



A University of Sussex PhD thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details



The Joint Modelling of Energy Forward Curves

Yufi Pak

Submitted for the Degree of Doctor of Philosophy
University of Sussex
June 2022

Declaration

I hereby declare that this thesis has not been and will not be submitted in whole or in part to another University for the award of any other degree.

Signature:

Yufi Pak

Acknowledgements

A PhD is often described as a journey – just like one’s life with a sequence of choices and encounters. I would like to thank my supervisors, Michael Coulon and Andreas Kaeck, who guided my academic journey with their expert knowledge and patience. Becoming their family friend enriched my non-academic journey also. Thinking of my journey at Sussex, I can’t thank enough my dearest friend Carol Alexander and her family for their love and friendship. Speaking of friendship, I was lucky to have had room 230 colleagues who inspired and motivated me to go through this rigorous journey. Last but not least, I would like to thank my parents who gave me a wonderful life.

Abstract

Analogous to yield curves in fixed income markets, commodity forward curves reflect market participants' views on future price levels and are the main inputs for pricing, risk management, and project evaluations. This thesis focuses on a factor-estimation method incorporating the joint dynamics of multiple commodity forward curves for a term structure model. First, we introduce PCA on PCA as the main tool to formulate the factor-volatility functions with commonality, extending the Heath-Jarrow-Morton (1992) model for multi-commodity modelling. The proposed factor estimation method is intuitive and easy to implement as a direct extension of ordinary PCA. We demonstrate the estimation procedure of common eigenstructures and analyse the loadings to give economic interpretations to the identified common factors. Second, we apply our common factor model for the pricing of commodity spread derivatives. Our contribution includes the option pricing formula when one of the underlying assets is denominated in foreign currency units, which has not been considered carefully in previous commodity literature. Third, we analyse hedge ratios and their effectiveness with and without assuming common factors across multiple forward curves. Empirical studies on each topic are carried out for European energy commodities or marine fuel products. We find that the existence of common factors lowers the option prices and improves the minimum-variance hedge where appropriate. These results have important implications for commodity producers, traders and financial institutions that trade highly dependent underlying assets: neglecting common factors may result in economic losses caused by mispricing financial contracts or mistreating inherent risks.

Contents

List of Tables	ix
List of Figures	xi
Overview	1
1 Classical Theory and Modelling Approaches	5
1.1 Economics of Commodities	5
1.2 Price Modelling	7
1.3 New Developments in Multi-commodity Models	11
2 Notes on Commodity Futures	15
2.1 Commodity Futures	15
2.2 Interpolation of Data	19
3 Principal Component Analysis on Multi-group Data	21
3.1 Introduction	21
3.2 Principal Component Analysis (PCA)	22
3.3 PCA on Multi-group Data	23
3.3.1 Conventional Approaches	23
3.3.2 PCA on PCA (PPCA)	25
3.4 Empirical Study	29
3.4.1 Data	29
3.4.2 Results	36
3.4.3 Interpretation	41
3.5 Summary	47
4 The Joint Modelling of Commodity Forward Curves	48
4.1 Introduction	48

4.2	Forward Curve Models	49
4.2.1	The PC Model	49
4.2.2	The PPC Model	50
4.3	Implementation	54
4.3.1	Parameter Estimation	54
4.3.2	Simulation	62
4.3.3	Results	64
4.4	Summary	71
5	Pricing Commodity Spread Derivatives	72
5.1	Introduction	72
5.2	Adopting Common Factor Structures in Margrabe (1978)	74
5.2.1	A Brief Sketch of the Margrabe Formula	74
5.2.2	Factor Volatilities in the Margrabe Formula	77
5.3	Exchange Options on Foreign to Domestic Spreads	79
5.3.1	Quanto Exchange Options	79
5.3.2	Cross-currency Exchange Options	82
5.4	Pricing	85
5.4.1	Results for Spark Spreads	86
5.4.2	Results for Spreads Including a Foreign Asset	90
5.5	Summary	97
6	Common Factor Hedge in the Shipping Market	98
6.1	Introduction	98
6.2	Bunker Price Hedge	100
6.2.1	Problem Formulation	100
6.2.2	Hedging from Consumers' Perspectives	102
6.3	Model	103
6.4	Empirical Study	106
6.4.1	Preliminary Analysis	106
6.4.2	Parameter Estimation	113
6.5	Cashflow Hedge	116
6.6	Summary	121
	Concluding Remarks	122

References	125
A Tables for Chapter 5	132

List of Tables

2.1	Contract specifications of commodity futures.	18
3.1	Quote patterns of UK gas, Dutch gas, and UK power futures at ICE.	33
3.2	Quote patterns of Dutch power futures at ICE.	34
3.3	Explanatory powers of PCA.	39
3.4	Cumulative explanatory powers of PPCA.	42
3.5	Ranking of category-wise contributions to PPCs.	44
3.6	Comparison of original and sparse PPC loading vectors.	46
4.1	Exponentially fitted model parameters for the PC(3) and PPC(12, 3, 1, 1) models.	57
4.2	Exponentially fitted model parameters for the PPC(12, 3, 2, 1) model.	58
4.3	Exponentially fitted model parameters for the PPC(12, 3, 3, 1) model.	59
4.4	PPC factor correlations.	60
4.5	Cross-correlation matrix of three PCs.	61
4.6	Five commodity pairs.	67
5.1	Five types of commodity spreads.	86
5.2	Option prices of spark spreads.	87
5.3	Standard deviations of spark spreads at option expiration.	87
5.4	The option prices of 360-day spark spread by various fuel efficiency rates.	88
5.5	The (at-the-money) option prices of 360-day spark spread.	90
5.6	Standard deviations of UK – Dutch spreads at option expiration.	92
5.7	Option prices of quanto exchange options.	92
5.8	Correlation coefficients between the spot FX rate and forward contracts.	94
5.9	The variance of cross-currency spreads at option expiration.	95
5.10	The prices of cross-currency exchange options.	96

6.1	Summary statistics of interpolated fuel prices and log returns.	109
6.2	The cumulative percentage of total variation explained by PCA.	110
6.3	Cross-correlation matrix of principal components.	112
6.4	The total variation explained by PPCA.	114
6.5	Fitted parameter values of the PC and PPC models.	115
6.6	Comparison of variance and hedge effectiveness between models.	117
A.1	Option prices of spark spreads.	133
A.2	Standard deviations of spark spreads at option expiration.	133
A.3	The option prices of 360-day spark spread.	133
A.4	Standard deviations of UK – Dutch spreads at option expiration.	134
A.5	Option prices of quanto exchange options.	134
A.6	The variance of cross-currency spreads at option expiration.	135
A.7	The prices of cross-currency exchange options.	136

List of Figures

3.1	Time series of the first-month and fourth-quarter prices.	30
3.2	The first-month and fourth-quarter log returns.	32
3.3	Interpolated forward curves with fixed time to maturities.	36
3.4	Term structure of volatilities.	37
3.5	The first three eigenvectors of fixed time-to-maturity forward curves by PCA.	38
3.6	Eigenvalues and eigenvectors by PPC on the cross-correlation matrix of PCs.	40
3.7	Eigenvalues and eigenvectors by PPC on the cross-covariance matrix of PCs.	41
4.1	Factor volatilities and exponentially fitted curves.	56
4.2	Simulated marginal distributions (local volatility fitting).	65
4.3	Simulated marginal distributions (local covariance fitting).	66
4.4	The means of simulated UK power - UK gas spread at the maturity of contracts (local volatility fitting).	67
4.5	Standard deviations of spread payoffs at the maturity of forward contracts.	68
4.6	Standard deviations of spread payoffs 90 days before the maturity of forward contracts.	69
4.7	Standard deviations of spread payoffs 90 days after the inception of forward contracts.	70
5.1	The option prices of UK and Dutch spark spreads.	87
5.2	The option price of 360-day spark spread by various fuel efficiency rates.	89
5.3	Spot FX rate and its rolling correlation with UK assets.	91
5.4	Quanto options with varying correlation coefficients between the spot FX rate and the UK assets.	93
5.5	The prices of cross-currency exchange options.	96
6.1	Interpolated forward curves with fixed time to maturities.	107
6.2	Log returns of 30-day and 510-day fixed time-to-maturity contracts.	108

6.3	Within-curve correlations of forward price log returns.	110
6.4	The first three eigenvectors of fixed time-to-maturity forward curves.	111
6.5	Time series of principal components in level.	112
6.6	PPC eigenvectors.	113
6.7	Standardised cashflows of monthly fuel consumption and hedge positions. . .	119
6.8	90% confidence intervals of hedged and unhedged cashflows by model. . . .	120

Overview

Pricing models evolve along with market transitions. For energy commodities, the market transition accelerated in the mid-2000s, triggered by the advent of technology, liberalisation of energy sectors, regulatory requirements, rising concerns for climate change worldwide, and so forth. In response to those shifts in energy markets, or perhaps foreseeing their structural changes in the near future, we study the joint modelling of commodity forward curves, for which the market boundaries have been eroding.

Until recently, the prices of energy commodities were regarded as market/commodity specific. Traditional pricing models consider the stochastic processes of marginal commodities separately and associate them by correlation. On the other hand, recent developments in energy markets, as seen in the extended pipeline connections in European gas markets and similar developments in electricity interconnectors, would have strengthened the linkage of price levels considerably due to tighter supply/demand conditions, implicitly suggesting that those substitutable commodities would be exposed to the same source of uncertainties rather than separate ones. Hence, our research studies the impact of common factors in the pricing and hedging of energy commodities.

This thesis is a self-contained piece of work, covering the topics of the fundamentals of commodities, pricing theories, estimation methods, and applications. Chapter 1, titled ***Classical Theory and Modelling Approaches***, revisits classical theories and pricing models for commodities to relate single-commodity models to multi-commodity models introduced in the recent literature. This literature review identifies three subgroups of multi-commodity models such that i) spot price models with common state variables, ii) cointegrated forward price models, iii) forward price models with commonality in their factor volatilities. This chapter provides clear positioning of our research that belongs to the third approach, which has the potential to overcome the general challenges of multi-commodity models that tend to be mathematically rigorous and computationally expensive.

Generally speaking, the analysis of commodity forwards and futures is not straightfor-

ward. The contracts are written on commodities (indexes) that refer to specific calendar days/months/years in a year. The data sets are discontinuous as the contracts expire after the period, requiring data interpolation or concatenation for mathematical modelling and statistical analysis. To address such matters, the sections in Chapter 2, titled ***Notes on Commodity Futures***, provide useful information to clarify the specifications of futures contracts and explain the rationale for the notation and interpolation method to analyse the data. Albeit short, this chapter can be considered essential supplementary material for this thesis.

Chapter 3, titled ***Principal Component Analysis on Multi-group Data***, serves as the foundation of the subsequent chapters. It suggests a factor estimation method for multi-commodity forward curves, Principal Component Analysis on Principal Component Analysis (PPCA), a two-stage procedure. In the first-stage PCA, the dimensionality of forward curves is reduced, and the orthogonal factors are identified at a marginal level. In the second-stage PCA, another PCA is applied to the cross-correlation matrix of principal components to find the common orthogonal axis, from which we define common latent factors. Working backwards from the common to marginal spaces, we derive an expression for marginal covariance matrices consisting of common and marginal eigenvalues/vectors. We apply the proposed factor estimation method to UK gas, Dutch gas, UK power, and Dutch power forward curves to quantify their common factor structures. We investigate the implications of the common latent factors with our own ranking algorithm and a shrinkage method in linear regression analysis by [Zou et al. \(2006\)](#). Interestingly, both approaches suggest that the most significant two common factors relate to the level and slope of the four forward curves, with the third common factor being unique to Dutch gas's curvature factor.

Chapter 4, titled ***The Joint Modelling of Commodity Forward Curves***, documents the model development procedure for a multi-commodity forward curve model using PPCA. The proposed PPCA-based model is an extension of [Heath et al. \(1992\)](#) consisting of a number of common orthogonal factors and one idiosyncratic factor that describe the comovement of closely-linked commodities. In addition, it considers commodity-specific dynamics via correlated idiosyncratic factors between forward curves. The model building starts with applying PPCA to estimate the common factors. Subsequently, the idiosyncratic factors and their correlations are calibrated to the covariance matrices estimated by PCA. Using a PCA-based model as the calibration target makes it possible for our PPCA-based model to have the same level of explanatory power as the benchmark model

for marginal forward curves, an attractive feature since PCA-based forward curve models have widely been used in practice. In the empirical section, we estimate the model parameters from the same European energy futures data and conduct simulation studies to analyse the joint and marginal distributions estimated by the models with a varying number of common factors. A general tendency of the PPCA-based models revealed by the results is that two or three common factors sufficiently reproduce the marginal distributions of forward curves compared to a PCA-based benchmark model with no common factor. Moreover, the PPCA-based models stabilise the dispersion of commodity spreads for some commodity pairs, as a result of enhanced common factor modelling.

In Chapter 5, titled ***Pricing Commodity Spread Derivatives***, we turn our attention to derivatives pricing by demonstrating the adaptability of the common factor structures to the [Margrabe \(1978\)](#) formula that gives the price of an option to exchange one asset for another. The complexity of the Margrabe formula is increased for the case when one of the underlying assets is denominated by foreign currency units. The latter type of options, named ‘quanto exchange options’ or ‘cross-currency exchange options’, accommodate both cases for exchange rate conversion: (i) fixed by contract in advance, or (ii) floating with the market FX rate. The empirical section presents the pricing results for power to gas (spark) and foreign to domestic spreads while investigating the impact of having common factors in the option pricing formula. As for spark spreads, our common factor model calculates the option prices generally lower than a non-common factor model, as the simulation outputs in the previous chapter indicated. It is also explained that the fuel efficiency of production facilities influences the common factors’ impacts on option prices; the price differences appear more significantly for at-the-money options. The results for foreign - domestic spreads show the negligible impact of FX adjustments on the quanto option prices. On the other hand, FX adjustments are found to be important for cross-currency exchange options as they directly affect the spread variance, therefore the option prices.

Chapter 6, titled ***Common Factor Hedge in the Shipping Market***, which is the last chapter in this thesis, focuses on a cross-hedge problem in maritime finance. In our study, a long position on a vessel fuel is cross-hedged by proxy fuel futures, for which more liquid and established markets are available. We construct this hedge problem from a consumer’s point of view and differentiate ourselves from the problem of portfolio managers who aim to maximise their profits by a dynamic hedge. We implement the same models introduced in earlier chapters to three subgroups of vessel fuels to calculate the minimum-

variance hedge ratios with and without common factors. Hedge effectiveness measures suggest the improved performance of the cross hedge with the presence of common factors. In addition, it is shown that the cross hedge achieves the most variance reduction when the underlying fuel contract is hedged only by one proxy fuel instead of multiple proxy fuels for the period under study. The confidence intervals of hedged cashflows are also in line with these observations; the single proxy hedge best estimates the price risk.

Finally, in the ***Concluding Remarks***, we summarise the contributions of our research and make suggestions for future research.

Chapter 1

Classical Theory and Modelling Approaches

1.1 Economics of Commodities

The main determinants of commodity prices are supply, demand, and inventory. The classical modelling of commodity prices stems from the *Theory of Storage*, which explains the relationship between the spot price, forward price, and level of inventory for storable commodities. When inventories are scarce, relatively small changes in production or consumption can significantly impact prices and cause high volatility, and vice versa when inventories are ample. Therefore, there exists an inverse relationship between the level of inventory and price volatility ([Fama and French, 1987](#); [Ng and Pirrong, 1994](#); [Geman and Nguyen, 2005](#)), related to market uncertainties.

To date, a number of authors have studied the Theory of Storage, extending the original work of [Kaldor \(1939\)](#). The classical theory investigates the reasons why producers of commodities store inventories even when the cost of storage is very high and introduces the term *convenience yield*, which is the benefit that belongs to the holder of commodities to meet unexpected demand or to take advantage of price rises in the future ([Working, 1949](#); [Brennan, 1958](#); [Telser, 1958](#)).

The fundamental cost-of-carry relationship summarises these concepts into a parsimonious equation, where the forward price is equal to the spot price compounded by the carry cost, including the cost of financing (interest rate), storage cost, and convenience yield. When the sum of interest rate and storage cost exceeds the benefit of holding a commodity, the commodity price in the future is higher than the current price of the commodity. Hence, the forward curve will form an upward sloping shape known as *contango*. By con-

trast, when the benefit of holding a commodity outweighs the interest rate plus storage cost, the forward curve will be downward sloping in shape, known as *backwardation*.

The *Theory of Normal Backwardation* argues that forward prices are downward biased estimates of the expected spot price in normal market regimes (Keynes, 1930; Hicks, 1946). This theory is built upon the recognition that the market is net short, meaning that the number of hedgers (producers) exceeds the number of speculators.¹ Therefore, the forward price should include a risk premium to compensate the risk-taking behaviour of speculators. On the contrary, Bouchouev (2012) argues that normal backwardation no longer holds in the current market environment for crude oil due to changes in the market microstructure. In his view, investors are the hedgers who seek mitigation of risks in the market. As a consequence, the *normal contango* appears in a market where investors pay the risk premium.

Generally speaking, the volatility of a commodity price is low when a contract matures in the distant future and is high when the contract matures in the near future since the price will be more susceptible to the arrival of news. On the other hand, the price of a distant maturity contract will not respond equally to the same news as the nearby contract, as the markets' expectations for increased production or availability of substitution will stabilise the price fluctuations in the long run. These characteristics of volatilities are known as the *Samuelson Effect* (Samuelson, 1965). The Samuelson Effect in volatility relates to the concept of mean reversion in the price dynamics of commodities in line with general economic theory: the price of goods eventually reaches equilibrium by supply/demand adjustments. Towards equilibrium, the distribution of a mean-reverting process stabilises, and the forward volatility decreases.

Bessembinder et al. (1995) conduct an empirical study to examine the existence of mean-reversion in energy, agricultural, and financial markets, finding the most substantial degree of mean-reversion in crude oil prices. By contrast, the authors find weak evidence of mean-reversion for financial markets. Geman (2009) notes the prominence of the Samuelson Effect in energy markets, in which the volatilities typically increase rapidly several months before the expiry of contracts.

Some commodities have periodic patterns in production or consumption known as seasonality, which could be regarded as an additional explanatory variable in the modelling of commodity price processes (Fama and French, 1987). While seasonality of commodities

¹Consumers can also be the hedgers since they enter into derivative contracts in order to hedge against adverse price movements in the purchase of commodities. However, the classical theory regards that the hedging needs are much more significant for commodity producers than for consumers.

may exhibit yearly, half-yearly, monthly, or half-monthly cycles (Milonas, 1991; Sørensen, 2002), the periodicity of seasonality for energy commodities may increase daily, hourly, or intra-hourly. For instance, the price of electricity depends on the day of a week (e.g. weekday or weekend) or time in a day (e.g. baseload or peakload). The price changes or returns of seasonal commodities fluctuate for periods where excess demand or supply shortage occurs. Therefore, the volatility may also exhibit seasonal patterns, as evidenced by Lucia and Schwartz (2002) in the study of the Nordic spot electricity market.

As for futures contracts, Swindle (2014) categorises two types of seasonal volatilities. One is the seasonal volatilities associated with the maturity month, e.g. the volatility of contracts maturing in January is always greater than the volatility of contracts maturing in August for the UK natural gas market, where the demand for heating increases during winter. By contrast, the second type of seasonality relates to the current time, e.g. the average volatility of futures contracts traded in January is consistently higher than the average volatility of futures traded in August.

Whether physical or financial, the prices of commodity contracts are linked to the underlying physical assets. Therefore, it is essential for us to consider the economics of commodities in the formulation of mathematical models.

1.2 Price Modelling

An early generation of commodity price models utilises the mathematical framework developed for equities, an extension of Black and Scholes (1973) and Merton (1973). While the models ensure the non-negativity of an underlying asset and benefit from tractable formulae in derivatives pricing, the models based on Geometric Brownian Motion (GBM) cannot describe the mean-reverting behaviour of commodity price dynamics. Later models such as the one-factor model of Schwartz (1997) circumvent the problem by specifying a model based on the mean-reverting lognormal process.

Gibson and Schwartz (1990) and Schwartz (1997) develop a two-factor model in which the convenience yield acts as a mean-reverting second factor in addition to the first factor following a GBM.² Further, Schwartz (1997) adds a risk-free interest rate into the three-factor version of his model, nesting the Vasicek (1977) model. Mathematically equivalent to Gibson and Schwartz (1990) and Schwartz (1997)'s two-factor model, however more intuitive to interpret, Schwartz and Smith (2000) propose a two-factor model in which two

²Gibson and Schwartz (1990) propose the preceding version of Schwartz (1997)'s two-factor model. The difference is that Schwartz (1997) provides an analytical solution whereas Gibson and Schwartz (1990) do not.

latent state variables represent the short-term and long-term dynamics of a commodity spot price. The underlying assumption is that the short-term factor fluctuates around the mean level due to temporary shocks, whereas the equilibrium level will evolve in the long run due to permanent changes in market fundamentals. Therefore, these short-term and long-term dynamics are modelled via an Ornstein-Uhlenbeck (OU) process and an Arithmetic Brownian Motion, respectively.

The convenience of the [Schwartz and Smith \(2000\)](#) type model is the flexibility in expressing the log spot price as the sum of state variables that have their own dynamics. It is also possible to assign more than one factor per state variable as in [Cartea and Figueroa \(2005\)](#), in which the authors account for a jump effect in the mean-reverting spot price of electricity with an OU process and a Poisson process. In addition, the additivity of the Gaussian state variables can incorporate the seasonality of underlying assets relatively easily into the model, for example, by using a deterministic seasonal component ([Sørensen, 2002](#); [Lucia and Schwartz, 2002](#)). [Manoliu and Tompaidis \(2002\)](#) apply the [Schwartz and Smith \(2000\)](#) model to the analysis of natural gas forward prices, in which the authors model the seasonal component by a periodic step function. Under a similar model specification, [Cartea and Williams \(2008\)](#) estimate the seasonal component by a second-order Fourier series to study the UK natural gas market in conjunction with the storage problem.

Despite several advantages that the [Schwartz and Smith \(2000\)](#) type models offer, a challenge arises in estimating the unobservable state variables and their parameter values, as well as the market price of risk. Moreover, these models suffer from the general drawback of spot price models that the model implied forward prices, which are obtained as the conditional expectation of the spot price under an appropriate pricing measure, are not consistent with the market price of forward contracts.

An alternative approach, namely the forward price modelling, overcomes these problems. The idea originates in the Heath–Jarrow–Morton ([Heath et al., 1992](#), HJM hereafter) framework for fixed income securities, which focuses on forward rates and describes the evolution of bond prices by their volatilities. In an analogous commodity theory, the forward price of a commodity is expressed as the sum of factor volatilities under the risk-neutral pricing measure.³

To adapt the HJM framework to commodities price modelling, one can find the

³Under the assumption of constant interest rates and absence of credit risk, the price of forwards and futures, and the associated pricing measures are equivalent. We shall impose these assumptions and use the two terminologies ‘forward’ and ‘futures’ interchangeably throughout this thesis.

functional form of factor volatilities endogenously through associated spot price models or estimate the functional forms from the market observed term structure of forward prices/volatilities. In the former approach, when a single factor follows a GBM in a spot price specification (e.g. [Black and Scholes, 1973](#)), the application of Ito's lemma gives a constant factor volatility in the corresponding forward price model. Alternatively, suppose a spot price model is a one-factor model in which the log of the factor follows an OU process (e.g. [Schwartz, 1997](#) one-factor model). In that case, the constant volatility in the forward price model is exponentially scaled by a decay parameter, which is the same parameter as the speed of mean reversion in the spot price model. The decay parameter controls the rate of decline in the magnitude of volatility with respect to time to maturity in the forward price model; the volatility increases when the time to maturity of a forward contract approaches zero and vice versa when the maturity is far ahead in the future. This result sheds light on the relationship between the mean reversion in spot price and the Samuelson Effect in forward volatility.

In the latter approach, [Gabillon \(1991\)](#) proposes a parametric model that segments the factor volatilities for the short-end and long-end of commodity forward curves. Moreover, his model incorporates the *split personality*, a characteristic known for the forward curves of energy commodities for which the factors driving the short-end and long-end of forward curves have little impact on each other. [Swindle \(2014\)](#) extends such approaches to the modelling of seasonal volatility and introduces a model with exponentially scaled double volatility functions explicitly separating the seasonal spot volatility and maturity month related volatilities.

In the HJM framework, the number of factor volatilities and the functional forms can vary depending on the model assumptions. In the meantime, it is likely that more than one volatility function, with at least one exponentially decaying parameterisation, are required to replicate the term structure of volatilities for energy commodities; see [Clewlow and Strickland \(2000\)](#) who provide the generalisation of volatility functions for commodity forward curves. The HJM type of forward curve model belongs to a family of term structure models that simultaneously drives the whole set of forward contracts by a relatively small number of factors. With that regard, Principal Component Analysis (PCA) has been a popular technique that effectively captures the covariance structure of multiple data series in a reduced dimension.

In the finance literature, PCA is first introduced by [Litterman and Scheinkman \(1991\)](#) in the context of factor hedging for a portfolio of U.S. treasury securities. From the shape of

eigenvector (factor loadings) plots, the authors name the first three principal components (PCs) as the *level*, *steepness*, and *curvature* factors,⁴ which explain approximately 98% of the total variation in the U.S. yield curve. Cortazar and Schwartz (1994) adapt the PCA to the modelling of commodity forward curves, finding that the results of PCA on copper data are similar to the patterns of principal component factors for the U.S. treasury bonds in Litterman and Scheinkman (1991). Following the HJM framework, they describe the dynamics of copper futures by three orthogonal factor volatilities in the valuation of copper-linked contingent claims.

The resemblance of the first three principal component factors seems to be persistent across different types of commodity forward curves. At the same time, Clewlow and Strickland (2000) remark that seasonal commodities tend to demand additional PCs to sufficiently reproduce the original data structure; For example, in their study, natural gas requires five factors to explain 99% of the total variation, compared to four factors for crude oil. Koekebakker and Ollmar (2005) observe that the forward curves require more than ten factors to explain over 95% of variation for the Nordic electricity data. In response to their findings, Borovkova and Geman (2008) argue that the PCA should be applied to seasonal commodities with caution since seasonality could distort the outputs of PCA. The authors introduce two ways of normalising data set in prior to PCA, by excluding the deterministic seasonality from either the original price or returns data. Carmona and Coulon (2014) demonstrate an alternative method that uses an instantaneous volatility to normalise a covariance matrix to which PCA is applied.

Without using PCA, Chiarella et al. (2009) develop a two-factor regime-switching forward curve model where the transition from one state to the other state is characterised via a finite-state Markov Chain. In their model, the state-dependent factor volatilities nest a deterministic seasonality component modelled by a truncated Fourier series. The authors apply the Markov Chain Monte Carlo (MCMC) for the parameter estimation and discover that the estimated factors behave similarly to PCA's level and slope factors. More recently, Thompson (2016) proposes a four-factor parametric model inclusive of a stochastic seasonality factor and three volatilities representing the typical shape factors of PCA. The model specification starts from identifying the number of orthogonal factors and their functional forms using PCA, followed by the transformation of the uncorrelated model to a correlated counterpart for which the factor volatilities are parametrically fitted to the implied volatility of the Henry Hub natural gas options. Although these two HJM

⁴Other common names of the steepness factor are the *tilt* and *slope* factors. The curvature factor is also known as the *bending* factor.

type models do not directly obtain the model parameters from PCA, the estimation results indicate the crucial roles of typical principal component factors in the volatility functions.

1.3 New Developments in Multi-commodity Models

From a modelling perspective, the complexity surges when the underlying asset increases from a single commodity to multiple commodities because it requires describing the cross-market dependence structures in a multi-dimensional space. However, many multi-commodity models follow the fundamental assumptions and structures of single commodity models that are introduced in the previous sections.

Several authors utilise the [Schwartz and Smith \(2000\)](#) type model, in which the spot price is an exponential affine function of state variables, and assume one or more state variables are shared between commodities. To estimate common factor structures of a multi-dimensional spot price model, [Cortazar et al. \(2008\)](#) perform a special variant of PCA by [Flury \(1988\)](#)⁵ that investigates the common eigenstructure of covariance matrices between data groups. Their model aims to improve the pricing of futures for a commodity using the information of other commodity futures for which liquid and longer maturity data are available. Their multi-commodity models outperform single commodity models for two pairs of crude oil and petroleum products when fitting the out-of-sample data.

[Frikha and Lemaire \(2013\)](#) propose a tailored model for the joint dynamics of UK gas and power prices where the deseasonalised log spot prices are described by the sum of common and non-common OU processes under the physical measure. Their model considers the statistical property of time series data for the auto-correlation and cross-correlation of gas and power, to which the rate of mean-reversion parameters is calibrated. The authors assign the normal-inverse Gaussian distribution to the non-common OU process to reproduce the leptokurtic distribution of data and compute the value of a gas-fired power plant by simulation. The results suggest that the cross-correlation of gas and power does not significantly change the value of the power plant; however, it reduces the Value-at-Risk (VaR).

When time series data are cointegrated, the processes may wander apart in the short run but are tied together in the long run. In the commodity literature, a number of authors acknowledge the existence of a long-term relative price equilibrium when the assets are the input/output in a production process (e.g. crude oil and jet fuel) or substitutes of each other (e.g. natural gas and coal), just as for the same reasons why the price of com-

⁵[Flury \(1988\)](#)'s common PCA is explained in the next chapter.

modities mean reverts. Since the concept of mean reversion closely relates to stationarity in time series analysis, one could interpret cointegration as the long-term mean-reverting relationship between a spread of commodities.

[Paschke and Prokopczuk \(2009\)](#) develop a cointegrated multi-commodity spot price model where n -dimensional log prices are jointly described by one common and $n - 1$ correlated state variables. A commodity's responsiveness to these latent variables depends on the factor loadings estimated by the Kalman filter. In analysing crude oil, heating oil, and gasoline prices, the authors identify one common stochastic factor that affects the three commodities to a similar degree and the other factor that affects the commodities to a varying degree. Their common factor model estimates the distribution of an oil refinery value to be significantly narrower than a non-common factor model.

[Nakajima and Ohashi \(2012\)](#) introduce a cointegrating state variable to the [Gibson and Schwartz \(1990\)](#) two-factor model adapting the original work of [Duan and Pliska \(2004\)](#) to commodity markets. In contrast to the [Gibson and Schwartz \(1990\)](#) model in which the stochastic convenience yield mean reverts to its equilibrium level, their multi-commodity model assumes the existence of long-term linear relationships across cointegrated commodities to which the prices mean-revert under the risk-neutral measure. Contrary to [Duan and Pliska \(2004\)](#) who argue that the cointegration is only relevant when the volatility is stochastic for equity options, [Nakajima and Ohashi \(2012\)](#) claim that cointegration does influence the price of commodity derivatives even if the volatility is deterministic.

Similarly, [Farkas et al. \(2017\)](#) extend the [Schwartz and Smith \(2000\)](#) model to the multivariate space where the latent state variables are not only correlated but also cointegrated across commodities. The authors derive an expression for driftless forward price dynamics under the risk-neutral measure where the matrix of cointegrating vectors adjusts for the decay parameters of factor volatilities. Their empirical analysis of weekly time series data suggests that one common factor drives the long-term log-price levels of crude oil, heating oil, and gasoline in agreement with [Paschke and Prokopczuk \(2009\)](#). In addition, it is indicated by the Monte Carlo simulation that their model calculates lower prices for the energy spread options compared to other models without cointegration, with the presence of an upward-sloping term structure of correlations between the energy commodities.

Not surprisingly, the multi-commodity spot price models inherit the general drawbacks of spot price models, such as estimating unobservable state variables for multiple assets, often accompanied by involved estimation procedures resulting in a large set of parameters. For example, [Paschke and Prokopczuk \(2009\)](#) report 58 parameters for their six-factor

model for three commodities. Likewise, [Nakajima and Ohashi \(2012\)](#) require 55 parameters for a model of three commodities with two linear relationships. Moreover, one may realise that the forward prices obtained as the conditional expectation of the cointegrated spot prices are not cointegrated under the risk-neutral measure in the sense that the stationary linear relationships among assets are no longer linear or stationary in the corresponding forward price models.

To address the last point, [Benth and Koekebakker \(2015\)](#) propose a class of pricing measures under which both the log of spot and forward prices are cointegrated. Their marginal spot price model is exponentially affine in two state variables driven by the Arithmetic Brownian Motion (ABM) for the shared non-stationary process and the Continuous-time AutoRegressive Moving Average (CARMA) for the commodity-specific stationary process. The authors derive a cointegrated forward curve model using the Musiela parameterisation and argue that the forward prices are asymptotically cointegrated when the time to maturities of contracts are fixed. By a distinct model formulation, [Casassus et al. \(2013\)](#) consider the linkage of commodities via convenience yield and their relative scarcity to introduce the feedback effect on the price equilibrium that corresponds to an error-correction mechanism of cointegration in their model, which stays cointegrated under the risk-neutral measure.

In a discrete-time framework, [Ohana \(2010\)](#) directly introduces an error correction mechanism to the drift of a forward curve model under the physical measure. His model can be regarded as a discrete-time version of the HJM model, which relates the local dependence structure to the short-term volatilities and correlations of forward curves using the Generalised AutoRegressive Conditional Heteroskedasticity (GARCH) model and Gaussian copula. It also associates the global dependence structure with the level and slope of forward curves to describe the long-term relationship between commodities. Econometric analysis reveals two global dependence structures for the US natural gas and heating oil forward curves. The author investigates the lead-lag relationship between the level and slope factors, finding the leading role of heating oil for the slope factors and the symmetric relationship for the level factors.

[Benmenzer et al. \(2007\)](#) translate earlier work of [Ohana \(2010\)](#)⁶ into a continuous-time framework. The forward curve model is transformed from the physical measure to the risk-neutral measure via the market price of risk, including a matrix of cointegrating vectors specified by a continuous-time counterpart of the Vector Error Correction Model

⁶[Ohana \(2010\)](#) is based on the PhD thesis of [Ohana \(2006\)](#).

(VECM). Under the risk-neutral measure, the model reduces to the HJM type model, where the correlation parameters represent the short-term dependence structure between factors. Thus, the cointegrating relationship disappears under the risk-neutral measure as required by no-arbitrage assumptions. [Edoli et al. \(2013\)](#) focus on the estimation of cross-commodity correlations extending [Benmenzer et al. \(2007\)](#). Although their model estimates the marginal term structures of WTI, Brent and Gasoil volatilities precisely, their study does not further investigate the joint distributions.

The aforementioned multi-commodity spot and forward price models typically define the dependence of multiple commodity prices with respect to common long-term factors, mainly by cointegration. However, the majority of these models do not preserve the stable long-term relationship of forward prices under the risk-neutral measure except for some particular cases, such as when the forward contracts have a fixed time to maturity, as discussed by [Benth and Koekebakker \(2015\)](#). Therefore, an alternative approach would be to include the dependence structures of commodities directly into a driftless forward price model under the risk-neutral measure. In other words, the forward curve modelling reduces to the simultaneous modelling of volatility functions for multiple commodities. To our knowledge, research in this direction of price modelling is yet scarce.

Previously, [Tolmasky and Hindanov \(2002\)](#) proposed a multi-commodity forward curve model where regular patterns of correlation matrices are used to define the volatility functions of seasonal and non-seasonal commodity forward curves. The authors argue that their eigenstructures should also be similar when the within-curve correlation structures are similar between a pair of commodities. Then, a correlation matrix could be factorised as a constant multiplication of other correlation matrices to stylise the cross-commodity dependence structures. Under a correlated representation of the HJM model, an empirical study by [Feron and Gruet \(2020\)](#) also finds similarities in the behaviour and number of factors required for six European electricity markets, which encourage the joint modelling of energy forward curves by the proposed approach.

Close to the spirit of [Tolmasky and Hindanov \(2002\)](#), we develop a new framework to estimate the volatility functions of [Heath et al. \(1992\)](#) by analysing common eigenstructures of multiple commodity forward curves. Compared to [Tolmasky and Hindanov \(2002\)](#), we do not impose strong assumptions on the functional form of correlation/covariance matrices to define the factor volatilities. Instead, we let the data speak for themselves if such common structures exist in the dispersion of commodity forward curves. The theories and factor estimation method are introduced in Chapter 3.

Chapter 2

Notes on Commodity Futures

2.1 Commodity Futures

The price of commodity futures is a function of several time variables. These include the pricing date t , last trading date L , price settlement date S or period $[S_1, S_2]$, and the delivery day T or period $[T_1, T_2]$. When and how the futures are settled and delivered (or not delivered) broadly vary depending on the contract specifications defined by commodity exchanges. However, the general order would be $t \leq L \leq S \leq T$ where T only happens if the contract requires physical delivery, otherwise it is financially settled.

