to analyse and rank such risks is fundamental not only to quantify and manage risk exposure but to be aware of the downside potential and decide on the risk appetite of the organisation so as to take further actions by reporting risk, establishing position limits and thresholds to potential losses. Nowadays there are several risk management tools to set some prerequisites and avoid recurrence of hedging cul-de-sacs. It is important to stress test assumptions for a wide range of market scenarios such as persisting backwardation and contango (to prevent extreme rollover losses), collapsing correlation of short vs. long term prices (which can lead to under- or over- hedging due to basis risk), extreme volatility or in general, worst case scenarios. VaR could also play a key role to assess the risk from physical vs. derivatives positions or even Cash flow at Risk (CFaR) which focuses on the operating cash flow during a period.

2.8 Term Structure of Futures Prices

The shape of the futures curve is of great interest to energy market participants since the pattern of futures prices at different maturity dates reflects market expectations integrating anticipated trends about current price and inventory levels, supply and demand schedules, OPEC behaviour, speculative activity, political involvement and possibly many other factors. An upward sloping futures curve is consistent with an expected spot price increase that compensates inventory owners for the cost of carrying inventories, i.e. warehousing costs and the interest foregone on the capital invested in storing the commodity (cost-of-carry). In finance theory, this condition is known as contango. Since futures prices are bound to converge to spot at delivery, in contango derivative prices will decrease until expiry, *ceteris paribus*. However, the cost-of carry relation (derived using standard no-arbitrage arguments) cannot effectively explain a downward sloping futures curve since lower expected future spot prices imply negative storage costs. Futures are said to be in backwardation or inverted when futures prices fall with maturity at a given point in time.

Several theories have been advanced to reconcile the issue of inverted markets with the two most widespread interpretations of the phenomenon put forth by the theory of normal backwardation (Keynes, 1930) and the theory of storage, (Kaldor, 1939; Working, 1949; Telser, 1958 and Brennan, 1958). It is useful to note beforehand that these are not mutually exclusive but rather complement each other. According to Keynes, backwardation is the result of the risk transferring function of futures markets and hedging pressures: agents involved in the physical commodity markets use futures to hedge their positions and unless hedging demand of the two sides of the market (buyers and sellers) is matched, risk cannot be transferred at zero cost. Hedgers will in effect have to induce speculators to bear their risk; that is, pay them a premium as compensation for this service/insurance. The second strand of literature identifies as the main determinant of storable commodity prices the inventories and introduces a fudge factor in the cost-of carry relationship, the *convenience yield*: consumption assets, such as oil, bestow benefits⁸ to inventory holders. If marginal convenience yield (net of storage costs) is high, the spot (prompt) price will exceed the futures (distant futures) price causing backwardation. Large convenience yields are a feature observed during low inventory periods where supply is rigid - so spot prices are high due to tight market conditions. Inversely, in periods of supply abundance spot prices fall and physical market participants are better off not having to pay the cost of storage. This shrinks convenience yields and the market switches to contango.

Figure 2.10 shows snapshots of the historical evolution of the futures curve for NYMEX heating oil against the spot price, spanning from January 2000 to December 2009. We can see strong diversity in shapes throughout time. First, different maturity futures can switch from backwardation to contango (and vice versa). Second, movements along the curve can be non-synchronised e.g. inverted markets in the short term and contangoed in the long end (and vice versa) or even display autonomy. Moreover, seasonality has a noticeable effect in the formation of heating oil curve indicating that agents allow for such cyclical behaviour in pricing. In the US, demand for heating oil always peaks in the winter, giving rise to a price premium for futures expiring then. The curve can shift upwards or downwards, rotate and/or swing from convex to concave, representing the differences in the dynamic behaviour of the

⁸ Holding stocks can absorb demand shocks, mitigate the risk of supply disruptions, smooth the refining process and circumvent frequent revisions in the production process. In addition, there is an embedded timing option attached regarding the timing of making the stocks available i.e. sell at high prices. For an extractive resource commodity such as crude oil, convenience yield is also associated with the benefits of holding reserves and leaving oil in the ground. Litzemberger and Rabinowitz (1995) estimated that 80–90% of the time the oil forward curve is in backwardation and this is attributed to the fact that backwardation is a necessary to keep production running. If extraction costs grow by no more than the interest rate while discounted futures prices are higher than the spot then all producers have the incentive to leave the oil in the ground and postpone production.

nearby and deferred contracts. To demonstrate this, consider the impact of a refining facility closing for a month and that of a sudden large oilfield discovery; the former will certainly not affect the long term prices as much as the latter and vice versa. Therefore, various short-, medium- and long-term risks interact in a complex way to determine the shape of the curve.




Figure 2.10: Evolution of Heating Oil Term Structure vs. Spot (New York)

A stylised fact of the futures curve is that variances and correlations between the nearest and subsequent futures decline with maturity (Samuelson, 1965); as futures contracts approach expiry, they are more sensitive to information due to offsetting positions to prevent delivery, inevitable convergence of future prices to spot and the stronger linkage of the short-term part of the curve to current demand and supply conditions. This effect is depicted in Figure 2.11. Grey columns are the estimated annualised volatilities of contracts with maturities from 1 up and to 18 months. The negative volatility-maturity relationship is also observable from the drawn lines that show the evolution of individual futures curves on arbitrary chosen consecutive days. Two sets of lines are presented. At the bottom there are typical inverted futures curves, each representing a random day during October 1999, whereas at the top there are typical contangoed futures curves for March 2009. In the displayed downward sloping curves, the nearby futures experience much larger fluctuations. For instance, 1 month to maturity contract lies in the range of \$21 - \$24 per barrel whereas 18 month contract fluctuates around the smaller range of \$19 - \$20/bbl (having a 3:1 ratio of absolute movements: \$3/bbl vs. \$1/bbl window). This also holds for the upward sloping curves. The nearest to maturity contract is in the wider range of \$38 - \$48/bbl whereas 18 month lies in the narrower window of \$53 - \$56/bbl (again, a 3:1 ratio: \$10/bbl vs. \$3/bbl).




Figure 2.11: Term Structure of Prices & Volatilities




Figure 2.12: Premium/Discount of 1- over 6- & 18- Month WTI Futures

In petroleum economics, storage and transportation costs are major factors in the pricing function of futures, forming complex term structures relative to the financial markets (see Alizadeh and Nomikos, 2004a on the impact of transportation costs). Consistent with the

theory of storage, the levels of oil volatility are directly linked to the futures curve. The role of the slope or the (interest and storage) adjusted calendar spread as a proxy for temporal deviations of demand and supply has been successfully tested and findings indicate that volatility increases with falling inventories i.e. under backwardation (see Ng and Pirrong, 1996 for refined petroleum products; Geman and Ohana, 2009 for oil and gas). In practice, the need to model the whole term structure simultaneously to optimise portfolios, measure and hedge risk is important. Figure 2.12 illustrates the dynamics of the WTI curve for three different deliveries in the form of premium or discount of 1- compared to 6- and 18-month contracts. Up to the summer of 2008, the market has historically been in backwardation, on average; nearby futures were above long term prices. After 2008, the curve altered, with a strong and relatively persistent contango.

Literature has shown (see for instance Cortazar and Schwartz, 1994; Clewlow and Strickland, 2000 and Tolmasky and Hindanov, 2002) that price curve movements can be distinguished to a parallel shift (level factor), a tilt (slope factor) and a twist (curvature factor) of the curve. Sensitivities regarding these underlying factors indicate a flat level factor capturing most of the variance, a slope factor with opposite signs at both ends of the term structure and a curvature factor having equal signs at both ends of the maturity spectrum, but an opposite sign in the middle. Effective use of futures contracts requires a thorough understanding of the risk factors determining futures prices and of the price sensitivities regarding these underlying factors. For energy market participants these implied risk factors that explain futures curve movements are essential from different perspectives. The term structure provides an information-rich framework that can be used as an input to value derivatives, to identify the risk-return profile of maintained physical portfolios (businesses maintain receivable accounts and liabilities across a wide range of maturities) and investment funds (investors trade in a wide range of maturities), to manage inventory making decisions on whether to “carry” or liquidate the maintained stock, to evaluate future demand growth and supply trends based on market expectations and to set up sound hedging strategies (for instance systematic backwardation/contango might lead to automatic rollover losses if the nearest to maturity contract is used for hedging; see for instance the Metallgesellschaft case study in section 2.7.3).

It is the aim of Chapter 6 to provide a rigorous multivariate statistical analysis of the main petroleum futures prices as well as economically meaningful petroleum spreads. In doing so, we will particularly look into the information content of the dependence structure between correlated petroleum futures curves which has not received much attention in the literature. We will particularly study the long-run equilibrium components for inter-commodity spreads’ risk factors and explore the information content derived from sophisticated regime switching

models. The goal will be to reveal possible disequilibria between the stochastic processes of same-nature risk factors (e.g. level of WTI vs. level of Brent), disaggregate their dynamics across different market conditions and examine how these can be utilised to exploit futures curve movements. We will present a unified approach, flexible enough to accommodate more than two cross-factors and tractable enough to forecast the whole term structure and derive risk measures such as volatility and value-at-risk. We will tackle exactly this issue in Chapter 6 and postpone the discussion to the corresponding sections.

2.9 Conclusion

In this chapter we described the structure of the oil markets. We presented the stylised facts surrounding the markets from the basic fundamental forces (demand-supply) to the power of cartels throughout time and the role of speculators. By reviewing historical fluctuations in oil prices since the inception of the industry in 1859, we offered a detailed outlook with chronological justification on how the main price drivers work and how the market has evolved today. Additionally we demonstrated important elements of energy risk and introduced the basic background to risk management issues such Value-at-Risk, optimum hedging and the term structure which will be the main topics of this thesis, analysed further in Chapters 4, 5 and 6, respectively. We also attempted to highlight the need of sound risk analysis and risk management strategies in the modern energy markets with the real recent hedging debacle of Metallgesellschaft. This case study, overall, summarises how adverse oil price shocks can lead to huge losses if speculation is inherent in a risk mitigating program.

The next chapter, *Regime Switching Models and Applications in Finance*, deals with an important class of econometric models, namely the Markov Regime Switching (MRS) models. Flexible to accommodate sudden changes in stochastic processes such as the dynamic adjustment of prices to equilibrium relationships, volatility and correlation, MRS models will be employed for empirical validation of our objectives in the field of energy risk. Our empirical applications in subsequent chapters will essentially confirm that this regime-switching specification turns out to exhibit many of the salient features of oil markets and will essentially present how to exploit the information content of such models. Chapter 3 will first, provide a short introduction regarding the use of MRS models; this will be followed by a survey regarding their development in the field of finance and energy economics; next, a detailed explanation of theoretical background and the estimation procedure and will be supplied by means of simple real life examples.

Chapter 3

Regime Switching Models and Applications in Finance

3.1 Introduction

Regardless of how sophisticated, models that attempt to describe the conditional distributions of oil prices, interdependencies within the market and the possibility of extreme phenomena, face a challenging task in capturing plausible scenarios, forecasting the path of the market and assessing the implications of such moves. Less than fifty years ago, Mandelbrot (1963) observed that the tails of the distributions of price changes are “*extraordinarily long*” with variances that “*vary in an erratic fashion*”. Arguing that there need not be any discontinuity between outliers and the rest of the distribution, Mandelbrot challenged the assumption of the Gaussian hypothesis, and put forward the stable-Lévy or stable Paretian family of distributions to model price changes, which includes the normal as a special case. As a by-product of this study, the famous volatility clustering phenomenon was first put in print. In this seminal work, the author also mentions the possibility of utilising, alternatively, a mixture of normal distributions, to tackle the issue. Subsequently, in a classic study of stock prices’ behaviour, Fama (1965) showed that the empirical distributions of daily price returns are usually highly peaked and heavy tailed; in fact, departures from normality are as predicted by Mandelbrot.

Non-normality, asymmetries and time-varying dependence are well-documented concerns in all commercial markets. As it also often happens, different segments of the data may favour different models since markets evolve and the underlying dynamics are generated by diverse mechanisms. Energy commodity markets show intricate behaviour and it frequently occurs that no model is likely to be precisely correct. Based on empirical evidence, researchers seek to discover a posteriori causal relationships, rationalise how the majority of the data works and reflect in the best possible way, the attributes of ever-changing petroleum economics and market conditions.

The focus of this thesis will be on explaining the prevalence and magnitude of different regimes in oil markets and whether these regimes trigger a change in the way risk should be perceived. Traditional models of energy markets rely on single state relationships being thus, relatively rigid to explain real world dynamics. Yet, more often than not, several incidents might produce such market shocks that their impact is capable of drastically altering the behaviour of the series' either permanently (structural break) or transitory (regime shift). The latter effects can be of different durations while depending on the nature of the episode these can occur on a regular (e.g. seasonality) or irregular (e.g. backwardation/contango) basis and can be highly persistent or very short-lived (jumps and spikes). Moreover, these events can repeatedly affect the market in a non-standard manner, creating risk factors of different shapes and forms which change through time in complicated ways and cause the market to switch back and forth among different processes. In such cases, it is natural to resort to nonlinear estimation methods for the temporal evolution in volatility dynamics and co-dependence across assets; in the presence of regimes, the application of linear models seems inadequate and non-informative, providing thus little insight into market patterns or the variability of prices, volatilities and correlations throughout time. Given the complexities met in the empirical validation of the energy markets, the thesis will attempt to offer a further perspective on energy risk by providing new framework for risk analysis aimed at assessing the significance of regime inference on practical considerations. The novelty of this approach, as applied to the petroleum industry, is to determine the information content derived from models of switching regime and assess their role and effectiveness in uncovering fundamental interactions, quantifying risk under different market conditions and, finally, evaluating the extent to which regimes convey relevant information on risk management objectives.

This chapter draws on the literature and estimation issues of MRS models. Particular interest will be given in petroleum markets; oil prices are assumed to be drawn from a Markovian system of alternating distributions. The subsequent section provides a brief introduction of MRS models. Next, a selective overview of the MRS literature in finance and energy economics is supplied. Some supplementary evidence regarding nonlinearities and evidence of structural changes in the oil markets is also presented. Next, the basic concepts behind MRS model calibration are laid down. This is followed by a detailed description of the process. Final section concludes. The usefulness of this methodology is demonstrated in real life applications in later chapters where various extensions are considered.

3.1.1 A Primer on Models of Changing Regime

Regime switching models comprise an important division of financial time series models and a functional approach to model nonlinearities. Ultimately, two broad nomenclatures of switching models exist, differing in the means of regime identification. The first category presupposes directly observable states and formulates regimes as a deterministic function of a known variable. This framework dates back to the switching regression of Quandt (1958) and the threshold autoregressive models of Tong (1978, 1983), soon after developed by Tsay (1989); further extended to smooth transition models (see also Chan and Tong, 1986 and Teräsvirta, 1994) which essentially allowed for the possibility of gradual rather than definite movement among regimes. Note that since only Markov Regime Switching (MRS) models are relevant for the context of this chapter and for the applications in the empirical part of the thesis, the latter framework will not hereafter be considered. MRS models constitute the second kind of switching nonlinear models. Although they can be traced back to Goldfeld and Quandt (1973), they were developed and popularised towards the end of the 1980's in finance and econometrics, by Hamilton (1989, 1990). Unlike the first category which inexorably requires auxiliary information or prior beliefs about why or how regime-switching is manifested within the data, the key feature of the Markovian structure is that regimes are assumed to be unknown; they are fully determined by a latent stochastic process and the transition probability matrix driving the motion among and within states.

Of the main advantages of MRS models is their ability to capture cyclical behaviour and unknown breaks. First, model parameters are functions of a hidden Markov chain and regime classification is based on optimal probabilistic inference. This way, instead of adhering to a strict pre-specified form on the junctures and the persistence of shifts, empirical data reveal their own structure without constraints. Second, discrete time MRS models are characterised by a number of distinct regimes within which different model parameters apply. Third, since the probabilities of each state change over time, model parameters become time-dependent. In essence, these probabilities play a weighting role in the switching scheme because it is not strictly required for a process to be in neither of the defined states e.g. it can alternatively be in-between e.g. in one state with probability 90% and in another with 10%; this way asymmetric behaviour across parameters and across time is also addressed. Forth, being an approach essentially based on mixture of distributions it can produce densities with nonstandard shapes and accommodate thus, some of the stylised facts of financial time series such as fat-tails and nonlinearities. Finally, due to their modularity, MRS models introduce certain flexibility, usually translated to improved 'fit' of real world data; specifying multi-state conditional means,

variances, correlations dynamics is therefore expected not only to be information-rich but also resourceful in every manner, such as enhancing forecast ability.

3.2 Literature Review

After the path-breaking paper of Hamilton (1989), the use MRS models for describing nonlinear behaviour of asset returns and nonlinear dependence among assets, has become widespread. By separating price trajectories into economically meaningful regimes, they allow for a great deal of flexibility in the parameterisation of conditional distributions. Related interest grows at a very fast pace, with a variety of contributions for several branches of the literature, as a sign of the highly multidisciplinary nature of regime-switching models. Therefore, they are extensively applied in various fields of finance and petroleum economics.

3.2.1 General Review

Hamilton (1989) developed a two-regime autoregression to model the post-war dating of business cycles. To derive a criterion for defining economic recessions, measuring their persistence and dealing with asymmetries in the business cycles, the author studied the real US GNP (Gross National Product) growth rate regime shifts between periods of recessions and expansions. He found distinct dynamics between the two states while the predicted dates of the turning points accurately matched the official dates set by the National Bureau for Economic Research. Hamilton's framework generated a remarkable amount of subsequent research. There is a wide range of papers analysing real business cycles, and turning points in a regime switching context such as Diebold and Rudebusch (1996) who combined the concepts of dynamic co-movement of economic variables throughout cycles and regime switching, Kim and Nelson (1998), Clements and Krolzig (2004) and so on. Others link financial stock market cycles to the evolution of economic conditions (Hamilton and Lin, 1996; Maheu and McCurdy, 2000). Issues such as financial crises (Coe, 2002 for the Great Depression) and contagion effects during crises (Guo et al., 2011 for the Great Recession) have also been the focus of various studies.

Regime switching models have also provided important insights in asset allocation (Ang and Bekaert, 2002 & 2004; Guidollin and Timmermann, 2008). Perez-Quiros and Timmermann (2000) show relatively high cyclical asymmetries in the risk-return profile of small firms compared to large ones, being particularly sensitive to recessions and monetary policy shocks. Using MRS based portfolios to capture regimes in the distribution of small and large stocks' returns they report substantial predictability during recessions. In addition,

Guidollin and Timmermann (2007) investigate optimal asset allocation under crash, slow growth, bull and recovery regimes for stock and bond returns and find that portfolio weights strongly depend on the state of the economy, verifying thus the economic implications of different regimes, in- and out- of sample. Their model captures both short- and long-term fluctuations in the joint stock-bond distribution with investors allocating more of their wealth to stocks as investment horizon increases only in the crash state; the more persistent bullish markets are associated with a decline in stock allocation as investment horizon increases. Moreover, they find that the utility cost of ignoring regimes is 3% p.a. at the short 1 month horizon while for longer annual investment horizons this falls to 1.3% p.a. Another substantial body of literature concentrates mainly on quantifying and forecasting portfolio risk such as Billio and Pelizzon (2000, 2003) who estimate regime switching Value-at-risk (VaR). Overall, they find that the regime switching model is superior at long horizons compared to simple Gaussian and multivariate GARCH specifications, suggesting a more conservative view of risk.

Furthermore, in view of Gray (1996), who successfully modelled the short term interest rate as a regime switching process, many papers studied the impact of regime shifts on the entire yield curve using dynamic term structure models. One of the earlier studies includes the one-factor, continuous time formulation of Naik and Lee (1997) who assumed constant market prices of regime-specific risk. Several others have extended the framework including Bansal and Zhou (2002) who considered regime-dependent market price of risk (see also Dai and Singleton, 2003). Apart from these models, the regime switching behaviour of interest rate term structure has also been studied in a cointegration scheme. Tillmann (2004) presents a contextual link of the yield curve in a switching error correction equilibrium model. The author discovers that the short-run adjustment of US interest rates and the term premium of long-term rates do experience regime changes while the underlying states are mainly triggered by Federal Reserve policies. Of course, the same econometric framework has been also applied to other markets. For instance, Clarida et al. (2003) find that a three-state MRS cointegration model for spot and forward exchange rates can outperform the random walk, especially for longer-term horizons.

MRS models have also penetrated the derivatives markets literature to explore and assess price discovery, market interrelationships, hedging and pricing. Sarno and Valente (2000) examine the existence of regime shifts in the relationship between spot and futures returns in the FTSE 100 and S&P 500 stock index futures markets and establish strong evidence of nonlinear mean reversion to the cost of carry spread. In the same markets, Alizadeh and Nomikos (2004b) argued that, in the view of these shifts, by allowing the minimum-variance hedge ratio to be dependent upon the state of the market, one may obtain more efficient forecasts; overall, their

results confirmed that by exploiting regime changes, market participants may be able to effectively improve hedging performance. Apart from the futures markets, switching models have also been employed in option valuation, after Naik (1993) addressed the issue of pricing and hedging contingent claims on assets that exhibit discontinuous volatility shifts. The author developed an analytical solution for European call options in terms of the integral of the Black and Scholes (1973) formula. A more general discrete model was suggested by Bollen et al. (2000) by using a lattice-based algorithm and simulation to price both European and American options. Duan et al. (2002) considers GARCH option pricing under regime switching proposing a lattice type approximation whereas Buffington and Elliott (2002) draws on risk-neutral valuation methods.

There is a plethora of studies prompted by the attractive features of MRS models and different extensions and perspectives have emphasised different views of their benefits and economic significance. Other empirical evidence includes Fong and See (2001) for commodity indices, Elliott and Hinz (2002) for portfolio and chart analysis, Mayfield (2004) for estimating the market risk premium, Mount (2006) for capturing the volatile behaviour of electricity markets and predicting spikes, Pelletier (2006) for regime switching dependence structures, Elliott et al. (2006) for MRS GARCH option pricing using Esscher transforms, Chung et al. (2007) for monetary and fiscal policy. Related literature continues to expand, in several modern topics like measuring hedge fund risk exposure (Billio, 2010), real options (Driffil et al., 2009) and in comparatively young markets such as the CO₂ emission allowances (Benz and Trück, 2009).

3.2.2 Evidence from the Oil Markets

Petroleum markets have always been at the core of economic research agenda mainly due to the far-reaching implications of oil price uncertainty on the economic and financial system. Numerous studies, such as Morana (2001), Giot and Laurent (2003), Sadorsky (2006) and Hung et al. (2008), have all well documented that energy markets do exhibit the properties that triggered the concern of Mandelbrot in 1963 i.e. non-normality, fat tailed distribution and volatility clustering. Energies are prone to sudden changes, not only in response to shifts in the fundamentals; such as periodic supply contractions or demand surges (for instance, due to the emergence of new large consumers like China and India, nowadays). Several events/episodes disrupt stability within the industry including, *inter alia*, exogenous geopolitical events, weather catastrophes, strikes, access to reserves and OPEC policies.

For instance, under different conditions such as backwardation/contango, the risk-return profile of oil is known to change dramatically: periods of low inventories are associated with higher volatility and reduced correlation in the term structure (Fama and French, 1987; Ng and Pirrong, 1994). Failure to account for these effects will inevitably lead to under- or over-estimation of volatility and correlation which in turn will have adverse repercussions on market participants' wealth in the form of non-robust hedge ratios, large contingent claims pricing errors etc. Pirrong (1997), in order to accommodate these features, calculated backwardation adjusted GARCH hedge ratios by incorporating the cost of carry relationship in the variance-covariance evolution. The author's findings indicated that minimum variance hedge ratios for crude oil were far less than the naïve 1:1 ratio in the period 1992 to 1994. Nevertheless, apart from the second moments, different market conditions are also expected to affect the price response function. For example, in equilibrium models, the sensitivity of price changes to the deviations from the long-run mean is not expected to be uniform under positive and negative errors. Ng and Pirrong (1996) were the first to report that the process of futures-spot price convergence of refined petroleum products is non-linear and asymmetric. In particular, the speed of adjustment is faster for large deviations from equilibrium and when the market is in contango¹. Fattouh (2009) examined the regime dynamics of the basis in the crude oil market to measure the effects of stocks and OPEC behaviour. The results showed that the basis-stocks relationship is nonlinear with higher basis elasticity when inventories are low. Contrary to what the theory of storage would predict, the author found that as the stocks increase the probability of staying in the contango regime decreases and this can be attributed to the role of OPEC in deciding output cuts in view of accumulation of excess stockpiles. The evidence presented above confirms that the behaviour of oil markets might not be described well when we assume a single underlying stochastic process. Backwardation and contango conditions are one simple example that can illustrate this phenomenon. For instance Fattouh (2010) using a threshold regime-switching model found evidence of non-linearity in the adjustment process of different quality crudes.

¹ Other relevant studies are Huang, Yang and Hwang (2009) and Ye et al. (2006) among others. Huang, Yang and Hwang (2009) investigate the dynamic interaction between the futures and spot prices for crude oil within three observed regimes classified according to the magnitude of the basis and find significant interaction when the basis is less or above a certain threshold value. Furthermore, Ye et al. (2006), in a simple regression framework include relative inventory levels to forecast WTI in the short run. This study supports the predictive power of inventories on prices especially after accounting for asymmetric price responses in high and low inventory periods. Allowing for nonlinearities (by including the squared low and high relative inventory levels) improved the fit of the data, especially during periods of large price swings (e.g. in July 2000 and in December 2002) and the forecast ability on an out-of-sample basis.

In reality, many are the events which might cause a changing market structure in the world oil market. Dvir and Rogoff (2010), taking a long-term view, test for changes in persistence and volatility of real crude oil prices for the period 1861-2009. They identify “*three epochs of oil*”: the persistent and volatile (28% vol. p.a.) period of 1861-1877, the less persistent and much less volatile (14% vol. p.a.) phase of 1878-1972, and finally, the 1973 onwards era where oil price behaviour revisited last century’s dynamics, however with relatively less pronounced volatility (23% vol. p.a.). The authors also observe a breakpoint in 1934 that coincides with the major discovery of the East Texas oilfield; up and to 1972 volatility was even lower than 1878-1933. 1972 is linked to the peak production of this oilfield which essentially ended US control over excess reserves giving OPEC the power to coordinate. The first and third transition points in 1878 and 1972 had two main differences in that first, in 1972 the oil industry was much bigger and second, economies were much more reliant on the use of oil. Both happened during years of expanding demand and overall economic growth as opposed to the 1934 breakpoint which actually occurred in a period of economic recession.

Wilson et al. (1996) looked at sudden changes in volatility². Employing an iterated cumulative sum-of-squares (ICSS) algorithm, the authors attempt to detect structural unconditional volatility changes in the NYMEX oil futures contract as well as a portfolio of oil and gas companies’ stock prices and the S&P index. Regarding the oil futures, 15 significant volatility changes were detected, between March 1983 and December 1992, 5 of which exceeded 100% in absolute terms. The most remarkable increase in daily volatility (239%) occurred in the period from November 1985, through December 1986, from 0.85% to 2.9%. Another significant upward change was observed during the eight day period following the invasion of Kuwait in the mid-January 1991 (213%) which was followed by the largest decrease (83%) in the study’s sample. Moreover, including this information in an ARCH framework the author finds a significant decrease in the persistence parameters.

More formally, Fong and See (2002, 2003) studied the temporal volatility dynamics in a combined GARCH and MRS setting. They found significant and distinct switches in the WTI futures market, mainly prompted by events that had profound influence in fundamentals. Within the high volatility state, an increase in backwardation is likely to increase regime persistence due to low inventories. In the same state, changes in the basis are more likely to affect futures volatility, consistent with the theory of storage. Moreover, GARCH persistence is significantly reduced showing that regime shifts dominate GARCH effects. In particular, the high variance

²Rather than the variance, other studies such as Maslyuk and Smyth (2008) tested the existence of structural breaks in a unit root context and found significant and meaningful breaks that may affect the intercept, the trend or both elements of the test.

state is associated with a six fold increase in the unconditional volatility whereas, within this state, volatility has no memory. In the low variance state, the GARCH persistence parameter is 0.50 compared to 0.99 of the single regime GARCH model. Out-of-sample forecasting comparisons favour regime switching models and the inferior performance of the simple GARCH is attributed to the presence of structural breaks; for the reason that these breaks make estimators reflecting essentially the persistence of volatility regimes rather than true volatility. Another study by Vo (2009) married the concept of regime switching with that of stochastic volatility to forecast the dynamics of WTI crude oil. The author finds that the simple MRS model captures better the in-sample dynamics in terms of mean absolute errors whereas out-of-sample, stochastic volatility with regime shifts is favoured.

Various studies have focused on revealing regime-dependent interrelationships and asymmetric effects. Noel (2007) and Lewis and Noel (2010) for instance examine the pricing and price response functions of retail gasoline prices. Other studies are focused on interrelationships between oil and the macroeconomy such as Raymond and Rich (1997). More recently, Cologni and Manera (2009) also employ an MRS analysis for the G-7 countries and show that net oil price increases and oil price volatility contribute to the output growth whereas their role in explaining recessions has not been steady over time. Additional evidence for the US business cycles is also provided by Clements and Krolzig (2002) who explore the role of oil in generating asymmetries (see also Holmes and Wang, 2003 for the UK). Finally, Aloui and Jammazi (2009, 2010) investigate the relationship of crude oil shocks and stock markets behaviour and Choi and Hammoudeh (2010) study the relative regimes and regime durations of the oil and equity markets as well as for industrial commodities.

3.3 Fundamental Concepts of Markov Processes

Before introducing the formulation of MRS models, it is helpful to briefly review the necessary groundwork behind hidden Markov models. The following subsections deal with two essential concepts to facilitate the designing and implementation of regime shifts in the modelling process, mixed distributions and Markov chains.

3.3.1 Mixture of Distributions

The main setback when approximating the distribution of a non-normal variables with the Gaussian is that we underestimate the probability of extreme price changes (fat-tails), we under- or over- estimate the probability that these may be positive or negative (positive and negative skewness, respectively), and we overestimate the probability of returns around the

mean (excess kurtosis). These features, as well as population heterogeneity can be efficiently addressed by hypothesising that price shocks are drawn from a mixture of several normal distributions that can have different moments and this conjecture can be helpful to determine the mean and variance changing process.

Fama (1965) argued that different components may be the result of recurrent features of financial series such as day-of-the-week effects. In addition, Kon (1984) suggested that the dissimilar distributions for a particular series arise due to diverse information signals; for instance, a mixture of three normals for the return distribution of stock prices might arise because of firm-specific information, market-specific information and noninformation. In the oil market similar conclusions can be drawn. For example, a mixture distribution might arise due to existence of two distinct market conditions, namely backwardation and contango, due to the dimorphic nature of supply for crude oil (low cost OPEC, high cost non-OPEC supply), due to uneven concentration of demand and supply around the world, global vs. regional effects, and so on, which may influence the underlying fundamentals and further alter price sensitivities, equilibria, risk-return profiles and dynamic linkages.

In the two component case, a random variable X_t , whose increments follow a mixture of two normal distributions, can be defined by ΔX_{1t} with probability π_1 and ΔX_{2t} with probability $\pi_2 = (1 - \pi_1)$, where ΔX_{1t} and ΔX_{2t} are independent normal random variables while $0 < \pi_1, \pi_2 < 1$ are the mixture coefficients (weights). In the general K component case, the sequence of a random variable ΔX_t with time index $t \in [1 \dots T]$ can be described by K probability density functions (pdf), denoted as $f_k(\Delta X_t | s_t = k)$ for $k = 1, \dots, K$. The mixing process can be defined as the discrete random variable s_t which determines the particular distribution each observation is drawn from, with an assigned probability of occurrence π_k :

$$s_t = \begin{cases} 1 \text{ with probability } \pi_1 \\ \dots \\ K \text{ with probability } \pi_K \end{cases} \quad (3.1)$$

where, the conditions $\sum_{k=1}^K \pi_k = 1$ and $\pi_k \geq 0$, for $k = 1, \dots, K$ apply. Let Ω_{t-1} be the information available up to time $t-1$. In terms of the conditional distribution of ΔX_t , we can write:

$$\Delta X_t | \Omega_{t-1} = \begin{cases} N(\mu_1, \sigma_1^2); s_t = 1 \\ \dots \\ N(\mu_K, \sigma_K^2); s_t = K \end{cases} \quad (3.2)$$

Note that the probabilities π_k 's are considered as weights to calculate the pdf of process ΔX_t as a linear combination of the state-dependent densities:

$$f(\Delta X_t | \Omega_{t-1}) = \sum_{k=1}^K \pi_k f_k(\Delta X_t | s_t = k, \Omega_{t-1}) \quad (3.3)$$

Under conditional normality, the pdf is:

$$f_k(\Delta X_t | s_t = k, \Omega_{t-1}) = \frac{1}{2\pi\sigma_k} \exp\left\{-\frac{1}{2} \frac{(\Delta X_t - \mu_k)^2}{\sigma_k^2}\right\} \quad (3.4)$$

The parameters of a mixture distribution model can be estimated using Maximum Likelihood (ML). In particular, given Eq. (3.3) and (3.4), the likelihood of the K components mixture of normals model, is equivalent to:

$$L(\boldsymbol{\theta}) = \sum_{t=1}^T f(\Delta X_t | \Omega_{t-1}; \boldsymbol{\theta}) = \sum_{t=1}^T \sum_{k=1}^K \pi_k \frac{1}{2\pi\sigma_k} \exp\left\{-\frac{1}{2} \frac{(\Delta X_t - \mu_k)^2}{\sigma_k^2}\right\} \quad (3.5)$$

where the vector of parameters to be estimated is $\boldsymbol{\theta} = (\mu_1, \dots, \mu_K; \sigma_1, \dots, \sigma_K; \pi_1, \dots, \pi_{K-1})$, with $\pi_K = 1 - \sum_{k=1}^{K-1} \pi_k$. The first and second moments of the mixture μ and σ^2 are a function of the respective components and can be expressed as:

$$\mu = E[\Delta X_t] = \sum_{k=1}^K \pi_k E[\Delta X_t | s_t = k] = \sum_{k=1}^K \pi_k \mu_k$$

$$\sigma^2 = Var(\Delta X_t) = E[\Delta X_t^2] - (E[\Delta X_t])^2 = \sum_{k=1}^K \pi_k E[\Delta X_t^2 | s_t = k] - \left(\sum_{k=1}^K \pi_k \mu_k\right)^2 =$$

$$= \sum_{k=1}^K \pi_k (\mu_k^2 + \sigma_k^2) - \left(\sum_{k=1}^K \pi_k \mu_k\right)^2 \tag{3.6}$$

Table 3.1: Fitted Finite Mixtures

	<u>WTI crude oil</u>				<u>Heating oil</u>				<u>Brent crude oil</u>				<u>Gas oil</u>			
	1-M	3-M	6-M	9-M	1-M	3-M	6-M	9-M	1-M	3-M	6-M	9-M	1-M	3-M	6-M	9-M
Panel A: Normal Distribution																
μ	0.038	0.039	0.040	0.041	0.038	0.037	0.037	0.038	0.040	0.041	0.042	0.042	0.038	0.037	0.037	0.039
σ	2.267	2.000	1.786	1.664	2.242	1.985	1.772	1.647	2.198	1.974	1.785	1.670	2.068	1.845	1.647	1.537
Panel B: Mixture of Two Normal Distributions																
μ_1	0.087	0.096	0.088	0.082	0.069	0.057	0.058	0.058	0.102	0.092	0.088	0.074	0.075	0.066	0.077	0.082
μ_2	-0.272	-0.240	-0.191	-0.156	-0.070	-0.006	-0.013	-0.012	-0.260	-0.136	-0.128	-0.093	-0.083	-0.069	-0.079	-0.084
σ_1	1.814	1.532	1.336	1.244	1.790	1.470	1.284	1.201	1.699	1.424	1.284	1.230	1.607	1.421	1.214	1.125
σ_2	4.092	3.476	3.145	2.925	3.358	2.829	2.578	2.436	3.756	3.221	2.961	2.854	3.130	2.909	2.521	2.323
π_1	0.864	0.832	0.828	0.826	0.775	0.695	0.702	0.718	0.828	0.777	0.785	0.808	0.766	0.786	0.747	0.736
π_2	0.136	0.168	0.172	0.174	0.225	0.305	0.298	0.282	0.172	0.223	0.215	0.192	0.234	0.214	0.253	0.264
Panel C: Mixture of Three Normal Distributions																
μ_1	0.205	0.192	0.150	0.127	-0.176	-0.070	-0.008	0.032	0.179	0.140	0.122	0.132	0.179	0.195	0.156	0.186
μ_2	0.036	0.034	0.047	0.045	0.135	0.123	0.105	0.077	0.032	0.018	0.042	0.024	-0.035	-0.048	-0.066	-0.096
μ_3	-0.519	-0.412	-0.372	-0.318	-0.457	-0.259	-0.297	-0.297	-0.405	-0.314	-0.328	-0.218	-0.054	0.114	0.224	0.262
σ_1	0.983	0.838	0.743	0.788	1.028	1.034	1.001	0.958	1.062	1.015	0.925	0.868	1.179	1.014	0.953	0.894
σ_2	2.165	1.904	1.672	1.580	2.155	1.947	1.839	1.817	2.128	2.012	1.775	1.637	2.240	1.872	1.727	1.584
σ_3	4.964	4.351	3.891	3.600	4.045	3.454	3.320	3.319	4.660	4.350	3.750	3.572	4.902	3.757	3.339	2.994
π_1	0.200	0.218	0.222	0.265	0.162	0.247	0.344	0.432	0.250	0.349	0.348	0.331	0.342	0.310	0.385	0.381
π_2	0.741	0.717	0.707	0.662	0.759	0.654	0.582	0.516	0.684	0.591	0.576	0.596	0.632	0.629	0.552	0.544
π_3	0.059	0.065	0.071	0.073	0.079	0.099	0.075	0.052	0.066	0.060	0.075	0.073	0.026	0.061	0.063	0.075

- Estimation period uses daily observations from June 1994 to December 2009.
- The table presents the parameters of fitting a normal and mixture of normals distribution to 1- month, 3- month, 6- month and 9- month to maturity petroleum futures of NYMEX WTI crude and heating oil and ICE Brent crude and gas oil.
- Let the petroleum returns be $\Delta \ln F(t,T) = \Delta X_t$, then the fitted distributions are $\Delta \ln F(t,T) \sim N_1\{E[\Delta X_t] = \mu, \text{Var}(\Delta X_t) = \sigma^2\}$, $\sim N_2\{E[\Delta X_t] = \mu_1 \pi_1 + \mu_2 \pi_2, \text{Var}(\Delta X_t) = \pi_1[\mu_1^2 + \sigma_1^2] + \pi_2[\mu_2^2 + \sigma_2^2] - \mu^2\}$ and $\sim N_3\{E[\Delta X_t] = \mu_1 \pi_1 + \mu_2 \pi_2 + \mu_3 \pi_3, \text{Var}(\Delta X_t) = \pi_1[\mu_1^2 + \sigma_1^2] + \pi_2[\mu_2^2 + \sigma_2^2] + \pi_3[\mu_3^2 + \sigma_3^2] - \mu^2\}$.
- Estimates are based on Maximum Likelihood (ML) estimation, by maximising Eq. (3.5) subject to the constraints that $\sum_{k=1}^K \pi_k = 1$ and $\pi_k \geq 0$, for $k = 1, \dots, K$.

Table 3.1 gives a comparison of the means and volatilities for different petroleum commodities futures price changes. The corresponding second moments of the two states in Panel B, show a twofold increase, on average, whereas the probability of the low variance state is much higher. Moreover, in Panel C, volatilities involve a threefold to fivefold increase between the two extreme cases, whereas the medium- volatility state is associated with higher

probability. What merits attention is that, as basic economic theory predicts (Samuelson *effect*, 1965; see Chapter 2, section 2.7), the futures prices become increasingly volatile as they approach maturity; an observation that holds even within regimes. More complex is the behaviour of probabilities of different contracts as they display some autonomy. Notice as well in Panel C that the probability of the low variance state seems to increase with maturity.




Figure 3.1: Fitted Mixture of Two Normals for WTI Crude Oil

Additionally, Figure 3.1 illustrates how the mixture distribution accommodates fatter tails and excess kurtosis. It displays the estimated mixed distribution (shaded areas) of 1- month WTI crude oil vs. the fitted simple normal distribution (dashed line). Looking at the resulting shape of the distribution it can be seen that this is asymmetric towards the left. This is accommodated by the fact that $\mu_1 \neq \mu_2$. If the difference of component means is not statistically significant i.e. $\mu_1 = \mu_2 = \dots = \mu_K$, the shape of the distribution is bound to be symmetric.

An interesting feature of mixture distribution models is that, ex-post, we can conduct inference about which state was more likely at each step t by obtaining a conditional probability that, the process was drawn from state k . Using the law of total probability and Bayes' rule³:

³Denote c the complementary event, say of A i.e. the event of "not A ". The *law of total probability* relates marginal probabilities with conditional probabilities. We can express the law of total probability as:

$$\pi_{kt|T} = \frac{\pi_k f_k(\Delta X_t | s_t = k, \Omega_{t-1})}{\sum_{k=1}^K \pi_k f_k(\Delta X_t | s_t = k, \Omega_{t-1})} = \frac{\pi_k f_k(\Delta X_t | s_t = k, \Omega_{t-1})}{f(\Delta X_t | \Omega_{t-1})} \quad (3.7)$$

Mixture distribution models can be considered as a special case of Markov models, with constant regime probabilities and a transition probability matrix of rank 1. For the case where the components' weights are obtained by hidden Markov models, more details will be provided in the next section.

3.3.2 Markov Chains

Defining any particular Markov chain requires a set of K possible states (state space) and a transition matrix which assigns probabilities to the state transitions. We will be interested in discrete time Markov chains that are described by a countable state space. Let s_t be a latent random variable with $s_t \in [1, \dots, K]$, $K \geq 2$ and finite and $t \in [1, \dots, T]$. The process will start in one of these states and then move successively among them. The first observation of such an order is called *initial state*. A *Markov chain* is a particular stochastic process with the distinctive property of restricted memory, in the sense that the current state contains as much information for the future as the whole history of the process⁴. This facet, the so called *Markov property*, can be expressed in mathematical form as:

$$\Pr(s_{t+1} = k | s_t, s_{t-1}, \dots, s_{t_0}) = \Pr(s_{t+1} = k | s_t) \quad (3.8)$$

The equation above illustrates that the future state s_{t+1} depends on the current state s_t alone and not on earlier consecutive realisations s_{t-n} , with $n \geq 1$. As for any stochastic process, probabilities must be assigned to the cascade of possible values. The conditional probabilities

$\Pr(B) = \Pr(A \cap B) + \Pr(A^c \cap B) = \Pr(B | A) \Pr(A) + \Pr(B | A^c) \Pr(A^c)$. *Bayes' theorem*, on the other hand, states that: $\Pr(A | B) \Pr(B) = \Pr(B | A) \Pr(A)$. Hence, combining these two rules:

$$\Pr(A | B) = \frac{\Pr(B | A) \Pr(A)}{\Pr(B | A) \Pr(A) + \Pr(B | A^c) \Pr(A^c)}$$

⁴ Note that, when this holds, the process follows a first-order Markov chain. It is possible to construct higher order chains by assuming that future states depend on the current and certain number of past states. It is also possible to construct zero-order Markov chains i.e. the processes are independent of both the current state and the whole history (Bernoulli processes). However these are beyond the scope of this chapter and therefore, are not hereafter discussed.

$\Pr(s_{t+1}=k|s_t)$ that the process will be in some state k , one step ahead at time $t+1$, are given by the transition kernel $\Psi = (p_{ij})$:

$$\Psi = \begin{pmatrix} p_{11} & p_{12} & \dots & p_{1K} \\ p_{21} & p_{22} & \dots & p_{2K} \\ \dots & \dots & \dots & \dots \\ p_{K1} & p_{K2} & \dots & p_{KK} \end{pmatrix}; \quad p_{ij} = \Pr(s_{t+1} = j | s_t = i) \geq 0 \quad (3.9)$$

where i, j are indices with $i, j \in k$. The main diagonal $p_{11}, p_{22}, \dots, p_{KK}$ gives the probability that state s_t will remain the same in the following period; off diagonal elements p_{ij} , give the transition intensities i.e. the probability that state i will be followed by state j , with $i \neq j$. Of course, $\sum_{j=1}^K p_{ij} = 1$ holds for every i because the events collected at each row of matrix Ψ constitute an exclusive and exhaustive partition of the space. That is, given the state at time t , the sum of the probabilities of transition to other states plus the probability of no change in the state must sum to 1.

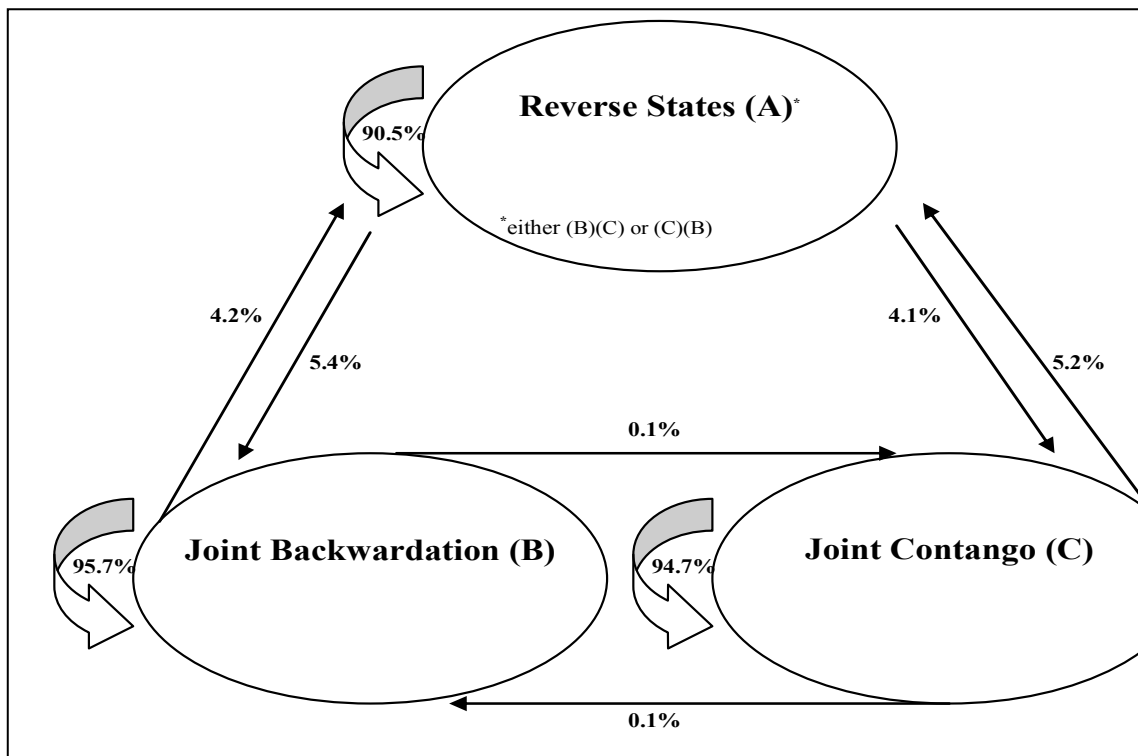


Figure 3.2: Illustration of the Structure of a Three-State Markov Process

To illustrate the application of the abovementioned definitions consider the states of backwardation and contango in the oil market. Figure 3.2 displays the state transitions of WTI crude oil and heating oil among three states i.e. both commodities being in backwardation (B), both commodities being in contango (C) and one commodity being in (B) while the other in (C). Calculations are based on actual NYMEX futures daily data from 1994 to 2009. Classification in regimes was carried out using the overall short-term to medium term slope of the term structure including the first 10 contract maturities. A negative (positive) slope is indicative of backwardation (contango). According to the figure, a joint crude-heating downward sloping futures curve has a 95.7% probability of repeating the following day. An interesting observation is that the lower probability of no transition p_{ii} (that is, equal to 90%) is assigned to the reverse state conditions (A) which means that this state is of a more transient nature. In addition, the probability of transition from (B) to (C) or vice versa is too low (0.1%) implying that the market will most certainly pass through (A) to transit from (B) to (C) and vice versa from (C) to (B).

Apart from Ψ , it is important to define the n-step transition kernel; that is, the conditional probabilities $\Pr(s_{t+n}|s_t)$ that the process will be in some state, n- steps ahead at time $t+n$. The Markov chain should obey the following relationship, known as *Chapman-Kolmogorov equation*:

$$\Pr(s_{t+n} | s_t) = \Pr(s_{t+n} | s_{t+n-1})\Pr(s_{t+n-1} | s_{t+n-2})\dots\Pr(s_{t+2} | s_{t+1})\Pr(s_{t+1} | s_t) \quad (3.10)$$

Therefore, the full set of forward transition probabilities can be expressed in matrix form as the n- power of Ψ , that is $\Psi^n = (p_{ij}^{(n)})$. In our example the 1-day and 2-weeks ahead transition probabilities are:

$$\Psi = \begin{pmatrix} 95.7 & 4.2 & 0.1 \\ 5.4 & 90.5 & 4.1 \\ 0.1 & 5.2 & 94.7 \end{pmatrix}; \quad \Psi^{(10)} = \begin{pmatrix} 70.9 & 23.8 & 5.3 \\ 30.5 & 47.4 & 22.1 \\ 8.3 & 28.0 & 63.7 \end{pmatrix}$$

Thus the probability that both crude oil and heating oil are inverted (in backwardation), given that they were inverted 2-weeks ago is $\Pr(s_{t+10} = B | s_t = B) = \Psi_{(1,1)}^{10} = 70.9\%$. Eq. (3.10) implies also that due to the Markov property, the distribution of a Markov chain is fully specified by its initial distribution $\Pr(s_{t_0})$ and the transition probability matrix as:

$$\begin{aligned} \Pr(s_t, s_{t-1}, \dots, s_{t_0}) &= \Pr(s_{t_0}) \Pr(s_{t_0+1} | s_{t_0}) \dots \Pr(s_t | s_{t-1}, \dots, s_{t_0}) = \\ &= \Pr(s_{t_0}) \prod_{\tau=t_0}^{t-1} \Pr(s_{\tau+1} | s_{\tau}) \end{aligned} \quad (3.11)$$

Moreover, the distribution of s_t for large t converges to a limit π_k irrespective of the initial point $\Pr(s_{t_0})$. This is called the unconditional probability of being in state k which essentially determines the long-term behaviour of each state. Let $\boldsymbol{\pi}$ represent the vector containing these unconditional probabilities i.e. $\boldsymbol{\pi} = (\pi_i \ \pi_j)$ with $\pi_i = \Pr(st=i)$. It can be shown that these are the solution to the system $\boldsymbol{\pi} = \boldsymbol{\pi} \boldsymbol{\Psi}$, subject to the constraint that $\sum_{k=1}^K \pi_k = 1$. In the two regime case this is:

$$\begin{aligned} \boldsymbol{\pi} = \boldsymbol{\pi} \boldsymbol{\Psi} &= (\pi_1 \ \pi_2) \begin{pmatrix} p_{11} & 1-p_{11} \\ 1-p_{22} & p_{22} \end{pmatrix} \Rightarrow \\ (\pi_1 \ \pi_2) &= (\pi_1 p_{11} + \pi_2 (1-p_{22}) \quad \pi_1 (1-p_{11}) + \pi_2 p_{22}) \end{aligned} \quad (3.12)$$

$$\text{Hence,} \quad \begin{cases} \pi_1 = \pi_1 p_{11} + \pi_2 (1-p_{22}) \\ \pi_2 = \pi_1 (1-p_{11}) + \pi_2 p_{22} \\ \pi_1 + \pi_2 = 1 \end{cases} \Rightarrow \begin{cases} \pi_1 (1-p_{11}) = (1-\pi_1)(1-p_{22}) \\ \dots \\ \dots \end{cases}$$

Finally, solving yields:

$$\pi_1 = \frac{1-p_{22}}{2-p_{11}-p_{22}}; \quad \text{or} \quad \Pr(st=i) = \frac{1-p_{jj}}{2-p_{ii}-p_{jj}} \quad (3.13)$$

3.4 The Baseline Markov Regime Switching Model

A Markov model in its basic first-order, K -state form is a multi-stochastic process, based on an underlying sequence of observations. It is a special class of dependent mixtures and consists of two processes: a latent K -state Markov chain that drives the regimes, and a state-dependent process of observations. Assume the price changes for oil futures, say ΔX_t follow the dynamics:

$$\Delta X_t = \mu_t + \varepsilon_t; \quad \varepsilon_t \sim iid \ N(0, \sigma^2) \quad (3.14)$$

where ε_t is a Gaussian white noise process and μ_t can represent, for instance, a simple drift $\mu_t = \mu$, a mean reversion process $\mu_t = \alpha (X_t - \mu)$, an autoregressive (AR) process $\mu_t = \alpha (\Delta X_{t-1} - \mu)$ or any regression for that matter. In general, μ_t is some predictable process. Assume now that we wish to assign different values to different subsamples. Say, for example, we want to capture the oil upward trending 1970-1981 years and the downward trending period of 1981 - 1987. We could include a dummy variable to represent this change and specify the conditional mean as $\mu_t = \mu + \alpha I_{t>t^*}$ with I the indicator function taking values of 1 for $[t^*, \dots, T]$ and zero otherwise. However, rather than claiming one abrupt structural change in the model we can specify a more general form that encompasses both specifications and many more. Mathematically:

$$\Delta X_t = \mu_{s_t,t} + \varepsilon_{s_t,t}; \quad \varepsilon_{s_t,t} \sim N(\mu_{s_t}, \sigma_{s_t}^2) \quad (3.15)$$

where s_t is a first order Markov chain defined by the probability law of Eq. (3.8) and having a transition probability matrix Ψ as in Eq. (3.9). If we consider a simple two regime process, the dynamics of the process ΔX_t and the transition kernel Ψ would be, respectively:

$$\Delta X_t = \begin{cases} \mu_{1,t} + \varepsilon_{1,t}; & \varepsilon_{1,t} \sim N(\mu_1, \sigma_1^2) \\ \mu_{2,t} + \varepsilon_{2,t}; & \varepsilon_{2,t} \sim N(\mu_2, \sigma_2^2) \end{cases} \quad (3.16)$$

$$\Psi = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} = \begin{pmatrix} p_{11} & 1-p_{11} \\ 1-p_{22} & p_{22} \end{pmatrix} \quad (3.17)$$

In this case ε_t has a conditional mixture of normals distribution with first and second moments given by a modification - regarding the mixing weights - of Eq. (3.6). If the shock that occurred at the breakpoint t^* is permanent, then we would expect that $p_{22}=1$, implying that the second state is an *absorbing state*; once the market enters that condition is not expected to revert back. However, Markov models have the flexibility to allow for $p_{KK} < 1$ since there are also transitory shocks that usually occur repeatedly. For instance, if the regime shift reflects a tightening of supply coupled with demand surge, it would be a sensible postulation to account for the possibility that the market will eventually absorb the imbalance at some point in time i.e.

$p_{22} < 1$. Every Markov chain that it is possible to move from every regime to every regime is called *ergodic* or *irreducible*.

