

ULRR

A market consistent gas storage modelling framework: valuation, calibration, & model risk

Item Type	Thesis
Authors	Kiely, Greg
Download date	2026-03-12 04:23:02
Item License	https://creativecommons.org/licenses/by-nc-sa/1.0/
Link to Item	https://hdl.handle.net/10344/5202

A Market Consistent Gas Storage Modelling Framework: Valuation, Calibration, & Model Risk

by

Greg Kiely

B.Sc, National University of Ireland, Galway, (1st Hons.)

M.Sc University of Limerick, (1st Hons.)

Submitted to the Department of Accounting and Finance, Kemmy Business School, University of Limerick in December 2015, in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Abstract

A typical natural gas derivatives book within an energy trading business, bank, or even large utility will generally be exposed to two broad categories of market risk. The first being outright price volatility, where contracts such as caps/floors, options and swing, will have a non-linear exposure to the variability of the gas price level. The second, although equally as prominent, is time-spread volatility where gas storage, take-or-pay contracts, and calendar spread options will be exposed to the realized variability of different time-spreads. Developing a market consistent valuation framework capable of capturing both risk exposures, and thus allowing for risk diversification within a natural gas trading book, is the primary goal of this thesis. To accomplish this, we present a valuation methodology which is capable of pricing the two most actively traded natural gas derivative contracts, namely monthly options and storage, in a consistent manner. The valuation of the former is of course trivial as the prices are set by the market, therefore the primary focus of this thesis is in obtaining market-based pricing measures for the purpose of storage valuation. A consistent pricing and risk management framework will, by definition, accurately reflect the cost of hedging both outright and spread volatility and thus our work can be viewed as a basis capable of incorporating the other less actively traded contracts listed above. Further, we develop a

methodology for estimating the model risk for general energy derivative pricing models. Such an analysis is a necessary pre-requisite to a model being used to manage the risk associated with a derivatives business.

We begin by introducing the modeling framework and valuation methodology used throughout this thesis. Whilst there is general agreement on the main drivers of storage value, namely the volatility and correlation of the forward curve, there is no industry standard approach to price modeling. Models proposed in the literature to date focus heavily on replicating the statistical properties of the gas price and fail to address both the desire of storage traders to monetize their gamma exposure through hedging in the vanilla options market and also the constraint that they mark these products to observable volatilities. The primary goal of this work is to demonstrate how one can attain general consistency with the natural gas options market using Lévy-driven Ornstein-Uhlenbeck (OU) processes rather than traditional Gaussian models, and also analyze the impact of model choice on storage value.

We provide a forward curve consistent Fourier-based pricing and calibration tool-kit which relies only upon knowledge of the conditional characteristic function of the underlying Lévy driven OU-process. Analytical solutions for such characteristic functions are generally not known and to date the only non-trivial example in the literature is the jump diffusion of Deng (2000). We derive a solution for the mean-reverting Variance Gamma process and demonstrate its effectiveness in both modeling the implied volatility smile and term structure present in the natural gas options market. One of the main benefits of choosing this process as the source of randomness in our modeling framework is parsimony. The model contains just a single parameter more than the more common mean-reverting diffusion model, and this parameter is entirely responsible for matching the implied volatility smile.

We next move on to extending these results to a multidimensional setting in order to further capture the rich dynamics of the underlying forward curve whilst still maintaining consistency with the options market. The proposed framework is thus a specific case of the general Cheyette, Cheyette (2001), model class uniquely specified to meet the needs of a storage trader. Whereas the single factor modelling framework relies entirely upon the implied volatility surface to determine the level of extrinsic value accruing to a storage position, this more general family of models allows one to utilize other sources of information to estimate the value. We demonstrate how a traditional PCA based analysis can be used in informing model specification and provide several examples of such. We extend the Fourier based pricing and calibration methods developed for the single factor models to a multidimensional setting, and for the latter derive an efficient implied moment based calibration routine which is independent of the number of factors in the underlying model. We go on to provide numerical examples of storage valuations under a range of multifactor model specifications and also, analysis on how the model implied forward curve dynamics impact the value of storage.

We finish by providing a storage model risk framework, with an emphasis on parameter risk, which

will aid in analyzing the risk inherent in adopting this innovative market consistent modeling framework. The proposed approach extends the current model risk literature by providing a methodology for incorporating both market based model calibration and statistical parameter estimation in a consistent manner. The potential benefit of such a framework will impact equally trading, risk and regulatory stakeholders within a storage business, from model validation through to deriving appropriate bid-offer levels. We provide detailed numerical examples, based upon the models specified previously, demonstrating how model prices can be adjusted to incorporate model risk and also how different models can be ranked depending upon the model risk implicit in their estimation.

Acknowledgment

I am very thankful to both my supervisors, Dr. Bernard Murphy and Dr. Mark Cummins, for their encouragement, guidance and support throughout. I am particularly grateful for their flexibility in making themselves available around my work schedule and for making the process of completing this research part-time as smooth and enjoyable as possible.

I would also like to thank both my parents, Dirk and Dolores Kiely, for their encouragement when I decided to embark upon my doctoral studies. Finally, I would like to thank my partner Pooja Ahluwalia for her support and advice (and patience!) throughout my research and particularly so when writing this thesis.

Contents

1	Introduction	10
2	Single Factor Storage Valuation	28
2.1	Lévy Driven Ornstein-Uhlenbeck Processes	28
2.2	Fourier Methods in Pricing Early Exercise Claims	36
2.3	Gas Storage Valuation	38
2.4	Empirical Analysis & Results	42
3	MultiFactor Storage Valuation	57
3.1	Markov Lévy Driven Forward Curve Models	57
3.2	Joint Model Calibration and Estimation	66
3.3	Multidimensional Valuation Algorithm	79
3.4	Empirical Analysis & Results	83
4	Storage Model Risk	90
4.1	Model Uncertainty and Parameter Risk	90
4.2	Storage Parameter Risk: Numerical Examples	101
5	Conclusion	118
	Bibliography	122
A	Chapter 2	132
A.1	Parseval's Theorem	132
A.2	Derivation of the Characteristic Function of the Mean Reverting Variance Gamma Process	134
A.3	MRVG Distribution Moments	134
A.4	Markov Implied Spot Price Model	137
A.5	Swaption Pricing Algorithm	140
A.6	Valuation Algorithm Error Analysis	141
A.7	Initial Storage Value Function	142

B Chapter 3	144
B.1 Forward Curve Consistent State Variable Dynamics	144
B.2 Derivation of Characteristic Function of the MRVG-3 Model	148
B.3 Derivation of the Characteristic Function of MRVG-2 Model	150
B.4 Distribution Moments	151
B.5 Log Forward Price	155
B.6 Forward Curve Covariance Function	158
B.7 Multi-Factor Swaption Pricing Algorithm	159
B.8 Moment Formulae	160
B.9 Valuation Convergence	163
C Chapter 4	165
C.1 Sample Parameter Covariance Matrix	165

List of Tables

2.4.1 MRVG & MRD Model Parameters	47
2.4.2 MRJD Model Parameters	47
2.4.3 Single Factor Valuation Results	49
2.4.4 MRD Convergence and Processing Times	50
2.4.5 MRJD Convergence and Processing Times	50
2.4.6 MRVG Convergence and Processing Times	51
2.4.7 Delta Results	51
2.4.8 Scenario 1: Parallel Shift	53
2.4.9 Scenario 2: Volatility Slope Shift	53
2.4.10 Scenario 3: Parallel Smile Shift	54
3.2.1 MRVG Model Parameters	73
3.2.2 MRVG-2x Model Parameters	74
3.2.3 MRVG-2 Model Historically Estimated Parameters	75
3.2.4 MRVG-2 Calibrated Parameters	76
3.2.5 MRVG-3x Historically Estimated Parameters	76
3.2.6 MRVG-3x Calibrated Parameters.	77
3.4.1 Two Factor Valuation Results	85
4.2.1 MRVG Model Parameters.	106
4.2.2 MRVG Storage Value Push Forward Distribution	107
4.2.3 MRJD Model Parameters	108
4.2.4 MRJD Storage Value Push Forward Distribution	109
4.2.5 MRVG-3x Model Parameters	109
4.2.6 MRVG-3x Conditional Storage Value Push Forward Distribution	110
4.2.7 MRVG-3x Parameter Deltas	111
4.2.8 MRVG-3x Calibration and Parameter Estimation Risk Push Forward Distribution.	113
B.9.1 Two Factor Convergence Results	164

List of Figures

2.4.1 NBP Forward Curve	44
2.4.2 NBP Volatility Surface	45
2.4.3 MRVG Calibration	46
2.4.4 MRJD Calibration	46
2.4.5 MRD Calibration	47
2.4.6 Scenario 1: Parallel shift at both maturities	52
2.4.7 Scenario 2: Shift at 6m maturity.	53
2.4.8 Scenario 3: Parallel shift of OTM strikes at both maturities.	54
3.1.1 NBP Gas Forward Curve Returns Correlation	61
3.1.2 Historical Volatility Function values.	62
3.2.1 Second Derivative of Moment Payoff Function	69
3.2.2 High Strike Variance Slope	70
3.2.3 NBP Vol Surface: 19th December 2012	72
3.2.4 MRVG-2 Calibrated 2 nd Volatility Function	75
3.2.5 MRVG-3x Calibrated 2 nd Volatility Function	77
3.2.6 MRVG-3x Distribution Moments.	78
3.2.7 MRVG-2x Distribution Moments.	78
3.4.1 NBP Forward Curve	84
3.4.2 MRVG-2x Initial Value Grid	86
3.4.3 MRVG-2x Implied Forward Curves	86
3.4.4 MRVG-3x Initial Value Grid	87
3.4.5 MRVG-3x Implied Forward Curves	88
4.2.1 MRVG-3x 2 nd Factor Volatility Function	104
4.2.2 MRVG Storage Value Push Forward Density	107
4.2.3 MRJD Storage Value Push Forward Density	108
4.2.4 MRVG-3x Calibration Error Storage Value Density	110
4.2.5 MRVG-3x Parameter Estimation Risk Storage Value Density	112
4.2.6 MRVG Storage Value Cumulative Distribution	114

4.2.7 MRJD Storage Value Cumulative Distribution	114
4.2.8 MRVG-3x Storage Value Cumulative Distribution	115
4.2.9 MRVG Storage Risk Adjusted Offers	116
4.2.10MRVG-3x Storage Risk Adjusted Bids	117
A.7.1Log Storage Value Function	143

Chapter 1

Introduction

Trading in natural gas has increased significantly across Europe in the past decade driven primarily by the development of liquid centralized trading hubs. These hubs act as a single trading point for the transfer of gas entering and exiting their respective market areas. The first and, until very recently, the most liquid of these trading hubs has been the National Balancing Point (NBP) in the United Kingdom. This has been followed by the creation of the Dutch Title Transfer Facility (TTF) and the German NCG and Gaspool hubs. In recent years trading hubs have been created in many other European countries although they still lag behind the more developed markets in terms of liquidity. A natural consequence of the existence of a market for spot gas trading was the development of a derivative market to manage future price risk. Forwards and Futures contracts delivering into the NBP market have been traded on an Over-the-Counter (OTC) basis since its creation, more recently exchange traded futures have grown significantly as a percentage of total NBP trades. As of 2012, trading on the Intercontinental Exchange's (ICE) NBP futures contracts have accounted for about one third of total NBP trading volume Heather (2012).

Forward contracts on natural gas are generally categorized into “prompt” and “curve” contracts. Prompt contracts refer to future deliveries within the current calendar month. Examples of these include, the Day-Ahead (DA) contract which delivers a flat hourly quantity of gas for each hour in the next business day and the Balance of Month (BOM) contract which delivers a flat hourly quantity of gas for each day remaining in the current calendar month. Curve contracts refer to gas deliveries occurring at any point beyond the end of the current calendar month. The majority of trading occurs on the Month Ahead, Front Quarter, and Front Season contracts. The Month Ahead refers to gas deliveries occurring each hour over the calendar month immediately following the current month. The Front Quarter contract refers to the calendar year quarter immediately following the current quarter, so for example, on the 1st November the Front Quarter would refer to the period 1st January to 31st March. The gas year is divided into two seasons, Winter which runs from October to March and Summer which runs from April to September. The Front Season contract on the 1st July would therefore refer to the period 1st October to 31st March. In tandem to the development of the forward market, trading in vanilla options has also

increased significantly, although predominantly on an OTC basis. These options typically settle against forward contracts, with options on monthly forward contracts generally being the most liquid products.

In addition to the vanilla options market, there exists an active market for structured gas derivatives traded predominantly on a bilateral or OTC basis. The characteristics of these contracts typically reflect the underlying dynamics of gas demand. For example, variable volume or swing contracts allow utilities/consumers to buy a certain volume of gas over a given delivery period which would be determined by their expected demand for gas. These contracts provide the additional flexibility to take more or less gas on a daily basis dependent upon the variability of realized demand versus forecast. Similarly, gas storage has become an increasingly prominent feature of the European gas market. Given the temperature dependency of natural gas heating demand, there is generally a large discrepancy between winter and summer demand and storage plays a crucial part in smoothing out seasonal supply-demand imbalances as well as managing the risk of supply disruption, and increasing market liquidity Le Fevre (2013); EFET (2009). Current European storage capacity stands at 101bcm, roughly one fifth of annual demand Hureau (2015). Allied to this is a quickly developing secondary market for storage products¹ which allows market players to adjust their seasonal flexibility to match portfolio growth. This additional liquidity has led to tighter bid-ask spreads and thus less room for error when modeling the value of these products. Storage contracts can refer to both capacity in physical storage units or “virtual” storage capacity, typically sold as a simplified tranche of physical storage. For virtual contracts the main parameters of the deal will be the total amount of gas that can be stored, generally referred to as the Working Gas Volume (WGV), and the injection/withdrawal rates and associated costs. The latter determine the number of days it takes to fill/empty the storage and thus the flexibility of the asset. Physical storage contracts would generally have more constraints on injection and withdrawal rates which would reflect the operating characteristics of the storage asset.

From a commercial point of view, storage assets allow a trader to buy gas to inject into storage during summer months when it is relatively cheap and withdraw and sell the gas the following winter when prices are higher, thereby collecting the price difference as profit. Given a liquid forward market, traders have the ability to lock in a base or intrinsic value of the storage asset by locking in prices for future traded volumes. As such, organizations considering investment in a storage facility are aware, to a high degree of certainty, what return they will receive. This valuation and trading approach is referred to as Intrinsic Valuation and, as is the case for standard financial products, gives a lower bound on the storage value. Although such a method for valuing and operating the storage facility has the obvious benefit of eliminating market risk it does so at a cost, that is, reducing completely the flexibility of the trader to adjust their planned injection/withdrawal schedule to respond to favourable market conditions. Of course the prices at which storage capacity is offered in the market will typically exceed the intrinsic value and

¹Although the majority is traded Over-The-Counter there are examples of exchange traded products, for example the recently launched Bergermeer contract on ICE.

thus traders require a method to determine what extrinsic value can be extracted from the asset through dynamically trading in the underlying forward and options market. For this reason, it is much more common to augment the intrinsic valuation approach with more advanced techniques.

The Intrinsic Basket of Spread Options method² captures some of the flexibility of the facility by equating its value to that of a, generally hypothetical, portfolio of calendar spread options rather than simply forward spreads, thus allowing the operator to monetize some of the optionality inherent in the storage asset. Instead of buying forward contracts on summer months and selling forward contracts on winter months, the trader would sell a spread option on winter and summer forward prices. If the option is exercised the trader can recoup any loss by immediately trading in the underlying forward contracts and using the storage to monetize the spread value. Operationally this method suffers from the lack of a liquid calendar spread options market and thus traders resort to replicating the portfolio by delta hedging with the underlying forwards, in which case the strategy is no longer risk free. Many extend the idea of replication of the storage unit via financial contracts using dynamic rather than static hedging strategies, one example being the Rolling Intrinsic strategy which calculates the intrinsic value every day over the life of the storage contract using Monte Carlo simulation. The portfolio is rebalanced each day if optimal to do so, and as is the case with the above methods, the injection/withdrawal schedule is determined by the financial positions. Similarly, one can value a rolling basket of spread options in order to further incorporate the flexibility of the facility. The most obvious difference between this rolling valuation method and the simple intrinsic approach is that it is model dependent and thus the obtained value is subject to model risk.

While all these methods have the obvious appeal of being hedge-based and therefore risk minimizing or even risk free, they fail to capture fully the optionality of the storage. Spot based optimization techniques apply a real options approach to derive the optimal spot market trading strategy which captures fully the operational flexibility and thus the full value of the facility. As such the value derived in this way dominates those of the aforementioned methods at the cost of greater variance in realized profit, although this can be reduced significantly by incorporating a delta hedging strategy. Similar to standard financial options, the motivation for such a hedging strategy is based upon replicating the value changes by continuously trading in the primary risk factors which affect the storage value, namely, the market forward curve. In theory the cumulative expected return and variance of a portfolio consisting of a storage asset and a delta hedging portfolio is zero. Of course, in reality continuous trading is impossible and therefore a hedging error is introduced where the variance of the error is primarily a function of the time between hedge rebalances.

Typically, spot optimization methods in gas storage valuation have utilized either the trinomial tree based “forest of trees” approach or Monte Carlo simulation. Manoliu (2004) introduces the concept of us-

²For a detailed overview of the standard approaches to storage valuation we direct the interested reader to Breslin et al. (2008).

ing layers of trinomial tree representations of the spot price evolution to solve for the storage value. Each layer represents a different potential storage level and thus the value can be derived by traversing backwards through time and evaluating the optimal storage level at each time-step for each spot price. This backward stochastic optimization approach is fundamental to all spot optimization methods. Boogert and De Jong (2008) utilize Least Square Monte Carlo techniques to derive the optimal control policy and storage value. By choosing appropriate basis functions at each time step they estimate the inventory dependent continuation value via regression on the future values. Due to the fact that it is simulation based, this method has the obvious benefit of being able to incorporate a wide range of spot price models and exotic constraints. However, the method can be quite slow while the choice of basis functions is somewhat subjective and thus cannot guarantee convergence to the true value.

Chen and Forsyth (2009) characterize the storage valuation as a stochastic control problem where the value function is the solution to the well known Hamilton-Jacobi-Bellman (HJB) equation. They apply Semi-Lagrangian time stepping methods (of which the “forest of trees” approach is a special case) to solve for the optimal control policy and storage asset value and prove convergence of their method to the viscosity solution of the HJB equation. This approach has the benefit of greater accuracy than the Monte Carlo method and although it is capable of incorporating jump processes in the spot price model it does suffer from the inability to extend the method easily to a multidimensional setting and handle a richer set of potential models. Felix and Weber (2012) propose a novel approach which combines simulated price paths and lattice based dynamic programming. They numerically construct recombining trees using the simulated paths and then solve for the storage value in the usual manner.

The spot optimization approach is applicable to all standard storage deals where the injection and withdrawal capacity is a function of the gas in store and/or time only. Certain storage contracts, which typically relate to a segment of a physical storage asset, constrain the injection and withdrawal capabilities of the owner based upon the level of gas in the storage asset. In such a scenario, the flexibility of the contract will be dependent upon the actions of other claimants on the physical storage capacity. This multi-agent optimization problem has been studied extensively by Holland and Walsh (2013) where the authors apply a game theoretic approach to the valuation problem. Other notable recent studies that examine gas storage valuation and modelling include Carmona and Ludkovski (2010), Lai et al. (2010), and Thompson et al. (2009). In this thesis we will focus solely on the spot optimization of standard storage contracts. In doing so we will utilize a range of Lévy driven models of the spot price specifically designed to allow for accurate calibration to market data and thus provide a link between storage value and the vanilla options market. These models will draw upon the rich literature of energy price modelling in capturing the statistical dynamics of energy market prices.

Traditionally, spot price modeling in commodity markets has adapted methods first applied in the interest rate market. Much of the early work in this area focused primarily on capturing the mean reverting nature of energy prices. Schwartz (1997) introduces a mean reverting diffusion model for the spot price and adds a second mean reverting stochastic factor, representing the “convenience yield”, which aids

in matching the observed forward curve for a range of commodities. Clewlow and Strickland (1999b) provide an alternative form of the mean reverting diffusion model which takes the current forward curve as an input to the spot price model. The authors go on to show how this model can be used to price a range of contingent claims including path-dependent options. These early models assumed normally distributed innovations and thus failed to capture the leptokurtotic nature of the statistical price returns. Attempts at overcoming this deficiency have relied heavily on the work of Barndorff-Nielsen and Shephard (2001) who introduced, in a financial context, Lévy driven Ornstein–Uhlenbeck processes, a family of semi-martingales which are both mean reverting and possess heavy tailed distributions. The primary focus of the authors is in modeling stochastic volatility which is traditionally characterized as mean reverting. They demonstrate how processes possessing this property can be constructed using general Lévy processes as the driving stochastic component.

One of the more notable applications of both Lévy driven Ornstein–Uhlenbeck and Lévy processes in the energy markets is given by Deng (2000). In this paper the author derives a mean reverting jump diffusion model with jumps following a double exponential distribution and provides an analytical solution for the characteristic function of the model. Benth et al. (2007) also model electricity prices in this manner and discusses derivative pricing under the assumption of an arithmetic, rather than the more common geometric, price process. The authors use a mean reverting Gamma process for modelling price spikes which only allows for positive spikes. This model also yields an analytical characteristic function although the assumption of strictly positive jumps is not entirely representative of the price dynamics in energy markets. As noted in Kjaer (2008), in general, the characteristic functions of mean reverting Lévy driven models are not available in analytical form and need to be calculated numerically. Thus, the model proposed by Deng (2000) is to date a special case of this model class which allows for both positive and negative jumps and has an analytical solution for the characteristic function. The importance of the having an analytical solution for the characteristic function is manifest by the fact that it greatly enhances the prospect of a model being used in practice as it allows one to utilize computationally efficient transform based methods in pricing.

Notable recent additions to the energy spot price modeling literature include Barndorff-Nielsen et al. (2010) who introduce the concept of Lévy semi-stationary processes which are capable of incorporating jumps and stochastic volatility. They also discuss a more general form for the auto-correlation structure of the spot price returns which encompasses the negative exponential associated with mean reverting processes. This latter aspect of their approach is interesting as it opens up more flexible forward curve volatility term structure forms albeit at the cost of Markovian single factor dynamics. Jaimungal and Surkov (2011) study a number of general mean reverting Lévy models with a focus on incorporating jumps in the spot price evolution. They go on to provide option valuation methods using the spot price models specified which utilize the analytically tractable form of the characteristic function given by Deng (2000). Luciano (2009), utilizes time-changed Brownian motions with Lévy subordinators in the model-

ing of the spark spread and demonstrates the ability of these models to provide a better fit to power and gas market option prices over a simple Gaussian model. Pinho and Madaleno (2011) comprehensively demonstrate the improved fit to the volatility smile present in electricity option markets for a number of Lévy processes relative to the Black Scholes model.

Li and Linetsky (2014) study Ornstein-Uhlenbeck processes subordinated by Lévy processes for the purpose of modelling commodity spot and futures prices. This class of model, which they term SubOU, allows one to incorporate jumps and stochastic volatility into single factor spot processes. The authors give general vanilla option pricing methods using eigenfunction expansions and calibrate an example model driven by an Inverse Gaussian subordinator to a number of commodity option volatility smiles. Their results demonstrate the ability of models of this class to match market option prices well within bid-offer. In Chapter 2 we will discuss the use of single factor Lévy driven mean reverting spot price models in the context of storage valuations. Specifically, we will demonstrate how one such model, the mean reverting Variance Gamma process, allows one to capture the observable dynamics of energy spot prices as well as those implied by the options market. A version of this model not specified to be consistent with the market forward curve has been studied previously by Tichy (2006). In this paper the author estimates the parameters of the model from historical Czech and Austrian spot power prices. The aim of the paper is to demonstrate the ability of the model to explain the skewness and kurtosis visible in the historical price series and as such the author does not provide a characteristic function for the process nor any discussion on derivatives pricing. One major contribution of this thesis is the detailed derivation of the forward curve consistent characteristic function of the spot price under the mean-reverting Variance Gamma process. This is of fundamental importance in pricing derivative contracts and thus providing a direct link between the model and the market.

In recent years, a predominate aspect of the storage valuation literature has been the need to accurately capture the gas forward curve calendar spread dynamics, particularly the pairwise correlations between injection/withdrawal month price returns. Given that traders typically view storage as a time-spread play, valuation methods should be able to reflect the extrinsic value arising due to the dynamics of these spreads. This is not purely an academic exercise, the p&l of ones delta hedged portfolio will be significantly influenced by the seasonal cross-gamma of your asset. This will generally be a positive given a less than perfect level of correlation between winter and summer months. Spot optimization methods which rely upon a Markovian model of the underlying price, dependent upon a single state variable, may in reality under price this extrinsic value. The second part of this thesis will focus on addressing this shortcoming by generalizing the storage valuation methodology presented in Chapter 2 to multidimensional forward curve models. Although, as we will demonstrate, this additional modelling flexibility does come at the cost of an increase in the computational effort required.

In their seminal paper Heath et al. (1992) derive a model for the entire forward interest rate curve ca-

pable of incorporating the individual volatilities and pairwise correlations of all points on the curve. The main drawback of their general approach is the non-Markovian nature of the resulting short rate which in practice prohibits the use of lattice based valuation algorithms. This general modeling approach was first applied to commodity forward curves by Clewlow and Strickland (1999a) in which the authors discuss the potential extension to lattice based algorithms for pricing early exercise claims by approximation of the full model by a low-dimension representation. Cheyette (2001) formalizes this approach with respect to interest rate derivatives where a number of additional state variables are added to the model along with volatility functions of an appropriate form such that the term structure of forward rate volatility can be reproduced to a high degree of accuracy and the implied short rate dynamics remain Markovian.

Eberlein and Raible (1999) extend the Heath-Jarrow-Morton (HJM) modeling approach to include Lévy driven processes and provide a proof that the resulting spot price models are Markovian if the volatility term structure is of the form found in the Vasicek³ or Ho-Lee models⁴. Gapeev and Kuchler (2006) formalize this assertion by providing the necessary conditions on the structure of the short rate volatility for the model to yield a Markovian short rate. Filipović and Tappe (2008) provide a proof of the existence and uniqueness of the solution of general Lévy driven term structure models, which they highlight had been missing from the literature. Eberlein and Kluge (2007) discuss the calibration of Lévy driven term-structure models to the market prices for caps. In testing the model, they suggest that a negative exponential plus a constant volatility term structure is suitable for replicating the shape of the implied volatility term structure. Andersen (2010) gives a detailed exposition of this general modeling approach in the context of commodity markets where the derived processes encompass stochastic volatility, jump, and regime-switching models. The author suggests the use of Principal Component Analysis in identifying the number of factors to use in a low dimensional Markov representation of the forward curve dynamics and also in determining the structure of the factor volatility term structure. Further, the author provides numerical evidence of the ability of a Markovian three factor model to accurately capture the historical dynamics of the US natural gas market, as well as reproducing the implied volatility term structure of single delivery options⁵.

Attempts at incorporating multifactor forward curve models in storage valuation include, Boogert and De Jong (2011) who proposed modeling the spot price using a three-factor Ornstein-Uhlenbeck process. The three factors were chosen to represent spot price volatility, forward curve volatility and winter-summer spread volatility. The optimization itself is carried out using Least Squares Monte Carlo, thus benefiting from the computational efficiency of Monte Carlo methods with respect to increases in dimensionality. However this does come at a cost, as the convergence of the algorithm deteriorates.

³See Vasicek (1977) for details.

⁴See Ho and Lee (1986) for details.

⁵While such products are available on an OTC basis, “bullet” options which deliver into monthly forwards/futures are far more liquid.

Bjerksund et al. (2008) proposes a six factor forward curve model which aims to replicate the historical dynamics of the UK gas market to a high degree of accuracy. The author utilizes the rolling intrinsic valuation method, presumably, to overcome the high dimensionality of the proposed model. Valuation results suggest a significant increase in storage value over a traditional single factor model combined with spot optimization. This result is significant as it highlights the trade-off between value, price model, and valuation algorithm. Whilst, spot optimization will always yield the highest value for a given price model, the actual price model choice is equally as important.