When we express the futures prices mathematically, $F(t, T)$ is probably the most common choice of functional form that we see in other literature. Unless otherwise stated by authors, it implies that the assumption of $S = T$ is in place, and often also $S_1 = S_2 = T_1 = T_2$. Therefore, other aspects of price determination processes such as the timing of price settlement and duration of delivery are often ignored to simplify the modelling framework, in which T is regarded as the *maturity* of the contract. In this thesis, we also denote the price of futures by $F(t, T)$ for distinct types of commodity products.

The top table in Table 2.1 introduces the product names and the trade venues for eight energy commodity futures that are used in the empirical analysis later in this thesis. Categories A to C classify the products by the last trading day, settlement method, delivery period, and contract type (physical or financial) as described at the bottom of the table.

In Chapters 3 – 5, we look at the European natural gas and electricity data that all belong to category A, in which $L = S < T_1 < T_2$ and the delivery takes place uniformly in the interval $[T_1, T_2]$ during the contract month. The exchanges offer trade venues for the month, quarter, season, and year (calendar) contracts throughout a year. Using the February 2020 contract of UK natural gas futures as an example, it is listed on the

exchange until 30th January 2020, and the floating price is settled on the same day. The delivery of the underlying asset is due every day between the 1st February and the 29th February 2020. One possible way of writing the futures would be $F(t, S, T_1, T_2)$. Instead, we set $S = T = [T_1 + T_2]/2$ and simplify the notation to be $F(t, T)$. In this way, we are treating the February 2020 contract as the price of natural gas futures that matures in the middle of February with immediate delivery. The same assumption is applied to the rest of category A products.

Chapter 6 deals with a hedging problem with financially settled fuel futures in the shipping market. In Table 2.1, category B products are the financial futures for ‘bunkers’ that are marine fuels used to run vessels, written on the daily price assessments of Platts.¹ For bunker fuels, Platts’s index reflects the market value of the fuels supplied one to eight days forward from the date of assessment.² By contrast, the category C product is the financial futures on the price of another futures contract: the first nearby physical contract of the NYMEX Light Sweet Crude Oil Futures (WTI).

Despite the name ‘futures’, both category B and C products are, in fact, calendar swaps. As a rule, the floating price of a calendar swap is determined as the arithmetic average of the underlying index for which the price settlement occurs during the contract month. If we choose to model the calendar swap as the arithmetic average, it would require separate calculations of the daily futures prices for the settlement period $[S_1, S_2]$ as is done in Lucia and Schwartz (2002). Instead, we approximate the nature of calendar swaps by setting $T = [S_1 + S_2]/2$. The replacement of $[S_1, S_2]$ by a single maturity T eases the computation of the price dynamics while adapting the delivery-time effect to the futures price (Benth and Koekebakker, 2015). For these reasons, we set the mid-point of the settlement period as the maturity of the category B products and denote the calendar swaps by $F(t, T)$.

The category C product requires further clarification to define the maturity date since the contract month of a calendar swap does not necessarily represent the value of the underlying commodity for the contract month. To see the point, let us compare the January 2020 contract of any European bunker fuel futures in category B and the WTI Financial Futures in category C. For the former contract, the floating price is determined as the arithmetic average of Platt’s daily index between 1st and 31st January 2020 in which the trade terminates the 31st January 2020. Hence, the January 2020 contract would

¹<https://www.spglobal.com/platts/en>

²It is one to eight days for the Rotterdam bunker fuels and three to seven days for the US Gulf Coast bunker fuel.

reflect the average price of the underlying spot asset between around the 2nd January and 8th February, including the time lag in the forward price assessment; see Footnote 2. By contrast, the January 2020 contract of category C futures represents the average price of the physical commodity in February 2020 and March 2020 weighted by trading days, as the underlying index is another futures contract; it alters from the first nearby to the second nearby futures during the settlement period.³ Swindle (2014) notes:

‘[For WTI swaps], the floating price is comprised of roughly a two-thirds weighting of one contract price and a one-third weighting of the following contract... While simple enough in concept, this causes a bit of headache in the design of risk systems in that a WTI swap for a single contract month will have nonzero delta exposure to two distinct futures prices.’

(Valuation and Risk Management in Energy Markets, p.18)

The compounded feature of calendar swaps, which is the average settled futures on another futures, makes the concept of the maturity date complicated. In our example, the underlying contracts of the January calendar swap are the February and March futures that have their own delivery schedule in the contract months, and the timing of delivery depends on the agreement between the buyer and seller of the contracts. Therefore, to make the problem simple enough, yet reasonable, we regard the middle of the next contract month as the maturity of the category C product, e.g., the maturity of the January calendar swap is mid-February.

³WIT physical futures (NYMEX product name: Crude Futures) typically expire 3 to 4 business days before the 25th calendar day of a month. In our example, it is the 21st January 2020. Therefore, the January contract of the WTI Financial Futures refers to the price of February physical futures between 1st–20th January and March physical futures between 21st–29th January.

Table 2.1: Contract specifications of commodity futures.

Category	Product Names	Exchange	Unit of Trading
A	UK Natural Gas Futures	ICE Futures Europe	GBp/1000 Therm
A	Dutch TTF Gas Futures	ICE Endex	EUR/MWh
A	UK Base Electricity Future (Gregorian)	ICE Futures Europe	GBP/MWh
A	Dutch Power Base Futures	ICE Endex	EUR/MWh
B	European FOB Rdam Marine Fuel 0.5% Barges (Platts) Futures	NYMEX	USD/1000 MT
B	European 3.5% Fuel Oil Barges FOB Rdam (Platts) Futures	NYMEX	USD/1000 MT
B	Gulf Coast HSFO (Platts) Futures	NYMEX	USD/1000 BBL
C	WTI Financial Futures	NYMEX	USD/1000 BBL

GBp...GB pence sterling, GBP...GB pound sterling, MWh...Megawatt hour, MT...Metric ton, BBL...Barrel

Category	Last Trading Day	Settlement ($S, [S_1, S_2]$)	Delivery	Type	Maturity Approximation for T
A	Two business days before the first calendar day of the delivery month, quarter, season, or calendar.	Takes place on the last trading day.	Every day during the contract period.	Physical	$T = (T_1 + T_2)/2$
B	The last business day of the contract month.	The arithmetic average of the daily price index during the contract month.	N.A	Financial	$T = (S_1 + S_2)/2$
C	The last business day of the contract month.	The arithmetic average of the first nearby physical futures price for each business day during the contract month.	N.A	Financial	$T \approx 1/12 + (S_1 + S_2)/2$ (Middle of next calendar month)

2.2 Interpolation of Data

The difficulty with commodity futures data arises from the fact that traded contracts expire. It is not possible to buy and hold a ‘shortest maturity futures’ without rolling the position each month. For instance, for the European energy futures that are introduced in the previous section, the first-to-expire futures in January refers to the February contract that requires the delivery of an underlying asset during the contract month. After the last trading day in January, the February contract is no longer available in the market, and the March contract becomes the first-to-expire contract.

The non-continuity of futures data introduces an extra complexity in the modelling of commodity price dynamics. As a remedy, many authors concatenate one price series on top of the other to produce continuous data. However, the data will be non-stationary, and the estimated parameters will be unstable as the volatility of each contract will rise towards the last trading day, as predicted by the Samuelson Effect. In addition, the statistical property of the data will be susceptible to the decision on which day to concatenate (roll) the data. The choice of roll day and its impact on the first two moments of price and return are discussed in detail in [Ma et al. \(1992\)](#).

One way of dealing with this issue in data analysis is to construct continuous time series data with a fixed time-to-maturity by the standard linear interpolation method.

$$f(t, t + \tau_i) = \frac{(T_{i+1} - (t + \tau_i))F(t, T_i) + ((t + \tau_i) - T_i)F(t, T_{i+1})}{T_{i+1} - T_i}, \quad T_i \leq t + \tau_i \leq T_{i+1} \quad (2.1)$$

where t is the current time, $\tau_i, i = 1, \dots, N$ are constant maturities typically corresponding to 30 days, 60 days etc., and $F(t, T_i)$ are the active futures contracts in the market with fixed maturity dates $T_i, i = 1, \dots, N$. In fact, $f(t, t + \tau_i)$ is the weighted average of two active futures contracts

$$f(t, t + \tau_i) = \alpha(t) F(t, T_i) + (1 - \alpha(t)) F(t, T_{i+1}) \quad (2.2)$$

where $\alpha(t) = \frac{T_{i+1} - (t + \tau_i)}{T_{i+1} - T_i}$. Thus, $f(t, t + \tau_i)$ can be seen as a portfolio of positions in two active futures if the contracts are infinitely divisible.

When we wish to ‘realise’ the returns of this portfolio, it requires the rebalancing of positions at a designated frequency. However, the returns calculated are not realisable since the weight $\alpha(t)$ changes at each observation. [Galai \(1979\)](#) proposes an alternative

data construction method for equity option indexes that offers continuity of data, as well as realisability of returns. The idea is to fix the weight $\alpha(t)$ between t and $t + 1$ to derive a new weight $\beta(t) = \frac{\alpha(t) \cdot F(t, T_i)}{F(t, T_{i+1})}$ to interpolate the returns of original data instead of prices

$$\tilde{r}(t) = \beta(t) r_i(t) + (1 - \beta(t)) r_{i+1}(t) , \quad (2.3)$$

where

$$\beta(t) = \frac{\alpha(t) \cdot F(t, T_i)}{F(t, T_{i+1})}$$

$$r_i(t) = \frac{F(t, T_i) - F(t - 1, T_i)}{F(t - 1, T_i)}.$$

One can retain the time series of fixed time-to-maturity contracts from the interpolated realisable returns utilising the general price-return relationship:

$$\tilde{f}(t, t + \tau_i) = \tilde{f}(t - 1, t - 1 + \tau_i)(1 + \tilde{r}(t)) \quad (2.4)$$

The question for us is which of Eq. (2.2) or Eq. (2.3) to use to interpolate our data. [Alexander et al. \(2013\)](#) argue that the [Galai \(1979\)](#)'s interpolation by return (or price differences where appropriate) is the only method that can recognise the value of money invested in the portfolio of assets. Hence it should be used for investment analysis, such as the calculation of hedge ratios.

While their arguments are relevant in constructing continuous data for the class of assets for which the maturity falls on a single trading day, we cannot directly apply the concepts to commodity futures data. As discussed in the previous section, it is inevitable for us to simplify the complex trading mechanism of commodity futures to define the 'price' and 'maturity'. Therefore, some degree of approximation has already been introduced in the futures price $F(t, T_i)$ and the precision that [Galai \(1979\)](#)'s method offers will be diminished as a consequence.

For that reason, we choose to use Eq. (2.2) in our analysis. Having said that, we remark that there is no single right or wrong method for interpolating commodity futures data. It all depends on the objective of the research and the property of the underlying processes that we want to incorporate in a modelling framework.

Chapter 3

Principal Component Analysis on Multi-group Data

3.1 Introduction

As discussed earlier, existing multi-commodity models are heavily dependent on mathematical assumptions accompanied by a large set of model parameters, most of which do not preserve the intended dependence structures of several commodities for their forward/futures prices under the risk-neutral measure. By taking an alternative route, we directly start from the risk-neutral measure and introduce the multi-commodity dependence structures to the factor volatilities of an established forward curve model. The main objective is to reduce the complexity of multi-commodity modelling caused by computationally expensive parameter estimation.

The factor estimation method to be presented in this chapter is a straightforward extension of PCA, which has been a widely adopted method in many fields in science since the pioneering work of [Pearson \(1901\)](#) and [Hotelling \(1933\)](#). PCA allows us to transform multi-dimensionally correlated data into a reduced set of orthogonal data, significantly decreasing the complexity of problems. Technically, it can be seen as an optimisation problem subject to constraints on the orthonormal eigenvectors, although one can attain identical solutions algebraically.

The main objective of Chapter [3](#) is to suggest a variation PCA to analyse more than one data set, and differentiate this approach from other conventional approaches to find the common eigenstructures in data. To start with, we briefly sketch the derivation of ordinary

PCA using the algebraic approach.¹ Subsequently, we introduce the base examples for the multi-group PCA commonly found in other finance literature. The empirical section applies our proposed method to the UK and Dutch energy commodities to investigate how the common eigenstructures relate to the latent factors that jointly drive their forward curves.

3.2 Principal Component Analysis (PCA)

Let $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) \in \mathbb{R}^{M \times N}$ be a matrix where $\mathbf{x}_i \in \mathbb{R}^{M \times 1}$ are column vectors of daily zero-mean (log) returns for a set of forward curve data with maturity $T_1, \dots, T_N$. M is a positive integer representing the number of observations in data. Define the sample covariance matrix of $\mathbf{X}$:

$$\boldsymbol{\Sigma} = M^{-1} \mathbf{X}' \mathbf{X} \quad (3.1)$$

PCA is based on the spectral decomposition of the symmetric matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$,

$$\boldsymbol{\Sigma} = \mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^{-1}, \quad (3.2)$$

where $\mathbf{V} \in \mathbb{R}^{N \times N}$ is the eigenvector matrix of $\boldsymbol{\Sigma}$ and $\boldsymbol{\Lambda} \in \mathbb{R}^{N \times N}$ is the diagonal matrix whose entries $\lambda_1, \lambda_2, \dots, \lambda_N$ are the eigenvalues of $\boldsymbol{\Sigma}$ corresponding to each column vector of unit length in $\mathbf{V}$. When $\boldsymbol{\Sigma}$ is a positive definite matrix and the eigenvalues are all distinct, $\mathbf{V}$ is an orthogonal matrix and $\mathbf{V}^{-1} = \mathbf{V}'$.

$$\boldsymbol{\Lambda} = \mathbf{V}' \boldsymbol{\Sigma} \mathbf{V} \quad (3.3)$$

In PCA, the eigenvalues (and corresponding eigenvectors) are ordered from the largest to smallest to define a new set of orthogonal data known as the principal components (PCs)

$$\mathbf{P} = \mathbf{X} \mathbf{V}, \quad (3.4)$$

where $\mathbf{P} = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N) \in \mathbb{R}^{M \times N}$ is a matrix consisting of $M \times 1$ orthogonal column vectors $\mathbf{p}_j$, and $\mathbf{V}$ is the reordered eigenvector matrix. Note that the variance of $\mathbf{P}$ is

$$\begin{aligned} M^{-1} \mathbf{P}' \mathbf{P} &= T^{-1} (\mathbf{X} \mathbf{V}')' \mathbf{X} \mathbf{V}' = \mathbf{V}' \boldsymbol{\Sigma} \mathbf{V} \\ &= \boldsymbol{\Lambda}. \end{aligned} \quad (3.5)$$

¹See Jolliffe (2011) for the general introduction of PCA, and Alexander (2008a) and Alexander (2008b) for various applications of PCA in finance.

The implication of Eq. (3.4) is that $\mathbf{V}$ linearly transforms the original returns data to orthogonal axes that maximise the total variation of the sample covariance matrix. Subsequently, one can obtain the principal component representation of original data reversing the relationship:

$$\mathbf{X} = \mathbf{P}\mathbf{V}' \quad (3.6)$$

PCA achieves dimension reduction when we use the first $n(< N)$ columns of $\mathbf{V}$ discarding the rest of the eigenvectors in Eq. (3.4). Similarly, we obtain the principal component approximation of the covariance matrix using the first n columns of $\mathbf{V}$ and $n \times n$ elements of $\mathbf{\Lambda}$ in Eq. (3.2). The cumulative sum of the first $n(\leq N)$ eigenvalues in (the reordered) $\mathbf{\Lambda}$ represents the explanatory power of the first n PCs, known as the total variation.

The convenience of PCA can be found in its orthogonal representation of original data, in which the key components represent approximations to the original data/covariance matrix while benefiting from dimension reduction. These advantageous features of PCA become notable when we analyse forward curves of multiple commodities, as we shall evidence in the following sections.

3.3 PCA on Multi-group Data

To extend PCA to more than one set of forward curve data (i.e. multi-group data), let us consider two $M \times N$ zero-mean (log) return matrices $\mathbf{X}_1$ and $\mathbf{X}_2$ that follow the same definitions as previously the single forward curve case. For the sake of simplicity, we set the number of forward curves to be two, assuming the same number of rows and columns for the matrices.

3.3.1 Conventional Approaches

Correlated PCA

The first approach performs separate PCA calculations for each data set and solely incorporates the dependence structures by correlation. Denote the cross-covariance matrix by $\mathbf{\Sigma}_{12}$. It gives the principal component representation of the cross-covariance matrix by the spectral decomposition Eq. (3.2):

$$\mathbf{\Sigma}_{12} = M^{-1}\mathbf{X}_1'\mathbf{X}_2 = \mathbf{V}_1\mathbf{\Lambda}_1^{1/2}\mathbf{R}_{12}\mathbf{\Lambda}_2^{1/2}\mathbf{V}_2' \quad (3.7)$$

In the equation above, $\mathbf{R}_{12} \in \mathbb{R}^{N \times N}$ denotes the cross-correlation matrix of PCs:²

$$\mathbf{R}_{12} = \begin{pmatrix} \rho_{11} & \rho_{12} & \cdots & \rho_{1N} \\ \rho_{21} & \rho_{22} & \cdots & \rho_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{N1} & \rho_{N2} & \cdots & \rho_{NN} \end{pmatrix} \quad (3.8)$$

For an approximation of the cross-covariance matrix, the first $n_k (< N)$ columns of $\mathbf{V}_k$, $n_k \times n_k$, elements of $\mathbf{\Lambda}_k$, and the corresponding elements of $\mathbf{R}_{12}$ are used for $k = 1, 2$ in Eq. (3.8). Note that the columns in $\mathbf{\Gamma}_k := \mathbf{V}_k \mathbf{\Lambda}_k^{1/2}$ are orthogonal within data sets; however, those columns are correlated by $\mathbf{R}_{12}$ between data sets. The extension to $k > 2$ data (matrices) is straightforward and involves the pair-wise estimation of the cross-correlation matrices $\mathbf{R}_{kl}$.

Joint PCA

The joint PCA follows the same procedure as the ordinary PCA described in Section 3.2 except that it applies the spectral decomposition to the sample covariance matrix of a pooled data set $\mathbf{Y} = (\mathbf{X}_1, \mathbf{X}_2) \in \mathbb{R}^{M \times 2N}$. In contrast to the correlation-based approach, the joint PCA technique seeks orthogonal axes that jointly maximise the total variation of the pooled covariance matrix. Hence, the eigenvalue, eigenvector, and the resulting principal components are shared by the pooled data set.

As is the case for ordinary PCA, the joint PCA approximates the original pooled data set (and the covariance matrix) when we retain $n (< 2N)$ columns of the common eigenvector matrix in the decomposition. In other words, one can approximate $\mathbf{X}_1$ and $\mathbf{X}_2$ from the joint PCA representation of $\mathbf{Y}$.

Common Principal Components

The correlated PCA and joint PCA are two contrasting approaches regarding the common eigenstructures of covariance matrices; the former method implies no commonality for the eigenstructures of multi-group data, whereas the latter method assumes perfect commonality in the eigenstructures of the pooled data set. However, the real-life data may only sometimes be ideally suited to either of these cases.

Flury (1988)'s hierarchical analysis, by contrast, classifies covariance matrices into five levels of similarities by their eigenstructures. The hierarchy of similarities ranges from

²This matrix is not symmetric, and the diagonal elements are not unity.

the equality (level 1) to arbitrary (level 5) with three additional levels of similarities: proportionality of covariance matrices (level 2), Common Principal Components (CPC, level 3), and partial Common Principal Components (pCPC, level 4). He argues that different types of PCA should be applied to multi-group data according to the similarity levels of the covariance matrices.

Level 1 is the equality where the eigenstructures are the same between covariance matrices. In level 2, the proportionality of covariance matrices assumes one or more covariance matrices are a constant multiple of other covariance matrices. In level 3, CPC assumes shared eigenvectors for all covariance matrices, however unique eigenvalues for each data group. In the subsequent pCPC at level 4, not all but some eigenvectors are assumed to be shared across covariance matrices. The last level is arbitrary, where covariance matrices do not share any eigenstructures. Therefore, PCA should be applied separately to covariance matrices.

In the finance literature, CPC and pCPC have been used in various applications. The examples include the calibration of model parameters for a lognormal forward rate model ([Alexander, 2002](#)), the modelling of implied volatility surface where options with different maturity groups share the common eigenstructures ([Fengler, 2006](#)), and the identification of common factors for the joint modelling of commodity prices ([Cortazar et al., 2008](#)).

Besides the contributions of [Flury \(1988\)](#)'s common principal component analysis, a group of researchers have identified the limitations, such as its heavy dependence on distributional assumptions, which could lead to inconsistent outcomes in the similarity analysis and non-convergence of solutions arising from the computational complexity in the parameter estimation procedure ([Houle et al., 2002](#); [Eslami et al., 2011](#)). Moreover, it may potentially misspecify the common factor structures in forward curve modelling, as the term structures of forward curves are typically characterised by the same shape factors, even between unrelated commodities.

3.3.2 PCA on PCA (PPCA)

As an alternative to previous approaches, we suggest a flexible estimation method to search for the common eigenstructures in several forward curves. The proposed method is a straightforward extension of the traditional PCA that we name *PCA on PCA* (PPCA). The concept of PPCA is finding the orthogonal axes which explain the common variation of PCs across data (instead of the common variation within data), motivated by the two-stage PCA originally introduced by [Alexander and Chibumba \(1996\)](#) in the context

of modelling a large cross-asset GARCH covariance matrix. By contrast, our research features in estimating unconditional volatilities for the HJM type forward curve models. We explain how to measure the explanatory power of PPCA and provide interpretations to the common principal components, neither of which are explored topics in the original literature.

PPCA on Covariance Matrix

Suppose that we have performed PCA on the sample covariance matrices of $\mathbf{X}_1$ and $\mathbf{X}_2$ as in Section 3.2, from which we obtain two $M \times N$ matrices of PCs $\mathbf{P}_1$ and $\mathbf{P}_2$. Suppose further that we select the first n_k ($< N$), $k = 1, 2$ significant PCs from each matrix.

To perform PPCA, define $\hat{\mathbf{P}}_1 \in \mathbb{R}^{M \times n_1}$ and $\hat{\mathbf{P}}_2 \in \mathbb{R}^{M \times n_2}$, concatenate the matrices in $\hat{\mathbf{P}} = (\hat{\mathbf{P}}_1, \hat{\mathbf{P}}_2) \in \mathbb{R}^{M \times n^*}$ where $n^* = n_1 + n_2$, and compute the covariance matrix of PCs $\hat{\Sigma} = M^{-1} \hat{\mathbf{P}}' \hat{\mathbf{P}} \in \mathbb{R}^{n^* \times n^*}$ to apply the spectral decomposition to the covariance matrix.

$$\hat{\Sigma} = \mathbf{U} \mathbf{\Omega} \mathbf{U}' \quad (3.9)$$

The *Principal components of Principal Components* (PPCs), $\mathbf{Q} \in \mathbb{R}^{M \times n^*}$, are defined as³

$$\mathbf{Q} = \hat{\mathbf{P}} \mathbf{U}, \quad (3.10)$$

where $\mathbf{U} \in \mathbb{R}^{n^* \times n^*}$ and $\mathbf{\Omega} \in \mathbb{R}^{n^* \times n^*}$ are the eigenvector and eigenvalue matrices of $\hat{\Sigma}$. Using the matrices, we obtain an expression for the PPC representation of PCs analogous to Eq. (3.6):

$$\hat{\mathbf{P}} = \mathbf{Q} \mathbf{U}' \quad (3.11)$$

PPCA achieves dimension reduction if we use the first m ($< n^*$) columns of $\mathbf{U}$ discarding the rest of the columns in Eq. (3.10).

Having calculated the PPCs, one can obtain the PPC representation of the original data by reversing the matrix operations; firstly, expressing the marginal PCs with respect to PPCs, and secondly, expressing the original data with respect to the PPC representation of PCs. To do so, we rely on the partition of the common eigenvector matrices as follows.

Denote the submatrices of $\mathbf{U}$ by $\mathbf{U}_1 \in \mathbb{R}^{n_1 \times n^*}$ and $\mathbf{U}_2 \in \mathbb{R}^{n_2 \times n^*}$ where the rows in $\mathbf{U}$ are partitioned by the number of PCs (n_k) used in the first stage PCA. Similarly, denote

³Assume that we have ordered the eigenvalues and corresponding eigenvectors from the largest to smallest in Eq. (3.9).

the submatrices of $\hat{\mathbf{P}}$ by $\hat{\mathbf{P}}_1 \in \mathbb{R}^{n^* \times n_1}$ and $\hat{\mathbf{P}}_2 \in \mathbb{R}^{n^* \times n_2}$. Then,

$$\begin{aligned}\mathbf{Q} &= \hat{\mathbf{P}}\mathbf{U} \\ &= \hat{\mathbf{P}}_1\mathbf{U}_1 + \hat{\mathbf{P}}_2\mathbf{U}_2\end{aligned}\tag{3.12}$$

and

$$\hat{\mathbf{P}}_k = \mathbf{Q}\mathbf{U}'_k.\tag{3.13}$$

Eq. (3.13) is the PPC representation of PCs. Recall that the principal component representation approximates the original data when $n_k < N$:

$$\mathbf{X}_k \approx \hat{\mathbf{P}}_k \hat{\mathbf{V}}'_k$$

Thus, we can find the PPC representation of the original data by combining these results all together for $k = 1, 2$

$$\mathbf{X}_k \approx \mathbf{Q}\mathbf{U}'_k \hat{\mathbf{V}}'_k,\tag{3.14}$$

as well as the PPC approximation of the covariance matrix:

$$\begin{aligned}\Sigma_k &\approx M^{-1}(\mathbf{Q}\mathbf{U}'_k \hat{\mathbf{V}}'_k)'(\mathbf{Q}\mathbf{U}_k \hat{\mathbf{V}}'_k) \\ &= M^{-1}(\hat{\mathbf{V}}'_k \mathbf{U}_k \mathbf{Q}' \mathbf{Q} \mathbf{U}_k \hat{\mathbf{V}}'_k) \\ &= \hat{\mathbf{V}}'_k \mathbf{U}_k \boldsymbol{\Omega} \mathbf{U}_k \hat{\mathbf{V}}'_k\end{aligned}\tag{3.15}$$

Note that $\boldsymbol{\Omega} \in \mathbb{R}^{n^* \times n^*}$ is indeed the variance of PPCs:

$$\begin{aligned}M^{-1}(\mathbf{Q}'\mathbf{Q}) &= M^{-1}(\hat{\mathbf{P}}\mathbf{U})'\hat{\mathbf{P}}\mathbf{U} \\ &= \mathbf{U}'(M^{-1}\hat{\mathbf{P}}'\hat{\mathbf{P}})\mathbf{U} \\ &= \mathbf{U}'\hat{\Sigma}\mathbf{U} = \boldsymbol{\Omega}\end{aligned}\tag{3.16}$$

The fewer $m (< n^*)$ we use in the second stage PCA, the greater the degree of approximation in Eqs. (3.14) and (3.15).

A number of remarks can be made from the results above. First, PPCs are defined as the linear combination of marginal PCs in the second line of Eq. (3.12), where $\mathbf{U}_k$, $k = 1, 2$, act as the loading matrix for the marginal PCs. Second, it is evident in Eq. (3.14) that marginal data set includes the same PPCs on the right-hand side of the equation. Likewise, the marginal covariance matrices share the same eigenvalue matrix in their PPC approximations as shown in Eq. (3.15). These results are particularly important for us to

justify why PPCs are regarded as common factors in our modelling approach.

PPCA on Correlation Matrix

Thus far, we have suggested the PPC matrix $\mathbf{Q}$ to be the common component in our approach. Nevertheless, there is a general tendency that the outputs of PCA are influenced by variables with large volatilities when the decomposition is applied to a covariance matrix (Alexander, 2008a). If this happens, the most important components we find through the second stage of PCA may not necessarily represent the common eigenstructures of multi-group data. A solution to this problem is to perform the second stage PCA on the cross-correlation matrix of PCs. In other words, we perform PPCA on the cross-covariance matrix of *standardised* PCs by dividing the original PCs by their standard deviations.

Denote the standardised $\hat{\mathbf{P}}$ by $\bar{\mathbf{P}} = (\bar{\mathbf{P}}_1, \bar{\mathbf{P}}_2) \in \mathbb{R}^{M \times n^*}$, where

$$\hat{\mathbf{P}}_k = \bar{\mathbf{P}}_k \hat{\mathbf{\Lambda}}_k^{1/2} \quad (3.17)$$

for $k = 1, 2$. By spectral decomposition, the covariance matrix of $\bar{\mathbf{P}}$ decomposes into

$$\bar{\mathbf{\Sigma}} = \bar{\mathbf{U}} \bar{\mathbf{\Omega}} \bar{\mathbf{U}}', \quad (3.18)$$

where $\bar{\mathbf{U}} \in \mathbb{R}^{n^* \times n^*}$ and $\bar{\mathbf{\Omega}} \in \mathbb{R}^{n^* \times n^*}$ are the eigenvectors and eigenvalues of $\bar{\mathbf{\Sigma}} \in \mathbb{R}^{n^* \times n^*}$. The derivation of (standardised) PPCs follows Eqs. (3.10) – (3.13) in basically the same way, except that $\bar{\mathbf{P}}$ replaces $\hat{\mathbf{P}}$. However, the differences emerge in Eq. (3.14) and Eq. (3.15) to express the original data and covariance matrices, as we need to *unstandardise* the PCs by $\hat{\mathbf{\Lambda}}_k^{1/2}$ based on the relationship in Eq. (3.17):

$$\mathbf{X}_k \approx \bar{\mathbf{Q}} \bar{\mathbf{U}}_k' \hat{\mathbf{\Lambda}}_k^{1/2} \hat{\mathbf{V}}_k' \quad (3.19)$$

$$\mathbf{\Sigma}_k \approx \hat{\mathbf{V}}_k \hat{\mathbf{\Lambda}}_k^{1/2} \bar{\mathbf{U}}_k \bar{\mathbf{\Omega}} \bar{\mathbf{U}}_k' \hat{\mathbf{\Lambda}}_k^{1/2} \hat{\mathbf{V}}_k' \quad (3.20)$$

The following equation links the decomposition of marginal covariance matrices by (i) PCA, (ii) PPCA on the cross-covariance matrix of PCs, and (iii) PPCA on the cross-correlation matrix of PCs. Note that $\mathbf{\Omega}$ and $\bar{\mathbf{\Omega}}$ are exactly the same for all marginal covariance matrices.

$$\mathbf{\Sigma}_k = \mathbf{V}_k \mathbf{\Lambda}_k \mathbf{V}_k' = \mathbf{V}_k \mathbf{U}_k \mathbf{\Omega} \mathbf{U}_k' \mathbf{V}_k' = \mathbf{V}_k \mathbf{\Lambda}_k^{1/2} \bar{\mathbf{U}}_k \bar{\mathbf{\Omega}} \bar{\mathbf{U}}_k' \mathbf{\Lambda}_k^{1/2} \mathbf{V}_k' \quad (3.21)$$

The equalities become approximations when we do not retain all components in the two-

stage PCA, and the degree of approximation increases as we move from the left to the right in the equation. On the other hand, the model better captures the commonality in eigenstructures towards the right of the equation.

Explanatory Powers of PPCA

The total variation typically measures the explanatory power of PCA. Many authors use the total variation as a rule-of-thumb method in financial modelling to determine the number of PCs required to analyse multivariate data. The interpretation is that if, for example, it is 90% for the n most significant PCs, they explain 90% of dispersion in the original data structure. At the same time, the rest of 10% can be regarded as the reconstruction error arising from omitting insignificant PCs, often left intentionally as ‘residuals’ or ‘noise’ components.

Mathematically, the ratio of cumulative eigenvalues of the first n ($\leq N$) most significant eigenvectors to the total sum of eigenvalues gives the total variation explained by PCA.

$$\text{Total Variation Explained (\%)} = 100 \times \frac{\sum_{i=1}^n \lambda_i}{\sum_{i=1}^N \lambda_i}, \quad (3.22)$$

where λ_i is the i -th largest eigenvalue. Or equivalently,

$$\text{Total Variation Explained (\%)} = 100 \times \frac{\text{tr}(\tilde{\Sigma})}{\text{tr}(\Sigma)}, \quad (3.23)$$

where $\tilde{\Sigma}$ denotes the PCA estimate of a covariance matrix. Eqs. (3.22) and (3.23) are invariant under the ordinary principal component transformation. Nonetheless, Eq. (3.22) cannot directly be applied to PPCA in which two types of eigenvalues $\Lambda_{\mathbf{k}}$ and Ω (or $\bar{\Omega}$) are compounded in the two-step matrix operations; see Eq. (3.21). For that reason, we use Eq. (3.23) to measure the explanatory powers of PPCA in the subsequent analysis, substituting $\tilde{\Sigma}$ by the PPCA estimate of a covariance matrix.

3.4 Empirical Study

3.4.1 Data

We apply PPCA to analyse the common factors in the forward curve dynamics of European energy commodities. The data set consists of ICE futures for UK gas and power, as well as Dutch gas and power introduced in Section 2.1. The data window covers two years, from 1st July 2016 to 29th June 2018, obtained through Bloomberg®.




Figure 3.1: Time series of the first-month and fourth-quarter prices. The data cover daily settlement prices between 1st July 2016 to 29th June 2018. The solid and dotted lines represent the front-month and fourth-quarter futures, respectively. All prices are expressed in EUR/MWh.

Fig. 3.1 depicts the time series of front-month contracts (solid lines) and fourth-quarter futures (dotted lines) for the UK and Dutch gas power, where the fourth-quarter futures are on average one year to maturity. The trading units of UK energy futures are converted into EUR/MWh to make a comparison with Dutch energy futures for illustrative purposes. Fig. 3.2 shows the log returns. In Fig. 3.1, the energy prices exhibit common movements, having peaks and troughs around the same periods. Moreover, the fluctuations of returns tend to synchronise, as observed in Fig. 3.2. The front-month futures are more volatile than the fourth-quarter contracts, as predicted by the Samuelson Effect. Note that the spikes in returns, most notably happening with the quarter futures in April and October, are caused by the rolling of the underlying contracts.⁴

The comovements of UK and Dutch gas futures are more noticeable than power futures, probably due to the pipeline connections of the physical markets. In addition, the price of

⁴Section 2.2 discusses the rationale for data interpolation to avoid this issue.


gas is susceptible to weather conditions, as seen in December 2017 when the front-month UK gas futures spiked with the arrival of news of the freezing weather known as ‘the Beast from the East’. During the observation period, the decline of supply inflow from the European gas markets caused the price rise in the UK power futures in late 2016; generally, price shocks in gas markets affect the price of power futures since natural gas is one of the primary fuels to generate electricity.

As for the trade mechanism of futures, the exchanges list month, quarter, season, and calendar (year) contracts simultaneously throughout the year for the UK and Dutch gas futures. The quarter contracts start from January, April, July, October, covering a three-month delivery period. Likewise, the season contracts start from April and October for a six-month delivery period. The year contracts start from January every year for the delivery of the entire calendar year. The UK power futures are listed similarly, except for the absence of year contracts. For the Dutch power futures, it is the season contracts that are not traded in the market.⁵ Table 3.1 illustrates the quote patterns of the UK gas, Dutch gas, and UK power futures for an eighteen-month window. Table 3.2 separately shows the quote patterns of Dutch power futures.


The exchanges set the maturities of these contracts five to six years ahead of today. However, the liquidity ceases at the short-end: typically after several months for month contracts, a few quarters for quarter contracts, a few seasons for season contracts, and a few years for calendar contracts. For example, using the monthly contracts alone would be insufficient to construct a forward curve with long enough maturity horizons. Instead, using calendar contracts only will lead to a forward curve with stepwise increments every twelve months, creating an artificially stale data series that is not very useful in time series analysis.

Due to the short liquidity horizons of listed contracts, it is plausible for us to use more than a single frequency of contract where necessary to construct a time series of forward curve data with the desired maturity horizon. Nevertheless, complexities arise when contract delivery periods overlap while the quote patterns shift every month. For instance, in March, when the nearby month contracts are listed for the first six months, the first and second quarter contracts and the first season contracts are also listed on the exchange. Similarly, the trading periods of the third and fourth quarters, the second season, and a part of first year contracts overlap during the month.

⁵The information is based on as of July 2018. The quote structures of European energy futures at ICE have changed since then. As of June 2022, quarter, season, year, and any period of month contracts can be traded for Dutch gas and power futures. For the details, see <https://www.theice.com/energy>.



(a) First month



(b) Fourth quarter

Figure 3.2: The first-month and fourth-quarter log returns.