A strength provided by this MRS model is that different regime-specific means for the mixture allow for time varying skewness and appropriate treatment of dynamic asymmetries; two statistically different means cause the mixture to be bimodal (Gray, 1996). Moreover, the regime-switching model supposes that s_t is unobserved and the time series is decomposed into two generating processes with different variances. These different variances are weighted by the conditional regime probabilities (defined in the next section) which are a function of t . Thus, even if state dependent variances are constant, the aggregate process will be time-dependent. To determine the timing of the states, the econometrician has to make inferences relying on the Markov probabilities. Furthermore, to determine the duration of the states i.e. how persistent each state is, we can use the transition probability matrix Ψ . The average expected duration of being in state 1 is calculated using the formula suggested by Hamilton (1989):

$$\sum_{k=1}^{\infty} k p_{11}^{k-1} (1 - p_{11}) = (1 - p_{11})^{-1} \quad (3.18)$$

Because of the latent nature of s_t , calibration of MRS models is rather demanding. Hamilton (1990) introduced the Expectation Maximisation (EM) algorithm, later refined by Kim (1994), Hamilton (1994) and Gray (1996). In general, estimation is based on Bayesian updating of the likelihood function using a recursive filter, based on the recursive nature of conditional regime probabilities.

3.4.1 Regime Inference and Maximum Likelihood

Assume θ contains all the MRS parameter estimates. Based upon θ estimated from data spanning through the time index $t \in [1 \dots T]$, three estimates about the unobserved state variable s_t , can be made. The first is the estimated probability that the unobserved state variable at time t equals k given the information set up to $t-1$, Ω_{t-1} ; this is the expected, predicted or *ex-ante probability* $\pi_{k|t-1} = \Pr(s_t=k|\Omega_{t-1})$. The ex ante probability is of particular interest in forecasting simulations. The second is the estimated probability that the unobserved state variable at time t equals k given all the information set up to t , Ω_t , with $t < T$; this is termed the *filter probability* $\pi_{k|t} = \Pr(s_t=k|\Omega_t)$. The third is the estimated probability that the unobserved state at time t equals k given the entire time index $t \in [1 \dots T]$ of the sample; this is termed the *smooth probability* $\pi_{k|T} = \Pr(s_t=k|\Omega_T)$. The smooth probability has been traditionally used to identify and establish the

timing of regime shifts. Although the econometrician observes directly ΔX_t he/she can only make inferences about the value of s_t based on ΔX_t .

To estimate any MRS specification, we can use the conditional distributions of each state and the assigned probabilities to integrate out the state variable. Rewriting ΔX_t in terms of the conditional distribution of $f(\Delta X_t | \Omega_{t-1})$ this can be done as:

$$\begin{aligned} f(\Delta X_t | \Omega_{t-1}) &= \sum_1^K f(\Delta X_t, s_t = k | \Omega_{t-1}) = \\ &= \sum_1^K f(\Delta X_t | s_t = k, \Omega_{t-1}) \Pr(s_t = k | \Omega_{t-1}) = \\ &= \sum_1^K f(\Delta X_t | s_t = k, \Omega_{t-1}) \pi_{kt|t-1} \end{aligned} \quad (3.19)$$

By conditioning on the regime at $t-1$, the ex-ante probability that the process is in state 1, is given by the transition probabilities and the filter probabilities at $t-1$, as:

$$\begin{aligned} \pi_{1t|t-1} &= \Pr(s_t = 1 | \Omega_{t-1}) = \sum_{k=1}^K \Pr(s_t = 1 | s_{t-1} = k) \Pr(s_{t-1} = k | \Omega_{t-1}) = \\ &= \sum_{k=1}^K \Pr(s_t = 1 | s_{t-1} = k) \pi_{kt-1|t-1} \end{aligned} \quad (3.20)$$

with $\Pr(s_{t-1}=k | \Omega_{t-1}) = \pi_{kt-1|t-1}$ being the filter probability at $t-1$. In the two regime case, it follows that:

$$\begin{aligned} \pi_{1t|t-1} &= \sum_{k=1}^2 \Pr(s_t = 1 | s_{t-1} = k) \Pr(s_{t-1} = k | \Omega_{t-1}) = \\ &= p_{11} \pi_{1t-1|t-1} + p_{21} \pi_{2t-1|t-1} = p_{11} \pi_{1t-1|t-1} + (1 - p_{22})(1 - \pi_{1t-1|t-1}) \end{aligned} \quad (3.21)$$

Following Hamilton (1994) and Gray (1996), by Bayes' rule the filter probability can be written as a function of the previous' step ex-ante probability $\Pr(s_{t-1} = k | \Omega_{t-2})$ i.e. $\pi_{1t-1|t-2}$:

$$\Pr(s_{t-1} = 1 | \Omega_{t-1}) = \frac{f(\Delta X_{t-1} | s_{t-1} = 1, \Omega_{t-2}) \Pr(s_{t-1} = 1 | \Omega_{t-2})}{\sum_{k=1}^K f(\Delta X_{t-1} | s_{t-1} = k, \Omega_{t-2}) \Pr(s_{t-1} = k | \Omega_{t-2})} \quad (3.22)$$

and conditional on the normality assumption:

$$f(\Delta X_{t-1} | s_{t-1} = k, \Omega_{t-2}) = \frac{1}{\sqrt{2\pi\sigma_{k,t-1}}} \exp\left\{-\frac{1}{2} \frac{(\Delta X_{t-1} - \mu_{k,t-1})^2}{\sigma_{k,t-1}^2}\right\} \quad (3.23)$$

Consequently, the ex-ante probability can be written as a simple recursive filter. In more general form:

$$\boldsymbol{\pi}_{t|t} = \boldsymbol{\Psi}^T \left[\frac{\boldsymbol{\pi}_{t|t-1} \odot \mathbf{f}_{t-1}}{\mathbf{1}^T (\boldsymbol{\pi}_{t|t-1} \odot \mathbf{f}_{t-1})} \right] \quad \boldsymbol{\pi}_{t|t-1} = \boldsymbol{\Psi}^T \boldsymbol{\pi}_{t-1|t-1} \quad (3.24)$$

where $\boldsymbol{\pi}_{t|t}$ and $\boldsymbol{\pi}_{t|t-1}$ denote vectors containing the probabilities of being in each regime at time t conditional on the observations up to time t and up to time $t-1$, respectively, \mathbf{f}_{t-1} is a vector of state dependent densities conditional on $t-1$ and \odot is the element-by-element multiplication.

Subject to the constraints that $\sum_{k=1}^K \pi_{kt|t-1} = 1$ and $0 \leq \pi_{kt|t-1} \leq 1$ and iterating the expressions in Eq. (3.24), the log-likelihood function $L(\boldsymbol{\theta})$ to be maximised using numerical optimisation methods is:

$$L(\boldsymbol{\theta}) = \sum_{t=1}^T \log f(\Delta X_t | \Omega_{t-1}) = \sum_{t=1}^T \sum_{k=1}^K \pi_{kt|t-1} \frac{1}{\sqrt{2\pi\sigma_{k,t}}} \exp\left\{-\frac{1}{2} \frac{(\Delta X_t - \mu_{k,t})^2}{\sigma_{k,t}^2}\right\} \quad (3.25)$$

Extensions of the above framework for the multivariate case are straightforward. For instance, if we are looking at the joint distribution of many variables, say collected in the vector $\Delta \mathbf{X}_t$, then, define:

$$\boldsymbol{\varepsilon}_{t,k} = \begin{pmatrix} \Delta X_{1t} - \mu_{1t,k} \\ \dots \\ \Delta X_{nt} - \mu_{nt,k} \end{pmatrix} \quad \boldsymbol{\Sigma}_{t,k} = \begin{pmatrix} \sigma_{11t,k}^2 & \dots & \sigma_{1nt,k} \\ \dots & \dots & \dots \\ \sigma_{1nt,k} & \dots & \sigma_{nnt,k}^2 \end{pmatrix} \quad (3.26)$$

Under the normality assumption Eq. (3.23) and (3.25) change to:

$$f(\Delta \mathbf{X}_t | s_t = k, \Omega_{t-1}) = \frac{1}{2\pi} |\boldsymbol{\Sigma}_{k,t}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \boldsymbol{\varepsilon}_{k,t}^T |\boldsymbol{\Sigma}_{k,t}|^{-1} \boldsymbol{\varepsilon}_{k,t} \right\} = \quad (3.27)$$

$$L(\boldsymbol{\theta}) = \sum_{t=1}^T \log f(\Delta \mathbf{X}_t | \Omega_{t-1}) = \sum_{t=1}^T \sum_{k=1}^K \pi_{kt|t-1} f(\Delta \mathbf{X}_t | s_t = k, \Omega_{t-1})$$

For more on multivariate MRS models please refer to Chapters 1 and 3. We now just present that in the bi-variate case we can write:

$$\begin{aligned} L(\boldsymbol{\theta}) &= \sum_{t=1}^T \log f(\Delta \mathbf{X}_t | \Omega_{t-1}) = \\ &= \sum_{t=1}^T \sum_{l=1}^T \frac{\pi_{kl|t-1}}{2\pi\sigma_{k,1t}\sigma_{k,2t}\sqrt{1-\rho_{k,12t}^2}} \exp \left\{ -\frac{1}{2(1-\rho_{k,12t}^2)} \left[\left(\frac{\varepsilon_{k,1t}}{\sigma_{k,1t}} \right)^2 + \left(\frac{\varepsilon_{k,2t}}{\sigma_{k,2t}} \right)^2 - 2\frac{\rho_{k,12t}\varepsilon_{k,1t}\varepsilon_{k,2t}}{\sigma_{k,1t}\sigma_{k,2t}} \right] \right\} \quad (3.28) \end{aligned}$$

In summary, estimation of MRS models is obtained by: First, calculating ex-ante probabilities using the transition matrix and the filter probabilities (Eq. 3.20). Second, calculate the densities, conditional on regime realisation (Eq. 3.23). Third, integrate out the state variable to obtain the unconditional (on regime) density (Eq. 3.19). Finally, update the probabilities (Eq. 3.20 to 3.22) and repeating the procedure for the next t .

Following estimation of MRS models - conditional on the specified model- and after obtaining the n step ahead forecasts of both the state-dependent variance-covariance and mean equations as well as step ahead forecasts of the regime probabilities, first and second moment forecasts of the overall process are obtained as:

$$\begin{aligned}
 \mu_{it+n} &= E_t [\Delta X_{it+n}] = \sum_{k=1}^K \pi_{k,t+n|t} \mu_{k,t+n} \\
 \sigma_{it+n}^2 &= E_t [\Delta X_{it+n}^2] - \left(E_t [\Delta X_{it+n}] \right)^2 = \\
 &= \sum_{k=1}^K \pi_{k,t+n|t} \left(\mu_{ik,t+n}^2 + \sigma_{k,t+n}^2 \right) - \left(\sum_{k=1}^K \pi_{k,t+n|t} \mu_{ik,t+n} \right)^2 \\
 \sigma_{ij,t+n} &= E_t [\Delta X_{it+n}, \Delta X_{jt+n}] - \left(E_t [\Delta X_{it+n}] E_t [\Delta X_{jt+n}] \right) = \\
 &= \sum_{k=1}^K \pi_{k,t+n|t} \left(\mu_{ik,t+n} \mu_{jk,t+n} + \sigma_{ijk,t+n} \right) - \left(\sum_{k=1}^K \pi_{k,t+n|t} \mu_{ik,t+n} \mu_{jk,t+n} \right)^2
 \end{aligned} \tag{3.29}$$

where i, j denotes the asset with $i \neq j$. Regime probabilities at $t + n$ $\pi_{k,t+n|t}$ can be obtained by utilising the estimates of the transition matrix at time t , Ψ_t , and the estimated regime probabilities at time t as $\pi_{k|t-1} \Psi^{t+n}$, where $\pi_{k|t-1}$ the vector of the ex-ante probabilities.

After maximising the log-likelihood function and obtain the parameter vector θ , smooth probabilities can also be estimated i.e. $\pi_{k|T} = \Pr(s_t = k | \Omega_T)$. The filter recursion before can be considered as a limited information technique since not all observations of the sample are used. Inference regarding the timing of regimes can be improved by utilising all the available information up to T . An efficient algorithm to calculate these probabilities has been developed by Kim (1994). Kim's smoothing algorithm can be considered as a backward iterating procedure on the previous recursive filter i.e. from $t = T$ to $t = 1$ which is based on the following relationship:

$$\begin{aligned}
 \Pr(s_t = k | \Omega_T) &= \sum_{j=1}^K \Pr(s_t = k, s_{t+1} = j | \Omega_T) = \\
 &= \sum_{j=1}^K \Pr(s_{t+1} = j | \Omega_T) \Pr(s_t = k | s_{t+1} = j, \Omega_T)
 \end{aligned} \tag{3.30}$$

We can conveniently assume now that $\Pr(s_t = k | s_{t+1} = i, \Omega_t) \approx \Pr(s_t = k | s_{t+1} = i, \Omega_T)$, on the basis of the Markov property. Hence:

$$\begin{aligned} \Pr(s_t = k | \Omega_T) &= \sum_{j=1}^K \Pr(s_{t+1} = i | \Omega_T) \Pr(s_t = k | s_{t+1} = i, \Omega_t) = \\ &= \sum_{j=1}^K \Pr(s_{t+1} = i | \Omega_T) \frac{\Pr(s_t = k, s_{t+1} = i, \Omega_t)}{\Pr(s_{t+1} = i | \Omega_t)} = \\ &= \sum_{j=1}^K \frac{\Pr(s_{t+1} = i | \Omega_T) \Pr(s_t = k | \Omega_t) \Pr(s_{t+1} = i | s_t = k)}{\Pr(s_{t+1} = i | \Omega_t)} \end{aligned} \quad (3.31)$$

Starting from T and iterating backwards we do know all the elements of the right hand-side of Eq. (3.27) and we perform this calculation for each t . To exemplify this procedure, consider that our dataset includes only 2 observations and follows a two state process, hence $t = \{1, 2\}$. The smooth probability $\Pr(s_1 = 1 | \Omega_2)$ for the process being in state 1, with $T = 2$ given, is:

$$\frac{\Pr(s_2 = 1 | \Omega_2) \Pr(s_1 = 1 | \Omega_1) \Pr(s_2 = 1 | s_1 = 1)}{\Pr(s_2 = 1 | \Omega_1)} + \frac{\Pr(s_2 = 2 | \Omega_2) \Pr(s_1 = 1 | \Omega_1) \Pr(s_2 = 2 | s_1 = 1)}{\Pr(s_2 = 2 | \Omega_1)}$$

Table 3.2: Simple MRS Models of Petroleum Spreads

$$\Delta Spread_t = \mu_{s_t} + \alpha_{s_t} Spread_{t-1} + \varepsilon_{s_t,t}; \quad \varepsilon_t \sim N(\mu_{s_t}, \sigma_{s_t}^2)$$

	μ_1	μ_2	α_1	α_2	σ_1	σ_2	ρ_{11}	ρ_{22}
Panel A: NYMEX Heating crack spread								
OLS	0.221		-0.0155		1.007		1	
MRS	0.189	0.289	-0.0136	-0.0190	0.716	1.506	0.989	0.971
Panel B: WTI-Brent Spread								
OLS	0.060		-0.0127		0.64			
MRS	0.024	0.109	-0.0062	-0.0203	0.370	1.067	0.970	0.925

- Estimation period uses daily observations of nearby futures from June 1994 to December 2009.
- The table presents the parameters of fitting an MRS model for the NYMEX 1:1 heating crack (Panel A) and the WTI-Brent intercrude spread (Panel B).

Consider the parameter estimates in Table 3.2 as an example. Using a simple mean reverting equation, we can see that first, high variance states are associated with greater speed of

mean reversion and second, high variance states are associated with less persistent regimes. The difference in the daily volatility between the regimes is for both spreads, more or less, 70 basis points, equivalent to around 11% on an annual basis.

Following the estimation of the above MRS model we also estimate and show the ex-ante and smooth regime probabilities for the NYMEX heating crack spread (1 barrel of crude vs. 1 barrel of heating). The result is plotted in Figure 3.3.




Figure 3.3: Ex-ante & Smooth Regime Prob. for NYMEX Heating Crack





Figure 3.4: MRS Model Volatility vs. Actual Returns for NYMEX Heating Crack

In addition, the Figure 3.4 is presented to demonstrate the evolution of volatility under the MRS models. Volatility is displayed as ± 3 standard deviations vs. the actual returns. Finally, regime classification (shaded area) according to the smooth probabilities is also plotted.

3.4.2 Path Dependency in Volatility

Due to the time varying nature of the regime probabilities, the overall conditional volatility of an MRS model implies that second moments are time varying, even if the within state variances are assumed constant. An alternative specification to model volatility has been the well-known GARCH framework which by definition is an ARMA process of the variance. The first to combine these two approaches in a unified framework are Hamilton and Susmel (1994) and Cai (1994). To accommodate within regime time variation they capture the volatility dynamics using ARCH family models. However, an important assumption of any Markov chain is that the state variable process does not depend on its history.



Source: Gray, S. F. (1996). *Journal of Financial Economics*. Vol. 42, Issue 1, pp 35.

Figure 3.5: Volatility Path-Dependency in the GARCH Model

Hamilton and Susmel (1994) and Cai (1994) noted that to avoid the conditional density to be dependent on the entire history, as is inherent in the GARCH structure the AR term need not be in the variance. The main drawback of their approach is that many lags of ARCH terms are needed in order to capture the volatility dynamics. Gray (1996) was the first to suggest a possible tractable method and offer a solution for preserving in a way the popular GARCH dynamics in a regime switching scheme. Essentially, his approach was to integrate the unobserved regime at each step by using the conditional expectation of the past variance rather

then the regime specific variances (see also Chapter 4). The conditional expectation of the past variance is given by Eq. (3.29). Thus, recursive estimation of the GARCH-like equation $\sigma_{k,t}^2 = \omega_k + A_k \varepsilon_{t-1}^2 + B_k \sigma_{t-1}^2$ becomes feasible by recombining the state-dependent variances at each time step. Dueker (1997) also uses another collapsing procedure, but he essentially adopts the same framework of Gray (1996). Klaassen (2002) adopts the same recombining method but utilises the probabilities at $t-1$ rather than $t-2$. An alternative approach is offered by Haas et al. (2004a, 2004b) who assume K independent GARCH processes where all $\sigma_{k,t}^2$ exist as latent variables (see also Chapter 4). Lee and Yoder (2007a) extend Gray's model to the bivariate case and fully solve the path dependency problem by developing a similar collapsing procedure for the covariance (see also Chapter 5).

3.5 Conclusion

In this chapter we reviewed some of the empirical evidence regarding the use of regime switching models in financial applications. This chapter addressed several aspects of an important class of econometric models, namely the Markov Regime Switching (MRS) models. Pioneered by Hamilton (1989) to model the evolution of business cycles, applications of MRS models for describing nonlinear behaviour of asset returns and nonlinear dependence among assets has expanded. Introducing a great amount of flexibility in modelling the conditional distributions of asset returns, they have been applied in various fields of finance and economics, from portfolio allocation and portfolio risk to forecasting and derivatives pricing. Furthermore, this chapter has also provided the theoretical background regarding some basic concepts behind estimation issues and inference in a regime switching setting. The fundamental frameworks that will be adopted in the subsequent empirical applications have been outlined.

Non-normality, asymmetries and time-varying dependence are well-documented features in energy markets. Modelling the petroleum price economic series with a Markov Regime Switching process in their conditional means and their conditional second moments permits us to consider many of the stylised facts that these markets exhibit. Our empirical applications in the three subsequent empirical chapters will essentially confirm that the Markovian formulation turns out to exhibit many of the salient features of petroleum markets. A common feature of the ensuing models that will be presented is that regime switches are driven by an unobserved latent variable driven by a Markov Chain. In doing so, we will consider univariate and multivariate models as well as Markov GARCH models.

Proper assessment of energy risk relies on models that reflect a number of important properties of the underlying assets which affect the performance of the participants' portfolios such as time-dependent volatility and heavy tails. Next, in Chapter 4: *Forecasting Petroleum Futures Markets Volatility: The Role of Regimes and Market Conditions*, we will focus on the short end of the futures curve. The purpose is to provide an in-depth analysis of forecasting volatility and Value-at-Risk for the more volatile nearest to expiry contracts. In doing so, as it will be seen, we will use both the information of the futures curve, regime switching models and another class of models, the family of Autoregressive Conditional Heteroscedasticity (ARCH) models, introduced by Engle (1982). What makes this a non-trivial exercise is the complex dynamics of petroleum commodities, compounded with the difficult task of simultaneously modelling the volatility as a GARCH and a Markov process. Finally, we will link the performance of models to the position of the futures curve to examine whether there is a tendency for the forecast errors to be better or worse under different market conditions, such as backwardation and contango. Before presenting our empirical evidence, Chapter 4 will first, provide a short introduction regarding volatility forecasting. This will be followed by some technical details on Mixture and Markov GARCH and next, a thorough explanation of the theoretical background and the estimation procedure will be supplied. The model will be fitted to daily historical futures prices from 1991 to 2008, providing strong statistical evidence, not only regarding the presence of regime shifts but also concerning the forecasting performance.

Chapter 4

Forecasting Petroleum Futures Markets Volatility: The Role of Regimes and Market Conditions

4.1 Introduction

In the volatile world of energy markets, quantifying and mitigating price risk presents a number of challenges due to the time-dependence in volatility, non-linear dynamics and heavy tails in the distribution of oil returns. Petroleum price volatility has always been at the core of economic research agenda not only because of its effect on the cash flows of oil-related businesses, but also due to the far-reaching implications of oil price uncertainty on the macroeconomy (Hamilton, 2003 and Chen and Chen, 2007) and the financial markets (Driesprong et al., 2008 and Aloui and Jammazzi, 2009). It is not surprising therefore that in the energy economics literature there is a plethora of empirical studies examining the issue of modelling volatility and risk management.

Traditionally, the family of Autoregressive Conditional Heteroscedasticity (ARCH) models - introduced by Engle (1982) - have been widely used to describe the conditional volatility of oil prices, due to their flexibility. However, empirical research suggests that in the presence of asymmetries, fat tails and time-dependent higher order moments, the standard Generalised ARCH model of Bollerslev (1986) is not appropriate and thus, numerous extensions have been developed in the literature either by assuming different distributions of the error structure or by adding asymmetric terms, such as leverage effects, in the variance process. Kang et. al (2009) for instance, compare the forecasting ability of different GARCH models in the WTI, Brent and Dubai crude oil futures markets and find that Fractionally Integrated GARCH processes provide more accurate volatility forecasts, concluding that persistence and long memory are essential elements of energy markets volatility. Agnolucci (2009) investigates the market volatility of WTI futures and finds that extensions of GARCH models with asymmetric effects and different error distributions out-perform implied volatility models'

predictive accuracy. Fan et al. (2008) show that the assumption of normality leads to underestimation of risk and GARCH models based on the Generalised Error Distribution (GED) produce more reliable forecasts compared to ordinary GARCH models. Hung et al. (2008) also highlight the importance of selecting the appropriate distribution in a GARCH context and find that crude oil and oil products' Value-at-Risk (VaR) is better captured by fat-tail distributions. Overall, the findings of this study imply that the assumption of fat tails plays an important role in VaR estimates since it directly affects the required quantiles. Costello et al. (2008) on the other hand, employ a GARCH filter and rely on historical simulations (semi-parametric GARCH) to forecast VaR whereas Huang, Yu, Fabozzi and Fukushima (2009) employ an alternative CAViaR (Conditional Autoregressive VaR) technique based on regression quantiles. Other studies testing different variants of GARCH models include Duffie et al. (2004), Sadorsky (2006), Cheong (2009) and Wei et al. (2010).

A major shortcoming of GARCH models is that they induce a high degree of persistence in shocks, that falsely implies high predictability but, in essence reflects regime shifts or structural breaks in the volatility process (Lamoureux and Lastrapes, 1990). This means that a regime-switching GARCH model may be more suitable for modelling volatility particularly in the energy markets where structural breaks are quite common¹. Another advantage of a regime GARCH process is its ability to deal with fat-tails (see Haas et. al, 2004a and 2004b for more details and derivation of higher moments of mixed normal distributions); this is very important for modelling volatility in the oil futures markets where demand shocks result in an asymmetrically higher volatility when the market is at the steep part of the supply stack.

In addition, oil market volatility is characterised by different dynamics under different market conditions. For instance, Fong and See (2002; 2003) document strong evidence of regime switching in the temporal volatility dynamics of oil futures, consistent with the theory of storage; an increase in backwardation is more likely to increase regime persistence in the high volatility state, due to low inventories. In the next chapter we will employ a Markov Regime Switching (MRS) approach for determining optimum hedge ratios in NYMEX energy futures markets. The findings will show that in a low variance regime, error correction coefficients are in accordance with convergence towards a long-run equilibrium relationship, while the high variance state is characterised by insignificant speed of adjustment coefficients, which

¹ See for instance Wilson et al. (1996). Employing an iterative cumulative sums-of-squares (ICSS) approach, they show evidence of sudden changes in the unconditional volatility of oil futures contracts. In particular, 15 significant volatility changes were detected from 1984 to 1992, whereas 5 of these exceeded 100% in absolute terms e.g. the eight day period following the invasion of Kuwait in 1991 was associated with a 213 percent upward change in the unconditional volatility.

effectively results in a widening of the basis thus explaining the high variance regime; hence, the adjustment process undergoes regime shifts and does not behave uniformly to shocks to equilibrium across different states. Another study by Vo (2009) combined the concept of regime switching with that of stochastic volatility to forecast the dynamics of WTI crude oil. The author finds that the simple MRS model captures better the in-sample dynamics in terms of mean absolute errors whereas out-of-sample, stochastic volatility with regime shifts is favoured.

Building on these studies, this chapter investigates the volatility dynamics for the NYMEX WTI crude and heating oil as well as the ICE Brent crude and gas oil futures contracts. In doing so, it contributes to the existing literature in a number of ways. First, we employ various volatility regime models, to accommodate some of the stylised features of the oil markets such as volatility clustering, non-normality, time-varying skewness and excess kurtosis. In particular, we consider the Mix (distribution) GARCH and the MRS GARCH models based on the mixed conditional heteroscedasticity models of Haas et. al (2004a) and Alexander and Lazar (2006) and the Markov model of Haas et. al (2004b), respectively. Our study is different from the above mentioned research in the sense that we provide a thorough empirical application of the provided framework in the energy markets. Although volatility modelling and forecasting in a regime framework has been widely documented in equity and foreign exchange markets (see Marcucci, 2005; Li and Lin, 2004; Giannikis et al., 2008), few studies have analysed in depth the nature of the volatility regimes of oil futures prices and the forecasting ability of those models in the specific market.

Second, we extend previous research by including the squared lagged basis of futures prices in the specification of the conditional variance in what is termed the GARCH-X model (Lee, 1994; Ng and Pirrong, 1996). A principal feature of the basis is that the time paths of spot and futures prices are influenced by the extent of deviations from their long-run equilibrium (Engle and Granger, 1987). As prices respond to the magnitude of disequilibrium then, in the process of adjusting, they may become more volatile. If this is the case then the inclusion of the basis term in the conditional variance specification may lead to the estimation of more accurate volatility forecasts. Examining different volatility components will enable us to investigate whether the dependence of volatility to the basis changes across different regimes and uncover how these asymmetries are transmitted. To the authors knowledge this is the first time that the GARCH –X methodology is tested in a regime volatility setting. Implementing such models allows us to draw some new interesting insights regarding the effect of disequilibrium and the persistence of volatility under different market conditions.

Third, we extend the above framework to a conditional extreme value theory (EVT) setting and use the estimated volatility models as filters, in order to combine the forecasts with EVT-based methods for quantile estimation and link the regime volatility background with tail estimation. From a risk management perspective, the tails of the conditional distributions of the models may contain important information that needs to be considered. Existing literature that addresses the issue is limited to the EVT-Switching ARCH model of Samuel (2008), applied in estimating VaR in the stock index market. In the oil market there is limited evidence on conditional EVT based VaR provided by Krehbiel and Adkins (2005) for the NYMEX complex and Marimoutou et al. (2009) for WTI and Brent crude oil.

Fourth, the forecasting performance of the proposed models is assessed and contrasted using a battery of forecast statistics which measure both the tracking errors from actual volatility measures, as well as the degree of volatility under or over-prediction. In addition, we evaluate the effectiveness of the proposed models in VaR applications for both long and short positions and this way, we provide robust evidence on the performance of the proposed volatility models. VaR forecasts are assessed by means of risk management loss functions and their relative performance is ranked using White's (2000) Reality Check.

Finally, volatility and VaR forecasts are tested across periods of backwardation and contango. Many authors (see Fama and French; 1987 and Geman and Ohana; 2008) have shown that price volatility has a negative correlation with inventory levels, in line with the theory of storage. Consequently, it is worth examining the performance of different models under conditions of backwardation and contango, since the risk-return profile of energy prices is known to change fundamentally, between the two different states.

The remainder of this chapter is organised as follows. Section 4.2 demonstrates the Regime GARCH models estimation procedure. In section 4.3, the data and their properties are discussed. This is followed by an evaluation of the proposed strategies in section 4.4. Finally, conclusions are given in the last section.

4.2 Methodology

To estimate the volatility models, the methodology used in this study follows the Mix-GARCH model of Haas et. al (2004a) and Alexander and Lazar (2006) and the MRS-GARCH model of Haas et. al (2004b). Both assume more than one individual component variances and differ in the way that they treat regime probabilities. For the former, what is important is the overall regime probability; for the latter, the probability of each observation belonging to any given regime is more important. However, both models assume that asset returns are generated

from different information distributions and in this regard, they can accommodate parameter shifts or switches among a finite number of regimes; this is expected to improve the performance of these models in financial applications, such as VaR. Besides, those models are also expected to address the issue of asymmetric behaviour not only across different market conditions (regimes) but also across short and long positions.

The GARCH model in its basic form is not tractable in the Markov framework because the conditional variance is a function of all past information, rather than a function of the current regime alone, thus violating the Markov property. Gray (1996) is the first to develop a tractable MRS-GARCH model where the conditional regime variance processes are a function of the conditional expectation of the overall variance. A similar approach is proposed by Dueker (1997) and Klaassen (2002). Haas et al. (2004b) argue that inferences about the variance process within the above setting are complicated by the fact that state dependent variances are conditioned on the aggregate variance - which in turn is a function of both regime probabilities and regime variances rather than own lagged values. Consequently, based on mixture of distribution models, Haas et al. (2004b) proceed to develop a framework that allows for different GARCH behaviour in different regimes whilst preserving the direct association of the GARCH parameters within each regime². In our analysis we use the latter formulation due to its flexibility and the ease of making straightforward inferences.

Let r_t represent daily observations of the log returns on the four petroleum commodities under study. Consider the following general conditional mean and variance dynamics of the form:

$$r_t = \mu_{s_t} + \varepsilon_{s_t,t}; \quad \varepsilon_{s_t,t} \sim N(0, h_{s_t,t}) \quad (4.1)$$

$$h_{s_t,t} = \omega_{s_t} + A_{s_t} \varepsilon_{s_t,t-1}^2 + B_{s_t} h_{s_t,t-1} + \varphi_{s_t} Z_{t-1}^2 \quad s_t = \{1, 2\}$$

with $\omega_{st} > 0$ and $\alpha_{st}, \beta_{st}, \varphi_{st} \geq 0$ to guarantee nonnegative variance. $\varepsilon_{st,t}$ is a Gaussian white noise process, $h_{st,t}$ the conditional variance and Z_t the basis, defined as the difference between the nearby and second nearby futures contracts. The state variable $s_t = \{1, 2\}$ describes the

² Consider a regime shift from low to high variance state. In Gray's (1996) model the variance dynamics are determined by the last period's overall variance which was effectively driven by the low volatility regime since in the previous period the low variance state prevailed. In contrast, Haas et al. (2004b) allow the variance dynamics to be directly determined by the current state (i.e. high variance regime) since the model implies two independent GARCH processes; when a regime shift occurs, this has an immediate impact on volatility.

state/regime that the system is in. In this setting, s_t is unobserved and follows a two-state, first order Markov process with the following transition probability matrix:

$$\hat{\mathbf{P}} = \begin{pmatrix} \Pr(s_t = 1 | s_{t-1} = 1) = p_{11} & \Pr(s_t = 1 | s_{t-1} = 2) = p_{21} \\ \Pr(s_t = 2 | s_{t-1} = 1) = p_{12} & \Pr(s_t = 2 | s_{t-1} = 2) = p_{22} \end{pmatrix} = \begin{pmatrix} 1 - p_{12} & p_{21} \\ p_{12} & 1 - p_{21} \end{pmatrix} \quad (4.2)$$

where p_{12} gives the probability that state 1 will be followed by state 2, p_{22} gives the probability that there will be no change in the state of the market in the following period given that we are in state 2 etc. These transition probabilities are assumed to remain constant between successive periods. Furthermore, assuming that the state dependent residuals follow a normal distribution³, the likelihood function for the entire sample is formed as a mixture of the probability distribution of the state variables as:

$$L(\boldsymbol{\theta}) = \sum_{t=1}^T \log f(\mathbf{X}_t; \boldsymbol{\theta}) \quad (4.3)$$

$$f(\mathbf{X}_t; \boldsymbol{\theta}) = \frac{\pi_{s_t=1,t}}{2\pi\sqrt{h_{s_t=1,t}}} \exp\left(-\frac{1}{2} \frac{\varepsilon_{s_t=1,t}^2}{h_{s_t=1,t}}\right) + \frac{\pi_{s_t=2,t}}{2\pi\sqrt{h_{s_t=2,t}}} \exp\left(-\frac{1}{2} \frac{\varepsilon_{s_t=2,t}^2}{h_{s_t=2,t}}\right)$$

where $\boldsymbol{\theta} = (\mu_{s_t=1}, \mu_{s_t=2}, \omega_{s_t=1}, \omega_{s_t=2}, A_{s_t=1}, A_{s_t=2}, B_{s_t=1}, B_{s_t=2}, \varphi_{s_t=1}, \varphi_{s_t=2}, p_{s_t=1|s_{t-1}=1}, p_{s_t=2|s_{t-1}=2})$ is the vector of parameters to be estimated, $\pi_{s_t,t}$ are the regime probabilities of being in regime s_t and are calculated recursively using Bayes rule, and $L(\boldsymbol{\theta})$ is maximised using numerical optimization methods, subject to the constraints that $\pi_{1,t} + \pi_{2,t} = 1$ and $0 \leq \pi_{1,t}, \pi_{2,t} \leq 1$.

As already mentioned, the Mix-GARCH model differs from the MRS-GARCH model described above in the definition of regime probabilities. For the Mix-GARCH what is important is the overall regime probability over the total sample. Vlaar and Palm (1993) and Palm and Vlaar (1997) were the first to suggest the Mix-GARCH model. Their formulation assumes that the state 2 variance process is given by $h_{s_t,t=2} = h_{s_t,t=1} + \zeta^2$, where ζ represents a scale parameter to be estimated. Another study by Lin and Yeh (2000) allows for each variance component to change through the intercept in the variance equation. In our analysis we use the

³ Note that we also tested the models under the distributional assumption of a Generalised Error Distribution (GED). Model parameters were found to be robust irrespective of the distribution chosen and results were similar to those reported in Table 4.2. Moreover, the high variance state was, as expected, associated with lower degrees of freedom i.e. fatter tails than the low variance state. GED distribution was preferred over alternative distributions, most notably student-t, due to its flexibility to accommodate both thin and fat tails, as emphasised in the energy economics literature by Fan et al. (2008) and Hung et al. (2008). Parameter estimates of the models under the GED distribution are presented in Appendix 4.A.

Mix-GARCH formulation of two component variances with no lagged-cross equation terms, based on Haas et al. (2004a). In this case Eq (1) still applies, only now $st=\{1,2\}$ does not represent an unobserved state variable but the number of component variances. Furthermore, the likelihood function of Eq (3) also applies, but vector θ reduces to $\theta = (\mu_{st=1}, \mu_{st=2}, \omega_{st=1}, \omega_{st=2}, A_{st=1}, A_{st=2}, B_{st=1}, B_{st=2}, \varphi_{st=1}, \varphi_{st=2}, \pi_{st=1})$, since for every t , $\pi_{st=1,t} = \pi_{st=1}$, in other words, the regime probabilities are assumed to be constant.

Summarising, an important feature of the Markovian formulation is time-variation of the model parameters due to the fact that state probabilities are a function of time. Furthermore, MRS models allow for the tendency that commodity markets exhibit when an event which caused volatility to reach high levels is followed by another similar event, i.e. persistence in the regimes; in economic terms, such behaviour has significant implications in derivatives pricing, among other things, because the switching mechanism - as reflected in the transition probability matrix - provides information on the current volatility state, the probability of switching to a different state and their respective expected durations. However, occasionally MRS models do not provide accurate forecasts on an out-of-sample basis. This may be due to parameter instability between in-sample and out-of-sample periods as well as uncertainty regarding the unobserved regime, as mentioned in Engel (1994) and Marsh (2000). Another reason may be the fact that in markets that exhibit extreme price spikes these might dominate the high variance state, making the latter short-lasting and rare (see also next chapter). All these issues may be addressed, using a more parsimonious parameterisation of the regimes such as the one provided by the Mix-GARCH specification. In fact, the Mix-GARCH can be considered a restricted version of the Markov model with the rank of transition probability matrix equal to one. Finally, in order to integrate the state dependent conditional variances and conditional mean equations, Gray's (1996) integrating method applies for both models:

$$\mu_t = \pi_{st=1,t} \mu_{st=1,t} + (1 - \pi_{st=1,t}) \mu_{st=2,t} \quad (4.4)$$

$$h_t = \pi_{st=1,t} (\mu_{st=1,t}^2 + h_{st=1,t}) + (1 - \pi_{st=1,t}) (\mu_{st=2,t}^2 + h_{st=2,t}) - (\mu_t)^2 \quad (4.5)$$

4.3 Description of the Data and Preliminary Analysis

The data set for this study comprises daily futures prices for four energy commodities: NYMEX WTI crude and heating oil and ICE Brent crude and gas oil. NYMEX futures cover the period from January 23, 1991 to December 31, 2008, and ICE futures from April 19, 1991 to December 31, 2008. All daily closing futures prices of 4,485 observations are obtained from

Datastream. Assuming 252 business days in a year, the first 3,225 observations are used for estimation of the models; out-of-sample analysis is carried out using the remaining 1,260 observations i.e. 5 years. In all cases, the nearest to expiry contract is used, rolling forward to the next nearby on the first business day of the delivery month in order to mitigate the impact of thin trading and expiration effects in the estimation and forecasting results.




Figure 4.1: Log-Prices of NYMEX WTI and Heating Oil Futures

Figure 4.1 displays the evolution of log-prices for the NYMEX WTI crude and heating oil futures markets. The impact of several economic and geopolitical events is evident in this graph. From 1991 until 1995 oil prices were relatively stable and this can be attributed to the restoration of Kuwait’s oil production after the Gulf war and overproduction from the OPEC countries, in combination with weak demand. In the period 1997 and 1998, we can notice a downward trend due to tension in the Middle East and the Asian crisis. Later on, in early 1999 OPEC cut down production and prices started to increase. In combination with the relatively low US stocks, the subsequent upward trend was fuelled by several other factors, such as the 9/11 attacks, the US military action in Iraq after 2003, North Korea’s missile launches, the conflict between Israel and Lebanon in 2006, and the Iranian nuclear brinkmanship. Changes in federal oil policies also contributed to the price increases until the July 2008 peak. Afterwards, prices declined steadily due to a drop in demand for oil and the global financial crisis. The ICE market displays identical dynamics, so the corresponding figures are not displayed (see also Chapter 2, section 2.4 for more on the evolution of oil prices).

Table 4.1: Summary Statistics & Unit Root Tests for NYMEX & ICE Petroleum Futures

	WTI Crude (CL)	Heating Oil (HO)	Brent Crude (CB)	Gas Oil (GO)
Panel A: Returns' Descriptive Statistics				
Annualised Mean (%)	3.326	3.226	3.276	2.822
Annualised Vol (%)	31.62	31.14	30.30	29.92
Skew	-0.419***	-0.454***	-0.226***	-0.272***
Kurt	4.322***	3.518***	3.704***	3.686***
J-B	2,604***	1,773***	1,871***	1,864***
Q(5)	8.221	5.596	13.33***	2.856
Q(10)	18.32**	12.92	31.59***	10.18
Q ² (5)	90.67***	89.21***	177.9***	218.7***
Q ² (10)	155.1***	135.0***	247.2***	334.7***
Panel B: Unit Root Tests				
Log-Levels				
PP	-1.789	-1.816	-1.795	-2.000
KPSS	2.595***	2.091***	2.740***	1.920***
Returns				
PP	-56.42***	-58.53***	-57.64***	-56.12***
KPSS	0.098	0.122	0.064	0.068
Panel C: Estimates of 1% & 5% empirical critical values for the oil futures standardised returns				
1% tail (left)	-2.678	-2.751	-2.649	-2.632
5% tail (left)	-1.612	-1.564	-1.629	-1.627
95 % tail (right)	1.562	1.571	1.586	1.546
99 % tail (right)	2.486	2.472	2.635	2.590

- In Sample period for the NYMEX futures is from January 23, 1991 to December 19, 2003 whereas for the ICE futures is from April 19, 1991 to January 30, 2004 (3,225 daily observations each). The remaining 1,260 daily observations are used for the out-of-sample tests.
- Skew and Kurt are the estimated centralised third and fourth moments of the data, denoted \hat{m}_3 and $(\hat{m}_4 - 3)$, respectively; their asymptotic distributions under the null are $\sqrt{T}\hat{m}_3 \sim N(0, 6)$ and $\sqrt{T}(\hat{m}_4 - 3) \sim N(0, 24)$.
- J-B is the Bera and Jarque (1980) test for normality of changes in log oil prices and the statistic is $\chi^2(2)$ distributed.
- Q(5) and Q(10) are the Ljung-Box (1978) Q statistics for the 5th and 10th order sample autocorrelation of the returns series, whereas Q²(5) and Q²(10) refer to the squared returns series. These tests are distributed as $\chi^2(5)$ and $\chi^2(10)$, respectively.
- PP is the Phillips and Perron (1988) unit root test, which tests the null hypothesis that the variable is non stationary, I(1), against the alternative that the variable is stationary, I(0). KPSS is the Kwiatkowski-Phillips-Schmidt-Shin (1992) test for unit roots, which tests the null hypothesis that the variable is I(0), against the alternative that the variable is I(1).
- The standardised return is defined as $(r_t - \text{Mean})/\text{SD}$ where r_t is the daily return at time t and SD the standard deviation. Note that the absolute of the 1% critical value of a standard normal distribution is 2.326 whereas the 5% critical value is 1.645.
- Asterisks ***, **, * indicate significance at 1%, 5% and 10% levels, respectively.

Table 4.1 reports the summary statistics of the return series as well as the unit root tests. Annualised mean returns for crude oil are higher than those of the corresponding petroleum product - within each market - consistent with the unconditional annualised volatilities which

also follow a similar pattern; that is, crude oil is more volatile than the corresponding petroleum product. In addition, NYMEX futures appear to be more volatile than ICE futures. The Ljung-Box (1978) Q statistic on the first five and ten lags of the sample autocorrelation function is significant only in the Brent crude oil market, at the 1% significance level. Engle's (1982) ARCH test, carried out as the Ljung-Box Q statistic on the squared series, indicates the existence of heteroscedasticity for all the return series. According to Phillips and Perron (1988) (PP) non-parametric unit root tests, performed on the log-levels and log-differences of all four petroleum futures, all futures prices' series under study follow unit root processes, while their first differences are stationary. Kwiatkowski et al. (1992) unit root tests (KPSS) confirm the results obtained from the PP test.

The coefficients of skewness and excess kurtosis indicate departures from normality for all the returns series. In particular, the observed negative skewness coefficients imply that long positions are associated with greater risk since more extreme losses are placed on the left side of the distribution of oil returns. The existence of fat-tails in the underlying series is also evidenced by calculating the empirical critical values of the standardised returns from the historical distributions. These imply that all futures returns series are fat-tailed relative to the 1% left and right tail regions, since the historical quantiles are greater in absolute value than the 1% critical value of standard normal distribution, i.e. 2.326. At the 5% tail regions the standardised returns are thin-tailed, since historical quantiles are less than the 1.645 critical value. Fat tails at the 1% regions imply that extreme events have higher probability of occurrence relative to the standard normal distribution.





Figure 4.2: NYMEX WTI Crude Oil Futures Log – Returns (standardised) Historical PDF vs. Standard Normal PDF



**Figure 4.3: Heating Oil # 2 Futures Log – Returns (standardised)
Historical PDF vs. Standard Normal PDF**

Left skewness and tail-fatness are confirmed graphically in Figures 4.2 and 4.3. These illustrate the historical distribution of standardised returns and their deviation from the corresponding normal probability density function. One potential explanation for the non-normality may be the existence of structural changes in the series (see Li and Lin, 2004) which can be captured by a regime model since these models assign different weights to different states of the market and, effectively, presuppose that sub-samples of the estimation period are drawn from different distributions.

4.4 Empirical results

This section presents the empirical results on the dynamics of the augmented regime volatility models of oil futures. First, the results of MRS- and Mix- GARCH models are presented; then, the out-of-sample forecasting performance of the proposed volatility models is compared to that of the benchmark restricted versions of those models, without the squared basis term; and, finally, the performance of the models is also assessed using risk management loss functions in VaR applications. Markov and Mixed distribution GARCH models are estimated assuming two regimes. The choice of a two-regime process is motivated by the fact that this model captures the dynamics of oil futures in a more efficient way and is intuitively appealing since these two regimes can be associated with periods of low and high volatility (see also Appendix 4.B). Table 4.2 presents the estimation results for the two regime GARCH-X models.

Table 4.2: Estimates of Switching GARCH-X Models for NYMEX & ICE Petroleum Futures

$$r_t = \mu_{st} + \varepsilon_{st,t}; \quad \varepsilon_{st,t} \sim N(0, h_{st,t})$$

$$h_{s_t,t} = \omega_{s_t} + A_{s_t} \varepsilon_{s_t,t-1}^2 + B_{s_t} h_{s_t,t-1} + \varphi_{s_t} Z_{t-1}^2 \quad s_t = \{1, 2\}$$

	WTI Crude Oil (CL)		Heating Oil #2 (HO)		Brent Crude Oil (CB)		Gas Oil (GO)	
	Mix- GARCH-X	MRS- GARCH-X	Mix- GARCH-X	MRS- GARCH-X	Mix- GARCH-X	MRS- GARCH-X	Mix- GARCH-X	MRS- GARCH-X
Panel A: Low Volatility Regime								
$E[\sigma_{1t}]$	19.66	19.70	20.95	21.02	18.01	18.03	18.64	19.12
$\mu_{0,st=1}$	0.0413 (0.034)	0.0358 (0.035)	0.0618 (0.034)*	0.0571 (0.035)*	0.0269 (0.033)	0.0239 (0.033)	0.0137 (0.035)	0.0139 (0.035)
$\omega_{st=1}$	0.0126 (0.005)***	0.0128 (0.005)***	0.012 (0.005)**	0.0114 (0.005)**	0.0089 (0.003)**	0.0087 (0.004)**	0.002 (0.002)	0.0019 (0.003)
$A_{st=1}$	0.0191 (0.004)***	0.0191 (0.004)***	0.0192 (0.004)***	0.0192 (0.004)***	0.0179 (0.004)***	0.018 (0.004)***	0.0201 (0.004)***	0.0231 (0.004)***
$B_{st=1}$	0.9621 (0.007)***	0.9608 (0.007)***	0.965 (0.006)***	0.9647 (0.006)***	0.9638 (0.007)***	0.9634 (0.007)***	0.9672 (0.005)***	0.9625 (0.006)***
$\varphi_{st=1}$	0.0044 (0.001)***	0.0046 (0.001)***	0.0016 (0.001)**	0.0016 (0.001)**	0.0049 (0.002)***	0.0049 (0.002)***	0.0007 (0.001)	0.0006 (0.001)
$\pi_{st=1}$	0.8456 (0.034)***	0.8243 -	0.8571 (0.037)***	0.8413 -	0.812 (0.040)***	0.8034 -	0.7764 (0.033)***	0.7641 -
p_{11}	- -	0.8043 (0.047)***	- -	0.8209 (0.049)***	- -	0.7949 (0.049)***	- -	0.7243 (0.047)***
Panel B: High Volatility Regime								
$E[\sigma_{2t}]$	41.99	41.87	43.09	42.73	37.57	37.57	34.98	35.29
$\mu_{0,st=2}$	-0.0918 (0.171)	-0.0273 (0.156)	-0.2069 (0.186)	-0.1496 (0.165)	0.0158 (0.132)	0.0342 (0.127)	0.0933 (0.148)	0.0888 (0.140)
$\omega_{st=2}$	0.7523 (0.538)	0.7782 (0.479)	1.0758 (0.730)	1.0652 (0.625)*	0.4436 (0.268)*	0.4546 (0.266)*	0.6363 (0.216)	0.6902 (0.205)
$A_{st=2}$	0.4362 (0.226)**	0.4804 (0.231)**	0.3651 (0.222)	0.4038 (0.231)*	0.363 (0.161)**	0.3779 (0.172)**	0.2939 (0.115)***	0.3257 (0.123)***
$B_{st=2}$	0.7441 (0.130)***	0.716 (0.126)***	0.7274 (0.145)***	0.7064 (0.137)***	0.7848 (0.089)***	0.7745 (0.094)***	0.7384 (0.089)***	0.7103 (0.088)***
$\varphi_{st=2}$	0.1137 (0.096)	0.1156 (0.087)	0.0825 (0.061)	0.0841 (0.055)	0.1002 (0.080)	0.1014 (0.078)	0.1207 (0.059)**	0.1295 (0.057)**
$\pi_{st=2}$	0.1544 (0.034)***	0.1757 -	0.1429 (0.037)***	0.1587 -	0.188 (0.040)***	0.1966 -	0.2236 (0.033)***	0.2359 -
p_{22}	- -	0.082 (0.051)	- -	0.0508 (0.055)	- -	0.1616 (0.069)**	- -	0.1068 (0.062)*
Panel C: Diagnostics								
LogLik	-6473.0	-6471.9	-6470.0	-6468.6	-6316.7	-6316.6	-6262.6	-6261.3
SBIC	13,035	13,041	13,029	13,034	12,722	12,730	12,614	12,619
$E[\sigma_t]$	24.49	25.06	25.38	25.75	23.00	23.21	23.32	23.95
$A_1 + B_1$	0.981	0.980	0.984	0.984	0.982	0.981	0.987	0.986
$A_2 + B_2$	1.180	1.196	1.092	1.110	1.148	1.152	1.032	1.036
Skew	-0.096	-0.098	0.004	0.012	0.007	0.008	0.722***	0.758***
Kurt	2.77***	2.66***	2.38***	2.31***	2.92***	2.89***	9.56***	9.86***
Q(10)	6.63	6.08	4.36	4.06	20.63**	20.33**	8.97	8.36
$Q^2(10)$	13.39	8.68	15.09	10.69	12.09	9.61	2.07	1.43

Figures in (·) are the estimated standard errors; LogLik is the Log-Likelihood function; SBIC is the Schwarz (1978) Bayesian Information Criterion; $E[\sigma_{1t}]$, $E[\sigma_{2t}]$ are the annualised unconditional volatilities in the low and high volatility states, respectively whereas $E[\sigma_t]$ is the corresponding figure for the aggregate variance process; $A_{st} + B_{st}$ is the regime specific degree of volatility persistence; MRS-GARCH-X models are the two regime GARCH-X models defined in Eq. (4.1) to (4.5); Mix-GARCH-X models are defined by the same equations but with a restricted transition probability matrix (with a rank equal to one) and constant mixing weights i.e. in Eq. (4.5), $\pi_{st=1,t} = \pi_{st=1}$ is a parameter to be estimated along with the other parameters of the model; See also notes in Table 4.1.

Regarding the coefficients of both the Mix- and MRS- GARCH equations the pattern is similar for all four commodities. First, there is marked asymmetry across regime variances, suggesting that the dynamics of the variance are different under the two regimes. The long term variances in the high volatility regimes are almost twice as large as the corresponding figures of the stable regime. Second, the degree of persistence in the variance, measured by the sum of $A_{s_t} + B_{s_t}$ coefficients for $s_t = 1, 2$, indicates that low variance states are characterised by lower persistence in volatility, whereas in the high variance state persistence increases. This is in line with other studies in the literature such as Gray (1996), Haas et al. (2004a; 2004b) and Alexander and Lazar (2006) and in the oil futures markets with Fong and See (2002). Also, all high volatility states are explosive; however, note that the overall variance process is covariance stationary in all cases.

In addition, as measured by the coefficients A_{s_t} , which show the sensitivity to shocks and B_{s_t} , which show the memory regarding market events, the low volatility state is associated with low sensitivity to shocks that nevertheless have long memory and die out very slowly; this is evidenced by the high values of the lagged variance coefficient, above 0.96 in all cases. On the other hand, in the high volatility regime (state 2) shocks that occur in the market tend to affect the variance more but die out much faster. This is confirmed visually in Figure 4.4 which displays the two volatility processes for the WTI crude oil market. The stable regime appears to be smooth, whereas in the high volatility state the process is more erratic. The overall volatility process, calculated using Eq. (4.5) is also presented in the graph as the line which lies between the two state variances.

Furthermore, the coefficients of the lagged squared basis⁴ are significant at the 1% level for the two crudes and at the 5% level for heating oil, in the low variance state. This can be attributed to the fact that under normal market conditions, the dynamics of the volatilities are expected to be more predictable and deviations from the equilibrium appear to have a certain degree of explanatory power on volatility. On the other hand, when the market is in the high volatility state, volatility movements occur mainly due to short-lived random shocks which are difficult to foresee, as is also shown in Figure 4.4. This is also evident from the estimated unconditional probability of being in the low variance regime, $\pi_{st=1}$, which is close to 80%, across all commodities, indicating that the high variance regime is of relatively short duration as opposed to the low variance regime which is the prevailing stable state. However, this does not

⁴ For the basis we performed a cointegration test using the Johansen (1988) procedure. The results of the λ_{trace} and λ_{max} statistics indicate in all cases that the components of the basis stand in a long-run relationship, at conventional significance levels. Moreover, parameter restriction tests on the cointegrating relationship indicate that there is a one to one relationship between the two nearest to expiry contracts, at 1% significance level.

hold in the ICE gas oil market, where the lagged squared basis term is significant at the 5% level only in the high variance state. This may be due to the fact that the specific market has the highest probability of occurrence for the high variance state as well as the lowest unconditional annualised volatility in the high variance state (35% p.a.). Also, note that although in all the other markets the state 2 volatility is more than double the state 1 volatility, in the gas oil case this does not hold; this may be another factor contributing to this since the regime dependent volatility dynamics appear to be relatively similar across the two regimes, compared to the other markets.




Figure 4.4: WTI Crude Oil Futures Volatility




Figure 4.5: WTI Crude Oil Regime Probabilities of being in the Stable Regime (MRS-GARCH-X Model)

Finally, considering the MRS models, Figure 4.5 demonstrates the evolution of the regime probability of being in the low variance state for the WTI crude oil market as estimated by the MRS-GARCH-X model. It can be observed that switching occurs frequently, however, the high variance state is short lived, in line with the estimated transition probabilities of the MRS models in Table 4.2 - the probability of staying in the high variance state, p_{22} , ranges between 8% and 16% - and other studies in the literature such as Haas et. al (2004b).