Parsons (2013) includes the long-term mean of the underlying price as an additional stochastic state variable and uses a forest of multidimensional trinomial trees in order to optimize the storage value. The tree can be constructed so as to attain consistency with the initial forward curve which is a prerequisite for storage valuation. Warin (2012) investigates the valuation and hedging of storage assets under a two factor forward curve model. The first factor is chosen to represent the short term volatility term structure, and the second, the long term volatility of the gas market. The author goes on to discuss in detail the calculation of delta hedging volumes and studies the effectiveness of the hedged portfolio under the proposed model.

In Chapter 3 we demonstrate how one can construct market consistent forward curve models which also capture the rich time-spread dynamics of the forward curve. We extend the Mean Reverting Variance Gamma process to a multifactor setting and in a similar manner to Andersen (2010) specify each factor to represent a latent principal component of the underlying forward curve covariance matrix. We go on to generalize many of the results of the single factor model to a multifactor setting including analysis of the model implied covariance structure and the analytical link between realized spot price and forward curve. As such, we follow the example of Clewlow and Strickland (1999b) in providing a full and detailed analysis of the properties of this class of models to provide practitioners with the necessary tools upon which a complete derivative pricing library can be based.

In all of the aforementioned literature on storage valuation the asset was valued using a model calibrated to the historical price distribution, thus making the implicit assumption that the future price distribution is equal to the recent past. Given the recent collapse in gas price volatility this assumption is likely to lead to a gross overestimation of the storage value. Boogert and De Jong (2011) propose a multifactor forward curve model incorporating summer/winter spread volatility. They calibrate their model to TTF gas price returns between 2008-2010. The calibrated spot volatility of the process exceeds 100% on an annual basis. The storage years following this calibration period have been characterized by diminishing levels of spot and spread volatility which may have led to traders overpaying for storage based on historical volatility. As we will show, the market implied distribution or market measure can be approximately recovered using an appropriate pricing model. As it is a forward looking estimate of the future distribution, it is a more appropriate measure to use when valuing and hedging a financial contract or asset. Further, it allows one to relate the value of one's storage asset directly to the implied volatility surface of liquid traded contracts, and consequently, it returns a price which is consistent with the cost of

hedging against changes in the level of market volatility. This has immense appeal from a practitioners perspective as it is typical for traders to express a view on the value of a non-linear contract in terms of changes to the implied volatility surface. We would therefore argue that the primary consideration when choosing a pricing model should be its ability to reproduce the price of available volatility hedges.

In gas markets the cost of hedging the volatility will typically be given by the price of monthly options on forwards. Although, depending upon market liquidity, daily options or even storage itself could be used. The difficulty this poses when calibrating a spot price model is in deriving a related swaption pricing function. In a single factor model, one can derive a reasonably quick valuation algorithm for model calibration. Clewlow and Strickland (1999b) provide an algorithm in the special case of a mean reverting diffusion. Their insight is that the swaption payoff is monotonic with respect to the spot price. Thus one can value the swaption by first finding the minimum spot price at which the payoff is positive and then taking the joint expectation of the payoff and the spot price exceeding this level in the usual manner. In the case of a mean reverting diffusion, the authors show how this reduces the pricing to a simple Black76⁶ type analytical form. We extend this approach to general single factor Lévy driven Ornstein-Uhlenbeck processes and demonstrate their ability in replicating the implied volatility surface for natural gas options. The pricing method is quick enough to allow for direct calibration to option prices across a range of strikes and maturities.

From a theoretical perspective we could extend such a numerical procedure for pricing swaptions under a multifactor model, however the resulting algorithm would in all likelihood be overly complex for practical calibration purposes. Further, as discussed in Guillaume and Schoutens (2013) the approach of minimizing pricing errors does present its own set of challenges, the most troublesome being the sensitivity of the fitted parameters to their initial estimates. The authors propose calibrating the model parameters to the moments of the market implied pricing measure. The resulting calibration problem has an algebraic solution and thus avoids the issues encountered when using a search function. Bakshi and Madan (2000) and Bakshi et al. (2003) were the first to discuss in detail the concept of market implied moments. The idea is to construct derivative contracts whose value function represents a specific moment of the underlying price returns' distribution. The authors provide market based, and thus model free, valuation formulae for the volatility, cubic, and quartic "derivative" contracts which can then be used to derive measures of market implied variance, skew, and kurtosis. This approach relies upon the existence of a liquid options market on the underlying which is being modelled, which in the case of the models presented in this thesis would be the spot or day-ahead gas price.

In most natural gas markets, single delivery options are traded in illiquid over the counter markets, there is however a liquid market for options on delivery periods over a calendar month which could be used to extract market information on the daily or spot process. Our contribution to this methodology involves estimating the implied moments of the monthly delivery forward price in order to calibrate an

⁶See Black (1976) for details.

instantaneous forward price model. This allows for efficient calibration regardless of the delivery tenor of the underlying forward contract as well as allowing for tenors of variable duration. The most appealing aspect of utilizing this approach in model calibration, as we will demonstrate in Chapter 3, is that the moments of the forward price can be easily derived for any reasonable forward curve model regardless of dimensionality.

For the purposes of valuing storage, this thesis will focus heavily on the use of the Fourier Transform so as to efficiently utilize the market consistent Lévy based price models developed throughout. Transform-based valuation methods allow one to incorporate much of the modeling flexibility inherent in Monte Carlo whilst also obtaining the greater efficiency of the PDE approach. The value is derived by repeatedly evaluating the continuation value at each point in the state space through numerical integration. The only restriction on the spot price model is that it is Markov, all one needs is the characteristic function of the transition density. The application of transform based methods in financial markets has seen increasing interest over the past two decades. The approach was first utilized by Heston (1993) to derive a quasi-analytical solution to option pricing under stochastic volatility by Fourier inversion of the characteristic function. Duffie et al. (2000) generalized this method to include the family of affine jump diffusions as well as a number of different option payoffs. Lewis (2001) derived single integral general Fourier transforms for a number of option payoffs on exponential Lévy processes.

Computational efficiency of the inversion method has been greatly enhanced through the use of the FFT algorithm. Carr and Madan (1999) apply this technique by taking the option price to be a function of the log strike and, after adjusting the transform via a dampening parameter to ensure square integrability in a manner similar to the approach of Lewis (2001), they are able to efficiently price options across a range of strikes. In a similar vein, Chourdakis (2004) utilizes the Fractional FFT for option pricing and demonstrates its flexibility in terms of grid choice, albeit at a cost of greater computational effort. Fang and Oosterlee (2008) introduce the COS method which they demonstrate to be considerably faster than the method of Carr and Madan (1999) for characteristic functions with exponential decay and for a small number of strikes. O'Sullivan (2005) extends the QUAD method of Andricopoulos et al. (2003) to incorporate the FFT, which allows one to efficiently price path dependent options by recursive integration of the option value. Lord et al. (2007) improves upon the convergence rate of this method by first deriving the Fourier transform of the convolution of the option's continuation value and log price density and then working backwards over a finite grid. Jackson et al. (2007) independently presented a similar method, Fourier time-stepping, for solving the partial differential equation associated with the value process of an early exercise claim. The difference between both methods is analogous to the difference between the martingale and PDE approach to valuing financial contracts, in that the actual numerical implementation will be almost identical in both cases.

Fang and Oosterlee (2009) present a method for pricing Bermudan options which is similar to that

of Lord et al. (2007), save for the fact that it uses the authors' previously developed COS method to perform the required integration. Further, the authors show how by solving for the exercise boundary directly the efficiency of the method can be improved by limiting the number of cosine series coefficients which require numerical evaluation. The authors state that although their method is half as fast as that of Lord et al. (2007) for the same number of grid points, it significantly outperforms when one looks for the same level of accuracy. This last claim comes with the condition that the characteristic function displays exponential decay however and, as pointed out by Kahl and Lord (2010), for characteristic functions with power law decay, such as the Variance Gamma process, the convergence rate of the COS method is geometric rather than exponential. This last point would seem to invalidate the latter's claim of greater speed for general Lévy processes. Černý and Kyriakou (2011) apply the convolution approach to the pricing of discretely monitored Asian options and notes the computational gain over and above Monte Carlo and Finite difference methods. Jaimungal and Surkov (2011) extend the Fourier time-stepping method to incorporate mean reversion in the log-price process. Their crucial insight being that the conditional characteristic function of the log-price can be converted to a form suitable for FFT evaluations by invoking the scaling property of the Fourier Transform.

Much of the recent work in this area has focused on improving the efficiency and stability of FFT methods. Carr and Madan (2009) expand on earlier work on saddle-point methods to improve the pricing of deep out of the money options. Ding (2010) combine a simple quadrature and the FFT approach of Carr and Madan (1999) in attempt to improve the stability of the integration at low strikes and remove the dependency of the solution on the dampening parameter. In a similar vein, Lord and Kahl (2007), discuss the choice of the contour of integration and conclude that finding the optimal contour is the most robust method for overcoming numerical instabilities. Joshi and Yang (2011) suggest using the Black-Scholes option price as a control variate when inverting European option prices using a Gaussian quadrature. They go on to analyze the dependency of the convergence of the valuation algorithm on the choice of contour and conclude that the computational cost of introducing the control variate is outweighed by a significant increase in accuracy. Haslip and Kaishev (2014) use spline interpolation of the complex integral representation of the option price to derive a closed form approximation, thereby significantly decreasing the computational effort required.

For the models developed in Chapter 3 we will of course need to apply these Transform based pricing techniques in a multidimensional setting. In this context, Beyna and Wystup (2011) investigate the characteristic functions of the Cheyette model class introduced by Cheyette (2001) and apply Fourier Transform techniques to the pricing of interest rate derivatives. In general, the pricing of early exercise claims under the Cheyette model class will require the addition of at least one state variable and thus lead to an increase in dimensionality. Benth et al. (2015) derive an explicit formula for exchange option prices under bi-variate Geometric Brownian Motion with stochastic volatility by reducing the problem to a single dimension via Esscher transform. For more complex payoff functions, the Fast Fourier Transform

algorithm has been utilized in a multidimensional setting, for example, by Dempster and Hong (2002) for the purpose of spread option valuation. Here, the authors use the two dimensional version of the algorithm to numerically evaluate the option value integral over the non-linear exercise region. Leentvaar and Oosterlee (2007) price multi-asset options using the Fast Fourier transform to evaluate the relevant multidimensional integrals. They introduce the parallel partitioning method in an effort to reduce the computational burden. This approach involves appropriately splitting the transform into segments which can be evaluated more efficiently in isolation. Ruijter and Oosterlee (2012) extend the Fourier cosine expansion method to a multidimensional price processes and provide numerical evidence of the efficiency of the method based on the pricing of rainbow options.

Against the backdrop of this literature, in Chapter 2 we utilize the Convolution Method of Lord et al. (2007) combined with that of Jaimungal and Surkov (2011) to construct a single factor FFT based storage valuation algorithm. We have chosen the Convolution Method over the COS method of Fang and Oosterlee (2009) due to the greater efficiency, as outlined above, of the former with respect to general Lévy models. We introduce an innovative extension to their approach which greatly increases the convergence rate and stability of the algorithm. Further, this new method removes the need for a dampening parameter, the choice of which, as discussed above, can be a major source of instability. In Chapter 3, we go on to generalize this methodology to an arbitrary number of dimensions in order to utilize the multifactor price models developed earlier in the chapter.

In the final part of this thesis (Chapter 4) we examine an often overlooked aspect in the storage valuation literature but a growing area of research in financial markets, that of model uncertainty. Unlike typical financial risk metrics which are based upon the assumption of a known probability distribution, model uncertainty takes a Knightian, Knight (1921), approach and deals with the ambiguity on the choice of probability distribution itself. In the wider economic literature, aversion to uncertainty is a major topic of interest, see Ellsberg (1961) and Epstein (1999) for recent examples. In a financial context, this aversion to uncertainty on the true nature of the probability distribution for future asset values can be mitigated by trading in related products. This one action both determines the probability distribution, or some subset of it, and eliminates ones exposure to it. Therefore, the aversion to uncertainty in the context of pricing a contingent claim is primarily a function of the availability of related products.

As discussed above, for gas storage the main driver of the extrinsic or option value is the expected future level of forward curve time-spread volatility which is not generally hedgeable. As such, a probabilistic model is required to estimate this value. The sensitivity of the valuation result and realized p&l to the uncertainty implicit in the model choice and parameters used adds an element of risk which should be accounted for when determining the prices to bid and offer storage capacity. In Chapter 4 we seek to fill a gap in the literature and derive a consistent and holistic approach to evaluating these contracts with respect to the uncertainty associated with parameter calibration/estimation. To that end, we consider both

calibration to market instruments and estimation from historical data. We will follow the recent example of the exotic options markets literature and propose adjusting storage valuations to account for model uncertainty using the recently developed theory of convex model uncertainty measures.

The potential benefit of such a framework will impact equally trading, risk and regulatory stakeholders within a storage business. Morini (2011) gives a comprehensive exposition of the various aspects of model risk analysis, including parameter risk and model specification. The author focuses on providing a practitioner based rather than purely theoretical introduction to the topic and highlights the potential usage across several business units. Further, the regulatory impetus for incorporating sound model validation and associated risk management structures is exemplified by the Office of the Comptroller of the Currency (OCC) 2000-16 Risk Bulletin on Model Validation issued in 2000 and updated in 2011. The scope of the bulletin covers all aspects of the model development process, from the conceptual soundness of the underlying theory, to the actual implementation and usage of the code. The OCC recommend independent auditing of each aspect wherever possible and ongoing comparison of the model output against reality.

In the 2011 report, the OCC outline two key sources of model risk; the first arising from potential errors during the design and implementation phase and the second arising from erroneous reporting of the model results. The three key processes identified for sound model validation are “a) independent review of the logical and conceptual soundness, b) comparison against other models, and c) comparison of model predictions against real world events”. As we will demonstrate, an objective yet informal method for ranking models based upon their robustness to calibration errors can be derived using the theory of model uncertainty measures and used to directly address the second point. Such a ranking methodology would also find an obvious use in model validation. It would allow risk teams to robustly quantify the additional risk a new model poses over an incumbent and advise on the risk/reward benefit of changing.

Early work on the topic of model risk includes Figlewski (1998) who gives three specific definitions of what he terms model risk; i) model misspecification ii) the presence of unobservable state variables and iii) the model dependency of hedging performance. The first two clearly come under the remit of risk control whilst the latter would be of chief concern to the trading desk utilizing the model to manage a product. In the context of interest rate derivatives, but equally applicable to energy markets, Gibson et al. (1998) identify a number of steps in model selection; i) characterization of the market and the factors driving the forward term structure, ii) model specification of the underlying factors, such as distribution assumptions and Markovian representation, and iii) the stability of parameter estimates from historical data. In Chapter 2 and Chapter 3, we will follow the first two steps when building appropriate models of the natural gas forward curve. In Chapter 4, the focus will be in quantifying the risk associated with the calibration of these models to both market and historical data.

Kerkhof et al. (2002) were the first to make a distinction between the terms “Parameter Risk” and

“Model Uncertainty”. The former results from uncertainty surrounding the choice of parameter(s) in a given pricing measure, for example the historical volatility associated with a single factor diffusion. Thus, it reflects the risk associated with the estimation/calibration of a pricing measure and the resulting impact on an asset value. The latter reflects the uncertainty on the pricing measure itself in incomplete markets, where several pricing models may be deemed acceptable based upon certain goodness-of-fit criteria such as replication of market option prices. Parameter risk can be categorized further into two distinct subsets, namely parameter estimation risk and calibration risk. Parameter estimation risk is generally associated with the historical estimation of model parameters and more specifically with the variability of these estimates. There is an unhedgeable realized p&l risk given by the difference between estimated and realized parameter values which should be accounted for when pricing a product. Calibration risk on the other hand is due to the uncertainty on the true value of the market calibrated parameter values, which are generally accepted if the model prices of benchmark market instruments lie within bid-offer. The p&l variation attributable to calibration risk is therefore driven by the selection of parameters from the set of all parameter combinations which yield acceptable model prices. In his seminal paper, Cont (2006) argues that the distinction between parameter risk and model uncertainty is largely irrelevant. Model uncertainty is determined by the uncertainty surrounding the choice of probability measure and since each possible parameter selection for a specific model yields a distinct probability measure, parameter risk is in fact a special case of model uncertainty.

Faced with the above definitions of parameter calibration and estimation risk one may conclude that the optimal choice of model is one which prices all benchmark instruments perfectly and requires no parameter estimation from historical data. Hull and Suo (2002) study the model risk associated with the Implied Volatility Function or Local Volatility model. This model takes the market price of European Options as input and therefore exhibits no calibration error. However, as the authors point out, the model typically misrepresents the future marginal distribution of the asset which leads to unrealistic pricing of certain exotic options, barrier options being the example used in this case. Thus, it is not necessarily the case that the optimal model is that which carries the lowest calibration error. A more complete view of model risk needs to take account of not only the mispricing of benchmark instruments, and the noisy effect this may have on p&l, but also incorporate any historical information which may shed light on the likely future physical measure. Dynamic Hedging in incomplete markets necessitates exposing oneself to the physical “real world” measure and thus the likely impact this may have on ones p&l must also be accounted for.

Of course capturing the dynamics of the real world whilst also maintaining consistency with the market will typically require a greater level of complexity in the chosen price model, which will in turn carry more parameter risk. In the context of storage pricing, where the main driver of the value is in the unhedgeable correlation structure of the forward market⁷, this trade-off between complexity with associated

⁷Storage asset values can be seen as a choosing the highest value forward curve time-spreads achievable through the

calibration risk, parameter estimation risk, and accuracy is even more stark. Each additional parameter which decorrelates the forward curve, and therefore adds value, will generally be estimated from, typically quite noisy, historical data whereas the level of forward curve volatility can be implied from the options market. We will show that the parameter risk metrics recently developed by Bannör and Scherer (2013b) can be unified to provide a more complete picture of model risk. This approach will allow traders to adjust their price levels to account for the potential p&l swings owing to uncertainty on the choice of model parameters. This would also allow risk control teams to attribute such p&l movements appropriately, which can only lead to an increase in confidence in a model.

The study of a such unified approach to calibration and parameter estimation risk is to date missing from the literature, where the focus has been to examine each aspect in isolation. Gupta et al. (2010), for example, address the impact of calibration risk and model uncertainty. For the former, they highlight the ill-posed nature of the calibration problem when using market data and the lack of a unique solution as the key drivers. For the latter, they propose a number of methods to quantify and account for model risk in asset prices including Bayesian averaging over the distribution of potential pricing measures, and also simply taking the worst-case. Similarly, Kerkhof et al. (2010) in proposing a methodology for accounting for model risk in capital reserves distinguish between parameter estimation risk and model misspecification risk, whilst refraining from elevating one source of model risk over another. What is clear from the literature is that the analytical methods available to estimate and rank these risks are the same regardless if they are being used to measure model uncertainty, calibration risk or parameter estimation risk.

In deriving probabilistic metrics of calibration risk, Deryabin (2012) utilizes the risk measures specified by Cont (2006) to develop bounds on the calibration risk dependent upon the choice of stochastic process and its calibration to benchmark instruments, but independent of the choice of model parameters. Bannör and Scherer (2013b) improve upon this by introducing the concept of calibration risk functionals for incorporating calibration risk into asset bid-offer levels. Their approach is similar to that of Bayesian averaging in the sense that one works with a distribution of acceptable pricing measures sharing a common specification. Further, the authors show how one can induce an asset value distribution from the calibration error given by a stochastic model. It is this novel approach which we will utilize in estimating the parameter risk associated with storage valuations.

In terms of measuring the parameter risk associated with real asset valuations, Bannor et al. (2013a) were the first to focus on this topic where they evaluated the effect of model risk on power plant valua-

storage injection and withdrawal capabilities. As the ranking and value of the most profitable time-spreads change the storage owner's optionality will ensure that the storage value also changes. This ability to change planned injection and withdrawals based upon observable time spreads is what drives the extrinsic value. Consequently, a model where each time-spread has zero volatility will return zero extrinsic value. For example, this would be the case for a forward curve model driven by a single factor with flat volatility term structure. At a minimum the volatility term structure needs to vary with forward curve maturity for a single factor model to generate any extrinsic value.

tions. The authors construct a value density based upon the sample covariance of the model parameters and show how different parameter estimation risk metrics can be derived using this density. In modelling the power plant they use a 17 parameter model of the spark spread and demonstrate the potentially large risk associated with varying specifications of the model. Their work highlights the increased exposure to model risk when looking at the more complex pricing measures typically associated with physical asset modeling.

Henaff et al. (2013) study the historically estimated parameter risk associated with storage valuation. The authors use the coherent model risk measure defined by Cont (2006) to calculate the parameter estimation risk associated with two proposed spot price models, both of which incorporate price spikes. They deviate from the methodology of Cont (2006) by replacing the calibration error as a criteria for model acceptance with a test which determines whether a potential set of model parameters returns a likelihood value “close” to that of the maximum likelihood value. This approach would appear to be a specific case of the general framework put forward by Bannor et al. (2013a), where model valuations whose underlying parameters fall outside a given confidence interval range are excluded from the risk metric. Our proposed framework for quantifying both calibration and parameter estimation risk for real assets can be seen as a unification of Bannor et al. (2013a) and Bannör and Scherer (2013b) to determine the risk of combined historical and market based parameter estimations. Further, and in the context of storage valuation, the framework improves upon that of Henaff et al. (2013) by incorporating market information into the wider model risk evaluation and thus bringing it more in line with the practical considerations facing a storage trading desk.

Finally, we set out the structure of the thesis and summarize for convenience the contributions mapped out above. In Chapter 2, we begin with a general overview of the use of Lévy driven Ornstein-Uhlenbeck processes in energy spot price modelling. We present the characteristic function of one such model, the mean reverting Variance-Gamma process, with a full derivation given in Section A.2 of the Appendix. We provide a rationale for using this model in fitting the implied volatility smile by comparing the process moments with the more common mean reverting diffusion model. We next present the dynamics of the implied spot price under a general single factor Lévy driven forward curve model, with a detailed derivation presented in Section A.4 of the Appendix. Using these results, we go on to present the forward curve consistent conditional characteristic function of the implied spot price model.

In order to value storage contracts under single factor Lévy driven forward curve models, we next present a detailed derivation of the Fourier based pricing algorithm using the methodology of Lord et al. (2007) and Jaimungal and Surkov (2011). Further, we go on to present a modification to this algorithm which removes the need for a dampening parameter and leads to an increase in valuation convergence. We finish the chapter with results of an example storage valuation. We first calibrate a number of single factor models to a strip of monthly NBP gas options using a Fourier based pricing algorithm derived in Section A.5 of the Appendix and then present valuation and Greek results for each model.

In Chapter 3, we focus on extending the results from the previous chapter to a multidimensional setting. We first give an overview of general Lévy driven multifactor forward curve models before proceeding to the presentation of a particular class of models where the main stochastic driver is the mean reverting Variance Gamma process. In Section B.1 of the Appendix we give a detailed derivation of the spot price factor dynamics for these models in a manner similar to the single factor derivation presented in Section A.4. Further, we present the conditional characteristic functions associated with each of these models with full derivations provided in Sections B.2 and B.3 of the Appendix. We go to present details of an implied moment based market calibration routine for general multifactor instantaneous forward curve models and investigate the effectiveness of a number of two factor models in matching the implied volatility smile for NBP gas. We go on to generalize the storage valuation algorithm presented in Chapter 2 to an arbitrary number of dimensions. We finish by presenting valuation results for two multifactor models and analyze the relationship between model specification and storage value.

Chapter 4 focuses on analyzing the parameter risk inherent in the storage valuation models presented in the earlier chapters. We begin with a detailed review of the model uncertainty literature and, in particular, previous work which focused on incorporating parameter risk within asset valuations. We identify and bridge a gap in the literature by presenting a parameter risk framework which can incorporate both market calibration risk and statistical parameter estimation risk in an unified manner. Using the models derived in the previous chapters, we finish with detailed numerical examples of the type of informative model risk analyses that can be carried using this proposed methodology. In Chapter 5 we summarize the main results of our research and discuss the potential for future research based upon our findings.

For the benefit of the reader, we will now enumerate the main contributions of this thesis to the storage and wider quantitative finance literature. In some instances, the innovations are to be found in the appendix material accompanying each chapter.

- Chapter 2
 - The derivation of the characteristic function for the mean reverting variance gamma (MRVG) spot price process along with the moments of the MRVG process.
 - The derivation of the general forward curve consistent characteristic function for mean reverting Lévy spot price processes under the Heath-Jarrow-Morton (HJM) framework.
 - We apply previous results in the interest rate market literature to energy forward curves to derive the implied log spot price dynamics under the HJM framework.
 - We present a swaption pricing algorithm for general single factor forward curve consistent mean reverting Lévy spot price processes.
 - We present an extension of the Fourier Transform based algorithm of Jaimungal and Surkov (2011) for pricing path dependent options under mean reverting processes which enhances

the convergence rate of the algorithm.

- Chapter 3

- The extension of the MRVG model to a multifactor setting under the Cheyette model framework.
- We derive the conditional characteristic functions for a number of multifactor model specifications.
- We present a detailed derivation of the implied spot price drift for each model specification and the conditional characteristic function of the resulting log spot price process.
- We present an innovative implied moments calibration technique applicable to any finite dimensional instantaneous forward curve model.
- We derive the spot factor implied forward curve for each model specification and use these results in deriving a general Fourier based swaption pricing algorithm.
- The extension of the Fourier based valuation algorithm for processes with state dependent increments to an arbitrary number of dimensions.

- Chapter 4

- We present a joint calibration-estimation risk measurement methodology, extending recent literature, which incorporates both market calibration and historical estimation risk within a meaningful distributional assessment of parameter risk
- We extend the emerging literature on model risk issues in energy markets, applying the above technique to the problem of natural gas storage valuation, using a flexible multifactor Mean Reverting Variance Gamma model specification that is both forward curve consistent and calibrated to market traded options
- We devise an accessible model selection technique based on our distributional assessment of parameter risk.

Chapter 2

Single Factor Storage Valuation

The focus of this chapter is in providing an introduction to the modelling framework to be utilized throughout the thesis. We will first present an overview of Lévy driven forward curve consistent spot price modelling. We provide a detailed derivation of the forward curve consistent conditional characteristic for one such model, the Mean Reverting Variance Gamma process, and demonstrate its effectiveness in replicating the implied volatility smile present in natural gas markets. We will then introduce a transform based storage valuation algorithm and discuss an innovative extension of the current literature which greatly enhances the efficiency of the method. We will finish with detailed numerical examples where we analyze the storage value implied by the NBP gas options market.

2.1 Lévy Driven Ornstein-Uhlenbeck Processes

Price discontinuities are an accepted reality for all financial markets and particularly relevant for energy markets which are frequently subject to price jumps, see for example, Nomikos and Andriosopoulos (2012) and Maslyuka et al. (2013) for recent empirical studies in this area. Failure to account for this stylized fact could lead to an undervaluation of any non-linear payoff, the effect of which could have far deeper consequences in asset valuations where large amounts of capital are allocated based upon real option type analysis. In the context of storage contracts, spikes, both positive and negative, in the spot price allow the owner to trade the day-ahead versus balance of month/month ahead time-spread, a strategy which can greatly enhance the value extracted from the asset. Recall from the introduction that the day ahead contract covers gas deliveries over the following business day, whereas the balance of month (BOM) covers the rest of the current month. Therefore, given a large negative decline in the day ahead price relative to the BOM price, a trader can buy the day ahead contract and inject the gas into the storage asset upon delivery. Simultaneously, he would sell the same volume on the BOM contract and withdraw the injected gas upon delivery. This ability to monetize relative price spikes is particularly relevant in the case of faster cycling storage contracts, typically defined as needing less than 60 days to fill/empty.

Similarly, mean reversion is another accepted characteristic of many commodity markets and should be included in any model which seeks to capture fully the behavior of the price dynamics and more importantly the dynamics of the term structure of implied volatility. Realized and implied forward price volatility for many commodities decreases with time to maturity, a phenomenon known as the “Samuelson Effect” Serletis (1992). Incorporating mean reversion into our price model allows us to capture this effect as it results in a stationary process which thus has an instantaneous volatility tending to zero as time to maturity increases. From a storage perspective, this model property has the effect of increasing the spread variance between different points on the forward curve depending upon their relative maturities. As this will imply greater time-spread variability, there is a direct link between the level of mean reversion and the extrinsic value of the storage contract.

We will begin now by giving a brief outline of the most general model which is capable of incorporating both jumps and mean reversion, namely an exponential Ornstein-Uhlenbeck (OU) process with a background driving Lévy process.