The time series of returns are calculated as the log differences of daily settlement prices for the period of 1st July 2016 to 29th June 2018 using the same price data as Fig. 3.1.

Table 3.1: Quote patterns of UK gas, Dutch gas, and UK power futures at ICE.

Jan	Month	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year			S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3	S3	S3
Feb	Month	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6
	Year		S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3	S3	S3	S3
Mar	Month	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season	Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6	Q6
	Year	S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3	S3	S3	S3	S3
Apr	Month	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year					S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S2	S3
May	Month	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6
	Year				S1	S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3
Jun	Month	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season	Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6	Q6
	Year				S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3	S3
Jul	Month	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year			S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3	S3	S3
Aug	Month	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6
	Year		S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3	S3	S3	S3
Sep	Month	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season	Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6	Q6
	Year	S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3	S3	S3	S3	S3
Oct	Month	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year			S1	S1	S1	S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3
Nov	Month	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6
	Year			S1	S1	S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3
Dec	Month	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season				Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y2	Y2	Y2	Y2	Y2	Y2

M...Month, Q...Quarter, S...Season, Y...Calendar. The row labels show traded contracts by calendar month. The column labels represent the contract months of futures for an eighteen-month window. The shaded areas represent a non-overlapping forward curves. Dark grey cells are where mapping is applied. Note that year contracts are not available for UK power futures.

Table 3.2: Quote patterns of Dutch power futures at ICE.

Jan	Month	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year												Y1	Y1	Y1	Y1	Y1	Y1	Y1
Feb	Month	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6
	Year											Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1
Mar	Month	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6	Q6
	Year										Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1
Apr	Month	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year									Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1
May	Month	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6
	Year							Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1
Jun	Month	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6	Q6
	Year						Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1
Jul	Month	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year			S1	S1	S1	S1	S1	S1	S2	S2	S2	S2	S2	S2	S3	S3	S3	S3
Aug	Month	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6
	Year				Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y2	Y2
Sep	Month	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6	Q6
	Year				Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y2	Y2	Y2
Oct	Month	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season			Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6
	Year			Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y2	Y2	Y2	Y2
Nov	Month	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6
	Year		Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y2	Y2	Y2	Y2	Y2
Dec	Month	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun
	Quarter	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13	M14	M15	M16	M17	M18
	Season		Q1	Q1	Q2	Q2	Q2	Q3	Q3	Q3	Q4	Q4	Q4	Q5	Q5	Q5	Q6	Q6	Q6
	Year	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y1	Y2	Y2	Y2	Y2	Y2	Y2

M...Month, Q...Quarter, Y...Calendar. The row labels represent contract types by calendar month. The column labels show the contract months for an eighteen-month window. The shaded cells follow the same definitions as in Table 3.1.

To resolve these issues, we construct non-overlapping forward curves month by month as illustrated by the shaded cells in Tables 3.1 and 3.2, based on the following principles. For instance, if we are in February, the delivery period of the second quarter contract (Q2) overlaps that of the fifth to seventh month (M5-M7) contracts starting in July. Then the following relationship should hold by no-arbitrage

$$\frac{M5 + M6 + M7}{3} = Q2 , \quad (3.24)$$

and we can rearrange the formula to obtain a theoretical price of the M7 contract (M7*), where the liquidity is thin for the actual contract:

$$M7^* = 3 \times Q2 - M6 - M5 \quad (3.25)$$

For other contracts, we utilise similar relationships to obtain

$$MQ^* = (3 \times Q2 - M6)/2 \quad (3.26)$$

$$Q4^* = (6 \times S2 - 3 \times Q3)/3 , \quad (3.27)$$

where MQ* and Q4* represent two-month equivalent Q2 and three-month equivalent S2 respectively, replacing the actual Q4 contract in the latter case.

We use listed contracts including the first to sixth months (M1-M6), second and third quarters (Q2-Q3), and second to fifth seasons (S2-S5) for the UK gas, Dutch gas, and UK power futures to map the data.⁶ Instead, we select the first to sixth months (M1-M6), second to sixth quarters (Q2-Q6), and first and second years (Y1-Y2) for the Dutch power futures to adapt different quote patterns.⁷

Once we obtain the non-overlapping forward curves, we linearly interpolate the data to construct continuous price data with fixed time to maturities by setting $\tau = (30, 60, 90, 120, 150, 180, 210, 240, 360, 540)$ days from the average time to maturity of the mapped contracts. Fig. 3.3 shows the interpolated forward curves in their original units of trading. We take the log differences of interpolated price series to calculate the returns and covariances for PPCA.

⁶Some liquid contracts such as Q1 and S1 are regarded as redundant despite their high volume of trades since these contracts are quoted parallel to more liquid and higher frequency contracts. For instance, in March, M1-M6 are prioritised over Q1, Q2, and S1.

⁷Although some contracts do not appear in the figures, they are necessary to calculate the theoretical futures over eighteen-month maturity contracts.




Figure 3.3: Interpolated forward curves with fixed time to maturities. Prices are shown in the original units with fixed time to maturities: 30, 60, 90, 120, 150, 180, 210, 240, 360, and 540-days for the period of 1st July 2016 to 29th June 2018.

Fig. 3.4 depicts the volatilities of log returns by time to maturity in days. The volatilities range between 0.1 and 0.4 and exhibit the Samuelson Effect. As explained in Section 1.1, the Samuelson Effect refers to the volatilities of nearby contracts being always higher than those of distanced maturity contracts. This is because a market shock will immediately cause a supply/demand imbalance and increase the trade volumes of short-term futures more significantly than for long-term futures. However, we observe an exception for Dutch power, for which the decay in volatilities is non-monotonic, indicating the market's distinctive supply and demand patterns from the rest of the energy markets.

3.4.2 Results

Marginal PCA

The primary aim of the first-stage PCA is to reduce the dimensionality of forward curve data. Our preliminary analysis indicates the minor impact of seasonality in the outputs of PCA. Therefore, we perform PCA on the covariance matrix of the interpolated forward




Figure 3.4: Term structure of volatilities.

The annualised standard deviation of log returns calculated from fixed time-to-maturity contracts.

curve data without deseasonalisation.⁸

Fig. 3.5 depicts the first three eigenvectors of the four energy forward curves. These eigenvectors represent the stylised shape factors of forward curves: the ‘level’, ‘slope’, and ‘curvature’, while the concavity in the second eigenvectors may introduce some degree of bending effects in the slopes.

The first eigenvectors are visually very similar between the forward curves, implying that these forward curves react very similarly to a shock in the European energy markets. Meanwhile, the loadings of UK power futures indicate their higher responsiveness to short-term shocks than others, especially at the shortest end of the term structure. As for the second eigenvectors, the slopes are relatively steeper for Dutch gas and power than UK gas and power. These results suggest a wider spread between the short-term and long-term contracts in a situation of contango/backwardation for the Dutch energy futures. The similarities of eigenvectors seem to diminish when we consider the first, second, third, $\dots$, N -th PCs in descending order.

Table 3.3 reports the cumulative eigenvalues. The first three PCs explain over 97%, 98%, 92%, and 90% of the total variation in the forward curves of UK gas, Dutch gas, UK power, and Dutch gas, respectively. Generally speaking, the higher the correlation of contracts, the lower the number of PCs required, and vice versa. The explanatory powers of PCs are typically lower for power than gas data given the same number of PCs due to the non-storability of the commodity that lowers the correlations between contracts with

⁸We conjecture that the relatively short data horizon and smoothing effect of interpolation contribute to the non-significant seasonality effects on our data. With regard to the latter point, Swindle (2014) reaches a similar conclusion in the application of PCA to US natural gas forward curve data.




Figure 3.5: The first three eigenvectors of fixed time-to-maturity forward curves by PCA.

The dotted line with squares, crosses, and triangles depict the eigenvectors associated with PC1, PC2, and PC3, respectively.

different maturities.

PPCA

The PPCA tries to find the common PC factors that contribute to the joint dynamics of closely related commodity forward curves. To define what are meant by common factors, let us recall the relationship between PPCs and marginal PCs in Eq. (3.12):

$$\begin{aligned}\mathbf{Q} &= \hat{\mathbf{P}}\mathbf{U} \\ &= \hat{\mathbf{P}}_1\mathbf{U}_1 + \hat{\mathbf{P}}_2\mathbf{U}_2\end{aligned}$$

For instance, if the elements in the first column of $\mathbf{U}_2$ were all zero, the first PPC (the first column in $\mathbf{Q}$) would solely be determined by the PCs of the first forward curve. Then, the first PPC would not be a common component to both forward curves. Instead, we would like to find an ideal situation where a subset of marginal PCs from all forward curves exhibits an approximately equal size of entries to a column vector of $\mathbf{U}$ (or corresponding columns in $\mathbf{U}_1$ and $\mathbf{U}_2$) to treat the PPC as a common factor.

In our analysis, we use three marginal PCs per forward curve (i.e. twelve PCs in total)

Table 3.3: Explanatory powers of PCA.

	n_k	1	2	3	4	5	6	7	8	9	10
Ug	Explained %	83.64	11.33	2.34	1.15	0.55	0.42	0.24	0.16	0.12	0.05
	Cumulative %	83.64	94.97	97.32	98.47	99.01	99.43	99.67	99.83	99.95	100.00
Dg	Explained %	88.93	7.41	1.71	1.07	0.33	0.23	0.15	0.08	0.05	0.02
	Cumulative %	88.93	96.34	98.06	99.13	99.47	99.70	99.85	99.93	99.98	100.00
Up	Explained %	75.63	11.77	4.97	2.79	1.83	1.30	0.85	0.49	0.23	0.14
	Cumulative %	75.63	87.40	92.36	95.15	96.99	98.29	99.14	99.64	99.86	100.00
Dp	Explained %	75.36	11.47	3.76	2.86	2.21	1.70	0.99	0.91	0.45	0.27
	Cumulative %	75.36	86.84	90.60	93.46	95.66	97.37	98.36	99.27	99.73	100.00

Ug...UK gas, Dg...Dutch gas, Up...UK power, Dp...Dutch power. n_k represents the number of PCs. The marginal and cumulative explanatory powers are shown in % for each forward curve.

to uniformly cover over 90% of total variations across data. Fig. 3.6 reveals the outputs of PPCA on the cross-correlation matrix of PCs, in which the top plot shows the eigenvalues along with the total variation explained (%) on the x-axis in descending order. In the same figure, the bottom plots show the corresponding eigenvectors that are the column vectors of $\bar{\mathbf{U}}$, generally called loading vectors. The black markers are the entries (loadings) of the j -th loading vector, where every set of three consecutive ticks indicates the first three marginal PCs of the UK gas, Dutch gas, UK power, and Dutch power, respectively.

From the magnitude of eigenvalues, we conjecture that the first three PPCs are the key factors responsible for the joint evolution of commodity forward curves. After the fourth PPCs, the eigenvalues decline sharply, and the eigenvectors tend to represent only one or two of the four forward curves. When we retain the first three PPCs, they cover 67% of the total variation in the cross-correlation structure, and the explanatory powers for original (marginal) data accumulate to 86.61% (UK gas), 87.97% (Dutch gas), 70.38% (UK power), 51.45% (Dutch power).

In contrast to Fig. 3.6, the eigenvalues in Fig. 3.7 highlight only one key component when we apply PPCA to the cross-covariance matrix of PCs, covering approximately 71% of the total variation in the cross-covariance structure. There is a tendency that the eigenvectors show the dominance of power futures since the volatilities of unstandardised PCs introduce idiosyncrasies of marginal forward curves to the common eigenspace. When one PPC is retained, the total variations of the marginal forward curves are: 77.82% (UK gas), 81.64% (Dutch gas), 60.66% (UK power), and 41.94% (Dutch power).

Table 3.4 summarises the cumulative % of the total variations explained by PPCA on the cross-correlation matrix (top) and cross-covariance matrix (bottom) of PCs for varying




Figure 3.6: Eigenvalues and eigenvectors by PPC on the cross-correlation matrix of PCs.

Ug...UK gas, Dg...Dutch gas, Up...UK power, Dp...Dutch power. PPCA is performed on the cross-correlation matrix of the first three PCs from each forward curve. The top figure shows the eigenvalues and total variation explained (%) on the second x-axis. The bottom subfigures show the corresponding eigenvectors, where every set of three consecutive black squares marks the loadings of marginal PCs for a forward curve.

numbers of PPCs shown in the first row.⁹ As expected, the explanatory powers of the first few PPCs are notably lower for power futures in both PPCA. This is because idiosyncratic components account for unique dynamics in the power forward curves.

We make a couple of remarks for the results. First, when we retain all PPCs ($m = 12$) from the second stage PCA, the explanatory powers are identical between PCA, PPCA on the cross-covariance matrix and PPCA on the cross-correlation matrix (of PCs). Second, the explanatory power of PPCA on the cross-correlation matrix is lower than that of PPCA on the cross-covariance matrix since PPCA on the cross-correlation matrix intentionally discards the idiosyncratic volatilities of PCs and the degree of approximation increases. However, in exchange, it improves the common principal component analysis without the

⁹Table 3.4 reports $n_k = 3$ PCs for $k = 1, \dots, 4$ due to limitations in space. However, Eq. (3.23) can be used to calculate the total variations by any combination of n_k PCs.




Figure 3.7: Eigenvalues and eigenvectors by PPC on the cross-covariance matrix of PCs.

Ug...UK gas, Dg...Dutch gas, Up...UK power, Dp...Dutch power. PPCA is performed on the cross-covariance matrix of the first three PCs from each forward curve. The top figure shows the eigenvalues and total variation explained (%) on the second x-axis. The bottom subfigures show the corresponding eigenvectors, where every set of three consecutive black squares marks the loadings of marginal PCs for a forward curve.

noise brought by idiosyncratic volatilities. When we compare PCA and PPCA in general, PCA explains the total variation of *marginal* data better than PPCA, while PPCA explains the *common* variation of multiple data better than PCA by the designated model structure in Eq. (3.21).

3.4.3 Interpretation

We perform PPCA on the cross-correlation matrix of PCs in order to extract the most critical common movements across data, from which we would like to obtain some meaningful interpretations for the common latent factors. Since the j -th PPC is defined as a linear combination of marginal PCs, it is crucial to study the elements in the j -th loading vector of $\bar{\mathbf{U}}$ to infer the implications of the common latent factors. However, the challenge arises when the loading vectors consist of many non-zero elements that are not insignificant. We

Table 3.4: Cumulative explanatory powers of PPCA.

Corr													
	m	1	2	3	4	5	6	7	8	9	10	11	12
Ug (%)		75.29	83.73	86.61	87.04	88.22	88.74	90.44	93.00	93.80	94.13	95.02	97.32
Dg (%)		79.43	85.16	87.97	88.10	88.64	88.80	89.72	90.37	92.68	95.96	96.36	98.06
Up (%)		60.69	68.50	70.38	72.42	74.12	75.45	81.12	85.08	87.21	91.69	92.20	92.36
Dp (%)		45.63	49.89	51.45	59.24	63.63	68.02	71.79	90.20	90.51	90.52	90.55	90.60
Cov													
	m	1	2	3	4	5	6	7	8	9	10	11	12
Ug (%)		77.82	83.62	88.25	91.78	92.65	94.11	95.40	96.30	96.40	96.90	97.17	97.32
Dg (%)		81.64	83.68	87.40	92.10	92.21	92.92	95.80	96.46	96.54	96.96	97.73	98.06
Up (%)		60.66	62.80	74.66	83.60	83.83	87.24	88.68	91.24	91.92	92.35	92.36	92.36
Dp (%)		41.94	68.33	78.63	79.30	86.46	87.29	87.49	87.92	89.95	90.57	90.60	90.60

Ug...UK gas, Dg...Dutch gas, Up...UK power, Dp...Dutch power. The cumulative total variations of PPCA on the cross-correlation matrix of PCs (top) and PPCA on the cross-covariance matrix of PCs (bottom) for varying numbers of PPCs (m), representing the explanatory powers of PPCs with respect to the dispersion of price returns.

handle this problem by two approaches: a categorical ranking and a shrinkage method.

In the first approach, the PPC loadings are classified into seven categories, $G = (\text{level, slope, curvature, gas, power, UK, Netherlands})$ indexed by $k = 1, \dots, 7$ for the shape factors, types of commodities, and location. Define the indicator matrix

$$\mathbb{1} = \begin{matrix} & \begin{matrix} G_1 \\ G_2 \\ G_3 \\ G_4 \\ G_5 \\ G_6 \\ G_7 \end{matrix} \end{matrix} \begin{pmatrix} 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 \end{pmatrix}, \quad (3.28)$$

for which the rows are labelled G_k for clarification. The (k, i) -th element equals 1 if $\bar{u}_{ij}$ (for any j) belongs to category G_k , and 0 otherwise. When $\bar{\mathbf{U}}$ is pre-multiplied by $\mathbb{1}$, it leaves the only relevant elements in $\bar{\mathbf{U}}$ with respect to G_k

$$\mathbf{C} = \mathbb{1} \bar{\mathbf{U}}, \quad (3.29)$$

which is a 7×12 matrix.¹⁰

¹⁰Recall that the dimension of $\bar{\mathbf{U}}$ is 12×12 in our empirical study.

Define $\bar{C}_{kj}$ as the categorical contribution to the j -th PPC

$$\bar{C}_{kj} = \frac{\sum_{j=1}^{12} |c_{kj}|}{n_{kj} \sum_{i=1}^{12} |\bar{u}_{ij}|}, \quad (3.30)$$

where c_{kj} are the elements in the k -th row of $\mathbf{C}$ and n_{kj} is the number of non-zero elements in the same row. In the equation above, the sum of absolute values of c_{kj} are divided by n_{kj} to calculate the average category-wise contribution relative to the j -th PPC expressed by $\sum_{i=1}^{12} |\bar{u}_{ij}|$. By ranking these categorical contributions in descending order, we nominate the most reasonable category to define the common factors.

Table 3.5 reveals the results. The shape factors are ranked at the top for all the PPCs, suggesting that the commodity forward curves share the common shape factors regardless of the type of commodities and location. Interestingly, the level and slope factors are ranked at the top for PPC1 and PPC2, respectively, while the remaining shape factors are ranked at the bottom. For more detailed interpretations, one may consider subcategories. For example, PPC2 could be a combination of the slope and gas factors from the top two factors in the ranking. In the case of PPC3 and PPC4, they may be curvature factors for the Netherlands, albeit for different commodities (gas or power).

It is also possible to combine categories when the top $\bar{C}_{jk}$ are non-declining in value. For example, for PPC5, the slope, curvature, and power factors are ranked as top three with loadings 0.1150, 0.1118, and 0.1037, respectively. In this case, PPC5 could be a common factor representing the slope and curvature of power forward curves. These interpretations are convincing when we refer back to Fig. 3.6 for the graphical representations of $\bar{u}_{ij}$ (black markers) in the eigenvector plots.

The second approach relies on the Sparse PCA (SPCA), which is a family of shrinkage methods in linear regression analysis (e.g. ridge and LASSO) introduced by Zou et al. (2006). SPCA writes PCA as a regression problem and shrinks the loadings of PCs towards zero if they are thought to be negligible. Hence, it improves the interpretability of PCA.

$$\begin{aligned} \arg \min_{\mathbf{A}, \mathbf{B}} \quad & \sum_{i=1}^M \|\mathbf{x}_i - \mathbf{A}\mathbf{B}'\mathbf{x}_i\|_2^2 + \sum_{j=1}^n \lambda_{1j} \|v_j\|_1 + \lambda_2 \sum_{j=1}^n \|v_j\|_2^2 \\ \text{s.t.} \quad & \mathbf{A}'\mathbf{A} = \mathbf{I}_n \end{aligned} \quad (3.31)$$

In the optimisation problem, the first term represents the reconstruction error between the original data $\mathbf{x}_i$ and the PC representation of data, where $\mathbf{x}_i$ denotes the i^{th} row in $M \times N$ data set. Both $\mathbf{A}$ and $\mathbf{B}$ are $M \times n$ matrices where n is the number of PCs. If

Table 3.5: Ranking of category-wise contributions to PPCs.

	PPC1		PPC2		PPC3		PPC4	
1	Level	0.2119	Slope	0.1730	Curvature	0.1728	Curvature	0.1619
2	UK	0.0846	Gas	0.0932	Gas	0.1048	Power	0.1382
3	Gas	0.0835	UK	0.0871	Netherlands	0.0844	Netherlands	0.0928
4	Power	0.0832	Netherlands	0.0796	UK	0.0823	UK	0.0738
5	Netherlands	0.0821	Power	0.0735	Slope	0.0633	Level	0.0595
6	Curvature	0.0199	Curvature	0.0624	Power	0.0619	Slope	0.0286
7	Slope	0.0182	Level	0.0146	Level	0.0139	Gas	0.0285

	PPC5		PPC6		PPC7		PPC8	
1	Slope	0.1150	Curvature	0.1214	Slope	0.1261	Level	0.1478
2	Curvature	0.1118	Power	0.1133	Power	0.0992	Power	0.1318
3	Power	0.1037	Slope	0.1113	UK	0.0926	Netherlands	0.0891
4	Netherlands	0.0925	Netherlands	0.0914	Level	0.0755	UK	0.0776
5	UK	0.0741	UK	0.0753	Netherlands	0.0740	Curvature	0.0580
6	Gas	0.0630	Gas	0.0533	Gas	0.0675	Slope	0.0442
7	Level	0.0233	Level	0.0173	Curvature	0.0484	Gas	0.0349

	PPC9		PPC10		PPC11		PPC12	
1	Curvature	0.1346	Level	0.1089	Slope	0.1848	Level	0.2015
2	Gas	0.1244	Gas	0.0981	Gas	0.1200	Gas	0.1308
3	UK	0.0941	UK	0.0962	UK	0.0981	UK	0.0936
4	Level	0.0940	Curvature	0.0763	Netherlands	0.0686	Netherlands	0.0730
5	Netherlands	0.0726	Netherlands	0.0704	Level	0.0538	Power	0.0358
6	Power	0.0423	Power	0.0685	Power	0.0467	Slope	0.0329
7	Slope	0.0214	Slope	0.0648	Curvature	0.0113	Curvature	0.0156

The ranking of $\bar{C}_{jk}$ by Eq. (3.30) for the seven categories: shape (level, slope, curvature), types of commodities (gas, power), and location (UK, Netherlands).

$\mathbf{A} = \mathbf{B}$, $\mathbf{B}$ is the eigenvector matrix, and it reduces to an ordinary PCA. In other cases, $\mathbf{B}$ is a sparse eigenvector matrix that is proportional to $\mathbf{A}$. The second term is the L1 penalty that suppresses less significant loadings for the j^{th} column (v_j) of the eigenvector matrix with a tuning parameter λ_{1j} . The larger the λ_{1j} , the more severe the penalty for the j^{th} PC loadings. The last L2 penalty term is necessary to have a unique solution when $N \gg M$.¹¹

In our study, we apply SPCA to the standardised PCs to obtain simplified eigenvectors, leaving the only significant elements in $\bar{\mathbf{U}}$. The ‘spca’ function of the R package ‘elasticnet’ (Zou and Hastie, 2020) is used to compute SPCA with the default value of $\lambda_2 = 10^{-6}$. We retain 1/3 of variables per eigenvector and set zero everywhere else; keeping a certain number of variables is an alternative way of setting the L1 penalty introduced by the

¹¹According to Zou et al. (2006), it is advisable to set a small positive number for λ to avoid collinearity in regression even when the number of variables (N) is smaller than the number of observations M .

authors.

Table 3.6 compares the loading vectors of the original PPCA (top) and sparse PPCA (bottom) for UK gas, Dutch gas, UK power, and Dutch power. We can see that SPCA selects the first marginal PCs for PPC1, the second marginal PCs for PPC2, and the third marginal PCs for PPC3 as the primary factors to explain the common eigenstructures in the bottom table. The retained loadings indicate the accountability of the shape factors in explaining the joint dynamics of commodity forward curves, in agreement with the previously obtained results by the categorical ranking. On the other hand, the variables may have weaker associations with PPCs where the loadings are close to zero, such as the loadings of Dutch power for PPC2 and UK power for PPC3. As for PPC4, PC3 of Dutch power dominates the loading vector with a remarkably high value of 0.98. Therefore, it could be a non-common factor representing the curvature of the Dutch power forward curve alone.

Lastly, we comment on the explanatory power of SPCA. The total variation of SPCA cannot be calculated in the usual way since the sparse eigenvectors are not orthogonal. Zou et al. (2006) propose the adjusted total variation using the QR decomposition, which is shown in the top two rows in the second half of Table 3.6. Note that, in the table, the (adjusted) total variations refer to the explanatory powers of the model for the cross-correlation of PCs, not the explanatory powers for original data. As a side-effect of the shrinkage method, the adjusted total variation does not sum to unity even if we account for all the components. Moreover, the explanatory power declines by approximately 6% compared to the original PPCA.

Table 3.6: Comparison of original and sparse PPC loading vectors.

Explained %	27.31	21.94	17.95	7.11	6.98	5.93	4.32	3.48	2.18	1.34	1.01	0.42
Cumulative %	27.31	49.26	67.21	74.32	81.31	87.24	91.56	95.04	97.22	98.57	99.58	100.00
	PPC1	PPC2	PPC3	PPC4	PPC5	PPC6	PPC7	PPC8	PPC9	PPC10	PPC11	PPC12
Ug PC1	0.5240	-0.0275	-0.0333	0.0723	0.0623	0.0042	-0.1852	-0.2704	0.1539	-0.1393	-0.1526	0.7352
PC2	0.0149	0.5198	-0.2035	-0.0601	-0.2957	-0.2416	-0.1877	-0.0256	0.0963	-0.1038	0.6904	0.0787
PC3	0.0250	0.1845	0.5745	-0.1050	0.1993	-0.1729	-0.0283	0.0972	0.6389	0.3682	0.0104	-0.0237
Dg PC1	0.5219	-0.0491	-0.0620	-0.0364	0.0057	-0.0092	-0.0950	-0.1325	0.3039	-0.4691	-0.0702	-0.6132
PC2	0.0174	0.5043	-0.2600	-0.0641	-0.2651	-0.0487	-0.3336	0.0042	-0.0362	0.2999	-0.6192	-0.1229
PC3	0.0597	0.2147	0.5182	-0.0379	0.2683	-0.3440	-0.2921	-0.0010	-0.5757	-0.2622	-0.0527	-0.0421
Up PC1	0.4928	0.0637	0.0031	0.1125	0.0855	0.0531	0.2523	-0.3394	-0.3262	0.5962	0.2215	-0.2015
PC2	-0.1112	0.4666	-0.0822	0.1170	0.2113	-0.1886	0.7134	-0.1909	0.0664	-0.2674	-0.2086	0.0663
PC3	-0.0108	0.1397	0.3999	-0.5065	-0.4365	0.4977	0.1738	-0.2584	-0.0842	-0.1360	-0.0208	0.0513
Dp PC1	0.4292	0.0163	-0.0471	-0.3020	-0.1167	-0.1108	0.3043	0.7580	-0.1257	0.0106	-0.0330	0.1097
PC2	0.0256	0.3654	-0.1194	0.0100	0.5621	0.6628	-0.1638	0.2285	-0.0078	-0.0517	0.1203	-0.0028
PC3	0.0897	0.1312	0.3230	0.7740	-0.3931	0.2308	0.0426	0.2324	-0.0039	-0.0855	-0.0165	-0.0112
Explained %	26.43	18.87	14.44	7.12	7.09	6.73	4.21	3.77	2.09	1.56	1.04	0.45
Cumulative %	26.43	45.30	59.73	66.86	73.95	80.68	84.89	88.67	90.76	92.32	93.36	93.81
	PPC1	PPC2	PPC3	PPC4	PPC5	PPC6	PPC7	PPC8	PPC9	PPC10	PPC11	PPC12
Ug PC1	-0.5545	0	0	0	0	0	0.0001	-0.1933	0	-0.3592	0	0.7246
PC2	0	-0.6354	0	0	0	0	0.1544	0	0	0	-0.7565	0
PC3	0	0	0.7663	-0.0190	-0.0024	0	0	0	0.6420	0	0	0
Dg PC1	-0.5739	0	0	0	0	0	0	-0.1952	0	-0.3989	0	-0.6890
PC2	0	-0.6614	0	0	0	0	0.3969	0	0	0	0.6365	0
PC3	0	0	0.6422	-0.0051	0	-0.0006	0	0	-0.7667	0	0	0
Up PC1	-0.5075	0	0	0	0	0	0	-0.1739	0	0.8437	0	-0.0140
PC2	0	-0.3985	0	0	0	0	-0.9048	0	0	0.0001	0.1499	0
PC3	0	0	-0.0005	-0.1812	0.7436	0.6437	0	0	0	0	0	0
Dp PC1	-0.3250	0	0	0	0	0	0	0.9457	0	0	0	0.0024
PC2	0	-0.0002	0	0	-0.6545	0.7561	0	0	-0.0014	0	0.0021	0
PC3	0	0	0.0178	0.9832	0.1370	0.1181	0	0	0.0083	0	0	0

Ug...UK gas, Dg...Dutch gas, Up...UK power, Dp...Dutch power. The top table reports the loadings of marginal PCs by PPCA on the cross-correlation matrix of PCs. The bottom table shows the sparse loadings by SPCA.

3.5 Summary

PCA is a traditional yet powerful tool that focuses on the key components contributing to the variations of data under orthogonal systems. In this chapter, we reviewed the extensions of PCA for more than one data group: correlated PCA, joint PCA, and [Flury \(1988\)](#)'s common PCA and its variations. Several authors have employed these methods for the analysis and modelling of financial data. Such examples include modelling yield curves, volatility surfaces of equity options, and commodity forward curves.

In practice, we tend to overlook the implicit assumptions of the simplest multi-group PCA (e.g. correlated PCA). On the other hand, more sophisticated models (e.g. CPC) have burdens in that they are heavily dependent on the model assumptions or are subject to high computational expenses, limiting the applicability of the multi-group PCA for practical use. Our proposed method, by contrast, is the generalisation of the two-stage PCA ([Alexander and Chibumba, 1996](#)), so-called PPCA, which is a distribution-free method that does not require heavy computation and finds the common eigenstructures of multi-group data only if they exist. We derived the succinct expression for the common eigenstructures and discussed why the second stage PCA should be performed on the cross-correlation matrix of PCs.

In the empirical study, the categorical ranking of PC loadings identified the common shape factors in the forward curves of UK gas, Dutch gas, UK power, and Dutch power. The numerical analysis of PC loadings by SPCA suggested similar interpretations for the first three common eigenvectors. We shall utilise these results in the formulation of a forward curve model in the next chapter.

Chapter 4

The Joint Modelling of Commodity Forward Curves

4.1 Introduction

The aim of Chapter 4 is to adopt the common factor structures by PPCA in the HJM model ([Heath et al., 1992](#)), a popular framework for modelling commodity forward curves amongst industry practitioners and academics. The convenience of the HJM model centres on its ability to describe the evolution of the whole term structure of forward prices by the sum of volatility functions, typically estimated by PCA; see [Clewlow and Strickland \(2000\)](#). Its extension for multi-commodity forward curves is proposed by [Tolmasky and Hindanov \(2002\)](#) who impose strong conditions, including that a correlation matrix is a constant multiplication of other correlation matrices. Their model assumption resembles level 2, the proportionality of covariance matrices, in [Flury \(1988\)](#)'s hierarchical similarity analysis. In Chapter 3, we argued that real-life data may not always fit into [Flury \(1988\)](#)'s stylised eigenstructures, suggesting an alternative approach to quantify the dependence structure of commodity forward curves by PPCA. In this chapter, we are going to demonstrate how to extend the HJM model for a single forward curve to multiple forward curves with the orthogonal representation of factor volatilities using PPCA.

As mentioned earlier, the PCA-based forward curve model has been used widely. Therefore, the potential users of the PPCA-based model may wish to achieve the same model performance as the PCA-based model at a marginal level while benefiting from enhanced common factor structures in the joint modelling of those marginal forward curves. We estimate additional factors from insignificant PPCs, representing the uniqueness of forward curves, and include them in the PPCA-based forward curve model to achieve the

objectives.

The empirical section conducts a simulation study to compare the distributional properties of the PCA and PPCA-based models with varying number of common factors and examine whether the inclusion of common factors improves the joint modelling of forward curves. The impacts are analysed before and at maturities of forward contracts to consider the time-to-maturity effect in the distributions of five commodity spreads.

4.2 Forward Curve Models

4.2.1 The PC Model

The following term structure model is the generalisation of [Heath et al. \(1992\)](#)

$$\frac{dF_k(t, T)}{F_k(t, T)} = \sum_{j=1}^{n_k} \sigma_{kj}(t, T) dZ_{kj}^{\mathbb{Q}}(t), \quad (4.1)$$

where $F_k(t, T)$, $k = 1, \dots, K$, represent the price of forward contracts¹ in the k -th forward curve of interest with maturity T , $\sigma_{kj}(t, T)$, $j = 1, \dots, n_k$, are the factor volatilities that govern the dynamics of the forward curves, and $Z_{kj}^{\mathbb{Q}}(t)$ are the $\mathbb{Q}$ -Brownian motions associated with the factors. These Brownian motions can be either correlated or uncorrelated, depending on the model assumptions. We shall denote uncorrelated Brownian motions by $Z(t)$ and correlated Brownian motions by $W(t)$ hereafter. Eq. (4.1) reduces to a single forward curve model when $k = 1$.

Suppose that PCA is used to estimate the orthogonal factor volatilities. Then, it gives another expression for Eq. (4.1) with the eigenstructures:

$$\begin{pmatrix} \frac{dF_k(t, T_1)}{F_k(t, T_1)} \\ \frac{dF_k(t, T_2)}{F_k(t, T_2)} \\ \vdots \\ \frac{dF_k(t, T_N)}{F_k(t, T_N)} \end{pmatrix} = \begin{pmatrix} v_{11} & v_{12} & \cdots & v_{1n_k} \\ v_{21} & v_{22} & \cdots & v_{2n_k} \\ \vdots & \vdots & \ddots & \vdots \\ v_{N1} & v_{N2} & \cdots & v_{Nn_k} \end{pmatrix} \begin{pmatrix} \sqrt{\lambda_1} & 0 & \cdots & 0 \\ 0 & \sqrt{\lambda_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sqrt{\lambda_{n_k}} \end{pmatrix} \begin{pmatrix} dZ_{k1}^{\mathbb{Q}}(t) \\ dZ_{k2}^{\mathbb{Q}}(t) \\ \vdots \\ dZ_{kn_k}^{\mathbb{Q}}(t) \end{pmatrix} \quad (4.2)$$

We denote this PCA-based model by PC(n_k), where n_k ($\leq N$) is the number of principal components used for a marginal forward curve model. Excluding the Brownian motions, the right-hand side of Eq. (4.2) is the square root equivalent of Eq. (3.2), $\mathbf{\Gamma}_k := \mathbf{V}_k \mathbf{\Lambda}_k^{1/2}$, where v_{ij} are the elements in the eigenvector matrix $\mathbf{V}_k$ and λ_j are the diagonal elements

¹In this chapter, we treat forward and futures, as well as the associated pricing measures interchangeably.

in the eigenvalue matrix $\mathbf{\Lambda}_k$.²

When we expand the right-hand side of the equation, it is easy to see that the j -th element in the i -th row of $\mathbf{\Gamma}_k$ corresponds to the volatility of the j -th orthogonal factor:

$$\begin{pmatrix} \frac{dF_k(t, T_1)}{F_k(t, T_1)} \\ \frac{dF_k(t, T_2)}{F_k(t, T_2)} \\ \vdots \\ \frac{dF_k(t, T_N)}{F_k(t, T_N)} \end{pmatrix} = \begin{pmatrix} \sqrt{\lambda_1} v_{11} & \sqrt{\lambda_2} v_{12} & \cdots & \sqrt{\lambda_N} v_{1n_k} \\ \sqrt{\lambda_1} v_{21} & \sqrt{\lambda_2} v_{22} & \cdots & \sqrt{\lambda_N} v_{2n_k} \\ \vdots & \vdots & \ddots & \vdots \\ \sqrt{\lambda_1} v_{N1} & \sqrt{\lambda_2} v_{N2} & \cdots & \sqrt{\lambda_N} v_{Nn_k} \end{pmatrix} \begin{pmatrix} dZ_{k1}^{\mathbb{Q}}(t) \\ dZ_{k2}^{\mathbb{Q}}(t) \\ \vdots \\ dZ_{kn_k}^{\mathbb{Q}}(t) \end{pmatrix} \quad (4.3)$$

Therefore, the i -th forward contract is expressed as the sum product of row vectors in $\mathbf{\Gamma}_k$ and the uncorrelated Brownian motions $Z_{kj}^{\mathbb{Q}}(t)$, which is the right-hand-side of Eq. (4.1). From another point of view, each column vector in $\mathbf{\Gamma}_k$ represents the term structure of volatilities for the j -th orthogonal factor.