Diagnostic tests for all models are also presented in Panel C of Table 4.2. Tests on the standardised residuals, $\varepsilon_t/(h_t)^{1/2}$, and standardised squared residuals $\varepsilon_t^2/(h_t)$, indicate that there are no significant signs of autocorrelation at the 1% significance level. Moreover, by comparing the unconditional and conditional coefficients of skewness in Table 4.1 and 4.2, we can note that both models achieve to eliminate the excess skewness for all the commodities with the exception of the ICE gas oil market. For the same group of commodities, there is a nominal reduction in the level of excess kurtosis which nevertheless still remains significantly different from zero in all cases. Finally, the negative of the Log-likelihood function is maximised for the MRS-GARCH models whereas the Schwarz Bayesian Information Criterion (SBIC) suggests that the Mix-GARCH models provide a more parsimonious representation of the volatility process.

4.4.1 Out-of-Sample Performance of Volatility Forecasts

In order to further examine the appropriateness of our volatility models, we test the performance of the proposed models in predicting the volatility of energy prices. The benchmark models considered in each case are two-regime MRS-GARCH and Mix-GARCH – with the restriction that the coefficient of the lagged basis term is zero - as well as single regime GARCH and GARCH-X models. This is done by estimating each model over the period January 1991 to December 2003, for NYMEX futures, and April 1991 to December 2003, for ICE futures, and leaving the last five years (1,260 daily observations) for out-of-sample forecasting. We perform one-step ahead forecasts of the state dependent variances of the regime models and obtain the one-step ahead forecast at time $t+1$ using Eq. (4.1)⁵. In order to reduce the computational burden of this process we update the parameters of the model once a month (every 20 business days).

⁵ In the case of the MRS-GARCH models, estimates of the transition matrix at time t , $\hat{\mathbf{P}}_t$, and the estimated regime probabilities, $\Pr(s_t = 1) = \hat{\pi}_{1,t}$ and $\Pr(s_t = 2) = \hat{\pi}_{2,t}$ are used to forecast regime probabilities at time $t+1$.

Since volatility is an unobserved variable, we compare the accuracy of out-of-sample volatility forecasts from different models against the realised squared returns. Forecast performance is assessed using the Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) which measure how close the variance estimates track the changes in the markets (see notes in Table 4.3 for further details on how these measures are calculated). However, none of the abovementioned metrics provide any information on the asymmetry of the prediction variance errors; that is, whether there is any difference between forecast errors when the model over-predicts or under-predicts the actual variance. This is an important forecast metric because, although we expect forecast errors to be unbiased on average, there might be occasions when a model produces small errors but consistently over-predicts or under-predicts the variance. Thus, we also look at the proportion of negative and positive forecast errors for each model, since a model with symmetric forecast errors should produce about 50% positive and 50% negative forecast errors, with similar means. For that, we use the Brailsford and Faff (1996) Mixed Mean Error statistic, which uses a mixture of positive and negative forecast errors with different weights:

$$MME(O) = \frac{1}{N} \left[\sum_{i=1}^U |\hat{\sigma}_{t+i}^2 - r_{t+i}^2| + \sum_{i=1}^O \sqrt{|\hat{\sigma}_{t+i}^2 - r_{t+i}^2|} \right] \quad (4.6)$$

$$MME(U) = \frac{1}{N} \left[\sum_{i=1}^U \sqrt{|\hat{\sigma}_{t+i}^2 - r_{t+i}^2|} + \sum_{i=1}^O |\hat{\sigma}_{t+i}^2 - r_{t+i}^2| \right] \quad (4.7)$$

The results are presented in panels A to D of Table 4.3. First of all it can be observed from the RMSE's and MAE's that the errors are smaller for the petroleum products compared to their corresponding crudes. The same, but with less pronounced effect, holds for the overall over- and under- prediction statistics, of the volatility forecasts. This is in line with the historical figures of volatility in Table 4.2, where volatilities for products seem to be lower. Second, we can note that percentage over-prediction occurs more often than under-prediction across all alternative volatility forecasting methods, and across all commodities. On average, all models over-predict volatility 70% of the time.

Table 4.3: Comparisons of Out-Of-Sample Forecasting Performance of Volatility Models

	<i>Overall Volatility Error Statistics</i>							<i>Backwardation Volatility Error Statistics</i>							<i>Contango Volatility Error Statistics</i>						
	RMSE	MAE	O (%)	MME (O)	U (%)	MME (U)	W-Sum	RMSE	MAE	O (%)	MME (O)	U (%)	MME (U)	W-Sum	RMSE	MAE	O (%)	MME (O)	U (%)	MME (U)	W-Sum
Panel A: NYMEX WTI Crude Oil																					
GARCH	10.61*	5.667***	74.8	3.716***	25.2	4.071	3.805***	6.840	4.899***	76.5	3.631***	23.5	3.321	3.558***	12.92	6.300***	73.5	3.786***	26.5	4.690	4.026***
Mix-GARCH	10.48	5.168***	68.6	3.055***	31.4	4.077	3.376**	6.743	4.379***	69.6	2.878***	30.4	3.392**	3.034***	12.77	5.820	67.7	3.201*	32.3	4.643	3.667
MRS-GARCH	10.72**	5.429***	69.2	3.337***	30.8	4.116	3.577***	7.012**	4.688***	70.0	3.254***	30.0	3.396**	3.297***	13.01*	6.041***	68.6	3.405***	31.4	4.711*	3.815**
GARCH-X	10.46	5.357***	71.6	3.323***	28.4	4.066	3.534***	6.761	4.568***	72.3	3.172***	27.7	3.354	3.222***	12.73	6.009***	71.0	3.447***	29.0	4.654	3.797***
Mix-GARCH-X	10.49	5.060	67.2	2.887	32.8	4.101	3.285	6.746	4.228	67.5	2.655	32.5	3.414**	2.902	12.78	5.748	67.0	3.079	33.0	4.669	3.604
MRS-GARCH-X	10.59	5.309***	68.3	3.233***	31.7	4.067	3.497***	6.943**	4.522***	68.2	3.013***	31.8	3.425**	3.144***	12.84	5.958**	68.4	3.416***	31.6	4.598	3.790**
Panel B: NYMEX Heating Oil # 2																					
GARCH	8.924	4.965	67.6	3.022	32.4	3.909	3.309	8.745	5.069	68.0	3.122	31.9	3.937	3.379	9.115	4.851	67.1	2.914	32.9	3.877	3.231
Mix-GARCH	8.909	5.045**	67.5	3.141***	32.5	3.896	3.386**	8.771	5.167*	67.5	3.220**	32.5	3.963	3.461**	9.057	4.912**	67.6	3.054***	32.4	3.822	3.303**
MRS-GARCH	9.095***	5.287***	67.5	3.452***	32.5	3.879	3.591***	9.011***	5.464***	67.6	3.590***	32.4	3.950	3.707***	9.186**	5.093***	67.4	3.301***	32.6	3.801	3.464***
GARCH-X	8.962	4.990	67.0	3.032	33.0	3.928*	3.328	8.805	5.119	67.8	3.151	32.2	3.965	3.413	9.130	4.848	66.1	2.901	33.9	3.887*	3.235
Mix-GARCH-X	8.946	4.963	66.9	2.988	33.1	3.939**	3.303	8.797	5.078	66.9	3.062	33.1	4.006**	3.374	9.105	4.837	66.9	2.908	33.1	3.866	3.225
MRS-GARCH-X	9.027**	5.157***	68.0	3.252***	32.0	3.917	3.465***	8.889**	5.283***	68.2	3.350***	31.8	3.968	3.547***	9.175*	5.018***	67.8	3.144***	32.2	3.861	3.375***
Panel C: ICE Brent Crude Oil																					
GARCH	10.06**	5.546***	77.6	3.835***	22.4	3.827	3.833***	6.489***	4.881***	79.6	3.768***	20.4	3.185	3.649***	12.12*	6.061***	76.1	3.887***	23.9	4.324	3.991***
Mix-GARCH	9.954	4.773*	68.4	2.824*	31.6	3.820	3.139	6.252	4.012**	69.6	2.634***	30.4	3.171	2.797**	12.06	5.362	67.5	2.971	32.5	4.324	3.411
MRS-GARCH	10.07**	5.019***	69.3	3.144***	30.7	3.805	3.347***	6.426*	4.261***	70.7	2.926***	29.3	3.191	3.004***	12.16*	5.605***	68.2	3.313***	31.8	4.281	3.621**
GARCH-X	9.847	5.096***	74.0	3.300***	26.0	3.778	3.424***	6.305	4.346***	76.2	3.100***	23.8	3.152	3.112***	11.89	5.676***	72.3	3.454***	27.7	4.262	3.678***
Mix-GARCH-X	9.860	4.680	67.1	2.713	32.9	3.806	3.073	6.255	3.873	67.3	2.434	32.7	3.183	2.679	11.93	5.305	66.9	2.928	33.1	4.288	3.378
MRS-GARCH-X	10.05*	4.909***	67.8	3.013***	32.2	3.786	3.262*	6.363	4.059**	68.4	2.654***	31.6	3.197	2.826**	12.16*	5.566**	67.3	3.291**	32.7	4.242	3.602*
Panel D: ICE Gas Oil																					
GARCH	7.475	4.642***	71.8	3.112***	28.2	3.467	3.212***	7.679	4.861***	71.7	3.248***	28.3	3.601	3.348***	7.291	4.449***	71.9	2.992***	28.1	3.349	3.092***
Mix-GARCH	7.470	4.470***	69.6	2.883***	30.4	3.460	3.058***	7.648	4.663***	69.2	2.988***	30.8	3.589	3.173**	7.310	4.299**	69.9	2.790***	30.1	3.346	2.957**
MRS-GARCH	7.553**	4.594***	69.8	3.023***	30.2	3.473	3.159***	7.731	4.820***	69.5	3.171***	30.5	3.597	3.301***	7.392**	4.395***	70.1	2.892***	29.9	3.364	3.033***
GARCH-X	7.457	4.511***	70.2	2.944***	29.8	3.460	3.098***	7.637	4.732***	70.4	3.098***	29.6	3.579	3.240***	7.294	4.316***	70.1	2.809***	29.9	3.354	2.972***
Mix-GARCH-X	7.490	4.344	66.7	2.684	33.3	3.494*	2.954	7.641	4.502	65.6	2.753	34.4	3.614	3.049	7.354	4.204	67.7	2.622	32.3	3.388	2.869
MRS-GARCH-X	7.521**	4.480***	68.3	2.873***	31.7	3.478	3.065***	7.696	4.678***	68.2	2.980***	31.8	3.609	3.180**	7.362*	4.306**	68.5	2.778***	31.5	3.362	2.962**

For the out-of-sample tests 1,260 volatility forecasts (5 years of data) are obtained by the rolling window forecasting scheme (3,225 in-sample observations at each step); See also Notes in Tables 4.1 & 4.2; The MRS- and Mix- GARCH models are restricted versions of Eq. (4.1). The restriction that applies for both is that $\phi_{st=1} = \phi_{st=2} = 0$; RMSE is the root mean squared error of each volatility forecast compared to the

realised squared demeaned returns, whereas MAE is the corresponding mean absolute error. These measures are calculated as $RMSE = \sqrt{\frac{\sum_{i=1}^N (\hat{\sigma}_{t+i}^2 - r_{t+i}^2)^2}{N}}$ and $MAE = \sum_{i=1}^N \frac{|\hat{\sigma}_{t+i}^2 - r_{t+i}^2|}{N}$, respectively, where N

represents the number of forecasts, MME(O) and MME(U) are Mixed Mean Error statistics (Brailsford and Faff, 1996) for comparisons of asymmetries in volatility forecasts; Mean Over (Under) Prediction is the average of forecast errors when predicted volatility is higher (lower) than the realised one (see Eq. 4.6 and 4.7). Percentage (%) U and O is the proportion of under prediction and over prediction, respectively, over the forecast period; All the error statistics are multiplied by 100; The column named W-Sum is the weighted summation of the Mean Over and Under Prediction error weighted according to the estimates % U and O, respectively; Asterisks ***, **, *, indicate that the loss function of the corresponding model is significantly higher than that of competing models at 1%, 5% and 10%, respectively; the p-values are provided from White's (2000) Reality Check using the stationary bootstrap of Politis and Romano (1994); The number of bootstrap simulations is set to 3,000 and the smoothing parameter is $q = 0.1$.

Furthermore, looking at the scale of over- and under- prediction errors, on average, mean under-prediction is higher than mean over-prediction, implying that all models fail to capture the large sudden jumps of volatility, which is nevertheless expected since jumps are due to random shocks that are very hard to predict. Asymmetric error statistics have important implications for different players of the energy markets. For instance, a regulatory body such as a bank (lender) that has financed a company's energy project (e.g. for oil exploration and extraction) may prefer a model which over-predicts risk since the company (borrower) would be required to allocate more funds for capital adequacy requirements. Conversely, energy companies, depending on their risk aversion, would prefer a model that 'efficiently' under-predicts risk, since this way they have to allocate fewer resources for future risks. Another example is in the pricing of oil options. In particular, under-prediction of volatility is undesirable to the writers of options since it leads to a downward bias of the option price.

Linking the shape of the forward curve to the magnitude of forecast errors has not been addressed in the context of commodities volatility forecasting. Our definition of backwardation (contango) market days is short-term, based on the two nearest to expiry futures contracts and it occurs when the second nearby futures price is less (greater) than the prompt month price. Therefore, we examine whether the performance of volatility forecasts differs over alternative market conditions and in what respect; in general, we would expect high volatility levels under backwardation and low volatility levels under contango since, due to the highly inelastic supply of oil in the short run, demand shocks usually cause price jumps. This applies to the whole forward curve since when supplies are short - in the case of inverted markets - correlations between spot and deferred futures prices decrease due to the abovementioned increases in futures and spot volatilities, whilst when the market is at full carry and inventories are high correlation increases (see Ng and Pirrong; 1994).

A striking result which can be observed in all markets but for gas oil, is that forecast errors are larger in periods of contango. This pattern contradicts the theory of storage and essentially implies that volatility dynamics under backwardation are more predictable. This finding directed us to divide the sample in sub-periods and examine in more depth the evolution of the errors throughout the out-of-sample period. Results are presented in Table 4.4. We can see that overall results are materially different for 2008. In contango, the futures contracts present a positive slope, which can normally be explained by interest and storage costs. Although strong contango has been evidenced in the second half of 2008, volatility levels rose significantly in that period (see also Figure 4.6 which shows that returns are much more erratic in 2008). It is possible that price formation during the turbulent year of 2008 was less dependent

on the fundamental drivers of supply and demand. The second half of 2008 saw a sudden drop in the price of oil (see as well Figure 4.1) and was a period of extraordinary market conditions characterised by a tightening in the availability of credit after the global economic downturn, postponement or cancellation of investments for the development of future petroleum production capacity, as well as low liquidity in commodity derivatives markets. It is possible that the combination of those factors, coupled with the fact that during that period there were selling pressures in the market from speculators who wanted to liquidate their positions, have interacted to link higher volatility levels with contango. Looking at Table 4.4, we can see that this has been indeed the case. Up to 2007 backwardation is associated with higher forecast errors as suggested by the theory of storage. In 2008 this picture is reversed and we can see the magnitude of errors is 2-3 times higher in contango. In addition, our results imply that backwardation related volatility is more or less the same in the two sub-periods; it is contango volatility that increased significantly in 2008. Overall, for the period 2004 to 2007, the annualised average volatility forecasts under backwardation are within the range of 30% to 37%, and under contango between 28% to 35%, whereas the corresponding figures for 2008 are 30% to 37% and 33% to 40%, corresponding to an average annualised increase of 500 basis points. Also, note that although petroleum products in 2004-2007 were more predictable than the two crudes, this is also reversed in 2008.

TABLE 4.4: Comparisons of out-of-sample forecasting performance of volatility models under different periods

	2004-2007												2008											
	Overall Statistics				Backwardation Statistics				Contango Statistics				Overall Statistics				Backwardation Statistics				Contango Statistics			
	MVol	RMSE	MAE	W-Sum	MVol	RMSE	MAE	W-Sum	MVol	RMSE	MAE	W-Sum	MVol	RMSE	MAE	W-Sum	MVol	RMSE	MAE	W-Sum	MVol	RMSE	MAE	W-Sum
Panel A: NYMEX WTI Crude Oil																								
GARCH	36.15	5.861**	4.353***	3.277***	37.03	6.562	4.834***	3.528***	35.42	5.224***	3.965***	3.075***	46.63	20.63	10.92***	6.243	35.60	7.804	5.149***	3.670***	38.82	27.43	16.08	9.266
Mix-GARCH	31.61	5.717	3.828**	2.742**	32.54	6.475	4.341*	3.028**	30.83	5.019	3.413	2.513	45.63	20.47	10.53*	6.049	31.25	7.671	4.520***	3.050*	34.84	27.22	15.90	9.186
MRS-GARCH	33.03	5.954***	4.084***	2.955***	34.39	6.804**	4.659***	3.294***	31.89	5.165**	3.618***	2.682***	46.09	20.81*	10.81***	6.215	33.61	7.750	4.794***	3.297***	37.44	27.69**	16.19	9.286
GARCH-X	33.79	5.740	4.036***	2.954***	34.56	6.475	4.518***	3.207***	33.16	5.066	3.647***	2.752***	45.72	20.38	10.64***	6.125	33.94	7.750	4.758***	3.280***	36.06	27.09	15.90	9.169
Mix-GARCH-X	30.78	5.708	3.744	2.663	31.31	6.477	4.220	2.916	30.35	4.998	3.358	2.463	43.98	20.48	10.33	6.021	31.01	7.682	4.257	2.850	33.09	27.24	15.76	9.215
MRS-GARCH-X	32.38	5.897**	4.007***	2.882***	33.30	6.721*	4.537***	3.179***	31.63	5.134**	3.578***	2.645***	46.52	20.52	10.51*	6.134	33.73	7.728	4.466*	3.011	35.50	27.29	15.93	9.154
Panel A: NYMEX Heating Oil #2																								
GARCH	33.58	7.360	4.410	3.046	34.98	8.986	5.185	3.435	31.76	4.633	3.456	2.566	38.35	13.47	7.187	4.483	34.14	7.297	4.441	3.083	35.86	16.39	9.054	5.449
Mix-GARCH	34.14	7.386	4.514*	3.136**	35.58	9.012	5.308*	3.540**	32.28	4.664	3.538**	2.641***	39.08	13.37	7.170	4.504	35.20	7.317	4.400	3.055	36.60	16.24	9.053	5.463
MRS-GARCH	35.51	7.592***	4.751***	3.336***	37.27	9.263**	5.600***	3.788***	33.21	4.797***	3.706**	2.782***	40.93	13.53*	7.430**	4.719**	37.08	7.490*	4.722**	3.282**	38.25	16.41	9.272**	5.690***
GARCH-X	33.71	7.407	4.461	3.084	35.26	9.050	5.259	3.487	31.70	4.649	3.479	2.584	37.47	13.50	7.105	4.460	35.63	7.330	4.357	3.025	36.34	16.41	8.973	5.431
Mix-GARCH-X	33.40	7.392	4.435	3.058	34.75	9.032	5.218	3.451	31.66	4.638	3.471	2.576	37.17	13.47	7.077	4.438	35.07	7.385	4.317	2.998	35.67	16.37	8.953	5.385
MRS-GARCH-X	34.72	7.492**	4.618***	3.218***	36.35	9.139*	5.430***	3.637***	32.61	4.737**	3.618***	2.702***	38.85	13.52*	7.312**	4.604*	36.91	7.382	4.484	3.101	37.57	16.44*	9.236*	5.611*
Panel A: ICE Brent Crude Oil																								
GARCH	36.64	5.439***	4.292***	3.356***	37.69	6.292**	4.873***	3.661***	35.73	4.606***	3.806***	3.102***	46.74	19.67	10.56*	6.106	36.96	7.406**	4.926***	3.588***	40.45	23.98	13.75	7.889
Mix-GARCH	30.30	6.484	3.460**	2.517***	31.44	6.043	4.007**	2.807**	29.31	4.172	3.002*	2.275*	43.69	19.77*	10.02	5.826	30.33	7.214	4.037	2.745	35.31	24.14*	13.41	7.742
MRS-GARCH	31.47	5.287**	3.668***	2.688***	32.98	6.270*	4.281***	3.033***	30.14	4.296**	3.155***	2.400***	46.56	19.87*	10.42*	6.149*	32.71	7.158	4.158**	2.869**	37.18	24.27*	13.96*	8.100*
GARCH-X	33.59	5.197*	3.807***	2.864***	34.48	6.089	4.342***	3.129***	32.82	4.311***	3.360***	2.641***	45.84	19.41	10.25	5.949	33.63	7.297	4.371***	3.033***	37.97	23.66	13.57	7.855
Mix-GARCH-X	29.31	5.099	3.358	2.423	30.13	6.039	3.869	2.683	28.60	4.152	2.931	2.207	43.78	19.55	9.970	5.908	29.45	7.249	3.892	2.684	33.91	23.84	13.40	7.768
MRS-GARCH-X	30.33	5.221*	3.529***	2.562***	31.38	6.185	4.077**	2.849**	29.43	4.252	3.071**	2.323**	46.88	19.91*	10.43*	6.231*	31.35	7.199	3.971	2.738	35.45	24.31*	14.08*	8.230*
Panel A: ICE Gas Oil																								
GARCH	33.49	6.514	4.229***	3.008***	35.15	7.650	4.914***	3.380***	31.87	5.250	3.597***	2.670***	37.29	10.47	6.295***	4.032***	35.65	7.805	4.623***	3.210***	37.43	12.07	7.530*	4.713
Mix-GARCH	32.10	6.484	4.052***	2.848***	33.94	7.610	4.745***	3.231**	30.31	5.232	3.412***	2.498***	35.98	10.53	6.141*	3.908	34.28	7.815	4.292	2.906	36.80	12.15	7.505*	4.673
MRS-GARCH	32.78	6.572*	4.178***	2.947***	34.83	7.691	4.893***	3.350***	30.76	5.334*	3.517***	2.578***	36.57	10.61	6.260***	4.016***	35.20	7.911	4.489**	3.072**	37.78	12.22	7.567**	4.744*
GARCH-X	32.63	6.478	4.105***	2.900***	34.51	7.607	4.802***	3.288***	30.80	5.223	3.461***	2.545***	35.98	10.49	6.137*	3.901	35.84	7.769	4.416**	3.023**	37.12	12.12	7.407	4.621
Mix-GARCH-X	30.73	6.484	3.928	2.740	32.26	7.606	4.578	3.105	29.25	5.237	3.327	2.414	34.31	10.60	6.009	3.819	33.81	7.790	4.158	2.794	34.79	12.26	7.375	4.629
MRS-GARCH-X	31.88	6.534	4.063***	2.850***	33.70	7.658	4.751***	3.229**	30.11	5.287	3.428***	2.504***	35.69	10.58	6.148*	3.941*	35.92	7.864	4.347	2.956	36.11	12.21	7.477	4.699

In this Table the out-of-sample period is divided into two sub-samples, 2004-2007 and 2008. The first period consists of 1,008 observations and the days that the WTI crude oil, heating oil, Brent crude oil and gas oil are in backwardation (contango) are 451 (557), 556 (452), 459 (549) and 484 (584), respectively. The second period consists 252 observations and the corresponding days of backwardation (contango) for the four commodities are 119 (133), 102 (150), 91 (161) and 107 (145), respectively; MVol is the average one-step ahead daily volatility forecast, annualised as $N^{-1} \left(\sum \sigma_{t+1} \right) \sqrt{252}$. See also notes in Table 4.3.

Finally, looking at the individual models across the different markets in Table 4.3, the GARCH-X model achieves the lowest RMSE's and it is dominant for the overall sample, with the exception of heating oil where the Mix-GARCH is preferred. In periods of backwardation the Mix-GARCH model achieves lower RMSE's in 2 out of 4 cases. The same holds for GARCH-X in contango. The better MAE is realised by the Mix-GARCH-X model across all markets and under different market conditions, with the exception of heating oil in periods of backwardation); this implies that, compared to the GARCH-X model, Mix-GARCH-X moves closer to the true volatility but some outliers result in a higher RMSE, since RMSE penalises large errors more than MAE. The forecast accuracy results seem to be more in favour of Mix-GARCH-X for balancing over- and under- prediction errors as given by the weighted sum of MME(O) and MME(U), across all commodities and market conditions. The same holds for the MME(O) statistics whereas results for under-prediction are mixed in periods of backwardation and contango, with the GARCH-X model appearing to be better in backwardation and the regime switching models better under contango (3 out of 4 cases), especially the augmented versions (2 out of those 3 cases). In Table 4.4, the overall dominance and consistency of Mix-GARCH-X model is confirmed across all markets, periods and market conditions. The only exception occurs in the heating oil market in 2004-2007 where the simple GARCH model is the best performer. However, in more volatile periods (2008) the GARCH model fails to provide the best forecasts and according to the weighted sum of MME(U) and MME(O) the Mix-GARCH and Mix-GARCH-X models are preferred.

Volatility forecast comparison using different loss functions is simply a historical measurement of how models would have performed in the out-of-sample period under study. Following Diebold and Mariano (1995) several papers have tested the hypothesis of equal predictive ability (see for instance Kang et al. 2009). However, considering only the nominal values of the loss function scores across models, results are prone to data snooping bias. In other words, by relying solely on the mean value of a statistical loss function it is difficult to refute that results would be qualitatively dissimilar in different periods or that they might be coincidental. Sullivan et al. (1999) and White (2000) proposed a new approach to handle such biases by approximating the empirical distribution of a performance measure. We consider the following relative loss differential:

$$fm_{k,t+1} = LF_{t+1}^k - LF_{t+1}^{benchmark} \quad (4.8)$$

where k represents the k^{th} model and LF is the corresponding loss function. The null hypothesis to be tested is $H_0 = \max\{E[fm_k]\} \leq 0$, i.e. there is no model better than the benchmark; a small

p-value indicates that there exists a model which provides superior forecasting results, based on a specific loss function. In the energy economics literature similar procedures have been applied to test volatility forecasts by Wei et al. (2010) and VaR forecasts by Huang, Yu, Fabozzi and Fukushima (2009). We use the stationary bootstrap (see Appendix 4.C) of Politis and Romano (1994) to obtain the average loss differential (Eq. 4.8) of each bootstrapped sample $\bar{f}m_k^*(b)$, based on 3,000 bootstrap simulations. The so called *bootstrap RC p-value* is obtained by comparing the observed statistic $T_n^{RC} = \max_k \left\{ N^{1/2} \left(\bar{f}m_k \right) \right\}$ with the quantiles of the empirical distribution of T_n^{RC*} . The simulated statistic T_n^{RC*} is calculated as:

$$T_n^{RC*} = \max_k \left\{ N^{1/2} \left(\bar{f}m_k^*(b) - \bar{f}m_k \right) \right\} \quad (4.9)$$

The superiority of the Mix-GARCH-X model is evident across all markets, market conditions and periods, especially in terms of MAE and the weighted sum of MME(U) and MME(O), as well as MME(O) alone (Tables 4.3 and 4.4). The results indicate that, overall, no model is significantly superior to the Mix-GARCH-X, at conventional significance levels. Even, in cases where the nominal value of the loss function statistic is not the lowest (see for instance the heating oil market in backwardation and the same market for 2004 -2007 in Tables 4.3 and 4.4, respectively), the performance of Mix-GARCH-X is not significantly different from that of alternative models. Thus, we can conclude that this model is not outperformed by other competing models, at conventional significance levels.

4.4.2 Evaluating the Predictive Performance of Value-at-Risk Forecasts

One of the most popular approaches for quantifying market risk is VaR, the computation of which is pivotal in risk management. VaR is the maximum expected loss in value of an asset or a portfolio of assets over a target horizon, given a specific confidence level $1-c$ ⁶. Then, conditional on the information set at t (Ω_t), VaR can be defined as the solution to the following expression:

⁶ c is typically chosen to be 1% or 5%. The confidence level reflects the degree of risk aversion of an investor since higher c is associated with lower number of violations of the maximum expected loss estimate.

$$\Pr(r_{t+1} \leq VaR_{t+1}^c | \Omega_t) = c \quad (4.10)$$

The fact that different VaR approaches are highly likely to yield significantly different risk estimates is a central concern in VaR inference and thus, the concept of VaR has attracted much attention from researchers leading to a wide spectrum of alternative estimation methods. This stresses the importance of backtesting in order to assess and monitor - on a continuous basis - the accuracy of competing VaR techniques.

Estimating VaR using the Mix- and MRS-GARCH model outlined above further allows for structural changes in the GARCH processes and overcomes some of the limitations that traditional GARCH models exhibit. First, the switching formulation improves on the autoregressive nature of GARCH-based VaR and ensures a better fit of the data as well as a better estimate of market risk by additionally conditioning on the state that the market is in. Second, the high volatility persistence imposed by single regime models decreases and the forecasting performance is expected to be better (see for example Cai, 1994 and Dueker, 1997). Consequently, one expects Mix- and MRS- based VaR to outperform the conventional VaR techniques. For instance, Li and Lin (2004) estimated the VaR of several stock indices using the model of Hamilton and Susmel (1994). Switching ARCH forecasts were found to be more accurate than ARCH and GARCH in terms of violation rate tests (especially at 99% confidence level) and were found to be superior in mitigating the non-normalities of the data. Another study by Marcucci (2005) applied GARCH and MRS-GARCH models to the S&P 100 stock index market in order to assess the predictive accuracy of volatility and VaR forecasts. The findings of the paper supported the superiority of those models at forecast horizons of less than one week.

To evaluate the performance of the estimated volatility models the one step ahead forecast of the VaR estimate at time $t+1$ is calculated using Eq. (4.11). For all calculations we consider interval forecasts with nominal coverage rates of 1% and 5% as well as 99% and 95% in order to account for both long and short positions. Let Φ_{t+1} be the inverse cumulative distribution function and $\mu_{st,t+1}$ and $h_{st,t+1}$ the mean and volatility forecasts in regime st , for $st = \{1,2\}$. Based on Eq. (4.10) the estimated VaR at time $t+1$ given all the available information up to t , at a specified tail probability level $c \in (0, 1)$ can be written as (see also Billio and Pellizon, 2000 and Marcucci, 2005):

$$VaR_{t+1}^c = \pi_{st=1,t+1} \left(\mu_{st=1,t+1} + \Phi_{t+1}(c) (h_{st=1,t+1})^{1/2} \right) + (1 - \pi_{st=1,t+1}) \left(\mu_{st=2,t+1} + \Phi_{t+1}(c) (h_{st=2,t+1})^{1/2} \right) \quad (4.11)$$

We also estimate VaR based on Historical Simulation (HS) as well as Extreme Value Theory (EVT). Regarding EVT, we employ the standard Peaks-Over-Threshold (POT) approach by fitting the GPD to a certain data set that exceeds a pre-set threshold, using maximum likelihood. Consider a threshold, say $u = r_{t+I,N}$ taken from the $(t+I)^{th}$ descending order statistic. Following McNeil and Frey (2000) as a tail we define the exceedances over the 90th percentile. For high thresholds (u), the distribution of the excess losses i.e. $y = r_t - u \geq 0$ can be approximated by the Generalised Pareto Distribution (GPD). In addition, EVT is also adopted to explicitly model the tails of the standardised residual distributions from the various GARCH models. Based on the *conditional EVT* methodology, after fitting the GARCH models to petroleum futures returns, the distribution of the excess negative shocks i.e. $y = \varepsilon_t - u \geq 0$ can also be approximated by the GPD to model the quantile implied by $\Phi_{t+1}(c)$ of Eq. (4.11) and derive VaR. Assuming scale and shape (tail index) parameters ϑ and ξ , respectively⁷:

$$\Phi_{\xi,\vartheta}(y) = \begin{cases} 1 - \left(1 + \xi \frac{y}{\vartheta}\right)^{-\xi^{-1}} & \text{if } \xi \neq 0 \\ 1 - e^{-\left(\frac{y}{\vartheta}\right)} & \text{if } \xi = 0 \end{cases} \quad (4.12)$$

Next, to formally assess the performance of the VaR estimates three tests are constructed: the likelihood ratio tests of unconditional coverage (LR_{UC}), conditional coverage (LR_{CC}) and independence (LR_{IND}), developed by Christoffersen (1998). LR_{UC} ⁸ tests the null hypothesis that the probability of realising a loss in excess of the forecasted VaR is statistically equal to the nominal confidence level c . VaR violations that occur more frequently than c % of the time imply that the VaR method used systematically underestimates the true level of risk, and vice-versa. However, as noted by Christoffersen (1998) and Lopez (1999) the power of this

⁷ $\Phi_u(y) = \Pr(r_t - u \leq y | R_t > u) = \frac{F_u(y+u) - F(u)}{1 - F(u)}$ This can be approximated by the GPD with the tail cumulative

distribution function: $\Phi_u(y) = 1 - (N_u / N)(1 + \xi y / \vartheta)^{-1/\xi}$, where N_u is the number of observations above the threshold. One-step ahead VaR at confidence level $1-c$ is then estimated as: $VaR_{t+1}^c = u + (\vartheta/\xi) \left\{ \left[(N / N_u) c \right]^{-\xi} - 1 \right\}$. Our results indicated that, in most cases, the tail index ξ of unconditional returns and standardised residuals series' is not significantly different from zero, implying that the limiting distributions of the four commodities are of the Gumbel type apart from some exceptions when GPD is applied to the unconditional returns where Fréchet distribution is more appropriate ($\xi > 0$).

⁸The likelihood ratio statistic is based on the assumption of a binomial distribution. Let n be the number of outcomes that fall outside the forecast interval, N the number of forecasts and \hat{c} the empirical level of coverage. Then, the statistic is expressed as: $LR_{UC} = -2 \log \left[\frac{c^n (1-c)^{N-n}}{\hat{c}^n (1-\hat{c})^{N-n}} \right] \sim \chi^2(1)$

test is small in distinguishing between close alternatives, particularly when the returns series are non-normal and/or exhibit volatility clustering. For instance, in periods of low volatility the interval VaR forecasts are expected to be relatively narrower compared to high volatility periods. Thus, tests for conditional coverage are also addressed since a model with correct unconditional coverage may have limited accuracy conditionally and thus, may not be able to capture the clustering of volatility ⁹.

The results of the above tests are presented in Table 4.5. First, we can see that unconditional EVT produces the most conservative VaR estimates in all markets whereas the lowest VaR forecasts are from the Mix-GARCH-X (as well as the Mix-GARCH-X models combined with EVT at the 5% and 95% tails) which implies that this model is the most efficient in terms of allocating capital reserves. At the 1% region for both long and short positions MRS-GARCH, GARCH-X, EVT-Mix-GARCH-X and EVT-MRS-GARCH-X are the models that pass all the tests (LR_{UC} , LR_{IND} and LR_{CC} , at 5% significance level). At the 5% VaR level this occurs for the Mix-GARCH and GARCH-X only whereas the two regime dependent GARCH-X specification fails to pass the tests for the short positions in the Brent crude oil market.

Turning next to the largest unexpected loss (UL), calculated as the average loss in excess of VaR violations, this occurs in the case of HS and EVT, which sometimes is more than double than the one calculated from the other models. For instance at the 99% tail for Brent crude oil, the improvement of the average unexpected loss of the EVT-Mix-GARCH-X model over HS and GARCH models is 200% ($=2.17/0.72-1$) and 23%, respectively. This implies that, conditional that a VaR violation occurs, an investor maintaining a \$1 million position on Brent futures is expected to lose \$21,680 on that day, based on the HS VaR estimates, whereas if the EVT-Mix-GARCH-X or Mix-GARCH-X method is employed this amount reduces to less than \$7,400, which is 3 times less.

⁹The conditional Coverage is a joint test of correct unconditional coverage and independent VaR exceptions against the alternative of a first order Markov process for the failures, thus, $LR_{CC} = LR_{UC} + LR_{IND}$. If n_{ij} for $i, j = \{0,1\}$ denotes the number of i 's followed by j 's in the failure process with $\{0,1\} = \{\text{success, failure}\}$, π_{ij} the probability that i is followed by j and $\hat{\pi} = (n_{01}+n_{11})/(n_{00}+n_{01}+n_{10}+n_{11})$, LR_{UC} and LR_{IND} are, respectively:

$$LR_{CC} = -2 \log \left[\frac{(1-\hat{\pi})^{N-n} \hat{\pi}^n}{(1-\hat{\pi}_{01})^{n_{00}} \hat{\pi}_{01}^{n_{01}} (1-\hat{\pi}_{11})^{n_{10}} \hat{\pi}_{11}^{n_{11}}} \right] \sim \chi^2(2) \quad \text{and}$$

$$LR_{IND} = -2 \log \left[\frac{(1-\hat{\pi})^{(n_{00}+n_{10})} (1-\hat{\pi})^{(n_{01}+n_{11})}}{(1-\hat{\pi}_{01})^{n_{00}} \hat{\pi}_{01}^{n_{01}} (1-\hat{\pi}_{11})^{n_{10}} \hat{\pi}_{11}^{n_{11}}} \right] \sim \chi^2(1)$$

Table 4.5: Value-at-Risk & Risk Management Loss Functions for Long & Short Positions

	Long Positions										Short Positions										Relative	
	1% VaR					5% VaR					99% VaR					95% VaR					Sum of QL	Rank
	\overline{VaR}	PF	UL	AQLF	QL	\overline{VaR}	PF	UL	AQLF	QL	\overline{VaR}	PF	UL	AQLF	QL	\overline{VaR}	PF	UL	AQLF	QL	QL'	Rank
Panel A: NYMEX WTI Crude Oil																						
HS	5.331	1.746**	1.589	8.301	8.127**	3.315	6.667**	1.461	31.32	26.42**	5.193	1.587	1.738	10.160	7.930**	3.291	6.587**	1.410	31.590	25.640**	1.224	10
EVT	7.097	0.556	1.586	2.298	8.000***	4.564	2.778**	1.663	14.40	27.55***	6.345	0.794	1.903	5.770	7.834***	4.422	3.095**	1.443	16.350	26.460***	1.236	11
GARCH	5.489	1.032	0.823	1.818	6.360	3.866	3.254**	1.293	10.34	23.65	5.591	0.556	1.251	1.351	6.264	3.968	3.413**	1.177	9.488	23.750	1.032	9
Mix-GARCH	4.891	1.349	0.990	2.960	6.248	3.444	5.238	1.047	12.79	22.82	4.984	1.270	0.785	1.619	5.959	3.538	6.190	0.937	11.670	23.380	1.002	1
MRS-GARCH	5.034	1.270	1.020	3.273	6.351	3.546	5.238	0.995	12.91	23.05	5.125	1.032	0.956	1.610	6.090	3.637	5.873	0.949	11.530	23.650	1.017	7
GARCH-X	5.177	1.032	1.040	2.109	6.272	3.647	4.127	1.150	11.12	23.09	5.268	0.952	0.881	1.342	6.085	3.738	4.127	1.159	10.300	23.360	1.011	4
Mix-GARCH-X	4.739	1.349	1.147	3.668	6.309	3.336	5.873	1.033	14.24	22.86	4.843	1.349	0.878	1.793	6.006	3.440	6.508**	0.988	12.760	23.520	1.008	3
MRS-GARCH-X	4.942	1.270	1.114	3.504	6.379	3.479	5.635	0.999	13.60	23.13	5.050	1.111	0.889	1.459	6.015	3.586	6.111	0.952	11.360	23.640	1.016	5
EVT-GARCH-X	5.957	0.397**	1.131	0.899	6.428	3.631	4.127	1.162	11.22	23.06	5.507	0.714	0.836	0.937	6.082	3.564	5.635	1.007	12.310	23.390	1.017	6
EVT-Mix-GARCH-X	5.418	0.794	1.063	1.917	6.284	3.335	5.794	1.044	14.11	22.83	5.069	1.190	0.735	1.253	5.921	3.287	7.222**	1.040	15.150	23.840	1.007	2
EVT-MRS-GARCH-X	5.660	0.714	1.085	1.842	6.457	3.477	5.635	0.996	13.47	23.11	5.264	0.952	0.773	0.998	5.977	3.421	7.222**	0.959	13.630	23.920	1.020	8
Panel B: NYMEX Heating Oil #2																						
HS	5.417	0.873	1.666	3.941	6.903**	3.285	5.873	1.296	19.83	24.19**	5.169	1.587	1.459	5.910	7.454*	3.316	7.460**	1.193	23.88	25.33*	1.100	10
EVT	7.141	0.397**	1.303	0.985	7.690***	4.582	2.381**	1.151	7.480	25.81***	6.244	1.032	1.075	2.564	7.322*	4.447	2.857**	1.379	10.54	26.02***	1.154	11
GARCH	4.958	1.111	1.090	2.262	6.200	3.495	5.079	0.950	10.81	22.46	5.033	1.825**	1.047	3.510	6.913	3.569	5.873	1.117	15.34	24.25	1.020	3
Mix-GARCH	5.024	1.032	1.147	2.266	6.240	3.541	4.206	1.072	10.40	22.38	5.101	1.667**	1.035	2.984	6.794	3.618	5.079	1.240	14.31	24.23	1.016	2
MRS-GARCH	5.193	1.032	1.121	2.158	6.381	3.661	3.730**	1.133	9.992	22.69	5.264	1.429	1.086	2.353	6.784	3.732	5.000	1.182	13.01	24.41	1.027	6
GARCH-X	5.177	0.952	0.718	0.913	5.892	3.647	4.127	0.998	7.758	22.51	5.268	1.429	1.167	3.792	6.904	3.738	4.762	1.192	13.57	24.21	1.007	1
Mix-GARCH-X	4.872	1.111	1.272	2.931	6.317	3.430	4.921	1.050	11.99	22.47	4.976	1.825**	1.098	3.694	6.949	3.534	5.952	1.148	15.66	24.35	1.027	7
MRS-GARCH-X	5.040	1.032	1.299	2.614	6.412	3.548	4.603	1.017	11.08	22.58	5.145	1.746**	0.991	2.783	6.844	3.653	5.873	1.068	13.93	24.38	1.029	8
EVT-GARCH-X	5.645	0.714	1.004	1.135	6.394*	3.414	5.079	1.040	12.12	22.51	5.280	1.349	1.125	2.718	6.765	3.465	6.587**	1.085	16.68	24.31	1.024	4
EVT-Mix-GARCH-X	5.497	0.714	1.194	1.526	6.382*	3.383	5.476	0.987	12.44	22.47	5.204	1.508	1.089	2.870	6.814	3.426	6.587**	1.139	17.26	24.47	1.026	5
EVT-MRS-GARCH-X	5.707	0.794	0.968	1.258	6.508*	3.504	4.683	1.041	11.47	22.55	5.361	1.349	1.064	2.102	6.765	3.537	6.429**	1.083	15.52	24.49	1.031	9
Panel C: ICE Brent Crude Oil																						
HS	5.102	1.825**	1.393	6.876	7.681**	3.283	5.714**	1.529	26.73	25.33**	5.201	1.270	2.168	10.56	7.919**	3.272	6.190	1.384	30.37	24.75**	1.259	10
EVT	6.765	0.556	1.490	2.115	7.628***	4.464	2.540**	1.636	11.91	26.66***	6.314	0.635	2.560	5.888	7.904***	4.370	2.302**	1.848	16.66	25.92***	1.284	11
GARCH	5.546	0.635**	0.856	1.148	6.126	3.905	2.778**	1.143	7.849	22.88	5.659	0.635	0.883	0.898	6.184	4.018	2.937**	1.184	8.268	23.39	1.069	9
Mix-GARCH	4.677	1.111	0.912	2.473	5.725	3.292	5.238	0.967	10.96	21.70	4.777	1.429	0.714	1.638	5.761	3.392	5.952	1.023	12.22	22.87	1.013	3
MRS-GARCH	4.851	1.032	0.945	2.451	5.862	3.417	5.079	0.912	10.06	21.90	4.943	1.111	0.790	1.348	5.785	3.508	5.635	0.988	11.07	22.93	1.023	6
GARCH-X	5.177	0.714	0.977	1.183	5.910	3.647	3.889	0.999	7.326	22.30	5.268	1.032	0.768	1.014	6.024	3.738	4.127	1.126	9.838	23.16	1.043	7
Mix-GARCH-X	4.538	1.429**	0.773	2.415	5.677	3.192	5.873	0.930	11.13	21.60	4.653	1.429	0.732	1.496	5.663	3.307	6.429**	1.021	12.21	22.92	1.006	1
MRS-GARCH-X	4.698	1.032	1.048	2.628	5.815	3.305	5.873	0.885	10.82	21.90	4.813	1.508	0.636	1.402	5.737	3.420	6.349**	0.958	11.52	23.00	1.020	4
EVT-GARCH-X	5.808	0.159**	2.015	0.668	6.163	3.597	3.730**	1.016	8.148	21.95	5.640	0.397**	0.974	0.518	5.991	3.639	4.683	1.011	9.678	22.75	1.045	8
EVT-Mix-GARCH-X	5.087	0.635	1.033	1.469	5.779	3.180	5.952	0.926	11.19	21.59	4.980	0.873	0.717	0.914	5.570	3.230	7.063**	1.006	13.46	23.08	1.008	2
EVT-MRS-GARCH-X	5.264	0.635	1.082	1.699	5.986	3.292	5.873	0.895	10.93	21.90	5.156	0.873	0.631	0.828	5.671	3.343	6.508**	1.012	12.64	23.12	1.026	6
Panel D: ICE Gas Oil																						
HS	5.132	1.190	1.018	2.612	6.383***	3.246	6.349**	1.072	15.00	23.23***	5.321	1.349	0.850	1.737	6.428	3.196	6.190	1.171	16.28	23.04*	1.057	10
EVT	6.763	0.317**	1.078	0.546	7.144***	4.462	1.905**	1.216	15.10	24.82***	6.216	0.397**	0.921	0.590	6.543	4.337	2.063**	1.401	5.667	24.38***	1.128	11
GARCH	4.906	1.032	0.921	1.534	5.895	3.453	5.159	0.957	9.117	22.40	5.012	1.032	1.251	2.298	6.264	3.559	4.286	1.070	9.494	22.19	1.011	3
Mix-GARCH	4.681	1.667**	0.758	1.781	5.983	3.295	6.032	0.989	10.72	22.64	4.779	1.270	1.111	2.686	6.151	3.394	5.079	1.056	10.85	22.14	1.012	5
MRS-GARCH	4.758	1.587	0.808	1.880	6.080	3.350	5.873	1.010	10.84	22.88	4.858	1.032	1.319	2.646	6.179	3.449	5.238**	1.002	10.50	22.30	1.022	8
GARCH-X	5.177	0.952	0.773	1.143	5.952	3.647	4.365	0.907	9.665	22.39	5.268	1.032	1.130	2.495	6.395	3.738	4.048	1.008	9.150	22.57	1.023	9
Mix-GARCH-X	4.463	2.302**	0.709	2.434	6.134	3.137	6.429**	1.059	12.86	22.69	4.593	1.508	1.087	3.353	6.193	3.266	5.794	1.041	12.35	22.17	1.021	7
MRS-GARCH-X	4.612	1.905**	0.759	2.086	6.096	3.242	6.190	1.028	11.84	22.77	4.740	1.508	0.946	2.865	6.128	3.371	5.476	1.017	11.06	22.23	1.018	6
EVT-GARCH-X	5.211	0.794	0.745	0.871	5.842	3.293	6.032	0.958	10.53	22.43	5.280	0.873	1.102	1.670	6.203	3.392	5.079	1.035	10.83	22.02	1.004	1
EVT-Mix-GARCH-X	4.898	1.190	0.776	1.243	5.862	3.101	6.746**	1.039	13.22	22.71	4.933	0.952	1.280	2.433	6.119	3.188	6.111	1.063	13.35	22.24	1.007	2
EVT-MRS-GARCH-X	5.042	1.111	0.770	1.041	5.937	3.205	6.270**	1.045	12.16	22.77	5.085	0.794	1.340	2.107	6.109	3.279	6.032	1.012	12.16	22.31	1.012	4

\overline{VaR} is the average forecasted Value-at-Risk for the out-of-sample period; PF presents the percentage of failures (violations) of each model; ** Asterisks in the PF column, indicate that the model fails to pass all the tests of unconditional coverage, independence and conditional coverage (see footnote 8 and 9 for more details); UL is the unexpected loss which is defined as the average loss in excess of the VaR estimate; AQLF is the Average Quadratic Loss Function; QL is the asymmetric Quantile Loss Function of Koenker and Bassett (1978); Regarding the asymmetric Quantile Loss Function (QL) we perform White's (2000) Reality Check using the stationary bootstrap of Politis and Romano (1994); All error statistics are multiplied by 100; ***, ** and * asterisks in the QL column indicate that the QL of the corresponding model is significantly higher than that of the competing models at 1%, 5% and 10%, respectively; The last column (QL') is a relative measure calculated as the sum of the QL functions at all confidence levels normalised according to the minimum QL, as shown in Eq. (4.15); See also the notes in Table 4.1.

Overall, results are mixed concerning which model is the best alternative. Nevertheless, all GARCH models perform better than HS and EVT. Moreover, the GARCH-X model is the only one that passes all the tests, and provides superior performance than the restricted unaugmented version. Regarding the magnitude of the percentage of failures there is not a specific pattern as to which model is superior. Finally, the augmented version of the regime models provide less conservative VaR estimates whereas if the residuals are filtered with EVT, more conservative estimates are obtained (with the exception of 99% tail).

Finally, Figure 4.6 depicts the excess losses of the 5% and 95% VaR from the single and two-regime GARCH-X models. Comparing the regime-switching models, it seems that the Mix-GARCH-X based VaR is smoother. This reflects that for the Mix-GARCH model the averaging between the two regimes is based on constant regime probabilities and, conditional on the fact that the MRS model produces accurate forecasts of the state that the market will be in, indicates that the latter model may capture more efficiently sudden changes in the volatility of the returns. On average, as shown in Table 4.5, the percentage of failures for the MRS models is most of the times lower and closer to the nominal coverage rates than the corresponding Mix-GARCH.




Figure 4.6: WTI Crude Oil 5% VaR Estimates for Long & Short Positions

4.4.2.1 Measuring Forecasting Performance with Risk Management Loss Functions

To provide a more informative insight into the economic benefits from different VaR strategies we also estimate risk management loss functions. In particular, following Lopez

(1999) and Sarma et al. (2003) we calculate the Average Quadratic Loss Function (AQLF), which considers the magnitude of the violations, and penalises more the large failures. If $I_{\{t\}}$ is the indicator function that takes a value of 1 when a return exceeds the VaR level (in absolute terms) i.e. $r_{t+1} < VaR_{t+1}^c$, the loss function becomes:

$$AQLF = \frac{1}{N} \sum_{i=1}^N (r_{t+i} - VaR_{t+i}^c)^2 I_{\{r_{t+i} < VaR_{t+i}^c\}} \quad (4.13)$$

Furthermore, following Koenker and Bassett (1978) we also employ another loss function, the predictive quantile loss (QL) which is based on quantile regression. The QL function penalises more heavily observations for which a violation occurs, and is actually a measure of fit of the predicted tail at a given confidence level. The objective is to minimise QL :

$$QL = \frac{1}{N} \sum_{i=1}^N |r_{t+i} - VaR_{t+i}^c| \left((1-c) I_{\{r_{t+i} < VaR_{t+i}^c\}} + c I_{\{r_{t+i} \geq VaR_{t+i}^c\}} \right) \quad (4.14)$$

The economic intuition behind the use of the QL is that capital charges should also be taken into account, hence, the capital forgone from overpredicting the true VaR should not be neglected. This latter loss function is asymmetric in view of the fact that underprediction and overprediction of VaR estimates have diverse implications. For instance, underprediction of risk might lead to liquidity problems and reoccurring underprediction causes insolvency. On the other hand, overprediction implies higher capital charges which, although are not a cause of bankruptcy, reflect the opportunity cost of keeping a high reserve ratio.

The results of the above loss functions are presented in Table 4.5. First the AQLF seems to be better for the EVT methods (unconditional and conditional) which is expected since this approach produces more conservative VaR estimates, as it has already been mentioned. Specifically, in 11 out of 16 cases the EVT approach is better than the alternatives in providing the lowest underprediction measures over VaR, and in 4 of those cases the EVT-GARCH-X model is better. In 6 cases the unconditional EVT is better but failing to pass the LR_{UC}, LR_{IND} and LR_{CC} tests. Isolating the models that do pass the latter tests, the GARCH-X model is better in 12 out of 16 cases (8 times the GARCH-X and 4 the EVT-GARCH-X), consistent with the MME(U) statistics in Table 4.3.