A Lévy process is a right-continuous stochastic process possessing stationary and independent increments characterized by an infinitely divisible distribution. The process itself is fully described by its associated triplet (γ, σ^2, μ) , where γ is the drift, σ the diffusion volatility, and $\mu(dj)$ is a Lévy measure, satisfying $\int_{\mathbb{R}/0} \min(j^2, 1) \mu(dj) < \infty$, which measures the relative frequency of different jump sizes, j . We refer the interested reader to Sato (2001) and Barndorff-Nielsen et al. (2012) for a thorough introduction to the theory of Lévy processes.

A Lévy-driven OU process $x(t)$ is generally defined as the unique solution to the stochastic differential equation

$$dx(t) = \alpha[\omega - x(t)]dt + dy(t) \quad (2.1.1)$$

where, α is the mean reversion rate of the process, ω is the long-term mean¹, and $y(t)$ is a Lévy process.

Integrating Equation 2.1.1 gives

$$x(t) = x(s) \exp(-\alpha(t-s)) + \omega(1 - \exp(-\alpha(t-s))) + \int_s^t \exp(-\alpha(t-c)) dy(c) \quad (2.1.2)$$

where $y(t)$ has the characteristic function $\Phi(z) = \exp(\varphi(z))$ $z \in \mathbb{C}$,

$$\varphi(z) = i\gamma z t - \frac{\sigma^2}{2} z^2 t + t \int_{-\infty}^{\infty} (\exp(izj) - 1 - izx1_{|j|<1}) \mu(j) dx$$

¹For simplicity, we have specified the long-term mean here as a constant. A seasonally varying mean level would be a far better reflection of the spot price behaviour in many energy markets. In Section 2.1.1 we incorporate this seasonal behaviour through the specification of forward curve consistent model dynamics.

The above is known as the Lévy-Khintchine Formula and it is a general representation of the characteristic function of a Lévy process. The characteristic function of the process $x(t)$ conditional on $x(s)$ and denoted $\Phi(z; x(s), s, t)$ is then given as²

$$\begin{aligned}
\Phi(z; x(s), s, t) &= E[\exp(izx(t)) | x(s)] \\
&= \exp(izx(s) \exp(-\alpha(t-s)) + iz\omega(1 - \exp(-\alpha(t-s)))) \times \\
&\quad E\left[\int_s^t \exp(-\alpha(t-c)) dy(c)\right] \\
&= \exp(izx(s) \exp(-\alpha(t-s)) + iz\omega(1 - \exp(-\alpha(t-s)))) \times \\
&\quad \exp\left(\int_s^t \varphi(z \exp(-\alpha(t-c))) dc\right)
\end{aligned} \tag{2.1.3}$$

For our purposes, we are interested in utilizing Lévy driven OU processes in modelling the natural gas spot price. We follow the approach of Lewis (2001) in writing the solution for the spot price, $S(t)$, driven by a general exponential OU Lévy process, as

$$S(t) = \exp(x(t))$$

where $x(t)$ represents the log spot price,

$$dx(t) = [\alpha(\omega - x(t)) - b(t)]dt + \sigma dW(t) + dJ(t) \tag{2.1.4}$$

$W(t)$ is a standard Brownian Motion, and $J(t)$ is a pure jump process defined such that

$$J(t) = \int_s^t \int_{\mathbb{R}/0} j(v(dj, dc) - \mu(j) dj dc)$$

subject to $\int_{\mathbb{R}} |j| \mu(j) dj < \infty$ and where $v(dj, dc)$ is a random measure, or jump measure, which counts the occurrence of jumps of different sizes.

The inclusion of the convexity adjustment $b(t)$ is to ensure that,

$$\begin{aligned}
E[S(t)] &= \exp(E[x(t)]) \\
&= \exp(x(0) \exp(-\alpha t) + \omega(1 - \exp(-\alpha t)))
\end{aligned}$$

²See Cont and Tankov (2004) Lemma 15.1 for proof.

which implies

$$\exp \left[\int_0^t \exp(-\alpha c) b(c) dc \right] = E \left[\exp \left(\int_0^t \exp(-\alpha t) \sigma dW(c) \right) \right] \times \exp \left(\int_0^t \int_{\mathbb{R}/0} \exp(-\alpha t) j(v(dj, dc) - \mu(j) dj dc) \right)$$

The solution to Eq. 2.1.4 is then

$$x(t) = x(s) \exp(-\alpha(t-s)) + (1 - \exp(-\alpha(t-s))) \omega + \int_s^t \exp(-\alpha(t-c)) b(c) + \int_s^t \sigma \exp(-\alpha(t-c)) dW(c) + \int_s^t \int_{\mathbb{R}/0} j \exp(-\alpha(t-c)) (v(dj, dc) - \mu(j) dj dc)$$

The conditional characteristic function is given by,

$$\begin{aligned} \Phi_{x(t)}(z; x(s), s, t) &= \exp \left(izx(s) \exp(-\alpha(t-s)) + (1 - \exp(-\alpha(t-s))) \omega + \int_s^t \exp(-\alpha(t-c)) b(c) \right) \\ &\quad \times \exp \left(\int_s^t \varphi(z \exp(-\alpha(t-c))) dc \right) \end{aligned} \quad (2.1.5)$$

One example of a mean reverting model which yields a characteristic function in closed form is that of Deng (2000). Initially developed to model electricity prices, the model has been utilized in general energy price modelling by Kjaer (2008) and Jaimungal and Surkov (2011). The model, which we will refer to as the MRJD model, is parametrized by $(\alpha, \sigma, \mu_1, \mu_2, \lambda_1, \lambda_2)$, where

- α is the mean reversion rate.
- σ is the diffusion volatility.
- λ_1, λ_2 represent the arrival rates of positive and negative jumps respectively.
- μ_1, μ_2 control the magnitudes of positive and negative jumps respectively.

The solution to the integral in Equation 2.1.5 is given by Kjaer (2008) as,

$$\begin{aligned} \int_s^t \varphi(z \exp(-\alpha(t-c))) dc &= \exp \left(\frac{-\sigma^2 z^2}{4\alpha} (1 - \exp(-2\alpha(t-s))) \right) \times \\ &\quad \left(\frac{1 - i\mu_1 z \exp(-\alpha(t-s))}{1 - i\mu_1 z} \right)^{\frac{\lambda_1}{\alpha}} \left(\frac{1 + i\mu_2 z \exp(-\alpha(t-s))}{1 + i\mu_2 z} \right)^{\frac{\lambda_2}{\alpha}} \end{aligned}$$

As discussed in the introduction, the implied volatility surface associated with many energy markets exhibits a well defined smile, inconsistent with the assumption of normally distributed returns, and suggests that a model where the underlying distribution is leptokurtotic is required to provide a better fit. To

that end, we will now proceed to the main contribution of this section, the derivation of the characteristic function of the Mean Reverting Variance Gamma (MRVG) model. In Section 2.4 we will use this model, along with the MRJD model, to reproduce the NBP gas implied volatility surface. The primary motivation for utilizing the MRVG model is the reduction in the number of parameters compared to the MRJD model, which is accomplished without comprising upon any of the model features. From a practitioners perspective, this translates into having one less parameter to hedge when managing the smile risk implied by the model.

The Variance-Gamma Process, $\{X(t), (\theta, \sigma, \nu)\}$, is an infinite activity³, finite variation pure jump process constructed as a time-changed Brownian Motion with activity rate driven by a Gamma process such that $X(t) = \theta\Gamma(t; 1, \nu) + \sigma W(\Gamma(t; 1, \nu))$ ⁴. The parameters θ and σ have the same meaning as in the diffusion context, namely the drift and volatility respectively. The Gamma time-change has unit mean, so that the expected increase in the process per unit time is 1 with variance ν . The associated Lévy measure, is given by,

$$\mu(j) dj = \frac{\exp\left(\frac{\theta j}{\sigma^2}\right)}{\nu|j|} \exp\left(-\frac{\sqrt{\frac{2}{\nu} + \frac{\theta^2}{\sigma^2}}|j|}{\sigma}\right) dj$$

The skewness of the distribution is controlled by θ . As will become evident in the empirical section however, the implied volatility smile is generally symmetric in log-strike and as such this parameter may be unnecessary. For the benefit of model parsimony, in what follows, we explicitly set θ to zero in which case the MRVG model will be parametrized by (α, σ, ν) , where,

- α is the mean reversion rate.
- σ is the process volatility.
- ν controls the variance of the jump magnitudes.

Similarly, the MRJD model parametrization can be reduced to $(\alpha, \sigma, \mu, \lambda)$ in order to enforce symmetry within the returns' distribution.

The conditional characteristic function of the log spot price $x(t)$ driven by the MRVG process is

$$\Phi_{x(t)}(z; x(s), s, t) = \exp\left(izx(s) \exp(-\alpha(t-s)) + iz\left(\omega - \frac{1}{\alpha} \kappa_{vg}(1)\right) (1 - \exp(-\alpha(t-s)))\right) \times E\left[\exp\left(\int_s^t iz \exp(-\alpha(t-c)) dX(c)\right)\right]$$

³A pure jump Lévy process is said to be infinitely active if $\int_{\mathbb{R}} \mu(x) dx = \infty$, the source of the divergence being the frequency of jumps close to zero.

⁴For a detailed exposition of the process we refer the interested reader to Madan et al. (1998).

with

$$E \left[\exp \left(\int_s^t iz \exp(-\alpha(t-c)) dX(c) \right) \right] = \exp \left(\int_s^t \varphi_{vg}(z \exp(-\alpha(t-c))) dc \right) \quad (2.1.6)$$

where $\varphi_{vg}(z)$ is the characteristic exponent of the Variance-Gamma process (with $\theta = 0$),

$$\varphi_{vg}(z) = -\frac{1}{v} \ln \left(1 + \frac{\sigma^2 v}{2} z^2 \right) \quad (2.1.7)$$

Solving Equation 2.1.6 gives⁵

$$\begin{aligned} \Phi_{x(t)}(z; x(s), s, t) &= \exp(izx(s) \exp(-\alpha(t-s))) \\ &\times \exp \left(iz \left(\omega - \frac{1}{\alpha} \kappa_{vg}(1) \right) (1 - \exp(-\alpha(t-s))) \right) \exp(A(z, s, t)) \end{aligned} \quad (2.1.8)$$

where

$$A(z, s, t) = \frac{1}{2v\alpha} [Li_2(l_1) - Li_2(l_0)]$$

$$l_1 = -\frac{\sigma^2 v}{2} z^2$$

$$l_0 = -\frac{\sigma^2 v}{2} z^2 \exp(-2\alpha(t-s))$$

$Li_2(z)$ is the dilogarithm function defined as, $\sum_{k=1}^{\infty} \frac{z^k}{k^2}$ $|z| < 1$ and $-\int_0^z \frac{\ln(1-t)}{t} dt$ by analytical continuation on $\mathbb{C} \setminus (1, \infty)$. For Equation 2.1.8 the branch cut along the positive real axis places a constraint on $Im(z)$, namely if $z = u + iw$, we require that $A(iw, s, t)$ is defined on the principal branch, which implies

$$|w| < \sqrt{\frac{2}{\sigma^2 v}}$$

The parametrization of the MRVG model will be familiar to those accustomed to working with the Mean Reverting Diffusion (MRD) process, where the additional parameter v has the intuitive impact of increasing the kurtosis of the underlying distribution and therefore attenuating the volatility smile. We can infer the effect of the model parameters on the distribution of the MRVG process by examining the central moments of the innovations $i(\tau) = \int_0^\tau \exp(-\alpha\tau) dX(c)$ over the time interval τ ⁶:

$$E[i(\tau)] = 0$$

⁵See Section A.2 in the Appendix for the derivation.

⁶The moments were calculated by differentiating the moment generating function, $m(u) = \exp \left(\int_s^t \varphi_{vg}(-iue^{-\alpha(t-c)}) dc \right)$, $u \in \mathbb{R}$. See Section A.3 in the Appendix for details.

$$E[(i(\tau) - E[i(\tau)])^2] = \frac{\sigma^2}{2\alpha} (1 - \exp(-2\alpha\tau)) \quad (2.1.9)$$

which is equal to the variance of the equivalent diffusion⁷. For the third and fourth moments it can be shown that,

$$\begin{aligned} E[(i(\tau) - E[i(\tau)])^3] &= 0 \\ E[(i(\tau) - E[i(\tau)])^4] &= \frac{3\sigma^4}{4\alpha^2} (1 - \exp(-2\alpha\tau))^2 + \frac{3\sigma^4\nu}{4\alpha} (1 - \exp(-4\alpha\tau)) \end{aligned} \quad (2.1.10)$$

where the equivalent diffusion has 4th central moment equal to $\frac{3\sigma^4}{4\alpha^2} (1 - \exp(-2\alpha\tau))^2$. Thus, the MRVG model adds an additional level of kurtosis, and related smile attenuation, of $\frac{3\sigma^4\nu}{4\alpha} (1 - \exp(-4\alpha\tau))$ relative to that of the MRD model.

2.1.1 Forward Curve Consistent Lévy Models

When constructing a market model it is obviously crucial that the model is capable of reproducing the prices of liquid traded instruments, particularly those which are likely to form part of a hedging portfolio. At the very least, the model should be consistent with the observed market forward curve⁸. Forward curve consistent models have underpinned fixed income modelling over the past number of decades beginning with the pioneering work of Heath, Jarrow, and Morton (1992), which simultaneously modelled the entire interest rate term structure. With respect to commodity markets, Clewlow and Strickland (1999b) applied the Heath-Jarrow-Morton framework to a mean-reverting Gaussian model of the forward curve in order to derive the equivalent Markovian spot price model. Interest rate term structure consistent Lévy models have been studied by Filipović and Tappe (2008) and Eberlein and Raible (1999) among others and the results can be applied to commodity markets without difficulty. A derivation of the forward curve consistent spot price dynamics for Lévy driven single factor models is provided in Section A.4 of the Appendix, where the approach taken mirrors that of Gapeev and Küchler (2006) for interest rates.

As will become clear in the next section, it is useful for our purposes if the derived spot price model is Markovian. In general, for a Lévy driven model for the log forward curve of the form,

$$df(t, T) = \left(-\frac{1}{2}\beta(t, T)^2 - \kappa_j(\gamma(s, T)) \right) dt + \beta(t, T) dW(t) + \int_{\mathbb{R}/0} \gamma(t, T) j\bar{\nu}(dj, ds)$$

where for brevity we have used the compensated jump measure, $\bar{\nu} = \nu(dj, ds) - \mu(dj)$, the spot price process will be Markovian if the continuously differentiable functions $\beta(t, T)$ and $\gamma(t, T)$ are separable

⁷By equivalent, we mean a mean-reverting diffusion with the same α and σ values.

⁸If the model misprices the forward curve then this immediately introduces a bias when valuing more complex instruments like options.

such that

$$\begin{aligned}\beta(t, T) &= \eta(t) \zeta(T) \\ \gamma(t, T) &= \theta(t) \zeta(T)\end{aligned}$$

subject to the following conditions

$$\begin{aligned}\int_0^t \beta(s, T) dW(s) &< \infty \\ \int_0^t \int_{\mathbb{R}/0} j(v(dj, ds) - \mu(j) dj ds) &< \infty\end{aligned}$$

Given the spot price dynamics specified in the previous section, an obvious choice for the functions $\beta(t, T)$ and $\gamma(t, T)$ are $\sigma \exp(-\alpha(T-t))$ and $\exp(-\alpha(T-t))$. In which case we can write the process for the log spot price as

$$\begin{aligned}dx(t) &= \left(\frac{\partial f(0, t)}{\partial t} - \frac{1}{2} \sigma^2 + \frac{1}{4} \sigma^2 (1 - \exp(-2\alpha t)) - \kappa_j (\exp(-\alpha t)) + \right. \\ &\quad \left. \alpha f(0, t) - \alpha \int_0^t \kappa_j (\exp(-\alpha(t-s))) ds - \alpha x(t) \right) dt + \sigma dW(t) + \int_{\mathbb{R}/0} j \bar{v}(dj, ds)\end{aligned}$$

or more concisely as,

$$dx(t) = (\omega(t) - \alpha x(t)) dt + dy(t)$$

which is of course of the form given by Equation 2.1.1 with time-varying mean log price level,

$$\begin{aligned}\omega(t) &= \frac{\partial f(0, t)}{\partial t} - \frac{1}{2} \sigma^2 + \frac{1}{4} \sigma^2 (1 - \exp(-2\alpha t)) \\ &\quad - \kappa_j (\exp(-\alpha t)) + \alpha f(0, t) - \alpha \int_0^t \kappa_j (\exp(-\alpha(t-s))) ds\end{aligned}$$

The conditional characteristic function of the log-price over the time period $(t-s)$ is then given by combining the above solution with Equation 2.1.3, which gives

$$\begin{aligned}\Phi_{x(t)}(z; x(s), s, t) &= \exp\left(izx(s) \exp(-\alpha(t-s)) + iz \int_s^t \omega(c) \exp(-\alpha(t-c)) dc\right) \quad (2.1.11) \\ &\quad \times \exp\left(\int_s^t \varphi(iz \exp(-\alpha(t-c))) dc\right)\end{aligned}$$

where, as before, $\varphi(z)$ is the characteristic exponent of the underlying stochastic driver. The integral

$$\int_s^t \omega(c) \exp(-\alpha(t-c)) dc$$

can be partially evaluated for time independent model parameters, see Section A.4 of the Appendix for details.

2.2 Fourier Methods in Pricing Early Exercise Claims

We will now provide an introduction on the use of the transform based method for the pricing of path dependent contingent claims which we will use in developing our storage valuation algorithm. The main benefit of utilizing Fourier based methods for valuing contingent claims is the flexibility given to the modeller when choosing a stochastic process to represent the evolution of the market variable(s) under study. No knowledge of the transition probability density of the process is required as one deals entirely with the characteristic function, which is known for many Lévy processes and can be derived for many exponentially affine Markov processes. The convention used in defining the Fourier Transform in financial applications generally differs from that used in engineering. For clarity we will explicitly give the definition of the Fourier Transform which will be used throughout. Given a function $f(x) \in \mathcal{L}^1$ we denote the Fourier Transform of f as $\hat{f}(z) : \mathbb{C} \rightarrow \mathbb{C}$ such that

$$\hat{f}(z) = \int_{-\infty}^{\infty} \exp(-izx) f(x) dx$$

and

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \exp(izx) \hat{f}(z) dz$$

A closely related concept in probability theory is that of the characteristic function of a random variable x , $\Phi(z) : \mathbb{C} \rightarrow \mathbb{C}$ given by

$$\Phi(z) = E[\exp(izx)]$$

For a good introduction to the characteristic function we refer the interested reader to Epps (1993).

In typical “real option” style analysis of asset valuation one deals with contingent claims which may depend upon the whole history of the underlying market over the lifetime of the asset and any previous actions by the owner of said asset. A straightforward financial market analogy is given by an American style option where the holder has the right to exercise at any point over the lifetime of the contract. Lord et al. (2007) present the Convolution Method as a transform based numerical scheme for valuing American and Bermudan style claims under exponential Lévy processes. This methodology will form the basis of the storage valuation algorithm presented in the next section. We have chosen the Convolution Method over the COS method of Fang and Oosterlee (2009) due to the fact that the Convolution Method is faster for an equal number of grid points. We leave a direct computational comparison of both methods for future research. Before proceeding directly to the storage valuation problem we first need to extend the method of Lord et al. (2007) to handle Markov processes without the restriction of independent increments. This will allow us to incorporate the interpolation technique used in Jaimungal and Surkov

(2011) to account for mean reversion when deriving the storage valuation algorithm. As discussed in the previous section, the tendency of spot prices to oscillate about a mean level is a stylized feature of many energy markets. Further, with respect to the forward curve models presented earlier in this chapter, this property acts to decorrelate the model implied forward curve returns at different maturities which is a key driver of the extrinsic value of a storage contract.

The main insight of the Convolution Method is that one can value a claim by backward induction over a finite grid, where the continuation value at each time step is calculated by repeated calls to a Fast Fourier Transform (FFT) algorithm. Under the assumption that the conditional log-price density $f(x_t|x_s)$ can be written as $f(x_t - x_s)$, which will be the case if the increments of the process are independent, the continuation value can be expressed as the convolution of the option payoff with the underlying probability density. In simple terms, the method can be summarized as deriving the optimal decision at each time-step and for a range of underlying prices based solely on any immediate cashflows and the expected future value of the contract/asset in question.

Unfortunately as discussed above, energy market price returns are typically modelled such that they are conditional upon the current price level. This makes the use of processes with independent stationary increments unsuitable and therefore prohibits the use of the Convolution Method as described by the authors. Our extension of Lord et al. (2007) utilizes Parseval's Theorem⁹ to derive the continuation value conditional upon the current price level. Thus, opening up a much wider range of exponential semi-martingales which may be used to model the underlying price.

By way of example, we will focus on the valuation of a Bermudian option. The storage deal introduced in the next section can be viewed as a generalization of this contract where multiple exercises are allowed. Given a maturity T , we discretize T into a strictly increasing sequence $\{t_0, t_1, \dots, T\}$, corresponding to exercise dates, where $t_j - t_{j-1} = \Delta t$. If the holder of the option exercises at some $t_j < T$, the exercise value $X(x_{t_j}, t_j)$ is received, where x_t represents the log-price of the underlying. Thus, we derive the value of the option today $V(x_{t_0}, t_0)$ by traversing backwards through time such that

$$\begin{aligned} V(x_T, T) &= X(x_T, T) \\ C(x_{t_j}, t_j) &= E_{t_j} [V(x_{t_{j+1}}, t_{j+1})] \\ V(x_{t_j}, t_j) &= \max \{C(x_{t_j}, t_j), X(x_{t_j}, t_j)\} \end{aligned}$$

for $t_j = T - \Delta t, \dots, 0$. Here, $C(x_{t_j}, t_j)$ denotes the continuation value of the option at time t_j , and is given by

$$C(x_{t_j}, t_j) = \int_{-\infty}^{\infty} V(x_{t_{j+1}}, t_{j+1}) f(x|x_{t_j}) dx$$

⁹A detailed summary of how the theorem applies to option pricing theory is given in Section A.1 of the Appendix.

which by Parseval's Theorem can be written as

$$C(x_{t_j}, t_j) = \frac{1}{2\pi} \int_{i\omega-\infty}^{i\omega+\infty} \hat{v}(z) \Phi(z; x_{t_j}, t_j, t_{j+1}) dz$$

where $z = u + iw$, $\Phi(z; x_{t_j}, t_j, t_{j+1})$ is the characteristic function of the conditional probability density $f(x|x_{t_j})$ and $\hat{v}(z)$ is the Fourier Transform of the option payoff.

In words, this convenient property is stating that the expected future value of our contract is equal to the integrated product of the characteristic function for the underlying log-price process and the Fourier Transform of its future value.

2.3 Gas Storage Valuation

We will begin by introducing the valuation problem with reference to the standard physical constraints and operating characteristics of a typical gas storage unit. We denote the current gas inventory level as $I \in [I_{min}, I_{max}]$. The amount of gas that can be injected or withdrawn from the storage asset in a given period is typically constrained and may be dependent upon both the time period and current inventory level. We denote the injection and withdrawal rates as $i(t, I)$ and $w(t, I)$ respectively. Given a valuation period of length T , we note the following constraints on the operation of the storage:

1. The allowed injection/withdrawal nomination times over the valuation period belong to a discrete set $\{t_j\}$.
2. For a given time step t_j and inventory level I , the range of attainable storage levels is given as

$$[\max(I - w(t_j, I), I_{min}), \min(I + i(t_j, I), I_{max})]$$

We assume that when operating a storage asset the objective is to maximize the expected discounted cashflows arising from ones injection/withdrawal policy. If we denote the log gas price at nomination time t_j as x_{t_j} , and the cash flow from moving to inventory level I^* from I as $\theta(x_{t_j}, I^*; I)$, then the value of the storage asset conditional upon the inventory level and log gas price is given by

$$V(x_{t_j}, I) = \sup_{I^*} \theta(x_{t_j}, I^*; I) + E[V(x_{t_{j+1}}, I^*) | x_{t_j}] \quad (2.3.1)$$

where I^* lies within the attainable range of storage levels surrounding I . The value function defined in this way satisfies the Bellman equation and thus can be solved via backward induction.

Valuation Methodology

We will now outline the proposed methodology for valuing the storage asset. As mentioned above we will solve for the initial value by backward induction, that is, starting at the end of the valuation period we will determine the most profitable injection/withdrawal strategy based upon the current gas price and inventory level. The most profitable strategy will be the one which maximizes the immediate cash flow due to our action plus all expected future cash flows. We will follow the approach outlined in Section 2.2 by computing the expected future cash flows numerically using the Fourier Transform. To reduce the notational burden in what follows we have assumed zero interest rates.

Beginning at the end of the valuation period at time T , the value function is given as

$$V(x_T; I) = \exp(x_T) w(T, I) \quad (2.3.2)$$

At the penultimate action date t , the value conditional upon a given gas price and inventory level is

$$V(x_t; I) = \sup_{I^*} \theta(x_t; I^*, I) + E[V(x_T, I^*) | x_t] \quad (2.3.3)$$

where $\theta(x_t, I^*, I) = \exp(x_t)(I - I^*)$, again, subject to the injection/withdrawal constraints. Therefore we first need to calculate the expected value of the storage unit conditional upon a given log gas price and inventory level,

$$E[V(x_T; I^*) | x_t] = \int_{-\infty}^{\infty} V(x; I^*) f(x | x_t) dx$$

If we define the Fourier transform of the value function as

$$\tilde{v}(z; I^*) = \int_{-\infty}^{\infty} \exp(-izx) V(x; I^*) dx \quad (2.3.4)$$

then by Parseval's Theorem we have

$$E[V(x_T; I^*) | x_t] = \frac{1}{2\pi} \int_{iw-\infty}^{iw+\infty} \tilde{v}(z; I^*) \Phi(z; x_t, t, T) dz \quad (2.3.5)$$

where $z = u + iw, \in \mathbb{C}$.

We re-emphasize that the use of Parseval's Theorem allows us to consider a wider range of models, which includes processes with non-independent increments such as the MRVG and MRJD models considered here. $\Phi(z; x_t, t, T)$ is the conditional characteristic function of the density $f(x_T | x_t)$. For mean reverting processes we can re-write $\Phi(z; x_t, t, T)$ as

$$\exp(izx_t \exp(-\alpha(T-t))) \psi(z; t, T)$$

In this case Equation 2.3.5 becomes

$$E[V(x_T; I^*) | x_t] = \frac{1}{2\pi} \int_{iw-\infty}^{iw+\infty} \exp(izx_t \exp(-\alpha(T-t))) \tilde{v}(z; I^*) \psi(z; t, T) dz \quad (2.3.6)$$

The inclusion of the $\exp(-\alpha(T-t))$ term in the exponent is unfortunate as it prohibits the direct application of the FFT algorithm. To overcome this problem we follow the example of Jaimungal and Surkov (2011). Firstly, we let $z' = z \exp(-\alpha(T-t))$ which allows us to write Equation 2.3.6 as

$$E[V(x_T; I^*) | x_t] = \frac{\exp(\alpha(T-t))}{2\pi} \int_{iw-\infty}^{iw+\infty} \exp(izx_t) \tilde{v}(\exp(\alpha(T-t))z'; I^*) \times \psi(\exp(\alpha(T-t))z'; t, T) dz' \quad (2.3.7)$$

We now need to evaluate $\tilde{v}(\exp(\alpha(T-t))z'; I_m)$. To do this we utilize the scaling property of the Fourier Transform. This states that if $\tilde{v}(z'; I)$ denotes the Fourier Transform $\mathcal{F}[V(x_T; I_m)]$, then

$$\tilde{v}(\exp(\alpha(T-t))z'; I^*) = \exp(-\alpha(T-t)) \mathcal{F}[V(\exp(-\alpha(T-t))x_T; I^*)] \quad (2.3.8)$$

To derive the value of $\tilde{v}(\exp(\alpha(T-t))z'; I^*)$ we must first interpolate from the payoff function $V(x_T; I^*)$ to find the values of $V(\exp(-\alpha(T-t))x_T; I^*)$, and then take the Fourier Transform. The above methodology is applied recursively at each nomination date, t_j , over the valuation period until the initial value is recovered at t_0 .

Discretization & Solution

From Equation 2.3.5 we see that calculating the expected future cashflows involves evaluating two integrals. The first, the Fourier transform of the future value function, the second, the product of the latter and the characteristic function of the conditional probability density of the log-gas price. To evaluate these integrals we first discretize the log gas price and transform variables. We consider sub-intervals of equal length over the truncated domain of integration, such that,

$$x \in [x_0, \dots, x_{N-1}]$$

where $x_n = x_0 + n\Delta x$ and N is the number of grid points in the log gas price domain.