Clewlow and Strickland (2000) parametrically fit the column vectors using exponential functions to express the Samuelson Effect in the term structure of volatilities. Thompson (2016) extends the approach by introducing a parameterised forward curve model that imposes the following functional form on uncorrelated three-factor volatilities

$$\begin{aligned} g_1(t, T, X, \alpha) &= X_1 e^{-\alpha_1(T-t)} + X_2 e^{-\alpha_2(T-t)} + X_3 \\ g_2(t, T, X, \alpha) &= X_4 e^{-\alpha_1(T-t)} + X_5 e^{-\alpha_2(T-t)} + X_6 \\ g_3(t, T, X, \alpha) &= X_7 e^{-\alpha_1(T-t)} + X_8 e^{-\alpha_2(T-t)} + X_9, \end{aligned} \quad (4.4)$$

where $g_j(\cdot)$, $j = 1, 2, 3$ replace $\sigma_{kj}(t, T)$ in Eq. (4.1). In his model, the volatility functions share the same exponents so that the linkage between the uncorrelated and correlated factor structures are preserved; see Thompson (2016).

4.2.2 The PPC Model

We extend the general PC(n_k) model by introducing common eigenstructures into the volatility functions. Denote the PPCA-based forward curve model by

$$\text{PPC}_k \left(\sum_{k=1}^K n_k, n_k, m, 1 \right), \quad (4.5)$$

where $\sum_{k=1}^K n_k$ is the total number of marginal principal components used to perform PPCA, and m is the number of common factors. The model includes one additional factor in order to account for idiosyncratic behaviour of a commodity, denoted by 1. For example,

²Note that the eigenvalues and vectors do not show the subscript k for notational simplicity.

when $m = 2$, $n_1 = 3$, and $n_2 = 2$ for $K = 2$, the marginal PPC models can be written as $\text{PPC}_1(5, 3, 2, 1)$ and $\text{PPC}_2(5, 2, 2, 1)$, although we shall simply denote the two models by $\text{PPC}(5, n, m, 1)$ if $n_1 = n_2$.

As for the forward curve model, the stochastic differential equation is described by $m + 1$ factors consisting of m common and one idiosyncratic factor volatilities

$$\frac{dF_k(t, T)}{F_k(t, T)} = \sum_{j=1}^m \sigma_{kj}(t, T) d\bar{Z}_j^{\mathbb{Q}}(t) + \sigma_{k, m+1}(t, T) dW_k^{\mathbb{Q}}(t), \quad (4.6)$$

where $\bar{Z}_j^{\mathbb{Q}}(t)$ denotes the shared Brownian motions between K forward curves. In Eq. (4.6), the Brownian motions $W_k^{\mathbb{Q}}(t)$ are commodity-specific and are correlated across forward curves. On the other hand, $W_k^{\mathbb{Q}}(t)$ and $\bar{Z}_j^{\mathbb{Q}}(t)$ are uncorrelated (across and within forward curves) by assumption. Below, we demonstrate the model building blocks by setting $n_k = 3$ (the number of PCs used in the first-stage PCA) and $m = 3$ (the number of PPCs used in the second-stage PCA on the cross-correlation matrix of PCs) for illustrative purposes.

Common Factors: the Backbone of the Model

We can write the eigenstructures of the m common factor volatilities in Eq. (4.6) as

$$\begin{pmatrix} \frac{dF_k(t, T_1)}{F_k(t, T_1)} \\ \frac{dF_k(t, T_2)}{F_k(t, T_2)} \\ \vdots \\ \frac{dF_k(t, T_N)}{F_k(t, T_N)} \end{pmatrix} = \begin{pmatrix} v_{11} & v_{12} & v_{13} \\ v_{21} & v_{22} & v_{23} \\ \vdots & \vdots & \vdots \\ v_{N1} & v_{N2} & v_{N3} \end{pmatrix} \begin{pmatrix} \sqrt{\lambda_1} & 0 & 0 \\ 0 & \sqrt{\lambda_2} & 0 \\ 0 & 0 & \sqrt{\lambda_3} \end{pmatrix} \begin{pmatrix} \bar{u}_{11} & \bar{u}_{12} & \bar{u}_{13} \\ \bar{u}_{21} & \bar{u}_{22} & \bar{u}_{23} \\ \bar{u}_{31} & \bar{u}_{32} & \bar{u}_{33} \end{pmatrix} \begin{pmatrix} \sqrt{\bar{\omega}_1} & 0 & 0 \\ 0 & \sqrt{\bar{\omega}_2} & 0 \\ 0 & 0 & \sqrt{\bar{\omega}_3} \end{pmatrix} \begin{pmatrix} d\bar{Z}_1^{\mathbb{Q}}(t) \\ d\bar{Z}_2^{\mathbb{Q}}(t) \\ d\bar{Z}_3^{\mathbb{Q}}(t) \end{pmatrix}, \quad (4.7)$$

where $\bar{u}_{ij}$ are the elements in the partitioned common eigenvector matrices $\bar{\mathbf{U}}_k$ and $\bar{\omega}_j$ are the common eigenvalues in $\bar{\mathbf{\Omega}}$. All forward curves share $\bar{\omega}_j$ and the uncorrelated Brownian motions. Notice that the right-hand side of Eq. (4.7) (without Brownian motions) is the square root equivalent of Eq. (3.20), $\mathbf{H}_k := \hat{\mathbf{V}}_k \hat{\mathbf{\Lambda}}_k^{1/2} \bar{\mathbf{U}}_k \bar{\mathbf{\Omega}}^{1/2}$. The implication of this model is that when there is a shock in the j -th common system, it is transmitted to marginal forward curves through the j -th factor volatilities; however, the responsiveness to the common shock may differ between the forward curves.

Idiosyncratic Components

For the estimation of the idiosyncratic factor volatility in Eq. (4.6), we consider an $N \times 4$ matrix

$$\begin{aligned} \begin{pmatrix} \frac{dF_k(t, T_1)}{F_k(t, T_1)} \\ \frac{dF_k(t, T_2)}{F_k(t, T_2)} \\ \vdots \\ \frac{dF_k(t, T_N)}{F_k(t, T_N)} \end{pmatrix} &= \begin{pmatrix} v_{11} & v_{12} & v_{13} \\ v_{21} & v_{22} & v_{23} \\ \vdots & \vdots & \vdots \\ v_{N1} & v_{N2} & v_{N3} \end{pmatrix} \begin{pmatrix} \sqrt{\lambda_1} & 0 & 0 \\ 0 & \sqrt{\lambda_2} & 0 \\ 0 & 0 & \sqrt{\lambda_3} \end{pmatrix} \begin{pmatrix} \bar{u}_{11}\bar{u}_{12}\bar{u}_{13}\bar{u}_{14}^* \\ \bar{u}_{21}\bar{u}_{22}\bar{u}_{23}\bar{u}_{24}^* \\ \bar{u}_{31}\bar{u}_{32}\bar{u}_{33}\bar{u}_{34}^* \end{pmatrix} \begin{pmatrix} \sqrt{\bar{\omega}_1} & 0 & 0 & 0 \\ 0 & \sqrt{\bar{\omega}_2} & 0 & 0 \\ 0 & 0 & \sqrt{\bar{\omega}_3} & 0 \\ 0 & 0 & 0 & \sqrt{\bar{\omega}_4^*} \end{pmatrix} \begin{pmatrix} d\bar{Z}_1^Q(t) \\ d\bar{Z}_2^Q(t) \\ d\bar{Z}_3^Q(t) \\ dW_k^Q(t) \end{pmatrix} \\ &= \begin{pmatrix} \eta_{11} & \eta_{12} & \eta_{13} & \eta_{14}^* \\ \eta_{21} & \eta_{22} & \eta_{23} & \eta_{24}^* \\ \vdots & \vdots & \vdots & \vdots \\ \eta_{N1} & \eta_{N2} & \eta_{N3} & \eta_{N4}^* \end{pmatrix} \begin{pmatrix} d\bar{Z}_1(t) \\ d\bar{Z}_2(t) \\ d\bar{Z}_3(t) \\ dW_k^Q(t) \end{pmatrix}, \end{aligned}$$

such that the right-hand side of the first equation includes additional elements $\bar{\mathbf{u}}_4^* = (\bar{u}_{14}^*, \bar{u}_{24}^*, \bar{u}_{34}^*)'$ and $\sqrt{\bar{\omega}_4^*}$. Therefore, our objective reduces to finding an optimal vector for $\boldsymbol{\eta}_4^* = (\eta_{14}^*, \eta_{24}^*, \dots, \eta_{N4}^*)'$ in the second equation.

We propose two estimation methods to obtain the vector. The first method, which we call *local volatility fitting*, sets the calibration target to be the square root variance of a forward curve. That is, the PPC model intends to mimic the term structure of volatilities estimated by the PC model. Since the (i, i) -th elements of the covariance matrix in the original system and the uncorrelated system have one-to-one relationships by construction,³ we can equate the PC estimate of volatilities (left) and PPC estimate of volatilities (right)

$$(\gamma_{i1}^2 + \gamma_{i2}^2 + \gamma_{i3}^2)^{\frac{1}{2}} = (\eta_{i1}^2 + \eta_{i2}^2 + \eta_{i3}^2 + \eta_{i4}^{*2})^{\frac{1}{2}}, \quad (4.8)$$

where $(\gamma_{i1}, \gamma_{i2}, \gamma_{i3})^{1/2}$ and $(\eta_{i1}, \eta_{i2}, \eta_{i3}, \eta_{i4}^*)^{1/2}$ are the norm of i -th row vectors in $\hat{\mathbf{\Gamma}}_k$ ($:= \hat{\mathbf{V}}_k \hat{\mathbf{\Lambda}}_k^{1/2}$) and $\mathbf{H}_k^*$ ($:= \hat{\mathbf{\Gamma}}_k \bar{\mathbf{U}}_k^* \bar{\mathbf{\Omega}}_k^{*1/2}$), respectively. Therefore, solving for η_{i4}^* , $i = 1, \dots, N$ in Eq. (4.8) gives the exact solution for the unknown vector $\boldsymbol{\eta}_4^*$. The advantage of this method is that it exactly replicates the term structure of volatilities estimated by the PC(3) model. On the other hand, the downside is that it only focuses on the diagonal entries of the covariance matrix, ignoring the covariance of forward contracts within a forward curve.

The second method, *local covariance fitting*, sets the calibration target to be the PC

³See Eq. (3.21).

estimate of an entire covariance matrix to circumvent the problem. In other words, it considers not only the diagonal elements but also the off-diagonal elements of the covariance matrix by minimising the Frobenius distance between two model covariance matrices to find the optimal vector $\boldsymbol{\eta}_4^*$

$$\arg \min_{\boldsymbol{\eta}_4} \sqrt{\sum_{i=1}^N \sum_{j=1}^N \|\boldsymbol{\Sigma}_{(i,j)}^{pc} - \boldsymbol{\Sigma}_{(i,j)}^{ppc}\|_1^2}, \quad (4.9)$$

where $\boldsymbol{\Sigma}^{pc}(i, j)$ and $\boldsymbol{\Sigma}^{ppc}(i, j)$, $i, j = 1, \dots, N$, are the (i, j) -th elements in the model covariance matrices by PCA and PPCA, respectively. Albeit an imperfect fit to the total variation, it better captures the covariance structures of forward contracts within a curve. Having obtained the complete $\mathbf{H}_k^*$ matrix by either method, the column vectors are fitted parametrically to the functional form of Eq. (4.4) to describe the time decaying factor volatilities.⁴

Factor Correlations

The idiosyncratic factor of the PPC model is derived from the ‘leftovers’ after extracting m common factors for each forward curve. Due to the estimation method, they may include residual cross-correlations. Thus, we impose the following model cross-covariance structure to complete the model specification.

Supposing two forward curves, the model cross-covariance structure is given as

$$\begin{aligned} \mathbb{Cov}\left(\frac{dF_1(t, T)}{F_1(t, T)}, \frac{dF_2(t, T)}{F_2(t, T)}\right) &= \mathbb{Cov}\left(\sum_{j=1}^m \sigma_{1j}(t, T) d\bar{Z}_j^{\mathbb{Q}}(t) + \sigma_{1,m+1}(t, T) dW_1^{\mathbb{Q}}(t), \right. \\ &\quad \left. \sum_{j=1}^m \sigma_{2j}(t, T) d\bar{Z}_j^{\mathbb{Q}}(t) + \sigma_{2,m+1}(t, T) dW_2^{\mathbb{Q}}(t)\right) \\ &= \left(\sum_{j=1}^m \sigma_{1j}(t, T) \sigma_{2j}(t, T) + \rho_{12} \sigma_{1,m+1}(t, T) \sigma_{2,m+1}(t, T)\right) dt, \end{aligned} \quad (4.10)$$

where

$$\begin{aligned} d\bar{Z}_j^{\mathbb{Q}}(t) d\bar{Z}_l^{\mathbb{Q}}(t) &= \begin{cases} dt, & \text{if } j = l \\ 0, & \text{otherwise} \end{cases} \\ dW_k^{\mathbb{Q}}(t) d\bar{Z}_j^{\mathbb{Q}}(t) &= 0 \\ dW_1^{\mathbb{Q}}(t) dW_2^{\mathbb{Q}}(t) &= \rho_{12} dt, \end{aligned} \quad (4.11)$$

⁴The number of volatility functions required for a PPC model is $m + 1$.

for $k = 1, 2$ and $j, l = 1, \dots, m$. Alternatively, it can be expressed by matrices

$$\Sigma_{12}^{ppc} = \mathbf{V}_1 \Lambda_1^{1/2} \bar{\mathbf{U}}_1^* \bar{\Omega}_1^{*1/2} \bar{\mathbf{R}}_{12} \bar{\Omega}_2^{*1/2} \bar{\mathbf{U}}_2^{*'} \Lambda_2^{1/2} \mathbf{V}_2', \quad (4.12)$$

where

Σ_{12}^{ppc} : Model implied cross-market covariance matrix

$\bar{\Omega}_k^*$: Model eigenvalue matrix with an additional element

$\bar{\mathbf{U}}_k^*$: Model eigenvector matrix with an additional vector

and the factor cross-correlation matrix

$$\bar{\mathbf{R}}_{12} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & 1 & 0 \\ 0 & \cdots & 0 & \rho_{12} \end{pmatrix} \in \mathbb{R}^{(m+1) \times (m+1)}. \quad (4.13)$$

The zero entries in $\bar{\mathbf{R}}_{12}$ impose the model assumption in Eq. (4.11). The Brownian motions of idiosyncratic factors are correlated by ρ_{12} which is the only unknown parameter in the model.

The estimation of ρ_{12} relies on a similar technique to that of idiosyncratic factors; we calibrate the unknown parameter as the residual component between the model cross-covariance matrix of the PC and PPC models, minimising the objective function

$$\begin{aligned} \arg \min_{\rho_{12}} & \sqrt{\sum_{i=1}^N \sum_{j=1}^N \|\Sigma_{12(i,j)}^{pc} - \Sigma_{12(i,j)}^{ppc}\|_1^2} \\ \text{s.t.} & \quad -1 \leq \rho_{12} \leq 1 \end{aligned} \quad (4.14)$$

where $\Sigma_{12(i,j)}^{pc}$ and $\Sigma_{12(i,j)}^{ppc}$, $i, j = 1, \dots, N$, are the (i, j) -th elements in the cross-covariance matrix estimated by the PC and PPC models.

4.3 Implementation

4.3.1 Parameter Estimation

The model parameters are estimated from the same data set as in Section 3.4.1. Fig. 4.1 depicts the model volatilities as a function of time to maturity, in which the column

labelled $m = 0$ refers to the PC(3) model with no common factor, and shows the three factor volatilities in one figure. Other columns show the volatility of common factors for the PPC(12, 3, m , 1) models for $m = 1, 2, 3$, where the $m + 1$ volatility functions represent the idiosyncratic factor, estimated either by the local volatility fitting (purple dots) or the local covariance fitting (green dots) together with the fitted lines. Tables 4.1 – 4.3 report the parameter values and root-mean-square error (RMSE). In the estimation procedure, we obtain the parameters of common factor volatilities first, then impose the β_1 and β_2 as constraints to calibrate the rest of the parameter values for the idiosyncratic factors. In this way, the parameters of common factors are unaffected by choice of calibration method for the idiosyncratic factors.

In Fig. 4.1, all models similarly calculate the term structure of volatilities for the first three factors. On the other hand, the PPC models smooth out the term structures overall, most notably when there is a hump in the shape of short-term volatilities, leading to smaller fitting errors than the PC model. The estimated decay parameters are similar between PC(3) and PPC(12, 3, 3, 1), indicating a high speed of mean reversion for both the first and second factors, while other PPC models suggest relatively slow mean-reversion rates for one or both of the factors.

In Fig. 4.1, the local volatility fitting estimates the idiosyncratic volatilities more consistently than the local covariance fitting, regardless of the number of common factors. Table 4.4 reports the model-implied factor correlations between UK gas (W_1), Dutch gas (W_2), UK power (W_3), and Dutch power (W_4) forward curves implied by the estimation method together with the Frobenius distance for the local covariance fitting. While the magnitude of the correlation is on average higher for the local covariance fitting, both methods estimate the sign and magnitude of correlation coefficients stably for $m = 2, 3$.

As for the PC(n_k) model, the factor dependence structure is typically modelled by the correlation of PCs across forward curves (cross-correlations). Table 4.5 shows the historical estimation of the cross-correlation matrix for the PC(3) model. The cross-correlation appears the strongest in the diagonal elements (93%, 86%, and 73% for PC1 – PC1, PC2 – PC2, PC3 – PC3, respectively) of the UK gas – Dutch gas submatrix, possibly due to the active trade flows between their physical gas markets. On the other hand, the diagonal elements (60%, 39%, and 20%) of the UK power – Dutch power submatrix reveal the weakest association of PCs albeit the same commodity, which probably reflects the underdevelopment of the cross-regional power trades during the observation period.




Figure 4.1: Factor volatilities and exponentially fitted curves.

The dots represent the model volatilities, and the lines are the fitted curves. The first column shows the result for the PC(3) model. The second to the last columns show the results for the PPC(12, 3, m , 1) model with $m = 1, 2, 3$.

Table 4.1: Exponentially fitted model parameters for the PC(3) and PPC(12, 3, 1, 1) models.

PC(3)												
	X1	X2	X3	X4	X5	X6	X7	X8	X9	α_1	α_2	RMSE
UK gas	-2.182	2.415	0.167	4.843	-4.353	-0.065	-10.620	10.450	-0.055	5.622	5.403	0.0134
Dutch gas	-1.664	1.808	0.169	2.354	-1.962	-0.043	-6.173	5.994	-0.037	6.601	6.228	0.0084
UK power	-7.560	7.744	0.123	-2.616	3.018	-0.075	-17.061	16.305	-0.036	9.998	9.522	0.0127
Dutch power	0.271	-0.120	0.158	0.563	-0.129	-0.052	-9.073	8.697	-0.035	8.137	7.667	0.0101
PPC(12, 3, 1, 1)												
Common	Y1	Y2	Y3						β_1	β_2	RMSE	
UK gas	0.179	0.104	0.130						6.169	1.359	0.0019	
Dutch gas	0.116	0.129	0.094						7.979	0.767	0.0008	
UK power	0.217	-0.364	0.484						4.042	-0.020	0.0032	
Dutch power	0.206	-0.256	0.246						1.666	0.375	0.0052	
Idiosyncratic	Local vol	I1	I2	I3	RMSE	Local cov			I1	I2	I3	RMSE
UK gas	0.192	-0.038	0.106	0.016				0.525	-0.124	-0.049	0.0162	
Dutch gas	0.184	-0.040	0.100	0.013				-0.047	-0.008	0.079	0.0086	
UK power	0.309	1.683	-1.662	0.012				0.447	3.280	-3.365	0.0125	
Dutch power	0.290	-0.455	0.366	0.025				0.084	-0.142	0.186	0.0118	

The parametric fit to the [Thompson \(2016\)](#)'s model in Eq. (4.4). X_i : fitted factor volatilities for the PC(3) model. α_i : shared exponents. Y_i : fitted common factor volatilities for the PPC(12, 3, 1, 1) model. I_i : fitted idiosyncratic factor volatilities reported for the local volatility fitting (left columns) and the local covariance (right columns). β_i shared exponents. The RMSE reports the fitting errors.

Table 4.2: Exponentially fitted model parameters for the PPC(12, 3, 2, 1) model.

PPC(12, 3, 2, 1)												
Common		Y1	Y2	Y3	Y4	Y5	Y6		β_1	β_2	RMSE	
UK gas		0.606	-0.373	0.150	2.595	-2.309	-0.042		2.890	2.567	0.0096	
Dutch gas		0.120	0.139	0.058	0.268	-0.095	0.003		5.031	0.389	0.0058	
UK power		0.854	-0.638	0.113	3.154	-2.889	-0.021		2.889	2.658	0.0088	
Dutch power		0.504	-0.411	0.114	1.694	-1.519	-0.021		2.852	2.616	0.0065	
Idiosyncratic	Local vol	I1	I2	I3	RMSE			Local cov	I1	I2	I3	RMSE
UK gas		0.755	-0.637	0.071	0.010				0.862	-0.768	0.063	0.0066
Dutch gas		0.143	-0.052	0.091	0.010				0.126	-0.069	0.097	0.0096
UK power		1.986	-1.756	0.062	0.009				2.506	-2.280	0.050	0.0072
Dutch power		1.685	-1.599	0.135	0.018				0.747	-0.723	0.119	0.0106

The parameters are fitted to Eq. (4.4). Y_i : fitted common factor volatilities. I_i : fitted idiosyncratic factor volatilities reported for the local volatility fitting (left columns) and the local covariance (right columns). β_i : the shared exponents. The RMSE reports the fitting errors.

Table 4.3: Exponentially fitted model parameters for the PPC(12, 3, 3, 1) model.

PPC(12, 3, 3, 1)													
Common		Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	β_1	β_2	RMSE
UK gas		-2.202	2.428	0.154	0.963	-0.610	-0.079	-9.418	9.118	-0.035	5.622	5.382	0.0120
Dutch gas		-1.797	1.930	0.153	0.154	0.095	-0.061	-4.812	4.528	-0.028	6.352	5.951	0.0077
UK power		-5.361	5.461	0.126	-6.004	6.160	-0.052	-8.945	8.459	-0.012	9.962	9.448	0.0099
Dutch power		-0.941	1.020	0.114	-1.124	1.310	-0.035	-3.580	3.329	-0.019	7.896	7.342	0.0065
Idiosyncratic	Local vol	I1	I2	I3	RMSE				Local cov	I1	I2	I3	RMSE
	UK gas	-0.108	0.209	0.062	0.004				-1.049	1.109	0.050	0.0018	
	Dutch gas	0.093	-0.009	0.057	0.004				-0.315	0.363	0.052	0.0029	
	UK power	-2.827	3.060	0.055	0.006				-2.862	3.106	0.033	0.0026	
	Dutch power	2.456	-2.217	0.121	0.008				1.304	-1.209	0.112	0.0054	

The parameters are fitted to Eq. (4.4). Y_i : fitted common factor volatilities. I_i : fitted idiosyncratic factor volatilities reported for the local volatility fitting (left columns) and the local covariance (right columns). β_i : the shared exponents. The RMSE reports the fitting errors.

Table 4.4: PPC factor correlations.

Local volatility															
Correlation															
$m = 1$					$m = 2$					$m = 3$					
	W_1	W_2	W_3	W_4		W_1	W_2	W_3	W_4		W_1	W_2	W_3	W_4	
W_1	1	0.141	-0.053	-0.336	W_1	1	0.200	-0.061	-0.448	W_1	1	0.194	-0.067	-0.540	
W_2		1	-0.241	-0.145	W_2		1	-0.300	-0.176	W_2		1	-0.352	-0.248	
W_3			1	-0.184	W_3			1	-0.235	W_3			1	-0.252	
W_4				1	W_4				1	W_4				1	
Local covariance															
Correlation															
$m = 1$					$m = 2$					$m = 3$					
	W_1	W_2	W_3	W_4		W_1	W_2	W_3	W_4		W_1	W_2	W_3	W_4	
W_1	1	-0.323	0.320	0.226	W_1	1	0.314	-0.097	-0.671	W_1	1	0.248	-0.093	-0.723	
W_2		1	-0.473	-0.252	W_2		1	-0.421	-0.245	W_2		1	-0.449	-0.316	
W_3			1	-0.233	W_3			1	-0.328	W_3			1	-0.334	
W_4				1	W_4				1	W_4				1	
Frobenius distance															
$m = 1$					$m = 2$					$m = 3$					
	W_1	W_2	W_3	W_4		W_1	W_2	W_3	W_4		W_1	W_2	W_3	W_4	
W_1		0.0437	0.0311	0.0555	W_1		0.0162	0.0164	0.0194	W_1		0.0054	0.0132	0.0144	
W_2			0.0222	0.0221	W_2			0.0138	0.0129	W_2			0.0128	0.0076	
W_3				0.0326	W_3				0.0140	W_3				0.0135	
W_4					W_4					W_4					

W_1 ...UK gas, W_2 ...Dutch gas, W_3 ...UK power, W_4 ...Dutch power. W_k , $k = 1, \dots, 4$, represents the Brownian motions of idiosyncratic factors. The top two panels report the model implied correlation coefficients for idiosyncratic factors by the factor estimation method. The last column reports the Frobenius distance for the local covariance fitting given by Eq. (4.14).

Table 4.5: Cross-correlation matrix of three PCs.

		Ug			Dg			Up			Dp		
		PC1	PC2	PC3	PC1	PC2	PC3	PC1	PC2	PC3	PC1	PC2	PC3
Ug	PC1	1	0	0	0.928	0.024	0.071	0.829	-0.233	-0.095	0.599	0.046	0.121
	PC2	0	1	0	0.005	0.857	0.078	0.052	0.565	0.048	0.084	0.324	0.054
	PC3	0	0	1	-0.033	-0.092	0.733	0.037	0.141	0.438	0.012	0.051	0.310
Dg	PC1	0.928	0.005	-0.033	1	0	0	0.771	-0.240	-0.072	0.677	0.008	0.058
	PC2	0.024	0.857	-0.092	0	1	0	0.056	0.490	0.033	0.071	0.422	0.026
	PC3	0.071	0.078	0.733	0	0	1	0.123	0.135	0.312	0.023	0.069	0.280
Up	PC1	0.829	0.052	0.037	0.771	0.056	0.123	1	0	0	0.596	0.113	0.188
	PC2	-0.233	0.565	0.141	-0.240	0.490	0.135	0	1	0	-0.113	0.393	0.049
	PC3	-0.095	0.048	0.438	-0.072	0.033	0.312	0	0	1	0.033	0.017	0.195
Dp	PC1	0.599	0.084	0.012	0.677	0.071	0.023	0.596	-0.113	0.033	1	0	0
	PC2	0.046	0.324	0.051	0.008	0.422	0.069	0.113	0.393	0.017	0	1	0
	PC3	0.121	0.054	0.310	0.058	0.026	0.280	0.188	0.049	0.195	0	0	1

Ug...UK gas, Dg...Dutch gas, Up...UK power, Dp...Dutch power.

The correlation coefficients are estimated from the time series of PCs calculated for the period of 1st July 2016 to 29th Jun 2018. The block diagonal elements are identity matrices representing the within-curve PC correlations.

4.3.2 Simulation

To analyse the effects of common factors in a forward curve model, we first explain the simulation procedure for the PPC model. Suppose that PCA is performed separately on the covariance matrix of forward contracts followed by PPCA on the cross-correlation matrix PCs, from which we obtain the parametrically fitted factor volatilities and correlations as reported above.

PPC Model Simulation

Outside the simulation loop:

1. Discretise Eq. (4.6).
2. Form a $K \times K$ matrix for the correlation of idiosyncratic factors⁵ and apply the Cholesky decomposition to the matrix.

For each time step t in a simulation loop:

3. Generate standard normal random variables $\bar{\varepsilon}_t = (\bar{\varepsilon}_1(t), \dots, \bar{\varepsilon}_m(t))'$, $\bar{\varepsilon}_j(t) \sim N(0, 1)$ for common factors.
4. Generate standard normal random variables $\varepsilon_t = (\varepsilon_1(t), \dots, \varepsilon_K(t))'$, $\varepsilon_k(t) \sim N(0, 1)$, for idiosyncratic factors.
5. Pre-multiply the lower-triangular Cholesky matrix to ε_t to correlate the random variables for idiosyncratic factors.

For each forward curve at t :

6. Calculate the common and idiosyncratic factor volatilities from fitted parameters.
7. Assign the random variables in $\bar{\varepsilon}_t$ and ε_t to the corresponding Brownian motions to simulate the forward prices.

Next, we describe the traditional multi-commodity modelling with the model specification Eq. (4.1), which is a no-common factor model whose joint price dynamics with other forward curves are typically captured by the correlated PCA described in Section 3.3.1. This approach requires the generation of correlated random variables to relate the factors

⁵Such as the 4×4 correlation matrices reported in Table 4.4.

between forward curves. Such random variables are obtained by applying the Cholesky decomposition to the cross-factor correlation matrix

$$\mathbf{Q} = \mathbf{L}\mathbf{L}', \quad (4.15)$$

where $\mathbf{L} \in \mathbb{R}^{n^* \times n^*}$ is the lower-triangular Cholesky matrix and

$$\mathbf{Q} = \begin{pmatrix} \mathbf{R}_{11} & \mathbf{R}_{12} & \cdots & \mathbf{R}_{1K} \\ \mathbf{R}_{21} & \mathbf{R}_{22} & \cdots & \mathbf{R}_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{R}_{K1} & \mathbf{R}_{K2} & \cdots & \mathbf{R}_{KK} \end{pmatrix} \in \mathbb{R}^{n^* \times n^*}. \quad (4.16)$$

In $\mathbf{Q}$, $\mathbf{R}_{kl} \in \mathbb{R}^{n_k \times n_k}$ are the cross-correlation matrices of PCs when $k \neq l$ and are identity matrices when $k = l$.

The main inputs of the PC model simulation, therefore, are the exponentially fitted parameters of factor volatilities and the cross-correlation matrix $\mathbf{Q}$. As is the case for the PPC model, the simulation of price paths involves discretisation of Eq. (4.1) and assignment of correlated random variables to the factor volatilities of marginal forward curves at time t . We summarise the simulation procedure as follows:

PC Model Simulation

Outside the simulation loop:

1. Discretise Eq. (4.1).
2. Apply the Cholesky decomposition to $\mathbf{Q}$.

For each time step t in a simulation loop:

3. Generate the standard normal random variables $\boldsymbol{\varepsilon}_t = (\varepsilon_1(t), \dots, \varepsilon_{n^*}(t))'$, $\varepsilon_j(t) \sim N(0, 1)$, for all factors at a time.
4. Pre-multiply the lower-triangular Cholesky matrix to $\boldsymbol{\varepsilon}_t$ to generate correlated random variables $\tilde{\boldsymbol{\varepsilon}}_t = (\tilde{\varepsilon}_1(t), \dots, \tilde{\varepsilon}_{n^*}(t))'$, $\tilde{\varepsilon}_j(t) \sim N(0, 1)$.

For each forward curve at t :

5. Calculate the factor volatilities from fitted parameters.
6. Assign the random variables in $\tilde{\boldsymbol{\varepsilon}}_t$ to the corresponding Brownian motions to simulate the forward prices.

Comparing the two simulation procedures, the most distinctive feature of the PPC model can be found in the generation of random variables; the same sources of uncertainty drive all forward curves simultaneously, unlike the traditional PC model in which correlated, yet different sources of uncertainties affect the forward curve dynamics. Hence, the PPC model would better suit an integrated market environment where commodities are exposed to the same risks. Moreover, when m is very small relative to $n^*(=\sum_{k=1}^K n_k)$, one can expect computational efficiency with the PPC model since it only requires $m + K$ factors in total, fewer than n^* for the PC model.

The Monte Carlo simulation computes fixed time-of-maturity contracts with $T=(30, 60, 90, 120, 150, 180, 210, 240, 360, 540)$ days for all forward curves. We set the starting values of the forward curves flat and equal to the average of the last observation in each data set. These values are 59.35 GBp/Therm, 21.94 EUR/MWh, 57.29 EUR/MWh, and 51.39 EUR/MWh for UK gas, Dutch gas, UK power and Dutch power, respectively. The model forward curves are simulated 10,000 times, assuming 252 trading days in a year.

4.3.3 Results

The PPC models should produce identical marginal distributions as the PC model by model design. Fig. 4.2 (local volatility fitting) and Fig. 4.3 (local covariance fitting) depict the simulated distribution of five representative contracts, for which the prices are transformed into log returns to visualise the distribution by box plots. The boxes indicate the middle 50% of the distribution, covering approximately 99% of the distribution with the whiskers. The x-axes display the number of common factors used in the PC(3) and PPC(12, 3, m , 1) models, where $m = 0$ refers to the PC(3) model with no common factor. These figures suggest that the overall distribution of simulated paths is consistent between the PC and PPC models regardless of the factor estimation method. However, the PPC model with $m = 1$ tends to underestimate the dispersion of log returns for short to mid-term contracts, especially for the UK and Dutch power data. The simulated distribution approaches the PC(3) benchmark as m increases to 3.

The subsequent analysis focuses on commodity spreads to investigate the impact of common factors on the joint price distributions. The general definition of a commodity spread refers to the price differential between two or more underlying assets with respect to commodity types, quality, maturity dates, or locations.⁶ In the case of two distinct

⁶The practical importance of commodity spreads and their application in the valuation of physical assets are discussed in Chapter 5.




Figure 4.2: Simulated marginal distributions (local volatility fitting).

The x-axes indicate the number of common factors. The boxes cover the middle 50% of the normal distribution with the mean values of log returns in red lines. The whiskers stretch to approximately $\pm 2.7\sigma$ from the mean to cover 99.3% of the distribution.

commodities, the time- t value of a commodity spread may be written as

$$V(t) = q_1 F_1(t, T) - q_2 F_2(t, T), \quad 0 \leq t \leq T. \quad (4.17)$$

Where necessary, the constants q_1 and q_2 adjust for differences in the currency or unit of measurements between commodities, and sometimes for fuel efficiencies.

Table 4.6 summarises five commodity pairs of our choice including cross-commodity (spread 1, 2, 5) and inter-regional (spread 3, 4, 5) spreads.⁷ As far as the inter-regional spreads are concerned, the units of UK assets are converted into EUR/MWh with the conversion factor of 1 Therm = 0.0293071 MWh and fixed foreign exchange rate of GBP/EUR = 1.13035 prevailing at the last observation in our data set. As for the calculation of power – gas spreads, the price of gas is multiplied by the heat rate: $h = 1/0.50$

⁷We exclude the Dutch power – UK gas spark spread from our study since the physical trade flow of natural gas was one way from the Netherlands to the UK before 2019.




Figure 4.3: Simulated marginal distributions (local covariance fitting).

The x-axes indicate the number of common factors. The boxes cover the middle 50% of the normal distribution with the mean values of log returns in red lines. The whiskers stretch to approximately $\pm 2.7\sigma$ from the mean to cover 99.3% of the distribution.


(= 1/fuel efficiency).⁸ When these adjustments are not necessary, q_1 , q_2 , or h are set equal to 1. The last two columns in Table 4.6 show the adjusted units for each pair of commodities. Except for the number of common factors in the models, we assume ceteris paribus to simulate i) the spot spread payoff at the maturity of the forward contracts, ii) the spread payoff of forward contracts Δ days before the maturity, iii) the spread payoff of forward contracts Δ days after the inception. In other words, T , $T - \Delta$, and $t + \Delta$ substitute t in Eq. (4.17) for the conditions i) to iii) respectively, where we set $\Delta = 90$ days.

⁸The heat rate is defined as the inverse of a fuel efficiency rate. The fuel efficiency represents the amount of fuel needed to generate one unit of electricity, and it varies depending on fuel types and power generation plants. We use a typical fuel efficiency rate of 50% for European energy data in reference to the Intercontinental Exchange website: <https://www.theice.com/products/67738090/UK-Spark-Spread> and <https://www.theice.com/products/81743160/Dutch-Spark-Spread-TTF>.

Table 4.6: Five commodity pairs.

	Asset 1	Asset 2	Ccy 1	Ccy 2	Unit 1	Unit 2	Ccy*	Unit*
1	UK power	UK gas	GBP	GBp	MWh	Therm	GBP	MWh
2	Dutch power	Dutch gas	EUR	EUR	MWh	MWh	EUR	MWh
3	UK gas	Dutch gas	GBp	EUR	Therm	MWh	EUR	MWh
4	UK power	Dutch power	GBP	EUR	MWh	MWh	EUR	MWh
5	UK power	Dutch gas	GBP	EUR	MWh	MWh	EUR	MWh

The table shows five commodity pairs with units of measurements and currencies for each underlying asset. *Ccy** (Currency) and *Unit** represent the adjusted units of measurements and currencies used for the spread payoff. For UK assets, GBp and GBP refer to pence sterling and pound sterling.