Table 4.6: Quantile Loss Across Different Market Conditions and Periods

	2004-2008										2004-2007			2008		
	<i>Backwardation Quantile Loss</i>					<i>Contango Quantile Loss</i>					<i>Relative Sum of Quantile Loss (QL')</i>					
	QL _{1%}	QL _{5%}	QL _{99%}	QL _{95%}	QL'	QL _{1%}	QL _{5%}	QL _{99%}	QL _{95%}	QL'	All	Back.	Cont.	All	Back.	Cont.
Panel A: NYMEX WTI Crude Oil																
HS	6.251	22.39	5.520	22.25	1.028	9.677***	29.75***	9.921***	28.44***	1.383	1.043	1.031	1.069	1.717	1.029	1.041
EVT	7.300***	25.33***	6.344***	22.48	1.145	8.579***	29.38***	9.064***	29.75***	1.315	1.217	1.155	1.296	1.334	1.177	1.092
GARCH	6.093	23.14*	5.686	21.48	1.029	6.581	24.07	6.740	25.62	1.045	1.042	1.029	1.068	1.053	1.037	1.054
Mix- GARCH	6.203	22.29	5.324	21.66	1.009	6.285	23.25	6.483	24.80	1.006	1.005	1.014	1.007	1.036	1.025	1.061
MRS - GARCH	6.241	22.17	5.420	21.72	1.014	6.441	23.78	6.643	25.24	1.029	1.017	1.033	1.015	1.060	1.040	1.079
GARCH - X	6.204	22.64	5.578	21.44	1.022	6.328	23.47	6.504	24.94	1.012	1.018	1.021	1.028	1.037	1.062	1.033
Mix-GARCH - X	6.210	22.29	5.486	21.93	1.020	6.391	23.32	6.436	24.83	1.009	1.006	1.021	1.003	1.058	1.017	1.070
MRS-GARCH - X	6.199	22.17	5.446	21.97	1.016	6.527	23.93	6.486	25.02	1.025	1.014	1.028	1.013	1.063	1.041	1.080
EVT-GARCH - X	6.232	22.61	5.673	21.57	1.029	6.590	23.44	6.420	24.88	1.018	1.035	1.027	1.060	1.014	1.044	1.045
EVT-Mix-GARCH - X	6.108	22.28	5.422	22.40	1.018	6.428	23.28	6.334	25.02	1.008	1.017	1.023	1.025	1.025	1.027	1.074
EVT-MRS-GARCH - X	6.241	22.16	5.429	22.38	1.022	6.635	23.89	6.431	25.19	1.029	1.031	1.033	1.044	1.036	1.051	1.083
Panel B: NYMEX Heating Oil #2																
HS	6.488	23.75	7.870	26.61*	1.071	7.356**	24.68	6.999	23.92	1.173	1.057	1.076	1.056	1.280	1.075	1.082
EVT	7.592***	25.84***	7.166	26.00	1.109	7.798***	25.76	7.493	26.04	1.249	1.166	1.097	1.314	1.160	1.069	1.057
GARCH	6.094	22.82	7.650	25.08	1.022	6.316	22.06	6.107	23.35	1.052	1.008	1.023	1.014	1.095	1.070	1.067
Mix- GARCH	6.137	22.62	7.480	25.14	1.016	6.351	22.11	6.044	23.24	1.050	1.007	1.015	1.024	1.082	1.034	1.067
MRS - GARCH	6.243	22.79	7.438	25.22	1.022	6.532*	22.57	6.070	23.54	1.068	1.017	1.023	1.038	1.097	1.037	1.069
GARCH - X	6.199	23.40	7.801	25.32	1.040	5.557	21.54	5.923	22.99	1.000	1.020	1.039	1.021	1.000	1.066	1.062
Mix-GARCH - X	6.066	22.57	7.838	25.40	1.028	6.592	22.37	5.978	1.012	1.061	1.011	1.028	1.012	1.120	1.065	1.089
MRS-GARCH - X	6.177	22.50	7.634	25.40	1.024	6.667*	22.66	5.981	23.27	1.068	1.015	1.029	1.022	1.112	1.040	1.084
EVT-GARCH - X	6.335	22.79	7.377	25.25	1.024	6.458	22.20	6.097	23.28	1.059	1.020	1.021	1.047	1.075	1.075	1.049
EVT-Mix-GARCH - X	6.255	22.57	7.546	25.67	1.028	6.521	22.37	6.015	23.17	1.059	1.019	1.025	1.040	1.087	1.058	1.088
EVT-MRS-GARCH - X	6.394	22.49	7.422	25.65	1.028	6.632*	22.62	6.047	23.22	1.069	1.028	1.030	1.054	1.079	1.050	1.085
Panel C: ICE Brent Crude Oil																
HS	5.651	21.46	5.826	22.00	1.051	9.254**	28.33**	9.541**	26.88*	1.414	1.069	1.054	1.095	1.775	1.100	1.052
EVT	6.979***	24.29***	6.247**	22.58	1.174	8.131**	28.49***	9.187**	28.51***	1.370	1.264	1.177	1.361	1.396	1.231	1.083
GARCH	5.778*	22.35**	5.856	21.38	1.061	6.395	23.29	6.438	24.94	1.079	1.094	1.064	1.134	1.058	1.086	1.045
Mix- GARCH	5.239	21.08	5.522	21.55	1.007	6.102	22.19	5.946	23.89	1.022	1.004	1.011	1.007	1.081	1.033	1.009
MRS - GARCH	5.403	21.12	5.517	21.53	1.015	6.218	22.50	5.993	24.01	1.034	1.017	1.025	1.019	1.085	1.053	1.033
GARCH - X	5.834**	21.98	5.873	21.53	1.062	5.969	22.55	6.141	24.43	1.035	1.068	1.066	1.080	1.031	1.107	1.072
Mix-GARCH - X	5.257	21.09	5.645	21.88	1.018	6.002	21.99	5.678	23.73	1.002	1.005	1.020	1.000	1.054	1.040	1.016
MRS-GARCH - X	5.300	21.09	5.559	21.82	1.015	6.214	22.53	5.874	23.92	1.028	1.010	1.022	1.008	1.090	1.066	1.028
EVT-GARCH - X	5.877**	21.48	5.837	21.26	1.053	6.385	22.31	6.111	23.91	1.043	1.076	1.054	1.107	1.016	1.084	1.053
EVT-Mix-GARCH - X	5.421	21.09	5.439	22.15	1.019	6.055	21.98	5.671	23.80	1.004	1.017	1.023	1.021	1.031	1.061	1.023
EVT-MRS-GARCH - X	5.561	21.09	5.470	22.03	1.026	6.316	22.52	5.828	23.97	1.030	1.030	1.035	1.034	1.063	1.087	1.044
Panel D: ICE Gas Oil																
HS	6.336	23.85	6.597	23.80*	1.049	6.424	22.68	6.279*	22.37	1.081	1.070	1.070	1.092	1.147	1.080	1.155
EVT	7.227***	25.36***	6.465	23.84*	1.097	7.070***	24.35***	6.611**	24.86***	1.172	1.178	1.106	1.282	1.100	1.114	1.061
GARCH	6.155	23.65	6.604	22.41	1.024	5.665	21.30	5.963	21.98	1.013	1.020	1.035	1.023	1.105	1.083	1.074
Mix- GARCH	6.214	23.56	6.558	22.42	1.024	5.778	21.82	5.791	21.88	1.015	1.011	1.025	1.015	1.139	1.049	1.068
MRS - GARCH	6.132	23.53	6.504	22.70	1.021	6.034	22.30	5.893	21.95	1.037	1.019	1.022	1.036	1.155	1.058	1.037
GARCH - X	6.159	23.44	6.652	22.63	1.026	5.769	21.45	6.168	22.53	1.034	1.052	1.045	1.079	1.047	1.043	1.126
Mix-GARCH - X	6.327	23.24	6.667	22.56	1.031	5.963	22.21	5.774	21.81	1.026	1.016	1.037	1.012	1.164	1.042	1.089
MRS-GARCH - X	6.319	23.34	6.434	22.62	1.023	5.900	22.27	5.857	21.89	1.029	1.014	1.027	1.019	1.157	1.038	1.068
EVT-GARCH - X	6.098	23.34	6.407	22.28	1.009	5.616	21.64	6.023	21.79	1.015	1.014	1.020	1.027	1.093	1.089	1.049
EVT-Mix-GARCH - X	6.076	23.28	6.366	22.71	1.011	5.672	22.20	5.901	21.83	1.019	1.006	1.017	1.014	1.133	1.052	1.055
EVT-MRS-GARCH - X	6.187	23.36	6.234	22.78	1.011	5.716	22.25	5.999	21.89	1.026	1.011	1.015	1.026	1.136	1.064	1.044

QL' is the normalised sum of the quantile-loss functions of Eq. (4.15). See also notes in Table 4.5.

Turning next to the more interesting results of the asymmetric Quantile Loss function the results are somewhat different and, as expected, the augmented versions perform better and are more able to capture the asymmetries penalised by the loss function. The GARCH-X model is superior in 7 out of 16 cases, of which 3 are EVT-based. The performance of the Mix-GARCH models is similar, since they outperform the other models in 7 out of 16 cases, of which 4 cases correspond to the Mix-GARCH-X model, of which 3 cases are EVT-based VaR. Overall, the conditional EVT VaR is preferred in 8 cases. In an attempt to construct a unified performance measure, Table 4.5 also presents the relative weighted sum of the QL functions. Because of the fact that at different confidence levels QL is not comparable (it is much higher at lower confidence levels) we propose a relative measure which is constructed by averaging the 4 quantile losses $QL_{1\%}$, $QL_{5\%}$, $QL_{95\%}$ and $QL_{99\%}$, normalising each individual loss by the minimum QL (best performer) at the corresponding confidence level as:

$$QL'_k = \frac{1}{4} \sum_c \frac{\bar{QL}_{c,k}}{\min\{\bar{QL}_{c,1}, \bar{QL}_{c,2}, \dots, \bar{QL}_{c,k}\}}, \text{ for } c = [0.01, 0.05, 0.95, 0.99] \quad (4.15)$$

where k represents the kth model. Clearly, the closer to one, the better the VaR performance of the model. A value of 1 indicates that the model is the best in all cases. The last column of Table 4.5 ranks all the models according to that measure and Mix-GARCH, GARCH-X, Mix-GARCH-X and EVT-GARCH-X perform better in the WTI crude oil, heating oil, Brent crude oil and gas oil, respectively. Even so, the most consistent model seems to be the EVT-Mix-GARCH-X model which is the second best in all markets but the heating oil in which ranks 5th. For instance, although the GARCH-X is superior in the heating oil market, it fails to be among the top three performers in the other commodities and ranks 9th in the ICE gas oil market.

Separate results for periods of backwardation and contango are presented in Table 4.6. The economic implications regarding the magnitude of the QL functions follow the discussion of Tables 4.3 and 4.4. Now, it seems that when the markets are backwardated the Mix-GARCH is the more safe choice, being best for all commodities except for gas oil. The MRS-GARCH is also consistent and accurate being second for all commodities and third for the gas oil market. Looking at the individual QL, out of 16 cases in backwardated markets, 10 times the augmented GARCH versions are selected, of which 8 belong to the two-regime augmented models and 7 fall into the category of the conditional EVT. In contango markets on the other hand, the augmented models are chosen in 12 cases whereas regime models 8 with the Mix-GARCH being the best in the WTI crude oil market and the Mix-GARCH-X in the Brent crude oil

market. As for the petroleum products simple GARCH and GARCH-X appear to be superior in the gas oil and heating oil markets, respectively.

Table 4.6 also presents the QL' function for the two sub-periods i.e. 2004-2007 and 2008. In the relatively lower volatility period of 2004-2007 the Mix-GARCH model is dominant, overall. In contango markets, the Mix-GARCH-X has performed better. In backwardation, again, the two regime models perform better with Mix-GARCH being best in the WTI crude oil and heating oil markets, whereas for the ICE Brent crude and gas oil markets the Mix-GARCH and EVT-MRS-GARCH-X are better, respectively. The more interesting period of 2008, is associated with less consistency, especially under contango. In general, the results indicate that more sophisticated models might perform better under both high and low volatility periods.

Finally, we also apply White's (2000) RC on the VaR-based QL function to test whether the quantile loss functions from each model are significantly different¹⁰. From Tables 4.5 and 4.6, the overall conclusion is that it is more difficult to achieve significance using the quantile loss; this is consistent with other studies in the literature, such as Bao et al. (2006). In our study, only the unconditional HS and EVT methods seem to perform significantly worse than the GARCH methods, irrespective of whether the latter follow a two regime process or whether they are augmented or combined with EVT. Exceptions are the GARCH for the 1% and 5% QL of Brent crude oil in periods of backwardation, as well as all the MRS based 1% QL of the NYMEX heating oil in periods of contango.

4.5 Conclusions

In this chapter we examined the performance of Regime Switching models in forecasting volatility and Value at Risk in the energy markets. Given that the excess kurtosis, skewness and volatility clustering are prominent features of oil price changes, both the Mix-GARCH and MRS-GARCH models are attractive candidates for modelling and forecasting risk. The rationale behind the use of these models stems from the fact that the volatility of these markets may be characterised by regime shifts and by allowing the second moments to be

¹⁰ Note that QL in this case is not differentiable due to the presence of the indicator function. However, according to supporting evidence by Sullivan and White (1998) the stationary bootstrap reality check delivers good approximations to the desired limiting distribution even when an indicator function is used. Instead, a smoothing function can be used as a proxy for QL as in Gonzalez-Rivera et. al (2004). However they find almost identical results of the smooth and nonsmooth version of QL, so we use the original form of the loss function. In addition when parameter estimation is involved, the impact of parameter estimation error is asymptotically negligible when the prediction period grows at a slower rate than the estimation period. Thus, in our study, we choose the out-of-sample period (1,260 observations) to be much smaller than the in-sample period (3,224 observations).

dependent upon the “state of the market”, one may obtain more efficient volatility and VaR estimates and hence, superior forecasting performance compared to the methods which are currently being employed. We apply those models in the NYMEX and ICE petroleum futures markets.

Our results indicate that the Mix-GARCH and MRS-GARCH models are better at capturing the persistence in volatility than the GARCH models, and also tend to perform better in an out-of-sample basis. Energy economists and financial analysts should consider these features in the modelling process of oil price volatility, since the Mix-GARCH-X model for all petroleum futures demonstrate better forecasting accuracy than the other models, in terms of balancing the under- and over-prediction of errors. This holds, irrespective of whether the market experiences backwardation or contango. Regarding the VaR application of the models, the augmented GARCH-X model is the most consistent one passing all the tests for all examined markets. Also, conditional EVT poses as a conservative alternative to VaR forecasts, thus, being more appropriate to risk averse investors. Further investigation, by employing risk management loss functions indicates that in both contango and backwardation periods, inclusion of the squared futures spread when predicting volatility is important. Overall, the magnitude of disequilibria is a factor that does have explanatory power in determining potential changes in oil price volatilities and, in addition, by identifying different volatility components for normal and highly volatile periods, market participants may benefit in terms of accurate quantification of risk as this is reflected in VaR forecasts.

In this chapter we presented the futures prices risk profile and focused on accurate oil volatility modelling. Price risk in the energy complex is likely to occur due to demand and supply changes, refinery capacity constraints, OPEC policy, regional and global economic activity and geopolitical risks, among others. The process of oil price risk measurement, as commonly applied to the measurement of exposures, involves evaluation of the size of exposures conditional on current market conditions; this is traditionally done based on the predictive probability distribution’s summary statistics such as VaR. However, risk managers must also design a framework for mitigating firm's daily exposure to price changes conditional on current market conditions. Next, in Chapter 5: *A Markov Regime Switching Approach for Hedging Petroleum commodities*, we will focus on this aspect of energy risk. When hedging price risk, the optimal proportion of the future contract that should be held to offset the cash position is called the optimal hedge ratio. This ratio is traditionally estimated by examining the ratio between the covariance between cash and futures prices and the variance of the price of futures. The classes of Markov and GARCH models represent two of the most important

techniques to model conditional volatility. In many practical applications it is necessary to adapt the procedures that were discussed above in the context of multivariate conditional heteroscedasticity. In doing so, as it will be seen, we will use the futures and spot long run equilibrium relationship, regime switching models and GARCH. The originality of this study is the challenge of simultaneously modelling the entire variance covariance matrix and linking the concept of disequilibrium (as measured by cointegration) with that of uncertainty (as measured by the conditional second moments) across high and low volatility regimes. Before presenting our empirical evidence, Chapter 5 will first, provide a short literature review on minimum variance hedging. This will be followed by some technical details on multivariate MRS cointegration GARCH models and next, a thorough explanation of the theoretical background and the estimation procedure will be supplied. The model will be successfully fitted to weekly historical futures and spot prices from 1991 to 2006.

APPENDIX 4.A: Alternative Distribution Assumptions

The results under the assumption of the Generalised Error Distribution (GED) are presented in Table 4.A.1. Although the degrees of freedom coefficient was found to be, overall, less than 2 in both regimes, in the low variance state it is not significantly different than 2, indicating that the assumed distribution is normal. Moreover, looking at the estimated models throughout the out-of-sample period, sometimes the low variance state had degrees of freedom of more than 2 (indicating thin tails). Note that, the GED distribution was preferred, over Student-t, because it is more flexible and can accommodate both fat and thin tails as opposed to Student-t which is only able to fluctuate from fat tailed distributions to normal. The fact that the GED may be a more appropriate distribution than Student-t in our empirical analysis is also supported by the evidence in the energy economics literature. There are several studies within the GARCH framework that test different distributional assumptions; for instance Agnolucci (2009), using data for WTI crude futures, reports that in GARCH models the GED assumption is superior to normal and Student-t and four of the five tested models have lower mean squared and absolute errors when assuming GED. Moreover, Hung et al. (2008) and Fan et al. (2008) note that because the Student-t distribution cannot deal simultaneously with fat-tails and leptokurtic distributions, it cannot appropriately capture the empirical distribution of oil prices. In the stock indices market, Marcucci (2005) also reports that MRS-GARCH models with GED outperform the MRS-GARCH under the Student-t distribution. As Haas et al. (2004b) conclude: *“it may often be the case that the specification of heavy tailed distributions in the context of regime switching models is avoidable...”*. The main reason why the corresponding forecasting results for the GED are not presented in this chapter is because, qualitatively, they fail either to add more information or to perform in any way better than the assumed distribution which is also why we present the results of the simpler case.

Table 4.A.1: Estimates of Switching GARCH-X Models for NYMEX & ICE Petroleum Futures Under the Assumption of Generalised Error Distribution

	WTI Crude Oil (CL)		Heating Oil #2 (HO)		Brent Crude Oil (CB)		Gas Oil (GO)	
	Mix- GARCH-X	MRS- GARCH-X	Mix- GARCH-X	MRS- GARCH-X	Mix- GARCH-X	MRS- GARCH-X	Mix- GARCH-X	MRS- GARCH-X
Panel A: Low Volatility Regime								
$E[\sigma_{1t}]$	21.65	20.66	25.84	21.62	19.24	18.86	14.34	18.70
$\mu_{0, st=1}$	0.0320 (0.032)	0.0299 (0.033)	0.0426 (0.032)	0.0431 (0.038)	0.0133 (0.033)	0.0153 (0.031)	-0.0016 (0.058)	0.0048 (0.056)
$\omega_{st=1}$	0.0155 (0.007)**	0.0153 (0.007)**	0.0138 (0.007)**	0.0100 (0.006)**	0.0091 (0.004)**	0.0094 (0.004)**	0.0009 (0.003)	0.0013 (0.003)
$A_{st=1}$	0.0221 (0.006)***	0.0218 (0.006)***	0.0215 (0.005)***	0.0206 (0.005)***	0.0209 (0.005)***	0.0212 (0.005)***	0.0181 (0.004)***	0.0200 (0.005)***
$B_{st=1}$	0.9610 (0.008)***	0.9606 (0.008)***	0.9647 (0.007)***	0.9675 (0.006)***	0.9661 (0.007)***	0.9655 (0.007)***	0.9730 (0.005)***	0.9703 (0.006)***
$\Phi_{st=1}$	0.0051 (0.002)***	0.0051 (0.002)***	0.0018 (0.001)**	0.0014 (0.001)*	0.0055 (0.002)***	0.0057 (0.002)***	0.0001 (0.001)	0.0001 (0.001)
$\pi_{st=1}$	0.8978 (0.072)***	0.8816 -	0.9205 (0.051)***	0.8577 -	0.8833 (0.082)***	0.8936 -	0.6635 (0.094)***	0.6736 -
p_{11}	- -	0.8726 (0.089)***	- -	0.8341 (0.074)***	- -	0.8952 (0.081)***	- -	0.6068 (0.138)***
$g_{st=1}$	1.8470 (0.876)**	1.8464 (0.773)**	2.4154 (1.248)*	1.6805 (0.105)***	1.6706 (0.737)**	1.7771 (0.842)**	2.0068 (0.201)***	1.9722 (0.187)***
Panel B: High Volatility Regime								
$E[\sigma_{2t}]$	47.35	42.82	63.39	38.81	37.39	37.36	21.59	28.53
$\mu_{0, st=2}$	-0.0388 (0.246)	0.0037 (0.223)	-0.2375 (0.249)	-0.1073 (0.213)	0.1077 (0.203)	0.0955 (0.204)	0.1102 (0.141)	0.0946 (0.142)
$\omega_{st=2}$	1.0049 (1.186)	0.8905 (0.778)	2.0527 (2.124)	1.1650 (0.848)	0.5561 (0.641)	0.5909 (0.529)	0.2967 (0.199)	0.4417 (0.241)*
$A_{st=2}$	0.6763 (0.857)	0.6200 (0.529)	0.8074 (1.097)	0.3340 (0.206)*	0.5857 (0.725)	0.5910 (0.610)	0.1319 (0.075)*	0.2086 (0.100)**
$B_{st=2}$	0.6909 (0.229)***	0.6770 (0.204)***	0.6832 (0.243)***	0.6686 (0.210)***	0.6943 (0.208)***	0.6957 (0.225)***	0.7564 (0.096)***	0.7346 (0.101)***
$\Phi_{st=2}$	0.1620 (0.224)	0.1450 (0.162)	0.1692 (0.215)	0.1042 (0.104)	0.1438 (0.207)	0.1392 (0.199)	0.0736 (0.053)	0.1121 (0.067)*
$\pi_{st=2}$	0.1022 (0.072)***	0.1184 -	0.0795 (0.051)***	0.1423 -	0.1167 (0.082)***	0.1064 -	0.3365 (0.094)***	0.3264 -
p_{22}	- -	0.0515 (0.065)	- -	0.0001 (0.094)	- -	0.1202 (0.105)	- -	0.1885 (0.123)
$g_{st=2}$	1.6675 (0.104)***	1.6790 (0.105)***	1.7011 (0.109)***	1.6580 (0.412)***	1.5976 (0.097)***	1.5993 (0.099)***	1.4124 (0.168)***	1.4285 (0.180)***
Panel C: Diagnostics								
LogLik	-6468.7	-6468.3	-6466.0	-6465.1	-6310.5	-6310.5	-6258.5	-6257.5
SBIC	13,042	13,050	13,037	13,043	12,726	12,734	12,622	12,628
$E[\sigma_t]$	25.50	24.364	30.597	24.822	22.152	21.607	17.159	22.411
$A_1 + B_1$	0.983	0.9824	0.9862	0.9881	0.9870	0.9867	0.9911	0.9903
$A_2 + B_2$	1.367	1.2970	1.4906	1.0026	1.2800	1.2867	0.8883	0.9432
Skewness	-0.191***	-0.194***	-0.089	-0.082	0.015	0.012	0.752***	0.743***
Kurtosis	5.830***	5.774***	5.457***	5.346***	6.089***	6.073***	13.33***	12.92***
$Q(10)$	6.320	5.716	4.260	3.957	22.04**	22.05**	9.835	8.459
$Q^2(10)$	16.53*	13.82	15.26	13.75	28.46*	29.79*	3.232	1.719

See Notes in Table 4.2; g_{st} denote the degrees of freedom parameter of the GED distribution.

APPENDIX 4.B: Tests of Two versus Three Regimes

The purpose of this appendix is to present some formal statistical tests and answer the question why a two-regime specification was preferred over a three regime specification. First, a three regime specification would have resulted in a rather large increase in the computational cost of the models, without necessarily a corresponding increase in model fit. Currently, we estimate two-regime processes involving 12 parameters for the MRS-GARCH-X specification: 2 for the mean equation, 8 for the variance (including the X term) and 2 for the transition probabilities. Increasing the number of regimes would result in the estimation of 9 additional parameters (1 mean, 4 variance and 4 transition probabilities parameters). The flexibility of the model would possibly increase, however, at the expense of over-parameterization. What is more, the assumption of a two-regime process is intuitively appealing since the data generating process is disaggregated into periods of high and low volatility. Haas et al. (2004b) estimate three regime models and find some signs of over parameterization in the exchange rate market. Sarno and Valente (2000) essentially show that the third regime only captures the effects of rolling over futures contracts: “Regimes 1 and 2 seem to characterize a very large fraction of the movements of the spot price and the futures price, with each regime being somewhat persistent but with a rather large number of switches over the sample.... Regime 3 is much less persistent and is likely to pick up outliers that do not fall within either Regime 1 or Regime 2”.

Table 4.B.1: Model Selection Criteria

Model	No of param.	WTI Crude Oil		Heating Oil		Brent Crude Oil		Gas oil	
		LogLik	SBIC	LogLik	SBIC	LogLik	SBIC	LogLik	SBIC
Panel A: 2 Regime Models									
Mix-GARCH	9	-6,484.10	13,040.90	-6,481.15	13,035.00	-6,334.32	12,741.34	-6,277.11	12,626.92
MRS-GARCH	10	-6,490.66	13,062.10	-6,480.81	13,042.40	-6,328.62	12,738.03	-6,276.50	12,633.79
Mix-GARCH-X	11	-6,472.99	13,034.83	-6,470.00	13,028.86	-6,316.67	12,722.21	-6,262.57	12,614.00
MRS-GARCH-X	12	-6,471.88	13,040.71	-6,468.59	13,034.12	-6,316.56	12,730.06	-6,261.26	12,619.47
Panel B: 3 Regime Models									
Mix-GARCH	14	-6,473.82	13,060.74	-6,473.04	13,059.17	-6,316.67	12,746.43	-6,269.28	12,651.66
MRS-GARCH	18	-6,463.91	13,073.23	-6,468.37	13,082.15	-6,314.53	12,774.48	-6,258.88	12,663.17
Mix-GARCH-X	17	-6,455.91	13,049.14	-6,451.07	13,039.47	-6,300.65	12,738.64	-6,257.57	12,652.47
MRS-GARCH-X	21	-6,459.82	13,089.28	-6,452.93	13,075.50	-6,301.23	12,772.11	-6,242.33	12,654.30

See notes in Table 4.2; SBIC is calculated as $2\text{LogLik} + N\log(T)$; T denotes the number of (in-sample) observations i.e. 3,224.

Table 4.B.1 presents the log-likelihood function value along with the Bayesian Information Criterion (BIC) for two- and three-regime models. As expected, the Log-Likelihood is improved in the 3 regime case; however the BIC suggests that the 2 regime models are more

parsimonious. We also compare the three regime models in terms of likelihood ratio tests based on Davies (1987) upper bound test. These tests are constructed as $2(\text{LogLik}_{\text{UNCON}} - \text{LogLik}_{\text{CON}})$, where $\text{LogLik}_{\text{UNCON}}$ and $\text{LogLik}_{\text{CON}}$ represent the unconstrained (3 regime) and the constrained (2 regime) maximum likelihood, respectively, and are distributed as $\chi^2(r)$ where r is the number of restrictions imposed. Due to the existence of nuisance parameters, LR statistics are adjusted according to the upper bound of Davies' (1987) test. Under the assumption that the Log-Likelihood function has a single peak i.e. $\Theta = 2M^{1/2}$, where M is the LR statistic, and denoting the gamma function as $\Gamma(\cdot)$ the p-value of the modified LR statistic is given by:

$$\Pr(\chi^2(r) > M) + \Theta M^{(r-1)/2} \exp^{-M/2} \frac{2^{-(r/2)}}{\Gamma(r/2)} \quad (4.B.1)$$

Overall results are mixed and show that at the 1% significance level, in 8 cases the two-regime model is preferred and in the remaining 8 cases the three-regime model is preferred. It is important to note here that there exist several econometric issues related to LR tests. Several studies have used Davies (1987) upper bound to test the number of regimes. However, mainly because the Markov model has a problem of nuisance parameters, it is possible that this test will only be valid if the null model is a linear model. In addition, the fact that innovations are heteroscedastic only complicates matters further. Therefore, Davies (1987) test seems to be rather weak because the distribution of the LR test probably differs from the χ^2 . Due to those issues, we additionally follow McLachlan (1987) and Rydén et al. (1998) and employ a bootstrap methodology, in order to estimate the p-values for the LR tests. As already mentioned, the main problem of performing standard Likelihood ratio tests in the context of regime switching models is the lack of identifiability in the presence of nuisance parameters and the fact that LR statistics do not follow the usual χ^2 distribution. Therefore, we approximate the empirical distribution of the LR statistic using the stationary bootstrap of Politis and Romano (1994) to resample the original data. Given the Maximum Likelihood parameter estimates, we feed the model with the bootstrapped samples and each time we obtain a simulated LR statistic. The number of bootstrap replications is set to 1,000 with a smoothing parameter of $q=0.1$, following Sullivan et al. (1999). Table 4.B.2 presents the maximum values of the simulated LR statistic along with the corresponding p-values constructed as the number of times the simulated statistic exceeds the observed one. The null hypothesis is that the two regime model describes the data better. We can see that it is only in the case of the augmented GARCH models in the

heating oil market that the hypothesis of a two regime model can be rejected. In all the other cases p-values indicate that a two-regime model is preferred at 1% significance level.

Table 4.B.2: Likelihood Ratio Tests – 2 vs. 3 Regimes

	Mix- GARCH	MRS- GARCH	Mix- GARCH-X	MRS- GARCH-X
Panel A: WTI Crude Oil				
LR stat	20.55	53.49	34.16	24.13
Davies LR Upper Bound	{0.020}	{0.000}	{0.000}	{0.077}
Max simulated LR stat	46.10	69.04	44.51	57.13
{p-value}	{0.257}	{0.021}	{0.012}	{0.244}
Panel B: Heating Oil				
LR stat	16.22	24.87	37.87	31.32
Davies LR Upper Bound	{0.091}	{0.033}	{0.000}	{0.007}
Max simulated LR stat	58.73	61.69	42.21	46.48
{p-value}	{0.786}	{0.527}	{0.002}	{0.002}
Panel C: Brent Crude Oil				
LR stat	35.30	28.18	32.04	32.66
Davies LR Upper Bound	{0.000}	{0.010}	{0.000}	{0.010}
Max simulated LR stat	54.94	106.1	63.87	60.86
{p-value}	{0.046}	{0.150}	{0.300}	{0.312}
Panel D: Gas Oil				
LR stat	15.66	35.25	9.992	27.88
Davies LR Upper Bound	{0.110}	{0.001}	{0.968}	{0.001}
Max simulated LR stat	75.84	77.47	33.57	81.74
{p-value}	{0.747}	{0.069}	{0.249}	{0.024}

- LR stat is the Likelihood Ratio statistic calculated as $2(\text{LogLik}_{3\text{REGIME}} - \text{LogLik}_{2\text{REGIME}})$; The first p-value is calculated assuming a χ^2 distribution with degrees of freedom equal to the number of restrictions. Rejection of the null hypothesis implies that restrictions imposed are not valid.
- Davies upper Bound is a correction to the p-value of the LR stat to accommodate that fact that the test might not be valid due to nuisance parameters.
- Max Simulated LR stat is the maximum value after using 1,000 bootstrap simulations of the original data and feeding each model with the new data to obtain the LR statistic.
- P-value is calculated as the number of simulated LR statistics that exceed the observed statistic.

APPENDIX 4.C: The Stationary Bootstrap

Here we present the algorithm that is used to implement the stationary bootstrap resampling technique of Politis and Romano (1994). The description of the algorithm here follows from Appendix C of Sullivan et al. (1999). The stationary bootstrap is calculated as follows: Given the original sample of T observations, $X(t)$, $t = \{1, \dots, T\}$, we start by selecting a “smoothing parameter”, $q = q_T$, $0 < q_T \leq 1$, $Tq_T \rightarrow \infty$ as $T \rightarrow \infty$, and from the bootstrapped series, $X(t)^*$, as follows:

1. At $t = 1$, select $X(1)^*$ at random, independently and uniformly from $\{X(1), \dots, X(T)\}$. Say for instance that $X(1)^*$ is selected to be the J th observation in the original series, $X(1)^* = X(J)$ where $1 \leq J \leq T$.
2. Increment t by 1. If $t > T$, then stop. Otherwise draw a standard uniform random variable U independently of all other random variables
 - (a) if $U < q$, then select $X(2)^*$ at random, independently from $\{X(1), \dots, X(T)\}$
 - (b) if $U > q$, then expand the block by setting $X(2)^* = X(J + 1)$, so that $X(2)^*$ is the next observation in the original series following $X(J)$. If $J + 1 > T$, then reset $J + 1$ to 1, so that the block continues from the final observation in the sample.
3. Repeat step 2 until we reach $X(T)^*$.
4. Repeat steps 1-3, 1000 times

Therefore, the stationary bootstrap re-samples blocks of varying length from the original data, where the block length follows a geometric distribution, with mean block length $1/q$. In general, given that $X(t)^*$ is determined by the J th observation $X(J)$, in the original series, then $X(t + 1)^*$ will be equal to the next observation in the block $X(J + 1)$ with probability $1 - q$ and picked at random from the original observations with probability q . Regarding the choice of q , a large value of q is appropriate for data with little dependence, and a smaller value of q is appropriate for data that exhibit serial dependence. The value of q chosen in our experiments is 0.1, corresponding to a mean block length of 10. This follows other studies in the literature, most notably Sullivan et al. (1999). Furthermore, we also perform sensitivity tests with different values of q , and find that the results presented in this section are not sensitive to the choice of q .

Chapter 5

A Markov Regime Switching Approach for Hedging Petroleum Commodities

5.1 Introduction

Oil price sensitivity to political and economic events, high levels of petroleum price volatility and growing concerns with respect to the security of energy supplies, all stress out the need for the development of a reliable and efficient petroleum price risk management programme. Hedging energy market price exposures involves locking in prices and margins (e.g. refining profit margins) to reduce cash flow uncertainty. Although the risk matrix of energy companies is high dimensional, containing foreign exchange risk, political and country risk, credit risk, regulatory risk, operational risk, force majeure etc., this chapter deals with the most important of all, that is, energy price risk.

Derivative markets allow market agents to reduce their price risk exposure. One parameter which is critical for the development of effective hedging strategies is the hedge ratio which provides the number of futures contracts to buy or sell for each unit of the underlying asset on which the hedger bears risk. Ederington (1979) derives hedge ratios that minimize the variance of the hedged portfolio, based on portfolio theory. Let ΔS_t and ΔF_t represent the price changes in spot and futures prices, respectively. Then, the minimum-variance hedge ratio is the ratio of the unconditional covariance between cash and futures price changes over the variance of futures price changes; this is equivalent to the slope coefficient, γ , in the following regression:

$$\Delta S_t = \mu + \gamma \Delta F_t + \varepsilon_t \quad \varepsilon_t \sim iid(0, \sigma^2) \quad (5.1)$$

The estimated R^2 of Eq. (5.1) represents the hedging effectiveness of the minimum variance hedge. However, the fact that many asset prices follow time-varying distributions suggests that the minimum variance hedge ratio should be time-varying (Kroner and Sultan, 1993) which in turn raises concerns regarding the risk reduction properties of hedge ratios based

on Eq. (5.1). To address this issue, a number of studies apply multivariate GARCH (Generalised Autoregressive Conditional Heteroscedasticity) (Engle & Kroner, 1995) models and derive time-varying hedge ratios directly from the estimated second moments (see for instance, Kroner and Sultan, 1993 and Kavussanos and Nomikos, 2000). The consensus from these studies is that GARCH hedge ratios change as new information arrives and, on average, tend to outperform, in terms of risk reduction, constant hedge ratios derived from Eq. (5.1). However, these gains are market specific and vary across different contracts while, occasionally, the benefits in terms of risk reduction seem to be minimal (Lien and Tse, 2002).

By allowing the volatility to switch stochastically between different processes under different market conditions, one may obtain more robust estimates of the conditional second moments and, as a result, more efficient hedge ratios compared to other methods such as GARCH models or OLS. For instance, as already mentioned in the previous chapter, a common feature of GARCH models is that they tend to impute a high degree of persistence to the conditional volatility which is generated by persistence in volatility regimes -in the presence of structural breaks (see Wilson et al., 1996 for oil futures) - rather than reflecting predictability (Lamoureux and Lastrapes, 1990; Fong and See, 2002, 2003). Alizadeh and Nomikos (2004b) examined the hedging effectiveness of FTSE-100 and S&P 500 stock index futures contracts, using MRS models for the estimation of dynamic hedge ratios. Allowing Eq. (5.1) to switch between two state processes, they provided evidence in favour of those models in terms of variance reduction and increase in utility both in- and out-of-sample. Similarly, Lee and Yoder (2007a) extend the univariate MRS-GARCH model of Gray (1996), to a state dependent multivariate GARCH model. They apply their model to the corn and nickel futures markets and they report higher, yet insignificant, variance reduction compared to OLS and single-regime GARCH hedging strategies. Similar results are obtained from the Lee and Yoder (2007b) MRS model of time varying correlation (MR-TVC-GARCH) as applied to the Nikkei 225 and Hang Seng index futures.

This chapter investigates the hedging effectiveness of the MRS models for the WTI Crude Oil, Unleaded Gasoline and Heating Oil futures contracts traded on NYMEX. In doing so, it contributes to the existing literature in a number of ways. First, we extend the univariate MRS model in the hedging literature by introducing, for the first time, a Regime Switching Vector Error Correction Model (VECM) with GARCH error structure, which includes in the mean equation the cointegrating relationship between spot and futures prices. Empirical evidence suggests that if spot and futures prices are cointegrated, omitting the equilibrium

relationship will lead to misspecification problems by underestimating the true optimal hedge ratio (see for instance Kroner and Sultan, 1993, Ghosh, 1993 and Lien, 1996).

Sarno and Valente (2000) provide a further dimension to the literature using a multivariate extension of the Markov Regime Switching (MRS) model proposed by Hamilton (1989) and Krolzig (1999). They find that the relationship between spot and futures is regime dependent and MRS models can explain this relationship better than simple linear models. Some preliminary evidence on the relationship between volatility and the long-run equilibrium of futures, as represented by the basis, was shown in the previous chapter; we saw that in the high volatility state, volatility movements occur mainly due to short-lived random shocks which are difficult to foresee whereas, in the low variance state, the dynamics of the volatilities are more predictable and deviations from the equilibrium appeared to have a certain degree of explanatory power on volatility. The inclusion of the error correction mechanism in the present chapter will enable us to examine whether the speed of adjustment of spot and futures prices to the long-run relationship changes across different regimes. The motivation for this stems from the fact that since the relationship between spot and futures prices changes over time, the adjustment to the equilibrium process should also be time-dependent. This in turn introduces an informative link between volatility and cointegration allowing for both time dependency and asymmetric behaviour across different states in the market. This chapter therefore is different from the Lee and Yoder (2007a) Switching BEKK study in the sense that our model also allows for switching in the error correction coefficients.

In addition, we evaluate the hedging effectiveness of the proposed model using both in- and out-of-sample tests. The performance of the MRS hedge ratios is compared to that of alternative hedge ratios generated from a variety of models that have been proposed in the literature and is assessed in terms of variance reduction, increase in utility and reduction in the value-at-risk for a given position. This way we provide robust evidence on the performance of the proposed hedging strategy. Finally, in addition to providing evidence on the statistical significance of the hedging performance from the competing models using White's (2000) Reality Check, we also address the issue of downside risk by examining whether the effects of mean-variance hedge ratios differ between long and short hedges.

The structure of this chapter is as follows. Section 5.2 presents the minimum-variance hedge ratio methodology and demonstrates the MRS-BEKK model estimation procedure. In section 5.3, the data and their properties are described. Section 5.4 discusses the empirical results. This is followed by an evaluation of the hedging effectiveness of the proposed strategies

in section 5.5; Section 5.6 describes the reality check for data snooping bias. Section 5.7 provides a note on downside risk and finally, conclusions are given in the last section.

5.2 Markov Regime Switching GARCH Models & Hedging

Market participants in futures markets choose a hedging strategy that reflects their individual goals and attitudes towards risk. The degree of hedging effectiveness in futures markets depends on the relative variation of spot and futures price changes as well as the hedge ratio. The hedge ratio that minimises the variance of the hedge portfolio is derived as the slope coefficient of spot price changes on futures price changes, as in Eq. (5.1). This can also be expressed as:

$$\gamma = \frac{\text{Cov}(\Delta S_t, \Delta F_t)}{\text{Var}(\Delta F_t)} \quad (5.2)$$

Therefore, the minimum variance hedge ratio of Eq. (5.2) is the ratio of the unconditional covariance between cash and futures price changes over the variance of the futures price changes.¹ Eq. (5.2) can also be extended to accommodate the *conditional* minimum-variance hedge ratio, $\gamma_{1,t}$, which is the time varying equivalent of the *conventional* hedge ratio γ_1 , in Eq. (5.1). This is believed to be more efficient in reducing the risk of a hedged position, because it is updated as it responds to the arrival of new information in the market. To estimate this dynamic hedge ratio, we employ an MRS VECM for the conditional means of spot and futures returns with a multivariate GARCH error structure. The conditional means of spot and futures returns are specified as:

$$\Delta \mathbf{X}_t = \sum_{i=1}^{p-1} \Gamma_{st,i} \Delta \mathbf{X}_{t-1} + \Pi_{st} \mathbf{X}_{t-1} + \boldsymbol{\varepsilon}_{st,t} \quad ; \quad \boldsymbol{\varepsilon}_{t,st} = \begin{pmatrix} \varepsilon_{S,s,t} \\ \varepsilon_{F,s,t} \end{pmatrix} \Big| \Omega_{t-1} \sim IN(0, \mathbf{H}_{st,t}) \quad (5.3)$$

¹ It can be shown that if expected futures returns are zero, i.e. if futures follow a martingale process $E_t(F_{t+1}) = F_t$ then, the minimum variance hedge ratio of Eq. (5.2) is equivalent to the utility-maximizing hedge ratio. A proof of this result is available at Kroner and Sultan (1993). The martingale assumption of futures returns implies that the expected returns from the hedged portfolio are unaffected by the number of futures contracts held, so that risk minimization becomes equivalent to utility maximization. The assumption of zero expected returns is also in line with the descriptive statistics presented in Table 5.1, which show that the unconditional futures returns have a mean of zero.

where $\mathbf{X}_t = (\Delta S_t \quad \Delta F_t)^\top$ is the vector of spot and futures returns, $\Gamma_{i,st}$ and Π_{st} are 2x2 state dependent coefficient matrices measuring the short- and long-run adjustment of the system to changes in \mathbf{X}_t , respectively, and $\boldsymbol{\varepsilon}_{st,t} = (\varepsilon_{S,s_t,t} \quad \varepsilon_{F,s_t,t})^\top$ is a vector of Gaussian white noise processes with time varying state dependent covariance matrix $\mathbf{H}_{st,t}$.

The following steps are involved in our analysis. First, assuming a single regime process, the existence of a stationary relationship between spot and futures prices, is investigated through the λ_{\max} and λ_{trace} statistics (Johansen, 1988) which test for the rank of Π . The rank of Π in turn determines the number of cointegrating relationships. In particular, if Π has a reduced rank, that is $\text{rank}(\Pi) = 1$, then there exists one cointegrating vector and the coefficient matrix Π can be decomposed as $\Pi = \boldsymbol{\alpha}\boldsymbol{\beta}'$, where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}'$ are 2x1 vectors. Using this factorisation $\boldsymbol{\beta}'$ represents the vector of cointegrating parameters and $\boldsymbol{\alpha}$ is the vector of error correction coefficients measuring the speed of convergence to the long-run steady state. The significance of incorporating the cointegrating relationship into the statistical modelling of spot and futures prices is emphasised in studies such as Kroner and Sultan (1993), Ghosh (1993), Chou et al. (1996) and Lien (1996); hedge ratios and measures of hedging performance may change sharply when this relationship is unduly ignored from the model specification.

The second step involves the introduction of Markovian regime shifts to the system. In order to reduce the computational burden, regime switching is allowed only through the error correction coefficients i.e. $\Pi_{st} = \boldsymbol{\alpha}_{s_t}\boldsymbol{\beta}'$. The unobserved state variable $s_t = \{1, 2\}$ follows a two-state, first order Markov process with the following transition probability matrix:

$$\hat{\mathbf{P}} = \begin{pmatrix} \Pr(s_t = 1 | s_{t-1} = 1) = p_{11} & \Pr(s_t = 1 | s_{t-1} = 2) = p_{21} \\ \Pr(s_t = 2 | s_{t-1} = 1) = p_{12} & \Pr(s_t = 2 | s_{t-1} = 2) = p_{22} \end{pmatrix} = \begin{pmatrix} 1 - p_{12} & p_{21} \\ p_{12} & 1 - p_{21} \end{pmatrix} \quad (5.4)$$

where p_{12} gives the probability that state 1 will be followed by state 2, p_{22} gives the probability that there will be no change in the state of the market in the following period etc. These transition probabilities are assumed to remain constant between successive periods.

Moreover, the conditional second moments of spot and futures returns are specified as a GARCH(1,1) model (Bollerslev, 1986). However, in the regime-switching framework, the GARCH model in its basic form would be intractable because both the conditional variance and the conditional covariance would be a function of all past information. Hamilton and Susmel (1994) and Cai (1994) solve the path dependency problem by eliminating the GARCH term. The main drawback of their model is that many lags of ARCH terms are needed in order to

capture the volatility dynamics. Gray (1996) suggests a possible formulation for the conditional variance process by using the conditional expectation of the variance. Lee and Yoder (2007a) extend Gray's model to the bivariate case and fully solve the path dependency problem by developing a similar collapsing procedure for the covariance. Following the augmented Baba et al. (1987) (henceforth BEKK) representation (see Engle and Kroner, 1995), the GARCH-like formulation of the variance/covariance matrix is:

$$\mathbf{H}_{st,t} = \boldsymbol{\omega}_{st}^T \boldsymbol{\omega}_{st} + \mathbf{A}_{st}^T \boldsymbol{\varepsilon}_{t-1} \boldsymbol{\varepsilon}_{t-1}^T \mathbf{A}_{st} + \mathbf{B}_{st}^T \mathbf{H}_{t-1} \mathbf{B}_{st} \quad (5.5)$$

for $s_t = \{1, 2\}$, where, $\boldsymbol{\omega}_{st}$ is a 2x2 lower triangular matrix of state dependent coefficients, \mathbf{A}_{st} and \mathbf{B}_{st} are 2x2 state dependent coefficient matrices restricted to be diagonal². In this formulation, the state dependent conditional variances are a function of the lagged values of both the lagged aggregated variances and aggregated error terms (after integrating the unobserved state variable) and $\mathbf{H}_{t,st}$ is positive definite for all t . In order to integrate the state dependent variances and residuals we use Gray's (1996) integrating method as adopted by Lee and Yoder (2007a). For instance, collapsing the variance and residuals of spot returns can be expressed as:

$$h_{ss,t} = \pi_{1,t}(\mu_{s,1,t}^2 + h_{ss,1,t}) + (1 - \pi_{1,t})(\mu_{s,2,t}^2 + h_{ss,2,t}) - [\pi_{1,t}\mu_{s,1,t} + (1 - \pi_{1,t})\mu_{s,2,t}]^2 \quad (5.6)$$

$$\varepsilon_{s,t} = \Delta S_t - [\pi_{1,t}\mu_{s,1,t} + (1 - \pi_{1,t})\mu_{s,2,t}] \quad (5.7)$$

where $h_{ss,t}$ is the aggregate spot returns' variance which is an element of the state independent variance matrix \mathbf{H}_t and $h_{ss,st,t}$ is the state dependent spot returns' variance for $s_t = \{1, 2\}$, an element of the state dependent variance matrix $\mathbf{H}_{t,st}$. $\mu_{s,st,t}$ is the state dependent mean equation of spot price changes, and $\pi_{st,t}$ the conditional regime probability that the process will be in a given state at a point in time.

Similar to the variance, the state dependent conditional covariance is a function of lagged aggregated covariance and lagged cross products of the aggregated error terms. Denoting $h_{sf,t}$ the aggregate and $h_{sf,st,t}$ the state dependent covariance, the unobserved state variable is integrated out as follows:

² Coefficient matrices \mathbf{A} and \mathbf{B} are restricted to be diagonal for a more parsimonious representation of the conditional variance (see Bollerslev et al. 1994). For a discussion of the properties of this model and alternative representations see also Engle and Kroner (1995).

$$h_{sf,t} = \pi_{1,t} [\mu_{s,1,t} \mu_{f,1,t} + h_{sf,1,t}] + (1 - \pi_{1,t}) [\mu_{s,2,t} \mu_{f,2,t} + h_{sf,2,t}] - [\pi_{1,t} \mu_{s,1,t} + (1 - \pi_{1,t}) \mu_{s,2,t}] [\pi_{1,t} \mu_{f,1,t} + (1 - \pi_{1,t}) \mu_{f,2,t}] \quad (5.8)$$

Under the specifications of Eq. (5.5), (5.6), (5.7) and (5.8) the MRS-BEKK model becomes path-independent because the variance/covariance matrix depends on the current regime alone and not on its entire history. Consequently, the Markov property for a first order Markov process is not violated and we can allow for a GARCH error structure. Finally, assuming that the state dependent residuals follow a multivariate normal distribution with mean zero and time varying state dependent covariance matrix $\mathbf{H}_{t,st}$, the likelihood function for the entire sample is formed as a mixture of the probability distribution of the state variables as:

$$f(\mathbf{X}_t; \boldsymbol{\theta}) = \frac{\pi_{1,t}}{2\pi} |\mathbf{H}_{1,t}|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \boldsymbol{\varepsilon}_{1,t}^T \mathbf{H}_{1,t}^{-1} \boldsymbol{\varepsilon}_{1,t}\right) + \frac{\pi_{2,t}}{2\pi} |\mathbf{H}_{2,t}|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \boldsymbol{\varepsilon}_{2,t}^T \mathbf{H}_{2,t}^{-1} \boldsymbol{\varepsilon}_{2,t}\right) \quad (5.9)$$

$$L(\boldsymbol{\theta}) = \sum_{t=1}^T \log f(\mathbf{X}_t; \boldsymbol{\theta}) \quad (5.10)$$

where $\boldsymbol{\theta}$ is the vector of parameters to be estimated and $\boldsymbol{\varepsilon}_{t,st}$ and $\mathbf{H}_{t,st}$ are defined in Eq. (5.3) and (5.5), respectively. $L(\boldsymbol{\theta})$ can then be maximised using numerical optimization methods, subject to the constraints that $\pi_{1,t} + \pi_{2,t} = 1$ and $0 \leq \pi_{1,t}, \pi_{2,t} \leq 1$.

Using the MRS specifications outlined above, the second moments of spot and futures returns are conditioned on the information set available at time $t-1$. Based on Eq. (5.2) the estimated hedge ratio at time t given all the available information up to $t-1$ can be written as:

$$\gamma_t^* \Big| \Omega_{t-1} = \frac{h_{sf,t-1}}{h_{ff,t-1}} \quad (5.11)$$

where $h_{sf,t-1}$ and $h_{ff,t-1}$ are calculated from the collapsing procedure as presented in Eq. (5.8) and (5.6), respectively.

Estimating the optimal hedge ratio using the MRS-BEKK model outlined above further allows for structural changes in the GARCH processes and overcomes some of the limitations that traditional GARCH models exhibit. First, by allowing the volatility equation to switch across different states, we relax the assumption of constant parameters throughout the estimation period thus improving the ‘fit’ of our model to the data. Second, the Markovian formulation improves on the autoregressive nature of GARCH-based hedge ratios and ensures a better estimate of the optimal hedge ratio by additionally conditioning on the state that the

market is in. Finally, by accounting for regime switching, the high volatility persistence imposed by single regime models decreases and the forecasting performance is expected to be better (see for example Lamoureux and Lastrapes, 1990; Cai, 1994 and Dueker, 1997). Consequently, one expects MRS hedge ratios estimated by the variance/covariance matrix to outperform the conventional hedging strategies.

5.3 Description of the Data & Preliminary Analysis

The data set for this study comprises weekly spot and futures prices for three energy commodities traded on NYMEX: WTI crude oil, Unleaded Gasoline and Heating oil, covering the period January 23, 1991 to December 27, 2006, resulting 832 weekly observations. Spot and futures prices are Wednesday prices; when a holiday occurs on Wednesday, Tuesday's observation is used in its place. The above dataset was obtained from Datastream and the Energy Information Administration (US Department of Energy) along with volume and open interest data. Data for the period January 23, 1991 to June 15, 2005 (752 observations) are used for the in-sample analysis; out-of-sample analysis is carried out using the remaining data for the period June 22, 2005 to December 27, 2006 (80 observations). In order to deal with thin trading and expiration effects, it is assumed that the hedger will switch contracts the next business day after trading activity has shifted from the nearest to the second nearest to maturity contract. Consequently, in all cases the nearest contract available is chosen as the appropriate hedging contract, and rolling over to the front month contract occurs the business day following the day that both trading volume and open interest exceed that of the nearest to expiry contract³.

Having constructed a continuous time series for the futures contracts prices, spot and futures prices are then transformed into natural logarithms. Summary statistics of the levels and return series are presented in Table 5.1, Panel A. Jarque-Bera (1980) tests indicate significant departures from normality for all the commodities and for both spot and futures returns. The Ljung-Box (1978) Q statistic on the first six lags of the sample autocorrelation function is significant for all spot/futures prices and spot returns revealing that serial correlation is present. Engle's (1982) ARCH test, carried out as the Ljung-Box Q statistic on the squared series, indicates the existence of heteroscedasticity for all the return series, with the exception of WTI futures. Finally, Phillips and Perron (1988) unit root tests on the levels and first differences indicate that spot and futures prices are first difference stationary.

³ For instance the November 2002 WTI futures contract expires on October, 22. The rollover to the December 2002 contract takes place on October 15 because open interest crossover between the two nearby contracts occurred on October 10 while volume crossover on October, 14.

Table 5.1: Summary Statistics, Unit Root & Cointegration Tests for Spot and Futures Prices of WTI Crude Oil, Unleaded Gasoline and Heating Oil # 2

Panel A: Descriptive Statistics												
	WTI light sweet crude oil				Unleaded Gasoline				Heating Oil # 2			
	Log Levels		% Returns		Log Levels		% Returns		Log Levels		% Returns	
	Spot	Futures	Spot	Futures	Spot	Futures	Spot	Futures	Spot	Futures	Spot	Futures
Mean (Weekly)	3.1329	3.1313	0.0723	0.0807	-0.4211	-0.3894	0.0762	0.0862	-0.4606	-0.4628	0.0699	0.0759
Vol (Weekly)	0.3220	0.3209	4.7890	4.3665	0.3180	0.3127	5.7356	4.8695	0.3284	0.3216	5.2694	4.4268
Skew	0.5997	0.6315	-0.6110	-0.9638	0.4409	0.5948	-0.1889	-0.4207	0.6589	0.7173	0.2695	-0.4255
Kurt	0.1872	0.2357	2.8064	5.2401	0.0371	0.0805	1.0019	2.2969	0.4183	0.5468	8.5998	2.5458
J-B	46.169	51.727	293.57	976.79	24.405	44.551	35.924	189.48	59.905	73.863	2326.4	225.77
Q(6)	4267.4	4289.6	15.684	5.5913	4084.5	4183.3	23.411	10.653	4167.6	4264.5	23.069	8.6217
Q ² (6)	4276.8	4297.5	23.921	7.7469	3748.9	3889.1	24.902	14.733	4017.9	4070.2	110.11	45.342
PP	-1.0591	-0.8972	-29.405	-28.550	-1.3475	-1.0522	-31.230	-28.697	-1.3995	-0.9490	-26.117	-27.353

Panel B: Cointegration Tests									
	Lags	H ₀ :	Statistic		Normalized CV (1 β ₁ β ₀)	LR test H ₀ : β ₁ = 1	Restricted CV		
			λ _{max} test	λ _{trace} test					
WTI light sweet crude oil	1	r=0	419.73	420.90	(1 -1.002 0.543)	1.2190	(1 -1 0.162)		
		r=1	1.1697	1.1697		[0.270]			
Unleaded Gasoline	1	r=0	71.212	73.339	(1 -1.004 2.982)	0.0711	(1 -1 3.173)		
		r=1	2.1270	2.1270		[0.790]			
Heating oil #2	1	r=0	101.44	102.55	(1 -1.014 -0.904)	1.2422	(1 -1 -0.237)		
		r=1	1.1091	1.1091		[0.265]			

- Sample period is from January 23, 1991 to June 15, 2005 (752 weekly observations).
- J-B is the Bera and Jarque (1980) test for Normality. The test follows a χ^2 distribution with 2 degrees of freedom.
- Q(6) and Q²(6) are Ljung-Box (1976) tests for 6th order autocorrelation in the level and squared series, respectively. The statistics are $\chi^2(6)$ distributed.
- PP is the Phillips and Perron (1988) unit root test. 1%, 5% and 10% critical values for this test are -3.4388, -2.8652 and -2.5689, respectively.
- Lags is the lag length of the unrestricted VAR model in levels. A VAR with p lags of the dependent variable can be reparameterized in a VECM with p-1 lags of first differences of the dependent variable plus the error-correction term. Lag length is based on Schwarz (1978) Information Criterion and the autocorrelation function of the estimated residuals from the VECM model.
- λ_{max} tests the null hypothesis of r cointegrating vectors against the alternative of r+1. The 5% critical values for H₀: r=0 and H₀: r=1 are 15.67 and 9.24, respectively. λ_{trace} tests the null hypothesis that there are at most r cointegrating vectors against the alternative that the number of cointegrating vectors is greater than r. The 5% critical values for H₀: r=0 and H₀: r=1 are 19.96 and 9.24, respectively. Critical values obtained from Osterwald-Lenum (1992).
- The LR tests the hypothesis that the cointegrating vector (β₂ β₁ β₀) is (1 -1 β₀). The statistic is $-T [\ln(1 - \hat{\lambda}_1^*) - \ln(1 - \hat{\lambda}_1)]$ where $\hat{\lambda}_1^*$ and $\hat{\lambda}_1$ denote the largest eigenvalues of the restricted and the unrestricted models, respectively. The statistic follows a $\chi^2(1)$ distribution. Figures in [] represent the corresponding p-values.