Similarly, we discretize the transform parameter $z' = u + iw$, such that

$$u \in [u_0, \dots, u_{M-1}]$$

where $u_m = u_0 + m\Delta u$ and M is the number of grid points in the transform variable domain.

We begin by evaluating Equation 2.3.8 for a general time step s to t

$$\begin{aligned}\tilde{v}(\exp(\alpha(t-s))z'; I^*) &= \exp(-\alpha(t-s)) \int_{-\infty}^{\infty} \exp(-iux + wx) V(\exp(-\alpha(t-s))x; I^*) dx \\ &\approx \exp(-\alpha(t-s)) \sum_{n=0}^{N-1} l_n \exp(-iu_k x_n + wx_n) V(\exp(-\alpha(t-s))x_n; I^*) \Delta x\end{aligned}$$

where we apply a composite trapezoidal rule by setting $l_n = \frac{1}{2}, n = 0, N-1$ and $l_n = 1$ elsewhere.

Discretizing Equation 2.3.6 and inserting the above gives us

$$\begin{aligned}E[V(x; I^*) | x_n] &= \frac{\exp(\alpha \Delta t)}{2\pi} \int_{iw-\infty}^{iw+\infty} \tilde{v}(\exp(\alpha(t-s))z'; I^*) \exp(izx_n) \psi(\exp(\alpha(t-s))z'; s, t) dz' \\ &\approx \frac{\exp(\alpha(t-s))}{2\pi} \sum_{m=0}^{M-1} l_m \tilde{v}(\exp(\alpha(t-s))(u_m + iw); I^*) \times \\ &\quad \exp(i(u_m + iw)x_n) \psi(\exp(\alpha(t-s))(u_m + iw); s, t) \Delta u \\ &= \frac{\Delta x \Delta u}{2\pi} \sum_{m=0}^{M-1} l_m \exp(im \Delta u n \Delta x + im \Delta u x_0 + iu_0 n \Delta x + iu_0 x_0 - w(x_0 + n \Delta x)) \times \\ &\quad \psi(\exp(\alpha(t-s))(u_m + iw); s, t) \sum_{n=0}^{N-1} l_n \exp(-im \Delta u n \Delta x - im \Delta u x_0 - iu_0 n \Delta x) \times \\ &\quad \exp(-iu_0 x_0 + w(x_0 + n \Delta x)) V(\exp(-\alpha \Delta t)(x_0 + n \Delta x); I^*)\end{aligned}$$

At each action date, we will solve for the expected future value function at each point in a two dimensional grid representing the log gas price and inventory level. In order to apply the Fast Fourier Transform algorithm the number of grid points in the log-gas price and transform variable domains must be equal, that is, $M = N$. Furthermore, the log-gas price and Fourier space steps Δx and Δu must be chosen such that

$$\Delta x \Delta u = \frac{2\pi}{N}$$

The trade-off between discretization of the log gas price and transform variable values is immediately apparent; for a fixed number of grid points the log gas price space step Δx , and the Fourier space step Δu are inversely proportional.

The effect of the dampening coefficient w , is more apparent when we move to the discrete approximation of the value integral where this methodology assumes the divergent value function is periodic. In such a scenario the end points are treated as discontinuities which results in the well known Gibb's Phenomenon Driscoll and Fornberg (2001). The effect of this on the convergence of the numerical scheme is

1. A failure to converge at the point of discontinuity.

2. Pointwise convergence elsewhere at a rate of $\mathcal{O}\left(\frac{1}{N}\right)$.¹⁰

In our case, the point of discontinuity lies at the extreme values of our truncated gas price domain and as such without dampening the value function, the continuation value at a given gas price will include an error which diminishes proportional to the distance of the gas price from the domain boundaries. The error accumulates at each time step and thus its affect is immediately evident when looking at the initial value of the contract at the extreme points of the gas price grid. By choosing an appropriate value of w , we can reduce this error significantly by decreasing the absolute size of the discontinuity. In general to minimize the error, w could be derived at each time step and inventory level from the extreme values of the future payoffs, V_0 and V_N , such that $V_0 \exp(-wx_0) \approx V_N \exp(-wx_N)$. From our empirical testing, a value of $w \approx -1$ works well in maximizing the rate of convergence of the at-the-money initial value for a range of parameter inputs. In addition to this, the absolute size of the discontinuity can be reduced by an appropriate choice of log gas price grid spacing which ensures a reasonable bound on the maximum of the log gas price.

However, we can eliminate the need for a dampening parameter altogether by utilizing the data flipping method outlined in Pan (1993) in the context of spectral analysis. The idea here is to create a mirror image of the non-periodic input function, in our case $V(\exp(-\alpha(t-s))x; I^*)$, such that the mid-point of the new function corresponds to the right boundary of the original. Formally, $V_2(\exp(-\alpha(t-s))x; I^*)$ is defined such that

$$\begin{aligned} V_2(\exp(-\alpha(t-s))x_n; I^*) &= V(\exp(-\alpha(t-s))x_n; I^*) \quad n = 0 \dots N-1 \\ V_2(\exp(-\alpha(t-s))x_n; I^*) &= V(\exp(-\alpha(t-s))x_{2N-1-n}; I^*) \quad n = N \dots 2N-1 \end{aligned}$$

The optimal exercise strategy and resulting storage value is still only evaluated on the original domain of integration, $x_n, n = 0 \dots N-1$. The effect of the additional grid points is to drastically increase the rate of decay of the payoff transform and therefore improve the convergence rate of the algorithm. Although this approach will require approximately twice the number of computations, there is the added benefit that the transform variate grid step is reduced by half for the same log-gas price grid. We refer to this method as FFT-m and will present valuation and convergence results for both approaches in the following section.

2.4 Empirical Analysis & Results

In this section we will value a typical storage contract using the aforementioned methodology. We will test the relative performance of three stochastic models of the gas price, namely, the commonly used Mean Reverting Diffusion (MRD), the Mean Reverting Jump-Diffusion (MRJD) and the Mean Reverting Variance Gamma Process (MRVG), in replicating the price of traded option contracts. We will then

¹⁰N here refers to the length of the Fourier Partial Sum.

use these calibrated models to value a standard storage contract and analyze the impact of model specification on the storage value. In a manner analogous to the pricing of exotic options in the financial markets, calibrating to the market pricing measure and utilizing it to price our storage deal means that we are effectively extrapolating the value of the storage contract from the prices of liquid tradeable assets. As discussed in the introduction, the benefit of this approach is that it allows one to relate the value of a storage contract to the market implied volatility surface and thus combine the volatility risk of the storage with those of a vanilla options book.

We will begin by describing the characteristics of the hypothetical storage contract under study as well as a description of the market data used when calibrating the price models. The most liquid option contracts available are against futures contracts with daily delivery over a specified period of time (typically monthly), and as such we need an appropriate pricing algorithm in order to accurately estimate the spot price dynamics. The price of such a contract can be easily recovered using the transform based swaption pricing algorithm outlined in Section A.5 of the Appendix.

The storage contract we have chosen to study has the following characteristics:

- Start Date: 19th December 2012
- End Date: 18th December 2013
- Initial Inventory Level: 0
- Final Inventory Level: 0
- Capacity: 29.3 GWh
- Max Injection/Withdrawal: 1.465 GWh per day.
- Underlying Gas Price: NBP (National Balancing Point) pence/therm

The above is a very simplistic example of a typical storage contract. Generally speaking, deals run through the storage year starting in April/May, however, we have opted to begin the storage deal immediately in this toy example in order to minimize the additional extrinsic value due to the maturity of the contract. As discussed, the valuation methodology can be applied to handle any path independent constraint. More complex constraints or features can be modeled by adding extra dimensions (both stochastic and deterministic) in a quite straightforward manner, although the computational performance of the algorithm will naturally deteriorate in this case. The NBP forward curve as of the 19th December 2012 is presented below (Figure 2.4.1), along with the corresponding intrinsic trading inventory levels. The strong seasonal behavior of the gas price is immediately evident where the winter months trade at a significant premium to the summer months. This behaviour is the primary driver of the intrinsic value of the storage contract.

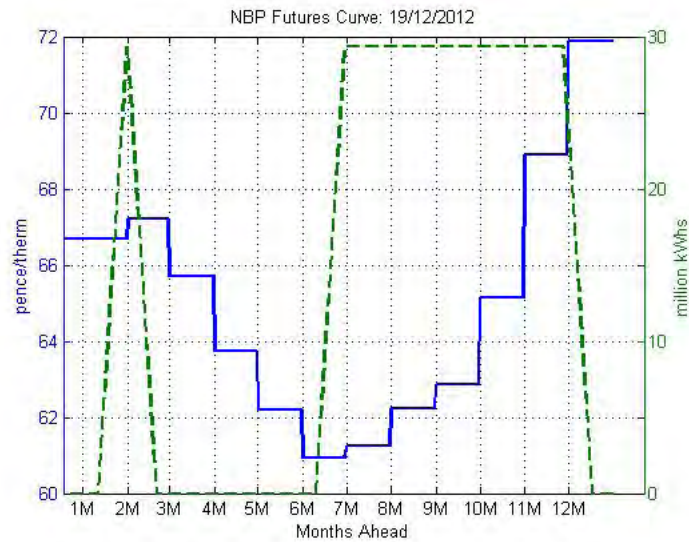


Figure 2.4.1: NBP Forward Curve

Futures closing prices (blue line) in pence per therm as of the 19th December 2012, source Bloomberg. The green dashed line displays the storage inventory levels (in million Kilowatt Hours) given by the intrinsic hedging strategy associated with this forward curve. The intrinsic value associated with this strategy is 10.983 pence/therm.

NBP futures contracts deliver daily over the contract period, which in our case is monthly. To create a daily forward curve we expanded the monthly curve piecewise flat. We decided against more involved interpolation methods primarily because we wished to conserve the intrinsic value inherent in the monthly forward curve. An intrinsic trading strategy involves locking in flat prices to buy/sell gas during delivery months in the future. When constructing a forward curve via an interpolation method other than piecewise flat certain days during a given month will either be above or below the current forward price for that month. This of course will lead to a potentially large discrepancy between the value that can be locked in by trading in the forward market and that which is returned by an optimal daily trading strategy when spot volatility is zero. This introduces an additional model dependent element to the derived value which for practical purposes should be avoided.

Exchange traded options on NBP are written on futures contracts and traded on the ICE exchange as well as over the counter. Figure 2.4.2 displays the volatility surface as of the 19th December 2012. The time to maturity or Samuelson effect is immediately apparent with implied volatility decreasing sharply with option expiry.




Figure 2.4.2: NBP Volatility Surface

Implied volatility levels as a function of time to maturity based upon closing option prices as of the 19th December 2012, source: Bloomberg. The surface displays a negative slope with respect to time to maturity and a pronounced smile with respect to moneyness, both of which support our choice of a Lévy driven mean reverting process as a price model.

We calibrated our models using both the June 2013 and December 2013 option chains in order to incorporate the effect of the mean reversion on the volatility term structure. We calibrated both models across all strikes using the sum of squared percentage errors between the market and model prices as the objective function. This metric was chosen over other norms, such as the sum of squared absolute pricing errors, as we wanted to give equal weight to pricing errors occurring at the extreme strikes. The results are summarized in Figures 2.4.3-2.4.5, and Tables 2.4.1-2.4.2 below.

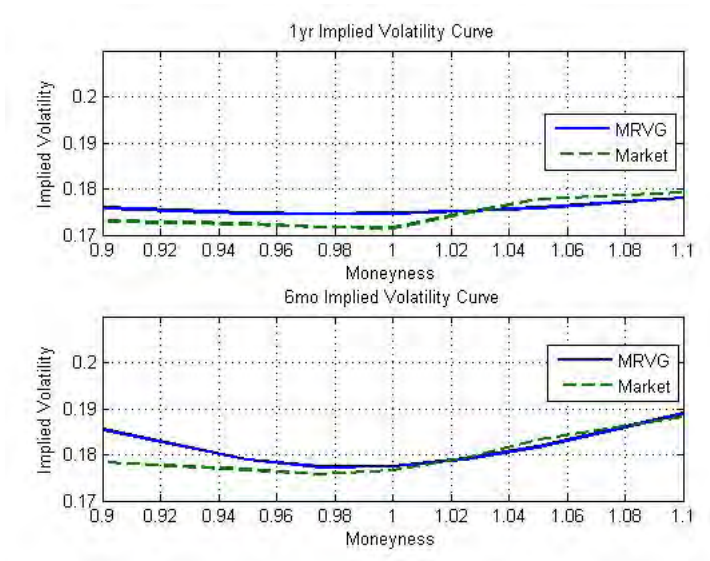


Figure 2.4.3: MRVG Calibration

The fitted MRVG and market implied volatility smiles for the 1 year and 6 month maturity options. The fitted implied volatilities were derived by calibrating the model to the option prices and then solving for the implied volatility associated with the fitted prices.

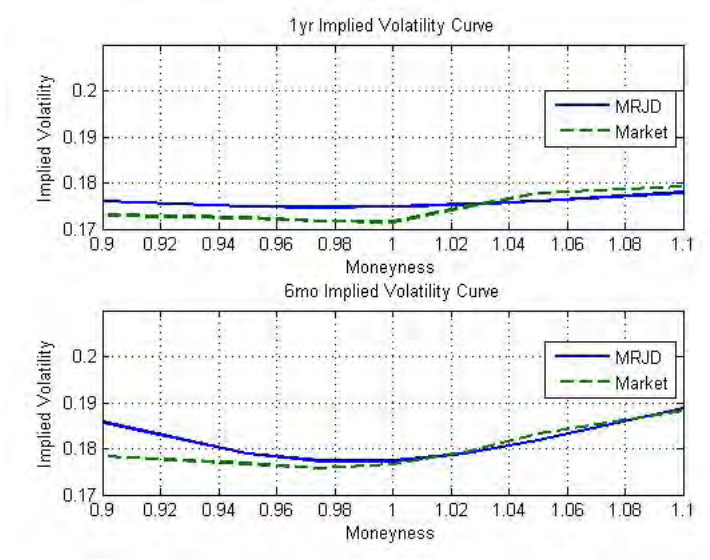


Figure 2.4.4: MRJD Calibration

The fitted MRJD and market implied volatility smiles for the 1 year and 6 month maturity options. The fitted implied volatilities were derived by calibrating the model to the option prices and then solving for the implied volatility associated with the fitted prices.

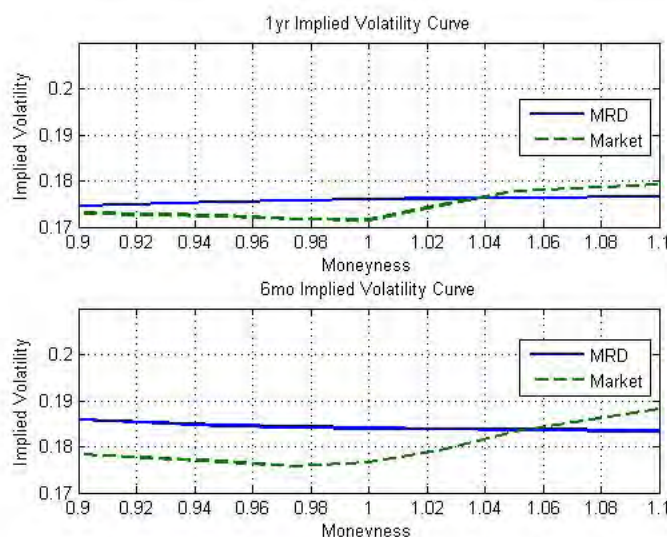


Figure 2.4.5: MRD Calibration

The fitted MRJD and market implied volatility smiles for the 1 year and 6 month maturity options. The fitted implied volatilities were derived by calibrating the model to the option prices and then solving for the implied volatility associated with the fitted prices.

Model	α	σ	ν	SSE
MRVG	0.2162	0.201	0.2560	0.16%
MRD	0.1079	0.1879	-	1.06%

Table 2.4.1: MRVG & MRD Model Parameters

The model parameters for both the MRVG and MRD models, estimated by calibrating to the market prices of options. Here, 'SSE' refers to the sum of squared percentage pricing errors, which was the metric used when fitting each model.

Model	α	σ	λ	μ	SSE
MRJD	0.2099	0.0334	8.7966	0.047	0.17%

Table 2.4.2: MRJD Model Parameters

The model parameters for the MRJD model estimated by calibrating to the market prices of options. Here, 'SSE' refers to the sum of squared percentage pricing errors, which was the metric used when fitting the model.

As is evident both graphically and via our informal goodness of fit measure, both the MRVG and MRJD model significantly outperform the MRD model in terms of matching both the implied volatility

smile and term structure. If we view this calibration procedure for what it is, a curve fitting exercise, this result is not surprising given the central moments for the MRVG process given in Section 2.1. Recall, that the 1st to 3rd central moments of the MRVG innovations are equal to the equivalent mean reverting diffusion innovations, however, the 4th central moment exceeds that of the equivalent MRD by $\frac{3\sigma^4\nu}{4\alpha}(1 - \exp(-4\alpha\tau))$. Thus, the extra degree of freedom provided by the parameter ν allows us to replicate the leptokurtotic nature of the market measure, as is evident by the pronounced smile in the implied volatility surface. Further, the model implied volatility surfaces from the MRVG and MRJD processes are almost identical. This confirms our intuition on the ability of the MRVG process to provide the same level of modelling flexibility as the MRJD process with the benefit of one less parameter. The model parameters are independent of time in this example. As mentioned previously, incorporating such a time dependency is straightforward from a theoretical perspective and would allow one to estimate the mean reversion rate independently of the implied volatility term structure, potentially from the prices of other storage contracts. Such an exercise is left for future research given the lack of availability of prices in the OTC virtual storage market, and for now we proceed by calibrating to the monthly swaption (or bullet) market only.

As noted in the introduction, the majority of the storage valuation literature has focused on Gaussian based models similar to the MRD model. Valuing our storage contract under both the MRVG and MRJD processes will mean using a pricing measure which is far more consistent with the market measure than would be the case under a simple diffusion. The valuation results are presented below (Table 2.4.3), for benchmarking purposes we have included the value derived under the MRD model using the “forest of trees” methodology¹¹. The number of grid points used in the FFT implementation is 4,026. The resulting log gas price grid was equal to $\{-6.9078, 6.0915\}$, which corresponds to a gas price range of $\{0.001, 442.07\}$. The dampening coefficient w was set to -1, this value was chosen to reflect the linear dependence of the initial value function on the gas price (see Section A.7 of the Appendix). The initial value of the storage was invariant for reasonable choices for w . We also include the value from the FFT-m method which, as outlined in Section 2.3, removes the need for a dampening parameter.

¹¹The tree based valuation was carried out using Lacima Analytics Storage Tool.

Initial Gas Price: £0.667	Intrinsic Value	Full Value	Extrinsic Value	% Δ Extrinsic Value
Tree: MRD	10.983	11.1013	0.1183	-
FFT: MRD	10.983	11.1013	0.1183	-0.024%
FFT-m: MRD	10.983	11.1013	0.1183	-0.024%
FFT: MRJD	10.983	11.2031	0.2201	186.09%
FFT-m: MRJD	10.983	11.2031	0.2201	186.09%
FFT: MRVG	10.983	11.2105	0.2275	192.34%
FFT-m: MRVG	10.983	11.2105	0.2275	192.34%

Table 2.4.3: Single Factor Valuation Results

The valuation results for each model are given here using both the FFT and FFT-m methods. For benchmarking purposes, we have also included the valuation results for the MRD model using the “forest of trinomial trees” valuation algorithm. All values are expressed in pence/therm.

Convergence Results

Tables 2.4.4-2.4.6 details the convergence of both the FFT and FFT-m algorithms with respect to the number of grid points for each model. The rate of convergence is noticeably greater for the FFT-m method, particularly for the MRVG model where it converges at 2^{10} log-gas prices compared to 2^{12} for the FFT method. This means that, although the FFT-m method is slower than the FFT method, it is far more efficient. The MRVG value converges in 17 seconds on average versus 47 seconds using the FFT method. Similar, although less dramatic, results can be seen for the MRJD model. For the MRD model however, the values converge at the same number of grid points and the FFT-m method takes almost twice as long on average.

The speed of the MRJD valuations dominates those of the MRVG model, particularly as the number of grid points increases. This is due primarily to costly calls to the dilogarithm function when evaluating the MRVG characteristic function. The function implementation has not been optimized and therefore we would hesitate to draw any firm conclusions on the relative speed of both models on the basis of this analysis. The valuations were carried out on an Intel Core i7 L640 processor and coded in Matlab. The main computational bottleneck for all three models is the interpolation over the value domain at each time-step. One obvious potential remedy would be to implement the algorithm in a more performance critical language.

	Grid Points: 2^N	FFT Value	Average Time (secs)	FFT-m Value	Average Time (secs)
MRD	4	6.1557	1.98	917.95	2.69
	5	8.2942	2.13	158.15	3.02
	6	10.6324	2.19	27.65	3.17
	7	11.0475	2.43	12.2484	3.36
	8	11.0989	2.90	11.1291	4.18
	9	11.1004	3.94	11.1014	6.31
	10	11.1013	6.34	11.1013	11.05
	11	11.1013	10.78	11.1013	19.08
	12	11.1013	20.36	11.1013	37.91

Table 2.4.4: MRD Convergence and Processing Times

Convergence results as a function of grid size, represented here by the function 2^N , for the MRD model using both the FFT and FFT-m valuation algorithms. The average time was taken across 5 runs using an Intel Core i7 L640 processor.

	Grid Points: 2^N	FFT Value	Average Time (secs)	FFT-m Value	Average Time (secs)
MRJD	4	0.00	3.26	1000.2	3.48
	5	3.6096	3.31	183.07	3.56
	6	8.8870	3.43	33.46	3.87
	7	10.4663	4.01	13.2989	4.31
	8	10.8930	4.58	11.5536	5.06
	9	11.0989	6.22	11.2333	7.57
	10	11.1867	9.47	11.2034	12.09
	11	11.2025	15.02	11.2031	22.20
	12	11.2031	27.46	11.2031	43.11

Table 2.4.5: MRJD Convergence and Processing Times

Convergence results as a function of grid size, represented here by the function 2^N , for the MRJD model using both the FFT and FFT-m valuation algorithms. The average time was taken across 5 runs using an Intel Core i7 L640 processor.

FFT	Grid Points: 2^N	FFT Value	Average Time (secs)	FFT-m Value	Average Time (secs)
MRVG	4	0.00	3.65	1003.9	4.21
	5	3.2114	4.30	184.49	4.33
	6	8.7484	4.71	33.82	4.79
	7	10.4356	5.50	13.3498	5.62
	8	10.9069	6.71	11.5514	7.53
	9	11.1004	9.70	11.2386	10.18
	10	11.1863	16.35	11.2105	17.31
	11	11.2088	25.84	11.2105	31.28
	12	11.2105	47.04	11.2105	66.15

Table 2.4.6: MRVG Convergence and Processing Times

Convergence results as a function of grid size, represented here by the function 2^N , for the MRVG model using both the FFT and FFT-m valuation algorithms. The average time was taken across 5 runs using an Intel Core i7 L640 processor.

Valuation Analysis

Although the extrinsic value of the storage asset is quite low, both the MRVG and MRJD models return values considerably higher than that of the MRD model. Table 2.4.7 lists the monthly delta values derived under each of the processes. Whilst the delta values for most of the months are reasonably consistent, there are large discrepancies in the Dec 12 and Jan 13 figures.

		MRD	MRJD	MRVG
Month	Price (pence/therm)	Delta	Delta	Delta
Dec 12	66.70	-0.3416	-0.3968	-0.3976
Jan 13	66.70	-0.6582	-0.5365	-0.4997
Feb 13	67.20	1.0000	1.0000	0.9999
Mar 13	65.69	0.0001	0.0002	0.0002
Apr 13	63.73	0.0001	0.0002	0.0003
May 13	62.18	0.0002	0.0003	0.0003
Jun 13	60.93	-0.9999	-0.9924	-0.9910
July 13	61.26	0.0002	-0.0068	-0.0082
Aug 13	62.23	0.0002	0.0006	0.0007
Sep 13	65.86	0.0002	0.0002	0.0001
Oct 13	65.13	0.0002	0.0004	0.0005
Nov 13	68.88	0.1503	0.1505	0.1505
Dec 13	71.86	0.8500	0.8500	0.8500

Table 2.4.7: Delta Results

Comparison of the value derivative with respect to monthly forward price (Delta) for the MRD, MRJD, and MRVG valuations. All values are expressed as a proportion of working gas volume.

As discussed above, the main benefit of using the market measure is that it allows one to relate the storage value to the implied volatility surface. This naturally allows one to measure the impact of changes in the volatility surface on the storage value and perform powerful scenario based risk analysis. By way of example, we have re-calibrated the MRVG model and revalued our contract under the following volatility surface scenarios:

- Scenario 1: A positive shift of 10% across all moneyness levels and maturities.
- Scenario 2: A positive shift of 10% across all moneyness levels with 6 months to maturity.
- Scenario 3: A positive shift of 10% at the extreme moneyness levels and 5% at adjacent levels across all maturities.

Scenario 1 reflects a general rise in market volatility, Scenario 2, a decrease in the slope of the volatility term structure, and Scenario 3, an increase in the curvature of the volatility smile across all maturities. The resulting implied volatility surfaces, model parameters, and contract values are detailed in Figure 2.4.6-2.4.8 and Tables 2.4.8-2.4.10.




Figure 2.4.6: Scenario 1: Parallel shift at both maturities

The above two plots display the fit of the MRVG model to the 1 year and 6 month maturity option implied volatility surfaces after applying parallel relative shifts to the original market volatility surface for both maturities and across all moneyness levels.

	Model Parameters		
α	0.2365	Storage Value	11.2343
σ	0.2241	% Δ Base	0.21%
ν	0.2932	% Δ Base Extrinsic	10.46%
SSE	0.14%		

Table 2.4.8: Scenario 1: Parallel Shift

The new MRVG model parameters after applying a parallel shift to the original market implied volatility surface and the resulting impact on storage values. Base refers to the initial value of the storage asset, and Base Extrinsic its extrinsic value, under the MRVG model.

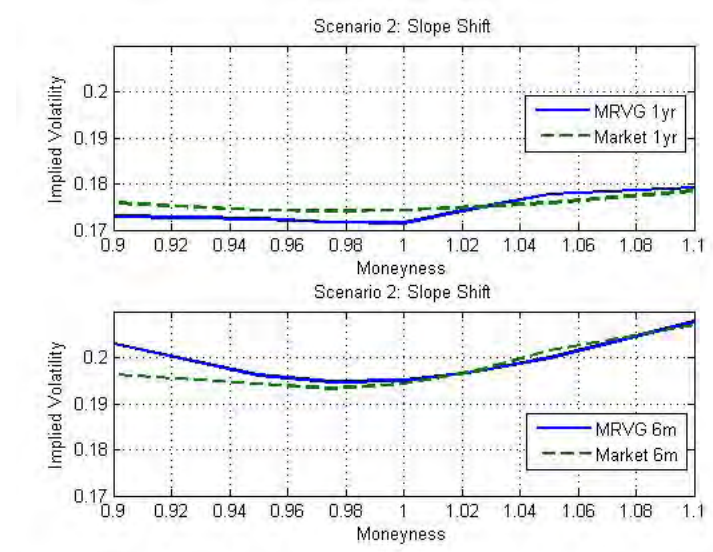


Figure 2.4.7: Scenario 2: Shift at 6m maturity.

The above two plots display the fit of the MRVG model to the 1 year and 6 month maturity option implied volatility surfaces after applying shifts to the original 6 month maturity market implied volatilities only.

	Model Parameters		
α	0.7083	Storage Value	11.8687
σ	0.2520	% Δ Base	5.87%
ν	0.2874	% Δ Base Extrinsic	289.31%
SSE	0.13%		

Table 2.4.9: Scenario 2: Volatility Slope Shift

The new MRVG model parameters after applying a shift to the original market implied volatility surface at the 6 month maturity only and the resulting impact on storage value. Base refers to the initial value of the storage asset, and Base Extrinsic its extrinsic value, under the MRVG model.