**Figure 4.4: The means of simulated UK power - UK gas spread at the maturity of contracts (local volatility fitting).**

The mean of simulated spread payoffs by the number of common factors, where the solid black squares indicate the PC(3) model with no common factor ($m = 0$).

As represented by Fig. 4.4 for UK power – UK gas (other spreads yield very similar plots), the mean values of simulated spread payoff stay constant across maturities, complying with the martingale property of the underlying processes. For case i), Fig. 4.5 depicts the standard deviations by contract maturities (in days) for local volatility fitting (left) and local covariance fitting (right). Except for UK gas – Dutch gas, the standard deviations are generally lower for the PPC models than the PC(3) model when local volatility fitting is used to estimate the idiosyncratic factors. When the number of common factors increases, the dominance of idiosyncratic factors decreases in the PPC model, which strengthens the power of common factors to lower the spread dispersion. On the right column, local covariance fitting follows a similar tendency for the PPC model. However, the standard deviations seem less sensitive to the number of common factors used in the model, providing stable outputs across maturities.




Figure 4.5: Standard deviations of spread payoffs at the maturity of forward contracts.

The left and right columns depict the standard deviation of simulated spread payoffs by the factor estimation method and the number of common factors, where the solid black squares indicate the PC(3) model with no common factor ($m = 0$).

The spread distributions exhibit more interesting patterns for forward payoffs. Fig. 4.6 depicts the standard deviation of case ii), where the time-to-maturity of the contracts is the x-axis minus 90 days (and therefore there is no distribution for 30, 60 and 90 day contracts). As seen in the left column, the PPC model with $m = 1$ increases the standard deviation of the forward spreads, while the models with $m = 2, 3$ are relatively close to the PC(3) model. On the other hand, in the right column, the PPC model by the local covariance fitting estimates lower standard deviations than the PC(3) model, and the gap between the common and non-common factor models widens for longer-term contracts.




Figure 4.6: Standard deviations of spread payoffs 90 days before the maturity of forward contracts.

The left (local volatility) and right (local covariance) columns depict the standard deviation of simulated spread payoffs 90 days before the contract maturities shown on the x-axis. There is no distribution for 30, 60 and 90 day contracts.

For case iii), the spread payoffs are calculated after 90 days from the starting time. Hence, the standard deviations accumulate the same amount of time and depict decay in their term structure as shown in Fig. 4.7. As suggested by case i), the PPC models lower the standard deviation of spot payoffs (90-day contracts) compared to the PC(3) models for both local volatility fitting and local covariance fitting. Meanwhile, the common factor models exhibit distinct patterns in different parts of the term structure for UK gas – Dutch gas, UK power – Dutch power, and UK power – Dutch gas due to the Samuelson Effect for each commodity: a) the widest gap for the spot payoff, b) no/less significant differences in the middle, c) widening gap at the long-end of the term structure. These tendencies




Figure 4.7: Standard deviations of spread payoffs 90 days after the inception of forward contracts.

The left (local volatility) and right (local covariance) columns depict the term structure of standard deviations by simulated forward spreads. The 30-day and 60-day contracts have expired.

appear mostly when the PPC model consists of more than two common factors plus one idiosyncratic factor estimated by the local covariance fitting.

4.4 Summary

This chapter presented the building blocks for the joint modelling of commodity forward curves, extending [Heath et al. \(1992\)](#). The proposed model, which consists of a few common factors and one idiosyncratic factor per forward curve, employs a data-driven factor estimation method by PPCA and imposes fewer assumptions on the eigenstructures of factor volatilities compared to [Tolmasky and Hindanov \(2002\)](#).

We suggested two alternative methods to estimate the idiosyncratic factor volatility: local volatility fitting and local covariance fitting. The PPC(12, 3, m , 1) model with a number of common factors $m = 1, 2, 3$ indicated that the PPC model with $m = 1$ is often insufficient to reproduce the marginal distributions of the PC(3) benchmark model. Irrespective of the factor estimation method by local volatility fitting or local covariance fitting, it is likely that the PPC model requires at least two common factors to adequately replicate the marginal distribution of forward curves for the European energy data analysed.

The second part of the empirical study analysed the joint distribution of forward curves for five commodity pairs, looking at mainly the standard deviations of the price spreads by the number of common factors. Due to the focused common factor modelling, the PPC models are expected to strengthen the dependence of related forward curves compared to the benchmark PC model, therefore lowering the standard deviation of commodity spreads. The effects of having common factors were insignificant when the local volatility fitting estimated the idiosyncratic factors. By contrast, the common factors consistently dampened the spread distributions in most cases when local covariance fitting is used to estimate the idiosyncratic factors, adding downward pressures to the standard deviations of the spot and forward payoffs.

These results highlighted the PPC models' notable feature in stabilising the dispersion of price spreads for closely linked commodities compared to the traditional PC model. In the next chapter, we shall investigate the impact of including common factors when pricing derivatives on multiple closely linked commodities.

Chapter 5

Pricing Commodity Spread Derivatives

5.1 Introduction

In the previous chapter, we proposed a new factor estimation method to introduce the dependence structures of commodities in the volatility functions of an established forward curve model. We demonstrated that the common factors in the PPC model potentially stabilise the fluctuations of commodity price spreads compared to a non-common factor model. As an extension of the modelling framework, Chapter 5 explores the impact of common factors on options written on the spread of two commodities with a zero strike, known as exchange options. We adapt the common factor structures to the [Margrabe \(1978\)](#) formula and investigate how our previous findings affect the option prices. Furthermore, we extend the option pricing formula for the trade where assets are denominated by different currency units in response to the recent developments in energy markets, in which the number of cross-currency transactions has risen due to improved trading infrastructure and transportation. To our knowledge, we are the first to extend [Margrabe \(1978\)](#) formula for foreign – domestic spreads on commodity prices.

In general, a spread refers to the price difference between two or more underlying assets or indexes, a vital measure for evaluating physical assets and managing risks in energy markets. Such examples include a spark spread for the price of power and natural gas representing the gross profit margin of gas-fired plants and a crack spread for the price of crude oil and its byproduct estimating the gross processing margin of petroleum products. A derivative contract on the spread of two underlying assets S_1 and S_2 with a fixed strike K is called a spread option. The option holder has the right to obtain the difference in

the asset values if it exceeds the strike price at expiry. Sometimes the strike price is set to zero for simplicity and analytical tractability, then it becomes an *exchange option* with payoff

$$V(T) = (S_1(T) - S_2(T))^+, \quad (5.1)$$

where the notation $(\cdot)^+$ indicates the positive spread.

The closed-form solution is derived in [Margrabe \(1978\)](#), who regards the pricing formula as an extension of [Black and Scholes \(1973\)](#) and [Merton \(1973\)](#). However, the main difference is that it includes two underlying assets that follow a bivariate lognormal distribution. The option can also be seen as a European call whose payoff is determined by a stochastic exercise price. Based on this interpretation, [Fischer \(1978\)](#) derives a closed-form formula using the capital asset pricing model, in which the fixed strike price in the Black-Scholes formula fluctuates over time. While [Margrabe \(1978\)](#) and [Fischer \(1978\)](#) give the no-arbitrage pricing of exchange options by portfolio replications, [Stapleton and Subrahmanyam \(1984\)](#) use risk-neutral pricing in a discrete-time setting based on the argument that assets should return a risk-free rate of return in equilibrium regardless of investors' risk appetite.

Due to its analytical tractability and convenience as a direct extension of the Black-Scholes formula, [Margrabe \(1978\)](#)'s approach has become the foundation for the valuation of exotic spread options in the finance literature.¹ The recent model by [Benth and Koekebakker \(2015\)](#) extends the Margrabe formula for a cointegrated forward price model that lowers the price of exchange options compared to a non-cointegrated model irrespective of the slope of forward curves and the sign of correlation parameters. A distinct approach by [Carmona et al. \(2013\)](#) incorporates the structural relationship of power and fuel in a pricing model to capture shifts in the merit order and fuel supply curves. The valuation of power plants shows that their model price diverges (resp. converges) from (to) [Margrabe \(1978\)](#)'s price depending on the marginal fuel and merit order in a market.

In contrast to these models, we introduce the dependence structures of underlying commodities directly into the model parameters of [Margrabe \(1978\)](#) without relying on too many assumptions. Furthermore, we extend the pricing formula for the exchange options whose underlying assets are denominated by different currencies in the context of increased cross-border transactions of energy commodities facilitated by market integration.

¹Those include the American exchange option of [Carr \(1995\)](#), the compound exchange option of [Carr \(1988\)](#), and the barrier exchange option of [Haug and Haug \(2002\)](#). Note that a compound exchange option gives the option buyer the right to buy an exchange option at T_1 which gives the payoff (5.1) at T_2 , ($t < T_1 < T_2$).

When a spread depends on two assets quoted in different currencies, it requires converting one asset's value into the other asset's currency unit to make the payoff under a unified measure. While the currency adjustment is vital for that commodity spreads, no other literature seems to have addressed the impact of the foreign exchange (hereafter the 'FX') component on the option pricing formula.² For an option written on a FX rate with a fixed strike, the pricing formula is developed by [Garman and Kohlhagen \(1983\)](#). Later, [Reiner \(1992\)](#) incorporates various FX structures, so-called the *quanto mechanics*, in the [Black and Scholes \(1973\)](#) formula. However, his list of quanto options does not consider the case when the strike price is stochastic, i.e. an exchange option on the spread of foreign and domestic assets.

Section 5.3 in this chapter introduces two payoff structures for options that enable the holder to exchange one domestic asset for another foreign asset at expiry. Therefore, our study contributes to the literature by combining [Garman and Kohlhagen \(1983\)](#) and [Margrabe \(1978\)](#) and adding two more variants of quanto mechanics to [Reiner \(1992\)](#). The derivation of closed-form formulae is centred on the change of numéraire, highlighting the model parameters associated with currency adjustment terms; it follows the risk-neutral pricing introduced in [Murakami \(2015\)](#).

5.2 Adopting Common Factor Structures in Margrabe (1978)

5.2.1 A Brief Sketch of the Margrabe Formula

Assume the Black-Scholes economy where the market is frictionless, complete, and free of arbitrage. Let $(\Omega, \mathfrak{F}, \mathbb{P})$ be a probability space. Denote $V(T)$ for an $\mathfrak{F}(T)$ -measurable random variable that represents the value of an exchange option at expiry T , where $\mathfrak{F}(t)$ is a collection of sigma algebras indexed by time. Eq. (5.1) gives the terminal payoff of the option with two $\mathfrak{F}(t)$ -measurable adapted stochastic processes $\{S_i(t)\}_{0 \leq t \leq T}$, $i = 1, 2$.

In [Margrabe \(1978\)](#), the dynamics of the underlying assets are described by two correlated geometric Brownian motions that are jointly lognormally distributed:

$$\begin{aligned}\frac{dS_1(t)}{S_1(t)} &= \mu_1 dt + \sigma_1 dW_1^{\mathbb{P}}(t) \\ \frac{dS_2(t)}{S_2(t)} &= \mu_2 dt + \sigma_2 dW_2^{\mathbb{P}}(t) \\ dW_1^{\mathbb{P}}(t)dW_2^{\mathbb{P}}(t) &= \rho dt ,\end{aligned}$$

²One exception is treating the FX rate as a fixed conversion factor such as q_1 and q_2 in Eq. (5.13) mentioned in [Benth and Koekebakker \(2015\)](#).

where ρ is the constant correlation between the $\mathfrak{S}(t)$ -adapted $\mathbb{P}$ -Brownian motions. The mean rates of return μ_i and volatilities σ_i are all constants for $i = 1, 2$.

The risk-neutral pricing of exchange options assumes the existence of an equivalent probability measure under which the discounted asset prices are martingales. When we choose S_2 as the numéraire, the value of an exchange option at t ($0 \leq t \leq T$) can be expressed as

$$\begin{aligned} \frac{V(t)}{S_2(t)} &= \mathbb{E}^{\mathbb{Q}_{S_2}} \left[\frac{V(T)}{S_2(T)} \middle| \mathfrak{S}(t) \right] \\ &= \mathbb{E}^{\mathbb{Q}_{S_2}} \left[\frac{(S_1(T) - S_2(T))^+}{S_2(T)} \middle| \mathfrak{S}(t) \right] \end{aligned} \quad (5.2)$$

$$V(t) = S_2(t) \mathbb{E}^{\mathbb{Q}_{S_2}} \left[\left(\frac{S_1(T)}{S_2(T)} - 1 \right)^+ \middle| \mathfrak{S}(t) \right], \quad (5.3)$$

where $\mathbb{E}^{\mathbb{Q}_{S_2}}[\cdot]$ denotes the expected payoff under the probability measure $\mathbb{Q}_{S_2}$. Eq. (5.3) is in fact the price of a European call written on the ratio of two tradable assets with a unit strike, therefore there exists an analytical solution similar to the [Black and Scholes \(1973\)](#) formula.

Following standard arguments, the stochastic differential equation of the underlying process is given by the Itô-Doeblin formula

$$d\left(\frac{S_1(t)}{S_2(t)}\right) = \frac{S_1(t)}{S_2(t)} \left((\mu_1 - \mu_2 - \rho_{12}\sigma_1\sigma_2 + \sigma_2^2)dt + \sigma_1 dW_1^{\mathbb{P}}(t) - \sigma_2 dW_2^{\mathbb{P}}(t) \right). \quad (5.4)$$

The mean and variance of the linear combinations of Brownian motions are

$$\begin{aligned} \mathbb{E}^{\mathbb{P}}[\sigma_1 W_1^{\mathbb{P}}(t) - \sigma_2 W_2^{\mathbb{P}}(t)] &= 0 \\ \text{Var}[\sigma_1 W_1^{\mathbb{P}}(t) - \sigma_2 W_2^{\mathbb{P}}(t)] &= \sigma_1^2 \text{Var}[W_1^{\mathbb{P}}(t)] - 2 \text{Cov}[W_1^{\mathbb{P}}(t) W_2^{\mathbb{P}}(t)] + \sigma_2^2 \text{Var}[W_2^{\mathbb{P}}(t)] \\ &= (\sigma_1^2 - 2\rho_{12}\sigma_1\sigma_2 + \sigma_2^2) t, \end{aligned}$$

from which one can define a new Brownian motion:³

$$W^{\mathbb{P}}(t) := \left(\frac{\sigma_1 W_1^{\mathbb{P}}(t) - \sigma_2 W_2^{\mathbb{P}}(t)}{\sqrt{\sigma_1^2 - 2\rho_{12}\sigma_1\sigma_2 + \sigma_2^2}} \right) \quad (5.5)$$

Thus, Eq. (5.4) can be rewritten as

$$d\left(\frac{S_1(t)}{S_2(t)}\right) = \nu \left(\frac{S_1(t)}{S_2(t)} \right) \left(\Theta dt + dW^{\mathbb{P}}(t) \right), \quad (5.6)$$

³The Lévy's theorem states that a continuous martingale with expected value zero and quadratic variation t is a Brownian motion. $W^{\mathbb{P}}$ satisfies these properties.

where

$$\begin{aligned}\nu &:= \sqrt{\sigma_1^2 - 2\rho_{12}\sigma_1\sigma_2 + \sigma_2^2} \\ \Theta &:= \frac{\mu_1 - \mu_2 - \rho_{12}\sigma_1\sigma_2 + \sigma_2^2}{\nu}.\end{aligned}$$

These results are used to find the new probability measure $\mathbb{Q}_{S_2}$ and the associated Brownian motion $W^{\mathbb{Q}_{S_2}}(t) := \Theta t + W^{\mathbb{P}}(t)$ by the [Girsanov \(1960\)](#) theorem with the Radon-Nikodim derivative

$$Z = \frac{d\mathbb{Q}_{S_2}}{d\mathbb{P}} = e^{-\frac{1}{2}\Theta^2 T - \Theta W^{\mathbb{P}}(T)}. \quad (5.7)$$

Therefore, the underlying process is a martingale under the $\mathbb{Q}_{S_2}$ -measure

$$d\left(\frac{S_1(t)}{S_2(t)}\right) = \nu \left(\frac{S_1(t)}{S_2(t)}\right) dW^{\mathbb{Q}_{S_2}}(t), \quad (5.8)$$

leading to the solution

$$\left(\frac{S_1(T)}{S_2(T)}\right) = \left(\frac{S_1(t)}{S_2(t)}\right) e^{-\frac{1}{2}\nu^2(T-t) + \nu(W^{\mathbb{Q}_{S_2}}(T) - W^{\mathbb{Q}_{S_2}}(t))}. \quad (5.9)$$

By substituting the right-hand side of Eq. (5.9) into Eq. (5.3), we can analytically solve the expected payoff of the exchange option similarly to solving the expected payoff in the derivation of the Black-Scholes formula for European options. The [Margrabe \(1978\)](#) price of an exchange option is

$$V(t) = S_1(t)\Phi(d_1) - S_2(t)\Phi(d_2) \quad (5.10)$$

$$\begin{aligned}d_1 &= \frac{\ln\left(\frac{S_1(t)}{S_2(t)}\right) + \frac{1}{2}\nu^2(T-t)}{\nu\sqrt{T-t}} \\ d_2 &= d_1 - \nu\sqrt{T-t},\end{aligned} \quad (5.11)$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution function. As a direct result of S_2 being the numéraire, the interest rate parameter r does not appear in the formula.⁴

For those assets paying continuous dividend or convenience yields, d_1 changes to

$$d_1 = \frac{\ln\left(\frac{q_1 S_1(t)}{q_2 S_2(t)}\right) + (\delta_2 - \delta_1 + \frac{1}{2}\nu^2)(T-t)}{\nu\sqrt{T-t}} \quad (5.12)$$

and Eq. (5.10) extends to

$$V(t) = e^{-\delta_1(T-t)} q_1 S_1(t) \Phi(d_1) - e^{-\delta_2(T-t)} q_2 S_2(t) \Phi(d_2), \quad (5.13)$$

⁴We remark that for the valuation of commodity exchange options, it is questionable to take the spot price S_2 as the numéraire since they are often non-transparent or unavailable, and the short-selling of assets is not possible. Hence, it is preferable to work with forwards and assume convergence of forward to spot via $F(T, T) = S(T)$.

where δ_i are the carry costs and the constants q_i are the conversion factors to adjust for differences in units and fuel efficiencies for $i = 1, 2$.

On the other hand, the prices of forwards or futures are inclusive of these carry costs.⁵ The price of an exchange option written on two forward contracts $F_1(t, T_1)$ and $F_2(t, T_2)$ with option expiry T ($t \leq T \leq \min(T_1, T_2)$) is therefore,

$$V(t) = e^{-r(T-t)} [q_1 F_1(t, T_1) \Phi(d_1) - q_2 F_2(t, T_2) \Phi(d_2)] \quad (5.14)$$

$$\begin{aligned} d1 &= \frac{\ln\left(\frac{q_1 F_1(t, T_1)}{q_2 F_2(t, T_2)}\right) + \frac{1}{2} \tilde{\nu}^2 (T - t)}{\tilde{\nu} \sqrt{T - t}} \\ d2 &= d1 - \tilde{\nu} \sqrt{T - t}, \end{aligned} \quad (5.15)$$

where $\tilde{\nu}^2(T - t)$ is the variance of $\ln\left(\frac{F_1(t, T_1)}{F_2(t, T_2)}\right)$ from t to T . This formula can be regarded as a variation of the Black (1976) formula for a European call option on forward contracts with a stochastic strike price.

5.2.2 Factor Volatilities in the Margrabe Formula

The generalisation of the variance term in Eq. (5.15) allows the pricing formula to flexibly adapt to the multi-factor dependence structure of the underlying processes. Consider Eqs. (4.1) – (4.2), where the i -th contract in the k -th forward curve is described by the sum of orthogonal factor volatilities σ_{kj} . From those expressions, one can think of three possible commodity spreads: (a) a calendar spread between two distinct forward contracts in the same forward curve (identical k , distinct i), (b) a calendar spread of cross-market forward contracts (both k and i are distinct), and (c) a cross-market spread of forwards with identical maturity dates (distinct k , identical i).

In order to express the three types of spreads flexibly (i.e. any market and any maturity), we modify the expression slightly to consider two arbitrary forward contracts for $k = 1, 2$

$$\frac{dF_k(t, T_l)}{F_k(t, T_l)} = \sum_{j=1}^n \sigma_{kj}(t, T_l) dZ_{kj}^{\mathbb{Q}}(t),$$

where T_l ($\geq t$) denotes forward maturities for the k^{th} forward curve. We assume the forward contracts are jointly lognormally distributed; however, any assumptions are yet imposed on the cross-market factor dependence structures.

⁵Hereafter, we treat the terminologies for forward and futures, as well as the associated pricing measures interchangeably to simplify our discussion.

Denote $\tilde{\mathbb{Q}}$ for an equivalent probability measure associated with the numéraire $\tilde{F}(t) := B(t)F_2(t, T_2)$,⁶ where $B(t)$ is the bank account such that $B(0) = 1$ and $B(T) = e^{rT}$. The value of an exchange option written on two forward contracts at expiry is then

$$\begin{aligned} \frac{V(t)}{B(t)F_2(t, T_2)} &= \mathbb{E}^{\tilde{\mathbb{Q}}} \left[\frac{V(T)}{B(T)F_2(T, T_2)} \middle| \mathfrak{F}(t) \right] \\ &= \mathbb{E}^{\tilde{\mathbb{Q}}} \left[\frac{(F_1(T, T_1) - F_2(T, T_2))^+}{B(T)F_2(T, T_2)} \middle| \mathfrak{F}(t) \right]. \end{aligned} \quad (5.16)$$

This pricing formula implies that we are effectively changing the probability from $\mathbb{P}$ to $\tilde{\mathbb{Q}}$ with the Radon-Nikodim derivative

$$\tilde{Z} = \frac{d\tilde{\mathbb{Q}}}{d\mathbb{P}} = \frac{d\tilde{\mathbb{Q}}}{d\mathbb{Q}} \cdot \frac{d\mathbb{Q}}{d\mathbb{P}}.$$

Therefore,

$$V(t) = e^{-r(T-t)} F_2(t, T_2) \mathbb{E}^{\tilde{\mathbb{Q}}} \left[\left(\frac{F_1(T, T_1)}{F_2(T, T_2)} - 1 \right)^+ \middle| \mathfrak{F}(t) \right]. \quad (5.17)$$

Solving Eq. (5.17) is just analogous to the derivation for the exchange option on non-dividend paying spot contracts. It requires the modelling of the distribution of $\frac{F_1(T, T_1)}{F_2(T, T_2)}$, or more specifically, the variance term that enters into the closed-form formula (5.15).

The application of the Itô-Doeblin formula to $\ln F_k(t, T_l)$ gives the marginal distribution of the forward contracts

$$\ln F_k(T, T_l) = \ln F_k(t, T_l) - \frac{1}{2} \int_t^T \sum_{j=1}^n \sigma_{kj}^2(u, T_l) du + \int_t^T \sum_{j=1}^n \sigma_{kj}(u, T_l) dZ_{kj}^{\mathbb{Q}}(u) \quad (5.18)$$

with the conditional mean and variance given by

$$\begin{aligned} \mu_k &= \ln F_k(t, T_l) - \frac{1}{2} \int_t^T \sum_{j=1}^n \sigma_{kj}^2(u, T_l) du \\ \sigma_k^2 &= \int_t^T \sum_{j=1}^n \sigma_{kj}^2(u, T_l) du, \end{aligned} \quad (5.19)$$

where σ_{kj} denotes the volatility of the j -th orthogonal factor. Since the distribution of

⁶This numéraire is a tradable asset, as the quantity $\tilde{F}(t) := B(t)F(t, T)$ can be replicated by a portfolio of investments in a bank account and a forward contract. To see the point, we write the time- t value of the forward contract $V(t)$ using the risk-neutral pricing formula $V(t) = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[S(T) - F(0, T)] = e^{-rT}(e^{rt}F(t, T) - e^{-rT}(e^{rt}F(0, T)))$, where $F(0, T)$ is the initial price (strike) of the forward contract. When $B(t) = e^{rt}$, we can rearrange and write the formula as $\tilde{F}(t) = e^{rT}V(t) + F(0, T)B(t)$. Hence, the replicating portfolio for the numéraire consists of e^{rT} amount of forward contracts and $F(0, T)$ amount of bank account both in long positions.

$\ln F_k(T, T_l)$ is jointly normal, the ratio is also normally distributed as

$$\ln \left(\frac{F_1(T, T_1)}{F_2(T, T_2)} \right) \sim \mathcal{N}(\tilde{\mu}, \tilde{\nu}), \quad (5.20)$$

where $\tilde{\mu} = \mathbb{E} \left[\ln \left(\frac{F_1(T, T_1)}{F_2(T, T_2)} \right) \right]$ and $\tilde{\nu} = \text{Var} \left[\ln \left(\frac{F_1(T, T_1)}{F_2(T, T_2)} \right) \right]$. Therefore, one can obtain the variance from the general result $\text{Var}[X - Y] = \text{Var}[X] + \text{Var}[Y] - 2 \text{Cov}[X, Y]$

$$\begin{aligned} \tilde{\nu}^2(t, T, T_1, T_2) = & \int_t^T \sum_{i=1}^n \sigma_{1i}^2(u, T_1) du + \int_t^T \sum_{j=1}^n \sigma_{2j}^2(u, T_2) du \\ & - 2 \int_t^T \left(\sum_{i=1}^n \sum_{j=1}^n \rho_{ij} \sigma_{1i}(u, T_1) \sigma_{2j}(u, T_2) \right) du, \quad (5.21) \end{aligned}$$

where i and j count the orthogonal factors for the first and second forward contracts.

The third term in Eq. (5.21) depends on the model and type of commodity spread. For instance, for a calendar exchange option on a single commodity, it represents the volatilities of orthogonal factors for the same forward curve ($k = 1$). When $i = j$ in the summations, both the PC and PPC models give the (i, j) -th factor correlations one and zero otherwise. This simplifies Eq. (5.21) into

$$\tilde{\nu}^2(t, T, T_1, T_2) = \sum_{j=1}^n \int_t^T (\sigma_{1j}(u, T_1) - \sigma_{2j}(u, T_2))^2 du. \quad (5.22)$$

On the other hand, such a simplification cannot be applied to a cross-commodity exchange option when $k \neq 1$. In the PC model, the whole covariance term remains since the cross-factor correlations are not necessarily zero when $i \neq j$. In the case of the PPC model, it imposes the model correlation structure as in Eq. (4.13). Therefore, not all but some of the cross-factor covariances cancel out in the third term of Eq. (5.21).

5.3 Exchange Options on Foreign to Domestic Spreads

5.3.1 Quanto Exchange Options

Let us first consider the terminal payoff of an option

$$V(T) = (\bar{X} S_1(T) - S_2(T))^+, \quad (5.23)$$

where S_1 denotes an asset quoted in a foreign currency and S_2 denotes an asset quoted in a domestic currency. At option expiry, the value of S_1 is converted into the domestic

currency units by a predetermined conversion rate $\bar{X}$, so that the payoff is made in the domestic currency irrespective of the market FX rate. Therefore, it enables the option holders to manage the FX risk effectively at the inception of trades. We shall call this type of options as *quanto exchange options* naming after other financial products with similar treatments for FX, including equity-linked FX options and quanto interest rate swaps. Note that our quanto exchange option is distinctive from an energy quanto option that refers to a volumetric option for energy commodities, where the payoff of an energy option is compounded by another option payoff typically associated with a temperature index representing demand on energy consumption. See [Benth and Koekebakker \(2015\)](#) for the analytical pricing of energy quanto options written on HDD⁷ and natural gas futures.

The risk-adjusted pricing formula of a quanto exchange option is⁸

$$\frac{V(t)}{S_2(t)} = \mathbb{E}^{\mathbb{Q}_{S_2}} \left[\left(\frac{\bar{X} S_1(T)}{S_2(T)} - 1 \right)^+ \middle| \mathfrak{F}(t) \right]. \quad (5.24)$$

We derive the closed-form expression without considering convenience yields and other conversion factors for simplicity. The rest of the model assumptions remain unchanged from standard exchange options. Solving the expectation in Eq. (5.24) requires the knowledge of the two underlying processes under the same probability measure, as they are quoted in different currency units. The two possible choices of numéraire pairs are $(B_f, \mathbb{Q}_f)$ and $(B_d, \mathbb{Q}_d)$ associated with S_1 and S_2 , respectively, where B_f and B_d are the foreign and domestic bank accounts. We choose to use the latter measure in the following derivation.

Let the *market* FX rate⁹ be described by a geometric Brownian motion

$$\frac{dX(t)}{X(t)} = \mu_X dt + \sigma_X dW_X^{\mathbb{P}}(t), \quad (5.25)$$

which is correlated with the domestic and foreign assets by ρ_{1X} and ρ_{2X} . We utilise the zero-drift conditions of the two martingale processes $\frac{S_1(t)X(t)}{B_d(t)}$ and $\frac{B_f(t)X(t)}{B_d(t)}$ to find the foreign asset dynamics under $\mathbb{Q}_d$. This leads to

$$\frac{dS_1(t)}{S_1(t)} = (r_f - \rho_{1X}\sigma_1\sigma_X)dt + \sigma_1 dW_1^{\mathbb{Q}_d}(t) \quad (5.26)$$

$$= (r_d - \alpha)dt + \sigma_1 dW_1^{\mathbb{Q}_d}(t), \quad (5.27)$$

⁷HDD is the acronym for Heating Degree Days. HDD is used as an index in the weather derivatives market to account for the number of days and degrees that exceeded a specified temperature during a contract period.

⁸See Footnote 4.

⁹Note that while this market FX rate is necessary for the change of numéraire, it is irrelevant to the computation of the expected payoff in Eq. (5.23), for which the FX rate is fixed by contract.

where $\alpha = r_d - r_f + \rho_{1X}\sigma_1\sigma_X$.

Thus, we can find the pricing formula of the foreign asset under the $\mathbb{Q}_{S_2}$ -measure analogously to the risk-adjusted pricing formula of an asset paying a continuous carry cost:

$$\frac{S_1(t)}{S_2(t)} = \mathbb{E}^{\mathbb{Q}_{S_2}} \left[\frac{S_1(T)e^{\alpha(T-t)}}{S_2(T)} \middle| \mathfrak{F}(t) \right] \quad (5.28)$$

$$\frac{S_1(t)e^{-\alpha(T-t)}}{S_2(t)} = \mathbb{E}^{\mathbb{Q}_{S_2}} \left[\frac{S_1(T)}{S_2(T)} \middle| \mathfrak{F}(t) \right] \quad (5.29)$$

Applying the Itô-Doeblin formula to the left-hand side of Eq. (5.29) and following the same derivation as for a standard exchange option, we obtain the closed-form solution for the quanto exchange option

$$V(t) = e^{-\alpha(T-t)} \bar{X} S_1(t) \Phi(d_1) - S_2(t) \Phi(d_2), \quad (5.30)$$

where

$$\begin{aligned} d_1 &= \frac{\ln \left(\frac{\bar{X} S_1(t)}{S_2(t)} \right) + (-\alpha + \frac{1}{2}\nu^2)(T-t)}{\nu\sqrt{T-t}} \\ d_2 &= d_1 - \nu\sqrt{T-t} \\ \nu^2 &= \sigma_1^2 + \sigma_2^2 - 2\rho_{12}\sigma_1\sigma_2. \end{aligned} \quad (5.31)$$

This is a special case of Eqs. (5.12) – (5.13), where $S_1 = \bar{X} S_1$, $\delta_1 = \alpha$, and $\delta_2 = 0$.

Our next aim is to derive the closed-form formula of quanto exchange options on forwards. To do so, we first find the quanto forward price $\tilde{F}_1(t, T_1) = \mathbb{E}^{\mathbb{Q}_d}[\bar{X} S_1(T_1)]$ from Eq. (5.27), where the probability measure has been transformed from $\mathbb{Q}_f$ to $\mathbb{Q}_d$ for S_1 . Since the stochastic differential equation can be treated similarly to the dynamics of a non-quanto asset paying a continuous dividend yield, we can solve Eq. (5.27) with the Itô-Doeblin formula:

$$S_1(T_1) = S_1(t) e^{(r_d - \alpha - \frac{1}{2}\sigma_1^2)(T_1-t) + \sigma_1(W_1^{\mathbb{Q}_d}(T_1) - W_1^{\mathbb{Q}_d}(t))} \quad (5.32)$$

By definition, $\tilde{F}_1(t, T_1) = \mathbb{E}^{\mathbb{Q}_d}[\bar{X} S_1(T_1)]$ and $\mathbb{E}[e^Y] = e^{\mathbb{E}[Y] + \frac{1}{2}\text{Var}[Y]}$ for $Y \sim \mathcal{N}(\mu_Y, \sigma_Y^2)$. Thus, the quanto forward price is

$$\begin{aligned} \tilde{F}_1(t, T_1) &= \bar{X} S_1(t) e^{(r_d - \alpha)(T_1-t)} \\ &= \bar{X} S_1(t) e^{(r_f - \rho_{1X}\sigma_1\sigma_X)(T_1-t)} \\ &= \bar{X} F_1(t, T_1) e^{-\rho_{1X}\sigma_1\sigma_X(T_1-t)}, \end{aligned} \quad (5.33)$$

where $F_1(t, T_1) = S_1(t)e^{r_f(T_1-t)}$ represents the standard (non-quanto) spot – forward cost-of-carry relationship for a single currency case. Note that the forward price $F_1(t, T_1)$ is given in foreign currency units, whereas $\tilde{F}_1(t, T_1)$ is given in domestic currency units. As seen in the third line of Eq. (5.33), the covariance of the foreign asset and the market FX rate results from the numéraire change applied earlier to the underlying foreign asset.

Eq. (5.33) facilitates the pricing of other quanto derivatives. As for quanto exchange options on forwards, we can utilise the cost-of-carry like relationship (5.33) with Eqs. (5.30) – (5.31) to obtain the closed-form formula, analogously transforming from the Black and Scholes (1973) formula to Black (1976) formula, and vice versa. The quanto exchange option formula on forwards with option expiry $T (< T_1)$ is therefore,

$$V(t) = e^{-r_d(T-t)} [\bar{X}F_1(t, T_1)e^{-\rho_{1X}\sigma_1\sigma_X(T_1-t)}\Phi(d_1) - F_2(t, T_2)\Phi(d_2)] , \quad (5.34)$$

where

$$d_1 = \frac{\ln\left(\frac{\bar{X}F_1(t, T_1)}{F_2(t, T_2)}\right) - \rho_{1X}\sigma_1\sigma_X(T_1 - t) + \frac{1}{2}\tilde{\nu}^2(T - t)}{\tilde{\nu}\sqrt{T - t}} \quad (5.35)$$

$$d_2 = d_1 - \tilde{\nu}\sqrt{T - t} .$$

In the formula, $\tilde{\nu}^2(T - t)$ is the variance of $\ln\left(\frac{F_1(T, T_1)}{F_2(T, T_2)}\right)$ from t to T , to which we can plug in a model variance introduced in Section 5.2.2. We remark that the quanto adjustment is only required for the drift term. Hence, the variance term remains unchanged from that of non-quanto exchange options in Eq. (5.15).

5.3.2 Cross-currency Exchange Options

The payoff structure of the second option type, called *cross-currency exchange options*, is similar to that of a quanto exchange option except that it lets the FX rate fluctuate until the expiry of the option. At expiry, the value of a foreign asset is converted into a domestic currency unit by the spot FX rate observed in the market.¹⁰

The payoff and the risk-adjusted pricing formula of this option are

$$V(T) = (X(T)S_1(T) - S_2(T))^+ \quad (5.36)$$

$$\frac{V(t)}{S_2(t)} = \mathbb{E}^{\mathbb{Q}_{S_2}} \left[\left(\frac{X(T)S_1(T)}{S_2(T)} - 1 \right)^+ \middle| \mathfrak{F}(t) \right] . \quad (5.37)$$

As before, we can obtain the closed-form solution of the exchange option by finding the

¹⁰The delivery dates of spot FX contracts vary between $t + 0$ day and $t + 3$ days in practice depending on currency pairs. For simplicity, we assume that the delivery date of a spot FX contract is $t + 0$ day in our discussion.

distribution of $\frac{X(T)S_1(T)}{S_2(T)}$ under the $\mathbb{Q}_d$ -measure.