Next, cointegration techniques are used to investigate the existence of a long run relationship between spot and futures price series. Johansen (1988) cointegration tests, presented in Table 5.1, Panel B, indicate that all physical commodity prices stand in a long-run relationship with the corresponding futures contracts. The normalized coefficient estimates of the cointegrating vector $\beta' = (\beta_2 \ \beta_1 \ \beta_0)$ represent this long-run relationship between the series. Furthermore the results of likelihood ratio tests on the hypothesis that there is a one-to-one relationship between spot and futures prices, that is that the cointegrating vector is the lagged basis: $H_0: \beta' = (1, -1, \beta_0)$, show that the null hypothesis cannot be rejected at conventional significant levels. Therefore, we use the restricted cointegrating vectors presented in Table 5.1, in the joint estimation of the conditional mean and the conditional variance.

5.4 Empirical Results

MRS models are estimated assuming two states (see also Chapter 4, Appendix 4.B). The choice of a two-state process is motivated by the fact that this model captures the dynamics of the spot and futures returns in a more efficient way and is intuitively appealing since these two states can be associated with periods of low and high volatility. On the other hand, Sarno and Valente (2000) use a three-state process to model spot-futures relationship in stock indices; nonetheless, in their study the third state seems to capture only jumps in the futures prices at the time of switching between contracts of different maturities and does not reflect fundamental changes in market conditions. Table 5.2 presents the single and two regime GARCH models.

Several observations merit attention. First, looking at the estimated MRS-BEKK models in Table 5.2, in the low variance regime ($s_t=1$), the speed of adjustment of spot and futures prices to their long-run relationship, measured by the $\alpha_{S,st=1}$ and $\alpha_{F,st=1}$ estimated coefficients respectively, are all negative. In the spot equation they are consistently negative and significant whereas in the futures equation they are either insignificant (Unleaded Gasoline) or of less magnitude than the coefficients of the spot equation (WTI and Heating oil). This means that in the low variance regime the estimated error correction coefficients are in accordance with convergence towards a long-run equilibrium relationship; that is, in response to a positive deviation at period $t-1$ (i.e. $S_{t-1} > F_{t-1}$), the spot price in the next period will decrease while the futures price will either be unresponsive or less responsive than spot prices thus restoring the long-run equilibrium. This can be attributed to the fact that petroleum spot prices are usually more sensitive to news since new information is automatically absorbed in the cash markets whereas in the futures markets the speed of adjustment to the available information is a function of several factors like maturity and liquidity.

Table 5.2: Estimates of Markov Regime Switching BEKK Hedge Ratios for NYMEX Energy Commodities

$$\Delta S_t = a_{S,st}(S_{t-1} - \beta_2 F_{t-1} - \beta_0) + \varepsilon_{S,st,t} \quad ; \quad \Delta F_t = a_{F,st}(S_{t-1} - \beta_2 F_{t-1} - \beta_0) + \varepsilon_{F,st,t} \quad ; \quad \varepsilon_t = \begin{pmatrix} \varepsilon_{S,t} \\ \varepsilon_{F,t} \end{pmatrix} \Big| \Omega_{t-1} \sim IN(0, \mathbf{H}_t)$$

$$\mathbf{H}_t = \begin{pmatrix} h_{SS,st,t} & h_{SF,st,t} \\ h_{SF,st,t} & h_{FF,st,t} \end{pmatrix} = \begin{pmatrix} \omega_{11,st} & \omega_{12,st} \\ 0 & \omega_{22,st} \end{pmatrix}^T \begin{pmatrix} \omega_{11,st} & \omega_{12,st} \\ 0 & \omega_{22,st} \end{pmatrix} + \begin{pmatrix} A_{11,st} & 0 \\ 0 & A_{22,st} \end{pmatrix}^T \varepsilon_{t-1} \varepsilon_{t-1}^T \begin{pmatrix} A_{11,st} & 0 \\ 0 & A_{22,st} \end{pmatrix} + \begin{pmatrix} B_{11,st} & 0 \\ 0 & B_{22,st} \end{pmatrix}^T \mathbf{H}_{t-1} \begin{pmatrix} B_{11,st} & 0 \\ 0 & B_{22,st} \end{pmatrix}$$

	West Texas Intermediate				Unleaded Gasoline				Heating Oil # 2			
	GARCH (BEKK)		MRS-BEKK		GARCH (BEKK)		MRS-BEKK		GARCH (BEKK)		MRS-BEKK	
Mean Equation												
$\alpha_{S,st=1}$	-0.8715	(0.144)***	-1.2146	(0.014)***	-0.2321	(0.034)***	-0.1765	(0.075)**	-0.3269	(0.111)***	-0.1958	(0.021)***
$\alpha_{F,st=1}$	0.1151	(0.132)	-0.2259	(0.004)***	-0.0760	(0.030)**	-0.0658	(0.051)	-0.1538	(0.084)*	-0.0359	(0.010)***
$\alpha_{S,st=2}$			-0.1671	(0.282)			-0.2923	(0.146)**			-0.0337	(0.181)
$\alpha_{F,st=2}$			0.2588	(0.300)			-0.0828	(0.102)			-0.4658	(0.127)***
Variance Equation												
$\omega_{11,st=1}$	1.8627	(0.271)***	3.8193	(0.145)***	2.7691	(0.500)***	3.3327	(0.140)***	2.5436	(0.372)***	2.2771	(0.137)***
$\omega_{12,st=1}$	2.2698	(0.270)***	3.8381	(0.145)***	1.5910	(0.227)***	3.1188	(0.135)***	2.4436	(0.326)***	2.0762	(0.144)***
$\omega_{22,st=1}$	2.7x10 ⁻⁶	(0.062)	-0.3307	(0.035)***	0.5541	(0.481)	3.0x10 ⁻⁶	(0.043)	1.6x10 ⁻⁶	(0.053)	0.3472	(0.006)***
$A_{11,st=1}$	0.3084	(0.155)**	-0.0938	(0.006)***	0.3073	(0.039)***	0.1916	(0.125)	0.3636	(0.047)***	0.2686	(0.015)***
$A_{22,st=1}$	0.2453	(0.260)	-0.0802	(0.005)***	0.2695	(0.045)***	0.1714	(0.100)*	0.1626	(0.039)*	0.2251	(0.007)***
$B_{11,st=1}$	0.8686	(0.042)***	1.0x10 ⁻⁵	(0.017)	0.8137	(0.059)***	0.4477	(0.047)***	0.8010	(0.060)***	0.7549	(0.034)***
$B_{22,st=1}$	0.8210	(0.077)***	3.2x10 ⁻⁶	(0.011)	0.8976	(0.028)***	0.3242	(0.083)***	0.8291	(0.085)***	0.7978	(0.032)***
$\omega_{11,st=2}$			5.2067	(0.781)***			4.8480	(0.419)***			9.4070	(0.607)***
$\omega_{12,st=2}$			4.2802	(0.818)***			3.4792	(0.402)***			7.7234	(0.791)***
$\omega_{22,st=2}$			-1.7314	(0.347)***			2.8075	(0.245)***			1.8x10 ⁻⁴	(0.986)
$A_{11,st=2}$			0.2766	(0.107)***			-0.3345	(0.165)**			0.3544	(0.202)*
$A_{22,st=2}$			0.6196	(0.145)***			-0.4611	(0.178)***			-0.2848	(0.137)**
$B_{11,st=2}$			0.5259	(0.247)**			0.9423	(0.059)**			0.9350	(0.077)***
$B_{22,st=2}$			0.2783	(0.183)			0.8873	(0.092)			0.2719	(0.164)*
Transition Probabilities												
p_{12}			0.2563	(0.029)***			0.3810	(0.070)***			0.0160	(0.005)***
p_{21}			0.7407	(0.049)***			0.7001	(0.050)***			0.6011	(0.074)***
Residuals Diagnostics												
LogLik	-3454.8		+390.2		-4055.3		+129.7		-3577.1		+223.7	
	Spot	Futures	Spot	Futures	Spot	Futures	Spot	Futures	Spot	Futures	Spot	Futures
SBIC	-3484.6	-3583.7	+353.8	+512.5	-4085.1	-4061.9	+93.3	+129.7	-3606.9	-3461.5	+187.3	+101.5
Skewness	-0.448***	-0.763***	-0.085	-0.151*	-0.204**	-0.282***	-0.061	-0.063	0.261***	0.269***	-0.077	-0.028
Kurtosis	2.160***	3.458***	-0.075	0.399**	0.868***	1.427***	-0.654***	-0.599***	2.810***	8.599***	-0.192	-0.047
J-B	171.4***	447.6***	1.084	7.867**	28.82***	73.75***	13.85***	11.75***	256.0***	2326***	1.898	0.168
Q(6)	3.032	3.222	5.795	4.436	16.44**	8.496	9.998	9.592	9.706	7.544	7.076	8.873
Q ² (6)	6.893	4.072	14.40**	15.17**	15.78**	5.601	7.152	4.552	10.79*	12.44*	20.78***	21.59***

*, ** and *** indicate significance at 1%, 5% and 10% respectively. Figures in () are the estimated standard errors; See also notes in Table 5.1.

Second, in the high variance regime, the same condition holds only in the Gasoline market. In the WTI market, both error correction coefficients are not significant. In the Heating oil market the results show that in response to a positive deviation the spot price in the next period will remain unresponsive while the futures price will decrease, leading thus to the differential between spot and futures prices to further deviate, which in turn explains the high variance state. This diverse behaviour that arises from assuming two states in the market is not reflected in the single regime VECM GARCH model where the error correction estimates are all in accordance with convergence towards the long-run equilibrium relationship, at 1% significance level. This suggests that the dynamics of the spot-futures relationship vary across the two states of the market; in other words, the adjustment process undergoes regime shifts and does not behave uniformly to shocks to equilibrium across different states but it is rather dependent on the state of volatility (high /low variance state).

Turning next to the conditional variance equation estimates, we can note an evident association between the degree of persistence ($A^2_{ii,st} + B^2_{ii,st}$, for $s_t = 1, 2$) and the state of the market; as expected, a high variance state is associated with high persistence in the variance and vice versa. This is in line with other studies in the literature such as Fong and See (2002) in the oil futures market. The only exception is the Heating oil market, where the low futures variance state incorporates longer memory; this can be attributed to the fact that the high variance regime occurs rarely, particularly at points when we have upward jumps of the basis. Visual inspection of the futures and spot prices shows that these jumps are caused solely by spot price spikes. As a result, the low variance state is dominant throughout the sample period and occasional spot price jumps are captured by the model as the high variance state. We can also note that overall volatility persistence is reduced compared to the single-regime GARCH model. In particular, the “low” state-dependent conditional variance is less sensitive to shocks which in turn have a short-lasting effect. On the other hand, in the high variance state volatility persistence is lower than the single regime GARCH model only in the WTI market⁴.

From the estimated transition probabilities p_{12} and p_{21} we can calculate the duration of being in each regime⁵. For instance in the case of WTI crude oil market the transition probabilities of MRS-BEKK (Table 5.2) are estimated as $p_{12} = 25.6\%$ and $p_{21} = 74.1\%$; these

⁴ We also estimate the MRS-BEKK model using time-varying transition probabilities, conditioned on inventory levels. The results indicated that neither the variation in transition probabilities nor the improvement in the log-likelihood function were significant. The variances of the hedged portfolios did not offer any significant improvement in terms of variance reduction. See Appendix 5.A for some estimation results.

⁵ The average expected duration of being in state 1 as suggested by Hamilton (1989) can be calculated as:

$$\sum_{i=1}^{\infty} i p_{11}^{i-1} (1 - p_{11}) = (1 - p_{11})^{-1} = (p_{12})^{-1}$$

indicate that the average expected duration of being in regime 1 is about 4 ($=1/0.256$) weeks compared to 1.3 ($=1/0.741$) weeks in regime 2. Thus, high variance states are less stable and are characterized by shorter duration compared to low variance states.




Figure 5.1: Smooth Regime Probabilities for WTI Crude Oil: Probability of being in the Low Variance State.

The “smooth” regime probability for the WTI crude oil market derived from the estimated MRS-BEKK model is presented in Figure 5.1⁶. This indicates the likelihood of being in state 1 (low variance state). We can see that state 1 is prevailing whereas the high variance state is short-lasting. For the WTI market until 1995 the low variance state can be attributed to the restoration of Kuwait’s production after the Gulf war and overproduction from the OPEC countries, in combination with relatively weak demand. The low variance state is then disturbed by bad weather conditions in the US and Europe as well as by tension in the Middle East and the Asian crisis in 1998. Similar results emerge when we consider the Unleaded Gasoline and Heating oil markets (graphs are not presented here). In the Unleaded Gasoline market for instance, the low variance regime is less persistent compared to WTI. This is expected since backwardation and supply shortages for light distillates are more frequent, due to constrained refining capacity and the fact that the level of production is also dependent on the quality of the

⁶Based upon the estimated parameter vector $\hat{\theta}$, estimated from data spanning the period $t=1$ to T , three estimates about the unobserved state variable s_t , can be made. The first is the estimated probability that the unobserved state variable at time t equals 1 given the observations 1 to $t < T$ and is termed the filtered probability about s_t . The second is the estimated probability that the unobserved state at time t equals 1 given the entire sample of observations from 1 to T , termed the “smooth” probability. The third is the estimated probability that the unobserved state variable at time $T+1$ equals 1 given observations 1 to T and is termed the expected or predicted probability about s_t . See Chapter 3, section 3.4.1 and Hamilton (1994) for further details.

crude. Even in periods of crude oil oversupply, constrained refining capacity may disturb the supply/demand dynamics of the refined products. For the heating oil market, on the other hand, the regimes seem to be more ‘distinct’ with the low variance state being dominant.

Finally, diagnostic tests of all models are also presented at the bottom of Table 5.3. Tests on the standardised residuals, $\varepsilon_t/(h_t)^{1/2}$, and standardised squared residuals $\varepsilon_t^2/(h_t)$, indicate that all models are well specified with no signs of autocorrelation.

5.5 Time Varying Hedge Ratios & Hedging Effectiveness

Following estimation of the MRS-BEKK models, smooth probability estimates are used to calculate an in-sample state-dependent hedge ratio for each market using Eq. (5.11), after integrating out the unobserved variable st as described in Eq. (5.6), (5.7) and (5.8). The in-sample OLS, GARCH and MRS-BEKK hedge ratios for the WTI market are presented in Figure 5.2. The variation in both time-varying hedge ratios indicates that the portfolio of spot and futures contracts must be revised frequently.




Figure 5.2: Constant OLS, VECM-GARCH and MRS: BEKK Hedge Ratios for WTI Crude Oil.

Figure 5.3 presents the basis for the crude oil market. We can note that when the basis is close to zero the market is in the low variance state (state 1). During these periods the hedge

ratio is higher and less volatile⁷. Similarly, when the market is in the high variance state (state 2) the basis is further away from zero. This indicates that there is a positive relationship between the volatility and the magnitude of the basis; this is consistent with the findings of other studies such as Lee (1994), Choudhry (1997) and Kavussanos and Nomikos (2000).




Figure 5.3: Basis for WTI Crude Oil.

To formally assess the performance of these hedges, portfolios implied by the computed hedge ratios each week are constructed and the variance of returns of these portfolios over the sample is calculated as:

$$Var(\Delta S_t - \gamma_t^* \Delta F_t) \quad (5.12)$$

where γ_t^* are the computed hedge ratios. To evaluate the hedging performance of the MRS models, we estimate hedge ratios based on the naïve model - by taking a futures position which exactly offsets the spot position (i.e. $\gamma_t^* = 1$), on the OLS model of Eq. (5.1), on a VECM (Engle & Granger, 1987; Johansen, 1988), as well as time varying hedge ratios generated from a VECM with GARCH error structure (denoted as GARCH). For benchmarking purposes, we also consider the use of univariate MRS models based on Eq. (5.1)⁸.

⁷ The relationship between the MRS-BEKK hedge ratio and the basis is also investigated by regressing the hedge ratio on the absolute value of the basis. The results indicate that the slope coefficient is significantly negative i.e. the more the basis deviates from zero (high variance), the lower the hedge ratio.

⁸ The univariate MRS specification, yields two hedge ratios, $\gamma_{1,1}$ and $\gamma_{1,2}$, which represent the minimum variance hedge ratios, given the state of the market. The optimal hedge ratio at any point in time can be expressed as: $\gamma_t^* = \pi_{1,t} \gamma_{1,1} + (1 - \pi_{1,t}) \gamma_{1,2}$, where $\pi_{1,t}$ and $\pi_{2,t} = 1 - \pi_{1,t}$ the regime probability of the market

The in-sample period is from January 23, 1991 to June 15, 2005. The in-sample portfolio variances for the three energy commodities are presented in Table 5.3, Panel A. The same table also presents the incremental variance improvement of the MRS-BEKK model against the other models. It can be seen that the MRS hedging strategies outperform the other models in terms of in-sample variance reduction (Panel A). Among the MRS models, the MRS-BEKK is the best model for both the WTI (3.8% - 8.4% improvement) and Heating oil markets (3.9% - 16.8% improvement); in the Unleaded Gasoline market the univariate MRS delivers better variance reduction compared to alternative strategies. Nevertheless, MRS-BEKK is the second best strategy (0.9%-4.6% improvement).

The in-sample performance of the alternative hedging strategies gives an indication of their historical performance. Since investors are more concerned with how well they can hedge their positions in the future, we mainly focus on the out of sample performance of the competing strategies. The out-of-sample period spans from June 22, 2005 to December 27, 2006 (1.5 years) and the assessment is implemented by estimating the models recursively, using only data up to the specific date.

In the case of the MRS-BEKK models, hedge ratios at time $t + 1$ are obtained using a three step procedure. First, estimates of the transition matrix at time t , $\hat{\mathbf{P}}_t$, and the estimated smooth regime probabilities at time t , $\Pr(s_t = 1) = \hat{\pi}_{1,t}$ and $\Pr(s_t = 2) = \hat{\pi}_{2,t}$ are used to forecast regime probabilities at time $t+1$, that is, $\pi_{1,t+1}^e$ and $\pi_{2,t+1}^e$:

$$\begin{pmatrix} \pi_{1,t+1}^e & \pi_{2,t+1}^e \end{pmatrix} = \begin{pmatrix} \hat{\pi}_{1,t} & \hat{\pi}_{2,t} \end{pmatrix} \begin{pmatrix} \hat{P}_{11,t} & \hat{P}_{12,t} \\ \hat{P}_{21,t} & \hat{P}_{22,t} \end{pmatrix} \quad (5.13)$$

Second, we perform one step ahead forecasts of both the variance-covariance of Eq. (5.5) and fitted state-dependent mean equations of Eq. (5.3). Third, by using the Eq. (5.6), (5.7) and (5.8) we integrate the state variable st at each step of the recursive estimation in order to obtain the one step ahead forecast of the optimal hedge ratio at time $t+1$ is computed using Eq. (5.11). The next week (June 29, 2005) the models are re-estimated with the new observation included in the dataset, and this exercise is repeated for every week in the out-of-sample period. For GARCH-based hedges, the model is re-estimated each week during the out-of-sample period and hedge ratios are generated by one-step ahead forecasts of the time varying variance-

being in state 1 and 2, respectively, at any point in time with $0 \leq \pi_{1,t}, \pi_{2,t} \leq 1$ (see Alizadeh and Nomikos, 2004b).

covariance matrix. In the case of VECM a different hedge ratio is obtained each week by re-estimating the model. Finally, for the univariate-based MRS hedges the one-step ahead optimal hedge ratio at time $t+1$, is calculated as the mean hedge ratio weighted by the forecasts of the regime probabilities.

Table 5.3, Panel B displays the results from the out-of-sample performance of the competing hedging strategies. The same table also presents the incremental variance improvement of the MRS-BEKK model against the other models. Looking at the results for both WTI and Unleaded Gasoline market, the highest reduction in the out-of-sample portfolio variance is achieved by the MRS-BEKK model. Compared to the OLS hedge the gain in variance reduction is 6.3% for WTI and 14.7% for Unleaded Gasoline. Regarding the Heating oil market, the greatest variance reduction is provided by the naive hedge whereas the MRS-BEKK model achieves almost the same level of variance reduction as the OLS hedge ratio. One possible explanation for this surprising result may be the fact that occasionally MRS models do not provide accurate forecasts on an out-of-sample basis. This may be due to parameter instability between in-sample and out-of-sample periods as well as uncertainty regarding the unobserved regime, as mentioned in Engel (1994) and Marsh (2000). Another reason may be attributed to the fact that in the specific market extreme spot price spikes are identified as the high variance state, making the latter short-lasting and rare.

Dynamic hedging strategies are more costly to implement than static hedges since they require frequent updating and rebalancing of the hedged portfolio. Consequently, hedging effectiveness is more appropriately assessed by considering the economic benefits from hedging using the hedger's utility function as in Kroner and Sultan (1993), and Lafuente and Novales (2003). If η is the degree of risk aversion ($\eta > 0$) of the individual investor and rp_{t+1} represents the returns from the hedged portfolio i.e. $\Delta S_{t+1} - \gamma^*_{t+1} \Delta F_{t+1}$, the relevant utility function employed is:

$$E_t U(rp_{t+1}) = E_t(rp_{t+1}) - \eta Var_t(rp_{t+1}) \quad (5.14)$$

Table 5.3: Hedging Effectiveness of MRS Against the Constant and Alternative Time-Varying Hedge Ratio Models

	WTI light sweet Crude Oil				Unleaded Gasoline				Heating Oil #2			
	Variance	Variance Improvement of MRSBEKK	Utility	Var _(5%) (\$)	Variance	Variance Improvement of MRSBEKK	Utility	Var _(5%) (\$)	Variance	Variance Improvement of MRSBEKK	Utility	Var _(5%) (\$)
Panel A: In- sample Hedging Effectiveness												
Unhedged	22.935	83.53%	-91.740	78,772.9	32.897	74.97%	-131.59	94,342.1	27.767	68.95%	-111.07	86,674.6
Naïve	3.9277	3.82%	-15.711	32,598.4	8.3128	0.94%	-33.251	47,424.3	9.0238	4.45%	-36.095	49,410.8
Constant	3.9276	3.81%	-15.710	32,598.0	8.3048	0.85%	-33.219	47,401.5	9.0145	4.35%	-36.058	49,385.3
VECM	3.9323	3.93%	-15.729	32,617.5	8.3073	0.88%	-33.229	47,408.6	9.0181	4.39%	-36.072	49,395.2
GARCH	4.1249	8.42%	-16.500	33,406.7	8.6316	4.60%	-34.526	48,325.1	10.362	16.79%	-41.448	52,947.9
MRS	3.9273	3.81%	-15.709	32,596.7	8.1256	-1.34%	-32.502	46,887.3	8.9699	3.88%	-35.880	49,263.0
MRS-BEKK	3.7778	-	-15.111	31,970.3	8.2345	-	-32.938	47,200.4	8.6221	-	-34.488	48,298.5
Panel B: Out-of- sample Hedging Effectiveness												
Unhedged	14.286	88.48%***	-57.144	62,170.2	71.163	82.25%***	-284.65	138,756.8	22.869	89.92%***	-91.476	78,659.4
Naïve	1.7484	5.89%**	-6.994	21,749.4	15.141	16.59%***	-60.564	64,003.6	2.2943	-0.48%	-9.177	24,914.5
Constant	1.7550	6.25%**	-7.020	21,790.4	14.796	14.65%***	-59.184	63,270.2	2.3052	-0.01%	-9.220	24,973.6
VECM	1.7420	5.55%**	-6.968	21,709.6	14.777	14.54%***	-59.108	63,229.6	2.3160	0.46%	-9.264	25,032.1
GARCH	1.7741	7.25%**	-7.096	21,908.7	13.710	6.38%***	-53.956	60,904.0	2.4467	5.78%	-9.787	25,728.7
MRS	1.7475	5.84%**	-6.990	21,743.8	14.284	11.59%***	-57.136	62,165.9	2.3331	1.19%	-9.332	25,124.3
MRS-BEKK	1.6454	-	-6.582	21,099.1	12.629	-	-50.516	58,453.7	2.3054	-	-9.222	24,974.7

- The in-sample period is from January 23, 1993 to June 15, 2005 (752 observations) whereas the out-of-sample period is from June 22, 2005 to December 27, 2006 (80 observations).
- Variance denotes the variance of the hedged portfolio. Note that the variance corresponds to logarithmic returns multiplied by 100 [Eq. (5.12)]. Figures in Bold denote the best performing model for each criterion.
- Variance Improvement of MRSBEKK measures the incremental variance reduction of the MRS-BEKK model versus the other models. This is estimated using the formula: $[\text{Var}(\text{Model}_i) - \text{Var}(\text{MRS-BEKK})] / \text{Var}(\text{Model}_i)$.
- Utility is the average weekly utility for an investor with a mean-variance utility function [Eq. (5.14)] and a coefficient of risk aversion of 4, using different hedging strategies.
- $\text{VaR}_{(5\%)}$ is the Value-at-Risk estimated using Eq. (5.15) with $\Phi(c)$ equal to the normal distribution 5% quantile i.e. -1.645.
- Asterisks (*, **, ***) in the column named “Variance Improvement of MRS-BEKK” indicate that the MRS-BEKK model outperforms the competing model at 1%, 5% and 10%, respectively; the p-values are provided from White’s (2000) Reality Check using the stationary bootstrap of Politis and Romano (1994).

Another way of considering the economic benefits from the proposed hedge is to look at the reduction in the Value-At-Risk (VaR) exposure, arising from the different hedging strategies. Assuming a normal distribution, if we denote as W_0 the initial value of the portfolio and $\Phi(c)$ the inverse of the standard Gaussian cumulative distribution function, the portfolio VaR is simply a constant multiple of the hedged portfolio standard deviation where the multiple is determined by the VaR confidence level $1-c$:

$$VaR = W_0 \left[E(rp_{t+1}) + \Phi(c) \sqrt{Var(rp_{t+1})} \right] \quad (5.15)$$

From Table 5.3, Panel B the average weekly variance of returns from the hedged position in the Unleaded Gasoline market is 14.8 when the constant hedge ratio is used and 12.6 when the MRS-BEKK model is used. Assuming that expected returns from the hedged portfolio are equal to zero and the degree of risk aversion is 4 then, on average, one obtains a weekly utility of $U(rp_{t+1}) = -4(14.8) = -59.2$ if the constant hedge ratio is used and $U(rp_{t+1}) = -4(12.6) = -50.5$ when the MRS-BEKK hedge ratio is used. Hence, by using the MRS-BEKK, hedgers in the market can benefit from an increase in the average weekly utility of 8.7 - y, over the constant hedge ratio, where y represents the reduced returns caused by the transaction costs incurred due to portfolio rebalancing; assuming transaction costs in the range of 0.01-0.02% (due to rebalancing), the MRS hedge would still result in an improvement in utility for an investor with a mean-variance utility function and $\eta = 4$. Similarly, results of the weekly VaR for a portfolio value of \$1m with 95% confidence level indicate that one obtains a weekly VaR = \$1m[- 1.65 (14.8)^{1/2}] = - \$63,230 if the constant hedge ratio is used and a VaR of \$1m[- 1.65 (12.6)^{1/2}] = - \$58,454 when the MRS-BEKK hedge ratio is used. Hence, by using the MRS-BEKK, hedgers in the market can benefit from a decrease in the average weekly VaR of \$4,776 over the OLS hedge, which results in an annualised decrease in VaR of \$34,440 or a decrease of 3.4% over the initial investment. Therefore, investors would prefer the MRS-based strategies to the constant strategy since the increase in utility and decrease in VaR more than offsets the higher transaction costs due to rebalancing.

5.6 Data Snooping Bias

Regardless of the encouraging results of the performance of the proposed MRS-BEKK hedging strategy, an important issue which has to be considered is that of data snooping. According to Sullivan et al. (1999) and White (2000) data snooping occurs when a dataset is

used more than once for data selection and inference purposes. In other words, using the same dataset frequently for testing different strategies may increase the probability of having satisfactory results purely due to chance or due to the use of posterior information rather than the superior ability of the competing strategies. In order to discount the possibility that the performance of the MRS-BEKK model may be due to data snooping bias we implement White's (2000) Reality Check (RC), in a similar way as was employed in Chapter 4. In doing so, we first construct a relative performance measure which can be defined as:

$$fm_{k,t+1} = \left[\Delta S_{t+1} - \hat{\gamma}_{k,t+1}^* \Delta F_{t+1} \right]^2 - \left[\Delta S_{t+1} - \hat{\gamma}_{MRS-BEKK,t+1}^* \Delta F_{t+1} \right]^2 \quad (5.16)$$

where k represents the k^{th} benchmark model and the expression in $[\cdot]$ is the loss function chosen; that is the squared out-of-sample portfolio return which in fact is an unbiased estimate of the true conditional variance (see for instance Andersen and Bollerslev, 1998). If the MRS-BEKK model outperforms the k^{th} model, the expected value of the performance measure will be positive. Therefore, we set the null such as that rejecting it would imply that the MRS-BEKK model is superior in terms of variance reduction compared to the competing hedging strategy. Mathematically:

$$H_0 : \max_k \{ E(fm_k) \} \leq 0 \quad (5.17)$$

Then, following White (2000), we can test the null hypothesis by obtaining the test statistic of the RC as $T_n^{RC} = \max_k (n^{1/2} \bar{f}m_k)$ where $\bar{f}m_k = n^{-1} \sum_{t=1}^n fm_{k,t}$ and n is the number of one-step ahead periods. In order to construct the test statistic, we use the stationary bootstrap technique of Politis and Romano (1994) to regenerate random paths of portfolio returns, whilst maintaining the distributional properties of the original series⁹. We then construct the loss function of Eq. (5.16) using the simulated portfolio returns which, in turn, generates a distribution of hedging statistics under the different hedging strategies. Let $\bar{f}m_k^*(b)$ represent the sample mean of the relative performance measure calculated from the b^{th} bootstrapped sample for $b = 1, \dots, B$. The RC p-value is obtained by comparing T_n^{RC} with the quantiles of the empirical distribution of T_n^{RC*} :

⁹ Politis and Romano (1994) method re-samples blocks of varying length from the original data, where the block length follows a geometric distribution, with a given mean block length.

$$T_n^{RC*} = \max_k \left\{ n^{1/2} \left(\bar{f}m_k^*(b) - \bar{f}m_k \right) \right\} \quad (5.18)$$

The null hypothesis that the variance improvement from the MRS-BEKK is not a better hedge from the other models is tested using White's RC with 1,000 bootstrap simulations and a smoothing parameter of $q=0.1$ (see Politis and Romano, 1994 for more technical details on the stationary bootstrap as well as Appendix 4.C). Results are reported in Table 5.3. The MRS-BEKK model provides significantly greater variance reduction in the WTI and Unleaded Gasoline markets across all models, at conventional significance levels. In the Heating oil market the MRS-BEKK model fails to outperform the competing strategies; however, when we invert the null and set the MRS-BEKK model as the benchmark in Eq. (5.16), we still cannot reject the null of no superior predictive ability of the naïve hedge over the MRS-BEKK model (p-value = 0.241).

5.7 Downside Risk Measures

Although variance reduction gives the overall picture about how well a hedging strategy performs, it does not consider whether there are any differences in the degree of hedging performance between long and short positions. The motivation for investigating this stems from both the pitfalls associated with variance as a measure of hedging effectiveness and the specific properties inherent in the MRS-BEKK model.

First, variance assigns the same weight to positive gains and negative losses. Under the assumption of either quadratic utility functions or multivariate elliptical distributed returns the investor is only concerned about the expected return and the standard deviation of the hedged portfolio¹⁰. However, in practice these assumptions are not likely to hold and a number of metrics have recently been proposed in the literature that are able to deal with possible asymmetries in the profiles of risk averse investors. For instance, Cotter and Hanly (2006) evaluate the hedging performance of short and long hedging positions based on Lower Partial Moments (LPM) and Value-at-Risk (VaR) estimates and find differences in terms of the best strategy compared to the traditional variance metric.

Second, it is of interest to test whether regime switching models are capable of adequately capturing the skewness and kurtosis typical of financial data and, if this is true, whether this can be used effectively to eliminate downside risk within the minimum-variance

¹⁰ That is because the hedged portfolio and the assets comprising the portfolio share the same distributional properties.

framework. Under the Markovian formulation, as specified in Eq. (5.3), (5.4) and (5.5), the dynamics of conditional means and variances imply that time-varying skewness and excess kurtosis are inherent in the model (see Haas et. al, 2004a for more details and derivation of higher moments of mixed normal distributions). Consequently, one would expect the MRS based hedge ratios to capture possible asymmetries that may affect short and long hedging positions differently.

In order to remove the effect of upside gains from the variance, the semi-variance metric is employed which acts as a measure for a downside risk averse investor, who is concerned about the variability of negative losses. Mathematically, this can be expressed as:

$$sv_{(-)} = \frac{1}{T} \sum_{i=1}^T \{\min(0, rp_{t+1} - u)\}^2 \quad (5.19)$$

This is equivalent to the second order lower partial moment (LPM) where the target return (threshold) u is set to zero in order to distinguish between positive and negative realised portfolio returns rp_{t+1} . A short hedging position is equivalent to selling futures contracts against the purchase of the underlying asset; hence the investor is concerned about negative semi-variance (the payoff of a short hedger is $rp_{t+1} = \Delta S_{t+1} - \gamma^*_{t+1} \Delta F_{t+1}$). Similarly, a long hedger is concerned about positive semi-variance, in which case in Eq. (5.19) we only consider the positive returns.

Table 5.4 presents the negative and positive semi-variance figures in Panel A and B, respectively where negative and positive semi-variance reflect the downside variation in the performance of short and long hedging strategies, respectively. Overall, the results indicate that the improvement in the semi-variance using the MRS-BEKK model is better in 4 out of 6 cases, thus supporting the suggested strategy. We also assess the different strategies using White's (2000) RC. The results in Table 5.4 illustrate that the MRS-BEKK is significantly better in the WTI market only for long hedgers (but not against GARCH) at 10% significance level. In the Unleaded Gasoline market the MRSBEKK model is significantly better than the competing models in hedging short positions (the improvement in semi-variance against the GARCH is 22%). For long hedgers the simple MRS model significantly outperforms the other models. In the Heating oil market, the MRS-BEKK is significantly better than the GARCH in short hedging positions but it is significantly outperformed by the VECM model. In long hedging positions the MRS-BEKK is not significantly better than the competing models, according to the Reality Check.

Table 5.4: Effectiveness Long/Short Hedging Positions of Markov Regime Switching Against the Constant and Alternative Time-Varying Hedge Ratio Models

	WTI light sweet Crude Oil				Unleaded Gasoline				Heating Oil #2			
	Semi-Variance	Semi-Variance Improvement of MRSBEKK	Semi-Utility	VaR _(5%) (\$)	Semi-Variance	Semi-Variance Improvement of MRSBEKK	Semi-Utility	VaR _(5%) (\$)	Semi-Variance	Semi-Variance Improvement of MRSBEKK	Semi-Utility	VaR _(5%) (\$)
Panel A: Short Hedgers Positions (Negative Semi-Variance)												
Unhedged	6.7703	87.16%***	-27.081	42,798.7	30.879	89.77%***	-123.516	91,402.7	10.469	87.38%***	-41.874	53,220.6
Naïve	0.8935	2.68%	-3.574	15,548.3	5.2580	39.95%***	-21.032	37,717.0	1.2902	-2.41%	-5.161	18,683.4
Constant	0.9011	3.49%	-3.604	15,614.1	5.1261	38.40%***	-20.504	37,241.0	1.2728	-3.80%	-5.091	18,557.0
VECM	0.8863	1.88%	-3.545	15,485.1	5.0341	37.27%***	-20.136	36,905.3	1.2724	-3.84%(+)	-5.090	18,554.1
GARCH	0.9257	6.05%	-3.703	15,825.4	4.0500	22.03%**	-16.200	33,102.0	1.4199	6.95%**	-5.680	19,600.0
MRS	0.8933	2.65%	-3.573	15,546.2	5.0575	37.57%***	-20.230	36,990.9	1.3046	-1.27%	-5.219	18,787.4
MRS-BEKK	0.8696		-3.479	15,338.9	3.1576	-	-12.631	29,228.5	1.3212	-	-5.285	18,906.5
Panel B: Long Hedgers Positions (Positive Semi-Variance)												
Unhedged	7.3477	89.72%***	-29.391	44,586.4	39.405	76.25%***	-157.62	103,253.1	12.115	92.08%***	-48.460	57,251.8
Naïve	0.8330	9.29%*	-3.332	15,012.6	9.7001	3.54%***	-38.800	51,228.9	0.9774	1.88%	-3.910	16,261.6
Constant	0.8319	9.17%*	-3.328	15,002.7	9.4908	1.41%**	-37.963	50,673.2	1.0054	4.62%	-4.022	16,492.9
VECM	0.8340	9.39%*	-3.336	15,021.2	9.5673	2.20%***	-38.269	50,877.1	1.0165	5.65%	-4.066	16,583.7
GARCH	0.8265	8.57%	-3.306	14,953.3	9.5698	2.22%	-38.279	50,883.7	0.9978	3.89%	-3.991	16,430.4
MRS	0.8324	9.22%*	-3.329	15,006.5	9.0516	-3.38%(+)	-36.206	49,486.9	1.0014	4.23%	-4.005	16,460.0
MRS-BEKK	0.7556		-3.023	14,298.2	9.3571	-	-37.429	50,315.0	0.9590	-	-3.836	16,107.8

- Results are presented for the out-of-sample period i.e. June 22, 2005 to December 27, 2006 (80 observations).
- Semi-Variance denotes the semi-variance of the hedged portfolio. [Eq. (5.19)].
- Semi-Variance Improvement of MRSBEKK measures the incremental semi-variance reduction of the MRS-BEKK model versus the other models. This is estimated using the formula: $[SVar(Model_i) - SVar(MRS-BEKK)]/SVar(Model_i)$.
- Semi-Utility is the average weekly semi-utility for an investor with a mean-variance utility function [Eq. (5.14)] and a coefficient of risk aversion of 4, using different hedging strategies.
- VaR_(5%) is the Value-at-Risk estimated by using Eq. (5.15) with $\Phi(c)$ equal to the normal distribution 5% quantile i.e. -1.645.
- A cross (+) indicates that the benchmark model outperforms the MRS-BEKK model at conventional significance levels; the p-values are provided from White's (2000) Reality Check using the stationary bootstrap of Politis and Romano (1994).
- See also notes in Table 5.3.

Table 5.4 reports also the semi-utility and asymmetric VaR calculated from Eq. (5.14) and (5.15), respectively by replacing variance with semi variance. Symmetric distributions would imply that utility is equal in both short and long positions $E_t U(rp_{t+1})^{(+)} = E_t U(rp_{t+1})^{(-)}$, which is not the case, at least in the Gasoline market. Moreover, with the downside risk expressed this way we can still use the quantiles of normal distribution (assuming that positive/negative returns follow half normal distribution) and calculate VaR estimates. It can be seen that these figures are reduced compared to the variance measure but the distributions have fat tails since the sum of short and long VaR is greater than the VaR estimated from the standard deviation.

5.8 Conclusions

In this chapter we examined the performance of hedge ratios generated from Markov Regime Switching models in the oil futures markets. The rationale behind the use of these models stems from the fact that the dynamic relationship between spot and futures prices may be characterized by regime shifts. This, in turn, suggests that by allowing the hedge ratio to be dependent upon the “state of the market”, one may obtain more efficient hedge ratios and hence, superior hedging performance compared to the methods which are currently being employed. We introduce a Markov Regime Switching Vector Error Correction model with GARCH error structure. This specification links the concept of disequilibrium (as measured by the error correction coefficients) with that of uncertainty (as measured by the conditional second moments) across high and low volatility regimes. The effectiveness of the MRS time-varying hedge ratios is investigated in the NYMEX WTI Crude Oil, Unleaded Gasoline and Heating Oil markets. The estimated models indicate that there is marked asymmetry in both conditional means and conditional volatilities under different market conditions. Moreover, all the MRS based hedge ratios appear to be higher when the volatility in the market is low, a finding that is in line with theory. In and out-of-sample tests indicate that by allowing the variances and the covariance of spot-futures returns to be state dependent, hedging effectiveness is significantly improved in most cases, as indicated by White’s (2000) Reality Check. Overall, the results indicate that using MRS models market agents may be able to obtain superior gains, measured in terms of both variance reduction and increase in utility. This finding holds even when we examine the downside risk and consider the asymmetric risk profile of long and short hedgers.

The next chapter, *Petroleum Term Structure Dynamics, Inter-Commodity Dependencies and the Role of Regimes*, deals with an important issue of petroleum market dynamics, the term structure of petroleum futures. As opposed to Chapters 4 and 5 which employed the nearest to

expiry and volume based futures contracts, Chapter 6 will deal with the entire forward curve. The objective of this study is to exploit the information content of the dependence structure of petroleum futures curves and describe inter-dependencies between petroleum commodities under different regimes. Before presenting our empirical evidence, Chapter 6 will first, provide a short introduction regarding the use of term structure models. This will be followed by some technical details on factor decomposition and next, a thorough explanation of the theoretical background and the estimation procedure will be supplied. A parsimonious regime switching model of correlated futures curves will then be presented where each state has its dynamic characteristics. The model will be successfully fitted to daily historical futures curves from 1994 to 2009, providing strong statistical evidence, not only regarding the presence of changes in regime but also concerning the specific properties of factor dynamics, mean reversion, co-integration and co-movement. After presenting the empirical results and some relevant and potentially useful applications in forecasting, as it will be seen, the model offers great improvement over appropriately assumed benchmarks.

APPENDIX 5.A : Time Varying Transition Probabilities

In this appendix, we explore whether inclusion of inventory levels in the MRS set up is able to provide any further improvement in the hedging performance perhaps by driving the switches between volatility regimes. Weekly and monthly data for the industrial inventories of crude oil and petroleum products for all OECD countries are available from the US Department of Energy - Energy Information Administration (EIA). Table 5.A.1 below displays the estimated MRS-BEKK model for the three petroleum futures markets with transition probabilities conditioned on the lagged inventory levels. These transition probabilities are estimated using the following logistic function:

$$p_{ij,t} = \frac{1}{1 + \exp(\phi_{i,1} + \phi_{i,2}IN_{t-1})} \text{ for } i \neq j \text{ and } i,j = \{1, 2\} \quad (5.A.1)$$

where $\phi_{1,1}$, $\phi_{1,2}$, $\phi_{2,1}$, $\phi_{2,2}$ are parameters to be estimated by maximum likelihood along with the other parameters of the model. IN_{t-1} represents the more recent update of inventories¹¹. Comparing the results between these models and the restricted versions (presented in Table 5.2) we can note that the coefficients of the conditional means and variances are very similar both in terms of magnitude and sign as well as in terms of statistical significance. In addition, transition probabilities do not seem to vary (since the estimated coefficients $\phi_{i,j}$ are insignificant) and, actually, the time varying transition probabilities are very close to the constant probabilities p_{12} and p_{21} as estimated by the restricted MRS-BEKK. This is also confirmed by looking at Figure 5.A.1 where we can see that the time-varying transition probabilities of WTI have a mean value very close to the unconditional ones; these are 0.256 and 0.741, respectively, according to Table 5.2. The same holds for Unleaded Gasoline probabilities, also depicted in the graph; however, for the particular product market we can observe a certain degree of variation, especially for p_{12} which fluctuates between 0.27 and 0.50 (this corresponds to the 0.381 probability of Table 5.2). What is more, the improvement in the Log-Likelihood is negligible and not statistically significant using standard Likelihood Ratio (LR) tests.

¹¹ In the analysis we use the logarithmic demeaned and detrended inventory levels so that the series reflects deviations from the normal inventory levels.

Table 5.A.1: MRS-BEKK models with Transition Probabilities Conditioned on Inventories

	West Texas Intermediate		Unleaded Gasoline		Heating Oil # 2		
Mean Equation							
$\alpha_{S,st=1}$	-1.2143	(0.113) ^{***}	-0.1716	(0.018) ^{***}	-0.1964	(0.027) ^{***}	
$\alpha_{F,st=1}$	-0.2268	(0.113) ^{**}	-0.0633	(0.016) ^{***}	-0.0370	(0.033)	
$\alpha_{S,st=2}$	-0.1585	(0.277)	-0.2959	(0.089) ^{***}	-0.0118	(0.209)	
$\alpha_{F,st=2}$	0.2703	(0.255)	-0.0837	(0.061)	-0.4707	(0.245) ^{**}	
Variance Equation							
$\omega_{11,st=1}$	3.8205	(0.127) ^{***}	3.3637	(0.222) ^{***}	2.3358	(0.281) ^{***}	
$\omega_{12,st=1}$	3.8392	(0.128) ^{***}	3.1566	(0.201) ^{***}	2.1321	(0.244) ^{***}	
$\omega_{22,st=1}$	-0.3324	(0.029) ^{***}	3.8×10^{-6}	(0.066)	0.3443	(0.062) ^{***}	
$A_{11,st=1}$	-0.0954	(0.033) ^{***}	0.1922	(0.118)	0.2783	(0.035) ^{***}	
$A_{22,st=1}$	-0.0815	(0.027) ^{***}	0.1709	(0.095) [*]	0.2252	(0.055) ^{***}	
$B_{11,st=1}$	6.7×10^{-7}	(0.010)	0.4391	(0.008) ^{***}	0.7441	(0.040) ^{***}	
$B_{22,st=1}$	8.6×10^{-7}	(0.012)	0.3074	(0.028) ^{***}	0.7899	(0.030) ^{***}	
$\omega_{11,st=2}$	5.1883	(0.686) ^{***}	4.8385	(0.415) ^{***}	9.3251	(1.247) ^{***}	
$\omega_{12,st=2}$	4.2739	(0.861) ^{***}	3.4693	(0.439) ^{***}	7.7145	(0.564) ^{***}	
$\omega_{22,st=2}$	1.7292	(0.336) ^{***}	2.8070	(0.322) ^{***}	-0.0002	(1.727)	
$A_{11,st=2}$	0.2786	(0.163) [*]	0.3402	(0.120) ^{***}	0.3608	(0.185) [*]	
$A_{22,st=2}$	0.6226	(0.200) ^{***}	0.4617	(0.170) ^{***}	-0.2764	(0.162) [*]	
$B_{11,st=2}$	0.5391	(0.081) ^{***}	0.9404	(0.043) ^{***}	0.9326	(0.072) ^{***}	
$B_{22,st=2}$	0.2852	(0.127) ^{**}	0.8870	(0.088) ^{***}	0.2478	(0.223)	
Transition Probabilities							
$P_{12}:$	$\phi_{1,1}$	1.0762	(0.166) ^{***}	0.4938	(0.239) ^{**}	4.0971	(0.865) ^{***}
	$\phi_{2,1}$	1.4053	(2.675)	4.1741	(2.465) [*]	1.4895	(3.505)
$P_{21}:$	$\phi_{1,2}$	-1.0678	(0.245) ^{***}	-0.8264	(0.309) ^{***}	-0.8535	(0.968)
	$\phi_{2,2}$	-1.6487	(4.600)	-1.8260	(3.962)	-4.2478	(3.932)
LogLik	3,064.4		3,924.1		3,351.5		
Compared to constant transition Prob.	+0.25		+1.68		+1.9		
LR-stat	0.50		3.36		3.80		
	[0.779]		[0.186]		[0.150]		
In-sample Variance of hedged portfolio	3.7793		8.2377		8.6416		
Compared to constant transition Prob. (%)	-0.040		-0.039		-0.226		
Compared to Naïve Hedge (%)	3.927		0.912		4.423		

See also notes in Table 5.1 and 5.2.

Finally, in-sample variances of the hedged portfolio indicate that the new model does not offer any significant improvement in terms of variance reduction. In fact, the variances of the hedged portfolios appear to be marginally higher than the ones obtained by the restricted MRS-BEKK model. Nevertheless, compared to the naïve hedge, there is still improvement of the range 0.9-4.4%. Summarising, it seems that inventories are not significant predictors of transition between regimes. This could be due to the fact that it is the expectations of inventories, rather than their actual realisations, that actually drive energy prices or alternatively, that a certain amount of inventory levels beliefs is captured by the regime

probabilities as unobservable component. More sophisticated models may be estimated, for instance the MRS-BEKK-X models (as a multivariate extension to the models presented in Chapter 4 with the conditional variances equations augmented by the inventory levels. However, given that our model is already complex, by including an exogenous variable in the conditional variance equation involves four additional parameters to be estimated and, therefore, this approach is not pursued here.

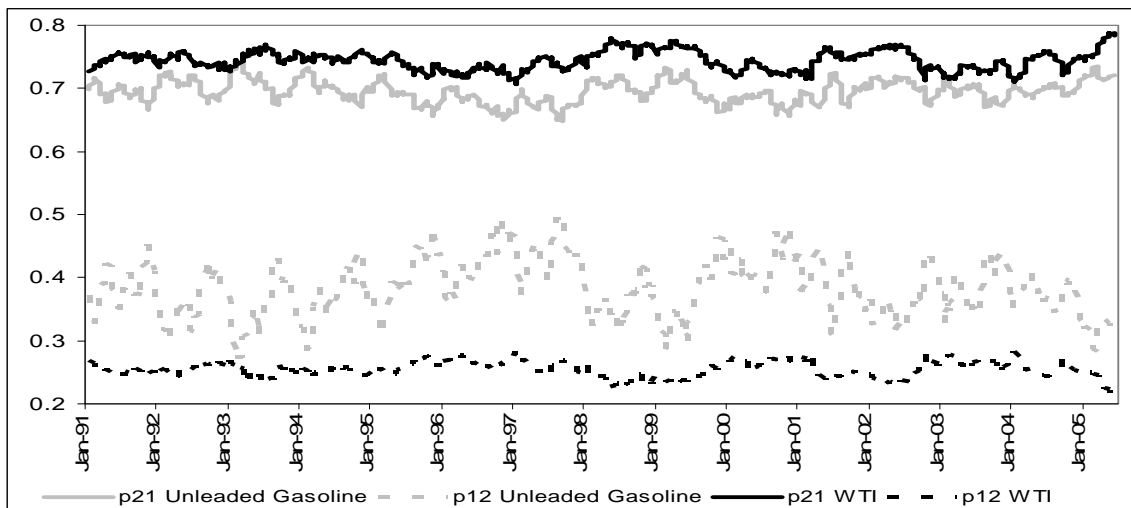


Figure 5.A. 1: Time Varying Transition Probabilities of WTI Crude Oil

APPENDIX 5.B: A Note on Seasonality and Hedging

Crude oil, unleaded gasoline and heating oil exhibit differences in the individual supply/demand fundamentals. First, crude oil is the leading commodity in international trade and its price is determined mainly by global economic conditions, thus, crude oil prices are not expected to show signs of strong seasonal patterns. Demand for petroleum products in the US, on the other hand, depends on local supply and demand conditions, as well as the availability of refining capacity near the centres of demand. Thus, crude oil distillates exhibit seasonality in their price behaviour with heating oil demand increasing in winter while Gasoline demand increasing in the summer. However, although seasonal patterns¹² are not modelled explicitly in the MRS framework, a certain degree of seasonal variation should be captured by the regime probabilities and regime dependent volatilities and correlations. That is one of the advantages of the Markov formulation, the fact that its flexible structure can capture unobserved components. Treating seasonality as an unobserved component is plausible (see for instance Harvey and Scott, 1994) since seasonality is not always obvious and irregular seasonality patterns in energy prices are not uncommon (seasonal breaks i.e. different seasonal behaviour throughout the years).

In order to remove seasonality we should be able to identify consistent peaks and troughs throughout a given time period (month or quarter) that can be explained by

¹² Seasonality in financial time series has traditionally been dealt with either by using deseasonalised data or by including additional explanatory variables in the model setting i.e. dummies or trigonometric functions. The limitations of the first approach have been reported in several studies. First, the question of which method to use is of crucial importance since different methods produce different results. For instance, Jaeger and Kunst (1990) examine the robustness of persistence to shocks in seasonally adjusted series and show that Census X-11 method overestimates the persistence of shocks compared to seasonal differencing or seasonal dummy adjustment. Second, Ghysels et al. (1996) demonstrate that seasonal adjustment is possible to introduce nonlinear behaviour in linear unadjusted series. Third, although smoothing eliminates some of the predictable seasonal variation, it is possible to either introduce artificial autocorrelation or retain a residual seasonal pattern (see also Ghysels, 1994). More recently, seasonally adjusted data have been also questioned in the regime switching setting by studies such as Frances and Paap (1999) and Luginbuhl and de Vos (2003). Both studies report longer and shallower recession periods when using seasonally adjusted data. In general, it is accepted that adjusted data can distort the information about the extent and timing of structural breaks causing problems in regime forecasting (see also Matas-Mir and Osborn; 2004). Finally, since seasonality is not constant, the effects throughout time are expected to be asymmetric across regimes and seasonally adjusted series are expected to result in loss of information across the identified regimes. Addressing the issue of seasonal adjustment is admittedly a controversial issue (see also Wallis; 1974 and Harvey and Jaeger; 1993) especially in the Markov switching framework which is employed in our study. Alternatively, the seasonal component can be a part of the model to be estimated. For instance, Fong and See (2003) modelled crude oil futures as a regime switching GARCH process and incorporated a dummy variable in the conditional variance equation to account for higher demand in winter seasons driven by heating oil. However, in a multivariate Markov setting such an analysis would complicate the estimation procedure since including dummy variables or sinusoidal/cosine functions would increase substantially the number of parameters to be estimated (over-parameterisation) and any gains would most likely be marginal.

fundamentals. The most common reasons for seasonality of petroleum commodities are weather and holidays but the pattern is not clear because factors such as OPEC policy, build-up of strategic reserves and in general, the politics surrounding the balance of demand and supply are key factors that affect the fundamentals in these markets. Seasonality in oil prices is rather a “hidden” function. Market anticipates seasonal prices by building stocks (subject to storage costs) but changes in demand can be at a great extent irregular. To gain an insight into the issue, we performed a test regressing the aggregate standardised, $\varepsilon_t/(h_t)^{1/2}$ and squared standardised, $\varepsilon_t^2/(h_t)$, residuals of the MRS-BEKK model on monthly and quarterly dummies. The results in Table 5.B.1 indicate that for all three markets the null hypothesis that the coefficients of the dummies are jointly zero cannot be rejected at conventional significance levels, indicating that there are no seasonal effects in the $\varepsilon_t/(h_t)^{1/2}$ series. However, at 1% significance level there is evidence of seasonality only in the Heating oil Market as indicated by the squared series.