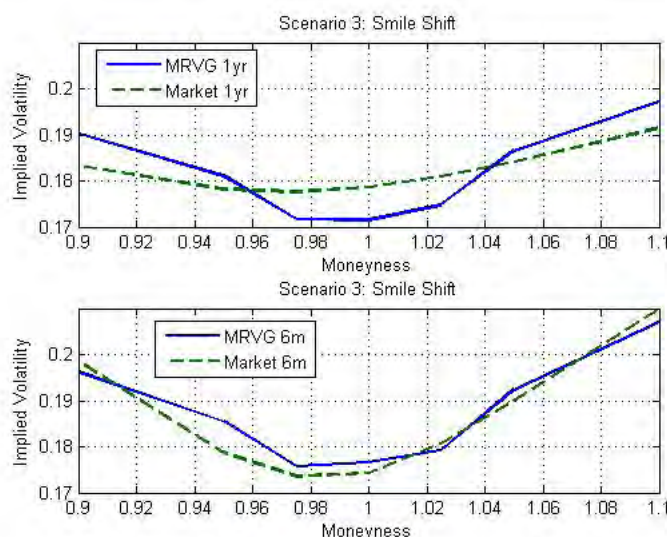


Figure 2.4.8: Scenario 3: Parallel shift of OTM strikes at both maturities.

The above two plots display the fit of the MRVG model to the 1 year and 6 month maturity option implied volatility surfaces after applying relative shifts to the extreme moneyness levels of the original market volatility surface for both maturities.

	Model Parameters		
α	0.3406	Storage Value	11.3283
σ	0.2338	% Δ Base	1.05%
ν	0.9026	% Δ Base Extrinsic	51.79%
SSE	0.96%		

Table 2.4.10: Scenario 3: Parallel Smile Shift

The new MRVG model parameters after applying a shift to the extreme moneyness levels of the original market implied volatility surface and the resulting impact on storage values. Base refers to the initial value of the storage asset, and Base Extrinsic its extrinsic value, under the MRVG model.

As we can see from Table 2.4.8, the positive shift in implied volatility levels under Scenario 1 results in a 10.46% increase in the extrinsic value of the storage deal. This intuitively makes sense as the shift in volatility implies an increase in the gas price variability throughout the life of the contract, allowing the holder to benefit from more lucrative trading opportunities.

The decrease in the slope of the volatility surface under Scenario 2 was driven by an increase in the 6m volatility levels so again one would expect the contract value to increase. In fact we see an increase of 290% over the original extrinsic value (Table 2.4.9). The 1 year implied volatility remaining constant has lead to the spread between winter and summer prices becoming more volatile, thus raising the potential

for more profitable levels at which one can buy and sell gas.

Finally under scenario 3, the increase in the curvature of the implied volatility smile across maturities has resulted in a 52% increase in the extrinsic value of the storage contract (Table 2.4.10), nearly $5\times$ the extrinsic value change due to shifting all levels of volatility. The change in the implied volatility values at extreme moneyness levels implies a greater cost to hedge against extreme price moves. The storage contract would of course benefit from extreme movements in price, particularly spikes at the front of the forward curve which fail to propagate further along the curve. Such scenarios allow the owner, particularly in the case of fast cycling storage, to provide/take gas to an under/over supplied market and benefit from the the large day-ahead versus balance of month/month ahead spreads available.

In conclusion, in this chapter we have presented a thorough introduction to forward curve consistent energy price modelling under Lévy driven Ornstein-Uhlenbeck processes. We gave a detailed exposition one such model, the mean reverting Variance Gamma process (MRVG), and demonstrated its ability to replicate the implied volatility smile present in the NBP gas options market. We then proceeded in deriving a Fourier based storage valuation algorithm and proposed an extension to the current literature, the FFT-m method, which enhances the convergence rate of the algorithm. Finally, we have shown how the MRVG model can be used in conjunction with the FFT-m valuation methodology to efficiently value storage deals and provide scenario based vega analysis on the storage value. In summary, the main contributions to the literature have been

- The derivation of the characteristic function for the mean reverting variance gamma (MRVG) spot price process along with the moments of the MRVG process. The latter results were used to rationalize the selection of the Variance Gamma process as a stochastic driver of our mean reverting model, specifically the relationship between the model parameters and the level of model kurtosis and resulting implied volatility smile attenuation.
- The derivation of the general forward curve consistent characteristic function for mean reverting Lévy spot price processes under the Heath-Jarrow-Morton (HJM) framework. We utilized previous work in the interest rate market literature to derive the implied log spot price dynamics and hence the general forward curve consistent conditional characteristic function of the log spot price process.
- The presentation of a swaption pricing algorithm for general single factor forward curve consistent mean reverting Lévy spot price processes where we extended the diffusion specific methodology of Clewlow and Strickland (1999b). The resulting algorithm allows for quick and efficient calibration to the swaptions market.
- The presentation of an extension of the Fourier Transform based algorithm of Jaimungal and Surkov (2011) for pricing path dependent options under mean reverting processes. Following the

approach of Pan (1993), we present an innovative improvement to the algorithm which greatly enhances the convergence rate and removes the need for a dampening parameter.

In the next chapter we will focus on extending many of the innovations presented thus far to a multifactor setting.

Chapter 3

MultiFactor Storage Valuation

In keeping with recent trends within the storage valuation literature, see Boogert and De Jong (2011) and Parsons (2013) for example, this chapter will focus primarily on extending the modelling framework developed in Chapter 2 to a multifactor setting. We will first present a family of models capable of replicating the complex covariance structure of the natural gas forward curve. While similar model specifications have been studied extensively in the wider energy and interest rate market literature, our approach allows for accurate calibration to market traded options and thus is the first class of multifactor storage models developed with the explicit intention of providing a link between the model and the options market. To that end, we also develop an innovative implied moments based calibration technique which allows for efficient calibration of general multifactor forward curve models. Finally, we extend the storage valuation algorithm presented in Chapter 2 to an arbitrary number of dimensions and present valuation results and analysis of an example storage deal.

3.1 Markov Lévy Driven Forward Curve Models

As stated in the introduction, a major development within the storage valuation literature in recent years has been the need to accurately capture the covariance structure of the commodity forward curve. In practice, there will be a combination of monthly and quarterly products available to monetize the storage value. The price return dependency between these products will fully determine the extrinsic value accruing to the owner of a storage asset. From a modelling perspective it is generally not feasible to model all of these variables individually and therefore a dimensionally reduced representation is chosen to approximate the curve dynamics. Further, in valuing complex path dependent deals, such as storage, choosing a Markov system representation greatly reduces the computational burden on the valuation algorithm.

In this section we will present a class of finite dimensional Markov models of the entire forward curve where the underlying state variables are driven by Lévy processes. The aim being to develop a price model capable of accurately capturing the time-spread volatility, which drives the storage value,

whilst maintaining consistency with the vanilla options market. One could argue that since the storage value is driven by spread volatility rather than outright volatility that consistency with the options market should be a secondary consideration. However, the business case for pricing these risks within a consistent framework should be viewed in the context of a wider and more general derivative book. Consider for example, the addition of a take-or-pay contract to a book containing storage and option positions. These contracts benefit from outright volatility similar to vanilla options and also time spread volatility. A pre-requisite for managing this book in practice would be a model which allows one to price all contracts in a consistent manner and quantify the benefit from the associated risk aggregation. The general modelling framework we will utilize in constructing our model(s) mirrors that of Andersen (2010). Here the author specifies a general jump diffusion model, treating the jump and diffusion components as separate state variables. In the interest of minimizing the dimensionality and maximizing the parsimony of our models we will derive our state variable equations under the assumption that both components can be captured by a single variable, thus generalizing further the work of Andersen (2010).

We demonstrate the flexibility of this framework through the specification of a number of models based upon the historical dynamics of the NBP gas forward market. Each model can be viewed as a generalization of the single factor mean reverting Variance-Gamma model set out in Chapter 2 and thus represent a unique family of models designed to reflect the rich dynamics of the forward market and also calibrate accurately to the option implied volatility surface. This contrasts the recent storage valuation literature, see Boogert and De Jong (2011) and Bjerksund et al. (2008) for example, where the focus has been on capturing the historical covariance structure of the forward curve using models which are typically driven by diffusion processes and, as demonstrated in Chapter 2, will thus fail to replicate the implied volatility smile observable in natural gas markets.

In order to price storage assets, and other derivative contracts, we present the forward curve consistent characteristic functions associated with our proposed forward curve models. Further, we provide an array of utility methods such as moment formulae and the factor implied forward price in the Appendix. These results should serve as a useful reference for practitioners, allowing them to utilize these models to price a wide range of potential derivative deals.

As in the single factor model derived previously, we begin by specifying the dynamics of the log-forward price $f(t, T)$

$$df(t, T) = \left(-\frac{1}{2} \beta(t, T) \beta(t, T)^\top - \kappa_j(\gamma(t, T)) \right) dt + \beta(t, T) d\vec{W}(t) + \int_{\mathbb{R}^K/0} \gamma(t, T) j \vec{v}(dj, dt) \quad (3.1.1)$$

where $\vec{W}(t)$ is now a K-dimensional Brownian Motion, $J(\vec{t}) = \int_{\mathbb{R}^K/0} j \vec{v}(dj, dt)$ a K-dimensional pure

jump process and κ_j is the cumulant function of the jump process. If the \mathcal{F}_t measurable¹ continuous differentiable functions $\beta(t, T)$ and $\gamma(t, T)$ can be separated in terms of time and maturity such that

$$\begin{aligned}\beta(t, T) &= \zeta(T) \eta(t) \\ \gamma(t, T) &= \zeta(T) \theta(t)\end{aligned}\tag{3.1.2}$$

with $\eta(t)$ and $\theta(t)$ both $f: \mathbb{R} \rightarrow \mathbb{R}^{K \times K}$ and $\zeta(T) f: \mathbb{R} \rightarrow \mathbb{R}^K$ then the evolution of the entire forward curve can be represented via a Markov dynamical system of K state variables. In which case we can write 3.1.1 as

$$df(t, T) = \sum_{k=1}^K dy^{(k)}(t, T)\tag{3.1.3}$$

where

$$\begin{aligned}dy^{(k)}(t, T) &= \left(-\frac{1}{2} \beta^{(k)}(t, T)^2 - \kappa_j \left(\gamma^{(k)}(t, T) \right) \right) dt + \beta^{(k)}(t, T) dW^{(k)}(t) \\ &\quad + \int_{\mathbb{R}/0} \gamma^{(k)}(t, T) j \bar{v}^{(k)}(dj, dt)\end{aligned}\tag{3.1.4}$$

and the superscript indicates the k^{th} element of a given vector.

In what follows we assume independence between the sources of randomness in our model. Imposing dependence amongst Lévy processes is well covered in the literature and is typically handled using a Copula, see for example Cont and Tankov (2004) for the general case and Luciano and Semeraro (2010) for the specific case of the Variance Gamma process. Marfè (2009) also study the multivariate Variance Gamma process and incorporate dependency using independent multivariate Gamma processes. As we will show, this restriction to independent processes does not prohibit us from capturing the inter-maturity pairwise dependency of forward curve returns to a high level of accuracy. Further, it does allow us to specify the effect of each state variable on the forward curve dynamics independently, which is in agreement with traditional Principal Component based analysis of forward curve movements. Utilizing PCA in multifactor model specifications is common in energy markets, for example see Clewlow and Strickland (1999a), Carmona and Coulon (2014). However, the approach is more established in the interest rate markets, where Litterman and Scheinkman (1991) were the first to identify the “Shift, Twist and Bend” characterization of the first three principal components.

The primary benefit of this modelling approach is the immense practical appeal of reducing the forward curve dynamics to a small number of distortions to the shape of the curve whilst still capturing the full complexity of the curve variability to a high level of accuracy. In contrast, multifactor models with complex interdependency between the state variables will typically not yield such an intuitive interpretation

¹ \mathcal{F}_t refers to the filtration defined on the probability space $(\Omega, \mathcal{F}, \mathbb{Q})$

of the impact of each latent variable on the forward curve dynamics and ultimately, the impact on derivative values. Driessen et al. (2003) present a range of multifactor diffusion models where the specification of each factor is chosen so as to replicate the sensitivity of the forward curve returns to the principal components. We intend to follow this data-driven approach to model specification when determining the structure of our class of multifactor models.

As in the single factor case, we need to derive the forward curve consistent dynamics of the log spot price process $x(t) = f(t, t)$,

$$x(t) = f(0, t) + \sum_{k=1}^K y^{(k)}(t) \quad (3.1.5)$$

where $y^{(k)}(t) = y^{(k)}(t, t)$.

Applying the same solution technique given for the single factor model in Section A.4 of the Appendix, we have the following²

$$\begin{aligned} dy^{(k)}(t) &= \left(-\frac{1}{2} \eta^{(k)}(t)^2 \zeta^{(k)}(t)^2 - \zeta^{(k)'}(t) \int_0^t \eta^{(k)}(s)^2 \zeta^{(k)}(s) ds - \kappa_j \left(\theta^{(k)}(t) \zeta^{(k)}(t) \right) \right. \\ &\quad \left. - \zeta^{(k)'}(t) \int_0^t \theta^{(k)}(s) \kappa_j' \left(\theta^{(k)}(s) \zeta^{(k)}(s) \right) ds \right. \\ &\quad \left. + \left(y^{(k)}(t) + \int_0^t \kappa_j \left(\theta^{(k)}(s) \zeta^{(k)}(s) \right) ds \right) \frac{\zeta^{(k)'}(t)}{\zeta^{(k)}(t)} dt \right. \\ &\quad \left. + \eta(t) \zeta(t) dW^{(k)}(t) + \int_{\mathbb{R}/0} \theta^{(k)}(t) \zeta^{(k)}(t) j \bar{v}^{(k)}(dj, ds) \right) \\ &= \left(\omega^{(k)}(t) + \frac{\zeta^{(k)'}(t)}{\zeta^{(k)}(t)} y^{(k)}(t) \right) dt + \eta(t) \zeta(t) dW^{(k)}(t) \\ &\quad + \int_{\mathbb{R}/0} \theta^{(k)}(t) \zeta^{(k)}(t) j \bar{v}^{(k)}(dj, dt) \end{aligned} \quad (3.1.6)$$

Let $\Sigma^{(k)}(t)$ denote the exponential of the anti-derivative of $-\frac{\zeta^{(k)'}(t)}{\zeta^{(k)}(t)}$. Integrating, we have

$$\begin{aligned} y^{(k)}(t) &= y^{(k)}(s) \frac{\Sigma^{(k)}(s)}{\Sigma^{(k)}(t)} + \int_s^t \omega(u) \frac{\Sigma^{(k)}(u)}{\Sigma^{(k)}(t)} du \\ &\quad + \int_s^t \frac{\Sigma^{(k)}(u)}{\Sigma^{(k)}(t)} \eta(u) \zeta(u) dW^{(k)}(u) + \int_s^t \frac{\Sigma^{(k)}(u)}{\Sigma^{(k)}(t)} \int_{\mathbb{R}/0} \theta^{(k)}(u) \zeta^{(k)}(u) j \bar{v}^{(k)}(dj, du) \end{aligned}$$

If the state variables are independent then the conditional characteristic function for the state vector $\vec{y} = \left[y^{(k)} \right]_{k=1}^K$ can be recovered easily as the product of each variable's characteristic function. Thus,

²Here, $f'()$ represents the derivative of $f()$.

$$E \left[\exp \left(i \vec{z}^T \cdot \vec{y}(t) \right) | \vec{y}(s) \right] = \Phi_{\vec{y}(t)} \left(\vec{z}; \vec{y}(s), s, t \right) = \prod_{k=1}^K \Phi_{y^{(k)}(t)} \left(z^{(k)}; y^{(k)}(s), s, t \right) \quad (3.1.7)$$

We will now show how the above framework can be used when constructing Markov price models by reference to the NBP Gas forward market. To provide some intuition on an appropriate specification for the functions $\beta^{(k)}(t, T)$ and $\gamma^{(k)}(t, T)$, we will perform a Principal Component Analysis on the historic relative maturity forward price returns for NBP Gas, a detailed overview of the methodology used is given in Section 2.2.5 of Carmona and Coulon (2014). Our data set covers the period 19th December 2009 - 19th December 2012. We begin by constructing a series of relative maturity returns beginning with the day-ahead contract and increasing in intervals of 30 days out to 1 year. These returns shall then be used to construct our forward curve covariance matrix. Figure 3.1.1 displays the inter-maturity correlations between each of our relative maturity contracts. The most obvious point to take from this plot is that the correlation coefficient between months decreases as the time between their delivery periods increases.

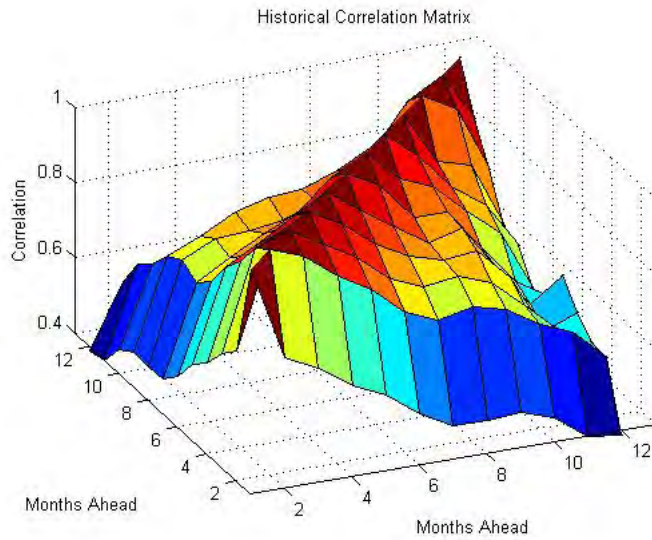


Figure 3.1.1: NBP Gas Forward Curve Returns Correlation

The correlation matrix of the daily returns on a set of relative maturity monthly contracts.

Performing a standard Principal Component Analysis shows that the first two principal components account for approximately 96% of the total variance of the forward curve. We define the factor volatility function for the forward curve as the sensitivity of a forward price log-return to a given principal component multiplied by the principal component's standard deviation. Formally, we have

$$v^{(k)}(\tau) = w^{(k)}(\tau) \sqrt{\lambda^{(k)}}$$

where $w^{(k)}(\tau)$ is the sensitivity of the relative maturity τ forward price return to the k^{th} principal component with variance $\lambda^{(k)}$. The motivation for defining the factor volatility curves in this manner is that, by definition of the eigenvector and eigenvalues, the variance of the τ maturity forward log returns is equal to the sum of squared τ maturity factor volatilities. This naturally implies a model of the log forward price equal to the sum of the independent factors, as given by Equation 3.1.5.

Thus, these historical volatility functions are analogous to the functions $\beta^{(k)}(t, T)$ and $\gamma^{(k)}(t, T)$ defined above and can be used to impose an appropriate functional form on our model which accurately reflects the curve dynamics. Figure 3.1.2 displays the values of both the 1st and 2nd historical volatility functions. The function values have been scaled by the sample spot volatility over the estimation period. The reason for this is to allow us, if required, to scale the functions by a spot volatility value obtained from calibrating to the options market.

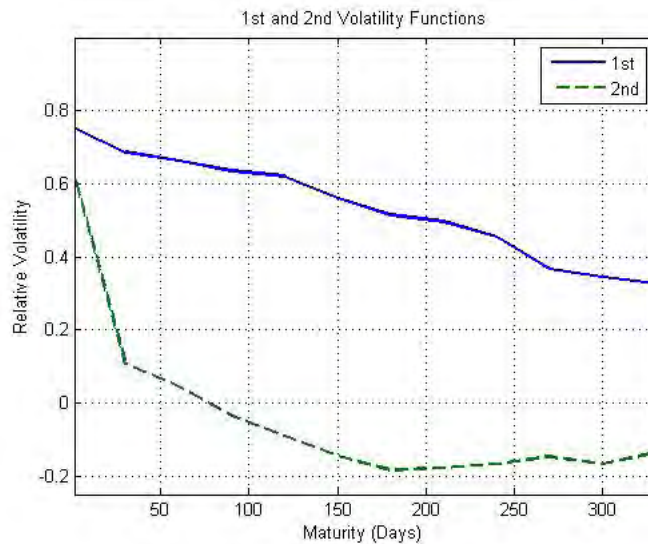


Figure 3.1.2: Historical Volatility Function values.

The first and second volatility functions refer to the scaled eigenvector values associated with the first two principal components of the NBP forward curve returns. They follow the “Shift” and “Twist” characterization typically observed in energy and interest rate markets.

The first volatility function which accounts for 75% of the total variance resembles the negative exponential volatility function associated with the traditional single factor model, thus,

$$\gamma^{(1)}(t, T) = b \exp(-\alpha(T - t))$$

The shape of the second volatility function, which accounts for a further 12% of the curve variability, explains the imperfect correlation between short and long maturities as it induces movements of similar

magnitude but with opposing directions at the near and far end of the curve. This curve can be approximated as negative exponential tending to a negative asymptote. The concept of negative volatility, whilst meaningless in the context of a single factor diffusion process, should be viewed in the context of the wider forward curve model. Essentially, what this factor is replicating is the simultaneous occurrence of positive(/negative) price returns at the short term maturities and negative(/positive) price returns at the longer term maturities. Formally, we have

$$\gamma^{(2)}(t, T) = \exp(-\varepsilon(T-t))(c_1 - c_2) + c_2$$

with $\alpha, \varepsilon > 0$. This forward volatility term structure is quite common in the literature, see Andersen (2010) and Cheyette (2001) for example.

Similar to the single factor MRVG model presented in Chapter 2, we will model the primary source of log forward curve variation using a Variance Gamma process³, $\{dX(t); (0, 1, \nu)\}$. The secondary source of variation is modeled using a Diffusion process, $\{dW(t); (0, 1)\}$. This will enable us to approximate the market smile for vanilla options whilst also enforcing an inter-maturity covariance structure on the model. The model is then given by the following system of stochastic differential equations (SDE)

$$\begin{aligned} dy^{(1)}(t, T) &= \left(-\kappa_{j(1)}(b \exp(-\alpha(T-t)) \sigma(t)) \right) dt + b \exp(-\alpha(T-t)) \sigma(t) dX(t) \\ dy^{(2)}(t, T) &= -\frac{1}{2} \left((\exp(-\varepsilon(T-t))(c_1 - c_2) + c_2) \sigma(t) \right)^2 dt \\ &\quad + (\exp(-\varepsilon(T-t))(c_1 - c_2) + c_2) \sigma(t) dW(t) \end{aligned}$$

with $E[dXdW] = 0$. Although the model has two stochastic drivers, due to the form of the second volatility function we need three factors in order to obtain a Markovian log spot price representation. This can lead to some confusion in the exact definition of a factor. For our purposes, a factor will always relate to a one dimensional stochastic process with a volatility function satisfying Equation 3.1.2. In what follows, the nomenclature used when identifying different models will reflect this definition of a factor.

We can specify an equivalent 3-dimensional system of SDE's representing the log-forward price dynamics as follows,

$$\begin{aligned} dy^{(1)}(t, T) &= \left(-\kappa_{j(1)}(b \exp(-\alpha(T-t)) \sigma(t)) \right) dt + b \exp(-\alpha(T-t)) \sigma(t) dX(t) \\ dy^{(2)}(t, T) &= -\frac{1}{2} \left((\exp(-\varepsilon(T-t))(c_1 - c_2) + c_2) \sigma(t) \right)^2 dt \\ &\quad + (\exp(-\varepsilon(T-t))(c_1 - c_2)) \sigma(t) dW(t) \\ dy^{(3)}(t, T) &= c_2 \sigma(t) dW(t) \end{aligned} \tag{3.1.8}$$

³Recall that the Variance-Gamma process is parametrized by (θ, σ, ν) , for our purposes we have explicitly set $\theta = 0$ and $\sigma = 1$.

The dynamics of $dy^{(i)}(t)$ for $i = 1, 2, 3$ are then given by Equation 3.1.6 as⁴,

$$\begin{aligned} dy^1(t) &= \left(\omega^{(1)}(t) - \alpha y^{(1)}(t) \right) dt + b\sigma(t) dX(t) \\ dy^2(t) &= \left(\omega^{(2)}(t) - \varepsilon y^{(2)}(t) \right) dt + (c_1 - c_2) \sigma(t) dW(t) \\ dy^3(t) &= c_2 \sigma(t) dW(t) \end{aligned} \quad (3.1.9)$$

Note that if the percentage of spot variance explained by the two principal components is given by s then we have $b = \sqrt{s - c_1^2}$. We shall refer to this model as MRVG-3. We can no longer write the characteristic function for our state vector according to Equation 3.1.7 as the second and third state variables are dependent. However, as they are driven by the same source of uncertainty we can write the characteristic function for $\vec{y}(t)$ with time independent parameters as,

$$\begin{aligned} \Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(s), s, t) &= \exp\left(iz^{(1)}y^{(1)}(s)\exp(-\alpha(t-s)) + iz^{(2)}y^{(2)}(s)\exp(-\varepsilon(t-s)) + iz^{(3)}y^{(3)}(s)\right) \times \\ &\exp\left(iz^{(1)}\int_s^t \omega^{(1)}(c)\exp(-\alpha(t-c))dc + iz^{(2)}\int_s^t \omega^{(2)}(c)\exp(-\varepsilon(t-c))dc\right) \times \\ &E\left[\exp\left(\int_s^t iz^{(1)}(\exp(-\alpha(t-c)))b\sigma dX(c)\right)\right] \times \\ &E\left[\exp\left(\int_s^t i\left(\left(z^{(2)}\exp(-\varepsilon(t-s))(c_1 - c_2) + z^{(3)}c_2\right)\sigma\right)dW(c)\right)\right] \end{aligned} \quad (3.1.10)$$

$$\text{for } \vec{z} \in S_y \cup \mathbb{C}^2, S_y := \left\{ z = u + iw; w \in \left(-\sqrt{\frac{2}{(b\sigma)^2 v}}, \sqrt{\frac{2}{(b\sigma)^2 v}} \right) \right\}$$

We can reduce the dimensionality of this model by explicitly setting $c_2 = 0$, in which case we will refer to it as MRVG-3x. This model approximates the second volatility function as a negative exponential tending to zero rather than the negative volatility asymptote associated with the MRVG-3 model. This is of course a misrepresentation of the nature of the second volatility function which will lead to an increase in the correlation between short and long maturity forward prices compared to the MRVG-3x model. However, the reduction in dimensionality is justified by the significant decrease in the computational burden when using the model to value complex derivative assets.

Alternatively, if we were to set $\alpha = \varepsilon$, so that the decay of both the first and second volatility functions are equal, the model reduces to the following two dimensional system of SDE's

$$\begin{aligned} dy^{(1)}(t, T) &= \left(-\kappa_j(\exp(-\alpha(T-t))b\sigma(t)) - \frac{1}{2}((\exp(-\alpha(T-t))(c_1 - c_2) + c_2)\sigma(t))^2 \right) dt \\ &\quad + \exp(-\alpha(T-t))[b\sigma(t)dX(t) + (c_1 - c_2)\sigma(t)dW(t)] \\ dy^{(2)}(t, T) &= c_2\sigma(t)dW(t) \end{aligned} \quad (3.1.11)$$

⁴The functions $w^{(i)}(t)$ for each model specification are given in Section B.1 of the Appendix

The dynamics of $dy^{(i)}(t)$ for $i = 1, 2$ are then

$$\begin{aligned} dy^{(1)}(t) &= \left(\omega^{(1)}(t) - \alpha y^{(1)}(t) \right) dt + b\sigma(t) dX(t) + (c_1 - c_2)\sigma(t) dW(t) \\ dy^{(2)}(t) &= c_2\sigma(t) dW(t) \end{aligned} \quad (3.1.12)$$

We shall refer to this model as MRVG-2. The characteristic function is given by

$$\begin{aligned} \Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(s), s, t) &= \exp\left(iz^{(1)}y^{(1)}(s)\exp(-\alpha(t-s)) + iz^{(2)}y^{(2)}(s)\right) \times \\ &\exp\left(iz^{(1)}\int_s^t \omega^{(1)}(c)\exp(-\alpha(t-c))dc\right) \times \\ &E\left[\exp\left(\int_s^t iz^{(1)}(\exp(-\alpha(t-c)))b\sigma dX(c)\right)\right] \times \\ &E\left[\exp\left(\int_s^t i\left(\left(z^{(1)}\exp(-\alpha(t-c))(c_1 - c_2) + z^{(2)}c_2\right)\sigma\right)dW(c)\right)\right] \end{aligned} \quad (3.1.13)$$

for $\vec{z} \in S_y \cup \mathbb{C}$.