Recall that

$$\begin{aligned}\frac{dS_1(t)}{S_1(t)} &= (r_f - \rho_{1X}\sigma_1\sigma_X)dt + \sigma_1 dW_1^{\mathbb{Q}_d}(t) \\ \frac{dS_2(t)}{S_2(t)} &= r_d dt + \sigma_2 dW_2^{\mathbb{Q}_d}(t),\end{aligned}$$

and the dynamics of the FX process are given by

$$\frac{dX(t)}{X(t)} = (r_d - r_f)dt + \sigma_X dW_X^{\mathbb{Q}_d}(t),$$

which is the well-known result for the pricing of foreign currency options by [Garman and Kohlhagen \(1983\)](#). The Brownian motions are assumed to be correlated by ρ_{1X} , ρ_{2X} , and ρ_{12} .

By applying the Itô-Doeblin formula to $\frac{X(T)S_1(T)}{S_2(T)}$ we obtain an expression for the price dynamics

$$\begin{aligned}d\left(\frac{X(t)S_1(t)}{S_2(t)}\right) &= \frac{X(t)S_1(t)}{S_2(t)} \left((\sigma_2^2 - \rho_{12}\sigma_1\sigma_2 - \rho_{2X}\sigma_2\sigma_X)dt \right. \\ &\quad \left. + \sigma_X dW_X^{\mathbb{Q}_d}(t) + \sigma_1 dW_1^{\mathbb{Q}_d}(t) - \sigma_2 dW_2^{\mathbb{Q}_d}(t) \right) \quad (5.38)\end{aligned}$$

and the variance of the linear combination of Brownian motions

$$\begin{aligned}\text{Var}[\sigma_X W_X^{\mathbb{Q}_d}(t) + \sigma_1 W_1^{\mathbb{Q}_d}(t) - \sigma_2 W_2^{\mathbb{Q}_d}(t)] &= (\sigma_X^2 + \sigma_1^2 + \sigma_2^2 \\ &\quad + 2\rho_{1X}\sigma_1\sigma_X - 2\rho_{2X}\sigma_2\sigma_X - 2\rho_{12}\sigma_1\sigma_2) t. \quad (5.39)\end{aligned}$$

Therefore, the closed-form solution of the option can be derived as it is done in Section 5.2.1. The drift term disappears after the change of numéraire. The difference is that the variance in the exchange option formula now includes the variance and covariance of the three assets

$$\nu_X^2 := \sigma_X^2 + \sigma_1^2 + \sigma_2^2 + 2\rho_{1X}\sigma_1\sigma_X - 2\rho_{2X}\sigma_2\sigma_X - 2\rho_{12}\sigma_1\sigma_2, \quad (5.40)$$

which replaces the variance of the standard Margrabe formula in Eqs. (5.10) – (5.11), where $S_1 = XS_1$.

As for the cross-currency exchange option written on forward contracts, the terminal payoff is

$$V(T) = (X(T)F_1(T, T_1) - F_2(T, T_2))^+. \quad (5.41)$$

At expiry, the value of a foreign forward contract is converted into a domestic currency unit by the spot FX rate observed in the market. The pricing formula takes the domestic bank account and the forward contract as the numéraire, as is the case for the standard Margrabe formula written on two forward contracts:

$$\frac{V(t)}{B_d(t)F_2(t, T_2)} = \mathbb{E}^{\mathbb{Q}} \left[\frac{(X(T)F_1(T, T_1) - F_2(T, T_2))^+}{B_d(T)F_2(T, T_2)} \middle| \mathfrak{F}(t) \right] \quad (5.42)$$

From previous derivations we learnt that what matters is the variance of $\ln \left(\frac{X(T)F_1(T, T_1)}{F_2(T, T_2)} \right)$ that enters into the Margrabe formula. The closed-form formula of the cross-currency exchange option is therefore,

$$V(t) = e^{-r_d(T-t)} [X(t)F_1(t, T_1)\Phi(d_1) - F_2(t, T_2)\Phi(d_2)] \quad (5.43)$$

$$d_1 = \frac{\ln \left(\frac{X(t)F_1(t, T_1)}{F_2(t, T_2)} \right) + \frac{1}{2}\tilde{\nu}_X^2(T-t)}{\tilde{\nu}_X\sqrt{T-t}} \quad (5.44)$$

$$d_2 = d_1 - \tilde{\nu}_X\sqrt{T-t},$$

where

$$\tilde{\nu}_X^2 := \tilde{\sigma}_1^2 + \tilde{\sigma}_2^2 - 2\tilde{\rho}_{12}\tilde{\sigma}_1\tilde{\sigma}_2 + \sigma_X^2 + 2\tilde{\rho}_{1X}\tilde{\sigma}_1\tilde{\sigma}_X - 2\tilde{\rho}_{2X}\tilde{\sigma}_2\tilde{\sigma}_X \quad (5.45)$$

and

$$\begin{aligned} \tilde{\sigma}_1 & : \text{Volatility of } F_1(t, T_1) \\ \tilde{\sigma}_2 & : \text{Volatility of } F_2(t, T_2) \\ \sigma_X & : \text{Volatility of } X(t) \\ \tilde{\rho}_{1X} & : \text{Correlation between } X(t) \text{ and } F_1(t, T_1) \\ \tilde{\rho}_{2X} & : \text{Correlation between } X(t) \text{ and } F_2(t, T_2) \\ \tilde{\rho}_{12} & : \text{Correlation between } F_1(t, T_1) \text{ and } F_2(t, T_2). \end{aligned}$$

It is worth noting that the variance of a cross-currency exchange option can be decomposed into

$$\text{Total variance} = \text{Spread variance} + \text{FX add-on} .^{11}$$

One can easily introduce the factor dependence structures, for example Eq. (5.21), for the spread variance. The add-on term consists of the variance and covariance terms associated with the spot FX rate.

Due to the additional uncertainty brought in by the FX components, the risk manage-

¹¹The two terms on the right-hand side correspond to the first three and second three terms in Eq. (5.45), respectively.

ment of option vega requires extra caution. The vega of a standard exchange option with respect to the volatility of the first underlying asset is

$$\frac{\partial V}{\partial \sigma_1} = \frac{\partial V}{\partial \nu} \cdot \frac{\partial \nu}{\partial \sigma_1} . \quad (5.46)$$

Geman (2009) discusses that the first quantity in the right-hand-side is the ordinary option vega that is always positive; however, the sign of the second quantity

$$\frac{\partial \nu}{\partial \sigma_1} = \frac{1}{\nu}(\sigma_1 - \rho_{12}\sigma_2) \quad (5.47)$$

can be positive or negative depending on the relative size of σ_1 and σ_2 , as well as the sign and magnitude of ρ_{12} . Thus, the sign of the vega is not always positive.

As for a cross-currency exchange option, Eq. (5.47) changes to

$$\frac{\partial \nu_X}{\partial \sigma_1} = \frac{1}{\nu_X}(\sigma_1 + \rho_{1X}\sigma_X - \rho_{12}\sigma_2) . \quad (5.48)$$

The sign of the option vega depends not only on the parameters associated with S_1^f and S_2^d but also on the parameters of spot FX rate, adding extra complexity to the risk factor. The same argument also applies to the vega of the cross-currency exchange option with respect to the second underlying asset.

5.4 Pricing

Table 5.1 updates Table 4.6 to display the type of commodity spreads in the first column. The type ‘Spark’ stands for the *spark spread* that is the general terminology for the power-gas spread often regarded as an approximation to the gross profit margin of a gas-fired plant. It can be interpreted as a per unit of power output given an amount of input fuel after adjusting for the fuel efficiency of a power generation plant.

We use the following formula to calculate a spark spread for the options, with the fuel efficiency of 50% and the conversion factor of 1 therm = 0.0293071 MWh:

$$\text{Spark spread (Ccy}^*/\text{MWh)} = \text{Power price (Ccy}^*/\text{MWh)} - \frac{\text{Gas Price (Ccy}^*/\text{MWh)}}{\text{Fuel efficiency}(\%)}$$

Ccy* refers to the payoff currency in Table 5.1. The fuel efficiency of 50.0% means that a generator spends 6,824 therms of gas to produce 1 kilowatt (= 3,412 therms) of electricity.

In Table 5.1, ‘xCcy’ represents the spreads that include a foreign asset as one of the

Table 5.1: Five types of commodity spreads.

Type	Asset 1	Asset 2	Ccy 1	Ccy 2	Unit 1	Unit 2	Ccy*	Unit*
Spark	UK power	UK gas	GBp	GBP	Therm	MWh	GBP	MWh
Spark	Dutch power	Dutch gas	EUR	EUR	MWh	MWh	EUR	MWh
xCcy	UK gas	Dutch gas	GBp	EUR	Therm	MWh	EUR	MWh
xCcy	UK power	Dutch power	GBP	EUR	MWh	MWh	EUR	MWh
xCcy/Spark	UK power	Dutch gas	GBP	EUR	MWh	MWh	EUR	MWh

The table shows the units of trades and currencies for each underlying asset. *Ccy** (currency) and *Unit** represent the adjusted units and currencies used to calculate the spread payoff. For UK assets, GBp and GBP refer to pence sterling and pound sterling.

underlying assets. We calculate quanto exchange options and cross-currency exchange options for these product types by the extended Margrabe formula presented in the previous sections. In the following analysis, we regard the UK forward contracts as foreign assets and convert the option payoff to EUR using the market quote for GBP/EUR.

The option prices are calculated for these five pairs of commodities using the same data set and fitted parameter values as the previous chapters,¹² in which the local covariance fitting estimates the idiosyncratic factor for the PPC model. In addition, some of the underlying assumptions are simplified, such as zero interest rates and flat forward curves for the initial prices: 59.35 GBp/Therm, 21.94 EUR/MWh, 57.29 EUR/MWh, and 51.39 EUR/MWh for UK gas, Dutch gas, UK power and Dutch power, respectively. All options expire five days before the maturities of the underlying forward contracts same as the market convention. The results are reported for forward contracts with maturities $T = (30, 90, 180, 360, 540)$ days due to limitations in space. Appendix A reports the full results, including the 60, 120, 150, 210, and 270-day contracts.

5.4.1 Results for Spark Spreads

The economic relationship of gas and power suggests that these two commodities are closely linked. Our guess is that the PPC model would better capture the comovement of these commodities, therefore tightening the spread more than the PC model and lowering the option prices. Fig. 5.1 depicts the model prices, and Table 5.2 reports the differences in the option prices to indicate how much the PPC price is below or above the PC benchmark. The standard deviations are separately shown in Table 5.3 by maturity. The results suggest that the gap between the two model prices widens monotonically as the contract maturities increase. However, for the UK spark spread, the PPC price is at most at an 18 pence

¹²See Section 3.4.1 and Section 4.3.




Figure 5.1: The option prices of UK and Dutch spark spreads.

The spreads are calculated with the fuel efficiency rate of 50%. The options expire five days before the maturities are shown on the x-axis. The y-axis shows the adjusted units. The unit of UK gas is converted to MWh using the conversion factor: 1 therm = 0.0293071 MWh.

Table 5.2: Option prices of spark spreads.

Contracts	30d	90d	180d	360d	540d
Panel A: Up-Ug (GBP/MWh)					
PC	16.814	16.866	16.903	16.949	16.998
PPC	16.791	16.793	16.798	16.807	16.819
diff	-0.023	-0.073	-0.105	-0.142	-0.179
diff %	-0.14	-0.43	-0.62	-0.84	-1.05
Panel B: Dp-Dg (EUR/MWh)					
PC	7.753	8.021	8.262	8.675	9.068
PPC	7.518	7.614	7.794	8.175	8.552
diff	-0.235	-0.408	-0.468	-0.499	-0.517
diff %	-3.03	-5.08	-5.67	-5.76	-5.70

The relative difference is given by the formula: $\text{diff \%} = 100 \times \frac{(\text{PPC price} - \text{PC price})}{\text{PC price}}$.

Table 5.3: Standard deviations of spark spreads at option expiration.

Contracts	30d	90d	180d	360d	540d
Panel A: Up-Ug					
PC	0.1478	0.1750	0.1871	0.1993	0.2100
PPC	0.0832	0.1155	0.1293	0.1421	0.1522
Panel B: Dp-Dg					
PC	0.1189	0.1477	0.1687	0.2004	0.2278
PPC	0.0637	0.0975	0.1240	0.1615	0.1913

The standard deviations are given by Eq. (5.21). The options are assumed to expire five days before the maturities of the underlying forward contracts.

(1.05%) discount to the PC price for the 540-day contract: the relative difference is rather small.

The above result is caused by the moneyness of the spark spread options. Generally speaking, moneyness is the price ratio of the underlying asset to its strike value, indicating by how much the option is currently in the money or out of the money. The strike price of an exchange option is not a fixed constant, yet the same concept applies, and fuel efficiency is the crucial parameter determining the moneyness of spark spread options. Given the initial prices and the fuel efficiency rates, both the UK and Dutch spark spreads happen to be in the money in our example. Therefore, the option values mainly consist of their intrinsic values, and the choice of stochastic models does not greatly influence the pricing results as expected.

We re-examine the model prices of spark spread by moneyness, similarly to computing the price of a European option for various strike prices. Table 5.4 reports the option prices of the 360-day contracts using varying fuel efficiency rates ranging from 30% to 60% by 5% intervals. These efficiency rates represent four typical generating facilities: gas turbine, steam generator, internal combustion, and combined cycle gas turbine (CCGT) in ascending order of fuel efficiency. In Fig. 5.2, the black y-axis and lines represent the option prices of the UK spark spread, and the blue y-axis and the lines are those for the Dutch spark spread.

Above 50% efficiency, the UK spark spread is deep in the money, and the option prices are indistinguishable between the two models. On the other hand, as the fuel efficiency goes down, the price gap widens gradually towards the direction of ‘at the money’ around 35% efficiency. At most, the PPC price is a 1.3 pound discount to the PC price. Likewise, the Dutch spark spread is at the money with the fuel efficiency of around 40%, where

Table 5.4: The option prices of 360-day spark spread by various fuel efficiency rates.

Efficiency	30%	35%	40%	45%	50%	55%	60%
Panel A: Up-Ug (GBP/MWh)							
PC	1.426	4.292	8.411	12.843	16.949	20.515	23.552
PPC	0.543	2.985	7.475	12.423	16.807	20.475	23.541
diff	-0.883	-1.307	-0.936	-0.420	-0.142	-0.041	-0.011
diff %	-61.93	-30.45	-11.13	-3.27	-0.84	-0.20	-0.04
Panel B: Dp-Dg (EUR/MWh)							
PC	0.193	0.960	2.731	5.450	8.675	11.944	14.983
PPC	0.050	0.482	1.962	4.709	8.175	11.683	14.868
diff	-0.142	-0.478	-0.769	-0.741	-0.499	-0.261	-0.115
diff %	-73.88	-49.80	-28.14	-13.60	-5.76	-2.19	-0.76

The prices of UK and Dutch spark spread options that depict Fig. 5.2.




Figure 5.2: The option price of 360-day spark spread by various fuel efficiency rates.

The black y-axis and lines show the price of UK spark spread in GBP/MWh for the PC (solid line) and PPC (dotted line) models. Similarly, the blue lines and the y-axis on the left-hand side show the PC and PPC prices for the Dutch spark spread.

we observe the maximum 77 cents difference between the two model prices. The options become further out of the money with lower efficiency rates and depreciate in value, notably for the PPC model. Table 5.5 updates Table 5.2 using the fuel efficiency that approximately gives at-the-money option prices, in which the average price differentials increase to 1.35 pounds (34.8%) for the UK spark spread and 85 cents (44.7%) for the Dutch spark spread.

Lastly, we note the implication of the PCA and PPC modelling approaches in conjunction with the valuation of power plants. Suppose that a utility company in the Netherlands wants to estimate the gross profit margin of a gas-fired power plant in one year. Suppose further that the company owns a portfolio of power generation plants in Europe, with diversified fuel sources to generate electricity. At the time T ($> t$), if the cost of natural gas exceeds the selling price of power, the company would not switch on the gas-fired plant, and use alternative production facilities or procure electricity in the spot market to meet their obligation to deliver electricity to the consumers. Therefore, the exchange option can mimic the payoff value of the gas-fired power plant with the optionality to switch on or off the generation facility.

For this example, Table 5.5 gives a rough estimate¹³ for the present value of the gas-

¹³For the sake of simplicity, we do not consider other operation and fixed costs in this example, hence setting $K = 0$. In practical applications, those costs need to be included in the valuation, which can be approximated by Kirk (1995), for example. In addition, electricity rates are generally lower at night and higher during peak hours. This example simplifies these features using a "base rate" per hour per day.

Table 5.5: The (at-the-money) option prices of 360-day spark spread.

Contracts	30d	90d	180d	360d	540d
Panel A: Up-Ug (GBP/MWh)					
PC	3.117	3.738	4.015	4.292	4.536
PPC	1.640	2.378	2.694	2.985	3.216
diff	-1.477	-1.360	-1.321	-1.307	-1.321
diff %	-47.38	-36.39	-32.91	-30.45	-29.11
Panel B: Dp-Dg (EUR/MWh)					
PC	1.157	1.697	2.104	2.731	3.283
PPC	0.270	0.780	1.250	1.962	2.551
diff	-0.886	-0.917	-0.855	-0.769	-0.733
diff %	-76.64	-54.05	-40.61	-28.14	-22.32

Re-calculated option prices with 35% and 40% fuel efficiency rates, respectively for the UK and Dutch spark spreads.

fired power plant with 40% fuel efficiency. Define the total notional amount:

$$Q_{TN} = \text{Capacity of plant} \times \text{Days in a contract period} \times \text{Delivery hours per day}$$

Consider a 500 MW generation capacity and 21 weekdays in a month. Using the results from Table 5.4, the total notional amount is 252,000 MWh ($= 500 \text{ MW} \times 21 \text{ days} \times 24 \text{ hours}$) and the present values of the month are

$$PV_{pc} = Q_{TN} \times 2.73 = 687,960 \text{ EUR}$$

$$PV_{ppc} = Q_{TN} \times 1.96 = 493,920 \text{ EUR}.$$

The difference in value (194,040 EUR) is mainly caused by the PPC model's common factor structures. This example illustrates the possibility that ignoring common factors results in overvaluing real assets.

5.4.2 Results for Spreads Including a Foreign Asset

Quanto Exchange Options

Quanto exchange options require three additional input parameters in addition to the standard exchange options: the quanto rate, the volatility of spot FX, and the correlation between the spot FX and the foreign asset. Although the quanto rate is purely contract-based and does not necessarily relate to a market FX rate, the market participants may refer to the forward FX rate in practice. We download the mid price of the spot GBP/EUR




Figure 5.3: Spot FX rate and its rolling correlation with UK assets.

Top: the trajectory of the spot GBP/EUR exchange rate between 1st July 2016 and 29th June 2018. Bottom: the historical estimation of correlation coefficients between the log returns of 30-day fixed time-to-maturity contracts and the spot GBP/EUR foreign exchange rate. Ug, Up, and FX refer to UK gas, UK power, and GBP/EUR spot FX rate respectively. The data covers a two-year period between July 2016 and June 2018. The rolling window starts on 23 June 2017 taking 252 observation points backwards at each day. The legend shows the minimum and maximum correlation estimates in brackets.

FX rate from Bloomberg® to estimate the FX parameters. The quanto rate is set at $\bar{X} = 1.1226$ in reference to the average forward FX rates over the maturity of the contracts on 29th June 2018.¹⁴

The top figure in Fig. 5.3 depicts the price trajectory of the spot FX between 1st July 2016 and 29th June 2018, during which the GBP/EUR stably fluctuates between 1.1 and 1.2, with the historical volatility estimate of $\sigma_X = 8.9\%$. According to our 252-day rolling window analysis, the correlation between the FX rate and the spot UK contracts¹⁵ is less than 20% in magnitude for both UK gas and UK power, as shown in the bottom of Fig. 5.3, where the legend reports the minimum and maximum correlations. Therefore, in this study, we set $\rho_{1X} = -0.2$ for both UK gas and UK power.

Tables 5.6 – 5.7 report the standard deviations and option prices. It shows that the PPC model calculates the option prices on average¹⁶ at 14 cents (11.0%), 68 cents (4.9%),

¹⁴It is the average forward FX (mid) rates with maturities $T = (30, 60, 90, 120, 150, 180, 210, 240, 360, 540)$ days as at the pricing day.

¹⁵We regard the 30-day contract as a proxy for the spot price.

¹⁶Including all other maturity contracts. See Table A.5

Table 5.6: Standard deviations of UK – Dutch spreads at option expiration.

Contracts	30d	90d	180d	360d	540d
Panel C: Ug-Dg					
PC	0.0632	0.0811	0.0932	0.1108	0.1265
PPC	0.0427	0.0645	0.0779	0.0921	0.1031
Panel D: Up-Dp					
PC	0.1715	0.2013	0.2163	0.2419	0.2661
PPC	0.0973	0.1362	0.1532	0.1778	0.1995
Panel E: Up-Dg					
PC	0.1515	0.1792	0.1912	0.2058	0.2195
PPC	0.0829	0.1196	0.1329	0.1456	0.1562

The model standard deviations are given by Eq. (5.21). The options are assumed to expire five days before the maturities of the underlying forward contracts.

Table 5.7: Option prices of quanto exchange options.

Contracts	30d	90d	180d	360d	540d
Panel C: Ug-Dg (EUR/MWh)					
PC	1.051	1.191	1.291	1.438	1.571
PPC	0.906	1.061	1.166	1.281	1.372
diff	-0.144	-0.130	-0.125	-0.157	-0.199
diff %	-13.75	-10.94	-9.69	-10.90	-12.67
Panel D: Up-Dp (EUR/MWh)					
PC	13.382	13.711	13.902	14.263	14.635
PPC	12.960	13.102	13.219	13.442	13.686
diff	-0.421	-0.609	-0.683	-0.821	-0.949
diff %	-3.15	-4.44	-4.91	-5.76	-6.48
Panel E: Up-Dg (EUR/MWh)					
PC	10.162	10.572	10.766	11.017	11.260
PPC	9.534	9.784	9.925	10.081	10.224
diff	-0.628	-0.789	-0.841	-0.936	-1.036
diff %	-6.18	-7.46	-7.81	-8.49	-9.20

The price of quanto option using the quanto factor of 1 GBP=1.1226 and the correlation coefficient of $\rho_{1X} = -0.2$ for UK gas and power. The fuel efficiency is 40% for the UK power – Dutch gas spark spread. The relative difference is given by the formula:

$$\text{diff \%} = 100 \times \frac{(\text{PPC price} - \text{PC price})}{\text{PC price}}.$$

and 84 cents (7.8%) discount to the PC model prices for the UK gas – Dutch gas, UK power – Dutch power, and UK power – Dutch gas, respectively, in which the calculation of the UK power – Dutch gas spark spread is based on a 40% fuel efficiency. For both UK power – Dutch power and UK power – Dutch gas, the difference in model prices widens




Figure 5.4: Quanto options with varying correlation coefficients between the spot FX rate and the UK assets.

The calculation of quanto options uses a hypothetical correlation parameter: $\rho_{1X} = -0.9, 0, 0.9$ replacing the historical estimates of correlation coefficients -0.2 . Other input parameters are unchanged from the previous calculation in Table 5.7

as the maturity of contracts increases. On the other hand, UK gas – Dutch gas does not follow the pattern; the gap between the two model prices narrows for the mid-term options.

Fig. 5.4 illustrates the impact of the quanto adjustment on option prices. In the figure, the correlation coefficient varies between -0.9 and $+0.9$, replacing the original value of -0.2 to calculate the quanto adjustment terms. As we can see in the figure, when the correlation (hence the quanto adjustment term) is negative, the option price is higher than the value of options with zero correlation (no quanto adjustment). Conversely, when the correlation is positive, the inclusion of the quanto adjustment term lowers the option price. Nevertheless, the impact of quanto adjustment on the options is limited since we find negligible differences even between the hypothetically high correlation and no correlation.

Cross-currency Exchange Options

We set $X(t) = 1.1304$ using the last observation in our data set and historically estimate the FX volatility $\sigma_X = 8.9\%$. The fuel efficiency of 40% is used for the calculation of the UK power – Dutch gas spark spread. Contrary to the model assumption of constant correlations, the empirical data suggests a term structure of correlation between the spot FX rate and forward contracts in Table 5.8. We use these time-dependent correlation coefficients to examine the relative impact of the correlation between the spot FX rate

Table 5.8: Correlation coefficients between the spot FX rate and forward contracts.

Commodities	Contracts									
	30	60	90	120	150	180	210	270	360	540
UK Gas (%)	-8.6	-10.2	-12.6	-13.9	-15.5	-15.6	-15.0	-15.8	-17.7	-18.5
Dutch Gas (%)	8.6	8.6	8.6	8.2	8.9	10.4	11.0	11.9	12.2	13.3
UK Power (%)	-4.5	-4.1	-9.3	-8.5	-9.3	-9.9	-10.8	-13.3	-16.6	-15.8
Dutch Power (%)	-0.3	2.7	2.7	5.1	5.9	7.1	9.7	11.2	9.3	8.4

The historical estimation of correlation coefficients between the log-returns of spot GBP/EUR FX rate and the underlying forward contracts for a 2-year estimation window from July 2016 to June 2018.

and the underlying forward contracts and the marginal volatilities on the option prices.

As explained earlier, the introduction of a stochastic FX rate can increase or decrease the overall variance for a spread of commodities. Table 5.9 reveals the decomposition of the total variance for the three commodity pairs that include a foreign asset. The FX add-on is mostly negative, while it is negligibly small and has little impact on short-term contracts. On the other hand, it imposes downward pressures on the spread variance of longer-dated contracts. Therefore, the diminishing total variance lowers the value of exchange options in Fig. 5.5 and Table 5.10, especially when the FX add-on accounts for a large proportion of the total variance as seen in the PPC prices of the UK gas – Dutch gas spread. Consequently, the option prices do not form a concave shape in Fig. 5.5, in contrast to the prices of quanto exchange options in Fig. 5.4.

Table 5.9: The variance of cross-currency spreads at option expiration.

Contracts	30d	90d	180d	360d	540d
Panel C: Ug-Dg					
PC					
Total var	0.0036	0.0052	0.0057	0.0068	0.0081
Spread var	0.0040	0.0066	0.0087	0.0123	0.0160
FX add-on	-0.0004	-0.0014	-0.0030	-0.0055	-0.0079
PPC					
Total var	0.0014	0.0029	0.0032	0.0032	0.0031
Spread var	0.0018	0.0042	0.0061	0.0085	0.0106
FX add-on	-0.0004	-0.0013	-0.0028	-0.0052	-0.0075
Panel D: Up-Dp					
PC					
Total var	0.0296	0.0402	0.0460	0.0551	0.0680
Spread var	0.0294	0.0405	0.0468	0.0585	0.0708
FX add-on	0.0002	-0.0003	-0.0008	-0.0034	-0.0028
PPC					
Total var	0.0097	0.0185	0.0232	0.0293	0.0382
Spread var	0.0095	0.0185	0.0235	0.0316	0.0398
FX add-on	0.0002	-0.0000	-0.0003	-0.0023	-0.0016
Panel E: Up-Dg					
PC					
Total var	0.0227	0.0311	0.0350	0.0376	0.0426
Spread var	0.0230	0.0321	0.0365	0.0424	0.0482
FX add-on	-0.0003	-0.0010	-0.0016	-0.0047	-0.0056
PPC					
Total var	0.0067	0.0135	0.0164	0.0173	0.0197
Spread var	0.0069	0.0143	0.0177	0.0212	0.0244
FX add-on	-0.0002	-0.0008	-0.0012	-0.0039	-0.0047

The table shows the breakdown of spread variance for cross-currency exchange options, where $\text{Total var} = \text{Spread var} + \text{FX add-on}$. In the table, *Total var* is the variance that enters into Eq. (5.44), *Spread var* is the standard variance of an exchange option, and *FX add-on* is the additional components that account for fluctuations in spot FX rates.

Table 5.10: The prices of cross-currency exchange options.

Contracts	30d	90d	180d	360d	540d
Panel C: Ug-Dg (EUR/MWh)					
PC	1.139	1.229	1.256	1.309	1.370
PPC	1.008	1.098	1.121	1.120	1.111
diff	-0.131	-0.130	-0.135	-0.189	-0.259
diff %	-11.54	-10.61	-10.78	-14.42	-18.90
Panel D: Up-Dp (EUR/MWh)					
PC	13.783	14.086	14.256	14.530	14.914
PPC	13.392	13.517	13.621	13.774	14.027
diff	-0.391	-0.568	-0.636	-0.755	-0.887
diff %	-2.84	-4.04	-4.46	-5.20	-5.95
Panel E: Up-Dg (EUR/MWh)					
PC	10.527	10.890	11.056	11.169	11.376
PPC	9.955	10.155	10.267	10.300	10.402
diff	-0.573	-0.736	-0.789	-0.869	-0.974
diff %	-5.44	-6.76	-7.14	-7.78	-8.56

The prices of cross-currency exchange options by Eq. (5.44) that depict Fig. 5.5. The relative difference is given by the formula: $\text{diff \%} = 100 \times \frac{(\text{PPC price} - \text{PC price})}{\text{PC price}}$.




Figure 5.5: The prices of cross-currency exchange options.

5.5 Summary

This chapter looked at the pricing of exchange options on forward contracts whose price dynamics are given by the [Heath et al. \(1992\)](#) model. Our study extended the [Margrabe \(1978\)](#) formula to generalise the option pricing formula by adjusting the variance of commodity spread to incorporate more sophisticated dependence structures for a pair of commodities. Moreover, it introduced two types of exchange options that involve cross-currency transactions: *quanto exchange option* and *cross-currency exchange option*. To our knowledge, no other commodity literature addresses the change of measure for cross-currency transactions within the [Margrabe \(1978\)](#) framework.

As the simulation study in the previous chapter implied, the PPC model calculated option prices to be lower than the PC model due to the enhanced common factor structures in the model variance. However, the reduction in option values varied between commodity pairs and maturities depending on their loadings to the common factors. It also depended on the fuel efficiency of spark spreads that determines the moneyness of the options; the impact of common factors appeared most significantly for at-the-money options. The valuation of a Dutch power plant provided an example in which omitting common factors led to the overpricing of real assets.

For the exchange options that involve FX transactions, the pricing exercise indicated that the FX adjustment is negligible for quanto exchange options. By contrast, the FX adjustment is non-negligible for cross-currency exchange options to which the relative size of the spread variance and FX add-on directly affects the spread variance in the pricing formula. Together with the remark on the option vega ([5.48](#)) that increases the complexity of risk management, our study raised awareness for latent risk factors that would have yet to be recognised by the model users and commodity traders in the international marketplace.

Chapter 6

Common Factor Hedge in the Shipping Market

6.1 Introduction

Bunker fuel refers to marine fuels classified by vessel types and chemical compositions. In January 2020, the International Maritime Organisation (IMO) imposed the mandatory requirement for bunker fuel with less than 0.5% (previously 3.5%) sulphur due to the rising concern on the environment and health caused by maritime emission. While ‘IMO 2020’ promoted the new listings of futures products on the major exchanges,¹ the new contracts suffered from illiquidity and price instability that is not uncommon in emerging markets. For the management of bunker price risk, this chapter studies the hedging of a low-sulphur bunker fuel with and without the assumption of shared common factors between proxy fuel futures. We analyse the bunker price hedge from shipowners’ view points in the market transition period.

Earlier work on hedging, as in [Johnson \(1960\)](#), [Stein \(1961\)](#) and among many others, implicitly or explicitly assumes the hedgers are commodity traders long on a physical (spot) position and take an appropriate number of short positions in a futures market. If the commodity price drops at a future time, the profit of selling the commodity decreases; however, the gain from the futures trade offsets this loss. The *minimum-variance* hedge ratio is the number of futures positions to trade to mitigate adverse price exposure when using portfolio variance as the risk metric. If the hedge aims to maximise the expected

¹Those include the Singapore Exchange (SGX), New York Mercantile Exchange (NYMEX), and Inter-continental Exchange (ICE).

return and minimise risk simultaneously, it is called the *mean-variance* hedge.² It is often the case that alternative futures are used when the attributes of the underlying and hedge contracts do not exactly match in practice (cross hedge).

For bunker risk management, the discussion of cross hedge is initiated by [Menachof and Dicer \(2001\)](#). They argue the superiority of a petroleum cross hedge against a conventional cost adjustment factor³ for Rotterdam bunker fuel at the time, with the rolling window estimate of minimum-variance hedge ratios. Later, [Alizadeh et al. \(2004\)](#) analyse the effectiveness of cross hedges for three locational bunker fuels (Rotterdam, Singapore, and Huston bunker fuels) using three petroleum futures (crude oil, gas oil, and heating oil futures). Their study finds a tendency that the constant hedge ratios perform better than dynamic hedge ratios in sample but the converse is the case for out-of-sample data.

While [Alizadeh et al. \(2004\)](#) regard shipowners as the hedgers whose aim is portfolio optimisation, [Wang et al. \(2018\)](#) regard ship liners⁴ as the consumers of bunker fuel to formulate the hedge problem. Their scenario-tree-based approach is designed to find an optimal short-term hedge strategy using a swap contract while dynamically updating a ship liner's cost function. Despite the differences in problem formulations, both studies employ a family of multivariate vector error correction GARCH (VECM-GARCH) models to account for the cointegrating relationship between the underlying and hedge contracts.

Although there is ample literature for determining optimal hedge ratios in finance, only a handful of studies focus on the hedging of bunker price risk. Given the scarcity of research in this area, firstly, our study addresses the consumption problem of shipowners who buy and consume bunker fuels rather than roll and reinvest the assets for portfolio management; these differences are often overlooked in the modern hedge literature. Secondly, it analyses the cross-hedge efficiency of bunker fuels, for which latent common factor models describe the price processes. Since our modelling approach accounts for the whole term structure of forward curves and their comovements when calculating hedge ratios, it introduces a new perspective to the existing literature.

The remainder of this chapter is organised as follows. Section 6.2 formulates a hedge problem for consumers of commodities and discusses how it differs from the hedge for portfolio managers. Section 6.3 re-introduces latent factor models and extends the framework for the calculation of minimum-variance hedge ratios. Section 6.4 performs empirical

²It is a well-known fact that these two hedge ratios are indifferent when the parameter of risk-aversion is infinite for the hedger or the expected return of the hedge instrument is zero.

³Bunker Adjustment Factor (BAF), also called fuel adjustment factor, is an additional surcharge to freight rates to adjust for bunker price fluctuations.

⁴A liner shipping company operates scheduled voyages on a sea route carrying passengers or cargo.

analysis and parameter fittings, followed by Section 6.5, in which the hedge effectiveness is evaluated. Section 6.6 summarises the results.

6.2 Bunker Price Hedge

6.2.1 Problem Formulation

Let us consider shipowners who earn money from lending their vessels to charterers for a specified route at a predetermined freight rate. Denote $S(T_i)$ for the freight charter rate in month T_i ($t < T_i$). We use the following formula to represent the economic activity of shipowners for the next N periods

$$V(t) = \sum_{i=1}^N e^{-r(T_i-t)} \mathbb{E}[\alpha S(T_i) - hB(T_i) - C], \quad (6.1)$$

where $B(T_i)$ denotes the fuel (bunker) price, C is the non-fuel costs such as port charges and canal dues, α and h adjust for the type of vessel, cargo size and fuel consumption rate, and r is the discount rate. For each period i , the cashflow excluding the discounting factor calculates a similar quantity to the Time-charter Equivalent (TCE) rate, which represents the net earnings of shipowners, typically in \$/day. In addition to the TCE components, we also consider other operating expenses and include them in C .⁵ On the cost side, the main source of uncertainty is the bunker price since it fluctuates dynamically from period to period.

Suppose that a shipowner enters into a hedge transaction in a futures market to stabilise the cost of fuel. If contracts are available for the same type of bunker fuel, taking a h long position in futures offsets a loss/gain in bunker procurement. However, it is not uncommon that the sulphur contents and viscosity of bunker fuels or the delivery point and maturity of futures are different from those of the underlying transaction. If that happens, the shipowner may need to rely on a cross hedge by buying n_i amount of alternative fuel futures that are highly correlated with the underlying bunker fuel, accepting the basis risks.