Table 5.B.1 : LR Tests on the Residuals of the MRS-BEKK Model

	WTI Crude Oil		Unl. Gasoline		Heating Oil # 2	
	Spot	Futures	Spot	Futures	Spot	Futures
Panel A: Standardized Residuals LR test						
Monthly Dummies	8.678	8.219	17.24	12.47	14.25	12.04
Quarterly Dummies	2.563	3.194	0.752	3.103	2.201	5.918
Panel B: Squared Standardized Residuals LR test						
Monthly Dummies	19.40*	17.54*	12.90	13.23	38.73***	32.87***
Quarterly Dummies	7.805*	5.906	3.191	5.962	19.58***	14.45***

Figures presented above are the statistics of a chi squared distribution with degrees of freedom equal to 11 (monthly dummies) and 3 (quarterly dummies). Asterisks *, ** and *** indicate significance at 10%, 5% and 1% significance levels, respectively.




Figure 5.B.1: Monthly Seasonal Components of Spot-Futures Weekly Volatilities

Moreover, Figures 5.B.1 and 5.B.2 present the estimated average weekly conditional volatilities and correlation estimates across the different months of the year from the MRS-

BEKK model. We see that, only in the Heating oil market there is a significant drop in volatilities and a corresponding rise in correlation during the summer months. Heating oil stocks tend to be highest in October and November and reach a minimum in the February - March when demand declines. June and July represent the summer fill season in anticipation of the colder weather ahead. The peak (trough) in the estimated volatilities (correlation) in February can be explained by the fact that when stocks normally reach the minimum levels at the same month, changes in demand (due to unexpected cold weather) result in more erratic price changes because supply is inelastic in the short-run.




Figure 5.B.2: Monthly Seasonal Components of Spot-Futures Correlations

Furthermore, regressing the regime probabilities on dummies, resulted the following equation in Heating oil market¹³:

$$\pi_{1,t} = \begin{matrix} 0.926 & +0.030Quarter_2 & +0.064Quarter_3 & +0.059Quarter_4 \\ (0.018)^{***} & (0.021) & (0.017)^{**} & (0.019)^{***} \end{matrix}$$

implying that the probability of being in the low variance regime is around 6% higher for the period June-November ($Quarter_2 + Quarter_3$) at 1% significance level. However, again, the adjusted R^2 indicates that only 3% of the regime probabilities can be explained by seasonal variations.

Summarising, although seasonality in petroleum prices is a stylised fact, this not of major concern in hedging. Myers and Thomson (1989) and Viswanath (1993) argue that the

¹³ In the WTI and Gasoline Markets regime the dummies as regressors of regime probabilities were not significant (at least at 5% significance level). This does not necessarily imply that seasonality has no impact in the regime probabilities. It may be that seasonal component is stochastic and the effect is differs from year to year.

predictable component of spot and futures price changes should be removed and this can be achieved by adding explanatory variables in the regression equation. A relevant study by Ederington and Salas (2007) examines the bias and efficiency of the OLS hedge ratio when spot price changes are partially predictable. Using the Henry Hub Natural Gas futures contract to hedge gas prices from 17 local gas hubs they find that incorporation of the futures - spot spread as an explanatory variable results in significant improvement in the hedging performance. However, although gas prices are highly seasonal, including seasonal dummies does not result in significantly higher gains¹⁴ since the seasonal component of returns is reflected by the basis. In our study, the mean equation of futures and spot prices also includes the futures – spot spread and seasonality is expected to be reflected in the cointegration relationship. Moreover, we also have to consider that adding parameters increases substantially the computational costs and the log likelihood function might become ill formatted. On the other hand we could use seasonal adjustment methods to overcome this difficulty but evidence suggests that seasonally adjusted data result in longer and shallower regimes. This in turn raises questions regarding regime forecasting. Finally, when a time series is dominated by irregular components, as the oil market does, it is not an easy task to identify and remove the seasonal pattern. Thus in our analysis we restrict ourselves in unadjusted data as this is the standard in the hedging literature.

¹⁴ The predictive power of seasonals was found to be significant only in cases where the futures-spot market was less connected.

Chapter 6

Petroleum Term Structure Dynamics, Inter-Commodity Dependencies and the Role of Regimes.

6.1 Introduction

Several studies in the energy economics literature indicate that oil markets around the world are interrelated and prices move together over time¹. The size of petroleum spreads is most prominently affected by transitory divergences between supply and demand, seasonal factors, transportation costs, convenience yields and the volatility of the underlying (Milonas and Henker, 2001). A shock that might affect a given pair of commodities will most probably have an asymmetric impact not only on each leg of the pair, but also across different maturities of the term structure. It is obvious that such non-parallel relative movements of correlated forward curves are important for the pricing of real assets such as power plants and might also create profitable investment opportunities. However to date, research in commodity futures term structure has primarily focused on the evolution of a single curve (Brennan and Schwartz, 1985; Gibson and Schwartz, 1990; Schwartz, 1997; Schwartz and Smith, 2000; Borovkova, 2006) whereas the issue of co-movement of multiple curves has received less attention despite the multi-asset nature of oil investments. For instance, oil companies involved in the management of physical assets, such as refineries, are mainly concerned with the relationship of crude oil and its distillates to optimise their operations.

Hence, little is known about the joint term structures of different commodities and their implied dependence. Two of the exceptions include Clewlow and Strickland (2000) and

¹ For example, Silvapulle and Moosa (1999) find that oil spot and futures prices react simultaneously to the arrival of new information to the market. Ewing and Harter (2000) provide evidence that Brent blend and Alaska North Slope crude oil prices move together over time and react similarly to shocks in the world oil market. Lin and Tamvakis (2001) investigate the information transmission mechanism between WTI and Brent futures and find that there are price and volatility spillovers between the two markets with the WTI market being the dominant in terms of information discovery.

Tomalsky and Hindanov (2002) on seasonal energy commodities. More recently, Ohana (2010), developed a model for the evolution of correlated forward curves for US natural gas and heating oil, based on the long/short term decomposition of Schwartz and Smith (2000), Manoliu and Tompaidis (2002) and Geman and Nguyen (2005), the study links the literature on correlated forward curves with the concept of cointegration and the results signify causal relations and stochastic volatility among the different shocks, whereas a bi-directional feedback effect is revealed in the formation of the long-term price. Building on this field of research, the present chapter investigates the co-movement and linkages of petroleum futures curves' factors. The objective is to develop a model and provide a new empirical framework not only to characterise the term structure of petroleum spreads but to test their predictive ability as well. The findings of this chapter have important implications and are of interest to oil and commodity traders, oil companies, refineries, and investment funds. For instance, if there are significant price discrepancies in the cross-market futures prices, due to say regional supply and demand imbalances, seasonal factors etc. then these departures should be reflected in the factor dynamics which will, consequently, signal anticipated trends.

This chapter contributes to the existing literature in several ways. First, based on the arbitrage-free evolution of the futures prices under the HJM framework (Heath et. al, 1992) our starting point is to perform Principal Components Analysis (PCA) (as in Cortazar and Schwartz, 1994) to derive sets of latent factors that drive the evolution of the individual forward curves of NYMEX and ICE petroleum futures. PCA is a powerful non-parametric tool that utilises all the available information to derive orthogonal factors that explain term structure fluctuations, eliminating thus the problem of collinearity. This provides an advantage over the approach of Ohana (2010) who proxied long- and short- term factors using arbitrary points on the forward curve by selecting a far futures contract (level) and the differential of a far and a nearby contract (slope).

Next, we introduce for the first time, a flexible multi-regime model of the joint evolution of futures curves factors' dependence. Regime switching models have been used in the energy economics literature in different contexts including studying the conditional volatility (Fong and See, 2002; 2003) or investigating the relationship of crude oil shocks and stock markets behaviour (Aloui and Jammazi, 2009; 2010). Also, in this thesis we have seen in Chapters 4 and 5, applications in risk measurement and hedging effectiveness, respectively. However, they have yet to be applied in studying the relationship between futures curves. By allowing for non-linearities in the term structure generating process, the conditional moments (means, volatilities and correlations) switch stochastically between different states

accommodating the dynamic relationship between pairs of commodity factors. To capture regime shifts in the conditional distributions of the commodities under study, we employ Markov Regime Switching (MRS) models. The rationale for the MRS approach to describe the risk factors lies in the fact that factor-specific dynamic features are typically inherited by asset returns. In the notion of the level-slope-curvature setting, all factors are allowed to switch independently, thus extending Bollen et al. (2000) model to the multivariate case. This way we effectively permit factor specific regimes to demonstrate diversity i.e. one being in the high volatility state and another in a low volatility state and hence, we can disaggregate the regimes as level, slope and curvature driven and study their interaction.

Third, we allow for nonlinear short run causality and mean reversion towards long run equilibrium. The motivation of a dynamic equilibrium correction regime switching model of the factor structure stems from the underlying economics of commodity markets. On the one hand, in a contango market, oil producers build inventories in the expectation of a rise in prices since higher future spot prices would compensate them for the total cost of carrying inventories. On the other hand, demand shocks and tight supplies raise the convenience yield and this will lead to negative sloped forward curves. The risk-return profile of the asset is known to change fundamentally, between these two different states i.e. low inventory levels lead to backwardation, high volatility and reduced correlation in the term structure (Fama and French, 1987; Ng and Pirrong, 1994) resulting concave or convex forward curves since additional units of inventory have uneven effects on different delivery dates. However, despite these short-run deviations supply and demand will eventually move towards a long-run equilibrium level. In a multivariate setting, the same should hold for deviations in the relative supply/demand function between two commodities. Due to refining capacity constraints, supply chain disruptions, seasonality, replenishing use of inventories and timing effect in production, among others, the simultaneous presence of contango and backwardation between the two curves is possible which in turn, may affect the adjustment pattern of the prices. However, these are transitory deviations and all these determinants may have different effects, creating diverse reactions to news and disproportional transmission mechanisms. It is the complex interaction of these mechanisms that we aim to capture with our model.

Finally, we evaluate the forecasting performance of these models using out-of-sample tests, in terms of statistical and risk management loss functions. MRS forecasts are compared to those from alternative models of correlated commodity curves either by modelling the individual contracts or by employing alternative specifications for the factor structure. Regarding the variance-covariance matrices we also use a GARCH Dynamic Conditional

Correlation (DCC) parameterisation, thus providing, robust evidence on the performance of the proposed framework. In addition, we refine the possible gains from using more sophisticated models by testing the statistical significance of the relative performance measures of the competing models employing White's (2000) Reality Check.

The chapter is structured as follows. Section 6.2 introduces multi-factor models of the forward curve dynamics, demonstrates the factor decomposition and derives the properties of the regime switching specification for the factors. In Section 6.3, the data along some preliminary results are discussed. This is followed by the empirical results from the model calibration. Section 6.5 carries out a numerical exercise of forecasting the whole term structure of both prices and risk. Finally, the last section concludes.

6.2 Methodology

Before specifying our model mathematically, for the purpose of clarity, we briefly review the concept of PCA in the framework of forward curves. We then formulate the MRS error correction models applied to test for dependency structures between correlated commodities.

6.2.1 Factor Decomposition

PCA is a procedure for extracting the systematic dynamics of correlated data in the form of orthogonal latent components, whilst making no ad hoc assumptions for their underlying process. After a spectral decomposition of the covariance matrix \mathbf{H} - so that $\mathbf{H} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, where T is the transpose operator, \mathbf{U} is the orthonormal matrix of eigenvectors and $\mathbf{\Lambda}$ a diagonal matrix of ordered eigenvalues - the resultant components are affine combinations of the original features and usually only a few are sufficient to mimic the volatility and correlation structure. The variance of each principal component is maximised so that each one portrays as large a part of the total variance as possible. Let $F(t, T)$ represent the futures daily prices at time t with delivery date T , k be the number of tradable contracts for $m=\{1, \dots, k\}$ and σ_X some volatility functions. We impose the following futures curve dynamics for each individual commodity i :

$$\Delta \ln F(t, T) \approx \frac{\Delta F(t, T)}{F(t - \Delta t, T)} = \sum_{m=1}^k \sigma_{X_m}(t, T, Q) \Delta X_m(t); \quad \Delta \ln F(t, T) \sim IN(0, \mathbf{H}_t) \quad (6.1)$$

$$\sigma_{X_m}(t, T, Q) = U_{X_m}(t - T) \sqrt{\lambda_{X_m}(t, Q)}$$

Applying PCA to the panel of log future price changes will result in k orthogonal factors (ΔX) describing the total variation in futures prices. We will retain the so-called level (ΔL), slope (ΔS), and curvature (ΔC)² risk factors, motivated by the fact that these shocks capture the futures curve dynamics efficiently; for example, Reisman and Zohar (2004) showed that principal components of higher order are relatively unstable. Besides, most term structures can be economically explained by these risk factors; see e.g. Litterman and Scheinkman (1991) for yield curves, Cortazar and Schwartz (1994) for copper, Borovkova (2006) for electricity and Clewlow and Strickland (2000) for crude oil and gas futures.

Notice that in Eq. (6.1), to address seasonal variations in volatility levels evidenced in the data (see Appendix 6.A, for instance), volatility functions are also dependent on the season, Q , as follows. Let \mathbf{H} to depend only on the season and the commodity and all \mathbf{H}_Q 's to share a unique eigenvector $k \times k$ matrix \mathbf{U} which jointly diagonalises \mathbf{H} . In practice, to find the volatility functions we will use the PCA on the unconditional covariance matrix \mathbf{H}^* of the standardised – by seasonal volatilities - historical returns. If we denote \mathbf{D}_Q the diagonal matrix containing the standardised seasonal volatilities of the observed variables and \mathbf{R} the unconditional correlation, each \mathbf{H}_Q can be decomposed to $\mathbf{D}_Q \mathbf{R} \mathbf{D}_Q$. Principal components ΔX are then the weighted average of the (seasonally) standardised price changes with weights given by \mathbf{U} , whereas the variance of the principal components i.e. eigenvalue λ , becomes the product of the standardised eigenvalues λ^* and the corresponding seasonal volatility of each contract contained in \mathbf{D}_Q . The k^{th} column of \mathbf{U} corresponds to the k^{th} eigenvalue (factor variance) where the latter are sorted as $\lambda_1(Q) > \dots > \lambda_k(Q)$, for each subperiod Q . The main advantage of using this formulation is that we can characterise the evolution of futures prices in a realistic fashion by considering the full set of historical data and capture the average variability throughout the sample period across

²Note that, to explain all the variance in the sample, all k principal components must be used. In any other case, where the dimension of the vector of latent factors is $i < k$ a $T \times k$ idiosyncratic component must be added, representing risk factors that have not been incorporated in the system. This implies that factor communalities (in other words the R^2 of a regression of the original series on the principal components) are less than one.

seasons³. The seasonal behaviour of energies has been documented by Girma and Paulson (1998), Tomalsky and Hindanov (2002) and Borovkova and Geman (2006), among others. For example, heating oil, experiences an upward (downward) pressure during winter (summer), whereas the storage cycle might not be able to absorb seasonal demand shocks, especially late in the peak demand season. Thus, a higher volatility is anticipated in winter, as opposed to the inventory build up period during summer. Crude oil demand, on the other hand, is derived by its products and their individual features might induce complex spill over effects. Although controls for seasonality could easily accommodate other frequencies, quarters were chosen for parsimony.

Moreover, factor variances are also assumed to evolve through time; therefore, \mathbf{H}_t also changes as a function of $\mathbf{U}\mathbf{\Lambda}_t\mathbf{U}^T$. Note that the time varying nature of the eigenvalues $\mathbf{\Lambda}_t$'s, or equivalently the second moments of the principal components, can be modelled independently of the PCA, following the estimation of the factor loadings – eigenvectors; hence, their time-dependent parametric form will be common to all futures series' volatilities. The core of this idea has been advocated in several forms. For instance, in the asset pricing framework, Engle et al (1990), adopt a two-step method in which static factors are extracted from the unconditional covariance matrix before being modelled as univariate ARCH processes. They note that assuming constant eigenvector and time-varying eigenvalue structure is a statistically convenient yet reasonable assumption which essentially implies constant relative riskiness and varying total riskiness. In short, this factor approach has a substantive motivation that produces a realistic variance-covariance structure.

To accommodate now two future curves, denote $\mathbf{\Lambda}_{12,t}$ the cross-commodity covariance matrix of the risk factors for commodities 1 and 2 and $\mathbf{H}_{12,t}$ the square matrix containing the cross-commodity covariances of futures returns both depending on time (as well as season; however, ignore seasonal parameterisation for notational convenience). Then, given that $\Delta \ln F_i(t,T) \sim \text{IN}(0, \mathbf{H}_{i,t})$ for $i=\{1,2\}$ the $(2k \times 2k)$ full covariance matrix of the original system of correlated factors \mathbf{V}_t (see also Alexander, 2008, vol. II, pp. 179-180) can be specified as:

³ Therefore in this chapter of the thesis, we standardised the returns by quarterly volatilities. An alternative approach would be to let both eigenvalues and eigenvectors to be seasonally dependent by performing PCA on seasonal blocks of futures returns; however, this specification would be less parsimonious compared to the one used. In addition some preliminary results showed that factor loadings were less stable in this case, indicating overfitting and the presence of noise in the factors; also, in this case, seasonality effects were not reduced to the same extent as was evidenced by the autocorrelation functions.

$$\begin{pmatrix} \Delta \ln F_1(t, T) \\ \Delta \ln F_2(t, T) \end{pmatrix}^T \sim IN(0, \mathbf{V}_t); \quad \mathbf{V}_t = \begin{pmatrix} \mathbf{H}_{1,t} & \mathbf{H}_{12,t} \\ \mathbf{H}_{12,t}^T & \mathbf{H}_{2,t} \end{pmatrix} = \begin{pmatrix} \mathbf{U}_1 \mathbf{\Lambda}_{1,t} \mathbf{U}_1^T & \mathbf{U}_1 \mathbf{\Lambda}_{12,t} \mathbf{U}_2^T \\ (\mathbf{U}_1 \mathbf{\Lambda}_{12,t} \mathbf{U}_2^T)^T & \mathbf{U}_2 \mathbf{\Lambda}_{2,t} \mathbf{U}_2^T \end{pmatrix} \quad (6.2)$$

6.2.2 Modelling the Information in the Term Structure

Once the shocks that determine futures prices fluctuations have been retrieved, we investigate the co-movement and volatilities-correlation structure of cross-commodity underlying factors. The discussion in this section suggests another perspective, unexplored by the present literature, on the linkages between forward curves that are linked by fundamental economic relationships such as the crack spread. The concept of cointegrated forward curves has been introduced by Ohana (2010) for heating oil and natural gas futures. What differentiates our approach - apart from the factor construction and the inclusion of the curvature factor - is that, first, we attempt to shed light on the relative short- and long-run dynamics across different regimes and second, our regime switching formulation allows the coefficients of short run causality and error correction mechanism to be time-varying i.e. state dependent.

Let $\Delta \mathbf{X}_t$ represent the $t \times 2$ vector of pairs of factor shocks (i.e. level, slope or curvature) across two commodities, $\mathbf{\Gamma}_{i,st}$ and $\mathbf{\Pi}_{st}$ the state dependent 2×2 coefficient matrices measuring, respectively, the short- and long-run adjustment of the system to changes in a specific factor \mathbf{X}_t and $\boldsymbol{\varepsilon}_{t,st}$ a vector of Gaussian white noise processes with state dependent covariance matrix $\boldsymbol{\Sigma}_{st}$. We employ the following MRS Vector Error Correction Model (VECM):

$$\Delta \mathbf{X}_t = \mathbf{v}_{st} + \mathbf{\Pi}_{st} \mathbf{X}_{t-1} + \sum_{i=1}^p \mathbf{\Gamma}_{i,st} \Delta \mathbf{X}_{t-i} + \boldsymbol{\varepsilon}_{t,st}; \quad \boldsymbol{\varepsilon}_{st,t} | \Omega_{t-1} \sim IN(0, \boldsymbol{\Sigma}_{st}) \quad (6.3)$$

Throughout the chapter, the focus is on the long-term equilibrium relationships across the same factors (e.g. heating oil level- crude oil level) as these exhibit the stronger linkages. Moreover, although orthogonality holds only unconditionally, we do not consider conditional correlation dynamics among orthogonal factors and/or cross-commodity cross-factors in order to maintain applicability and avoid more complex structures. Besides, the proposed decomposition is very versatile and, as it will be shown, it is possible to blend the MRS models for the levels, slopes and curvatures, making our MRS framework capable of handling multiple models and dealing with heterogeneity in preferences. For example, a market agent might want to consider only one risk factor whereas a more risk-averse agent might prefer to include more.

Therefore, the proposed model provides, in our view, the most parsimonious representation of correlated forward curves.

The following steps are involved in our analysis. First, following Krolzig (1999) and assuming a single regime process, the existence of a stationary relationship between the cross-commodity factors is examined, through the λ_{\max} and λ_{trace} statistics (Johansen, 1988) which test for the rank of $\mathbf{\Pi}$, once each factor's path has been derived by setting some arbitrary initial value $X_{t_0=0}$, as $X_t = X_{t_0=0} + \sum_{t_0}^t \Delta X_t = X_{t-1} + \Delta X_t$. If $\mathbf{\Pi}$ has a reduced rank, that is $\text{rank}(\mathbf{\Pi}) = 1$, then there exists one cointegrating vector and $\mathbf{\Pi}$ can be decomposed to $\mathbf{\alpha}\mathbf{\beta}'$, where $\mathbf{\beta}$ represents the vector of cointegrating parameters and $\mathbf{\alpha}$ the vector of error correction coefficients measuring the speed of convergence to the long run mean. However, the VECM of Equation (6.3) provides a framework for valid inference only in the presence of $I(1)$ variables. Should $I(0)$ processes for the factors emerge, we adjust $\mathbf{\Pi}$ to a diagonal matrix $\mathbf{\Pi}^* = \text{diag}(\alpha_1, \alpha_2)$ representing the univariate mean reversion rates. Thus, the model reduces to a modified Vector Autoregression (VAR) with a different equilibrium mechanism for each process. In this case, the Augmented Dickey Fuller (1979, 1981) regression is employed, further augmented by cross commodity factor lags, to allow for causality feedback.

The second step involves the introduction of Markovian regime shifts to the system. Regime switching is allowed in all coefficient matrices \mathbf{v}_{st} , $\mathbf{\Gamma}_{i,st}$, $\mathbf{\Pi}_{st} = \mathbf{\alpha}_{st}\mathbf{\beta}'$ as well as the second moments $\mathbf{\Sigma}_{st}$. Extending the Bollen et al. (2000) specification⁴ to the bi-variate case, we permit an independent two-state first order MRS process for each factor. This way, although each equation of the system (Eq. 6.3) follows a two-state *self-directed* process, the joint system can be characterised by a four-state, first order *combined* Markov process with constant transition probabilities p_{ij} , implied by the individual probability matrices. To illustrate this, consider the transition probability matrix \mathbf{P} for each element X_i of factor \mathbf{X} (X_i being either L or S or C, for the respective commodities) which is given as:

$$\mathbf{P} = \begin{pmatrix} p_{11} & 1 - p_{11} \\ 1 - p_{22} & p_{22} \end{pmatrix}; \quad p_{ij} = \Pr(s_{t+1} = j | s_t = i) \quad (6.4)$$

⁴ Bollen et al. (2000) define a four regime univariate model with two independent state processes (thus, limiting the number of the specified parameters compared to the unrestricted four regime case): one for the mean and one for the variance. In our case, however, we allow each process of the system to be governed by high-low volatility regimes which occur independently from the states of the second process.

where $st = \{1, 2\}$, p_{11} and p_{22} give the probability that state st will remain the same in the following period and $p_{12}=1-p_{11}$ and $p_{21}=1-p_{22}$ give the probability that state 1 will be followed by state 2 and 2 by 1, respectively. Consequently, two local factor-specific regimes are defined for each pair of commodities, implying that each factor follows an independent two-state process, motivated by the fact that specifying periods of low and high volatility is intuitively appealing. To account for the possibility that factor specific regimes across commodities might display diversity i.e. one being in the high volatility state and another in a low volatility state, we link these two processes by estimating the joint transition probability matrix Ψ based on the transition matrices of the two factors, say, \mathbf{P} and \mathbf{Q} , resulting the following four state process:

$$\Psi = \mathbf{P} \otimes \mathbf{Q} = \begin{pmatrix} p_{11}q_{11} & (1-p_{11})q_{11} & p_{11}(1-q_{11}) & (1-p_{11})(1-q_{11}) \\ (1-p_{22})q_{11} & p_{22}q_{11} & (1-p_{22})(1-q_{11}) & p_{22}(1-q_{11}) \\ p_{11}(1-q_{22}) & (1-p_{11})(1-q_{22}) & p_{11}q_{22} & (1-p_{11})q_{22} \\ (1-p_{22})(1-q_{22}) & p_{22}(1-q_{22}) & (1-p_{22})q_{22} & p_{22}q_{22} \end{pmatrix} \quad (6.5)$$

where \otimes denotes the Kronecker product. This way we obtain the $2^2 \times 2^2$ Ψ_L , Ψ_S and Ψ_C matrices, corresponding to the level, slope and curvature set of factors, respectively. The elements in the off-diagonal of the transition matrix Ψ denote the probabilities of a regime switch, while the elements in the main diagonal reflect the probability that the same state will be maintained. Focusing on the main diagonal of Ψ in Eq. (6.5), the upper left and lower right element of Ψ i.e. $\Psi(1,1)$ and $\Psi(4,4)$, show the probability that both commodity factors are jointly in state 1 and 2, respectively; for clarification define state 1 (2) as the local factor-specific low (high) variance state. Similarly, element $\Psi(2,2)$ shows the probability of the first factor being in the high and the second being in the low variance state and $\Psi(3,3)$ the probability of the two factors being in the low and high variance states, respectively. Hence the new constructed combined regimes are four: low-low, high-low, low-high and high-high variance states. Next, given that the pairs of the log factor changes evolve according to the process defined in Eq. (6.3), for each pair of commodity factors the density function for each regime is used to construct the likelihood function $f(X_t; \theta)$. This can be formed as a mixture of the probability distribution of the state variables, with θ being the vector of parameters to be estimated, including Ψ . The weights of the mixture of the distributions are the conditional regime probabilities which are estimated recursively along with the likelihood function as shown in Hamilton (1994) and Gray (1996), using the Markov property and Bayes rule as (see also Chapter 3, section 3.4.1):

$$\boldsymbol{\pi}_{t|t} = \boldsymbol{\Psi}^T \left[\frac{\boldsymbol{\pi}_{t|t-1} \odot \mathbf{f}_{t-1}}{\mathbf{1}^T (\boldsymbol{\pi}_{t|t-1} \odot \mathbf{f}_{t-1})} \right] \quad (6.6)$$

$$\boldsymbol{\pi}_{t|t-1} = \boldsymbol{\Psi}^T \boldsymbol{\pi}_{t-1|t-1}$$

where $\boldsymbol{\pi}_{t|t}$ and $\boldsymbol{\pi}_{t|t-1}$ denote vectors containing the probabilities of being in each regime at time t conditional on the observations up to time t (filtered) and up to time $t-1$ (ex-ante), respectively, \mathbf{f}_{t-1} is a vector of state dependent densities conditional on the observations up to time $t-1$ and \odot symbolises the element-by-element multiplication. Then, given Eq. (6.3) and subject to the constraints $\pi_{1,t} + \pi_{2,t} + \pi_{3,t} + \pi_{4,t} = 1$ and $0 \leq \pi_{1,t}, \pi_{2,t}, \pi_{3,t}, \pi_{4,t} \leq 1$ - where $\pi_{st,t}$ are elements of the ex-ante $\boldsymbol{\pi}_{t|t-1}$ probability matrix with $\pi_{st,t} = \Pr(st=i|\Omega_{t-1})$ - iterating the expressions in Eq. (6.6), the log-likelihood function $L(\boldsymbol{\theta})$ to be maximised using numerical optimisation methods is:

$$f(\mathbf{X}_t; \boldsymbol{\theta}) = \sum_{i=1}^4 \frac{\pi_{i,t}}{2\pi} |\boldsymbol{\Sigma}_i|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \boldsymbol{\varepsilon}'_{t,i} \boldsymbol{\Sigma}_i^{-1} \boldsymbol{\varepsilon}_{t,i}\right) \quad (6.7)$$

$$L(\boldsymbol{\theta}) = \sum_{t=1}^T \log f(\mathbf{X}_t; \boldsymbol{\theta})$$

Finally, following the estimation of the MRS models, global joint probabilities that are commodity pair-specific (e.g. heating oil crack spread) can also be constructed in a similar way to Eq. (6.5) as $\boldsymbol{\Psi}^* = \boldsymbol{\Psi}_L \otimes \boldsymbol{\Psi}_S \otimes \boldsymbol{\Psi}_C$. This results in a versatile multi-regime system that enhances the pertinence of our model; although common forces drive oil market regimes, we do not lose valuable information such as correlation break downs that arise from the possibility that pairs of commodities might not be at the same state e.g. one might be in backwardation and the other in contango. For instance, in Chapter 5, we have seen that in the unleaded gasoline market, the low variance regime was less persistent compared to WTI and this is attributed to the fact that light distillates volatility is vulnerable to the quality of the crude, constrained refining capacity and prone to frequent backwardations and supply shortages.

6.3 Data Description and Preliminary Analysis

The data set for this study comprises daily closing prices for NYMEX WTI (CL), NYMEX heating (HO), ICE Brent (CB) and ICE gas oil (GO) futures, from June 27, 1994 to December 31, 2009. All prices are obtained from Datastream. We use constant maturity futures, constructed by linear interpolation from the market prices of traded contracts⁵. Closer examination of volume and open interest data lead us to consider a block of 10 contracts for each commodity from 1 up to 10 months to maturity⁶. Our choice of constant maturity contracts ensures that all prices are measured at the same point in time and we avoid problems associated with thin trading and expiration effects that might complicate inference regarding the volatility functions (Eq. 6.1). This way, therefore, we deal with three main concerns that can potentially cause estimation issues. First, we avoid discontinuities arising from the limited life span of individual contracts. Second, we mitigate the problems that nonstationary volatilities impose which may be due to increased demand for offsetting positions as contracts approach maturity or rolling futures forward to prevent delivery (Samuelson, 1965). And finally, we effectively address the issue of futures-spot convergence at expiry; for instance, a continuous futures series in a backwardated market will inevitably experience a downward trend near delivery which might distort the results.

Results of the PCA on the correlation matrix of the normalised by seasonal volatilities futures returns (see also section 6.2.1 and Eq. 6.1) of each individual futures curve are presented in Panel A of Table 6.1. Three factors are adequate to explain more than 99% proportion of prices fluctuations. Because the purpose of the chapter is to describe commodity interdependencies we also present the results after combining the factors to explain petroleum market spreads rather than individual petroleum term structures; we can observe that the explanatory power of three factors now is less. What merits attention is that the importance of the first factor diminishes (with a lower bound of 77% in the WTI-Brent spread) as opposed to an increase in the importance of the second and third factors (with an upper limit of 12.4% for the slope and 3.2% for the curvature, both occurring for the WTI-Brent spread). The relative contribution of each factor across seasons is also presented.

⁵ Prices of constant maturity futures are calculated by averaging near and distant contracts, weighted according to their respective days from maturity. This way, we obtain an actual price for exactly x days prior to expiry; in fact, the resulting series is a “perpetual” contract with constant rolling delivery date (see Pelletier, 1983).

⁶ 1 contract of crude oil is 1,000 barrels whereas those of heating and gas oil are 42,000 US gallons and 100 tonnes, respectively. A barrel is equivalent to 42 gallons or 7.45 (in line with ICE calculations) tonnes. Prices are converted to \$US/bbl and transformed to natural logarithms for the ensuing analysis.

For all commodities, the level has the lower explanatory power during winter, whereas the importance of slope and curvature factors increases that same period. Regarding the spreads, although no clear-cut results are obtained, there are larger variations in the relative contribution values especially in the crack spread series. For instance in the HO-CL (GO-CB) spread, the contribution of the level factor is maximised during summer (spring) and minimised during winter (summer) with a difference of around 1,000 (800) basis points. Overall, these results indicate that the performance of PCA differs across seasons thus, further justifying the use of seasonal volatilities.




Figure 6.1 Seasonality Adjusted Weights of Principal Components for Heating Oil & WTI Crude Oil.

This can also be confirmed by looking at the factor loadings of NYMEX heating and crude oil for the winter and summer seasons, presented in Figure 6.1. It is clear that in winter months there is an upward shift in the level, a clockwise rotation in the slope and an increased convexity in the curvature. This indicates that a shock of the same magnitude has greater impact on prices during winter, illustrated also in Panel B of Table 6.1 which represents the annualised unconditional futures returns volatilities by season across maturity. The seasonal pattern appears fairly consistent and the amplitude of volatility varies across quarters for all commodities. Summer months (June-August) are associated with relatively lower volatility whereas the latter reaches its peak either in autumn (September-November) or winter (December-February). Seasonal demand and storage might be some of the reasons; various studies have used different approaches to remove this seasonal behaviour such as Clewlow and Strickland (2000) and Borovkova and Geman (2006), just to mention a few. For comparison we also perform PCA on the panel of futures prices without seasonal controls. In terms of the proportion of variance explained results are qualitatively similar, yet, the seasonal behaviour of the futures prices is markedly evident in the factor process, thus justifying the use of the approach presented in the chapter. Finally, Panel B, Table 6.1 shows that the volatility of shorter term contracts exceeds

that of the more distant contracts, thus, shocks in the market are expected to spread out gradually (Samuelson effect) along the futures curve.

Descriptive statistics of the principal components log-returns are reported in Table 6.1, Panel C. Annualised unconditional volatilities for the crude oil level factor are slightly higher than those of the corresponding petroleum product - within each market - whereas the opposite is observed for the slope and curvature. The former are in the range of 49% p.a., the latter 6.5% and 2.5% p.a., respectively, whereas, ICE short run factors appear more volatile than NYMEX. Note that these annualised volatilities are actually the eigenvalues λ (elements of diagonal matrix Λ) as measured over the whole sample period; using Eq. (6.1) and Eq. (6.2), the whole covariance matrix of the actual futures contracts can be replicated since the factor log-returns are just linear portfolios of the futures log returns. As their variance decreases so does their importance in explaining futures curve movements. The Ljung-Box (1978) Q statistic on the first five and twenty lags of the sample autocorrelation function is significant in all cases, with the exception of gas oil slope. Engle's (1982) ARCH test, carried out as the Ljung-Box Q statistic on the squared series, indicates the existence of time-varying heteroscedasticity.

The coefficients of skewness and excess kurtosis indicate departures from normality which is also confirmed by the Jarque-Bera (1980) test. In particular, the coefficient of skewness is negative for all level factors which means that long positions involve greater risk since large negative level shocks are more likely than positive ones. Regarding the slope factors negative shocks that will essentially push short term prices down and long term prices up are more likely in all cases apart from WTI. Similarly, regarding the curvature factors positive curvature shocks that will drive both short- and long- term prices up and medium term prices down are more likely in the NYMEX market, while in the ICE market, curvature shocks are balanced in terms of skewness, at the 5% significance level.

6.3.1 Unit Root and Co-integration Results

To determine the order of integration of petroleum principal components, we perform Augmented Dickey-Fuller (1979, 1981) and Phillips and Perron (1988) non-parametric unit root tests. Results (Panel A, Table 6.2) show that both level and slope series follow unit root processes while their first differences are stationary, rendering support for the use of VECM to capture the short run dynamics and long run trends. In contrast, all curvature factors are $I(0)$ variables, hence the use of VAR is more appropriate. Figure 6.2 plots the estimated WTI and heating oil sets of factors. It seems that while they drift apart in the short run, factors move

together in the long run, pair-wise. Long run co-movement can be attributed to common driving forces, such as the prevailing global oil market conditions, while differences in the short run dynamics to temporary supply/demand imbalances caused by seasonality, refining capacity constraints etc. The mean reverting behaviour of the curvature factor is also obvious.




Figure 6.2 Estimated Level (top), Slope (2nd, middle) and Curvature (3rd, bottom) factors prices for WTI crude oil and Heating oil

Table 6.2: Unit Root & Johansen Cointegration tests for Petroleum Futures Factors

Panel A: Unit Root Tests

	Augmented Dickey Fuller			Phillips Perron		
	Level	Slope	Curvature	Level	Slope	Curvature
CL	-1.923	-2.185	-4.491***	-1.903	-2.083	-3.172**
HO	-1.732	-1.927	-4.955***	-1.733	-1.847	-4.375***
CB	-1.803	-1.996	-3.460***	-1.818	-1.906	-6.300***
GO	-1.651	-1.696	-3.317**	-1.595	-1.715	-5.151***
ΔCL	-65.31***	-54.11***	-56.94***	-65.43***	-54.02***	-56.80***
ΔHO	-66.23***	-55.14***	-23.87***	-66.31***	-54.97***	-61.93***
ΔCB	-66.19***	-62.61***	-37.95***	-66.23***	-62.70***	-97.43***
ΔGO	-62.81***	-61.59***	-50.15***	-62.85***	-61.58***	-70.68***

Panel B: Johansen Cointegration Tests & Mean Reversion Rates

	Lags	H ₀ :	λ Statistics		Adjust. Coefficients		CV: (1 β ₁ β ₀)	Rolling λ Statistics	
			λ _{max}	λ _{trace}	a ₁	a ₂	Normalised	λ _{max} test	λ _{trace} test
Level Factors Pair wise Cointegration									
HO-CL	1	r=0	32.51***	36.08***	-1.2984	0.2597	(1 -1.107 -0.089)	[18.9 27.2]	[19.7 33.4]
		r=1	3.578	3.578	(0.679)*	(0.782)		{0.1%}	{2.1%}
GO-CB	4	r=0	24.95***	27.78***	-1.6328	0.7522	(1 -1.195 -0.071)	[18.4 26.1]	[19.9 29.1]
		r=1	2.833	2.833	(0.612)***	(0.679)		{0.8%}	{2.2%}
CL-CB	2	r=0	43.45***	47.17***	0.5988	3.6512	(1 -0.950 0.018)	[20.3 26.1]	[21.7 29.3]
		r=1	3.721	3.721	(1.617)	(1.603)**		{0.0%}	{0.0%}
HO-GO	3	r=0	55.01***	57.77***	-2.4590	5.5925	(1 -0.884 -0.007)	[27.7 65.3]	[29.6 66.6]
		r=1	2.758	2.758	(1.285)*	(1.168)***		{0.0%}	{0.0%}
Slope Factors Pair-wise Cointegration									
HO-CL	1	r=0	37.08***	40.34***	-0.8513	0.8933	(1 -1.148 -0.016)	[21.8 26.9]	[23.3 32.3]
		r=1	3.259	3.259	(0.267)***	(0.263)***		{0.0%}	{0.0%}
GO-CB	1	r=0	52.72***	55.50***	-1.7678	0.9514	(1 -1.266 -0.004)	[23.8 32.8]	[26.3 34.6]
		r=1	2.779	2.779	(0.333)***	(0.318)***		{0.0%}	{0.0%}
CL-CB	2	r=0	36.37***	40.24***	-1.2101	1.4913	(1 -0.995 0.007)	[16.9 26.7]	[18.7 31.3]
		r=1	3.875	3.875	(0.539)**	(0.545)***		{6.3%}	{2.9%}
HO-GO	4	r=0	39.37***	42.23***	-1.6184	1.8876	(1 -0.903 -0.004)	[21.6 30.8]	[22.8 34.9]
		r=1	2.866	2.866	(0.518)***	(0.530)***		{0.0%}	{0.0%}

Lags is the lag length of the unrestricted VAR model in levels chosen on the basis of Schwarz Information Criterion (1978); λ_{max} tests the null hypothesis of *r* cointegrating vectors against the alternative of *r*+1. λ_{trace} tests the null hypothesis that there are at most *r* cointegrating vectors against the alternative that the number of cointegrating vectors is greater than *r*. Critical values obtained from Osterwald-Lenum (1992); Figures in (·) are standard errors, which are calculated using a Newey-West (1987) correction for serial correlation and heteroscedasticity; Asterisks *, ** and *** indicate significance at 10%, 5% and 1%, respectively; β' = (1 β₁ β₀) are the coefficient estimates of the cointegrating vector where the coefficient of X_{1,t-1} is normalised to be unity, β₁ is the coefficient of X_{2,t-1} and β₀ is the intercept term; The rolling cointegration tests are conducted by applying the Johansen multivariate approach to rolling daily ten-year sub-samples. The first trace test statistic is obtained by using observations from the beginning of the sample period through to the 2530th observation. The next test statistic is obtained by using data from the second observation through to the 2531st observation, and so on, until the last observation was used; Numbers in [·] correspond to the 10 % confidence interval of the λ_{max} and λ_{trace} statistics throughout the rolling period; Numbers in {·} correspond to the % of times that the null hypothesis of no cointegration cannot be rejected at 10% significance level using critical values obtained from Osterwald-Lenum (1992) i.e. 13.75 and 17.85 for the λ_{max} and λ_{trace} statistics, respectively.

Pair wise Johansen (1988) cointegration tests for the levels and slopes (Panel B, Table 6.2), assuming a single regime process, indicate that all stand in a long-run relationship. Therefore, these pairs evolve in close proximity to one and other and any deviations signal disequilibria, which will eventually be restored. Note that correct specification of the deterministic components in the VECM is important because the asymptotic distributions of the cointegration test statistics are dependent upon the presence of trends and/or constants. In our case likelihood ratio tests indicated that the intercept term should be restricted to lie on the cointegrating space - in Eq (3) for a single regime process *v* is a 2x1 zero matrix- hence, the vector series becomes $X_{t-1} = (X_{1,t} \ X_{2,t} \ 1)$, with a cointegrating vector $(1 \ \beta_1 \ \beta_0)$, where the coefficient of X_{1,t-1} is normalised to be unity, β₀ is the intercept term, and β₁ is the coefficient on

$X_{2,t}$. The normalised coefficient estimates of the cointegrating vector $(1 \beta_1 \beta_0)$ represent the long-run relationship which can be regarded as a spread e.g. $L_{1,t} - \beta_1 L_{2,t} - \beta_0$. All the error coefficients have the correct sign (negative for the first leg of the spread and positive for the second) implying that in response to a positive deviation from their long-run mean at period $t-1$, i.e. when $L_{1,t-1} - \beta_1 L_{2,t-1} > \beta_0$, $L_{1,t}$ will decrease and $L_{2,t}$ will increase the following period to restore balance. This adjustment process is not uniform across factors. For the level factors it is primarily driven by the refined products, while for the slopes both crude and refined products move in response to disequilibrium. For instance, looking at the level factor for the crude - refined product pairs, equilibrium is restored following the adjustment of petroleum products. In the inter-crude market, WTI is non responsive to the differential. This is expected since the US reflects by far the largest oil consumer and importer of crude oil and this introduces a high degree of sensitivity to the US oil prices, which perhaps makes the WTI market dominant in terms of information discovery (Lin and Tamvakis, 2001). In the inter-product market the estimates of the error correction coefficients, in terms of magnitude and significance, indicate that both heating and gas oil prices move to adjust equilibrium for the both level and slope equations, at 10% significance level. Overall, regarding the slope factors, there is a two-way feedback relationship in all cases, at 5% significance level.

Finally, to discount the possibility that convergence is sample specific we use rolling cointegration tests that explicitly take into account the possibility that two or more series may be more integrated during some periods but less so or not at all during other periods. These tests are conducted by applying the Johansen (1988) procedure to rolling ten-year sub-samples (i.e. using a moving window of 2,530 daily observations). The 90% confidence intervals from the rolling cointegration tests are reported in the last two columns of Panel B, Table 6.2 along with an associated failure rate i.e. the percentage of sub-samples that the null hypothesis of cointegration is rejected. Overall, cointegration is confirmed over the sub-samples and the highest rejection percentage occurs in the WTI-Brent slopes spread (=6.3%), where only 78 sub-samples out of 1,250 are not cointegrated.

6.4 Empirical Results

Having identified the cointegration and mean reversion properties of the data, MRS Models are employed next to investigate the dynamic relationships between the petroleum futures factors, as described in section 6.2.2 (Eq. 6.3, 6.4 and 6.5). Given the results of the previous section for the level and slope shocks, an MRS-VECM is specified whereas curvature

shocks are linked through a VAR-X (in first differences) with each equation augmented by the level of the dependent variable, setting Π (Eq. 6.3) equal to a diagonal matrix Π^* . Individual shocks follow an independent two-state process by disaggregating each data generating process to periods of low and high volatility. In order to accommodate the possibility that one factor might be in the high volatility state and the other in a low volatility state and thus get a more realistic representation of the correlation structure, we link these two processes by maximising the joint log-likelihood function of Eq. (6.7); log-likelihood is maximised each time by taking the pairs of same factors (L₁-L₂, S₁-S₂ and C₁-C₂). In the ensuing analysis four cases are considered; the NYMEX heating crack (HO-CL), the ICE gas oil crack (GO-CB), the inter-crude (WTI-Brent) and the inter-product (HO-GO) spreads. Results are presented in Table 6.3.

Several points are worth noting. First, looking at the estimated regime switching intercepts across the two regimes results are mixed but the coefficients display asymmetries in several ways. For instance, when considering gas oil slope shocks in the inter-product spread equations this asymmetry is manifested as a sign change; another example is the WTI crude oil level shocks in the inter-crude equations where the intercept becomes much higher in the high variance state. Overall, intercepts are larger in the high variance state in terms of absolute values (19 out of 24 cases). Note that for the MRS models of Eq. (6.3) we modify the cointegrating vector to $(I \ \beta_l)$ in order to allow switching in the equilibrium means; removing β_0 from the cointegrating vector, this is now incorporated in the system vector of intercepts v_{st} so that the results in Table 6.3 display the aggregate switching. This way, intercepts depend on both the drifts and the equilibrium mean of the system; shifts in the intercept term of the system can be decomposed into changes in the drift $E[\Delta X_t | st] = \delta_{st}$ and equilibrium mean $E[\beta' X_t | st] = \mu_{st}$ as $v_{st} = \delta_{st} + a_{st} \mu_{st}$, where $\delta_{st} = \beta_{\perp} (\alpha_{st}' \perp \beta_{\perp})^{-1} \alpha_{st}' \perp v_{st}$ and $\mu_{st} = - (\beta' \alpha_{st})^{-1} [\beta' v_{st}]$, \perp denoting the orthogonal complement. Given the above property, it is clear that for the $I(1)$ series of levels and slopes (see Table 6.2), intercepts depend on all endogenous variables because they share a system-specific common long run mean; therefore, it is not unusual to observe diversity across same-commodity shocks' intercepts of different systems i.e. the gas oil level in the ICE crack MRS-VECM implies a low (high) volatility state intercept of 0.042 (0.184) whereas for the inter-market (HO-GO) equation the corresponding figure is 0.082 (-0.023). Such variation across same-commodity factors does not occur for curvatures because intercepts depend only on the individual corresponding shock, that is, $\alpha_{1,st} [C_{1,t} - \mu_{1,st}]$ and $\alpha_{2,st} [C_{2,t} - \mu_{2,st}]$; the equilibrium mean is not common because C 's are already $I(0)$ processes.

Turning next to the estimated speeds of adjustment a_{st} to the long-run equilibrium level, overall, level and slope error correction coefficients' signs are consistent with the corresponding single regime estimates (α_1 and α_2), given in Table 6.2. Moreover, they are in line with theory having the correct sign across both regimes, in the sense that the first is negative, the second positive and at least one of them is statistically significant, consistent with convergence towards the long-run equilibrium. An exception is the WTI-Brent equation for the levels where both error correction coefficients are positive in the high variance state but only the coefficient for Brent is statistically significant. It is also interesting to note that in the high volatility state for both crack spread markets, crude oil leads the information discovery process; this can be explained by the fact that first, crude oil prices are determined by the worldwide supply and demand and constitute a much larger market as opposed to refined products where regional supply/demand dynamics are important and second, crude is the single most important production cost for those products affecting their price formation accordingly (see also Alizadeh and Nomikos, 2008). However, although for the NYMEX crack this holds globally - in both low and high variance states - this is not the case for the ICE crack levels since under the low volatility state gas oil leads Brent; this can in turn be attributed to the fact that crude oil demand is derived from refined products. Note though that this apparent contradiction may be explained considering that Europe's gasoil market relative to crude is much more important than the heating market in the US. For instance, the relative volumes of futures contracts of crude/heating in NYMEX was around 5:1 as opposed to less than 2:1 in ICE for crude/gasoil. In the inter-product market both commodities constitute goods with seasonal and capacity constrained flow of supply and are vulnerable to supply disruptions. As such, in the high variance state, a two way feedback effect is observed. Nevertheless, heating oil is not responsive to the differential in the low variance state; perhaps this reflects the transient nature of temporary imbalances in the European market, which are more vulnerable to extreme weather conditions (Milonas and Henker, 2001), the faster response of heating to the larger market of crude so that gasoil seems to follow the discovery process at a slower pace, or even different inventory policies in a way that the US efficiently accumulates stocks to deal with temporary demand shocks.

For the slope factors, error correction coefficients are all larger in the high variance state implying that high volatility is associated with faster reversion to the long-run equilibrium. Here we can see a more clear pattern compared to the Level factors. First, in the low volatility state, a one way (long-run) feedback effect is detected with the equilibrium relationship having explanatory power only on refined products and not on the crudes (for the cracks) and the US

market leading the European market. Second, in the high volatility regime, a two way feedback effect is reported across all cases. As for the curvature factors, all mean reversion rates are significant - apart from the heating oil high volatility state process in both NYMEX crack and inter-product equations - and, overall, high volatility state is associated with faster reversion to the long-run factor-specific (rather than common as for levels and slopes) mean. Finally, short run causality tests, carried out as a likelihood ratio (LR) statistic on the lagged cross-factor terms indicate that in almost all cases there is a two way feedback in the short run dynamics. Notable exceptions are the slope shocks where crude oil (in case of the cracks) and heating oil (in case of the inter-product spread) lead the markets in the short run.

Table 6.3 also reports annualised volatilities and correlation across regimes. First, regime dependent volatilities from the estimated models indicate that there is marked asymmetry in all cases. Regarding the level factors, volatilities of the high variance state are overall 40%-70% higher compared to the low variance state, apart from the WTI-Brent equation where the percentage increase drops to 10%. The corresponding figures for slopes and curvatures are 95%-120% and 140%-200%, respectively. Therefore, each regime clearly differentiates two distinct market dynamics for the volatility of the underlying process. As for the state dependent correlations, it seems that when one market is in the low and the other in the high variance state, it is more likely to observe, on average, lower correlations (10 out of 12 cases). This holds even when we compare the average correlations of high-low and low-high variance states with the low-low and high-high variance states individually (both times 10 out of 12 cases) whereas high correlations seem to be a feature of the low variance state (9 out of 12 cases). Overall, both level and slope shocks display strong dependence structure with correlation bounds of 0.35-0.99 and 0.26-0.81, respectively, whereas curvature shocks correlation is not as strong, being between 0.08-0.35. Correlations are positive and significant, at the 1% significance level; only for the curvature shocks of ICE crack and inter-crude spread we observe a zero state dependent coefficient. Furthermore, a far stronger co-movement is noted in the NYMEX crack and WTI-Brent spreads than in the ICE crack and heating-gas oil spreads, while the largest variations in correlation measures occur in the inter-product level shock pair with a spread of 0.56 (0.38-0.94). For the slopes and curvatures, inter-crude shocks involve the highest spread of 0.26 (0.54-0.8) and 0.33 (0.06-0.39), respectively. Curvature correlations diverge across regimes by a factor of 3 (heating-gas oil) to a factor of nearly 8 (ICE crack, 0.04-0.34). Overall, results suggest that there is variation in the volatilities and correlation processes across regimes.

The findings up to now indicate that caution should be taken when making inferences about the dynamics of the adjustment to the long-run equilibrium, because convergence and direction of causality cannot be known a priori. Results indicate differences not only between factors but between markets as well. For instance crack spreads in the US and European markets present dissimilar dynamics across regimes, indicating different pass through mechanisms, an effect not accommodated by linear models (e.g. in Table 6.2, Panel B). In addition, the estimates for the slopes present an interesting facet which can be attributed to the theory of storage. Under normal market conditions (i.e. low volatility regime) products will respond to the crude oil slope in order to restore the long-run equilibrium of relative backwardation/contango. On the other hand, in volatile periods the crude will also respond to the product's slope. This is expected since if crude oil inventories are low, products' inventories will also deplete after some time period subject to product-specific inventories. If both enter extreme backwardation then demand for products will play a key role since small changes in demand, for e.g. residential needs, will drive both crude and refined products' prices up in view of constrained capacity and time lag of production. Another interesting feature is that correlations, at least for level factors, are much stronger in the US crack (0.88-0.96 for levels) and the inter-crude spreads (0.83-0.99) as opposed to the European crack (0.65-0.80) and the inter-product market (0.73-0.94). For the crack spread markets this difference can be attributed to differences in market structure. First, as already mentioned, European petroleum products are more susceptible to extreme weather conditions. Second, the European market is more dependent on middle distillates, such as gas oil, and consumption has been historically higher than in the U.S where light distillates, such as gasoline, play a prominent role. Also, another reason as to why gas oil seems to display this relative autonomy might be that gas oil market has been growing at a faster pace than the US heating market for the last 5 years - futures volume and open interest data have surpassed the corresponding NYMEX heating oil figures. On the whole, MRS specification results in a more rich structure which produces economically meaningful results concerning the factor dynamics, being flexible enough to accommodate several fundamentals across diverse market conditions.

The estimates of the transition probabilities in Table 6.3, p_{11} and p_{22} , imply that the individual factor-specific regimes are fairly persistent. For instance, looking at the curvature factor for the ICE crack, p_{11} is 0.955 for gas oil and 0.985 for Brent whereas p_{22} is 0.868 and 0.985, respectively. Therefore, the probabilities of a low-to-high (p_{12}) volatility regime shift are $1-0.955=4.5\%$ and $1-0.985 = 1.5\%$ for gas oil and Brent, respectively. Similarly, the corresponding probabilities of a high-to-low (p_{21}) volatility regime shift are 13.2% and 1.5%, respectively. In Table 6.4, Panel A we also calculate the unconditional regime probabilities as

well as the duration⁷ of being in each regime. For the individual factor-specific regimes these are fairly high for the low variance state, nearly 70% on average, and relatively lower for the high variance state, nearly 30% on average.