Similar to the MRVG-3x model, the MRVG-2 model has the extremely desirable property that there are only two state variables as opposed to the three required for the MRVG-3 model. This of course comes at a cost of constraining the second volatility function and by extension the inter-month correlation implied by the model. We can reduce the MRVG-2 model further by constraining the parameters b, c_1 and c_2 such that $b = 1, c_1 = c_2 = \frac{\sigma_L(c)}{\sigma(c)}$ where $\sigma_L(c)$ is the long-term volatility of the spot price. This model, which we will refer to as MRVG-2x, can be calibrated to the options market alone and as such should be viewed as an extension of the MRVG model presented in Chapter 2. For completeness, the characteristic function is given by

$$\begin{aligned} \Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(s), s, t) &= \exp\left(iz^{(1)}y^{(1)}(s)\exp(-\alpha(t-s)) + iz^{(2)}y^{(2)}(s)\right) \times \\ &\exp\left(iz^{(1)}\int_s^t \omega^{(1)}(c)\exp(-\alpha(t-c))dc\right) \times \\ &\exp\left(-iz^{(2)}\frac{1}{2}\int_s^t \sigma_L^2 dc\right) \times \\ &E\left[\exp\left(\int_s^t iz^{(1)}(\exp(-\alpha(t-c)))\sigma dX(c)\right)\right] \times \\ &E\left[\exp\left(\int_s^t iz^{(2)}\sigma_L dW(c)\right)\right] \end{aligned} \quad (3.1.14)$$

3.2 Joint Model Calibration and Estimation

We will now present results of a joint market based calibration and historical estimation of the class of models introduced in the previous section. We first discuss calibrating general multidimensional forward curve models to the market for natural gas swaptions. For our purposes we focus on monthly options, however, the proposed calibration methodology can easily be applied to swaptions of varying delivery periods. As such, this approach can be seen as a generalization of the methodology of Guillaume and Schoutens (2013) where the authors focus on single delivery options. The proposed approach is particularly appealing in that the computational effort is independent of the number of dimensions in the underlying model.

As discussed in Section 3.1, the shape of our factor volatility functions will be estimated directly from the eigenvalue decomposition of the forward curve covariance matrix. This estimation method is a deliberately simplified, almost an heuristic, approach. This again is to reflect practical usage of these models, where unobservable characteristics of the model are typically anchored around historical estimates and then adjusted to attain general consistency with the market. As argued in the previous section, the motivation for utilizing historical model estimation is the need to accurately forecast the cost of dynamically hedging time spread risk through trading in the forward market. Similarly, market based model calibrations are needed to incorporate the cost of hedging the outright volatility or “Vega” risk. Utilizing both approaches allows one to incorporate estimates of the cost of hedging against the primary risk factors affecting natural gas derivatives. Alternatively, if one wished to estimate a statistical “fair value” for a given asset using historical data alone, the Kalman Filter is generally considered the optimal approach for estimating the parameters of general exponentially affine factor models. We refer the interested reader to Duan and Simonato (1999) and De Jong (2000) for an introduction to the use of the Kalman Filter in this context. More recently, Dempster and Tang (2011) discuss the dependency of the measurement error specification on the parameter estimates and provide evidence of serial correlation in the measurement errors in both commodity and interest rate markets.

3.2.1 Market Implied Moments

Before proceeding to the calibration of the general class of models described above we will first review, in general, the proposed method for calibrating these models to market. Due to the complexity of the models it will generally not be feasible to utilize a swaption pricing algorithm to calibrate to monthly swaptions. Therefore we propose using a variation on the implied moments technique of Guillaume and Schoutens (2013). Using the formula of Bakshi and Madan (2000) any twice differentiable payoff

function on a price F_T can be expressed as

$$\begin{aligned} v(F_T) &= v(K_0) + v'(K_0)(F_T - K_0) \\ &\quad + \int_{K_0}^{\infty} v''(K)(F_T - K)^+ dK \\ &\quad + \int_0^{K_0} v''(K)(K - F_T)^+ dK \end{aligned}$$

and therefore by taking expectations we have

$$\begin{aligned} E[v(F_T)] &= v(K_0) + v'(K_0)(F_t - K_0) \\ &\quad + \int_{K_0}^{\infty} v''(K)C(F_t, K, T) dK \\ &\quad + \int_0^{K_0} v''(K)P(F_t, K, T) dK \end{aligned} \quad (3.2.1)$$

Defining the payoff function as $v(F_T) = \ln\left(\frac{F_T}{F_t}\right)^n$, the formula for the n^{th} market moment is then given by Guillaume and Schoutens (2013) as

$$\begin{aligned} E\left[\ln\left(\frac{F_T}{F_t}\right)^n\right] &= \ln\left(\frac{K_0}{F_t}\right)^n + N \ln\left(\frac{K_0}{F_t}\right)^{n-1} \left(\frac{F_t}{K_0} - 1\right) + \\ &\quad \int_{K_0}^{\infty} \frac{n}{K^2} \left[(n-1) \ln\left(\frac{K}{F_t}\right)^{n-2} - \ln\left(\frac{K}{F_t}\right)^{n-1} \right] C(F_t, K, T) dK \\ &\quad + \int_0^{K_0} \frac{n}{K^2} \left[(n-1) \ln\left(\frac{K}{F_t}\right)^{n-2} - \ln\left(\frac{K}{F_t}\right)^{n-1} \right] P(F_t, K, T) dK \end{aligned} \quad (3.2.2)$$

which they discretize using a trapezoidal rule such that

$$\begin{aligned} E\left[\ln\left(\frac{F_T}{F_t}\right)^n\right] &= \ln\left(\frac{K_0}{F_t}\right)^n + n \ln\left(\frac{K_0}{F_t}\right)^{n-1} \left(\frac{F_t}{K_0} - 1\right) + \\ &\quad \sum_{i=1}^M \Delta K_i \frac{n}{K_i^2} \left[(n-1) \ln\left(\frac{K_i}{F_t}\right)^{n-2} - \ln\left(\frac{K_i}{F_t}\right)^{n-1} \right] Q(K_i) \end{aligned} \quad (3.2.3)$$

where

- K_0 is the at-the-money strike or $\sup i : K_i \leq F_t$

- $Q(K_i)$ is the mid-price of the option with strike K_i chosen such that

$$\begin{cases} Q(K_i) = P(K_i) & : K_i < K_0 \\ Q(K_i) = \frac{P(K_i) + C(K_i)}{2} & : K_i = K_0 \\ Q(K_i) = C(K_i) & : K_i > K_0 \end{cases}$$

- ΔK_i is the option spacing defined as

$$\begin{cases} \Delta K_1 = K_2 - K_1 \\ \Delta K_i = \frac{K_{i+1} - K_{i-1}}{2} & \forall i \neq 1 \text{ or } M \\ \Delta K_M = K_M - K_{M-1} \end{cases}$$

One drawback of this approach is that there is of course no guarantee that ones model will perfectly reproduce the price of traded options and thus you will in all likelihood bear some model risk. To overcome this, Guillaume and Schoutens (2013) propose a weighting scheme which places more importance on the lower moments and provides a better fit to the original implied volatility curve.

Although it would be desirable to derive the implied moments using only market quoted options, in practice there maybe be valid reasons for interpolating a volatility surface from the existing quotes or a subset thereof, for example, liquidity and arbitrage concerns. If certain products are unlikely to contain trustworthy information relating to the implied pricing measure they need not form part of the implied moment calculation, and similarly, if the mid-prices at certain strikes admit an arbitrage then they too should be excluded from the calculation.

Further, restricting the number of option prices to only those quoted in the market may have a detrimental effect on the accuracy of the numerical integration needed to derive the option implied moments. The trapezoidal rule used above, which approximates the option price surface as a piecewise linear function, would yield an error proportional to the the difference in strike levels and inversely proportional to the number of strikes. Whilst this issue may be somewhat negligible for call and put price functions, the second derivative of the moment payoff function displays a much greater level of non-linearity particularly at higher moments as Figure 3.2.1 demonstrates

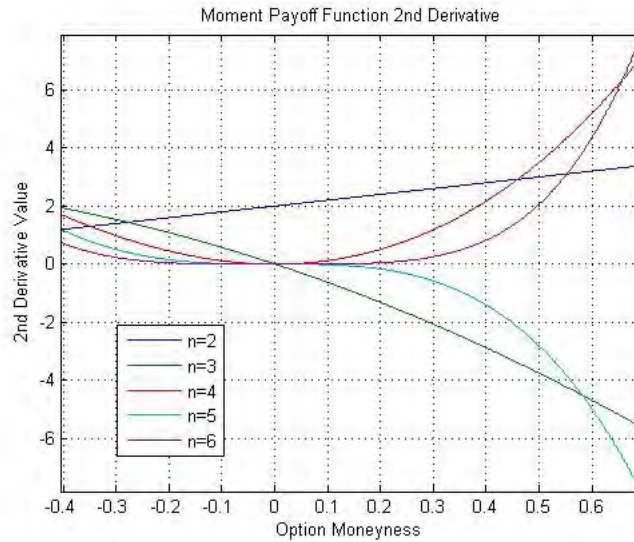


Figure 3.2.1: Second Derivative of Moment Payoff Function

The value of the second derivative of the n^{th} moment payoff function given within the summation in Equation 3.2.3. The second derivative becomes highly non-linear for moments higher than 3.

This issue can be easily overcome by interpolating over the option prices and increasing the granularity of the summation in Equation 3.2.3. Homescu (2011) gives a detailed summary of the literature surrounding implied volatility surface construction and interpolation of call prices / implied volatilities. Piecewise polynomial or spline interpolations are capable of meeting the requirements of no call spread or calendar spread arbitrage, although the latter condition is rarely a concern for energy markets. One benefit of interpolating over implied volatility versus prices even if both are arbitrage free, is that relative value is judged on an implied volatility basis and visual inspection can be used to verify the interpolated values, something which is far more difficult when looking at a call price function. Further, extrapolation may often be required for many of the same reasons highlighted above, although this will obviously introduce an element of model risk into calibration. The easiest choices of course are to extrapolate flat or linearly, where the latter at least contains some market based information. However, we know from the moment formula for extreme strikes of Lee (2004) that the slope of the implied integrated variance as a function of moneyness has a limit determined by the number of finite moments. It is therefore consistent to assume the number of finite moments is at least equal to the number of moments you wish to imply for the purposes of calibration. The slope formula for high strikes is given by

$$\beta_r = 2 - 4 \left(\sqrt{\tilde{p}^2 + \tilde{p}} - \tilde{p} \right)$$

and for lower strikes

$$\beta_l = 2 - 4 \left(\sqrt{\tilde{q}^2 + \tilde{q}} - \tilde{q} \right)$$

where

$$\tilde{p} = \sup \{p : E [F^{1+p} < \infty]\}, \quad \tilde{q} = \sup \{q : E [F^{-q} < \infty]\}$$

Flat extrapolation is equivalent to finite moments of all order whereas linear interpolation with a slope equal to 2 implies no finite moments greater than one. Figure 3.2.2 displays the behavior of the upper slope value at different levels of \tilde{p} .

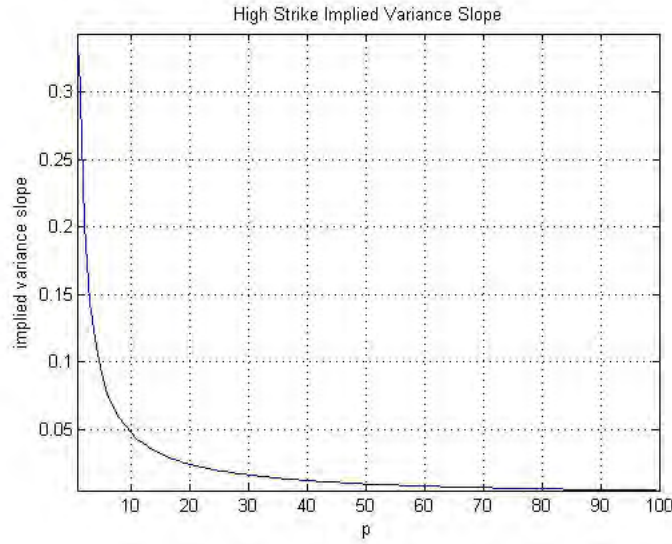


Figure 3.2.2: High Strike Variance Slope

The relationship between the slope of the implied variance of option prices at very high strike prices (y-axis) and the number of finite moments greater than 1 (x-axis) of the underlying probability distribution.

In the context of natural gas markets, where option prices are generally most liquid on monthly products, the implied moments approach allows us to calibrate an instantaneous forward price model to the monthly implied volatilities without the need for a computationally expensive pricing algorithm. As long as the moments for the instantaneous forward price, $F(t, T)$, are known, the model parameters can be fitted to the implied moments of the monthly forward price such that

$$E [F(t, T_1, T_2)^n] = E \left[\left(\sum_{j=0}^{N-1} \frac{1}{N} F(t, T_1 + j\Delta t) \right)^n \right] \quad (3.2.4)$$

where t is the expiry date of the monthly option delivering over N periods beginning at time T_1 .

The right hand side can be implied from the options market using a simplified version of Equation 3.2.2,

$$E[F^n] = K_0^n + nK_0^{n-1}(F_t - K_0) + \int_K^\infty n(n-1)K^{n-2}C(F_t, K, T) dK + \int_0^K n(n-1)K^{n-2}P(F_t, K, T) dK$$

which can be discretized as before to give

$$E[F^n] = K_0^n + nK_0^{n-1}(F_t - K_0) + \sum_{i=1}^M \Delta K_i n(n-1)K_i^{n-2}Q(K_i)$$

For the purposes of comparison it may be preferable to scale these values by the n^{th} power of the underlying forward price, in which case the above can be written as

$$\frac{1}{F_0^n} E[F^n] = \frac{K_0^n}{F_0^n} + n \frac{K_0^{n-1}}{F_0^n} (F_0 - K_0) + \frac{1}{F_0^n} \sum_{i=1}^M \Delta K_i n(n-1)K_i^{n-2}Q(K_i) \quad (3.2.5)$$

Equation 3.2.4 can be easily solved for any Markov forward curve model as demonstrated in Section B.8 of the Appendix.

3.2.2 Calibration Example: NBP Monthly Options market

We will now demonstrate the implied moments calculation, taking as input the 6 month and 1 year maturity NBP monthly volatility surface as of the 19th December 2012. We begin by interpolating/extrapolating the volatility surface over a set of equally spaced strikes. We have chosen linear interpolation of the variance to fill in the missing points within the original surface and also linear extrapolation of the variance for points outside. For the 1 year maturity the slopes at the very low and very high strikes were 0.0034 and 0.0111 respectively. These values correspond to 144 inverse moments and 45 moments in the underlying distribution. Similarly, the slopes for the 6 month maturity were 0.0051 and 0.0175, which correspond to 98 inverse moments and 28 moments respectively. In order to calculate the implied moments put and call option price surfaces were required and these are shown in the first row of Figure 3.2.3. The second row of Figure 3.2.3 shows the extrapolated volatility smiles for the 6 month and 1 year options.

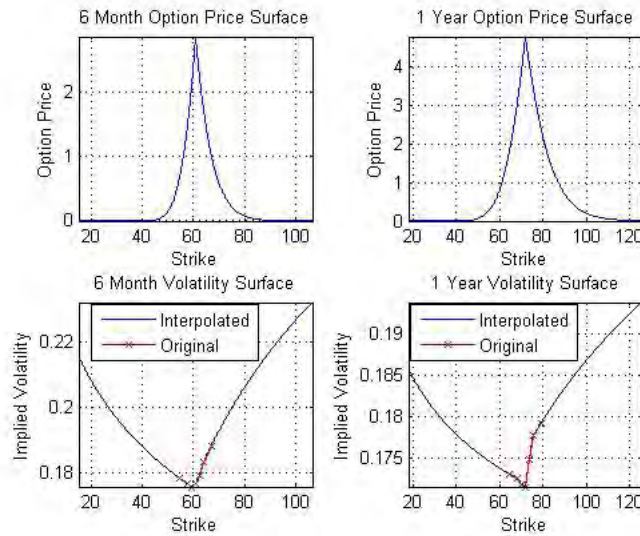


Figure 3.2.3: NBP Vol Surface: 19th December 2012

The top two panels above display the interpolated option price function at the 6 month and 1 year maturities. For strikes greater than the current forward price, the function returns the call price and conversely, the put price for strikes below the current forward price. The bottom two panels display the associated implied volatility functions for both maturities along with the original implied volatility values.

We will now proceed to demonstrating the calibration procedure for three potential spot price models as set out previously: MRVG-2x, MRVG-2 and MRVG-3x. To fit the implied moments we need the monthly forward price moments under each model. For a general exponential model of the form

$$F(t, T) = F(0, T) \exp(\vec{y}(t, T))$$

the monthly price is given by

$$F(t, T_1, T_2) = E \left[\frac{1}{N} \sum_{j=0}^{N-1} F(0, T_1 + j\Delta t) \exp(\vec{y}(t, T_1 + j\Delta t)) | y(\vec{0}) \right]$$

for N delivery days between times T_1 and T_2 . We wish to calibrate the 2^{nd} to 4^{th} moments. The general formulae are presented below and the exact formulae for each model of interest are given in Section B.8 of the Appendix.

$$\begin{aligned}
M_2 &= \frac{1}{F(0, T_1, T_2)^2} E \left[\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} F(0, T_1 + i\Delta t) F(0, T_1 + j\Delta t) \times \right. \\
&\quad \left. \exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t)) |y(\vec{0}) \right] \\
M_3 &= \frac{1}{F(0, T_1, T_2)^3} E \left[\frac{1}{N^3} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} F(0, T_1 + i\Delta t) F(0, T_1 + j\Delta t) F(0, T_1 + k\Delta t) \times \right. \\
&\quad \left. \exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t) + y(t, T_1 + k\Delta t)) |y(0) \right] \\
M_4 &= \frac{1}{F(0, T_1, T_2)^4} E \left[\frac{1}{N^4} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} F(0, T_1 + i\Delta t) F(0, T_1 + j\Delta t) F(0, T_1 + k\Delta t) F(0, T_1 + l\Delta t) \times \right. \\
&\quad \left. \exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t) + y(t, T_1 + k\Delta t) + y(t, T_1 + l\Delta t)) |y(\vec{0}) \right]
\end{aligned}$$

Unlike the approach specified in Guillaume and Schoutens (2013) for simple payoffs, we will fit the moments using an optimization algorithm. This is due to the nature of the moment formula and the lack of a straightforward algebraic solution. We have used the Simplex algorithm in order to minimize the objective function in question. To benchmark the approach we begin by calibrating the MRVG model, which we have calibrated previously to the same data set by matching the underlying option prices.

Method	α	σ	ν	SSE(option prices)
Option Prices	0.2162	0.201	0.2560	0.16%
Moments	0.1668	0.1926	0.1626	0.75%

Table 3.2.1: MRVG Model Parameters

The model parameters for the MRVG model estimated using the implied moments technique and also by calibrating to the market prices of options. Here, 'SSE' refers to the sum of squared percentage pricing errors, which was the metric used when fitting the model.

As we can see from Table 3.2.1, the fit in terms of option prices is worse, although this isn't surprising given that the optimization is no longer attempting to reproduce the option prices directly. Essentially, the moment matching methodology in this setting gives one freedom in choosing a potentially more appropriate multifactor forward curve model which is generally consistent with the options market at the cost of accuracy in reproducing mid-prices in agreement with the market⁵. Moving on to the MRVG-2x model, Table 3.2.2 displays the fitted parameters, alongside the MRVG model and includes the sum of squared percentage error between the model and market implied moments which was used as the objective function to be minimized when optimizing the model parameters.

⁵Of course this doesn't necessarily mean the model will offer an arbitrage; this would only occur if the resulting option prices lay outside bid-offer.

Model	α	σ	ν	σ_L	SSE(implied moments)	SSE(option prices)
MRVG	0.1668	0.1926	0.1626	-	1.4299×10^{-8}	0.75%
MRVG-2x	0.3130	0.1479	0.5653	0.1238	5.7964×10^{-9}	0.85%

Table 3.2.2: MRVG-2x Model Parameters

The model parameters for the MRVG and MRVG-2x models estimated using the implied moments technique. Here, 'SSE(implied moments)' refers to the sum of squared percentage error between the model and market moments, which was the metric used when fitting the model. 'SSE(option prices)' refers to the sum of squared percentage errors between the model and market option prices.

As you would expect, the mean reversion rate is now much higher as the model has the additional freedom to replicate the variability of the forward curve using the long term volatility. The overall fit to the implied moments is also much better, again due to the additional degree of freedom gained by the inclusion of the second factor. The sum of squared percentage errors on the option prices is of similar magnitude to the MRVG case at 0.85%.

Given the lack of a liquid market for time-spread options, we will now outline a methodology for calibrating the MRVG-2 and MRVG-3x models using a combination of historical and market data. Central to this approach is the idea outlined in Section 3.1 of firstly capturing the covariance structure through the principal components of the forward curve returns and secondly, expressing the historical volatility functions of each factor relative to the spot volatility. This allows us to re-scale any of the factors we wish to estimate from history by a market implied spot price volatility. In general, the approach can be enumerated as follows

1. Construct the covariance matrix of historical forward curve relative maturity returns.
2. Scale the returns by the historical spot price volatility.
3. Perform Eigenvalue Decomposition on the covariance matrix.
4. Select the number of factors you wish to model and derive a relative volatility function for each.
5. Select the factors which you can calibrate using market inputs, typically, this will only be the first.
6. Estimate the remaining model parameters by matching the shape of the historical volatility functions for the relevant factors.

In the case of the MRVG-2 model, we will estimate σ and ν by calibrating to the NBP monthly options market implied moments. Therefore $\{b, \alpha, c_1, c_2\}$ need to be estimated from the intercept of the first volatility function and the shape of the second volatility function, given by $c_1 \exp(-\alpha\tau) + c_2$. We can estimate these values by fitting the model volatility function to the historical second volatility function

curve presented in Figure 3.1.2. The results are given in Table 3.2.3 and the fit to the second volatility function is presented in Figure 3.2.4.

Parameter	Value
b	0.7511
α	12.734
c_1	0.6254
c_2	-0.1354

Table 3.2.3: MRVG-2 Model Historically Estimated Parameters

The parameters have been estimated by fitting the second volatility function of the MRVG-2 model, given by $c_1 \exp(-\alpha\tau) + c_2$, to the historical function values derived through spectral decomposition of the forward curve returns' covariance matrix.

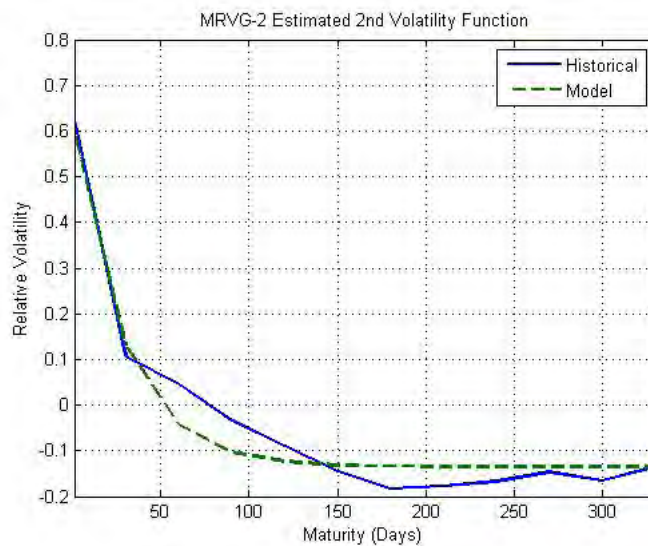


Figure 3.2.4: MRVG-2 Calibrated 2nd Volatility Function

The fit between the model and historically estimated second volatility function for the MRVG-2 model.

The results for the market calibration of ν and σ are given in Table 3.2.4 below. The fit to the implied moments, measured by the sum of squared percentage errors, was 0.293021×10^{-3} , which is considerably higher than the fit for the other models. This is reflective of the fact that we are imposing certain priors on the forward curve distribution dynamics which are not reflected in the terminal distribution. The sum of squared percentage errors of the model to the market option prices is 35.82% which is unreasonably high and so we can discard this model on that basis.

Parameter	Value
σ	0.9676
ν	0.6976
SSE(implied moments)	0.293021×10^{-3}
SSE(option prices)	35.82%

Table 3.2.4: MRVG-2 Calibrated Parameters

The model parameters for the MRVG-2 model estimated using the implied moments technique. Here, 'SSE(implied moments)' refers to the sum of squared percentage error between the model and market moments, which was the metric used when fitting the model. 'SSE(option prices)' refers to the sum of squared percentage errors between the model and market option prices.

For the MRVG-3x model, we will again estimate α, σ and ν by calibrating to the market implied moments. This leaves $\{b, \varepsilon, c_1\}$ to be estimated from the shape of the second volatility function, in this case approximated by $c_1 \exp(-\varepsilon\tau)$. The fitting strategy here is to mimic the shape of the volatility function, however, rather than decaying to a negative value we assume it tends to zero. The results are given in Table 3.2.5 and the fit to the second volatility function is presented in Figure 3.2.5.

Parameter	Value
b	0.7511
ε	12.734
c_1	0.6254

Table 3.2.5: MRVG-3x Historically Estimated Parameters

The parameters have been estimated by fitting the second volatility function of the MRVG-3x model, given by $c_1 \exp(-\varepsilon\tau)$, to the historical function values derived through spectral decomposition of the forward curve returns' covariance matrix.

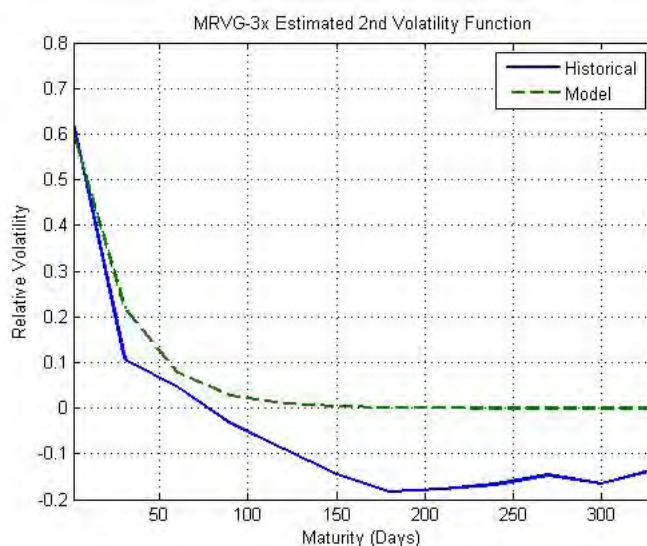


Figure 3.2.5: MRVG-3x Calibrated 2nd Volatility Function

The fit between the model and historically estimated second volatility function for the MRVG-3x model.

The results for α , σ and ν are given in Table 3.2.6 below, the fit to the market moments was similar to that of the MRVG model at 1.6608×10^{-8} . The fit to option prices is better than both the MRVG and MRVG-2x models with a SSE value of 0.74%.

Parameter	Value
α	0.1418
σ	0.2518
ν	0.1675
SSE(implied moments)	1.6608×10^{-8}
SSE(option prices)	0.74%

Table 3.2.6: MRVG-3x Calibrated Parameters.

The model parameters for the MRVG-3x model estimated using the implied moments technique. Here, 'SSE(implied moments)' refers to the sum of squared percentage error between the model and market moments, which was the metric used when fitting the model. 'SSE(option prices)' refers to the sum of squared percentage errors between the model and market option prices.

On the basis of this analysis we will utilize the MRVG-3x and MRVG-2x models in pricing an example storage deal. Using these calibrated parameters we can visually analyze the instantaneous log-forward price moments versus time to maturity (Figures 3.2.6 and 3.2.7). As one would expect from the fitted parameters, the instantaneous volatility of the MRVG-3x model is greater than that of the MRVG-2x. The

model kurtosis, given by the 4th moment divided by the square of the 2nd, is much higher in the MRVG-2x model for short maturities and decays quite rapidly in both models as time to maturity increases. This is a useful model feature which reflects the greater spike risk at the front of the curve.