Denote by $\tilde{F}(t, T_i)$ the price of alternative fuel futures at time t for a contract month T_i and $J_i = \{d \in \mathbb{N} \mid 1 \leq d \leq 30\}$ for the set of trading days in the contract month. Since bunker futures are typically settled as the arithmetic average of the underlying index

⁵Due to the high costs involved (hence rare occurrence), we ignore the optionality in the formula associated with the possibility of ‘laying up’ the ship. Otherwise, Eq. (6.1) can be written as $V(t) = \sum_{i=1}^N e^{-r(T_i-t)} \mathbb{E}[\max(\alpha S(T_i) - hB(T_i) - C, X)]$, where X represents the layup cost.

during a contract month,⁶ we can write the settlement price as

$$\frac{1}{M_i} \sum_{d \in J_i} \tilde{B}(d) ,$$

where M_i is the size of J_i and $\tilde{B}(d)$ is the price of the alternative fuel on day d . Therefore, the present value of the hedged cash flow is

$$\tilde{V}(t) = \sum_{i=1}^N e^{-r(T_i-t)} \mathbb{E} \left[\alpha S(T_i) - hB(T_i) - C + n_i \left(\frac{1}{M_i} \sum_{d \in J_i} \tilde{B}(d) - \tilde{F}(t, T_i) \right) \right] , \quad (6.2)$$

for which we impose the following simplifying assumptions: i) risk neutrality of agents, ii) absence of credit risk, iii) absence of transaction costs, vi) independence of freight rates and bunker prices, v) constant interest rate, and vi) absence of timing risk. As $\mathbb{P} = \mathbb{Q}$ by i), the expected payoff of the hedge cashflow is zero, and the hedge only affects its variance and not the mean. Thus, the minimum-variance hedge ratio is indeed the optimal hedge ratio (Benninga et al., 1984; Myers and Thompson, 1989).⁷ Moreover, the elimination of timing risk leads to the single-day settlement of futures (instead of arithmetic average) and the simultaneous occurrence of the underlying and hedge cashflows:

$$\tilde{V}(t) = \sum_{i=1}^N e^{-r(T_i-t)} \mathbb{E} [\alpha S(T_i) - hB(T_i) - C + n_i (\tilde{B}(T_i) - \tilde{F}(t, T_i))] \quad (6.3)$$

Hence, as far as the month-by-month hedge⁸ is concerned, our problem reduces to minimising the variance of cashflows by trading futures according to the optimal hedge ratio

$$n_i^* = h \cdot \frac{\text{Cov}[B(T_i), \tilde{B}(T_i)]}{\text{Var}[\tilde{B}(T_i)]} , \quad (6.4)$$

or more generally

$$n_i^* = h \cdot \frac{\text{Cov}[F(s, T_i), \tilde{F}(s, T_i)]}{\text{Var}[\tilde{F}(s, T_i)]} , \quad (6.5)$$

where $F(s, T_i) = \mathbb{E}^{\mathbb{Q}}[B(T_i)]$ is the futures price of the underlying asset at $s(\leq T_i)$.

⁶They are known as calendar swaps. See Section 2.1.

⁷The first assumption enables us to focus on the impact of model selection on the optimal hedge ratios without concerning the risk preference of agents.

⁸Here, month-by-month hedge means hedging a cashflow in month T_i with a futures contract maturing in the same month.

6.2.2 Hedging from Consumers' Perspectives

While Eq. (6.4) calculates the hedge ratio directly from the distribution of price levels, other literature typically regards the changes in (log) prices as the risk and calculates the hedge ratio for portfolio optimisation problems. The second approach is inappropriate to hedge bunker price risks since shipowners are the consumers of bunker fuel and not the portfolio managers who continuously hold positions in the underlying asset. Below, we construct a minimum-variance hedge from the perspective of portfolio management to make the point.

Suppose that a commodity is held in storage between t_0 and t_1 . Then, the price change in this position is $\Delta B = B(t_1) - B(t_0)$. If the stored commodity is hedged by n short futures, the derivatives price also varies between the same period. Regarding the profit and loss of the futures, the change can be written as $e^{-r(T-t_1)}F(t_1, T) - e^{-r(T-t_0)}F(t_0, T)$, where the discount factors $e^{-r(T-t_i)}$, $i = 0, 1$ adjust for the time value of money for the cashflows since the payment occurs at T . When the time interval is small or the interest rate is low, $e^{-r(T-t_0)} \approx e^{-r(T-t_1)}$, and we can replace the discount factors with $D(t_0, T) := e^{-r(T-t_0)}$. Therefore, the expected change in the hedged portfolio is

$$\begin{aligned}\mathbb{E}[\Delta \tilde{V}] &= \mathbb{E}[\Delta B] - n\mathbb{E}[e^{-r(T-t_1)}(F(t_1, T) - F(0, T)) - e^{-r(T-t_0)}(F(t_0, T) - F(0, T))] \\ &\approx \mathbb{E}[\Delta B] - n D(t_0, T) \mathbb{E}[\Delta F].\end{aligned}\tag{6.6}$$

The minimisation of $\mathbb{V}\text{ar}[\Delta \tilde{V}]$ with respect to n gives the minimum-variance hedge ratio:⁹

$$n^* = \frac{1}{D(t_0, T)} \frac{\text{Cov}[\Delta B, \Delta F]}{\mathbb{V}\text{ar}[\Delta F]}\tag{6.7}$$

Ederington (1979) notes that the mean-variance hedge can be interpreted as the portfolio allocation between the spot and futures markets.¹⁰ In commodity markets, the portfolio allocation involves holding commodity stocks at t_0 and the sale and repurchase of futures contracts to hedge the physical position over the hedge period unless the commodity is delivered (Stein, 1961).

By contrast, our bunker hedge does not start from holding an underlying asset (or

⁹In fact, Eq. (6.7) and the hedge ratio by the variance and covariance of price levels are identical when we consider a one-step hedge. Since $B(t_0)$ and $F(t_0, T)$ are known quantities at t_0 , $\mathbb{V}\text{ar}(\Delta F) = \mathbb{V}\text{ar}(F(t_1, T))$ and $\text{Cov}(\Delta B, \Delta F) = \text{Cov}(B(t_1), F(t_1, T))$.

¹⁰Alternatively, when hedging a forward position with another (alternative) forward/futures, Eq. (6.6) can be written as

$$n^* = \frac{D(t_0, T_1)}{D(t_0, T_2)} \frac{\text{Cov}[\Delta F, \Delta \tilde{F}]}{\mathbb{V}\text{ar}[\Delta \tilde{F}]},\tag{6.8}$$

where ΔF and $\Delta \tilde{F}$ represent the change in the underlying and hedge positions, respectively.

short futures) at t_0 or carrying over the commodity during the hedge period. Instead, the shipowner buys bunker fuel at t_1 in the physical market and consumes the fuel for immediate use. For an N -period hedge, the fuel cost accumulates, and the expected costs sum to

$$\sum_{i=1}^N \mathbb{E}[B(t_i)] = \mathbb{E}[B(t_1) + B(t_2) + \dots + B(t_N)], \quad (6.9)$$

for which the variance is minimised using futures contracts. On the other hand, the expected price change in the portfolio approach is

$$\begin{aligned} \sum_{i=1}^N \mathbb{E}[\Delta B(t_i)] &= \mathbb{E}[B(t_N) - B(t_{N-1}) + B(t_{N-1}) - B(t_{N-2}), \dots, +B(t_1) - B(t_0)] \\ &= \mathbb{E}[B(t_N) - B(t_0)]. \end{aligned} \quad (6.10)$$

In an extreme situation, when two prices are the same at t_0 and t_N , the price change in the portfolio approach is zero. However, this is not the case in the fuel consumption problem, where the sum of all prices enters the calculation: the consumption problem is *path-dependent*.

The path-dependency feature in the consumption problem highlights the reason why the hedge ratios by price change cannot capture the inherent characteristics of our bunker price hedge. In other words, while $\text{Var}[\Delta V]$ is an accurate representation of risk, $\text{Var}[\Delta B]$ fails to capture the cumulative exposure of monthly spot bunker purchases given by Eq. (6.3). In general, the formulation of the hedge ratio (6.7) is only applicable for market participants who are capable of continuously rebalancing their positions, which is not always feasible in commodities markets due to liquidity, transaction costs, or physical operations involved. As seen in the bunker consumption problem, the hedge ratio needs to be obtained directly from the distribution of level data to account for the cumulative effect in the price exposure when dealing with a sequence of future cash flows over a non-small time interval.

6.3 Model

This section outlines the structures of common and non-common factor models by PCA and PPCA for the calculation of hedge ratios. For some parts, we shall provide references to the relevant equations and sections in the previous chapters instead of re-introducing the notations and formulas.

Notation

Denote the non-common factor model by $\text{PC}(n_k)$, where n_k is the number of principal components used in the forward curve model. Similarly, denote the PPC model by $\text{PPC}(\sum_{k=1}^K n_k, n_k, m, 1)$, where $\sum_{k=1}^K n_k$ is the total number of marginal principal components used to perform the second-stage PCA, m is the number of common components, and 1 refers to the additional idiosyncratic component: see Section 4.2.2. We set $n_k = n$ for $k = 1, \dots, K$ for simplicity.

Fuel Price Dynamics

The forward curve dynamics of bunker futures are described by the term structure model (4.1):

$$\frac{dF_k(t, T)}{F_k(t, T)} = \sum_{j=1}^n \sigma_{kj}(t, T) dZ_{kj}^{\mathbb{Q}}(t), \quad k = 1, \dots, K,$$

where $F_k(t, T)$, $k = 1, \dots, K$, represent the price of forward contracts in the k -th forward curve with maturity T , $\sigma_{kj}(t, T)$, $j = 1, \dots, n$, are the factor volatilities that govern the dynamics of the forward curves, and $Z_{kj}^{\mathbb{Q}}$ are the $\mathbb{Q}$ -Brownian motions associated with the factors. These Brownian motions can be either correlated or uncorrelated, depending on the model assumptions.

Calculation of Factor Hedge Ratios

Section 5.2.2 shows that solving the multi-factor model gives us the marginal distribution of log processes $\ln F_k(s, T) \sim \mathcal{N}(\mu_k, \sigma_k^2)$, $t < s \leq T$, with

$$\begin{aligned} \mu_k &= \mathbb{E}[\ln F_k(s, T)] = \ln F_k(t, T) - \frac{1}{2} \int_t^s \sum_{j=1}^n \sigma_{kj}^2(u, T) du \\ \sigma_k^2 &= \text{Var}[\ln F_k(s, T)] = \int_t^s \sum_{j=1}^n \sigma_{kj}^2(u, T) du, \end{aligned}$$

and the covariance

$$\text{Cov}[\ln F_1(s_1, T_1), \ln F_2(s_2, T_2)] = \int_t^{\min(s_1, s_2)} \left(\sum_{i=1}^n \sum_{j=1}^n \rho_{ij} \sigma_{1i}(u, T_1) \sigma_{2j}(u, T_2) \right) du,$$

where $t < s_i < T_j$, $\forall i, j \in \{1, 2\}$, for which the forward curves are driven by the same number of orthogonal factors. Since the computation of the hedge ratio (6.5) requires the distributions of level price at some future date s_k , we cannot directly substitute the

variance and covariance of $\ln F_k(s_k, T_k)$ into the formula. Instead, we rely on the following general results to obtain the variance and covariance of lognormal processes.

Let $Y_k = e^{X_k}$ be a lognormally distributed random variable where $X_k \sim \mathcal{N}(\mu_k, \sigma_k^2)$. Then the mean and variance of Y_k are given by

$$\begin{aligned}\mu_{Y_k} &= \mathbb{E}[Y_k] = e^{\mu_k + \frac{1}{2}\sigma_k^2} \\ \sigma_{Y_k}^2 &= \mathbb{V}\text{ar}[Y_k] = \left[e^{\sigma_k^2} - 1 \right] e^{(2\mu_k + \sigma_k^2)},\end{aligned}\tag{6.11}$$

and the covariance of two lognormally distributed random variables is

$$\begin{aligned}\mathbb{C}\text{ov}[Y_1, Y_2] &= \mathbb{E}[Y_1 Y_2] - \mathbb{E}[Y_1] \mathbb{E}[Y_2] \\ &= e^{(\mu_1 + \mu_2)} \cdot e^{\frac{1}{2}(\sigma_1^2 + \sigma_2^2 + 2\rho_{12}\sigma_1\sigma_2)} - e^{\mu_1 + \frac{1}{2}\sigma_1^2} \cdot e^{\mu_2 + \frac{1}{2}\sigma_2^2} \\ &= e^{(\mu_1 + \mu_2)} \cdot e^{\frac{1}{2}(\sigma_1^2 + \sigma_2^2)} \left[e^{\rho_{12}\sigma_1\sigma_2} - 1 \right].\end{aligned}\tag{6.12}$$

From the relationship above, the variance and covariance of $F_k(s, T)$, $k = 1, 2$ can be found analytically for μ_k and σ_k :

$$\mathbb{V}\text{ar}[F_k(s, T)] = (F_k(t, T))^2 \left[\exp \left\{ \int_t^s \sum_{j=1}^n \sigma_{kj}^2(u, T) du \right\} - 1 \right] \tag{6.13}$$

$$\begin{aligned}\mathbb{C}\text{ov}[F_1(s_1, T_1), F_2(s_2, T_2)] &= (F_1(t, T_1) F_2(t, T_2)) \\ &\times \left[\exp \left\{ \int_t^{\min(s_1, s_2)} \left(\sum_{i=1}^n \sum_{j=1}^n \rho_{ij} \sigma_{1i}(u, T_1) \sigma_{2j}(u, T_2) \right) du \right\} - 1 \right]\end{aligned}\tag{6.14}$$

Hence, one can obtain a succinct expression for the minimum-variance hedge ratio from Eqs. (6.13) and (6.14):¹¹

$$n_i^* = \frac{h F_1(t, T_1)}{F_2(t, T_2)} \cdot \frac{\exp \left\{ \int_t^{\min(s_1, s_2)} \left(\sum_{i=1}^n \sum_{j=1}^n \rho_{ij} \sigma_{1i}(u, T_1) \sigma_{2j}(u, T_2) \right) du \right\} - 1}{\exp \left\{ \int_t^{s_2} \sum_{j=1}^n \sigma_{2j}^2(u, T_2) du \right\} - 1} \tag{6.15}$$

The covariance structure of Eq. (6.15) depends on the model of our choice; see Section 5.2.2 for the explanations.

Parameter Estimation

Both the PC and PPC volatility functions are exponentially fitted by the functional form of Eq. (4.4); see Thompson (2016). The calibration results are evaluated by the matrix-

¹¹ $F_2(t, T_2)$ is the hedge instrument in Eq. (6.15).

wise root-mean-square errors (m-RMSE) between the model volatility σ_{ij} and fitted factor volatility $\tilde{\sigma}_{ij}$, where the number of estimation points M equals to the number of maturities multiplied by the number of factors in a model: $M = N \times n$ for a PC model or $M = N \times (m + 1)$ for a PPC model.

$$\text{m-RMSE} = \sqrt{\frac{\sum_{i=1}^N \sum_{j=1}^n (\sigma_{ij} - \tilde{\sigma}_{ij})^2}{M}} \quad (6.16)$$

6.4 Empirical Study

6.4.1 Preliminary Analysis

The hedger in this study is a shipowner who operates in the TC2 clean-tanker route between Rotterdam and the US Atlantic coast. For the calculation of hedged cashflows per day, the constant parameters are determined by the following formulas for Eq. (6.3)

$$\begin{aligned} \alpha &= \frac{\text{Cargo size} \times \text{Flat rate}}{\text{Number of days}} \times \frac{1}{100} \\ h &= \frac{\text{Fuel consumption}}{\text{Number of days}} \\ C &= \frac{\text{Other costs}}{\text{Number of days}} + \text{Daily operating expenses} , \end{aligned}$$

for which we substitute typical values representing the TC2 route:¹²

Cargo size	: 37,000 MT
Flat rate	: 14.73 (\$/MT)
Fuel consumption (per voyage)	: 644 MT
Other costs (per voyage)	: \$ 88,884
Operating expenses (per day)	: \$ 6,000
Number of days (per voyage)	: 28

Therefore, $\alpha = 194.65$, $h = 23$, and $C = 9,174.43$. A flat rate is published by the Worldscale Association¹³ in January every year as a reference charter rate (\$/MT) for a standardised model ship running a particular route. On the other hand, a market freight rate is quoted in the % value of the flat rate known as the Word Scale (WS). It fluctuates day to day, reflecting the supply and demand on the route. We refer to the year 2019's

¹²One can easily extend the analysis for other tanker routes by altering these values.

¹³<https://www.worldscale.co.uk/>




Figure 6.1: Interpolated forward curves with fixed time to maturities.

All prices are shown in \$/1000 MT with fixed time to maturities: $\tau_i = (30, 60, 90, \dots, 540)$ days for the period of 19th November 2019 to 20th October 2020.

flat rate and fix the market freight rate at WS 150 since the freight rate and bunker price are independent by assumption, and therefore it does not influence the hedge ratio (6.15).

The fuel data consists of bunker futures that belong to categories B and C in Section 2.1, whose names are denoted by the Bloomberg® ticker codes in the subsequent analysis: European FOB Rotterdam Marine Fuel 0.5% (RMI), European 3.5% Fuel Oil Barges FOB Rotterdam (LW), Gulf Coast HSFO (WT), and WTI Financial (VH). We shall use their short names for labelling also: Rotterdam 0.5% (RMI), Rotterdam 3.5% (LW), US Gulf coast 3.0%, and WTI crude. These price data are linearly interpolated with a 30-day interval for every forward curve to construct the time series of 18 fixed time-to-maturity contracts, where $\tau_i = (30, 60, \dots, 540)$ days in Eq. (2.1), to cover 231 trading days between 19th November 2019 and 20th October 2020.

Figure 6.1 visualises the interpolated RMI, LW, WT, and VH prices in \$/1000 MT,¹⁴ in which the forward curves exhibit similar patterns. The price of VH was above \$350 at

¹⁴For WT and VH, the trade units are adjusted from \$/barrel to \$/metric tonnes (MT) using the conversion factor of 1 barrel = 6.7 MT.




Figure 6.2: Log returns of 30-day and 510-day fixed time-to-maturity contracts.

a premium to LW and WT until the world's demand for crude oil plummeted sharply due to the first outbreak of the COVID-19 pandemic in March 2020. The forward prices of LW and VH were in steep contango between early April and May 2020 until they reverted to the pre-crisis level in June 2020. On the other hand, the nearby price of RMI was above \$500 in the introductory phase of the new IMO regulation in January 2020. Subsequently, it jumped downwards in February 2020, followed by a gentle price rise towards \$300 in later months.

Fig. 6.2 depicts the log returns of the 30-day and 510-day contracts. The fluctuation of returns synchronises across forward curves in a stressful period, for example, between March and May 2020. Overall, the 30-day contracts exhibit greater volatility with log returns in the range of -36% to $+24\%$ compared with -18% to $+12\%$ for 510-day contracts. Table 6.1 reports the summary statistics of price levels and log returns.¹⁵ Notice that the 30-day volatilities are smaller than the 90-day volatilities due to the average settling method of the bunker futures that restrains radical price movements in the contract month. In fact, the term structure peaks at 60 days with 68%, 84%, 84%, and 90% volatilities for RMI, LW, WT, and VH.

Fig. 6.3 investigates the within-curve correlations of forward price log returns, where the lower-triangle of a heatmap represents the off-diagonal elements of a correlation matrix. The first columns of the heatmaps highlight the tendency that the farther the maturities, the weaker the correlations due to the split personality of energy commodities.¹⁶ However,

¹⁵The table only reports the statistics for $T_i = \{30, 90, \dots, 510\}$ with a 60-day interval due to limitations in space, but the patterns are similar for the other maturities.

¹⁶It is said that the split personality is caused by different fundamentals affecting the short and long-end of forward curves. Pilipovic (1998) argues that it is the current storage conditions that determine short-term prices, whereas it is the expected long-term supply conditions that determine long-term prices.

Table 6.1: Summary statistics of interpolated fuel prices and log returns.

RMI									
Maturities	30	90	150	210	270	330	390	450	510
Price									
Min	168	167.15	193.35	210.60	221.40	233.21	244.37	253.64	262.48
Max	575.61	556.94	540.05	525.16	511.29	499.07	485.93	473.68	462.18
Mean	333.74	337.73	341.25	344.05	346.07	348.13	349.93	351.23	352.35
SD	108.50	102.84	94.31	86.77	79.74	72.89	66.70	61.13	55.71
Return(%)									
Min	-19.73	-20.82	-18.25	-17.12	-15.96	-13.06	-12.51	-12.80	-13.63
Max	12.84	12.96	12.75	12.40	11.88	12.11	12.02	11.59	11.97
Mean	-0.17	-0.18	-0.18	-0.17	-0.16	-0.15	-0.14	-0.14	-0.13
Volatility	57.09	64.08	58.02	54.43	51.80	50.38	48.96	46.81	46.09
LW									
Maturities	30	90	150	210	270	330	390	450	510
Price									
Min	101.99	100.20	117.15	126.53	134.91	143.95	151.98	159.05	165.31
Max	276.20	277.80	276.10	273.16	269.21	267.12	268.02	268	267.36
Mean	213.95	214.67	216.86	219.63	222.25	224.95	227.85	230.63	233.03
SD	46.53	44.06	39.57	36.90	34.57	32.26	30.49	29.09	28.02
Return(%)									
Min	-33.88	-33.39	-31.31	-28.80	-26.84	-24.25	-22.07	-19.30	-17.48
Max	21.21	18.96	16.95	15.93	14.87	13.67	12.42	11.31	11.20
Mean	0.08	0.12	0.10	0.08	0.07	0.06	0.05	0.03	0.03
Volatility	73.70	77.70	69.85	64.98	61.17	57.18	53.52	49.83	47.45
WT									
Maturities	30	90	150	210	270	330	390	450	510
Price									
Min	105.19	99.76	114.90	125.96	135.00	144.25	153.16	160.60	167.37
Max	302.44	299.36	295.81	292.32	288.57	284.95	285.42	286.29	288.70
Mean	229.50	227.25	228.81	231.28	233.78	236.42	239.42	242.34	244.92
SD	52.08	49.78	45.46	42.44	39.98	37.69	35.88	34.47	33.32
Return(%)									
Min	-36.03	-33.52	-30.78	-28.25	-25.95	-23.41	-21.24	-19.17	-17.56
Max	24.36	19.18	18.10	16.90	15.62	14.31	13.18	12.43	11.80
Mean	0.04	0.11	0.09	0.07	0.06	0.05	0.03	0.02	0.01
Volatility	77.50	79.31	71.64	65.97	61.57	57.41	53.59	50.44	48.02
VH									
Maturities	30	90	150	210	270	330	390	450	510
Price									
Min	101.91	132.86	160.67	176.01	187.13	195.71	203.08	207.90	212.12
Max	422.84	419.89	413.12	405.22	397.24	390.41	383.71	378.22	373.39
Mean	278.23	284.25	288.62	290.92	292.50	293.70	294.58	295.29	295.96
SD	79.87	72.14	65.22	60.06	55.57	51.92	48.79	46.24	43.96
Return(%)									
Min	-30.09	-31.10	-24.67	-22.55	-20.60	-18.81	-17.31	-16.02	-14.85
Max	21.89	19.56	16.74	14.08	12.10	10.42	9.32	8.29	7.39
Mean	-0.14	-0.12	-0.11	-0.10	-0.09	-0.08	-0.08	-0.08	-0.07
Volatility	71.70	75.39	62.83	55.55	50.30	46.14	42.86	40.25	38.11

all combinations of contracts in a single forward curve are highly correlated, approximately from 60% to perfect correlation.

In the next analysis, PCA is performed individually to the covariance matrices of fuel forward curves. Table 6.2 reports the cumulative and marginal % of total variations




Figure 6.3: Within-curve correlations of forward price log returns.

Table 6.2: The cumulative percentage of total variation explained by PCA.

		1	2	3	4	5	6	7	8	9	10
RMI	Cumulative %	93.61	97.24	98.95	99.51	99.77	99.87	99.93	99.96	99.98	99.98
	Marginal %	93.61	3.63	1.71	0.57	0.26	0.10	0.05	0.04	0.01	0.01
LW	Cumulative %	94.41	97.71	99.51	99.79	99.93	99.97	99.98	99.99	99.99	100.00
	Marginal %	94.41	3.30	1.80	0.28	0.13	0.04	0.02	0.01	0.00	0.00
WT	Cumulative %	94.17	97.96	99.58	99.83	99.95	99.97	99.99	99.99	100.00	100.00
	Marginal %	94.17	3.80	1.61	0.25	0.12	0.03	0.01	0.01	0.00	0.00
VH	Cumulative %	91.65	97.66	99.66	99.91	99.96	99.99	100.00	100.00	100.00	100.00
	Marginal %	91.65	6.02	2.00	0.25	0.05	0.03	0.01	0.00	0.00	0.00

The reported values are in % as the cumulative sum of first ten eigenvalues in the PCA.

captured by the ten most significant principal components, in which the explanatory power of the first two components exceeds 97% for all the fuels. The marginal increments are similar between RMI, LW, and WTI, whereas the additional contribution declines sharply after the third component.

Fig. 6.4 depicts the eigenvectors corresponding to the first three largest eigenvalues. The first eigenvectors seem visually identical, indicating the presence of a common level factor for the forward curves. In the meantime, the factor loading of 30-day contracts dominates a large proportion in the second and third factor loadings. The shape of the third eigenvectors partially duplicates the second eigenvectors, and they do not form a






Figure 6.4: The first three eigenvectors of fixed time-to-maturity forward curves.

typical quadratic shape as the ‘curvature’ factor, but still reflect the case that the shortest- and longest-end of the curve move in the same direction.

Fig. 6.5 shows the history of the first two PCs in level. In the top figure, the level factors outline the comovement of original price trajectories in reduced dimensions (see Fig. 6.1), highlighting the divergence of RMI – VH and LW – WT after the market crash in April 2020. On the other hand, the slope factors move steadily during the same observation period in the bottom figure. In absolute value terms, the correlation coefficients vary between 87.9% and 99.5% for the first principal components, 36.2% and 93.1% for the second principal components, and remain below 11% for cross-components (P_1 and P_2 of different fuels), as reported in Table 6.3.

Table 6.3: Cross-correlation matrix of principal components.

		RMI		LW		WT		VH	
		PC1	PC2	PC1	PC2	PC1	PC2	PC1	PC2
RMI	PC1	1	0	0.879	0.069	0.882	-0.092	0.887	-0.014
	PC2	0	1	-0.088	0.501	-0.109	-0.468	-0.106	-0.362
LW	PC1	0.879	-0.088	1	0	0.995	-0.039	0.917	0.002
	PC2	0.069	0.501	0	1	-0.033	-0.931	0.053	-0.581
WT	PC1	0.882	-0.109	0.995	-0.033	1	0	0.922	0.039
	PC2	-0.092	-0.468	-0.039	-0.931	0	1	-0.062	0.555
VH	PC1	0.887	-0.106	0.917	0.053	0.922	-0.062	1	0
	PC2	-0.014	-0.362	0.002	-0.581	0.039	0.555	0	1

**(a) First principal components (level factors)****(b) Second principal components (slope factors)****Figure 6.5: Time series of principal components in level.**

The principal component representation of log returns are transformed into levels using the relationship $p_t = p_{t-1} e^{\tilde{x}_t}$, where $\tilde{x}_t$ is the principal component representation of daily return on day t , and p_t is the principal component in level. The starting value (p_0) is set equal to the first observation in the price data.

6.4.2 Parameter Estimation

The fuels are divided into three panels to investigate the impact of common factors in our bunker cross-hedge problem. In panel 1, the price risk of the low-sulphur bunker fuel (RMI) is hedged with the high-sulphur bunker fuel (LW) contract that refers to the same delivery point but with better liquidity. In addition to the base case, panels 2 and 3 consider WT and VH data that may potentially improve or diminish the performance of the cross hedge. The underlying and hedge contracts remain the same in all three panels.

Underlying contract: RMI


Hedged by: LW

Data used:


Panel 1: RMI, LW

Panel 2: RMI, LW, WT

Panel 3: RMI, LW, WT, VH




(a) Panel 1: RMI, LW



(b) Panel 2: RMI, LW, WT

Figure 6.6: PPC eigenvectors.

The plots of common eigenvectors. The subfigure titles refer to the column numbers in the common eigenvector matrix $\bar{\mathbf{U}}$ in Eq. (3.18). Every set of two consecutive ticks indicates the marginal PC loadings for a forward curve.



(c) Panel 3: RMI, LW, WT, VH

Table 6.4: The total variation explained by PPCA.

		1	2	3	4	5	6	7	8
Panel 1	Cumulative %	47.09	84.73	97.09	100.00				
	Marginal %	47.09	37.64	12.36	2.91				
Panel 2	Cumulative %	47.45	85.80	96.32	98.80	99.92	100.00		
	Marginal %	47.45	38.35	10.52	2.48	1.13	0.08		
Panel 3	Cumulative %	46.91	81.22	89.47	96.10	97.97	99.20	99.95	100.00
	Marginal %	46.91	34.31	8.24	6.64	1.86	1.23	0.75	0.05

We perform PPCA on the cross-correlation matrix (6.3) using two PCs from each fuel forward curve to find the common eigenstructures. Fig. 6.6 reveals the common eigenvectors consisting of $2K$, $K = 2, 3, 4$, marginal eigenvectors for panels 1 – 3, respectively, for which the forward curve model can be denoted by $\text{PPC}(2K, 2, m, 1)$.¹⁷ In these figures, every set of two consecutive ticks indicates the marginal PC loadings for a forward curve. Noticeably, the factor loadings are equally significant among the first and second PCs, while the unequal factor loadings become apparent after the third common eigenvectors in panels 2 – 3. The latter observation suggests that the ‘common eigenvectors’ may no longer be common to all the forward curves but instead dominated by factors from one or two of the commodities. Table 6.4 reports PPCs’ explanatory powers by the eigenvalues. The first two PPCs capture approximately 85% of the cross-market correlation structure for panels 1 and 2, while slightly below 81% for panel 3. Adding the third PPC increases the explanatory power to 97%, 96%, and 91%; however, its marginal contribution is less than one-third of the second principal component.

In accordance with the results above, we suggest two common factors (i.e. $m = 2$) for the $\text{PPC}(2K, 2, m, 1)$, $K = 2, 3, 4$, models to compare the hedge performances with a non-common factor model, $\text{PC}(2)$, in the subsequent analysis. The formulation of the forward curve models follows Section 4.2. We employ the local covariance fitting for the estimation of idiosyncratic factors and parametrically fit the factor volatilities using the functional form in Eq. (4.4).

Table 6.5 reports the estimated parameters by model. The fitted parameters $X_1 - X_3$ and $X_4 - X_6$ correspond to the 1st and 2nd PC factor volatilities. The PPC models report $Y_7 - Y_9$ for the idiosyncratic factors in addition to $Y_1 - Y_3$ and $Y_4 - Y_6$. According to m-RMSE, the exponential parameterisation works better for the PPC models than the PC model.

¹⁷Recall the notation $\text{PPC}(\sum_{k=1}^K n_k, n_k, m, 1)$, where n_k , m , 1 represent the number of PC factors, common factors, and idiosyncratic factors, respectively. We use two principal components per forward curve to perform PPCA. Therefore, $\sum_{k=1}^K n_k = 2K$.

Table 6.5: Fitted parameter values of the PC and PPC models.

PC		X1	X2	X3	X4	X5	X6				α_1	α_2	m-RMSE
	RMI	0.363	-18.363	0.480	5.068	0.480	-0.274				0.848	54.052	0.0082
	LW	0.584	-2.964	0.249	3.785	0.249	-0.155				0.963	29.022	0.0072
	WT	0.600	-2.161	0.231	3.064	0.231	-0.144				1.071	24.845	0.0062
	VH	0.836	-2.819	0.352	1.601	0.352	-0.114				2.703	20.743	0.0160
Panel 1		Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	β_1	β_2	m-RMSE
PPC(4, 2, 2, 1)	RMI	0.308	-16.309	0.347	0.433	2.587	-0.203	0.123	-3.695	0.062	0.889	52.673	0.0061
	LW	0.580	-3.195	0.293	0.206	3.931	-0.148	0.136	-0.952	0.079	0.925	30.921	0.0056
Idio.Corr	RMI/LW	-0.965											
Panel 2		Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	β_1	β_2	m-RMSE
PPC(6, 2, 2, 1)	RMI	0.292	-20.029	0.349	0.350	2.381	-0.156	0.173	-6.044	0.080	0.942	55.309	0.0054
	LW	0.584	-2.929	0.305	0.226	3.810	-0.156	0.092	-0.603	0.056	0.950	29.490	0.0058
Idio.Corr	WT	0.579	-2.367	0.339	-0.254	-2.834	0.117	0.100	-0.242	0.047	1.058	25.181	0.0050
	RMI/LW	-0.921											
	RMI/WT	-0.891											
	LW/WT	0.814											
Panel 3		Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	β_1	β_2	m-RMSE
PPC(8, 2, 2, 1)	RMI	-0.287	23.588	-0.351	-0.329	-2.835	0.149	0.201	-7.069	0.071	0.950	57.277	0.0053
	LW	-0.584	2.898	-0.301	-0.223	-3.792	0.154	0.096	-0.933	0.075	0.947	29.619	0.0058
	WT	-0.575	2.398	-0.338	0.252	2.801	-0.114	0.118	-0.267	0.054	1.053	25.333	0.0049
	VH	-0.738	4.251	-0.344	0.263	2.072	-0.093	-0.205	1.151	-0.094	2.418	27.381	0.0115
Idio.Corr	RMI/LW	-0.625	LW/WT	0.870									
	RMI/WT	-0.620	LW/VH	0.529									
	RMI/VH	0.236	WT/VH	0.505									

The top panel shows the exponentially fitted parameter values for the volatility curves of the PC(2) model. The factor correlation matrix is shown separately in Table 6.3. Panels 1 – 3 report the results of the PPC($2K$, 2, 2, 1) models for $K = 2, 3, 4$. The rows under *Idio.Corr* show the factor correlations. The last column reports the matrix-wise root-mean-square errors (m-RMSE) by Eq. (6.16).

The same table reports the model implied idiosyncratic factor correlations at the bottom of each panel. In absolute values, RMI/LW is close to perfect correlation in panel 1 (-97%). However, the correlation declines when other fuels are added in panels 2 (-92%) and 3 (-63%). As for other fuel pairs, the estimation results reflect differences in fuel composition and product of origin. For example, the high sulphur fuels LW/WT are highly correlated in panels 2 (81%) and 3 (87%), while RMI/VH are weakly correlated (24%) in panel 3, being the most diverse fuels in terms of sulphur contents and location.

6.5 Cashflow Hedge

The hedge aims to achieve month-by-month variance reduction for the bunker purchase cashflows in Eq. (6.3). The cashflows are assumed to occur on the 15th day of month T_i , $i = 1, \dots, 18$, which is multiplied by 30 to convert the fuel consumption from daily to monthly. The minimum-variance hedge ratios are given by Eq. (6.15), where $s_1 = s_2$ and $T_1 = T_2$ in our case.

The hedge performances are evaluated by two measures: the *Hedge Effectiveness* (HE) of Ederington (1979) and Relative Effectiveness (RE). The first measure quantifies the % reduction made in the variance of hedged positions:

$$HE(T_i) = \frac{\text{Var}[C(T_i)] - \text{Var}[\tilde{C}(T_i)]}{\text{Var}[C(T_i)]}, \quad (6.17)$$

where $C(T_i)$ and $\tilde{C}(T_i)$ denote the unhedged and hedged cashflows in period T_i . We use the same formula and denote the hedged cashflows of the PC and PPC models by $C(T_i)$ and $\tilde{C}(T_i)$ to determine the RE. A positive RE implies the superiority of a PPC model over a PC model in risk reduction, and vice versa.

Table 6.6 reports the results of the cashflow hedge for nine representative months.¹⁸ It also compares the hedge effectiveness with the naïve hedge ratio, which is $n_i = h = 23$ in Eq. (6.3), as the simple one-to-one hedge may perform better than more sophisticated but computationally expensive hedge strategies (Alexander et al., 2013). We round down the minimum-variance hedge ratios to the nearest integer; the hedged cashflows are subject to position risk, which is the unhedged risk arising from omitting the optimal number of futures contracts below decimal points. The RE is reported for the naïve and minimum-variance hedges separately at the bottom.

In Table 6.6, the mean-variance hedge achieves several points higher HE than the

¹⁸The statistics include all eighteen months' cashflows in the following analysis.

Table 6.6: Comparison of variance and hedge effectiveness between models.