Table 6.4: Unconditional Probabilities & Expected Duration

	<u>HO-CL</u>			<u>GO-CB</u>			<u>CL-CB</u>			<u>HO-GO</u>		
	<i>L</i>	<i>S</i>	<i>C</i>	<i>L</i>	<i>S</i>	<i>C</i>	<i>L</i>	<i>S</i>	<i>C</i>	<i>L</i>	<i>S</i>	<i>C</i>
Panel A: Factor-Specific Regimes												
$\pi_{1,st=1}$	0.774	0.708	0.772	0.747	0.534	0.745	0.547	0.568	0.763	0.596	0.627	0.591
	{15.2}	{5.4}	{6.2}	{6.1}	{11.8}	{4.4}	{5.5}	{9.6}	{4.5}	{50.0}	{8.4}	{15.3}
$\pi_{1,st=2}$	0.226	0.292	0.228	0.253	0.466	0.255	0.453	0.432	0.237	0.404	0.373	0.409
	{4.4}	{2.2}	{1.8}	{2.1}	{10.3}	{1.5}	{4.6}	{7.3}	{1.4}	{33.9}	{5.0}	{10.6}
$\pi_{2,st=1}$	0.738	0.773	0.766	0.534	0.626	0.490	0.786	0.728	0.776	0.749	0.696	0.747
	{16.1}	{13.6}	{7.7}	{10.0}	{5.4}	{13.2}	{22.5}	{6.6}	{5.8}	{40.8}	{9.5}	{8.0}
$\pi_{2,st=2}$	0.262	0.227	0.234	0.466	0.374	0.510	0.214	0.272	0.224	0.251	0.304	0.253
	{5.7}	{4.0}	{2.4}	{8.7}	{3.2}	{13.7}	{6.1}	{2.5}	{1.7}	{13.7}	{4.1}	{2.7}
Panel B: Joint Factor-Specific Regimes												
$\pi_{FG,st=1,1}$	0.572	0.547	0.592	0.399	0.334	0.380	0.430	0.414	0.592	0.446	0.437	0.441
	{7.9}	{3.9}	{3.5}	{3.9}	{3.7}	{3.4}	{4.5}	{4.0}	{2.6}	{22.5}	{4.5}	{5.3}
$\pi_{FG,st=2,2}$	0.059	0.066	0.053	0.118	0.174	0.125	0.097	0.118	0.053	0.102	0.113	0.104
	{2.5}	{1.5}	{1.1}	{1.7}	{2.5}	{1.4}	{2.7}	{1.9}	{0.8}	{9.8}	{2.3}	{2.2}
Panel C: Global Regimes												
$\pi_{G1,st=1}$		0.185			0.051			0.105			0.086	
		{1.6}			{1.3}			{1.2}			{2.3}	
$\pi_{G2,st=2}$		0.000			0.003			0.001			0.001	
		{0.6}			{0.7}			{0.5}			{1.1}	
p_{G11}		0.872			0.844			0.837			0.912	
p_{G22}		0.649			0.694			0.625			0.814	

This table reports the unconditional probabilities for each regime as implied from the estimated transition probability matrices of Table 3; $\pi_{i,st}$ is the unconditional individual factor specific probability of being in regime st , where subscripts i denotes the commodity i.e. 1 for the first and 2 for the second leg; $\pi_{FG,st}$ is the unconditional joint-factor specific global regime probability of two factors being simultaneously in either the low ($st=1$ & $st=1$) or high ($st=2$ & $st=2$) volatility regimes i.e. L-L, S-S and C-C. Obviously $1 - (\pi_{FG,1,1} + \pi_{FG,2,2})$ will give the probability of two factors being in different states i.e. either low-high or high-low states. $\pi_{Gi,st}$ is the unconditional global regime of all six factors corresponding to a specific spread being simultaneously in volatility regime i (hence, global regime); p_{Gii} is the (transition) probability of all six factors staying in volatility regime i , in the following period; Numbers in $\{\cdot\}$ display the expected weekly durations of being in regime st .

Looking at the joint factor-specific regime probabilities in Panel B, the low variance state is more stable lying in the range of 33%-59%, as opposed to 5%-17% for the high variance state; these represent the joint unconditional probabilities of both factors being in either $s_t = 1$ or 2. Note that in each case, unconditional probabilities are the solution to $\boldsymbol{\pi} = \boldsymbol{\pi} \boldsymbol{\Psi}$ (see Chapter 3, section 3.3.2, pp. 62), where $\boldsymbol{\pi}$ the vector of the equilibrium state probabilities e.g. for levels this would be $\boldsymbol{\pi}_L = \boldsymbol{\pi}_L \boldsymbol{\Psi}_L$. The joint unconditional probability that the components of a pair of factors are at different regimes is 33%-50% and this is the area where of the lower relative

⁷ The average expected duration of being in state 1 is calculated using the formula suggested by Hamilton (1989):

$$\sum_{i=1}^{\infty} i p_{11}^{i-1} (1 - p_{11}) = (1 - p_{11})^{-1}$$

correlations are observed; this is implied from the table as $1 - (\pi_{FG,st=1,1} + \pi_{FG,st=2,2})$. Next, after the Ψ^* matrices are calculated as $\Psi_L \otimes \Psi_S \otimes \Psi_C$ (see section 6.2.2), the upper left elements of the *global* transition probability matrices are equivalent to $(p_{11,1} p_{11,2})_L (p_{11,1} p_{11,2})_S (p_{11,1} p_{11,2})_C$ and vary between 0.84-0.91; that is, the probability of staying in the pair-specific global low volatility regime denoted as p_{G11} in Panel C. Similarly the lower right element of global Ψ^* lies between 0.62-0.81, showing that low volatility regime is associated with greater persistence. Overall, local factor-specific regimes (Panel A) are more persistent with expected duration ranging from more than a month to nearly a year. For joint factor-specific regimes (Panel B) this falls to a maximum of 6 months and for global regimes (Panel C) persistency falls radically to a maximum of 2.3 weeks for the low and 1.1 for the high variance state. Consequently, the multi-regime factor model allows for frequent transitions since the probabilities of being in one regime fall as we include more factors, in other words, it is more difficult for each of the underlying processes to coincide within the same regime. Therefore, a one factor model is characterised by 4 regimes, a two factor by 4^2 and a three factor by 4^3 regimes (given Eq. 6.5). Overall, high variance states are less stable and are characterised by much shorter duration compared to low variance states, consistent with other studies in the literature (see also Chapters 4 and 5).

For exposition purposes, the “smooth” regime probabilities for the principal components’ processes derived from the estimated MRS model are presented in Figures 6.3 and 6.4 for the NYMEX crack and WTI-Brent cases⁸. These indicate the likelihood of a pair of factors being in the low variance state with the shaded areas in the graphs identifying the periods when the same factors are in the high variance state. Note that the actual regimes for each pair of factors are four and thus, state probabilities do not add to one.

⁸ Based upon the estimated parameter vector $\hat{\theta}$, estimated from data spanning the period $t=1$ to T , “smooth” probability is the estimated probability that the unobserved state at time t equals 1 given the entire sample of observations from 1 to T . See Chapter 3, section 3.4.1 and Hamilton (1994).




Figure 6.3: NYMEX Heating Crack Regime Smoothed Probabilities for Level, Slope & Curvature Factors




Figure 6.4: WTI – Brent Spread Regime Smoothed Probabilities for Level, Slope & Curvature Factors

Figure 6.5 plots the volatility of the NYMEX crack spread for 1- month, 2- months and 9- months to maturity futures, estimated from the factor MRS model of Eq. (6.2); in this case, Λ becomes a function of the estimated state dependent volatilities. Time-variation arises from estimating the lower (low variance state) and upper (high variance state) bounds of variance and weighting these by the estimated regime probabilities. Evident is the presence of the Samuelson (1965) effect where shorter maturity contracts volatility is always at the top. NYMEX Crack spreads appear to be more volatile but volatility persists and has longer memory whereas the WTI-Brent market (Figure 6.7) is associated with more noise, increases in the volatility are in the form of jumps and shocks die out fast. This is also manifested in the time evolution of correlations. Short term spreads are associated with lower correlation between their components e.g. 1- month vs. 9- month NYMEX cracks, as observed in Figure 6.6. This holds for all spreads (not presented; for illustration we also present the WTI-Brent case in Figure 6.8). Medium term spreads, on the other hand, are more balanced in terms of correlation, being higher (on average) with less temporary decreases of magnitude e.g. 1- month NYMEX crack correlation falls to nearly 0.8, whereas the minimum value for the 3- month spread is never below 0.85. On the other hand, longer term spreads' correlation lies in-between the medium and short term spreads' correlations.




Figure 6.5: NYMEX Heating Crack 1, 3 & 9 Month Futures Spread Volatilities




Figure 6.6: NYMEX Heating Crack Correlations of 1, 3 & 9 Month Futures




Figure 6.7: WTI- Brent 1, 3 & 9 Month Futures Spread Volatilities

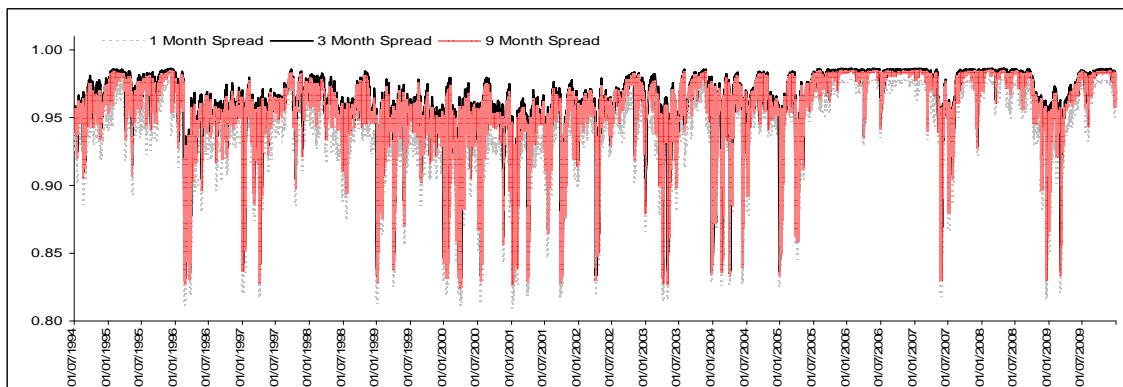


Figure 6.8: WTI- Brent Correlations of 1, 3 & 9 Month Futures

Finally, diagnostic tests of all models are presented in Table 6.5. Tests on the standardised residuals and standardised squared residuals indicate that most of the autocorrelation in the principal component series has been removed; however, some signs still remain e.g. Brent slope and curvature factors, at 1% significance level. By comparing the conditional and unconditional coefficients of skewness (in Table 6.1), we can note a nominal reduction in the levels of excess skewness and kurtosis. Moreover, linearity tests, using likelihood ratio (LR) statistics are also reported in Table 6.5. P-values indicate in each case the rejection of the linear model in favour of a nonlinear alternative. The same holds for the restricted two state regime switching model. Due to the existence of nuisance parameters, both LR tests are adjusted according to Davies (1987). Notably, the LR test of a regime-dependent intercept and heteroscedasticity model versus the full MRS model of Eq. (6.3) implies that the model can be reduced for the level and slope factors of NYMEX and ICE cracks. However, in the following section, to allow for a richer interaction between factors, we use all models to obtain the out-of-sample forecasts.

Table 6.5: Model Diagnostics

	<u>HO-CL</u>			<u>GO-CB</u>			<u>CL-CB</u>			<u>HO-GO</u>		
	<i>L</i>	<i>S</i>	<i>C</i>	<i>L</i>	<i>S</i>	<i>C</i>	<i>L</i>	<i>S</i>	<i>C</i>	<i>L</i>	<i>S</i>	<i>C</i>
LogLik	-14,832	-2,451.2	5,904.9	-17,087	-2,990.5	4,180.0	-13,391	-1,664.3	6,305.8	-16,849	-2,599.1	4,136.3
LR1	[0.271]	[0.119]	[0.002]	[0.233]	[0.113]	[0.000]	[0.011]	[0.005]	[0.000]	[0.000]	[0.000]	[0.000]
LR2	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
LR3	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Skew	-0.011	-0.088**	0.069*	-0.063	0.079**	-0.249***	-0.166***	0.075*	0.113***	-0.101**	-0.121***	0.037
Kurt	1.283***	1.974***	4.753***	1.850***	2.154***	9.348***	2.481***	2.059***	6.493***	1.169***	1.889***	4.674***
J-B	259.3***	618.0***	3,559***	541.3***	734.3***	13,793***	986.5***	670.8***	6,643***	221.6***	570.8***	3,439***
Q(5)	2.507	2.880	57.30***	3.195	3.574	8.722	4.752	1.868	4.873	0.932	3.978	43.79***
Q(20)	23.97	29.32*	121.3***	21.74	20.83	36.68**	39.24**	18.65	22.22	21.44	29.13*	106.2***
Q ² (5)	102.0***	118.0***	172.4***	97.48***	56.67***	25.38***	426.8***	59.78***	66.27***	69.14***	118.5***	193.5***
Q ² (20)	318.4***	226.2***	210.0***	172.5***	125.9***	25.95	1,371***	170.4***	91.54***	176.7***	218.6***	230.2***
Skew	-0.150***	0.135***	0.089*	-0.171***	-0.001	0.457***	-0.115***	-0.017	0.617***	0.090**	-0.035	-0.226***
Kurt	1.338***	2.321***	6.386***	1.777***	2.203***	12.56***	2.539***	2.332***	14.39***	2.584***	2.181***	5.442***
J-B	296.2***	859.3***	6,423***	515.4***	763.9***	24,951***	1,023***	855.9***	32,860***	1,056***	749.3***	4,694***
Q(5)	4.956	1.213	5.201	0.557	7.573	16.82***	2.817***	13.62**	9.484*	4.911	1.697	7.298
Q(20)	32.06**	18.34	22.08	26.15	25.25	39.21***	36.61**	35.51**	39.98***	26.44	23.54	29.58*
Q ² (5)	142.1***	37.99***	61.65***	92.94***	52.91***	26.42***	427.0***	56.04***	24.21***	32.50***	77.93***	23.22***
Q ² (20)	393.8***	137.1***	85.59***	237.4***	99.97***	45.27***	1344***	114.2***	41.53***	61.85***	144.7***	31.26*

LR1 is a test statistic of the null hypothesis of a regime-dependent intercept and heteroscedasticity model versus an MRS model where all coefficients are subject to regime switching. LR2 tests the null of a 4-regime model versus a 2-regime model whereas LR3 is the linearity test of a 4-regime model against a single regime alternative. These tests are constructed as $2(LL_{UNCON}-LL_{CON})$, where LL_{UNCON} and LL_{CON} represent the unconstrained and the constrained maximum likelihood respectively and are distributed as $\chi^2(r)$ where r is the number of restrictions imposed. Due to the existence of nuisance parameters, LR2 and LR3 are adjusted and they represent the upper bound of Davies' (1987) bound test. Under the assumption that the LL function has a single peak i.e. $\Theta = 2M^{1/2}$ and denoting the gamma function as $\Gamma(\cdot)$ the p -values of LRstat for $M = \{LR2, LR3\}$ are given by:

$$\Pr(\chi^2(r) > M) + \Theta M^{(r-1)/2} \exp^{-M/2} \frac{2^{-(r/2)}}{\Gamma(r/2)}.$$

6.5 Forecasting the Futures Curve Dynamics

Oil term structure evolution has important implications in the fields of energy risk management and derivatives pricing. To the authors' knowledge, the issue of predictability of the dynamics of oil price curves has received surprisingly little attention. Most papers deal with forecasting the very short end of the futures curve (for instance, see Sadorsky, 2002 for return forecasts and Chapter 4 for volatility forecasts) where most of the liquidity is concentrated. An exception is Chantziara and Skiadopoulos (2008), who by applying PCA on the futures curve of petroleum futures attempt to forecast the term structure by utilising lags of the estimated principal components; however, they find poor forecasting performance results. In contrast to the oil market, literature is more extensive in forecasting yield curves. Diebold and Li (2006), extract the level, slope and curvature factors and extend the *Nelson–Siegel yield curve* to a dynamic model, able to generate encouraging prediction results, especially for longer horizons (for more recent related studies the reader is referred to Moench, 2008 and Yu and Zivot, 2010).

In this section of the chapter we attempt to test the validity of the MRS futures curve model on an out-of-sample basis. While the model can in principle be employed to analyse interrelationships of correlated petroleum futures curve dynamics, we also provide evidence on

the usefulness of our approach in forecasting returns, variance-covariance matrices and risk management downside risk measures. Thus, to provide a more informative insight into the economic benefits and the appropriateness of our framework, we perform an out-of-sample comparison to study and corroborate the predictability of futures curves evolution, based on the suggested framework. For this reason, we estimate each model over the period 1994 to 2004, leaving the last five years (1,250 daily observations) for out-of-sample forecasting. Historical estimators of the PCA loadings, the co-integration equation coefficients and the model-specific parameters are updated on a quarterly basis (every 63 business days), using a rolling window of 10 years (2,530 daily returns). To provide robust evidence we investigate the performance of four models. The MRS models we test can be classified into two categories denoted as I-MRS (independent MRS) and R-MRS (restricted MRS). I-MRS stands for the case where each factor is modelled as an independent two-state first order MRS process, resulting a four regime MRS-VECM for the level and slope factors and a four-regime MRS-VAR-X for the curvature; this is the model that was tested empirically in section 6.4. R-MRS stands for the restricted case where each pair of factors is modelled as a two-state first order MRS process. For both I-MRS and R-MRS, apart from the unrestricted models where switching is permitted to all parameters of the model (see for instance Table 6.3 for the unrestricted I-MRS models), we also estimate two restricted versions (see also footnote of Table 6.6), one with switching permitted only in the intercepts v_{st} and the variance covariance matrix Σ_{st} and one with switching intercepts v_{st} , equilibrium adjustment coefficients α_{st} and variance covariance matrix Σ_{st} (i.e. no switching in the short run dynamics of the system, in which case $\Gamma_{i,st}$ of Eq. (6.3) is regime independent). In addition, we obtain forecast results from a linear single-regime version of the MRS models, denoted as F-DCC. Regarding the second moments, we employ a Dynamic Conditional Correlation (DCC) GARCH(1,1). Hence, once the time varying factor Σ_t is estimated from the DCC specification, the full futures returns' variance-covariance matrix V_t is obtained using Eq. (6.2). Considering all factor models, note that each time we obtain a set of three outcomes, using either 1-, 2- or 3- factors in obtaining the forecast; for instance, in the I-MRS case, these are denoted as I-MRS(1), I-MRS(2) and I-MRS(3), respectively. Finally, the natural benchmark that we employ (denoted as AR-DCC) is also a DCC model applied directly to the futures returns (20 series), filtered using an AR-GARCH model. Of course, for the returns forecasting results we also report the simple Random Walk's performance for completeness.

6.5.1 Forecasting Petroleum Spreads

The first experiment we put forth is to examine whether there is an improvement in forecasting the term structure of petroleum futures spreads using the suggested framework. To evaluate the forecasting error we use the root mean squared error (RMSE) of contemporaneous petroleum futures spreads. Let N be the number of out-of-sample observations (1,250), T the number of contract maturities (10), ho the forecast horizon and superscript e denote the forecasted value. The RMSE metric used, can be represented as:

$$\sqrt{\frac{1}{NT} \sum_{t=1}^N \sum_{T=1}^T \left[\left\{ \Delta \ln F_1(t+ho, T) - \Delta \ln F_2(t+ho, T) \right\} - \left\{ \Delta \ln F_1(t+ho, T)^e - \Delta \ln F_2(t+ho, T)^e \right\} \right]^2} \quad (6.8)$$

Table 6.6 presents this aggregate RMSE between the actual and forecasted spreads, for the entire out-of-sample period and also between the short term end of the spreads' futures curve and the longer term end. The first involves only the first five maturities under study whereas the latter the more distant months from six months to expiry, up to ten. Results of the forecast performance statistics for each model across the different forecast horizons 1-day, 1-week, 2-weeks and 1- month ahead are also reported.

Table 6.6: Root Mean Squared Errors, Forecasting the term structure of contemporaneous spreads

		HO-CL	% Gain/Loss against RW	GO-CB	% Gain/Loss against RW	CL-CB	% Gain/Loss against RW	HO-GO	% Gain/Loss against RW
Panel A: 1-Day ahead forecasts									
<i>Overall :</i>	Random Walk	0.8008		1.7790***		0.3532		1.7111***	
	AR-DCC	0.8007	0.01	1.7260***	3.07	0.3618*	-2.38	1.6535***	3.49
	F-DCC	0.8050	-0.52	1.5750	12.95	0.3704***	-4.65	1.3991	22.30
	R-MRS	0.8056**	-0.60	1.5743	13.00	0.3634	-2.81	1.3957	22.60
	I-MRS	0.8049	-0.51	1.5718	13.18	0.3612	-2.21	1.3954	22.63
<i>Short Term End:</i>	Random Walk	0.8896		1.8820***		0.4032		1.7847***	
	AR-DCC	0.8906	-0.11	1.8407***	2.25	0.4132*	-2.41	1.7442***	2.32
	F-DCC	0.8971*	-0.84	1.6779	12.17	0.4242***	-4.94	1.4622	22.05
	R-MRS	0.8979**	-0.92	1.6768	12.24	0.4170	-3.29	1.4578	22.42
	I-MRS	0.8969*	-0.81	1.6736	12.45	0.4144**	-2.68	1.4550	22.66
<i>Long Term End:</i>	Random Walk	0.7008		1.6696***		0.2947		1.6343***	
	AR-DCC	0.6994	0.21	1.6031***	4.15	0.3017*	-2.31	1.5575***	4.93
	F-DCC	0.7009	-0.01	1.4650	13.97	0.3074**	-4.11	1.3329	22.61
	R-MRS	0.7013	-0.07	1.4646	14.00	0.3004	-1.89	1.3306	22.82
	I-MRS	0.7009	-0.01	1.4630	14.13	0.2986	-1.31	1.3331	22.59
Panel B: 1-Week ahead forecasts									
<i>Overall :</i>	Random Walk	1.7250		2.2933***		0.7396**		1.8235***	
	AR-DCC	1.7206	0.25	2.2259***	3.03	0.7396**	-0.01	1.7560***	3.84
	F-DCC	1.7175	0.43	2.0950	9.46	0.7306	1.23	1.7124	6.49
	R-MRS	1.7167	0.48	2.0922	9.61	0.7302	1.28	1.7080	6.76
	I-MRS	1.7160	0.52	2.0955	9.44	0.7296	1.36	1.7054	6.92
<i>Short Term End:</i>	Random Walk	2.0146		2.5060***		0.8960*		1.8971***	
	AR-DCC	2.0103	0.21	2.4468***	2.42	0.8971*	-0.12	1.8484***	2.63
	F-DCC	2.0084	0.31	2.2946	9.21	0.8849	1.26	1.7833	6.38
	R-MRS	2.0070	0.38	2.2909	9.39	0.8848	1.26	1.7778	6.71
	I-MRS	2.0053	0.46	2.2932	9.28	0.8853	1.21	1.7772	6.74
<i>Long Term End:</i>	Random Walk	1.3758		2.0586		0.5396**		1.7468***	
	AR-DCC	1.3711	0.34	1.9804	3.95	0.5380*	0.30	1.6584*	5.33
	F-DCC	1.3661	0.71	1.8743	9.83	0.5334	1.16	1.6383	6.62
	R-MRS	1.3660	0.71	1.8726	9.93	0.5325	1.33	1.6353	6.82
	I-MRS	1.3668	0.65	1.8771	9.67	0.5301	1.79	1.6304	7.14
Panel C: 2 Weeks ahead forecasts									
<i>Overall :</i>	Random Walk	2.1756		2.4336***		0.9624		1.9110**	
	AR-DCC	2.1764	-0.04	2.3790**	2.30	0.9629	-0.05	1.8475	3.44
	F-DCC	2.1618	0.64	2.2966	5.97	0.9485	1.46	1.8764	1.84
	R-MRS	2.1620	0.63	2.2939	6.09	0.9481	1.51	1.8694	2.23
	I-MRS	2.1586	0.79	2.3006	5.78	0.9491	1.40	1.8642	2.51
<i>Short Term End:</i>	Random Walk	2.5439		2.6924**		1.1640		1.9998*	
	AR-DCC	2.5456	-0.07	2.6460*	1.75	1.1662	-0.19	1.9548	2.31
	F-DCC	2.5313	0.49	2.5456	5.77	1.1498	1.23	1.9709	1.47
	R-MRS	2.5300	0.55	2.5416	5.93	1.1489	1.31	1.9611	1.98
	I-MRS	2.5234	0.81	2.5454	5.77	1.1482	1.38	1.9523	2.44
<i>Long Term End:</i>	Random Walk	1.7307		2.1438**		0.7054		1.8178***	
	AR-DCC	1.7300	0.04	2.0780*	3.17	0.7030	0.34	1.7335	4.86
	F-DCC	1.7144	0.95	2.0170	6.29	0.6909	2.10	1.7769*	2.30
	R-MRS	1.7168	0.81	2.0160	6.34	0.6912	2.06	1.7729*	2.53
	I-MRS	1.7179	0.74	2.0263	5.80	0.6952	1.47	1.7718*	2.59
Panel D: 1-Month ahead forecasts									
<i>Overall :</i>	Random Walk	3.0534		2.9752**		1.2624		2.0381	
	AR-DCC	3.0657	-0.40	2.9307	1.52	1.2667	-0.34	1.9863	2.61
	F-DCC	3.0284	0.82	2.8838	3.17	1.2367	2.08	2.1079**	-3.31
	R-MRS	3.0309	0.74	2.8836	3.18	1.2360	2.13	2.0983**	-2.87
	I-MRS	3.0226	1.02	2.8849	3.13	1.2352	2.21	2.0928**	-2.61
<i>Short Term End:</i>	Random Walk	3.5618		3.3639		1.5489*		2.1398	
	AR-DCC	3.5766	-0.41	3.3295	1.03	1.5566*	-0.49	2.1055	1.63
	F-DCC	3.5378	0.68	3.2644	3.05	1.5256	1.53	2.2339**	-4.21
	R-MRS	3.5367	0.71	3.2613	3.14	1.5242	1.63	2.2184*	-3.54
	I-MRS	3.5185	1.23	3.2568	3.29	1.5006	3.22	2.2048*	-2.95
<i>Long Term End:</i>	Random Walk	2.4414		2.5275		0.8877		1.9311	
	AR-DCC	2.4505	-0.37	2.4683	2.40	0.8866	0.13	1.8596	3.85
	F-DCC	2.4139	1.14	2.4446	3.39	0.8552	3.81	1.9738**	-2.16
	R-MRS	2.4216	0.82	2.4482	3.24	0.8558	3.73	1.9710**	-2.02
	I-MRS	2.4276	0.57	2.4573	2.85	0.8941	-0.71	1.9744**	-2.19

The out-of sample data include 1,250 observations i.e. 5 years of data ending on 31 December 2009; For any given day, squared errors are calculated as the sum of the squared errors of 1 Month, up to 10 Month petroleum spreads (overall); for the *Short Term End* case we include only squared errors up to the 5th Month to maturity spreads, whereas for the *Long term End* case we included the squared errors for the contracts after the 5th Month and up to the 10th Month prior to expiry. Numbers in bold indicate the best performing model; The benchmark models are a simple Random Walk (RW) and an Autoregressive Process (AR) of the individual futures returns (20 time series) with an overall optimum lag order 2 –according to both the Schwarz information criterion (1978) and the autocorrelation function of futures returns; R-MRS and I-MRS are the 2- and 4- regime models, described in section 6.2. For economy of space we report the best performing model of each set. For instance I-MRS includes 9 models depending on the regime switching parameters and the number of factors utilised i.e. MSIH, MSICH and MSIACH each of them either a 1-, 2- or a 3- factor model; Asterisks *, **, *** indicate that the RMSE of the corresponding model is significantly higher than the competing models at 1%, 5% and 10%, respectively; the p-values are provided from White's (2000) Reality Check using the stationary bootstrap of Politis and Romano (1994). The number of bootstrap simulations is set to 3,000 and the smoothing parameter is $q = 0.1$.

First, we can observe that by allowing the factor structure to follow a Markov specification forecasts are overall improved. Looking at the results for the entire sample period, it is only in the NYMEX crack spread and inter-crude spread markets for the short term 1- day ahead forecasts that the RW model achieves better performance. However, RMSE values are marginally better at 0.5% and 2.2%, respectively. On the other hand, in the ICE crack and inter-product market, the I-MRS model achieves an improvement of 13.18% and 22.63%, respectively. We can also note that, forecast errors increase with the forecast horizon, reflecting the fact that uncertainty regarding future prices increases as well. NYMEX crack and inter-crude spreads display threefold increase in forecast errors from 1-day to 1- month horizon, whereas the corresponding increase for the ICE crack is less than twofold, and HO-GO seems to be the least affected with less than 20% increase in forecast errors. This is also reflected in the relative performance compared to RW, indicating the difficulties in forecasting longer term prices; for example, the 13.18% RMSE improvement of the I-MRS over the RW model in the GO-CB case is reduced at the 1- month horizon to 3.13%. On the whole, Markov models (I-MRS and R-MRS) are better in 12 out of the 16 cases. Comparing now the short and long term end of the futures spreads curve, the above effect is more pronounced for the more volatile, prompt months' spreads, where RMSE's are 15%-27% higher in the 1 month horizon, compared to the 1-day ahead forecasts. It is only the HO-GO spread that involves less than 2% increase in the RMSE's in the short term end of the futures curve. Results of the two sub-cases are consistent with the overall period, however, we can note that longer maturity months are associated with smaller forecast errors; this is expected since volatility increases as we approach expiry (Samuelson, 1965; see also Table 6.1, Panel B). Finally, I-MRS and R-MRS are better in 13 out of the 16 cases of the short term end contracts and 7 out of 16 cases of the long term end contracts, highlighting their ability to capture better, markets that exhibit higher volatility.

Finally, we also investigate whether modelling the factor structure leads to more accurate predictions of future prices using formal statistical tests because by considering only the nominal values of the RMSE scores across models, results are prone to data snooping bias. For that, we assess whether the forecasting performance of the competing models is equally accurate, employing White's (2000) reality check⁹ and the stationary bootstrap of Politis and

⁹ Forecast comparison is a historical measurement of how models would have performed in the out-of-sample period. However, by relying solely on the mean value of a statistical loss function it is difficult to refute that results would be qualitatively dissimilar in different periods or that they might be coincidental. Sullivan et al. (1999) and White (2000) proposed an approach to handle such biases by approximating the empirical distribution of a performance measure. Consider the loss differential: $fm_{k,t+1} = LF_{t+1}^k - LF_{t+1}^{benchmark}$, where k represents the k^{th} model and LF is the corresponding loss function. The null hypothesis to be tested is $H_0 = \max\{E[fm_k]\} \leq 0$, i.e. there is no model better than the benchmark; a small p-value

Romano (1994) (see Appendix 4.C for more technical details on bootstrap simulations). Results indicate that only in 3 cases out of 48, the I-MRS model is found to significantly underperform the competing models at 5% significance level. That is, the HO-GO 1-month ahead forecasts for the overall and longer term contracts, where AR-DCC is significantly better, and the CL-CB 1-day ahead forecasts for the prompt months contracts where Random Walk is better. On the contrary, at the same significance level, RW and AR-DCC are found to provide poor relative forecasts in 19 and 12 cases out of 48, respectively.

6.5.2 Forecasting the Variance Covariance Matrix

The second experiment we put forth is to examine whether there is any improvement in forecasting the full variance-covariance matrix of the petroleum futures curves' components. Since variance-covariance matrix is unobserved, the proxy that we use to compare the accuracy of out-of-sample forecasts from different models is the realised variance-covariance matrix, denoted as \mathbf{RV}_t . For each date t in the out-of-sample period, each element of $\mathbf{RV}_{ij,t+1}$ is calculated as $(r_{i,t+1} - E_t[r_{i,t+1}])(r_{j,t+1} - E_t[r_{j,t+1}])$, so that $\mathbf{RV}_{11,t+1}$ is $(r_{1,t+1} - E_t[r_{1,t+1}])^2$, $\mathbf{RV}_{12,t+1}$ is $(r_{1,t+1} - E_t[r_{1,t+1}])(r_{2,t+1} - E_t[r_{2,t+1}])$ and so on, where r is the actual realised return. Thus, realised variances are the squared demeaned returns whereas realised covariances are the cross products of the realised demeaned returns. Next, having defined a proper proxy for the true variance-covariance matrix, to formally assess the performance of the conditional second moments estimates we use the following set of loss functions to summarise the information of the multidimensional matrices of forecast errors:

$$LF_{t+1,p} = \|\mathbf{V}_{t+1} - \mathbf{RV}_{t+1}\|^p = \sqrt{\sum_{i,j} |\mathbf{V}_{ij,t+1} - \mathbf{RV}_{ij,t+1}|^p}; \quad i, j = 1, \dots, 20; \quad p = 1, 2 \quad (6.9)$$

$$LF_{t+1,EIGEN} = \sqrt{\lambda_{\max} \left[(\mathbf{V}_{t+1} - \mathbf{RV}_{t+1})(\mathbf{V}_{t+1} - \mathbf{RV}_{t+1})^T \right]} \quad (6.10)$$

indicates that there exists a model which provides superior forecasting results, based on a specific loss function. We use the stationary bootstrap of Politis and Romano (1994) to obtain the average loss function of each bootstrapped sample $\bar{f}m_k^*(b)$, based on 3,000 bootstrap simulations. The so called *bootstrap RC p-value* is obtained by comparing the observed statistic $T_n^{RC} = m a_x \{N^{1/2}(\bar{f}m_k)\}$ with the quantiles of the empirical distribution of T_n^{RC*} . The simulated statistic T_n^{RC*} is calculated as: $T_n^{RC*} = m a_x \{N^{1/2}(\bar{f}m_k^*(b) - \bar{f}m_k)\}$ (see also Chapters 4 and 5 for similar applications in volatility loss functions and hedged portfolios, respectively).

$$LF_{t+1,TRACE} = \left[\log \left(\frac{\sqrt{\text{trace}(\mathbf{V}_{t+1}^T \mathbf{V}_{t+1})}}{\sqrt{\text{trace}(\mathbf{R}\mathbf{V}_{t+1}^T \mathbf{R}\mathbf{V}_{t+1})}} \right) \right]^2 \quad (6.11)$$

The first two loss functions are *Frobenius distances* between the actual and the forecasted variance covariance matrix. For $p = 2$ the loss function is the natural extension of the RMSE to the multivariate case and is defined as the square root of element-wise squared differences. For $p=1$ the loss function resembles the MAE (Mean Absolute Error) and is defined as the square root of element-wise absolute differences. We employ both functions, denoted as LF_{RMSE} and LF_{MAE} , respectively. Note that for each t we consider only the triangular of the variance-covariance matrix; this is because every variance covariance matrix is symmetric, thus, taking into account the full matrix would lead to double counting the errors of the covariances. The next loss function of Eq. (6.10) is the square root of the largest eigenvalue of the matrix containing the squared forecast errors; this is the known as the *Hermitian distance*. The last loss function of Eq. (6.11) measures the proportional loss as the difference between the trace of the forecasted to that of the realised variance-covariance matrix; this was first introduced by Moskowitz (2003); see also Laurent et al. (2009) for more on the employed loss functions.

Results are presented in Table 6.7. The two loss functions based on the Frobenius distances, namely LF_{RMSE} and LF_{MAE} are consistent and direct us each time to the same model; that is, the I-MRS model. The only exception is the WTI-Brent case where the restricted version the Markov model i.e. the R-MRS is the best alternative. Similar are the results based on the LF_{EIGEN} apart from the ICE crude-product market which shows that the 1- factor DCC is slightly better, whereas the LF_{TRACE} supports the 1- factor R-MRS model for the CL-HO and CL-CB markets and the 1- factor DCC for the GO-CB and HO-GO markets. A noteworthy observation is that 1- factor models are adequate to forecast the true variance-covariance matrix, consistent with the initial PCA results (Table 6.1) where the first factor explains more than 97% of the variation of each individual futures curve. Including more factors only marginally changes the results and 2- and 3- factor models are ranked exactly next to each of the 1-factor models e.g. (in the HO-CL case, best model according to LF_{RMSE} is the 1 factor I-MRS, second and third best the 2- and 3- factor I-MRS).

Table 6.7: Forecasting the Variance Covariance Matrix of correlated futures curves

	LF _{RMSE}	LF _{MAE}	LF _{EIGEN}	LF _{TRACE}	LF _{MME(U)}	% U	LF _{MME(O)}	% O	LF _{W-sum}
Panel A: HO-CL									
AR-DCC	4.4885*	1.7995***	9.1041*	3.2483**	12.1352***	32.40	8.6392	67.60	9.7720*
F- DCC(1)	4.4921**	1.7935***	9.1113**	3.2327**	11.3286***	32.69	8.8725***	67.31	9.6755
F- DCC(2)	4.4943**	1.7940***	9.1132**	3.2336**	11.4755***	32.69	8.8357**	67.31	9.6985
F- DCC(3)	4.4944**	1.7940***	9.1132**	3.2336**	11.4858***	32.68	8.8333**	67.32	9.7003
R-MRS(1)	4.4186	1.7602	8.9589	3.1630	10.7869	33.92	9.0514***	66.08	9.6401
R-MRS(2)	4.4211	1.7607	8.9612	3.1641	10.9372	33.91	9.0038***	66.09	9.6593
R-MRS(3)	4.4213	1.7608	8.9612	3.1641	10.9469	33.90	9.0011***	66.10	9.6607
I-MRS(1)	4.3960	1.7504	8.9132	3.1780	10.6847	34.14	9.0797***	65.86	9.6277
I-MRS(2)	4.3984	1.7509	8.9154	3.1791	10.8468	34.14	9.0206***	65.86	9.6441
I-MRS(3)	4.3986	1.7509	8.9154	3.1791	10.8593	34.14	9.0171***	65.86	9.6460
Panel B: GO-CB									
AR-DCC	4.1908*	1.7215***	8.2127	2.2177	12.1458***	30.74	8.3419	69.26	9.5111
F- DCC(1)	4.1172	1.6862	8.0899	2.1578	11.3561	32.55	8.6030**	67.45	9.4992
F- DCC(2)	4.1185	1.6865	8.0908	2.1582	11.4728	32.55	8.5692**	67.45	9.5142
F- DCC(3)	4.1186	1.6866	8.0908	2.1581	11.4776	32.55	8.5682**	67.45	9.5151
R-MRS(1)	4.2854***	1.7453***	8.4277***	2.3663***	12.5149***	29.63	8.4190	70.37	9.6328
R-MRS(2)	4.2875***	1.7457***	8.4291***	2.3669***	12.6723***	29.63	8.3890	70.37	9.6582*
R-MRS(3)	4.2877***	1.7458***	8.4291***	2.3670***	12.6797***	29.63	8.3879	70.37	9.6594*
I-MRS(1)	4.1149	1.6814	8.1020	2.2258*	11.2828	32.27	8.6293**	67.73	9.4856
I-MRS(2)	4.1166	1.6818	8.1031	2.2263*	11.4285	32.24	8.5925**	67.76	9.5068
I-MRS(3)	4.1167	1.6818	8.1031	2.2263*	11.4393	32.23	8.5899**	67.77	9.5083
Panel C: CL-CB									
AR-DCC	4.5687*	1.7938***	9.3165*	4.1141***	12.1458***	31.49	8.3419	68.51	9.5399
F- DCC(1)	4.5851**	1.7890***	9.3477**	4.1006***	11.3206***	31.94	8.6067***	68.06	9.4736
F- DCC(2)	4.5873**	1.7895***	9.3500**	4.1017***	11.4351***	31.95	8.5691**	68.05	9.4847
F- DCC(3)	4.5875**	1.7895***	9.3500**	4.1017***	11.4399***	31.95	8.5678**	68.05	9.4853
R-MRS(1)	4.4319	1.7133	9.0367	3.8813	10.1732	36.01	9.0029***	63.99	9.4243
R-MRS(2)	4.4338	1.7138	9.0387	3.8825	10.3378	35.98	8.9461***	64.02	9.4468
R-MRS(3)	4.4340	1.7138	9.0387	3.8825	10.3485	35.98	8.9435***	64.02	9.4490
I-MRS(1)	4.5133	1.7433*	9.2027	4.0494**	10.6505*	34.00	8.8318***	66.00	9.4502
I-MRS(2)	4.5153	1.7437*	9.2048	4.0504**	10.7848**	33.99	8.7958***	66.01	9.4720
I-MRS(3)	4.5154	1.7438*	9.2048	4.0504**	10.7939**	33.99	8.7932***	66.01	9.4733
Panel D: HO-GO									
AR-DCC	4.0737***	1.7226***	7.9691**	2.1235**	11.3924***	31.63	7.5947	68.37	8.7959***
F- DCC(1)	3.9717	1.6811**	7.7964	2.0409	8.9466**	33.54	8.3933***	66.46	8.5788
F- DCC(2)	3.9730	1.6814**	7.7975	2.0413	9.0593***	33.52	8.3491***	66.48	8.5872
F- DCC(3)	3.9731*	1.6814**	7.7975	2.0413	9.0757***	33.52	8.3432***	66.48	8.5887
R-MRS(1)	3.9453	1.6634	7.7719	2.0656	8.5563	34.17	8.6124***	65.83	8.5933
R-MRS(2)	3.9470	1.6638	7.7734	2.0662	8.7268	34.16	8.5367***	65.84	8.6016
R-MRS(3)	3.9472	1.6638	7.7734	2.0662	8.7528	34.15	8.5248***	65.85	8.6027
I-MRS(1)	3.9302	1.6588	7.7428	2.0690	8.6102	34.10	8.5221***	65.90	8.5521
I-MRS(2)	3.9319	1.6591	7.7440	2.0695	8.7602	34.08	8.4560***	65.92	8.5596
I-MRS(3)	3.9320	1.6592	7.7440	2.0695	8.7782	34.08	8.4478***	65.92	8.5604

- For the out-of-sample tests 1,250 forecasts (5 years of data) of the variance-covariance matrix are obtained by the rolling window forecasting scheme (2,531 in-sample observations at each step); Numbers in () indicate the 1-, 2- and 3- factor model; See also notes in Table 6.6.
- The regime switching models presented in this table, both R-MRS and I-MRS are the models where all coefficients are subject to regime switching.
- LF_{RMSE} and LF_{MAD} are given in Eq. (6.9), for $p = 2$ (RMSE) and $p = 1$ (MAE), respectively; LF_{EIGEN} and LF_{MTR} are given in Eq. (6.10) and Eq. (6.11), respectively; Mixed error loss functions of over- and under- prediction i.e. LLF_{MME(O)}, LLF_{MME(U)} are both given in Eq. (6.12); (% U and O is the proportion of under prediction and over prediction, respectively, over the forecast period. All the error statistics are the average value of the loss functions defined in the above equations. All are rescaled for exposition purposes (e.g. RMSEs with a multiple of 10^3 and MAE of 10^2). The column named W-Sum is the weighted summation of the Mean Over and Under Prediction error according to the estimates % U and O, respectively.
- Asterisks *, **, *** indicate that the loss function of the corresponding model is significantly higher than the competing models at 1%, 5% and 10%, respectively; the p-values are provided from White's (2000) Reality Check using the stationary bootstrap of Politis and Romano (1994). The number of bootstrap simulations is set to 3,000 and the smoothing parameter is $q = 0.1$.

Results of White (2000) reality check are also presented in Table 6.7 in form of asterisks. We can see that the I-MRS is never significantly outperformed at 1% significance level. On the other hand, AR-DCC produces larger errors in all cases, at 10% significance level

indicating the benefits of dynamic factor models; F-DCC is also not very consistent - only in the GO-CB case performs well - indicating the benefits of modelling the factor as Markov processes.

Because none of the abovementioned metrics provide any information on the asymmetry of the prediction variance errors; that is, whether there is any difference between forecast errors when the model over-predicts or under-predicts the actual variance we employ an extension of the Brailsford and Faff (1996) Mixed Error statistics to the multivariate case. This uses a mixture of positive and negative forecast errors with different weights. This is an important forecast metric because, although we expect forecast errors to be unbiased on average, there might be occasions when a model produces small errors but consistently over-predicts or under-predicts the conditional second moments. Thus, we also look at the proportion of negative and positive forecast errors for each model, since a model with symmetric forecast errors should produce about 50% positive and 50% negative forecast errors, with similar means. The corresponding loss functions employed are:

$$\begin{aligned}
 ME_{t+1}(U) &= \sum \left(\sqrt{\sum_{i,j}^O |V_{ij,t+1} - RV_{ij,t+1}|} + \sum_{i,j}^U |V_{ij,t+1} - RV_{ij,t+1}| \right); \\
 ME_{t+1}(O) &= \sum \left(\sqrt{\sum_{i,j}^U |V_{ij,t+1} - RV_{ij,t+1}|} + \sum_{i,j}^O |V_{ij,t+1} - RV_{ij,t+1}| \right); \quad i, j = 1, \dots, 20 \quad (6.12)
 \end{aligned}$$

Results are now more informative. The proportion of over- and under- prediction is similar across all models and markets: 29%-35% of all models underpredict second moments whereas overprediction ranges between 65%-71% in line with Chapter 4. Furthermore, looking at the scale of over- and under- prediction errors, it can be seen that, on average, mean under-prediction is higher than mean over-prediction, implying that all models fail to capture the large sudden jumps of volatility, which is nevertheless expected since jumps are due to random shocks that are very hard to predict. Overprediction results support the same models as the previous metrics, whereas if someone is only interested in underprediction then the errors are significantly minimised by full modelling of the 20 individual contracts series each time as an AR-DCC-GARCH model (involving 82 parameters). Asymmetric error statistics have important implications for different players of the energy markets. For instance, a regulatory body such as a bank (lender) that has financed a company's energy project (e.g. for oil exploration and extraction) may prefer a model which over-predicts risk since the company (borrower) would be

required to allocate more funds for capital adequacy requirements. Conversely, energy companies, depending on their risk aversion, would prefer a model that ‘efficiently’ under-predicts risk, since this way they have to allocate fewer resources for future risks. However, the more balanced models that produce the optimum trade-off of over- and under- prediction are the 1-factor I-MRS for the two crude-refined product and the inter-product market whereas in the inter-crude market, the 1-factor R-MRS is more accurate.

6.5.3 An Application to Value-at-Risk

Having obtained the full variance covariance matrix forecasts for each pair of petroleum futures we finally examine the practical relevance and usefulness of our findings in estimating risk management measures, in particular Value-at-Risk. In doing so, we consider portfolios consisting of 20 futures contracts each time. We assume that we hold a portfolio $\mathbf{W} = [\mathbf{W}_1, \mathbf{W}_2]$ of futures, with $\mathbf{W}_1 = [w_{1,1}, w_{1,2}, \dots, w_{1,10}]$ being the position in the i^{th} futures contract of the first commodity of the pair and $\mathbf{W}_2 = [w_{2,1}, w_{2,2}, \dots, w_{2,10}]$ the position in the second commodity of the pair. We consider four different portfolios with constant weights throughout time. Each portfolio consists of 10 long positions and 10 short positions. The first portfolio is an *equally weighted portfolio* of spreads: we assign a weight of 10% to each of the maturities (i.e. $w_{1,1}=\dots=w_{1,10}=10\%$ for the long leg of the spread and $w_{2,1}=\dots=w_{2,10}=-10\%$ for the short leg of the spread) and the weight vector is $\mathbf{W} = [\mathbf{W}_1=\mathbf{1}_{(1 \times 10)} \mathbf{W}_2=-\mathbf{1}_{(1 \times 10)}]/10$. The second portfolio is a *slope portfolio* constructed by taking opposite positions in the long and short end part of the term structure of the spreads: the weight vector is $\mathbf{W}_1 = [\mathbf{1}_{(1 \times 5)} \mathbf{-1}_{(1 \times 5)}]$ and $\mathbf{W}_2 = [\mathbf{-1}_{(1 \times 5)} \mathbf{1}_{(1 \times 5)}]$. The third is a *curvature portfolio* of spreads constructed by taking the same positions in the long and short end part of the term structure of the spreads, but opposite in the medium part: the weight vector consists of $\mathbf{W}_1 = [\mathbf{1}_{(1 \times 3)} \mathbf{-1}_{(1 \times 4)} \mathbf{1}_{(1 \times 3)}]$ and $\mathbf{W}_2 = [\mathbf{-1}_{(1 \times 3)} \mathbf{1}_{(1 \times 4)} \mathbf{-1}_{(1 \times 3)}]$. The last is an arbitrary *calendar portfolio* of spreads where the weight vector consists of $\mathbf{W}_1 = [1, -1, 1, -1, 1, -1, 1, -1, 1, -1,]$ and $\mathbf{W}_2 = -\mathbf{W}_1$. Note that in each case we examine both long and short positions.

VaR is one of the most popular approaches for quantifying market risk, defined as the maximum expected loss in value of an asset or a portfolio of assets over a target horizon, given a specific confidence level $1-c$. Then, conditional on the information set at t (Ω_t), VaR can be defined as the solution to $\Pr(r_{t+1} \leq VaR_{t+1}^c | \Omega_t) = c$, where r_{t+1} are the actual returns of each examined portfolio. VaR forecasts are obtained using the forecasted matrix \mathbf{V}_t as $VaR_{t+1}^c = \mathbf{W}\boldsymbol{\mu}_{t+1} + \Phi_{t+1}(c)\sqrt{\mathbf{W}\mathbf{V}_{t+1}\mathbf{W}^T}$, where $\boldsymbol{\mu}$ is a column vector consisting of the daily return forecasts. Note that by assigning weights and constructing commodity portfolios we

obtain portfolio returns with different distributional properties, hence, instead of assuming a normal distribution we use the filtered unconditional historical simulation quantile $\Phi_{t+i}(c)$. This way, $\Phi_{t+i}(c)$ is the same across models and we compare the VaR forecasts purely on the forecasted \mathbf{V} , and $\boldsymbol{\mu}$ matrices. Comparisons are made on the basis of the likelihood ratio test of unconditional coverage. LR_{UC} ¹⁰ tests the null hypothesis that the probability of realising a loss in excess of the forecasted VaR is statistically equal to the nominal coverage rate c . VaR violations that occur more frequently than c % of the time imply that the VaR method used systematically underestimates the true level of risk, and vice-versa. Furthermore, following Koenker and Bassett (1978) we also employ a loss function, the predictive quantile loss (QL) which is based on quantile regression (similar to Chapter 4, section 4.4.2.1). The QL function penalises more heavily observations for which a violation occurs, and is actually a measure of fit of the predicted tail at a given confidence level. The objective is to minimise QL :

$$QL = \frac{1}{N} \sum_{i=1}^N (r_{t+i} - VaR_{t+i}^c) \left((1-c) I_{\{r_{t+i} < VaR_{t+i}^c\}} + c I_{\{r_{t+i} \geq VaR_{t+i}^c\}} \right) \quad (13)$$

The economic intuition behind the use of the QL is that capital charges should also be taken into account, hence, the capital forgone from overpredicting the true VaR should not be neglected. This latter loss function is asymmetric in view of the fact that underprediction and overprediction of VaR estimates have diverse implications. For instance, underprediction of risk might lead to liquidity problems and reoccurring underprediction causes insolvency. On the other hand, overprediction implies higher capital charges which, although are not a cause of bankruptcy, reflect the opportunity cost of keeping a high reserve ratio.

Results are presented in Table 6.8. We report the VaR violation rates for the four portfolios at the 1%, 2.5%, 5% and 10% level for both long and short positions. In contrast to the variance-covariance forecasts, results are mixed. There does not seem to be much difference across models, or consistency. Performance is associated with the specific portfolios. All models perform rather good across all confidence levels. An exception is the slope portfolio - for all models - and the inter-crude related portfolios.

¹⁰ Let n be the number of outcomes that fall outside the forecast interval, N the number of forecasts and \hat{c} the empirical level of coverage. Then, the statistic is expressed as: $LR_{UC} = -2 \log \left[\frac{c^n (1-c)^{N-n}}{\hat{c}^n (1-\hat{c})^{N-n}} \right] \sim \chi^2(1)$

However, we can observe that violation rates mainly arise from overprediction of volatility rather than underprediction, implying that all forecasts of portfolios volatilities are rather conservative. For instance, in the slope portfolio of the HO-GO spread, all models fail to pass the LR_{UC} test. At the low level of 1% VaR for both long and short positions the maximum number of violations is 0.2% from the AR-DCC model (i.e. 2-3 violations only). Similar is the performance at lower confidence levels e.g. 90% where again actual losses exceed the AR-DCC based VaR at 10% only 6.2% of the time (76-78 exceedances instead of the theoretical 125 [=10% \times 1,250]). Turning next to the quantile loss functions the results are not at all the same as we would have predicted based on Tables 6.6 and 6.7. The main competitors now are the AR-DCC and factor F-DCC models, where the latter is better 66% of the time. However, a closer look reveals that in fact, models are marginally different in nominal values.

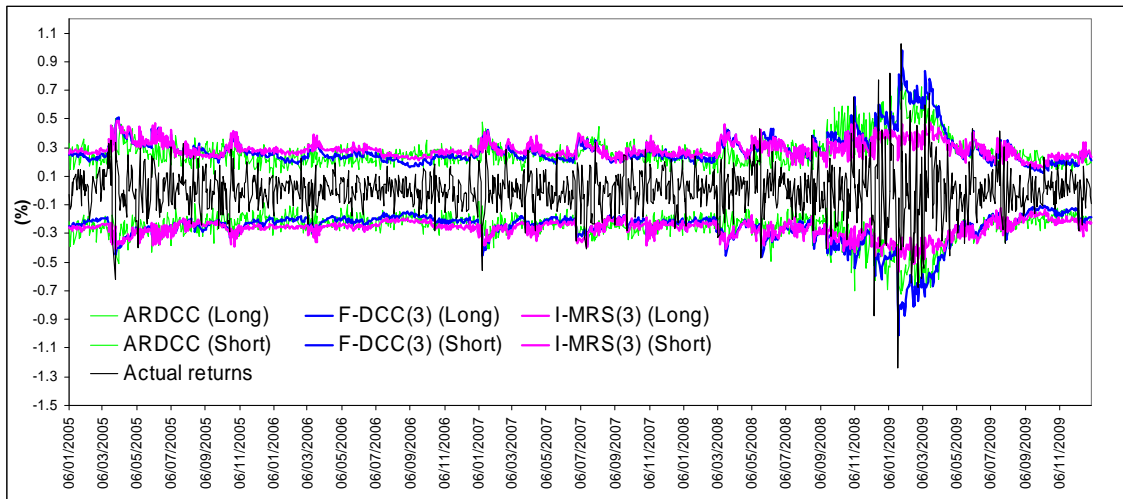


Figure 6.9: NYMEX Crack 5% VaR Estimates for the Equally Weighted Portfolio (Long & Short Positions)

Finally, Figure 6.9 and 6.10 depict the excess losses of the 5% and 95% VaR from the AR-DCC, F-DCC(3) and I-MRS(3) models for the *equally weighted portfolio* (Table 6.8, Panel A) and the *calendar portfolio* (Table 6.8, Panel D). Comparing the AR-DCC and F-DCC(3) model, it seems that the estimates are very similar. Another observation that can be made is that in highly volatile periods e.g. the last quarter of 2008 and the first of 2009, GARCH models are more responsive to sudden market changes than the I-MRS with relatively higher average VaR estimates. This may be due to parameter instability in the specific period as well as uncertainty regarding the unobserved regime, as mentioned in Engel (1994) and Marsh (2000). However, as shown in Table 6.8, the percentage of failures for the I-MRS models is more accurate for the

long positions on the equally weighted and calendar portfolios that are plotted in Figures 6.9 and 6.10 whereas the AR-DCC is the best alternative for short positions.