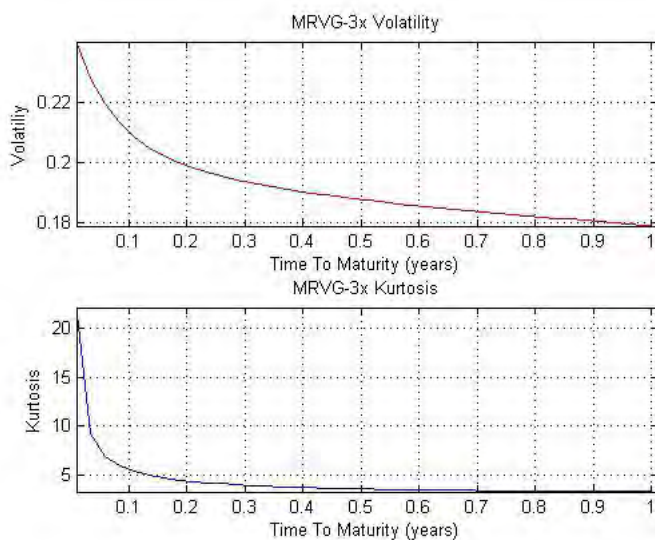


Figure 3.2.6: MRVG-3x Distribution Moments.

The annualized volatility and Kurtosis of the MRVG-3x model using the fitted model parameters.




Figure 3.2.7: MRVG-2x Distribution Moments.

The annualized volatility and Kurtosis of the MRVG-2x model using the fitted model parameters.

3.3 Multidimensional Valuation Algorithm

We will now outline another key contribution of this chapter, the generalization of the single factor FFT based valuation algorithm to an arbitrary number of dimensions. Although the emphasis here is on storage valuation, the algorithm could easily be altered to price other path dependent options such as take-or-pay contracts. Further, valuations which depend on multiple underlying assets could also be easily incorporated within this framework. One relevant example of interest being storage contracts with location flexibility e.g. a German storage with the option to inject/withdraw to both NCG and Gaspool.

As already outlined in the single factor case, the storage value is derived through backward stochastic dynamic programming where the value function is given by⁶

$$V(x_{t_{j-1}}(\vec{y}_{t_{j-1}}), I) = \sup_{I^*} \theta(x_{t_{j-1}}(\vec{y}_{t_{j-1}}), I^*; I) + E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}] \quad (3.3.1)$$

where $\vec{y}_t \in \mathbb{R}^K$. We therefore simply need to solve the expectation in Equation 3.3.1 at each time-step and for each combination of (x_{t_j}, I^*) ,

$$\begin{aligned} E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}] &= \int_{\mathbb{R}^n} V\left(f(0, t_j) + \sum_{k=1}^K y_{t_j}^{(k)}, I^*\right) f(\vec{y}_{t_j} | \vec{y}_{t_{j-1}}) d\vec{y}_{t_j} \\ &= \int_{\mathbb{R}} \dots \int_{\mathbb{R}} V\left(f(0, t_j) + \sum_{k=1}^K y_{t_j}^{(k)}, I^*\right) \times \\ &\quad f\left(y_{t_j}^{(1)}, y_{t_j}^{(2)}, \dots, y_{t_j}^{(K)} | y_{t_{j-1}}^{(1)}, y_{t_{j-1}}^{(2)}, \dots, y_{t_{j-1}}^{(K)}\right) dy_{t_j}^{(1)} dy_{t_j}^{(2)} \dots dy_{t_j}^{(K)} \end{aligned}$$

where the application of Fubini's Theorem requires the mild technical restriction that

$$E[|V(x_{t_j}, I^*)| | \vec{y}_{t_{j-1}}] < \infty$$

Applying Parseval's Theorem we have,

$$E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}] = \left(\frac{1}{2\pi}\right)^K \int_{\mathbb{C}} \dots \int_{\mathbb{C}} \tilde{v}(\vec{z}; I^*) \Phi_{\vec{y}_{t_j}}(\vec{z}; \vec{y}_{t_{j-1}}, t_{j-1}, t_j) dz^{(1)} dz^{(2)} \dots dz^{(K)} \quad (3.3.2)$$

where

$$\tilde{v}(\vec{z}; I^*) = \int_{\mathbb{R}^n} \exp(-i\vec{z}^T \cdot \vec{y}_{t_j}) V\left(f(0, t_j) + \sum_{i=1}^K y_{t_j}^{(i)}, I^*\right) d\vec{y}_{t_j}$$

⁶For brevity, we denote the function $x_{t_j}(\vec{y}_{t_j})$ by the shorthand x_{t_j} throughout the remainder of this section.

with $\vec{z} \in \mathbb{C}^n : \vec{u} + i\hat{w}$. The vector \hat{w} is defined such that

$$\int_{\mathbb{R}^n} \left| \exp(\vec{w}^\top \cdot \vec{y}_{t_j}) V \left(f(0, t_j) + \sum_{k=1}^K y_{t_j}^{(k)}, I^* \right) \right| d\vec{y}_{t_j} < \infty$$

Now, noting from Equation 3.1.10 and Equation 3.1.13 that the conditional characteristic function is of the form

$$\Phi_{\vec{y}_{t_j}} = \exp(i\vec{a}\vec{z}^\top \cdot \vec{y}_{t_{j-1}}) \exp(\Psi(\vec{z}, t_{j-1}, t_j))$$

where $\vec{a}\vec{z}$ is the element wise multiplication of the vector \vec{z} with

$$\vec{a} = \begin{bmatrix} a^{(1)}(t_{j-1}, t_j) \\ a^{(2)}(t_{j-1}, t_j) \\ \cdot \\ \cdot \\ a^{(K)}(t_{j-1}, t_j) \end{bmatrix}$$

Therefore, we can write Equation 3.3.2 as

$$E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}] = \left(\frac{1}{2\pi} \right)^K \int_{\mathbb{C}} \dots \int_{\mathbb{C}} \tilde{v}(\vec{z}; I^*) \exp(i\vec{a}\vec{z}^\top \cdot \vec{y}_{t_{j-1}}) \exp(\Psi(\vec{z}, t_{j-1}, t_j)) dz^{(1)} dz^{(2)} \dots dz^{(K)}$$

Applying the substitution $\vec{z}' = \vec{a}\vec{z}$ gives us

$$\begin{aligned} E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}] &= \left(\prod_{k=1}^K \frac{1}{a^{(k)}(t_{j-1}, t_j)} \right) \left(\frac{1}{2\pi} \right)^K \int_{\mathbb{C}} \dots \int_{\mathbb{C}} \tilde{v}\left(\frac{\vec{z}'}{a}; I^*\right) \\ &\quad \times \exp(i\vec{z}'^\top \cdot \vec{y}_{t_{j-1}}) \exp\left(\Psi\left(\frac{\vec{z}'}{a}, t_{j-1}, t_j\right)\right) dz'^{(1)} dz'^{(2)} \dots dz'^{(K)} \end{aligned} \quad (3.3.3)$$

Now as in the single dimension case, we can apply the scaling property of the Fourier Transform to derive

$$\tilde{v}\left(\frac{\vec{z}'}{a}; I^*\right) = \left(\prod_{k=1}^K a^{(k)}(t_{j-1}, t_j) \right) \mathcal{F} \left[V \left(f(0, t_j) + \sum_{k=1}^K a^{(k)}(t_{j-1}, t_j) y_{t_j}^{(k)}, I^* \right) \right] \quad (3.3.4)$$

We will now discuss the solution of the valuation problem using the multidimensional Fast Fourier Transform (FFT) algorithm. Similar to the discussion in Section 2.3, we will focus on discretizing the expectation given in Equation 3.3.3 and recasting it in a form suitable for evaluation using the FFT.

Discretization & Solution

To solve for Equation 3.3.3 we first truncate and discretize the domains of \vec{y} and \vec{z}' such that

$$y^{(k)} \in [y_0^{(k)}, \dots, y_{N^{(k)}-1}^{(k)}] \quad \forall k$$

where $y_n^{(k)} = y_0^{(k)} + n^{(k)} \Delta y^{(k)}$ for $n^{(k)} = 1, \dots, N^{(k)}$, and $N^{(k)}$ is the number of grid points for the k^{th} factor. Similarly, for each $z'^{(k)} = u^{(k)} + iw^{(k)}$ we have

$$u^{(k)} \in [u_0^{(k)}, \dots, u_{M^{(k)}-1}^{(k)}] \quad \forall k$$

where $u_m^{(k)} = u_0^{(k)} + m^{(k)} \Delta u^{(k)}$ and $M^{(k)}$ is the number of grid points. Applying a composite trapezoidal product rule Equation 3.3.4 then becomes

$$\begin{aligned} \tilde{v} \left(\frac{\vec{z}'}{a}; I^* \right) &= \left(\prod_{k=1}^K a^{(k)}(t_{j-1}, t_j) \right) \sum_{n^{(1)}=0}^{N^{(1)}-1} l_{n^{(1)}}^{(1)} \dots \sum_{n^{(K)}=0}^{N^{(K)}-1} l_{n^{(K)}}^{(K)} \exp \left(-iz'^{(1)} y_{n^{(1)}}^{(1)} \dots - iz'^{(K)} y_{n^{(K)}}^{(K)} \right) \times \\ &V \left(f(0, t_j) + \sum_{k=1}^K a^{(k)}(t_{j-1}, t_j) y_{n^{(k)}}^{(k)}, I^* \right) \Delta y^{(K)} \dots \Delta y^{(1)} \end{aligned}$$

where $l_{n^{(k)}}^{(k)} = \frac{1}{2}, n^{(k)} = 0, N^{(k)} - 1$ and $l_{n^{(k)}}^{(k)} = 1$ elsewhere.

Now, each term $\exp \left(-iz'^{(k)} y_{n^{(k)}}^{(k)} \right)$ can be expanded such that

$$\begin{aligned} \exp \left(-iz'^{(k)} y_{n^{(k)}}^{(k)} \right) &= \exp \left(-iu_0^{(k)} n^{(k)} \Delta y^{(k)} - im^{(k)} \Delta u^{(k)} n^{(k)} \Delta y^{(k)} - iw_0^{(k)} y_0^{(k)} \right. \\ &\quad \left. - im^{(k)} \Delta u^{(k)} y_0^{(k)} + w^{(k)} n^{(k)} \Delta y^{(k)} + w^{(k)} y_0^{(k)} \right) \end{aligned}$$

Terms which are independent of $n^{(k)}$ can be factored out to give

$$\begin{aligned} \tilde{v} \left(\frac{\vec{z}'}{a}; I^* \right) &= \exp \left(\sum_{k=1}^K -iu_0^{(k)} y_0^{(k)} + w^{(k)} y_0^{(k)} \right) \exp \left(\sum_{k=1}^K -im^{(k)} \Delta u^{(k)} y_0^{(k)} \right) \left(\prod_{k=1}^K a^{(k)}(t_{j-1}, t_j) \right) \times \\ &\sum_{n^{(1)}=0}^{N^{(1)}-1} l_{n^{(1)}}^{(1)} \dots \sum_{n^{(K)}=0}^{N^{(K)}-1} l_{n^{(K)}}^{(K)} \exp \left(\sum_{k=1}^K -im^{(k)} \Delta u^{(k)} n^{(k)} \Delta y^{(k)} \right) \times \\ &\exp \left(\sum_{k=1}^K -iu_0^{(k)} n^{(k)} \Delta y^{(k)} + w^{(k)} n^{(k)} \Delta y^{(k)} \right) \times \\ &V \left(f(0, t_j) + \sum_{k=1}^K a^{(k)}(t_{j-1}, t_j) y_{n^{(k)}}^{(k)}, I^* \right) \Delta y^{(K)} \dots \Delta y^{(1)} \end{aligned}$$

Similarly, we discretize Equation 3.3.3 to give

$$\begin{aligned}
E [V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}] &= \left(\frac{1}{2\pi}\right)^K \left(\prod_{k=1}^K \frac{1}{a^{(k)}(t_{j-1}, t_j)}\right) \sum_{m^{(1)}=0}^{M^{(1)}-1} l_{m^{(1)}}^{(1)} \dots \sum_{m^{(K)}=0}^{M^{(K)}-1} l_{m^{(K)}}^{(K)} \tilde{v}\left(\frac{\vec{z}'}{a}; I^*\right) \times \\
&\quad \exp\left(iz'_{m^{(1)}y_{n^{(1)}}}^{(1)} + \dots + iz'_{m^{(K)}y_{n^{(K)}}}^{(K)}\right) \Psi\left(\frac{\vec{z}'}{a}, t_j, t_{j-1}\right) \Delta u^{(K)} \dots \Delta u^{(1)} \\
&= \left(\frac{1}{2\pi}\right)^K \left(\prod_{k=1}^K \frac{1}{a^{(k)}(t_{j-1}, t_j)}\right) \sum_{m^{(1)}=0}^{M^{(1)}-1} l_{m^{(1)}}^{(1)} \dots \sum_{m^{(K)}=0}^{M^{(K)}-1} l_{m^{(K)}}^{(K)} \times \\
&\quad \exp\left(\sum_{k=1}^K iu_0^{(k)} n^{(k)} \Delta y^{(k)} + iu_0^{(k)} y_0^{(k)}\right) \exp\left(\sum_{k=1}^K -w^{(k)} n^{(k)} \Delta y^{(k)} - w^{(k)} y_0^{(k)}\right) \times \\
&\quad \exp\left(\sum_{k=1}^K im^{(k)} \Delta u^{(k)} n^{(k)} \Delta y^{(k)}\right) \exp\left(\sum_{k=1}^K im^{(k)} \Delta u^{(k)} y_0^{(k)}\right) \times \\
&\quad \Psi\left(\frac{\vec{z}'}{a}, t_{j-1}, t_j\right) \tilde{v}\left(\frac{\vec{z}'}{a}; I^*\right) \Delta u^{(K)} \dots \Delta u^{(1)} \\
&= \left(\frac{1}{2\pi}\right)^K \left(\prod_{k=1}^K \frac{1}{a^{(k)}(t_{j-1}, t_j)}\right) \exp\left(\sum_{k=1}^K iu_0^{(k)} n^{(k)} \Delta y^{(k)} + iu_0^{(k)} y_0^{(k)}\right) \times \\
&\quad \exp\left(\sum_{k=1}^K -w^{(k)} n^{(k)} \Delta y^{(k)} - w^{(k)} y_0^{(k)}\right) \times \\
&\quad \sum_{m^{(1)}=0}^{M^{(1)}-1} l_{m^{(1)}}^{(1)} \dots \sum_{m^{(K)}=0}^{M^{(K)}-1} l_{m^{(K)}}^{(K)} \exp\left(\sum_{k=1}^K im^{(k)} \Delta u^{(k)} n^{(k)} \Delta y^{(k)}\right) \Psi\left(\frac{\vec{z}'}{a}, t_{j-1}, t_j\right) \times \\
&\quad \exp\left(\sum_{k=1}^K im^{(k)} \Delta u^{(k)} y_0^{(k)}\right) \tilde{v}\left(\frac{\vec{z}'}{a}; I^*\right) \Delta u^{(K)} \dots \Delta u^{(1)}
\end{aligned}$$

Expanding $\tilde{v}\left(\frac{\vec{z}'}{a}; I^*\right)$ gives us

$$\begin{aligned}
E [V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}] &= \left(\prod_{k=1}^K \frac{\Delta u^{(k)} \Delta y^{(k)}}{2\pi}\right) \exp\left(\sum_{k=1}^K iu_0^{(k)} n^{(k)} \Delta y^{(k)} - w^{(k)} n^{(k)} \Delta y^{(k)}\right) \times \\
&\quad \sum_{m^{(1)}=0}^{M^{(1)}-1} l_{m^{(1)}}^{(1)} \dots \sum_{m^{(K)}=0}^{M^{(K)}-1} l_{m^{(K)}}^{(K)} \exp\left(\sum_{k=1}^K im^{(k)} \Delta u^{(k)} n^{(k)} \Delta y^{(k)}\right) \Psi\left(\frac{\vec{z}'}{a}, t_{j-1}, t_j\right) \times \\
&\quad \sum_{n^{(1)}=0}^{N^{(1)}-1} l_{n^{(1)}}^{(1)} \dots \sum_{n^{(K)}=0}^{N^{(K)}-1} l_{n^{(K)}}^{(K)} \exp\left(\sum_{k=1}^K -im^{(k)} \Delta u^{(k)} n^{(k)} \Delta y^{(k)}\right) \times \\
&\quad \exp\left(\sum_{k=1}^K -iu_0^{(k)} n^{(k)} \Delta y^{(k)} + w^{(k)} n^{(k)} \Delta y^{(k)}\right) V\left(f(0, t_j) + \sum_{k=1}^K a^{(k)}(t_{j-1}, t_j) y_{n^{(k)}}^{(k)}, I^*\right)
\end{aligned}$$

As in the single factor case, in order to utilize the multidimensional Fast Fourier Transform to solve for

the above we need to enforce the restriction that $M^{(k)} = N^{(k)}$ for $k = 1, \dots, K$ and also

$$\Delta u^{(k)} \Delta y^{(k)} = \frac{2\pi}{N^{(k)}} \quad \forall k$$

3.4 Empirical Analysis & Results

We now proceed to valuing an example storage deal under both the calibrated MRVG-3x and MRVG-2x models. We will begin by presenting the valuation results and then analyze the value difference attributable to the differences in model specification. Value convergence results are presented in Section B.7 of the Appendix . To allow comparison with the single factor models, we will value the storage deal specified in Chapter 2. To recap, the deal parameters are

- Start Date: 19th December 2012
- End Date: 18th December 2013
- Initial Inventory Level: 0
- Final Inventory Level: 0
- Capacity: 29.3 GWh
- Max Injection/Withdrawal: 1.465 GWh per day.
- Underlying Gas Price: NBP (National Balancing Point) pence/therm

The valuation date is the 19th December 2012 and the forward curve for NBP Gas on this date is given in Figure 3.4.1 . This curve yields an intrinsic value of 10.983 pence/therm, where we have for simplicity ignored the effect of discounting.




Figure 3.4.1: NBP Forward Curve

Futures closing prices (blue line) in pence per therm as of the 19th December 2012, source Bloomberg. The green dashed line displays the storage inventory levels (in million Kilowatt Hours) given by the intrinsic hedging strategy associated with this forward curve. The intrinsic value associated with this strategy is 10.983 pence/therm.

Valuation Analysis

Initial Gas Price: £0.667	Intrinsic Value	Full Value	Extrinsic Value	% Δ Extrinsic Value
MRVG	10.983	11.2105	0.2275	-
MRVG-2x	10.983	11.9642	0.9812	431.3%
MRVG-3x	10.983	15.1477	4.1647	1830.64%

All values are expressed in pence/therm

Table 3.4.1: Two Factor Valuation Results

The valuation results for both two factor models along with the results of the single factor MRVG model.

Table 3.4.1 presents the results for the two multifactor models along with the results for the single factor MRVG model. The valuations for the MRVG-2x and MRVG-3x models were carried out using the FFT-m approach outlined in Chapter 2, with grid sizes of 512×512 and 256×2048 respectively. For these valuations, the factor domains (in log-space) were approximated as $[-3.4539, 3.4471]$. The MRVG-2x value is 11.9642 pence/therm with 0.9812 pence/therm extrinsic value, this represents a 431% increase in extrinsic value over the single factor MRVG valuation. The MRVG-3x model value is much higher at 15.1477 pence/therm with 4.1677 pence/therm extrinsic value, a 1831% increase over the MRVG model. This increase highlights the sensitivity of the storage asset to the model implied seasonal decorrelation, which is highest in the case of the MRVG-3x model.

Figure 3.4.2 displays the initial value for the MRVG-2x model for different levels of the individual factors. These graphs are helpful in understanding the differences between the models and identifying the drivers of the extrinsic value in each case. We can see that for the MRVG-2x model, the value reaches a maximum when both factors are at their positive extremes, which is also where the initial gas price is at a maximum for this sample of the value grid. Figure 3.4.3 displays the implied forward curve given by the factor values at these extreme points. At $y_1, y_2 = [0.4, 0.4]$ the whole curve and therefore the initial intrinsic value would increase due to the proportional shifts to the forward curve caused by the second factor and to a lesser extent the first factor. At $y_1, y_2 = [-0.4, 0.4]$, the curve is anchored at the short end by the decrease in the first factor which causes an increase in the summer/winter spread and therefore an increase in intrinsic value.




Figure 3.4.2: MRVG-2x Initial Value Grid

The two dimensional storage value grid for the MRVG-2x model. The value is expressed in pence/therm and displayed as function of the underlying factor levels.

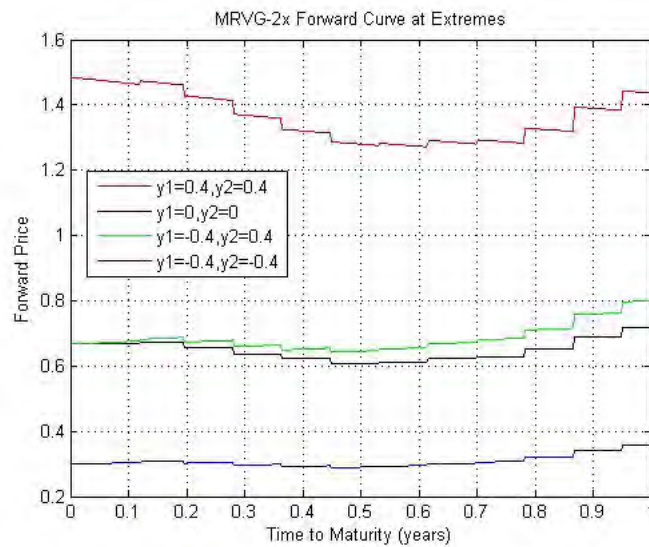


Figure 3.4.3: MRVG-2x Implied Forward Curves

The factor implied forward curves for the MRVG-2x model at each of the extreme points of the storage value grid presented in Figure 3.4.2.

For the MRGV-3x model both factors contribute significantly to the value of the storage asset through their effect on the forward curve as the individual factor levels change. The value is maximized in Figure

3.4.4 when the first factor is at its maximum and the second factor is at its minimum at $y_1, y_2 = [0.4, -0.4]$. The model implied forward curve at this level is displayed in Figure 3.4.5. We see that the changes in both factors cancel each other out initially, so that the initial gas price remains unchanged. The positive change in the first factor would persist along the forward curve whereas the negative change in the second factor would decay quite quickly given their respective mean reversion rates. The net effect is a large increase in the forward curve relative to the prompt gas price and therefore a large change in the intrinsic value. A similar effect is seen when both factors are at their minimum values at $y_1, y_2 = [-0.4, -0.4]$. Here the large mean reversion rate of the second factor leads to a much larger decrease in the forward curve at the front end, leading to more profitable injection opportunities.

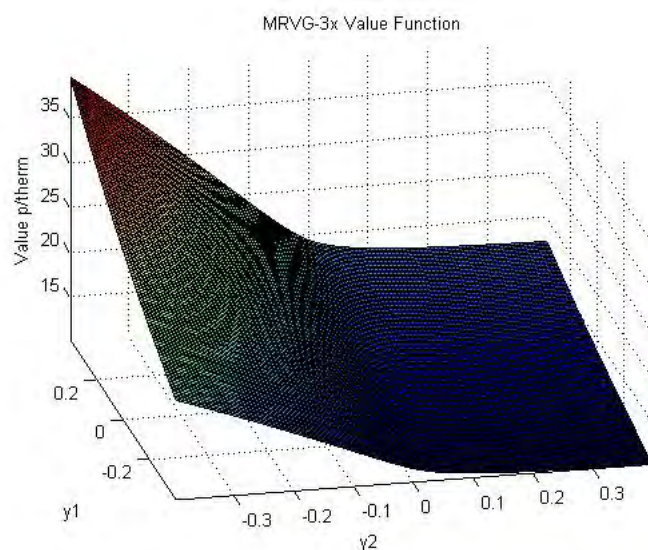


Figure 3.4.4: MRVG-3x Initial Value Grid

The two dimensional storage value grid for the MRVG-3x model. The value is expressed in pence/therm and displayed as function of the underlying factor levels.

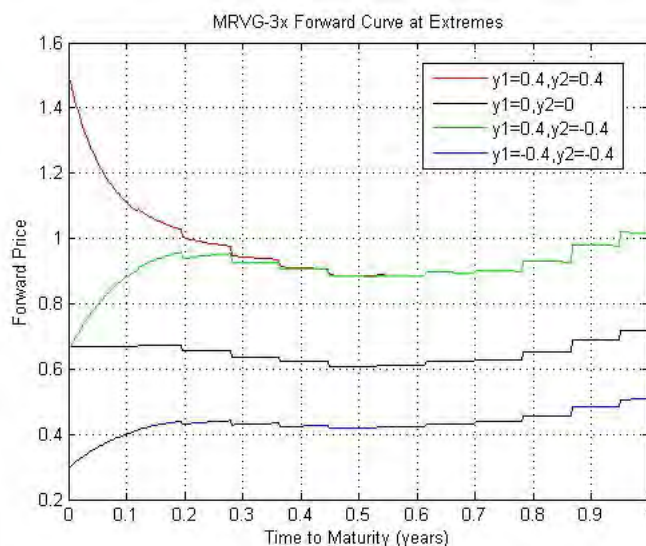


Figure 3.4.5: MRVG-3x Implied Forward Curves

The factor implied forward curves for the MRVG-3x model at each of the extreme points of the storage value grid presented in Figure 3.4.4.

It is clear from the above results that incorporating an accurate representation of the forward curve dynamics in conjunction with a market based calibration of the general level of curve variability leads to substantially greater storage values than market calibration alone. This is of course due entirely to market incompleteness with respect to time-spread optionality. As such, a model which is capable of representing the dynamics observable from historical returns gives traders at least some comfort in the level of extrinsic being bid/offered. Not only does the MRVG-3x model meet this requirement, the parsimony of the model allows one to adjust ones price levels significantly by adjusting a single parameter, ϵ , without significantly impacting the models ability to calibrate to market. Given the growing liquidity observable in the over-the-counter storage and 100% take-or-pay markets however, this requirement of utilizing historical information may not be necessary for much longer. In such a situation, a model which can easily and quickly be calibrated to market prices for storage and take-or-pay contracts would allow a trader to infer the price of time-spread optionality directly from these products. Again, the MRVG-3x, where the entire curve correlation structure is controlled by two parameters, would be considerably more suited to this purpose than a market model of the forward curve with a full correlation matrix.

In conclusion, in this chapter we aimed to present an unique class of multifactor forward curve models that were reflective of the statistical dynamics of the natural gas forward curve and allowed for accurate calibration to the options market. To that end, we first conducted a Principal Component Analysis of the historical relative maturity NBP gas forward curve returns in order to inform our model specification. We presented a three factor model with a volatility specification given by the shape of the eigenvector values for the first two principal components as a function of relative maturity. Given concerns over the

potential computational burden of using three stochastic factors, we went on to introduce a range of two factor models which approximated the volatility specification of the three factor model. The conditional characteristic functions for this family of models were derived in Sections B.2 and B.3 of the Appendix, along with a range of utility methods which allowed us to further analyze the dynamics of each model. We introduced an implied moment based calibration technique in order to calibrate these models to the options market and went on to present results of the joint calibration and statistical estimation of the model parameters. In order to value storage contracts we extended the valuation algorithm presented in Chapter 2 to a multidimensional setting and finished with the presentation of valuation results and analysis using two of the proposed multifactor models. In summary, the main contributions to the literature have been

- The extension of the MRVG model to a multifactor setting under the Cheyette model framework.
- The derivation of the conditional characteristic functions for a number of multifactor model specifications.
- The presentation of a detailed derivation of the implied spot price drift for each model specification and the conditional characteristic function of the resulting log spot price process.
- The presentation of an innovative implied moments calibration technique applicable to any finite dimensional instantaneous forward curve model. This was an extension of the work of Guillaume and Schoutens (2013) specifically tailored to the calibration of energy forward curve models but equally applicable to interest rate models.
- The derivation of a number of utility methods for each model specification including the forward curve covariance function and the moments of each process, which we used to analyze the model implied factor dynamics.
- The derivation of the spot factor implied forward curve for each model specification and using these results to derive a general Fourier based swaption pricing algorithm.
- The extension of the Fourier based valuation algorithm for processes with state dependent increments to an arbitrary number of dimensions.

Given the innovative model development work carried out in Chapters 2 and 3, we feel that a detailed analysis surrounding the model risk inherent in our proposed modelling framework is important at this juncture, particularly for the models to take traction in both industry and academia. We address this in the next chapter, where a selection of the models set out thus far are analyzed under an innovative calibration and parameter estimation risk framework, with a particular focus on the impact of model risk on the storage valuation results presented in Chapters 2 and 3.