Maturities		30	90	150	210	270	330	390	450	510
PC(2)										
Variance	Unhedged	1.87	9.95	17.66	24.29	29.99	35.16	39.77	44.27	48.53
	Naïve	1.82	5.22	8.21	10.52	12.36	13.98	15.38	16.76	18.17
	MV	1.68	4.76	7.20	9.26	11.03	12.67	14.17	15.72	17.27
Hedge Ratio	Naïve	23	23	23	23	23	23	23	23	23
	MV	12.40	32.91	33.36	32.36	31.30	30.35	29.39	28.47	27.79
HE (%)	Naïve	2.80	47.48	53.50	56.71	58.80	60.24	61.33	62.13	62.56
	MV	10.38	52.18	59.21	61.88	63.24	63.98	64.36	64.49	64.42
Panel 1: PPC(4, 2, 2, 1)										
Variance	Unhedged	1.51	9.46	17.11	23.71	29.39	34.52	39.09	43.51	47.68
	Naïve	0.97	3.71	6.03	7.70	8.93	9.94	10.70	11.38	12.05
	MV	0.97	2.57	4.01	5.22	6.26	7.19	8.01	8.81	9.59
Hedge Ratio	Naïve	23	23	23	23	23	23	23	23	23
	MV	22.25	38.76	37.88	36.28	34.88	33.73	32.63	31.62	30.89
HE (%)	Naïve	35.36	60.82	64.77	67.53	69.60	71.20	72.62	73.85	74.73
	MV	35.40	72.85	76.55	77.98	78.68	79.18	79.52	79.75	79.89
RE (%)	Naïve	46.57	29.08	26.59	26.80	27.70	28.89	30.41	32.15	33.72
	MV	42.09	46.02	44.31	43.62	43.19	43.25	43.52	43.95	44.47
Panel 2: PPC(6, 2, 2, 1)										
Variance	Unhedged	1.68	9.62	17.25	23.84	29.50	34.61	39.13	43.50	47.59
	Naïve	1.40	4.18	6.59	8.36	9.68	10.77	11.60	12.32	13.03
	MV	1.36	3.32	4.91	6.25	7.41	8.43	9.32	10.19	10.97
Hedge Ratio	Naïve	23	23	23	23	23	23	23	23	23
	MV	16.66	36.58	36.45	35.16	33.92	32.85	31.82	30.85	30.15
HE (%)	Naïve	16.74	56.50	61.80	64.95	67.19	68.89	70.37	71.67	72.62
	MV	19.54	65.51	71.53	73.77	74.90	75.65	76.18	76.57	76.94
RE (%)	Naïve	22.96	19.92	19.73	20.55	21.67	22.98	24.61	26.49	28.32
	MV	19.26	30.29	31.82	32.47	32.84	33.46	34.23	35.17	36.48
Panel 3: PPC(8, 2, 2, 1)										
Variance	Unhedged	1.99	9.97	17.62	24.22	29.88	34.99	39.51	43.87	47.94
	Naïve	1.77	4.67	7.19	9.06	10.47	11.63	12.53	13.33	14.10
	MV	1.70	3.88	5.64	7.11	8.35	9.47	10.44	11.38	12.27
Hedge Ratio	Naïve	23	23	23	23	23	23	23	23	23
	MV	15.54	36.03	35.98	34.72	33.50	32.46	31.44	30.48	29.78
HE (%)	Naïve	10.95	53.10	59.18	62.61	64.98	66.76	68.29	69.61	70.59
	MV	14.21	61.09	67.98	70.63	72.04	72.94	73.58	74.05	74.40
RE (%)	Naïve	2.86	10.54	12.39	13.88	15.31	16.80	18.53	20.47	22.41
	MV	-1.50	18.50	21.67	23.18	24.24	25.24	26.34	27.58	28.93

The table shows the results of naïve and minimum-variance hedges by model, in which MV stands for the minimum-variance hedge. While HE (%) represents the Ederington's effectiveness measure given a model, RE (%) represents the relative performance of a PPC model with respect to the PC(2) model for the same type of hedges.

naïve hedge overall, for most of which the difference appears notably at the shortest end. The minimum-variance hedge offers on average 3.91% more variance reduction than the naïve hedge for PC(2). For PPC models, the mean-variance hedge is on average 8.06%,

6.71% and 6.08% more efficient than the naïve hedge for panels 1 – 3, respectively. The cashflow hedging works ineffectively at 30-day irrespective of model and hedge types. The HE increases gradually after 30 day towards longer-term cashflows, and this tendency is persistent between panels.

The RE quantifies the gain/loss in variance reduction brought by common factor approach. It is only negative for the minimum-variance hedged 30-day cashflow in panel 3, meaning that the PC(2) model outperforms PPC(8, 2, 2, 1) model for the cashflow hedging in month 1. The RE is always positive for the rest of cashflows; however, it declines as the number of underlying assets increases in panels 2 and 3. The results of panel 1 highlights the ability of our common factor approach to improve bunker price hedging when RMI and VH are modelled together. On the other hand, it is unlikely that the inclusion of WT and VH in panels 2 and 3 add any values to the common factor hedge.

Fig. 6.7 visualises the minimum-variance hedge by the PC(2) and PPC(4, 2, 2, 1) models for monthly cashflows for purchasing of bunker fuels shown in shaded bars. In the figure, the monthly cashflows (6.3) are standardised by $h = 23$ for an illustration. The dark and light grey bars represent hedge positions by model, and the black line indicates the initial level of the forward curve for LW. It shows that the PPC model calculates greater cashflows than the PC model based on the minimum-variance hedge ratios in Table 6.6. However, when we decompose the model covariances, the two models give similar levels of marginal volatilities; hence, the implied model correlation is causing this behaviour. As a result, when PPC models capture the comovement of contracts better than the benchmark model via enhanced model correlation structures, the hedge ratios, which are the ratio of the covariance to the variance of hedge instrument, increase in size, which better offsets the risks in monthly cashflows.

To conclude our analysis, we construct the confidence intervals for the expected hedged cashflows $\tilde{C}(T_i)$ and unhedged cashflows $C(T_i)$ by month. With the assumption of an *i.i.d* normal distribution for the cashflows, which are the differences of two lognormally distributed random variables, the $100(1 - \alpha)\%$ confidence intervals can be approximated as

$$CI_{(1-\alpha)} = [\mu_i - Z_{\alpha/2}\sigma_i/\sqrt{df}, \quad \mu_i + Z_{\alpha/2}\sigma_i/\sqrt{df}] ,$$

where

$$\mu_i = \mathbb{E}[C(T_i)], \quad \sigma_i = \sqrt{\text{Var}[C(T_i)]}, \quad Z_{\alpha/2} := \Phi^{-1}(1 - \alpha/2),$$

Φ is the standard normal cumulative distribution function, and df is the degree of free-




Figure 6.7: Standardised cashflows of monthly fuel consumption and hedge positions.


The cashflows are standardised by $h = 23$ for Eq. (6.3). Shaded bars: monthly cash outflows for the purchase of RMI. Dark grey (PC(2)) and light grey (PPC(4, 2, 2, 1)) bars: the monthly cash inflows from the hedge positions. Black line: the initial level of forward curve for LW.

dom.¹⁹


Fig. 6.8 depicts the confidence intervals, in which $\alpha = 0.1$ and the expected cashflows are expressed in thousands. As seen in Fig. 6.8a, the PPC(4, 2, 2, 1) model gives narrower 90% confidence intervals than the PC(2) model, precisely estimating the uncertainties associated with the expected cashflows. On the other hand, when WT and VH are included for the PPC(8, 2, 2, 1) model in Fig. 6.8b, the confidence intervals widen and approach the PC(2) model's intervals.²⁰ In agreement with our earlier observation, these results suggest that our common factor approach improves the management of bunker price risk for RMI. However, including other fuel data improves neither the reduction nor estimation of the price risk.

¹⁹Replace $C(T_i)$ by $\tilde{C}(T_i)$ for hedged cashflows.

²⁰The PPC(6, 2, 2, 1) model produces a similar figure, whose 90% confidence intervals lie in between the PPC(4, 2, 2, 1) and PPC(8, 2, 2, 1) models' intervals.



(a) Panel 1: PC(2) vs PPC(4, 2, 2, 1)



(b) Panel 3: PC(2) vs PPC(8, 2, 2, 1)

Figure 6.8: 90% confidence intervals of hedged and unhedged cashflows by model.

6.6 Summary

This chapter explored another application of PPCA for the hedging of bunker price risk in the shipping market. The cross hedge aimed to reduce inherent risks associated with the price fluctuation of the Rotterdam low-sulphur bunker fuel (RMI) using the high-sulphur bunker fuel (LW) futures, for which a more liquid and established market was available. From a modelling perspective, we proposed a closed-form formula for the minimum-variance hedge ratio that can accommodate the common factor structures of closely related fuel futures in the expression. In addition, we argued why the calculation of hedge ratios for commodity consumers should be distinguished from a portfolio management approach that typically derives the hedge ratios based on the (log) differences of asset prices. The proposed closed-form formula enables easy computation of minimum-variance hedge ratios for log-normal processes in line with the argument.

We conducted an empirical study to find the best-performing hedge between three sub-groups of marine fuels, comparing the $PC(2)$ and $PPC(2K, 2, 2, 1)$ models for $K = 2, 3, 4$ to analyse the effect of common factor modelling in the cross hedge. The results indicated that the presence of common factors improves the hedge effectiveness and estimation of price risk compared to a non-common factor model. The effectiveness measures suggested that the $PPC(4, 2, 2, 1)$ model, which assumes the existence of common factors between RMI and LW, achieved the most variance reduction. On the other hand, other common factor models, including the US Gulf coast high-sulphur fuel (WT) and West Texas Intermediate crude (VH) futures, did not beat the base-case model. These results hinted the weak association of RMI with these proxy futures at the introductory phase of the IMO 2020 regulation.

Concluding Remarks

Advances in infrastructure and technology have accelerated the cross-market trading of energy commodities in recent years. When markets are affected by the same supply/demand factors, this would strengthen the price dependence and comovement, which requires multi-commodity price modelling. This thesis focused on developing a new factor estimation method for the joint modelling of multi-commodity forward curves.

To begin with, it revisited the recent literature on multi-commodity price modelling by classifying other authors' approaches into three main categories: i) spot price models where one or more state variables are shared among commodities, ii) forward price models including a shared error-correction mechanism in the drift under $\mathbb{P}$, iii) forward price models where the commonality of commodities is embedded in the factor volatilities under $\mathbb{Q}$. The spot price models of i) tend to result in large parameter sets, with the general drawback of the model-implied forward price disagreeing with the market price. On the other hand, the error-correction models of ii) cannot preserve the long-term equilibrium relationship of the underlying assets under the risk-neutral measure. Therefore, our research went in the direction of iii) in conjunction with the conventional risk-neutral pricing theories for practical applicability and derivatives pricing, for which the literature seemed scarce.

We selected PCA on PCA (PPCA) as the main tool to quantify common eigenstructures among several dispersion matrices for commodity term structure modelling, extending the previous work of [Alexander \(2002\)](#). The empirical sections analysed four forward curves of UK and Dutch natural gas and power prices with the proposed factor estimation method. The results highlighted the ability of PPCA to identify common factors when they are likely to exist. When common factors are unlikely to exist, it shows the idiosyncratic patterns of the factors via a distribution-free, data-driven method. We investigated the qualitative aspects of common eigenvectors using a unique ranking algorithm, and the sparse PCA of [Zou et al. \(2006\)](#). The results indicated the dominance of the level and slope factors as the two most important common factors for the European energy forward

curves.

PPCA is a parsimonious yet intuitive and useful tool with much potential in finance applications. As for the price modelling (the ‘PPC models’ in the context of forward curve modelling), we adapted the common factor structures to the volatility functions of [Heath et al. \(1992\)](#) with a few common factors and one idiosyncratic factor per forward curve. The simulation study showed that the PPC model with two or three common factors is sufficient to mimic the marginal distribution of UK and Dutch energy forward curves at a similar level to a benchmark model, while the idiosyncratic factor by local covariance fitting provided desirable outcomes for the spread of commodities to stabilise the fluctuation in the joint distributions.

The joint modelling framework is applied to the pricing of exchange options by [Margrabe \(1978\)](#), in which we found the tendency of the PPC models to lower the option prices compared with a non-common factor model due to the enhanced common factor structures in the model variance. In addition, we derived two variants of the [Margrabe \(1978\)](#) formula when one of the underlying assets is denominated in foreign currency units. The option pricing formula for exchanging foreign to domestic assets has yet to be considered carefully in the previous commodity literature. Hence, we shed a spotlight on the subject, showing the derivation of the FX adjustment terms for both cases: when the FX rate is fixed or stochastic. The impact of the FX factor appeared more significantly for the latter type of options than the former depending on the relative size of the spread variance and FX add-on.

In the last chapter, we demonstrated another usage of PPCA for common factor hedging in maritime finance, deriving a minimum-variance hedge ratio that appropriately reflects the risk profile of a commodity consumption problem. We compared three subgroups of fuels in the empirical analysis to find the most effective cross hedge. Our findings suggested a cross-hedge can be improved by including common factors for most seemingly related fuel products. However, increasing the number of hedge instruments did not do any better than the most simple cross-hedge, in which an underlying fuel contract is hedged by a proxy fuel future, due to the weak association of the underlying and other proxy fuel futures. While our current study looked at a simple one-to-one hedge, exploring this topic in multi-dimensional space would be interesting, assigning appropriate weights to all the proxy fuels included in PPCA to calculate the hedge ratios.

This thesis revealed important findings for the participants of commodity markets, financial institutions, and their regulators since ignoring common factors may cause economic

losses due to mispricing assets or mistreating inherent risks in transactions. In future, one may extend our research by examining the impact of common factors on the pricing of other multi-asset derivatives, such as basket options or fixed income term structure models, where interest rates are exposed to the same macroeconomic shocks. Else, the PPC model could be improved by estimating/calibrating the factor volatilities from/to implied volatilities instead of historical volatilities. However, this approach depends on market development, as many commodities' implied volatilities are still unobservable in today's markets, as is the implied correlation of spread options that would be especially insightful for those models.

Bibliography

- Alexander, C. (2002). Common correlation structures for calibrating the libor model. *ISMA Centre Finance Discussion Paper*, (2002-18).
- Alexander, C. (2008a). *Market risk analysis, practical financial econometrics*, volume 2. John Wiley & Sons.
- Alexander, C. (2008b). *Market risk analysis, quantitative methods in finance*, volume 1. John Wiley & Sons.
- Alexander, C. and Chibumba, A. (1996). Multivariate orthogonal factor garch models. *Discussion Papers in Mathematics. Universitv of Sussex*.
- Alexander, C., Prokopczuk, M., and Sumawong, A. (2013). The (de) merits of minimum-variance hedging: application to the crack spread. *Energy Economics*, 36:698–707.
- Alizadeh, A. H., Kavussanos, M. G., and Menachof, D. A. (2004). Hedging against bunker price fluctuations using petroleum futures contracts: constant versus time-varying hedge ratios. *Applied Economics*, 36(12):1337–1353.
- Benmenzer, G., Gobet, E., and Jérusalem, C. (2007). Arbitrage free cointegrated models in gas and oil future markets. arXiv preprint arXiv:0712.3537.
- Benninga, S., Eldor, R., and Zilcha, I. (1984). The optimal hedge ratio in unbiased futures markets. *Journal of Futures Markets*, 4(2):155–159.
- Benth, F. E. and Koekebakker, S. (2015). Pricing of forwards and other derivatives in cointegrated commodity markets. *Energy Economics*, 52:104–117.
- Bessembinder, H., Coughenour, J. F., Seguin, P. J., and Smoller, M. M. (1995). Mean reversion in equilibrium asset prices: Evidence from the futures term structure. *Journal of Finance*, 50(1):361–375.

- Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3(1-2):167–179.
- Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–659.
- Borovkova, S. and Geman, H. (2008). Forward curve modelling in commodity markets. *Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy*, pages 9–32.
- Bouchouev, I. (2012). Inconvenience yield, or the theory of normal contango. *Quantitative Finance*, 12(12):1773–1777.
- Brennan, M. J. (1958). The supply of storage. *American Economic Review*, 48(1):50–72.
- Carmona, R. and Coulon, M. (2014). A survey of commodity markets and structural models for electricity prices. In *Quantitative Energy Finance*, pages 41–83. Springer.
- Carmona, R., Coulon, M., and Schwarz, D. (2013). Electricity price modeling and asset valuation: a multi-fuel structural approach. *Mathematics and Financial Economics*, 7(2):167–202.
- Carr, P. (1988). The valuation of sequential exchange opportunities. *Journal of Finance*, 43(5):1235–1256.
- Carr, P. (1995). The valuation of american exchange options with application to real options. *Real options in capital investment: models, strategies and applications*.
- Cartea, A. and Figueroa, M. G. (2005). Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, 12(4):313–335.
- Cartea, A. and Williams, T. (2008). UK gas markets: The market price of risk and applications to multiple interruptible supply contracts. *Energy Economics*, 30(3):829–846.
- Casassus, J., Liu, P., and Tang, K. (2013). Economic linkages, relative scarcity, and commodity futures returns. *Review of Financial Studies*, 26(5):1324–1362.
- Chiarella, C., Clewlow, L., Kang, B., et al. (2009). Modelling and estimating the forward price curve in the energy market. *Quantitative Finance Research Centre, University of Technology Sydney. Working Paper*.

- Clewlow, L. and Strickland, C. (2000). *Energy derivatives: pricing and risk management*. Lacima Publ.
- Cortazar, G., Milla, C., and Severino, F. (2008). A multicommodity model of futures prices: Using futures prices of one commodity to estimate the stochastic process of another. *Journal of Futures Markets*, 28(6):537–560.
- Cortazar, G. and Schwartz, E. S. (1994). The valuation of commodity contingent claims. *Journal of Derivatives*, 1(4):27–39.
- Duan, J.-C. and Pliska, S. R. (2004). Option valuation with co-integrated asset prices. *Journal of Economic Dynamics and Control*, 28(4):727–754.
- Ederington, L. H. (1979). The hedging performance of the new futures markets. *Journal of Finance*, 34(1):157–170.
- Edoli, E., Tasinato, D., and Vargiolu, T. (2013). Calibration of a multifactor model for the forward markets of several commodities. *Optimization*, 62(11):1553–1574.
- Eslami, A., Kohler, A., El Qannari, M., and Bougeard, S. (2011). General overview of methods of analysis of multi-group datasets. In *HDSDA*, pages 108–123.
- Fama, E. F. and French, K. R. (1987). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, 60(1).
- Farkas, W., Gourié, E., Huitema, R., and Necula, C. (2017). A two-factor cointegrated commodity price model with an application to spread option pricing. *Journal of Banking & Finance*, 77:249–268.
- Fengler, M. R. (2006). *Semiparametric modeling of implied volatility*. Springer Science & Business Media.
- Feron, O. and Gruet, P. (2020). Estimation of the number of factors in a multi-factorial Heath-Jarrow-Morton model in electricity markets. hal-02880824f.
- Fischer, S. (1978). Call option pricing when the exercise price is uncertain, and the valuation of index bonds. *Journal of Finance*, 33(1):169–176.
- Flury, B. (1988). *Common principal components & related multivariate models*. John Wiley & Sons, Inc.
- Frikha, N. and Lemaire, V. (2013). Joint modelling of gas and electricity spot prices. *Applied Mathematical Finance*, 20(1):69–93.

- Gabillon, J. (1991). *The term structures of oil futures prices*. Oxford institute for energy studies.
- Galai, D. (1979). A proposal for indexes for traded call options. *Journal of Finance*, 34(5):1157–1172.
- Garman, M. B. and Kohlhagen, S. W. (1983). Foreign currency option values. *Journal of International Money and Finance*, 2(3):231–237.
- Geman, H. (2009). *Commodities and commodity derivatives: modeling and pricing for agriculturals, metals and energy*. John Wiley & Sons.
- Geman, H. and Nguyen, V.-N. (2005). Soybean inventory and forward curve dynamics. *Management Science*, 51(7):1076–1091.
- Gibson, R. and Schwartz, E. S. (1990). Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance*, 45(3):959–976.
- Girsanov, I. V. (1960). On transforming a certain class of stochastic processes by absolutely continuous substitution of measures. *Theory of Probability & Its Applications*, 5(3):285–301.
- Haug, E. G. and Haug, J. (2002). Knock-in/out margrabe. *Wilmott Magazine*, December, pages 38–41.
- Heath, D., Jarrow, R., and Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica: Journal of the Econometric Society*, 60:77–105.
- Hicks, J. R. (1946). Value and capital, 1939. *Mathematical Appendix*, pages 311–2.
- Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24(6):417.
- Houle, D., Mezey, J., and Galpern, P. (2002). Interpretation of the results of common principal components analyses. *Evolution*, 56(3):433–440.
- Johnson, L. L. (1960). The theory of hedging and speculation in commodity futures. *Review of economic studies*, 27(3):139–151.
- Jolliffe, I. T. (2011). *Principal component analysis*. Springer-Verlag New York, 2 edition.

- Kaldor, N. (1939). Speculation and economic stability. *Review of Economic Studies*, 7(1):1–27.
- Keynes, J. M. (1930). *A treatise on money: in 2 volumes*. Macmillan & Company.
- Kirk, E. (1995). Correlation in the energy marketss. In *Managing Energy Price Risk*, pages 71–78. Risk Publications and Enron.
- Koekebakker, S. and Ollmar, F. (2005). Forward curve dynamics in the nordic electricity market. *Managerial Finance*, 31(6):73–94.
- Litterman, R. and Scheinkman, J. (1991). Common factors affecting bond returns. *Journal of Fixed Income*, 1(1):54–61.
- Lucia, J. J. and Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, 5(1):5–50.
- Ma, C. K., Mercer, J. M., and Walker, M. A. (1992). Rolling over futures contracts: A note. *Journal of Futures Markets (1986-1998)*, 43(1):203.
- Manoliu, M. and Tompaidis, S. (2002). Energy futures prices: term structure models with kalman filter estimation. *Applied Mathematical Finance*, 9(1):21–43.
- Margrabe, W. (1978). The value of an option to exchange one asset for another. *Journal of Finance*, 33(1):177–186.
- Menachof, D. A. and Dicer, G. N. (2001). Risk management methods for the liner shipping industry: the case of the bunker adjustment factor. *Maritime Policy & Management*, 28(2):141–155.
- Merton, R. C. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science*, pages 141–183.
- Milonas, N. T. (1991). Measuring seasonalities in commodity markets and the half-month effect. *Journal of Futures Markets (1986-1998)*, 11(3):331.
- Murakami, H. (2015). *Introduction to Martingale Approach*. Kindai Kagakusha.
- Myers, R. J. and Thompson, S. R. (1989). Generalized optimal hedge ratio estimation. *American Journal of Agricultural Economics*, 71(4):858–868.
- Nakajima, K. and Ohashi, K. (2012). A cointegrated commodity pricing model. *Journal of Futures Markets*, 32(11):995–1033.

- Ng, V. K. and Pirrong, S. C. (1994). Fundamentals and volatility: Storage, spreads, and the dynamics of metals prices. *Journal of Business*, 67(2):203–230.
- Ohana, S. (2006). *Deux contributions en gestion des risques de matieres premieres*. PhD thesis, Paris 9.
- 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.
- Paschke, R. and Prokopczuk, M. (2009). Integrating multiple commodities in a model of stochastic price dynamics. *Journal of Energy Markets*, 2(3).
- Pearson, K. (1901). Liii. on lines and planes of closest fit to systems of points in space. *London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572.
- Pilipovic, D. (1998). *Energy risk: Valuing and managing energy derivatives*. McGraw-Hill New York.
- Reiner, E. (1992). Quanto mechanics. *Risk*, 5(3):59–63.
- Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *IMR; Industrial Management Review (pre-1986)*, 6(2):41.
- Schwartz, E. and Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911.
- Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52(3):923–973.
- Sørensen, C. (2002). Modeling seasonality in agricultural commodity futures. *Journal of Futures Markets*, 22(5):393–426.
- Stapleton, R. C. and Subrahmanyam, M. G. (1984). The valuation of multivariate contingent claims in discrete time models. *Journal of Finance*, 39(1):207–228.
- Stein, J. L. (1961). The simultaneous determination of spot and futures prices. *American Economic Review*, 51(5):1012–1025.
- Swindle, G. (2014). *Valuation and Risk Management in Energy Markets*. Cambridge University Press.

- Telser, L. G. (1958). Futures trading and the storage of cotton and wheat. *Journal of Political Economy*, 66(3):233–255.
- Thompson, M. (2016). Natural gas storage valuation, optimization, market and credit risk management. *Journal of Commodity Markets*, 2(1):26–44.
- Tolmasky, C. and Hindanov, D. (2002). Principal components analysis for correlated curves and seasonal commodities: The case of the petroleum market. *Journal of Futures Markets: Futures, Options, and Other Derivative Products*, 22(11):1019–1035.
- Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188.
- Wang, Y., Meng, Q., and Tan, Z. (2018). Short-term liner shipping bunker procurement with swap contracts. *Maritime Policy & Management*, 45(2):211–238.
- Working, H. (1949). The theory of price of storage. *American Economic Review*, 39(6):1254–1262.
- Zou, H. and Hastie, T. (2020). Elastic-net for sparse estimation and sparse pca. *URL* <http://www.stat.umn.edu/~hzou>.
- Zou, H., Hastie, T., and Tibshirani, R. (2006). Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286.

Appendix A

Tables for Chapter 5

Table A.1: Option prices of spark spreads.

(Full reports of Table 5.2)

Contracts	30d	60d	90d	120d	150d	180d	210d	270d	360d	540d
Panel A: Up-Ug (GBP/MWh)										
PC	16.814	16.844	16.866	16.882	16.893	16.903	16.911	16.927	16.949	16.998
PPC	16.791	16.792	16.793	16.795	16.796	16.798	16.800	16.803	16.807	16.819
diff	-0.023	-0.052	-0.073	-0.087	-0.097	-0.105	-0.112	-0.124	-0.142	-0.179
diff %	-0.14	-0.31	-0.43	-0.52	-0.57	-0.62	-0.66	-0.73	-0.84	-1.05
Panel B: Dp-Dg (EUR/MWh)										
PC	7.753	7.916	8.021	8.110	8.189	8.262	8.332	8.470	8.675	9.068
PPC	7.518	7.560	7.614	7.672	7.732	7.794	7.857	7.984	8.175	8.552
diff	-0.235	-0.357	-0.408	-0.439	-0.457	-0.468	-0.476	-0.487	-0.499	-0.517
diff %	-3.03	-4.51	-5.08	-5.41	-5.58	-5.67	-5.71	-5.75	-5.76	-5.70

The prices of UK and Dutch spark spread options that depict Fig. 5.1. The relative difference is given by the formula: $\text{diff \%} = 100 \times \frac{(\text{PPC price} - \text{PC price})}{\text{PC price}}$.

Table A.2: Standard deviations of spark spreads at option expiration.

(Full reports of Table 5.3)

Contracts	30d	60d	90d	120d	150d	180d	210d	270d	360d	540d
Panel A: Up-Ug										
PC	0.1478	0.1659	0.1750	0.1805	0.1843	0.1871	0.1895	0.1937	0.1993	0.2100
PPC	0.0832	0.1049	0.1155	0.1216	0.1259	0.1293	0.1321	0.1365	0.1421	0.1522
Panel B: Dp-Dg										
PC	0.1189	0.1374	0.1477	0.1558	0.1626	0.1687	0.1744	0.1852	0.2004	0.2278
PPC	0.0637	0.0845	0.0975	0.1076	0.1162	0.1240	0.1311	0.1441	0.1615	0.1913

The model standard deviations are given by Eq. (5.21). The options are assumed to expire five days before the maturities of the underlying forward contracts.

Table A.3: The option prices of 360-day spark spread.

(Full reports of Table 5.5)

Contracts	30d	60d	90d	120d	150d	180d	210d	270d	360d	540d
Panel A: Up-Ug (GBP/MWh)										
PC	3.117	3.530	3.738	3.864	3.949	4.015	4.070	4.164	4.292	4.536
PPC	1.640	2.137	2.378	2.517	2.616	2.694	2.757	2.858	2.985	3.216
diff	-1.477	-1.393	-1.360	-1.347	-1.333	-1.321	-1.313	-1.305	-1.307	-1.321
diff %	-47.38	-39.47	-36.39	-34.85	-33.76	-32.91	-32.26	-31.35	-30.45	-29.11
Panel B: Dp-Dg (EUR/MWh)										
PC	1.157	1.501	1.697	1.853	1.986	2.104	2.216	2.429	2.731	3.283
PPC	0.270	0.568	0.780	0.954	1.108	1.250	1.382	1.628	1.962	2.551
diff	-0.886	-0.933	-0.917	-0.899	-0.877	-0.855	-0.834	-0.801	-0.769	-0.733
diff %	-76.64	-62.18	-54.05	-48.52	-44.19	-40.61	-37.62	-32.97	-28.14	-22.32

Re-calculated option prices with 35% and 40% fuel efficiency rates, respectively for the UK and Dutch spark spreads.

Table A.4: Standard deviations of UK – Dutch spreads at option expiration.
(Full reports of Table 5.6)

Contracts	30d	60d	90d	120d	150d	180d	210d	270d	360d	540d
Panel C: Ug-Dg										
PC	0.0632	0.0751	0.0811	0.0857	0.0897	0.0932	0.0964	0.1024	0.1108	0.1265
PPC	0.0427	0.0566	0.0645	0.0701	0.0744	0.0779	0.0808	0.0858	0.0921	0.1031
Panel D: Up-Dp										
PC	0.1715	0.1916	0.2013	0.2075	0.2121	0.2163	0.2205	0.2291	0.2419	0.2661
PPC	0.0973	0.1237	0.1362	0.1433	0.1486	0.1532	0.1576	0.1659	0.1778	0.1995
Panel E: Up-Dg										
PC	0.1515	0.1700	0.1792	0.1847	0.1883	0.1912	0.1938	0.1987	0.2058	0.2195
PPC	0.0829	0.1078	0.1196	0.1258	0.1299	0.1329	0.1355	0.1399	0.1456	0.1562

The model standard deviations are given by Eq. (5.21). The options are assumed to expire five days before the maturities of the underlying forward contracts.

Table A.5: Option prices of quanto exchange options.
(Full reports of Table 5.7)

Contracts	30d	60d	90d	120d	150d	180d	210d	270d	360d	540d
Panel C: Ug-Dg (EUR/MWh)										
PC	1.051	1.143	1.191	1.229	1.262	1.291	1.317	1.367	1.438	1.571
PPC	0.906	1.001	1.061	1.104	1.138	1.166	1.189	1.230	1.281	1.372
diff	-0.144	-0.142	-0.130	-0.125	-0.124	-0.125	-0.128	-0.137	-0.157	-0.199
diff %	-13.75	-12.40	-10.94	-10.18	-9.82	-9.69	-9.70	-10.04	-10.90	-12.67
Panel D: Up-Dp (EUR/MWh)										
PC	13.382	13.596	13.711	13.787	13.846	13.902	13.958	14.077	14.263	14.635
PPC	12.960	13.038	13.102	13.147	13.184	13.219	13.254	13.326	13.442	13.686
diff	-0.421	-0.558	-0.609	-0.640	-0.662	-0.683	-0.704	-0.751	-0.821	-0.949
diff %	-3.15	-4.10	-4.44	-4.64	-4.78	-4.91	-5.05	-5.33	-5.76	-6.48
Panel E: Up-Dg (EUR/MWh)										
PC	10.162	10.429	10.572	10.659	10.719	10.766	10.810	10.893	11.017	11.260
PPC	9.534	9.679	9.784	9.847	9.891	9.925	9.955	10.008	10.081	10.224
diff	-0.628	-0.750	-0.789	-0.813	-0.828	-0.841	-0.855	-0.885	-0.936	-1.036
diff %	-6.18	-7.19	-7.46	-7.62	-7.72	-7.81	-7.91	-8.13	-8.49	-9.20

The price of quanto option using the quanto factor of 1 GBP=1.1226 and the correlation coefficient of $\rho_{1X} = -0.2$ for UK gas and power. The fuel efficiency is 40% for the UK power – Dutch gas spark spread. The relative difference is given by the formula: $\text{diff \%} = 100 \times \frac{(\text{PPC price} - \text{PC price})}{\text{PC price}}$.

Table A.6: The variance of cross-currency spreads at option expiration.
(Full reports of Table 5.9)

Contracts	30d	60d	90d	120d	150d	180d	210d	270d	360d	540d
Panel C: Ug-Dg										
PC										
Total var	0.0036	0.0048	0.0052	0.0056	0.0056	0.0057	0.0061	0.0065	0.0068	0.0081
Spread var	0.0040	0.0056	0.0066	0.0073	0.0080	0.0087	0.0093	0.0105	0.0123	0.0160
FX add-on	-0.0004	-0.0009	-0.0014	-0.0018	-0.0024	-0.0030	-0.0031	-0.0040	-0.0055	-0.0079
PPC										
Total var	0.0014	0.0024	0.0029	0.0032	0.0032	0.0032	0.0035	0.0035	0.0032	0.0031
Spread var	0.0018	0.0032	0.0042	0.0049	0.0055	0.0061	0.0065	0.0074	0.0085	0.0106
FX add-on	-0.0004	-0.0008	-0.0013	-0.0017	-0.0023	-0.0028	-0.0030	-0.0038	-0.0052	-0.0075
Panel D: Up-Dp										
PC										
Total var	0.0296	0.0370	0.0402	0.0427	0.0445	0.0460	0.0471	0.0498	0.0551	0.0680
Spread var	0.0294	0.0367	0.0405	0.0430	0.0450	0.0468	0.0486	0.0525	0.0585	0.0708
FX add-on	0.0002	0.0003	-0.0003	-0.0003	-0.0005	-0.0008	-0.0015	-0.0026	-0.0034	-0.0028
PPC										
Total var	0.0097	0.0157	0.0185	0.0206	0.0220	0.0232	0.0240	0.0258	0.0293	0.0382
Spread var	0.0095	0.0153	0.0185	0.0205	0.0221	0.0235	0.0248	0.0275	0.0316	0.0398
FX add-on	0.0002	0.0004	-0.0000	0.0000	-0.0001	-0.0003	-0.0008	-0.0018	-0.0023	-0.0016
Panel E: Up-Dg										
PC										
Total var	0.0227	0.0287	0.0311	0.0332	0.0343	0.0350	0.0356	0.0364	0.0376	0.0426
Spread var	0.0230	0.0289	0.0321	0.0341	0.0355	0.0365	0.0375	0.0395	0.0424	0.0482
FX add-on	-0.0003	-0.0003	-0.0010	-0.0009	-0.0011	-0.0016	-0.0020	-0.0031	-0.0047	-0.0056
PPC										
Total var	0.0067	0.0115	0.0135	0.0152	0.0160	0.0164	0.0168	0.0171	0.0173	0.0197
Spread var	0.0069	0.0116	0.0143	0.0158	0.0169	0.0177	0.0184	0.0196	0.0212	0.0244
FX add-on	-0.0002	-0.0002	-0.0008	-0.0006	-0.0008	-0.0012	-0.0016	-0.0025	-0.0039	-0.0047

Total var = Spread var + FX add-on. *Total var* is the variance that enters into Eq. (5.44), *Spread var* is the variance of an ordinary exchange option, and *FX add-on* accounts for fluctuations in spot FX rates.

Table A.7: The prices of cross-currency exchange options.
(Full reports of Table 5.10)

Contracts	30d	60d	90d	120d	150d	180d	210d	270d	360d	540d
Panel C: Ug-Dg (EUR/MWh)										
PC	1.139	1.208	1.229	1.249	1.251	1.256	1.278	1.293	1.309	1.370
PPC	1.008	1.071	1.098	1.121	1.121	1.121	1.138	1.138	1.120	1.111
diff	-0.131	-0.137	-0.130	-0.128	-0.131	-0.135	-0.139	-0.155	-0.189	-0.259
diff %	-11.54	-11.34	-10.61	-10.24	-10.45	-10.78	-10.90	-12.00	-14.42	-18.90
Panel D: Up-Dp (EUR/MWh)										
PC	13.783	13.991	14.086	14.158	14.211	14.256	14.290	14.371	14.530	14.914
PPC	13.392	13.466	13.517	13.560	13.592	13.621	13.639	13.683	13.774	14.027
diff	-0.391	-0.525	-0.568	-0.598	-0.618	-0.636	-0.651	-0.688	-0.755	-0.887
diff %	-2.84	-3.75	-4.04	-4.23	-4.35	-4.46	-4.55	-4.79	-5.20	-5.95
Panel E: Up-Dg (EUR/MWh)										
PC	10.527	10.785	10.890	10.982	11.028	11.056	11.082	11.118	11.169	11.376
PPC	9.955	10.082	10.155	10.219	10.251	10.267	10.281	10.292	10.300	10.402
diff	-0.573	-0.702	-0.736	-0.763	-0.778	-0.789	-0.801	-0.826	-0.869	-0.974
diff %	-5.44	-6.51	-6.76	-6.95	-7.05	-7.14	-7.23	-7.43	-7.78	-8.56

The prices of cross-currency exchange options that depict Fig. 5.5. The relative difference is given by the formula: $\text{diff \%} = 100 \times \frac{(\text{PPC price} - \text{PC price})}{\text{PC price}}$.