Figure 6.10: NYMEX Crack 5% VaR Estimates for the Calendar Portfolio (Long & Short Positions)

6.6 Conclusions

In this chapter we examined linkages in pairs of the main underlying orthogonal factors explaining the variation of petroleum futures. Using non-linear equilibrium adjustment models, we examined the short and long run relationships between the extracted factors in respective petroleum futures, namely the components of the NYMEX and ICE crack spreads (heating and gas oil) and the inter-commodity spreads (inter-crude and inter-product). We find evidence in favour of the existence of a long-run relationship between level and slope factors, however, curvatures are found to be mean-reverting to commodity-specific equilibria. Specifying flexible dynamic regime switching evolution equations for the respective factors changes, we introduced a new functional multi-regime model driven by Markov dynamics. The rationale behind the use of these models stems from the fact that first, the dynamics of correlated futures curves should be inherent in the common factors explaining the price variation and second this relationship may be characterised by regime shifts, suggesting that by allowing the data generating process to be dependent upon the “state of the market”, one may obtain more efficient estimates. Results indicate that each regime clearly differentiates two distinct market dynamics for both the conditional mean and the volatility of the underlying process. Moreover, it seems that when one market is in the low and the other in the high variance state, it is more likely to observe lower correlations. A far stronger co-movement is noted in the NYMEX crack and WTI-Brent than

that of ICE crack and heating-gas oil. Overall, both level and slope shocks display strong dependence structure (0.35-0.99 and 0.26-0.81, respectively), whereas curvature shocks are correlated to a lesser extent (0.08-0.35). While our multi-regime framework is primarily designed to aid our understanding of nonlinear behaviour and inter-dependencies in the factor structure of correlated futures curves we also provide evidence on the predictive ability of such models in forecasting the conditional first and second moments as well as in forecasting popular risk measures such as portfolio Value-at-Risk. Results from these exercises indicate that the multi-regime factor MRS models can sometimes achieve significant gains compared to competing models. Overall, the resulting model is very promising, providing a very practical policy analysis tool to market participants for identifying and timing the possible states that the market of combined futures curves is in, as well as forecasting large covariance matrices and estimating risk metrics for large portfolios.

In this chapter, we provided a thorough analysis of linkages and analysis of risk in a multivariate multi-regime switching framework. This essay completes the empirical part of the thesis. Overall, there seem to be some insightful benefits from assuming regime switching behaviour of petroleum dynamics. We examined linkages, interdependencies and risk attitudes, volatilities and VaR as well as optimum hedging. Summary of findings, formal conclusions and potential directions for further future research will be given in chapter 7, *Concluding Remarks and Future Research* which concludes the thesis.

APPENDIX 6.A: Factor Seasonality and Auto-correlogram

In this Appendix we briefly illustrate the effects of allowing the resulting eigenvalues of PCA factor decomposition, to be dependent on seasonal volatilities. Note that, in terms of the proportion of variance explained when performing traditional PCA results are qualitatively the same, yet, the periodic behaviour of the futures prices is markedly reflected in the factor process, as confirmed by the sample autocorrelation functions (ACF). The seasonal features of futures prices were mainly absorbed by higher order components, mostly the curvature. Figure 6.A.1 plots the autocorrelation function (ACF) of heating oil (left panel) and WTI crude oil (right panel) before (grey) and after (black) the adjustment for seasonality. Figure 6.A.1 clearly verifies that for heating oil curvature a big proportion of seasonality is mitigated using Eq. (6.1) whereas crude oil seasonal behaviour is not evident in either case.




Figure 6.A.1: ACF of the 3rd Factor (Curvature) for Heating (left) & WTI Crude Oil (right), before and after the Adjustment for Seasonality.

Chapter 7

Concluding Remarks and Future Research

7.1 Summary and Conclusions

With the fast growing energy sector, providing robust, innovative and resourceful financial solutions is crucial for the success of any oil business. Petroleum commodity prices are determined by supply and demand but these forces are driven by complex interactions, from events in the Middle East and country-specific energy security policies to climatic conditions and speculative money flows. As a result, the specific properties of oil prices call for reliable and consistent models. In turn, these models are an essential tool for market participants to comprehend the evolution of prices, volatilities, correlations and economic relationships with the aim to develop efficient risk measurement schemes and devise sound risk management strategies.

Some of the major empirical properties of the price series data employed in Chapters 4, 5 and 6 (as well as Chapters 2 and 3 to a smaller degree) illustrate that petroleum prices' distributions are heavy tailed and asymmetric while volatilities and correlations are time-dependent. These stylised facts of the energy markets motivate the modelling of change and the use of advanced quantitative techniques to describe their conditional distributions. The present work dealt with the modelling of change in the context of established concepts in energy economics. In particular, the focus was on explaining regime switching behaviour in oil markets, since, more often than not, market shocks alter the properties of the series' in various ways; some shocks are persistent, some transitory, some regular and some irregular.

One potential setting that involves regime changes in the mechanism that generates oil prices is the transition from backwardation to contango market conditions and vice versa. The theory of storage (Kaldor, 1939; Working, 1949; Telser, 1958 and Brennan, 1958) asserts that during low inventory periods, where supply is exhausted, due to tight market conditions spot prices are high and delivery in the future is priced at a discount. Moreover, empirical findings show that volatilities (correlations across the term structure) increase (decrease) with falling inventories i.e. under backwardation (see for example Ng and Pirrong, 1996 for refined

petroleum products and Geman and Ohana, 2009 for oil and gas as well as Fama and French, 1987 and Ng and Pirrong, 1994). Inversely, in periods of supply abundance i.e. under contango, spot prices fall and delivery in the future is priced at a premium- whereas volatilities are expected to be lower due to the flexible supply conditions. These different types of regimes serve as motivation for our regime switching approach to characterise oil prices. Driven by several complexities in the empirical validation of the energy markets, this thesis offered an alternative viewpoint on energy risk, providing a new framework for risk analysis that concentrated on practical applications.

Regime switching models are designed to capture cyclical behaviour and unknown breaks. Model parameters are functions of a hidden Markov chain and empirical data reveal their own structure. These models are characterised by a pre-specified number of distinct regimes within which different model parameters apply, whereas the probabilities of each state change over time and model parameters become time-dependent. Models of changing regime demonstrated the potential to benefit energy market participants in many applications. Three major research themes were carried out: a) to compare the different modelling techniques' ability to characterise and accurately predict the time-varying nature of oil price risk in a Value-at-Risk context, b) to explore the practical relevance of state-dependent time varying hedge ratios and c) to derive the most important energy risk factors and provide a risk analysis framework of correlated futures curves. The next section reviews the main findings throughout the thesis. Next, potential future directions to continue research are pointed out.

7.1.1 Risk Measurement

Chapter 4 addressed the concept of oil price risk and how to deal with non-normality, non-linearity, and non-constant conditional second moments in the risk measurement process. Traditionally, the family of GARCH models (Bollerslev, 1986) has been widely used to describe conditional volatility. Nevertheless, empirical evidence suggests that those models are rigid to accommodate the modelling complexities that energy markets exhibit. For instance, they induce a high degree of persistence in shocks, implying, falsely, highly predictability (Lamoureux and Lastrapes, 1990). For a robust estimation of the volatilities and the quantiles of the returns distributions, regime switching models were employed. By allowing the second moments to be dependent upon the state of the market, the volatility and VaR forecasts obtained were more efficient.

Results indicated a longer duration low volatility state, associated with low sensitivity to market shocks that die out very slowly, and a transitory high volatility state with shocks that affect the variance more but die out faster; this implies that the regime-based models are

superior at capturing persistence in volatility. Moreover, our volatility modelling framework extended previous research by including the squared deviations of futures from their long-run equilibrium as represented by the lagged basis (Lee, 1994; Ng and Pirrong, 1996) because as prices respond to the magnitude of disequilibrium then, in the process of adjusting, they become more volatile. The findings implied that in the low volatility state, the dynamics of the volatilities are more predictable; in the high volatility state, volatility changes mainly due to short-lived random shocks. Finally, market participants should consider regime behaviour in the modelling process, since augmented regime volatility models for all petroleum futures demonstrated improved forecasting accuracy under both periods of backwardation and contango. These models were also combined with Extreme Value Theory which posed as a conservative alternative in forecasting VaR, thus, being more apt to risk averse investors. Overall, the magnitude of disequilibria is a factor that does have explanatory power in determining potential changes in oil price volatilities and by identifying different volatility components in different periods, market participants may benefit in terms of accurate quantification of risk.

7.1.2 Risk Management

While the risks faced by industry are various and differ throughout the sectors of the industry, - from upstream to downstream - price risk is universal to all. Chapter 5 addressed the concept of hedging oil price risk. Oil price risk management has always been a vital part of the successful operation of oil-related businesses. A key parameter in devising effective futures hedging strategies is the hedge ratio. Traditionally, hedge ratios are estimated to minimise the variance of the hedged portfolio (Ederington, 1979). To allow for time-dependency in the hedging decision, GARCH models have been widely used (Kroner and Sultan, 1993). This chapter extended Chapter 4 and presented a multivariate regime error correction GARCH model to investigate the hedging effectiveness of petroleum futures.

The regime dependent conditional variances uncovered a link between persistence and the state of the market, consistent with the results of Chapter 4 as well as other studies in the petroleum economics literature such as Fong and See (2002). Overall, the high variance state is associated with high variance-covariance persistence and low duration and vice versa. Furthermore, combining the concept of disequilibrium (as measured by the error correction coefficients) with that of uncertainty (as measured by the conditional second moments) across high and low volatility regimes, the regime error correction GARCH model illustrated that the dynamics of the spot-futures relationship do not behave uniformly to shocks to equilibrium across different states. For instance, only in the low variance regime the speeds of adjustment of

spot and futures prices to their long-run relationship were in accordance with convergence towards a long-run equilibrium relationship; that is, equilibrium is primarily restored by the spot price as cash markets are more sensitive to news while futures depend on several factors like maturity and liquidity. Regarding the regime-dependent hedge ratios, these were found to be higher when the volatility in the market is low. Overall, the forecasting results indicated that regime dependent hedge ratios may be able to offer superior gains to market agents, measured in terms of both variance reduction and increase in utility. These findings held even when we examined the downside risk and considered the asymmetric risk profile of long and short hedgers.

7.1.3 Term Structure of Correlated Curves

The last empirical part of the thesis, Chapter 6, dealt with an important issue in petroleum market dynamics, correlated petroleum futures curves. Little is known about the joint term structures of different commodities and their implied dependence. Exceptions are Clewlow and Strickland (2000), Tomalsky and Hindanov (2002) and Ohana (2010). The price and volatility pattern across prompt and deferred contracts, as well as the correlation term structure, have been a cause of concern for market participants. This chapter exploited the information content of the dependence structure of petroleum futures curves and described interdependencies between petroleum commodities under different regimes. Employing a flexible multivariate error correction multi-regime framework we extended Chapters 4 and 5 and assessed the forecast ability of those models.

After decomposing the individual petroleum futures curves to the main risk factors i.e. level, slope and curvature we employed regime switching model for pairs of factors. All factors were allowed to switch independently, extending Bollen et al. (2000) model to the multivariate case; this way we effectively permitted factor specific regimes to demonstrate diversity i.e. one being in the high volatility state and another in a low volatility state, and hence, we disaggregated the regimes as level, slope and curvature driven and studied their interaction. Each regime clearly differentiated two distinct market dynamics for both the conditional mean and the volatility of the underlying process. Moreover, when one market was in the low and the other in the high variance state, it seemed more likely to observe low correlations. A far stronger co-movement was noted in the US crack and intercrude spreads than that of European crack and inter-product spreads. Results indicated that both level and slope shocks displayed strong dependence structure, sharing a common equilibrium relationship (in comparable pairs e.g. level with level) whereas curvature shocks were correlated to a lesser extent and were mean-reverting and stationary. While our multi-regime framework was primarily designed to aid our

understanding of regime behaviour and co-movement in the factor structure of correlated futures curves, the model was found useful when forecasting the conditional first and second moments as well as in risk measures such as portfolio Value-at-Risk. Results from these exercises indicated that the multi-regime factor MRS model can sometimes achieve significant gains compared to competing approaches. Overall, the resulting model is very promising, providing a very practical policy analysis tool to market participants for identifying and timing the possible states that the market of combined futures curves is in, as well as forecasting large covariance matrices and estimating risk metrics for large portfolios.

7.2 Directions for Further Research

All three studies in this thesis reveal the complex evolution of the conditional volatilities and correlation dynamics in petroleum markets. Furthermore the last two chapters reveal dependencies in two correlated assets, such as spot-futures, or in two correlated futures curves. The experiments considered indicate that efficiency improves and regime switching models serve for a better understanding of the conditional distribution of petroleum returns whereas the gains from using regime models are translated to enhanced forecasting ability. Future research should therefore be devoted to the development of models that allow for more realistic dynamics and new experiments can be set up with more flexible formulation.

First, given the results of Chapter 4, 5 and 6, an interesting extension would be to add a second part in each of these analyses and study and compare the potential pros and cons of specifying observed regime models rather than latent state models. Although unobserved state models let the data speak for themselves it might be beneficial for both academics and practitioners to reveal the specific behaviour of the price, volatility and correlation process under pre-specified regimes. Backwardation and contango will serve as an ideal alternative regime identification process to discover the particular backwardation-contango GARCH dynamics, backwardation-contango hedge ratios and backwardation-contango factor structure. For this reason we suggest the use of either, a model that will include dummies to differentiate the parameter estimates in periods of backwardation and contango or, alternatively, the family of smooth transition models. This can be applied to all parts of the thesis.

Second, within the setting of Chapters 4 and 5, multivariate extensions will be more appealing to market participants and more challenging on an academic level, since the risk measurement and risk management process in practice concerns portfolios rather than specific assets. For instance, oil companies involved in the management of physical assets, such as refineries, are mainly concerned with upstream-downstream commodities interrelationships to plan and optimise their operations. Also, investors without commercial interest in the real asset,

such as hedge funds, are genuinely interested in investment strategies that contain more than one assets. For instance, Haigh and Holt (2002) estimate time-varying hedge ratios for an energy trader exposed to the crack spread. Another study by Börger et al. (2009) examine risk measures and implications to risk management for economically meaningful energy related portfolios using data for crude oil, electricity, coal and CO₂ emission allowances; they mainly focus on gas and coal fired power plants. Our regime switching framework could be adopted to study the information content of regimes in risk measures and hedge ratios of such portfolios. Given that the number of parameters increases substantially with the inclusion of additional assets, a solution to this can be a regime switching copula model which will explicitly model the dependence as a regime switching process - rather than the whole set of marginal distributions.

Third, a challenging extension would be to apply the proposed regime switching framework in markets which entail non-linear payoffs i.e. options. Conventional option pricing models of commodity prices that rely on the Geometric Brownian Motion assumption are too simplistic and more sophisticated approaches may benefit market participants in terms of market understanding as well as the division of sound trading strategies and risk management techniques. The Markov GARCH class of models might be able to improve the pricing and hedging performance. Under the GARCH setup markets are incomplete and a finite number of risk-neutral densities exist. Duan (1995) described a technique, the Local Risk Neutral Valuation to approximate prices in a GARCH option valuation setting. In addition Badescu et al. (2008) studied the pricing function of options under normal mixture GARCH processes and further applied Esscher transforms and the Girsanov principle; this model is a restriction of Haas, Mittnik and Paolella (2004b) model and it would be interesting to investigate their relative performance. Finally, a more demanding study would be to extend the above into spread options such as the exchange-traded European Calendar options of the petroleum markets. Then, based on Duan and Pliska (2004)¹ and Duan and Theriault (2007) who propose a co-integration GARCH volatility framework we can also derive cointegration risk premia within the local risk neutral valuation scheme and extend to regime switching approaches (see also Duan et al., 2002).

¹ This study was the first to apply a discrete time approach to model the multivariate dynamics of spread options using co-integration and GARCH volatility. They approximate European option prices using Monte Carlo techniques and further examine the Greeks and the sensitivity of the estimated prices to the inclusion of co-integration. Moreover, they provide a diffusion limit for co-integrated systems with constant volatility, which constitutes a complete market model as opposed to the discretised version which becomes complete only after applying the Local Risk Neutral Valuation (LRNV) technique, described in Duan (1995).

References

- Agnolucci, P., 2009. Volatility in crude oil futures: A comparison of the predictive ability of GARCH and implied volatility models. *Energy Economics*, 31 (2), 316-321.
- Alexander, C., 2008. Market risk analysis, *Practical financial econometrics*, vol. 2, John Wiley & Sons.
- Alexander, C. & Lazar, E., 2006. Normal mixture GARCH(1,1): Applications to exchange rate modelling. *Journal of Applied Econometrics*, 21 (3), 307-336.
- Alizadeh, A.H., Kavussanos, M.G. & Menachof, D.A., 2004. Hedging against bunker price fluctuations using petroleum futures contracts: Constant vs. time varying hedge ratios. *Applied Economics*, 36 (12), 1337-1353.
- Alizadeh, A.H. & Nomikos, N.K., 2004a. Cost of carry, causality and arbitrage between oil futures and tanker freight markets, *Transportation Research Part E: Logistics and Transportation Review*, 40 (4), 297-316.
- Alizadeh, A.H. & Nomikos, N.K., 2004b. A Markov regime switching approach for hedging stock indices. *Journal of Futures Markets*, 24 (7), 649-674.
- Alizadeh, A.H., & Nomikos, N.K., 2008. Performance of statistical arbitrage in petroleum futures markets. *Journal of Energy Markets*, 1 (2), 3-33.
- Alizadeh, A.H., Nomikos, N.K. & Pouliasis, P.K., 2008. A Markov regime switching approach for hedging energy commodities. *Journal of Banking and Finance*, 32 (9), 1970-1983.
- Aloui, C. & Jammazi, R., 2009. The effects of crude oil shocks on stock market shifts behaviour: A regime switching approach. *Energy Economics*, 31 (5), 789-799.
- Andersen, T.G & Bollerslev, T., 1998. Answering the sceptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39 (4), 885-905.
- Ang, A. & Bekaert, G., 2002. International asset allocation with regime shifts. *The Review of Financial Studies*, 15 (4), 1137-1187.
- Ang, A. & Bekaert, G., 2004. How regimes affect asset allocation? *Financial Analysts Journal*, 60 (2), 86-99.
- Baba, Y., Engle, R.F., Kraft, D. & Kroner, K.F., 1987. Multivariate simultaneous generalized ARCH. Unpublished manuscript, University of California, San Diego.
- Badescu, A., Kulperger, R. & Lazar, E., 2008. Option valuation with normal mixture GARCH models. *Studies in Nonlinear Dynamics & Econometrics*, 12 (2), Article 5, 1-40.

- Bansal, R. & Zhou, H., 2002. Term structure of interest rates with regime shifts. *Journal of Finance*, 57 (5), 1997-2043.
- Bao, Y., Lee, T.H. & Saltoglu, B., 2006. Evaluating predictive performance of value-at-risk models in emerging markets: A reality check. *Journal of Forecasting*, 25 (2), 101–128.
- Benz, E. & Trück, S., 2009. Modelling the price dynamics of CO2 emission allowances. *Energy Economics*, 31 (1), 4-15.
- Bera, A. & Jarque, C., 1980. Efficient tests for normality, heteroscedasticity, and serial dependence of regression residuals. *Economic Letters*, 6 (3), 255-259.
- Billio, M., 2010. Dynamic risk exposures in hedge funds. *Computational Statistics and Data Analysis*, forthcoming.
- Billio, M. & Pelizzon, L., 2000. Value at risk: A multivariate switching regime approach. *Journal of Empirical Finance*, 7 (5), 531–554.
- Billio, M. & Pelizzon, L., 2003. Volatility and shocks spillover before and after EMU in European stock markets. *Journal of Multinational Financial Management*, 13 (4-5), 323-340.
- Black, F. & Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81 (3), 637-654.
- Bollen, N.P.B., Gray, S.F. & Whaley, R.E., 2000. Regime switching in foreign exchange rates: Evidence from currency option prices. *Journal of Econometrics*, 94 (1-2), 239-276.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics*, 31 (3), 307–327.
- Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH models. In: Engle, R.F. & McFadden, D.L. (Eds.), *Handbook of econometrics*, vol. 4, 2959-3038. Elsevier Science.
- Börger, R., Cartea, Á., Kiesel, R. & Shindlmayr, G., 2009. Cross-commodity analysis and applications to risk management. *Journal of Futures Markets*, 29 (3), 197-217.
- Borovkova, S., 2006. Detecting market transitions and energy futures risk management using principal components. *The European Journal of Finance*, 12 (6-7), 495-512.
- Borovkova, S. & Geman, H., 2006. Seasonal and stochastic effects in commodity forward curves. *Review of Derivatives Research*, 9 (2), 167-186.
- Brailsford, T.J. & Faff, R.W., 1996. An evaluation of volatility forecasting techniques. *Journal of Banking and Finance*, 20 (3), 419-438.
- Brennan, M. J., 1958. The supply of storage. *American Economic Review*, 48 (1), 50-72.
- Brennan, M. J. & Schwartz, E. S., 1985. Evaluating natural resource investments. *Journal of Business*, 58 (2), 135-157.

- Buffington, J. & Elliott, R.J., 2002. American options with regime switching. *International Journal of Theoretical and Applied Finance*, 5 (5), 497-514.
- Cai, J., 1994. A Markov model of switching regime-ARCH. *Journal of Business and Economic Statistics*, 12 (3), 309-316.
- Chan, W.S. & Tong, H., 1986. On estimating thresholds in autoregressive models. *Journal of Time Series Analysis*, 7 (3), 179-190.
- Chantziara, Th. & Skiadopoulos, G., 2008. Can the dynamics of the term structure of petroleum futures be forecasted? Evidence from major markets. *Energy Economics*, 30 (3), 962-985.
- Chen, K.C., Sears R.S. & Tzang, D., 1987. Oil prices and energy futures. *Journal of Futures Markets*, 7 (5), 501-518.
- Chen, S.S. & Chen, H.C., 2007. Oil prices and real exchange rate. *Energy Economics*, 29 (3), 390-404.
- Cheong, C.W., 2009. Modeling and forecasting crude oil markets using ARCH-type models. *Energy Policy*, 37 (6), 2346-2355.
- Chincarini, L.B., 2007. The Amaranth debacle: A failure of risk measures or a failure of risk management? *Journal of Alternative Investments*, 10 (3), 91-104.
- Choi, K. & Hammoudeh, S., 2010. Volatility behaviour of oil, industrial commodity and stock markets in a regime-switching environment. *Energy Policy*, 38 (8), 4388-4399.
- Chou, W.L., Denis, K.K.F. & Lee, C.F., 1996. Hedging with the Nikkei index futures: The conventional versus the error correction model. *The Quarterly Review of Economics and Finance*, 36 (4), 495-505.
- Choudhry, T., 1997. Short run deviations and volatility in spot and futures stock returns: Evidence from Australia, Hong Kong, and Japan. *Journal of Futures Markets*, 17 (6), 689-705.
- Christoffersen, P., 1998. Evaluating interval forecast. *International Economic Review*, 39 (4), 841-864.
- Chung, H., Davig, T. & Leeper, E.M., 2007. Monetary and fiscal policy switching. *Journal of Money, Credit and Banking*, 39 (4), 809-842.
- Clarida, R.H., Sarno, L., Taylor, M.P. & Valente, G., 2003. The out-of-sample success of term structure models as exchange rate predictors: A step beyond. *Journal of International Economics*, 60 (1), 61-83.
- Clements, M.P. & Krolzig, H.M., 2002. Can oil shocks explain asymmetries in the US Business Cycle? *Empirical Economics*, 27 (2), 185-204.

- Clements, M P. & Krolzig, H.M., 2004. Can regime-switching models reproduce the business cycle features of US aggregate consumption, investment and output? *International Journal of Finance and Economics*, 9 (1), 1-14.
- Clewlow, L. & Strickland, C., 2000. *Energy derivatives: Pricing and risk management*. Lacima Publications.
- Coe, P.J., 2002. Financial crisis and the great depression: A regime switching approach. *Journal of Money, Credit and Banking*, 34 (1), 76-93.
- Cologni, A. & Manera, M., 2009. The asymmetric effects of oil shocks on output growth: A Markov-switching analysis for the G-7 countries. *Economic Modelling*, 26 (1), 1-29.
- Cortazar, G. & Schwartz, E.S., 1994. The valuation of commodity-contingent claims. *Journal of Derivatives*, 1 (4), 27-39.
- Costello, A., Asem, E. & Gardner, E., 2008. Comparison of historically simulated VaR: Evidence from oil prices. *Energy Economics*, 30 (5), 2154-2166.
- Cotter, J. & Hanly, J., 2006. Reevaluating hedging performance. *Journal of Futures Markets*, 26 (7), 677-702.
- Crowder, W. & Hamed, A., 1993. A cointegration test for oil futures market efficiency. *Journal of Futures Markets*, 13 (8), 933-941.
- Culp, C. & Miller, M., 1994. Hedging a flow of commodity deliveries with futures: Lessons from Metallgesellschaft. *Derivatives Quarterly*, 1 (1), 7-15.
- Culp, C. & Miller, M., 1995a. Metallgesellschaft and the economics of synthetic storage. *Journal of Applied Corporate Finance*, 7 (4), 62-76.
- Culp, C. & Miller, M., 1995b. Hedging in the theory of corporate finance: A reply to our critics. *Journal of Applied corporate Finance*, 8 (1), 121-127.
- Dai, Q. & Singleton, K., 2003. Term structure dynamics in theory and reality. *The Review of Financial Studies*, 16 (3), 641-678.
- Davies, R.B., 1987. Hypothesis testing when a nuisance parameter is present only under the alternative. *Biometrika* 74 (1), 33-43.
- De Santis, R.A., 2003. Crude oil price fluctuation and Saudi Arabia's behaviour. *Energy Economics* 25 (2), 155-173.
- Dickey, D. & Fuller, W., 1979. Distribution of the estimates for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74 (366), 427-431.
- Dickey, D. & Fuller, W., 1981. Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica*, 49 (4), 1057 - 1072.

- Diebold, F.X. & Li, C., 2006. Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130 (2), 337-364.
- Diebold, F.X. & Mariano, R.S., 1995. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13 (3), 253-263.
- Diebold, F.X. & Rudebusch, G.D., 1996. Measuring business cycles: A modern perspective. *The Review of Economics and Statistics*, 78 (1), 67-77.
- Driesprong, G., Jacobsen, B. & Maat, B., 2008. Striking oil: Another puzzle? *Journal of Financial Economics*, 89 (2), 307-327.
- Driffil, J., Kenc, T. & Sola, M. 2009. Real options with priced regime-switching risk. Department of Economics Working Papers 2009-09, Universidad Torcuato Di Tella.
- Duan, J-C., 1995. The GARCH option pricing model. *Mathematical Finance*, 5 (1), 13-32.
- Duan, J-C. & Pliska, S.R., 2004. Option valuation with co-integrated asset prices. *Journal of Economic Dynamics and Control*, 28 (4), 727-754.
- Duan, J-C., Popova, I. & Ritchken, P., 2002. Option pricing under regime switching. *Quantitative Finance*, 2 (2), 1-17.
- Duan, J-C. & Theriault, A., 2007. Co-integration in crude oil components and the pricing of crack spread options. Working Paper.
- Dueker, M.J., 1997. Markov switching in GARCH processes and mean reverting stock market volatility. *Journal of Business and Economic Statistics*, 15 (1), 26-34.
- Duffie, D., Gray, S.F. and Hoang, P.H., 2004. Volatility in energy prices. In: Kaminski, V. (Eds.), *Managing energy price risk: The new challenges and solutions* (3rd ed.). Risk Books, 539-571.
- Dvir, E. & Rogoff, K.S., 2010. The three epochs of oil. Harvard University.
- Ederington, L.H., 1979. The hedging performance of the new futures markets. *Journal of Finance*, 34 (1), 157 - 170.
- Ederington, L.H. & Salas, J., 2008. Minimum variance hedging when spot price changes are partially predictable. *Journal of Banking and Finance*, 32 (5), 654-663.
- Edwards, F. & Canter, M. 1995. The collapse of Metallgesellschaft: Unhedgeable risks, poor hedging strategy, or just bad luck? *Journal of Futures Markets*, 15 (3), 211-264.
- Elliott, R.J. & Hinz, J., 2002. Portfolio optimization, hidden Markov models, and technical analysis of P&F-charts. *International Journal of Theoretical and Applied Finance*, 5 (4), 385-399.
- Elliott, R.J., Siu, T. K. & Chan, L., 2006. Option pricing for GARCH models with Markov switching. *International Journal of Theoretical and Applied Finance*, 9 (6), 825-841.

- Engel, C., 1994. Can the Markov switching model forecast exchange rates? *Journal of International Economics*, 36 (1-2), 151-165.
- Engle, R.F., 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of the United Kingdom inflation. *Econometrica*, 50 (4), 987-1007.
- Engle, R.F. & Granger, C.W.J., 1987. Cointegration and error correction: Representation, estimation and testing. *Econometrica*, 55 (2), 251-276.
- Engle, R.F. & Kroner, K.F., 1995. Multivariate simultaneous generalized ARCH. *Econometric Theory*, 11 (1), 122-150.
- Engle, R.F., Ng, V. & Rothschild, M., 1990. Asset pricing with a factor ARCH covariance structure: Empirical estimates for treasury bills. *Journal of Econometrics*, 45 (1-2), 213-237.
- Ewing, B. & Harter, C., 2000. Co-movements of Alaska North Slope and UK Brent crude oil prices. *Applied Economic Letters*, 7 (8), 553-558.
- Fama, E.F., 1965. The behaviour of stock-market prices. *Journal of Business*, 38 (1), 34-105.
- Fama, E.F. & French, K.R., 1987. Commodity futures prices: Some evidence on forecast power, premiums and the theory of storage. *Journal of Business*, 60 (1), 55-73.
- Fan, Y., Zhang, Y.J., Tsaic, H.T. & Wei, Y.M., 2008. Estimating 'value at risk' of crude oil price and its spillover effect using the GED-GARCH approach. *Energy Economics*, 30 (6), 3156-3171.
- Fattouh, B., 2009. Basis variation and the role of inventories: Evidence from the crude oil market. Oxford Institute of Energy Studies, WPM38.
- Fattouh, B., 2010. The dynamics of crude oil price differentials. *Energy Economics*, 32 (2), 334-342.
- Ferderer, J. P., 1996. Oil Price volatility and the macroeconomy. *Journal of Macroeconomics*, 18 (1), Winter, 1-26.
- Fleming, J. & Ostdiek, B, 1999. The impact of energy derivatives on the crude oil market. *Energy Economics*, 21 (2), 135-167.
- Fong, W.M. & See, K.H., 2001. Modelling the conditional volatility of commodity index futures as a regime switching process. *Journal of Applied Econometrics*, 16 (2), 133-163.
- Fong, W.M. & See, K.H., 2002. A Markov switching model of the conditional volatility of crude oil prices. *Energy Economics*, 24 (1), 71-95.
- Fong, W.M. & See, K.H., 2003. Basis variations and regime-shifts in the oil futures market. *The European Journal of Finance*, 9 (5), 499-513.

- Franses, P. H. & Paap, R., 1999. Does seasonality influence the dating of business cycle turning points? *Journal of Macroeconomics*, 21 (1), 79-92.
- Geman, H. & Nguyen, V.N., 2005. Soybean inventory and forward curves dynamics. *Management Science*, 51 (7), 1076-1091.
- Geman, H. & Ohana, S., 2009. Forward curves, scarcity, and price volatility in oil and natural gas markets. *Energy Economics*, 31 (4), 576-585.
- Ghosh, A., 1993. Hedging with stock index futures: Estimation and forecasting with error correction model. *Journal of Futures Markets*, 13 (7), 743-752.
- Ghysels, E., 1994. On the economics and econometrics of seasonality. In: Sims, C.A. (Eds.), *Advances in econometrics, Sixth World Congress*, vol. 1, 257-316. Cambridge University Press.
- Ghysels, E., Granger, C.W.J. & Siklos, P.L., 1996. Is seasonal adjustment a linear or nonlinear data filtering process? *Journal of Business and Economic Statistics*, 14 (3), 374-386.
- Giannikis, D., Vrontos, I.D. & Dellaportas, P., 2008. Modelling nonlinearities and heavy tails via threshold normal mixture GARCH Models. *Computational Statistics & Data Analysis*, 52 (3), 1549-1571.
- Gibson, R. & Schwartz, E.S., 1990. Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance*, 45 (3), 959-976.
- Giot, P. & Laurent, S., 2003. Market risk in commodity markets: A VaR approach. *Energy Economics*, 25 (5), 435-457.
- Girma, P.B. & Paulson, A.S., 1998. Seasonality in petroleum futures spreads. *Journal of Futures Markets*, 18 (5), 581-598.
- Goldfeld, S.M. & Quandt, R.E., 1973. A Markov model for switching regressions. *Journal of Econometrics*, 1 (1), 3-15.
- Gonzalez-Rivera, G., Lee, T.H. & Mishra, S., 2004. A reality check based on option pricing, utility function, value-at-risk and predictive likelihood. *International Journal of Forecasting*, 20 (4), 629-645.
- Gray, S.F., 1996. Modelling the conditional distribution of interest rates as regime switching process. *Journal of Financial Economics*, 42 (1), 27-62.
- Griffin, J.M., 1985. OPEC behaviour: A test of alternative hypotheses. *The American Economic Review*, 75 (5), 954-963.
- Guidolin, M. & Timmermann, A., 2007. Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31 (11), 3503-3544.

- Guidolin, M. & Timmermann, A., 2008. International asset allocation under regime switching, skew and kurtosis preferences. *The Review of Financial Studies*, 21 (2), 889-935.
- Guo, F., Chen, C.R. & Huang, Y.S., 2011. Markets contagion during financial crisis: A regime-switching approach. *International Review of Economics & Finance*, 20 (1), 95-109.
- Gülen, S.G., 1996. Is OPEC a cartel? Evidence from cointegration and causality tests. *The Energy Journal*, 17 (2), 43-58.
- Haas, M., Mittnik, S. & Paolletta M.S., 2004a. Mixed normal conditional heteroscedasticity. *Journal of Financial Econometrics*, 2 (2), 211-250.
- Haas, M., Mittnik, S. & Paolletta M.S., 2004b. A new approach to Markov-switching GARCH models. *Journal of Financial Econometrics*, 2 (4), 493-530.
- Haigh, M.S. & Holt, M.T., 2002. Crack spread hedging: Accounting for time-varying volatility spillovers in the energy futures markets. *Journal of Applied Econometrics*, 17 (3), 269-289.
- Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time series and business cycle. *Econometrica*, 57 (2), 357-384.
- Hamilton, J.D., 1990. Analysis of time series subject to changes in regime. *Journal of Econometrics*, 45 (1-2), 39-70.
- Hamilton, J.D., 1994. *Time series analysis*. Princeton, NJ: Princeton University Press.
- Hamilton, J.D., 2003. What is an oil shock? *Journal of Econometrics*, 113 (2), 363-398.
- Hamilton, J.D. & Lin, G. 1996. Stock market volatility and the business cycle. *Journal of Applied Econometrics*, 11 (5), 573-593.
- Hamilton, J.D., & Susmel, R. 1994. Autoregressive conditional heteroscedasticity and changes in regime. *Journal of Econometrics*, 64 (1-2), 307-333.
- Harvey, A.C. & Jaeger, A., 1993. Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics*, 8 (3), 231-247.
- Harvey, A.C. & Scott, A., 1994. Seasonality in dynamic regression models. *The Economic Journal*, 104 (427), 1324-1345.
- Haushalter, G.D., 2000. Financing policy, basis risk, and corporate hedging: Evidence from oil and gas producers. *Journal of Finance*, 55 (1), 107-152.
- Heath, D., Jarrow, R. & Morton, A., 1992. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60 (1), 77-105.
- Holmes, M.J. & Wang, P. 2003. Oil price shocks and the asymmetric adjustment of UK output: A Markov-switching approach. *International Review of Applied Economics*, 17 (2), 181-192.
- Holton, G.A., 2003. *Value-at-risk: Theory and practice*. Academic Press, Elsevier Science.

- Hotelling, H., 1931. The Economics of exhaustible resources. *Journal of Political Economy*, 39 (2), 137-175.
- Huang, B.N., Yang, C.W. & Hwang M.J., 2009. The dynamics of a nonlinear relationship between crude oil spot and futures prices: A multivariate threshold regression approach. *Energy Economics*, 31 (1), 91-98.
- Huang, D., Yu, B., Fabozzi, F.J. & Fukushima, M., 2009. CAViaR-based forecast for oil price risk. *Energy Economics*, 31 (4), 511-518
- Hung, J.C., Lee, M.C. & Liu, H.C., 2008. Estimation of value-at-risk for energy commodities via fat-tailed GARCH models. *Energy Economics*, 30 (3), 1173-1191.
- Jaeger, A. & Kunst, R.M., 1990. Seasonal adjustment and measuring persistence in output. *Journal of Applied Econometrics*, 5 (1), 47-58.
- Jammazi, R. & Aloui, C. 2010. Wavelet decomposition and regime shifts: Assessing the effects of crude oil shocks on stock market returns. *Energy Policy*, 38 (3), 1415-1435.
- Johansen, S., 1988. Statistical analysis of cointegrating vectors. *Journal of Economic Dynamics and Control*, 12 (2-3), 231-254.
- Johnson, L.L., 1960. The theory of hedging and speculation in commodity futures. *The Review of Economic Studies*, 27 (3), 139-151.
- Kaldor, N., 1939. Speculation and economic stability. *The Review of Economic Studies*, 7 (1), 1-27.
- Kang, S.H., Kang S.M. & Yoon, S.M., 2009. Forecasting volatility of crude oil markets. *Energy Economics*, 31 (1), 119-125.
- Kaufmann, R.K., Bradford, A., Belanger, L.H., McLaughlin, J.P & Miki Y., 2008. Determinants of OPEC production: Implications for OPEC behaviour. *Energy Economics*, 30 (2), 333-351.
- Kaufmann, R.K., Dees, S., Karadeloglou, P. & Sanchez, M., 2004. Does OPEC matter? An econometric analysis of oil prices. *The Energy Journal* 25 (4), 67-90.
- Kaufmann, R.K. & Ullman, B., 2009. Oil prices, speculation and fundamentals: Interpreting causal relations among spot and future prices. *Energy Economics*, 31 (4), 550-558.
- Kavussanos, M.G. & Nomikos, N.K., 2000. Hedging in the freight futures markets. *Journal of Derivatives*, 8 (1), 41-58.
- Keynes, J.M., 1930. *A treatise on Money: The applied theory of money*. vol. 2, MacMillan.
- Kim, C.J., 1994. Dynamic linear models with Markov-switching. *Journal of Econometrics*, 60 (1-2), 1-22.

- Kim, C.J. & Nelson, C. R., 1998. Business cycle turning points, a new coincident index, and tests of duration dependence based on a dynamic factor model with regime switching. *The Review of Economics and Statistics*, 80 (2), 188-201.
- Klaassen, F., 2002. Improving GARCH volatility forecasts with regime-switching GARCH. *Empirical Economics*, 27 (2), 363-394.
- Kleit, A.N., 2001. Are regional oil markets growing closer together? An arbitrage cost approach. *The Energy Journal*, 22 (2), 1-15.
- Koenker, R. & Bassett, G., 1978. Regression quantiles. *Econometrica*, 46 (1), 33-50.
- Kon, S.J., 1984. Models of stock returns: A comparison. *Journal of Finance*, 39 (1), 147-165.
- Krehbiel, T. & Adkins, L.C., 2005. Price risk in the NYMEX energy complex: An extreme value approach. *Journal of Futures Markets*, 25 (4), 309-337.
- Krichene, N., 2006. World crude oil markets: Monetary policy and the recent oil shock. *International Monetary Fund (IMF) Working Paper No. 06/62*.
- Krolzig, H.M., 1999. Statistical analysis of cointegrated VAR processes with Markovian regime shifts. Unpublished manuscript, Department of Economics, University of Oxford.
- Kroner, K.F. & Sultan, J., 1993. Time-varying distributions and dynamic hedging with foreign currency futures. *Journal of Financial and Quantitative Analysis*, 28 (4), 535-551.
- Kwiatkowski, D., Phillips, P.C.B., Schmidt, P. & Shin, Y., 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54 (1-3), 159-178.
- Lafuente, J. & Novales, A., 2003. Optimal hedging under departures from the cost-of-carry valuation: Evidence from the Spanish stock index futures market. *Journal of Banking & Finance*, 27 (6), 1053-1078.
- Lamoureux, C.G. & Lastrapes, W.D., 1990. Persistence in variance, structural change, and the GARCH model. *Journal of Business & Economic Statistics*, 8 (2), 225-234.
- Laurent, S., Rombouts, J. & Violante, F., 2009. On Loss Functions and Ranking Forecasting Performances of Multivariate Volatility Models. *Cirano discussion paper*, 2009-45.
- Lee, H.T. & Yoder J.K., 2007a. A bivariate Markov regime switching GARCH approach to estimate time varying minimum variance hedge ratio. *Applied Economics*, 39 (10), 1253-1265.
- Lee H.T. & Yoder J.K., 2007b. Optimal hedging with a regime-switching time-varying correlation GARCH model. *Journal of Futures Markets*, 27 (5), 495-516.
- Lee, K., Ni, Shwan & Ratti, R.A., 1995. Oil shocks and the macroeconomy: The role of price variability. *The Energy Journal*, 16 (4), 39-56.

- Lee, T.H., 1994. Spread and volatility in spot and forward exchange rates. *Journal of International Money and Finance*, 13 (3), 375-383.
- Lewis, M. & Noel, M. 2010. The speed of gasoline price response in markets with and without Edgeworth cycles. *The Review of Economics and Statistics*, forthcoming.
- Li, M.-Y. L. & Lin H.-W.W., 2004. Estimating value-at risk via Markov switching ARCH models: An empirical study on stock index returns. *Applied Econometrics Letters*, 11 (11), 679-691.
- Lien, D., 1996. The effect of the cointegration relationship on futures hedging: A note. *Journal of Futures Markets*, 16 (7), 773-780.
- Lien, D. & Tse, Y.K., 2002. Some recent developments in futures hedging. *Journal of Economic Surveys*, 16 (3), 357-396.
- Lin, B.H. & Yeh, S.K., 2000. On the distribution and conditional heteroscedasticity in Taiwan stock prices. *Journal of Multinational Financial Management*, 10 (3-4), 367-395.
- Lin, S.X. & Tamvakis, M.N., 2001. Spillover effects in energy futures markets. *Energy Economics*, 23 (1), 43-56.
- Lin, S.X. & Tamvakis, M.N., 2010. OPEC announcements and their effects on crude oil prices. *Energy Policy*, 38 (2), 1010-1016.
- Litterman, R. & J. Scheinkman, 1991. Common factors affecting bond returns. *Journal of Fixed Income*, 1 (1), 54-61.
- Litzenberger, R. & Rabinowitz, N., 1995. Backwardation on oil futures markets: Theory and empirical evidence. *Journal of Finance*, 50 (5), 1517-1545.
- Ljung, M. & Box, G., 1978. On a measure of lack of fit in time series models. *Biometrika*, 65 (2), 297-303.
- Lopez, J.A., 1999. Methods for evaluating value-at-risk estimates. *Federal Reserve Bank of New York, Economic Policy Review*, 2, 3-17.
- Luginbuhl, R. & de Vos, A., 2003. Seasonality and Markov switching in an unobserved component time series model. *Empirical Economics*, 28 (2), 365-386.
- Maheu, J.M. & McCurdy, T.H., 2000. Identifying bull and bear markets in stock returns. *Journal of Business and Economic Statistics*, 18 (1), 100-112
- Mandelbrot, B., 1963. The variation of certain speculative prices. *Journal of Business*, 36 (4), 394-419.
- Manoliu, M. & Tompaidis, S., 2002. Energy futures prices: Term structure models with Kalman filter estimation. *Applied Mathematical Finance*, 9 (1), 21-43.

- Marcucci, J., 2005. Forecasting volatility with regime switching GARCH models. *Studies in Nonlinear Dynamics & Econometrics*, 9 (4), Article 6, 1-53.
- Marimoutou, V., Raggad, B. & Trabelsi, A., 2009. Extreme value theory and value at risk: Application to oil market. *Energy Economics*, 31 (4), 519-530.
- Markowitz H.M., 1952. Portfolio selection. *Journal of Finance*, 7 (1), 77-91.
- Marsh, I.W., 2000. High frequency Markov switching models in the foreign exchange market. *Journal of Forecasting*, 19 (2), 123-134.
- Maslyuk, S. & Smyth, R. 2008. Unit root properties of crude oil spot and futures prices. *Energy Policy*, 36 (7), 2591-2600.
- Matas-Mir, A. & Osborn, D.R., 2004. Does seasonality change over the business cycle? An investigation using monthly industrial product series. *European Economic Review*, 48 (6), 1309-1332.
- Mayfield, E.S., 2004. Estimating the market risk premium. *Journal of Financial Economics*, 73 (3), 465-496.
- McLachlan, G.J., 1987. On bootstrapping the likelihood ratio test statistic for the number of components in a normal mixture. *Applied Statistics*, 36 (3), 318-324.
- McNeil, A.J. & Frey, R., 2000. Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach. *Journal of Empirical Finance*, 7 (3), 271-300.
- Mello, A. & Parsons, J., 1995. The maturity structure of a hedge matters: Lessons from the Metallgesellschaft debacle. *Journal of Applied Corporate Finance*, 8 (1), 106-120.
- Milonas, N. & Henker, T., 2001. Price spread and convenience yield behaviour in the international oil market. *Applied Financial Economics*, 11 (1), 23-36.
- Moench, E., 2008. Forecasting the yield curve in a data-rich environment: A no-arbitrage factor-augmented VAR approach. *Journal of Econometrics*, 146 (1), 26-43.
- Moosa, I.A., 2002. Price discovery and risk transfer in the crude oil futures market: Some structural time series evidence. *Economic Notes by Banca Monte dei Paschi di Siena SpA*, 31 (1), 155-165.
- Morana, C., 2001. A semi-parametric approach to short-term oil price forecasting. *Energy Economics*, 23 (3), 325-338.
- Moskowitz, T.J., 2003. An analysis of covariance risk and pricing anomalies. *The Review of Financial Studies*, 16 (2), 417-457.
- Mount, T.D., Ning, Y. & Cai, X., 2006. Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Economics*, 28 (1), 62-80.

- Myers, R.J. & Thompson, S.R., 1989. Generalized optimal hedge ratio estimation. *American Journal of Agricultural Economics*, 71 (4), 858-868.
- Naik, V., 1993. Option valuation and hedging strategies with jumps in the volatility of asset returns. *Journal of Finance*, 48 (5), 1969-1984.
- Naik, V. & Lee, M.H., 1997. Yield curve dynamics with discrete shifts in economic regimes: Theory and estimation. Working paper, University of British Columbia.
- Narayan, P. & Narayan, S., 2007. Modelling oil price volatility. *Energy Policy*, 35 (12), 6549-6553.
- Newey, W. & West, K., 1987. A simple positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55 (3), 703-708.
- Ng, V. & Pirrong, S.C., 1994. Fundamentals and volatility: Storage, spreads, and the dynamics of metals prices. *Journal of Business*, 67 (2), 203-230.
- Ng, V. & Pirrong, S.C., 1996. Price dynamics in refined petroleum spot and futures markets. *Journal of Empirical Finance*, 2 (4), 359-388.
- Noel, M., 2007. Edgeworth price cycles, cost based pricing and sticky pricing in retail gasoline markets. *The Review of Economics and Statistics*, 89 (2), 324-334.
- Nomikos, N.K. & Pouliasis, P.K., 2011. Forecasting petroleum futures markets volatility: The role of regimes and market conditions. *Energy Economics*, 33 (2), 321-337.
- Ohana, S., 2010. Modeling global and local dependence in a pair of commodity forward curves with an application to the US natural gas and heating oil markets. *Energy Economics*, 32 (2), 373-388.
- Osterwald-Lenum, M., 1992. A note with the quantiles of the asymptotic distribution of the ML cointegration rank test statistics. *Oxford Bulletin of Economics and Statistics*, 54 (3), 461-472.
- Palm, F.C. & Vlaar, P.J.G., 1997. Simple diagnostic procedures for modeling financial time series. *Allgemeines Statistisches Archiv*, 81, 85-101.
- Pelletier, D., 2006. Regime switching for dynamic correlations. *Journal of Econometrics*, 131, 445-473.
- Pelletier, R., 1983. Contracts that don't expire aid technical analysis. *Commodities*, March, 71-75.
- Perez-Quiros, G. & Timmermann, A., 2000. Firm size and cyclical variations in stock returns. *Journal of Finance*, 55 (3), 1229-1262.
- Phillips, P.C.B. & Perron, P., 1988. Testing for a unit root in time series regressions. *Biometrika*, 75 (2), 335-346.

- Pindyck, R.S., 1991. Irreversibility, uncertainty, and investment. *Journal of Economic Literature*, 29 (3), 1110-1148.
- Pindyck, R.S., 2004a. Volatility and commodity price dynamics. *Journal of Futures Markets*, 24 (11), 1029-1047.
- Pindyck, R.S., 2004b. Volatility in natural gas and oil markets. *Journal of Energy and Development*, 30 (1), 1-21.
- Pirrong, S. C., 1997. Metallgesellschaft: A prudent hedger ruined, or a wildcatter on NYMEX? *Journal of Futures Markets*, 17 (5), 543-578.
- Politis, D.N. & Romano, J.P., 1994. The stationary bootstrap. *Journal of the American Statistical Association*, 89 (428), 1303-1313.
- Quandt, R.E., 1958. The estimation of the parameters of a linear regression system obeying two separate regimes. *Journal of the American Statistical Association*, 53 (284), 873-880.
- Raymond, J.E. & Rich, R.W., 1997. Oil and the macroeconomy: A Markov state-switching approach. *Journal of Money, Credit and Banking*, 29 (2), 193-213.
- Reisman, H. & Zohar, G., 2004. Short term predictability of the term structure. *Journal of Fixed Income*, 14 (3), 7-14.
- Rydén, T., Teräsvirta, T. & Åsbrink, S., 1998. Stylized facts of daily returns series and the hidden Markov model. *Journal of Applied Econometrics*, 13 (3), 217-244.
- Sadorsky, P., 1999. Oil price shocks and stock market activity. *Energy Economics*, 21 (5), 449-469.
- Sadorsky, P., 2002. Time-varying risk premiums in petroleum futures prices. *Energy Economics*, 24 (6), 539-556.
- Sadorsky, P., 2003. The macroeconomic determinants of technology stock price volatility. *Review of Financial Economics*, 12 (2), 191-205.
- Sadorsky, P., 2006. Modeling and forecasting petroleum futures volatility. *Energy Economics*, 28 (4), 467-488.
- Samuel, Y.M.Z., 2008. Value at risk and conditional extreme value theory via Markov regime switching models. *Journal of Futures Markets*, 28 (2), 155-181.
- Samuelson, P.A., 1965. Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 6 (2), 41-49.
- Sarma, M., Thomas, S. & Shah, A., 2003. Selection of VaR models. *Journal of Forecasting*, 22(4), 337-358.
- Sarno, L., & Valente, G., 2000. The cost of carry model and regime shifts in stock index futures markets: An empirical investigation. *Journal of Futures Markets*, 20 (7), 603-624.

- Schwartz, E.S., 1997. The Stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52 (3), 923-973.
- Schwartz, E.S. & Smith, J.E., 2000. Short-term variations and long term dynamics in commodity prices. *Management Science*, 46 (7), 893-911.
- Schwarz, G., 1978. Estimating the dimension of a model. *Annals of Statistics*, 6 (2), 461-464.
- Schwarz, T.V. & Szakmary, A.C., 1994. Price discovery in petroleum markets: Arbitrage, cointegration, and the time interval of analysis. *Journal of Futures Markets*, 14 (2), 147–167.
- Silvapulle, P. & Moosa, I.A., 1999. The relationship between spot and futures prices: Evidence from the crude oil market. *Journal of Futures Markets*, 19 (2), 175-193.
- Smith, C. & Stulz, R., 1985. The determinants of firms' hedging policies. *Journal of Financial and Quantitative Analysis*, 20 (4), 391-405.
- Smith, J.L., 2005. Inscrutable OPEC? Behavioural tests of the cartel hypothesis. *The Energy Journal*, 26 (1), 51-82.
- Stein, J.L., 1961. The simultaneous determination of spot and futures prices. *The American Economic Review*, 51 (5), 1012-1025.
- Sullivan, R., Timmermann, A. & White, H., 1999. Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance*, 54 (5), 1647-1691.
- Sullivan, R. & White, H., 1998. Finite sample properties of the bootstrap reality check for data-snooping: A Monte Carlo assessment. QRDA, LLC Technical Report, San Diego.
- Telser, L.G. 1958. Futures trading and the storage of cotton and wheat. *Journal of Political Economy*, 66 (3), 233- 55.
- Teräsvirta, T, 1994. Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association*, 89 (425), 208-218.
- Tillmann, P., 2004. Cointegration and regime-switching risk premia in the U.S. term structure of interest rates. Bonn Econ Discussion Papers, University of Bonn.
- Tolmasky, C. & Hindanov, D., 2002. Principal components analysis for correlated curves and seasonal commodities: The case of the petroleum market. *Journal of Futures Markets*, 22 (11), 1019-1035.
- Tong, H., 1978. On a threshold model. In: Chen, D.H.C. (Eds.), *Pattern Recognition and Signal Processing*, Amsterdam: Sijhoff and Noordoff, 101-141.
- Tong, H., 1983. *Threshold models in nonlinear time series analysis*. Springer-Verlag.
- Tsay, R.S., 1989. Testing and modelling threshold autoregressive processes. *Journal of the American Statistical Association*, 84 (405), 231-40.

- Viswanath, P.V., 1993, Efficient use of information, convergence adjustments, and regression estimates of hedge ratios. *Journal of Futures Markets*, 13 (1), 43-53
- Vlaar, P.J.G. & Palm, F.C., 1993. The message in weekly exchange rates in the European monetary system: Mean reversion, conditional heteroscedasticity, and jumps. *Journal of Business and Economic Statistics*, 11 (3), 351–360.
- Vo, M.T., 2009. Regime-switching stochastic volatility: Evidence from the crude oil market. *Energy Economics*, 31 (5), 779-788.
- Wallis, K.F., 1974. Seasonal adjustment and relations between variables. *Journal of the American Statistical Association*, 69 (345), 18-31.
- Wei, Y., Wang, Y. & Huang, D., 2010. Forecasting crude oil market volatility: Further evidence using GARCH-class models. *Energy Economics*, 32 (6), 1477-1484.
- Weiner, R. J., 1991. Is the world oil market “One Great Pool”? *The Energy Journal*, 12 (3), 95-107.
- White, H., 2000. A reality check for data snooping. *Econometrica*, 68 (5), 1097-1126.
- Wilson, B., Aggarwal, R., & Inclan C., 1996. Detecting volatility changes across the oil sector. *Journal of Futures Markets*, 16 (3), 313-330.
- Working, H., 1949. The theory of the price of storage. *American Economic Review*, 39 (6), 1254-1262.
- Ye, M., Zyren, J. & Shore, J., 2006. Forecasting short-run crude oil price using high- and low-inventory variables. *Energy Policy*, 34 (17), 2736-2743.
- Yu, W-C. & Zivot, E., 2010. Forecasting the term structures of treasury and corporate yields using dynamic Nelson-Siegel models. *International Journal of Forecasting*, 27 (2), 579-591.