Chapter 4

Storage Model Risk

4.1 Model Uncertainty and Parameter Risk

In deriving appropriate risk measures for quantifying parameter risk the literature to date has drawn heavily upon the theory of coherent and convex risk measures. In their seminal work Artzner et al. (1999) introduce the theory of coherent risk metrics as a response to the perceived shortcomings of traditional risk metrics used in practice in the financial industry. Percentile based risk metrics, such as VaR, can seriously misrepresent risk in the case of discontinuous payoffs. Further, VaR may fail to reflect the benefits of diversification for non-Gaussian underlyings. The authors propose an axiomatisation of risk measurement so as to ensure that “rules of thumb” such as this and certain other desirable properties are met. Föllmer and Schied (2002) extend this approach by introducing the notion of convex risk measures defined by weakening the diversification property of coherent risk measures.

This general framework was quickly identified by Cont (2006) as a means of overcoming one of the primary goals of quantitative finance, the valuation and hedging of derivatives in incomplete markets. Earlier work by Figlewski (1998) and Gibson et al. (1998), amongst others, highlighted the dependency of exotic option prices on the choice of stochastic model in the absence of a static replicating portfolio and also the ambiguity surrounding the “correct” model selection. This uncertainty and the resulting impact on asset value came to be identified under the term “model risk”. Kerkhof et al. (2002) sub-categorize this risk into two distinct areas, the first owing to the choice of model specification from an array of acceptable models, and generally referred to as model uncertainty. The second arising from either the calibration of a particular model to market prices or the estimation of model parameters from historical data, and referred to generally as parameter risk. In order to make clear the distinction between these two elements of parameter risk, we will refer to the risk owing to model calibration to current market prices as calibration risk and the parameter risk associated with estimation from historical data as parameter estimation risk.

In their practitioner orientated work Gupta et al. (2010) provide a detailed introduction to the topic of calibration risk, including the key risk drivers and potential remedies. For example, the authors demonstrate how Bayesian averaging over potential parameter values can be used to both infer the level of calibration risk and also “smooth” the inverse problem associated with market calibrations. This approach improves upon the “point” estimates for calibration risk given by Cont (2006) by associating a probability, albeit subjective, with a given parameter choice. Deryabin (2012) further extends this distributional view of calibration risk by introducing upper and lower bounds on the risk which rely only upon the specification of the underlying model and its calibration to market data. This work was further extended by Bannör and Scherer (2013b) who showed how meaningful probabilistic metrics of the calibration and parameter estimation risk inherent in a given model can be constructed using empirical densities derived from the calibration errors and sample parameter variance respectively. This approach marked a step-change in the quantification of calibration risk by providing a complete probabilistic framework for its assessment.

The main contribution of this chapter is the unification of their distinct methods of quantifying calibration and parameter estimation risk within a single risk metric. The benefit of this approach will be demonstrated with reference to storage asset valuations using the models developed in the earlier chapters. The same authors addressed parameter estimation risk with respect to real assets in an earlier paper, Bannor et al. (2013a), where they focus on power plant valuations. Henaff et al. (2013) present a similar approach in estimating the parameter estimation risk on a storage asset. Our work is thus the first to analyze the calibration risk associated with real asset valuations and further to incorporate both calibration and parameter estimation risk in a consistent manner.

We will begin with a brief review of the theory of coherent and convex risk measures and introduce their application in the field of parameter risk analysis. This will serve the dual purpose of providing the necessary context to set up the more specific discussion to follow in the latter half of this chapter, as well as demonstrating the strong theoretical footing upon which the notions of parameter estimation and calibration risk are based.

4.1.1 Coherent and Convex Risk Measures

Artzner et al. (1999) define a coherent risk measure as a map $\rho : \mathcal{X} \rightarrow \mathbb{R}$, where \mathcal{X} is a function space existing on the sample space Ω , possessing the following properties

- Subadditivity: $\rho(X + Y) \leq \rho(X) + \rho(Y)$ for $X, Y \in \mathcal{X}$
- Positive Homogeneity: $\rho(\lambda X) = \lambda \rho(X)$ for a given scalar $\lambda \geq 0$
- Monotonicity: If $X \leq Y$, then $\rho(X) \geq \rho(Y)$

- Translation Invariance: $\rho(X + m) = \rho(X) - m$ for $m \in \mathbb{R}$

The monotonicity property states the rather obvious fact that if the future value of an asset is dominated by another then it is relatively more risky. The translation invariance property is designed to capture the existence of a risk free asset. The subadditivity property ensures that diversification does not add risk, and the positive homogeneity property states that risk must increase linearly with position size. The latter two properties were relaxed by Föllmer and Schied (2002) to allow for non-linear changes in risk when the size of a given position increases and replaced with the convexity property

$$\rho(\lambda X + (1 - \lambda)Y) \leq \lambda\rho(X) + (1 - \lambda)\rho(Y) \quad \lambda \in [0, 1]$$

This merely restates the diversification principle as the risk on a weighted sum of two assets being less than or equal to the weighted sum of the individual asset risks. The authors then formally define the notion of a convex risk measure as a map $\rho : \mathcal{X} \rightarrow \mathbb{R}$ which satisfy the convexity, monotonicity, and translation invariance properties. Assuming that the measure is normalized, that is $\rho(0) = 0$, and using the translation invariance property, the risk of a position X can be covered by adding $\rho(X)$ units of the risk free asset (i.e. cash) to the position. The authors use this property to define a set of acceptable positions which require no additional capital,

$$A := \{X \in \mathcal{X} | \rho(X) \leq 0\}$$

which yields an alternative definition of the convex risk measure as the current cost of the remedial asset allocation, or hedging instrument, needed to make an unacceptable risk exposure acceptable

$$\rho(X) := \inf \{m \in \mathbb{R} | m + X \in A\}$$

The acceptance set A associated with a convex risk measure possesses the following properties

- A is convex and non-empty
- If $X \in A$ and $Y \geq X$ then $Y \in A$
- If $X \in A$ then $\{\lambda \in [0, 1] | \lambda X + (1 - \lambda)Y \in A\}$ is closed in $[0, 1]$

Further, a convex risk measure has the representation

$$\rho(X) = \sup_{\mathbb{Q} \in \mathcal{Q}} (E_{\mathbb{Q}}[-X] - \alpha(\mathbb{Q}))$$

where \mathcal{Q} is a set of probability measures on Ω and $E^{\mathbb{Q}}$ denotes expectation on the probability space $(\Omega, \mathcal{F}, \mathbb{Q})$. The function $\alpha : \mathcal{Q} \rightarrow \{\mathbb{R} \cup \infty\}$ penalizes potential probability measures based on certain criteria inherent to the metric ρ , an obvious example which we will focus on in more detail later is the

mispricing of benchmark instruments.

4.1.2 Model Uncertainty

We will now give a brief overview of the methods for incorporating model uncertainty, which encompasses parameter risk, into asset valuations. The exact approach taken will of course differ depending upon the information that is used to estimate the parameters of the pricing model, however, the general methodology is similar for both market calibrated and historical parameter estimates. In his seminal paper Cont (2006) extends the theory of convex risk measures in order to derive derivative valuations adjusted for model uncertainty. As argued by the author¹, the distinction between parameter uncertainty and model uncertainty is largely irrelevant. Model uncertainty is determined by the uncertainty surrounding the choice of probability measure \mathbb{Q} from a family of potential measures \mathcal{Q} . Since every parameter combination for a given model specification yields a distinct probability measure, it is clear that parameter uncertainty is in fact embedded in the definition of model uncertainty. In quantifying parameter uncertainty/risk, Cont (2006) focuses on the model's ability to replicate the prices of a set of benchmark derivative values. The author denotes the payoffs on these benchmark derivatives as $(H_i)_{i \in I}$ and their associated prices by $C_i^* \in [C_i^{bid}, C_i^{ask}]$, where the values C_i^{bid}, C_i^{ask} denote the market quoted prices. The existence of a set of arbitrage free pricing models \mathcal{Q} is assumed, such that,

$$\forall \mathbb{Q} \in \mathcal{Q}, i \in I \ E_{\mathbb{Q}}[|H_i|] < \infty, \ E_{\mathbb{Q}}[H_i] \in [C_i^{bid}, C_i^{ask}]$$

Then given a set of payoffs

$$\mathcal{C} = \left\{ H, \sup_{\mathbb{Q} \in \mathcal{Q}} E_{\mathbb{Q}}[|H|] < \infty \right\}$$

the author defines a model uncertainty measure as a mapping $\mu : \mathcal{C} \rightarrow [0, \infty[$ which possesses the following properties:

For benchmark instruments, model uncertainty reduces to uncertainty on market value:

- $\forall i \in I$ we have $\mu(H_i) \leq |C_i^{ask} - C_i^{bid}|$.

Model uncertainty cannot be reduced through hedging with the underlying:

- $\forall \phi \in \mathcal{S}$, $\mu\left(X + \int_0^T \phi_t ds_t\right) = \mu(X)$ where \mathcal{S} is the set of admissible trading strategies in the underlying asset s_t

However, static hedging with tradeable options may lead to reductions in model risk:

- $\forall u \in \mathbb{R}^k$, $\mu\left(X + \sum_{i=1}^k u_i H_i\right) \leq \mu(X) + \sum_{i=1}^k |u_i| (C_i^{ask} - C_i^{bid})$

¹See Remark 4.1 Cont (2006).

and finally, the diversification property:

- $\mu(\lambda X + (1 - \lambda)Y) \leq \lambda\mu(X) + (1 - \lambda)\mu(Y)$

Cont (2006) gives the following example of a coherent model risk measure by first, defining

$$\begin{aligned}\bar{\pi}(X) &= \sup_{\mathbb{Q} \in \mathcal{Q}} E_{\mathbb{Q}}[X] \\ \underline{\pi}(X) &= \inf_{\mathbb{Q} \in \mathcal{Q}} E_{\mathbb{Q}}[X] \\ &= -\bar{\pi}(-X)\end{aligned}$$

so that all of the pricing models $\mathbb{Q} \in \mathcal{Q}$ return a price within $[\underline{\pi}(X), \bar{\pi}(X)]$. Then a coherent measure of model risk is given by

$$\mu_{\mathcal{Q}} = \bar{\pi}(X) - \underline{\pi}(X)$$

The idea of utilizing only pricing models capable of reproducing the prices of benchmark instruments may be overly restrictive in practice. Cont (2006) addresses this shortcoming by introducing the notion of a convex measure of model uncertainty based primarily upon extending \mathcal{Q} to a wider class of pricing models and then penalizing each model by its pricing error, such that we now redefine

$$\begin{aligned}\bar{\pi}(X) &= \sup_{\mathbb{Q} \in \mathcal{Q}} \{E_{\mathbb{Q}}[X] - \|C^* - E_{\mathbb{Q}}[H]\|\} \\ \underline{\pi}(X) &= \inf_{\mathbb{Q} \in \mathcal{Q}} \{E_{\mathbb{Q}}[X] - \|C^* - E_{\mathbb{Q}}[H]\|\} \\ &= -\bar{\pi}(-X)\end{aligned}$$

where $H = (H_i)_{i \in I}$ and C^* is the vector of associated market prices. Different choices for the pricing error term given by the authors include

$$\|C^* - E_{\mathbb{Q}}[H]\|_{\infty} = \sup_{i \in I} |C_i^* - E_{\mathbb{Q}}[H_i]|$$

$$\|C^* - E_{\mathbb{Q}}[H]\|_1 = \sum_{i \in I} |C_i^* - E_{\mathbb{Q}}[H_i]|$$

$$\|C^* - E_{\mathbb{Q}}[H]\|_p = \left[\sum_{i \in I} |C_i^* - E_{\mathbb{Q}}[H_i]|^p \right]^{\frac{1}{p}}$$

$$\|C^* - E_{\mathbb{Q}}[H]\|_{1,w} = \sum_{i \in I} w_i |C_i^* - E_{\mathbb{Q}}[H_i]|, w_i \geq 1$$

Under this approach $\bar{\pi}(-X)$ is now a convex risk measure, with associated penalty function

$$\begin{aligned}\alpha(\mathbb{Q}) &= \|C^* - E_{\mathbb{Q}}[H]\| \text{ if } \mathbb{Q} \in \mathcal{Q} \\ &= \infty \text{ if } \mathbb{Q} \notin \mathcal{Q}\end{aligned}$$

Whereas the risk metrics given by Cont (2006) provide point estimates of the model uncertainty, Deryabin (2012) introduces the metric model uncertainty measure and derives upper and lower bounds on the calibration risk for a specific model based on the calibration errors associated with different parameter choices. The author begins by defining a region B_ρ of “acceptable” parameter combinations² θ within a parameter constellation Λ_0 , such that,

$$B_\rho = \{\theta, d_p(\theta, \theta_0) \leq \rho\}$$

where, using the notation of Cont (2006) given above,

$$d_p(\theta_1, \theta_2) = \|E_{\theta_1}[H] - E_{\theta_2}[H]\|_p$$

for some integer p , and

$$\theta_0 = \min_{\theta} d_p(C^*, \theta)$$

Similarly to Cont (2006), the author goes on to give a model uncertainty measure defined as

$$\mu(X, B_\rho) = \bar{\pi}(X, B_\rho) - \underline{\pi}(X, B_\rho)$$

where

$$\begin{aligned}\bar{\pi}(X, B_\rho) &= \sup_{\theta \in B_\rho} E_\theta[X] \\ \underline{\pi}(X, B_\rho) &= \inf_{\theta \in B_\rho} E_\theta[X]\end{aligned}$$

The author derives a lower bound for $\mu(X, \Lambda_0)$ based on the initial calibrated parameters θ_0 , and given by

$$\underline{\mu}(X, \Lambda_0) = (r_+ + r_-) \|\nabla_{\theta} \pi(X, \theta_0)\| + \mathcal{O}(r_+^2 + r_-^2), \quad r_+, r_- \rightarrow 0$$

where r_+ and r_- are obtained through shifting the gradient operators, ∇_{θ} and $-\nabla_{\theta}$ respectively, from the initial parameters θ_0 to the boundaries of Λ_0 . Similarly, an upper bound for $\mu(X, \Lambda_0)$ is given by,

$$\bar{\mu}(X, \Lambda_0) = 2R \|\nabla_{\theta} \pi(X, \theta_0)\| + \mathcal{O}(R^2), \quad R \rightarrow 0$$

²Note that Deryabin (2012) denotes a parameter combination as λ , we have chosen to use θ to remain consistent with the notation used throughout this chapter.

where, again using the notation of Cont (2006),

$$R^2 = \sum_{i \in I} w_i (C_i^* - E_{\theta_0} [H_i])^2$$

Bannör and Scherer (2013b) focus on parameter uncertainty as a specific case of model uncertainty. The authors introduce the concept of a risk functional, Γ , for incorporating parameter risk into derivative prices. This approach mirrors that of Cont (2006) in the sense that ones bid and offer are now defined by $[-\Gamma(-X), \Gamma(X)]$. The authors go on to show how these functionals can be used to derive push-forward densities for exotic option prices induced by the risk associated with the model parametrization. It is this methodology which we will employ when constructing storage value densities induced by both calibration and parameter estimation risk. Γ has certain desirable properties which we present here in a non-technical fashion³:

- Order Preservation: If one payoff dominates another then it should not incur a greater model risk premium.

$$X \geq Y \Rightarrow \Gamma(X) \geq \Gamma(Y)$$

- Diversification: Combining of payoffs exposed to model risk should not increase ones model risk premium

$$\Gamma(\lambda X + (1 - \lambda)Y) \leq \lambda \Gamma(X) + (1 - \lambda)\Gamma(Y) \quad \lambda \in [0, 1]$$

- Model Independence: If the payoff's value is model independent then there should be no model risk premium incorporated in its price

$$\Gamma(X) = E_{\mathbb{Q}} [X] \quad \forall \mathbb{Q} \in \mathcal{Q}$$

where \mathbb{Q} is one of a family of potential models \mathcal{Q} with an associated probability distribution R .

The functional Γ generated by normalized, law invariant⁴, convex risk measures ρ satisfy the above properties such that

$$\Gamma(X) := \rho(\mathbb{Q} \mapsto E_{\mathbb{Q}} [X])$$

These functionals can thus be used to rank different models based upon the relative levels of parameter risk they carry. Of course the distribution, R , of the measures \mathbb{Q} will generally be unknown. However, in many cases the estimators of the models will possess asymptotic normality in which case the *delta method*, as described by Bannör and Scherer (2013b), can be used. Specifically, the delta method provides that for an asymptotically normal sequence of (consistent) estimators $\{\theta_N\}_{N \in \mathbb{N}}$ of the true model parameters $\theta_0 \in \Theta$ with positive definite covariance matrix Σ (whereby $\Sigma^{-1} = I(\theta)$ is the Fisher Information matrix),

³See Properties 5, Bannör and Scherer (2013b) for a technical treatment.

⁴ ρ is said to be law invariant if for two claims X, Y with probability measures $P^X = P^Y$, then $\rho(X) = \rho(Y)$ holds.

if $\theta \mapsto E_\theta[X]$ is continuously differentiable and the gradient $\nabla E_\theta \neq 0$ then, as $N \rightarrow \infty$,

$$\sqrt{N} (E_{\theta_N}[X] - E_{\theta_0}[X]) \longrightarrow \mathcal{N} \left(0, (\nabla E_{\theta_0}[X])' \cdot \Sigma \cdot \nabla E_{\theta_0}[X] \right) \text{ weakly,}$$

where $E_{\theta_N}[X]$ is the expected value of the payoff using the sample parameter estimates and $E_{\theta_0}[X]$ is the expected value under the true model parameters, and $\mathcal{N}(\cdot, \cdot)$ is the normal distribution.

Bannör and Scherer (2013b) introduce the $\theta_N * AVaR$ risk captured functional generated by the Average Value at Risk (AVaR) convex risk measure. The offer price of a claim X under this functional is then approximated by

$$\theta_N * AVaR(X) \approx E_{\theta_0}[X] + \frac{\varphi(\Phi^{-1}(1-q))}{q\sqrt{N}} \sqrt{(\nabla E_{\theta_0})^\top \Sigma \nabla E_{\theta_0}}$$

where φ and Φ represent the standard normal density and cumulative distribution respectively. Thus the mid price is adjusted by an amount

$$\frac{\varphi(\Phi^{-1}(1-q))}{q\sqrt{N}} \sqrt{(\nabla E_{\theta_N}[X])^\top \Sigma \nabla E_{\theta_N}[X]} \quad (4.1.1)$$

where $q \in (0, 1)$ and

$$\nabla E_{\theta_N}[X] = \frac{E_{\theta_{N+\varepsilon}}[X] - E_{\theta_{N-\varepsilon}}[X]}{2\varepsilon} \quad (4.1.2)$$

Bannör and Scherer (2013b) go on to apply risk functionals to the errors associated with model calibration in a manner similar to the approach of Deryabin (2012). The procedure involves first calculating the calibration errors for a set of model parameter combinations within a region B_ρ , as defined above. The next step is to construct the probability distribution R of the model parameter combinations based upon their calibration errors. Bannör and Scherer (2013b) do so by defining a transformation function h , which maps the error function $\varepsilon(\theta)$ to the real line and possesses the following properties:

- h is decreasing. This is to ensure that parameter combinations which yield a higher total error are given less likelihood.
- $\int h(\varepsilon(\theta)) d\theta = 1$. This is simply a normalization to ensure the function meets the definition of a probability measure.

One example of such a transformation function is the normal transformation function

$$h_\lambda^N(t) := c \exp \left(- \left(\frac{t-t^*}{\lambda} \right)^2 \right) \quad t \geq t^* \quad (4.1.3)$$

with scaling parameters $\lambda > 0$ and $c > 0$ chosen so that the function returns a probability distribution centered on the minimum calibration error t^* . Bannör and Scherer (2013b) go on to show how one can

utilize this density to construct a push-forward density of exotic option prices induced by the calibration error associated with a given model. They present empirical results based on the pricing of Asian and Barrier options under a number of models calibrated to equity index option prices and demonstrate the dependency of parameter risk upon both the model and the asset being valued.

4.1.3 Energy Model Risk Analysis

To date, the model risk literature in relation to energy market assets is quite sparse. Bannor et al. (2013a) were the first to focus on this area where they investigate the parameter risk associated with a power plant valuation model with model parameters estimated from historical data. In the case of gas storage valuations, Henaff et al. (2013) took a similar approach in estimating the parameter risk associated with two proposed price models, which again were estimated using only historical information. We extend the work of Henaff et al. (2013) by analyzing the storage model parameter risk with reference to models which are both forward curve consistent and calibrated to market traded options. This proposed approach is much more in line with the valuation and model risk methodologies utilized in the exotic options space and is far more reflective of how an energy trading business manages their real asset optionality. However, market calibration alone is generally not sufficient to capture the rich dynamics observable in many energy markets. Therefore we propose an extension of the risk functional approach of Bannör and Scherer (2013b) which allows us to estimate the combined calibration and parameter estimation risk, utilizing both market and historical data in a consistent manner.

As a specific example, realistic models of the natural gas forward curve cannot be calibrated to benchmark instruments alone due to the lack of a liquid time-spread options market and thus the correlation structure is typically estimated from historical data and then approximated by a suitable model specification. Storage valuation models are particularly sensitive to the model implied correlation structure and thus highly exposed to the parameter risk associated with its estimation. However, the overall level of volatility implied by the model is constrained to be estimated from the market so as to attain consistency with the products used to hedge volatility risk, either directly or as part of a larger derivatives book. This situation is the primary motivating factor behind our innovative approach as it requires a combination of both the parameter distribution arising from historical estimation and the error function distribution from market calibration in deriving a risk adjusted price. The utility of this framework would benefit business units across the deal life-cycle. Beginning with the initial implementation of a model, this approach would allow a model validation team to evaluate several models with potentially differing reliance on historical and market data in a consistent manner. In terms of calibration risk, it would allow the traders and product control teams to separately identify the risk due to market calibrations from the risk due to historical parameter estimation and attribute a suitable amount of capital in reserve to mitigate against both.

We begin by denoting a general model of the forward curve $F(t, T; \theta_m, \theta_h)$, dependent upon market calibrated parameters θ_m and parameters θ_h estimated from historical data. For asset valuations, θ_h can be heuristically characterized as the model information not reliably available from the market. The most obvious way to proceed would be to assume no knowledge of the parameter estimates sample distribution as per the market calibration approach outlined above. However, this would ignore the information we have in relation to the distribution of θ_h inherent in the estimation of the parameters. To that end, we must generalize the transformation function, given by Equation 4.1.3, such that it is capable of incorporating this additional useful information. We begin by redefining our parameter set θ to be the union of θ_m and θ_h , and the transformation function as

$$h(\varepsilon(\theta)) = h_m(\varepsilon(\theta) | \theta_h) p(\theta_h) \quad (4.1.4)$$

The function $h(\varepsilon(\theta))$ has the same properties as before but is now decomposed into the product of $h_m(\varepsilon(\theta) | \theta_h)$; the error term density conditional upon the historically estimated parameters θ_h , and the sampling error density $p(\theta_h)$, associated with the historical estimation. As highlighted by Bannor et al. (2013a), we know that the maximum likelihood estimators of the parameters θ_h calculated from a sample of N observations are asymptotically normally distributed⁵,

$$\sqrt{N}(\theta_h^N - \theta_h) \sim N(0, \Sigma)$$

This gives us a suitable function $p(\theta_h)$, namely the Gaussian density, which, once scaled such that it integrates to unity, can be used to give relative weight to the calibration induced probabilities. From an implementation point of view, the steps taken are very similar to Bannor et al. (2013a). We need to create a push forward density from the market option pricing error density, however, we will then apply the delta method described by the authors in conjunction with the sampling error density to construct a joint market and historical parameter risk induced value density. The following section sets this out in detail.

Storage Model Risk

In the case of storage valuations the parameter of interest, which is typically estimated from historical data, is the covariance matrix of relative maturity forward curve returns. With respect to combined market calibration and historical estimation, the correlation matrix or scaled covariance matrix is more appropriate as the absolute volatility levels can be inferred from the market. Further for computational purposes, the dimensionality of the forward curve model is typically reduced through spectral decomposition. In this class of dimensionally reduced models, the historically estimated parameters will relate to the structure of selected eigenvectors of the covariance matrix. A detailed summary of this PCA based approach to energy forward curve modelling is given by Clewlow and Strickland (1999a) and Carmona and Coulon

⁵The same is true of any asymptotically normal estimators, such as Bayes estimators, of θ_0 .

(2014), while the approach is also well established in the interest rate market, see for example Driessen et al. (2003) and Litterman and Scheinkman (1991).

In natural gas markets it is typical for such a decomposition to yield two to three factors which account for a large proportion of the overall forward curve variability. If, for example, one was to plot the eigenvector relating to the first factor, that is the factor which contributes the most to the overall variance, it will typically display a mild exponential decay. Of course, the values of the eigenvector represent the sensitivity of forward curve returns at each tenor to changes, in this somewhat abstract, stochastic factor. This behavior is similar to that given by traditional single factor Ornstein-Uhlenbeck models of the forward curve and thus the first factor can be modeled as such. Similar model specifications can be constructed to capture the sensitivity of forward curve returns to the other selected factors. We should emphasize at this point that the only reason for modeling factors in this manner is for computational efficiency in approximating the estimated covariance matrix.

Of course, such an approximation will naturally imply that the forward curve returns covariance matrix implied by the model will differ from the historical sample covariance matrix. It is our contention that this approximation error should not directly impact the risk associated with a particular estimation rather it should be noted as a deficiency of the model during the specification phase of model development. It is of course distinct from parameter risk by definition, as the model parameters relate only to the eigenvector values which impact the approximating model volatility function.

In order to quantify the parameter estimation risk, we begin by defining the covariance matrix of M relative maturity forward curve returns to be $C \in \mathbb{R}^{M \times M}$, with spectral decomposition

$$C = WDW^{-1}$$

and denote the dimensionally reduced approximation of C as C_e . Further, let us denote the model approximation of C_e as C_a with approximation error given by $C_e - C_a = \mu_\theta$, which as discussed above we take as given. Denote the covariance generating function associated with this model as $g(\theta_h)$ such that $g(\theta_h) = C_a$. We then have

$$\sqrt{N} (C_e^N - C_e) \sim N(0, \Sigma)$$

where Σ is the inverse Fisher Information matrix associated with the sample covariance matrix C_e^N estimated from a sample of size N . Therefore

$$\sqrt{N} (C_a^N - C_a) \sim N(0, \Sigma)$$

and so

$$\sqrt{N} (g(\theta_h^N) - g(\theta_h)) \sim N(0, \Sigma)$$

Applying the delta method we then have

$$\begin{aligned} \sqrt{N} (g^{-1}g(\theta_h^N) - g^{-1}g(\theta_h)) &\sim N\left(0, (\nabla g^{-1})^\top \Sigma \nabla g^{-1}\right) \\ \sqrt{N} (\theta_h^N - \theta_h) &\sim N\left(0, (\nabla g^{-1})^\top \Sigma \nabla g^{-1}\right) \end{aligned}$$

For a storage payoff given by X , applying the delta method again gives the sampling error induced storage value density

$$\sqrt{N} \left(E_{\theta_h^N} [X] - E_{\theta_h} [X] \right) \sim N\left(0, (\nabla E_{\theta_h})^\top (\nabla g^{-1})^\top \Sigma \nabla g^{-1} \nabla E_{\theta_h}\right) \quad (4.1.5)$$

This last result gives the storage value distribution induced by the uncertainty, represented by Σ , of the forward curve covariance matrix. The relationship between the two can be understood as first, weighting the matrix Σ by the sensitivity of the model parameters to the forward curve covariance matrix represented by ∇g^{-1} and then secondly, weighting the result by the sensitivity of the storage value to the model parameters, represented by ∇E_{θ_h} . The steps required to derive the parameter risk variance are summarized below, further details can be found in Section C.1 of the Appendix

1. Derive the sample covariance matrix, C , of the relative maturity forward curve returns.
2. Using the sample Fischer Information Matrix associated with C , derive the sample parameter covariance matrix, Σ .
3. Derive the gradient of the model parameter estimation function, ∇g^{-1} . The estimation function takes the matrix C as input and returns a vector of model parameters.
4. Derive the sample model parameter covariance matrix given by, $(\nabla g^{-1})^\top \Sigma \nabla g^{-1}$.
5. Finally, using the storage value parameter sensitivities ∇E_{θ_h} , derive the parameter risk variance as per Equation 4.1.5.

4.2 Storage Parameter Risk: Numerical Examples

We will now present numerical results utilizing three of the forward curve models detailed in Chapters 2 and 3. These provide canonical examples of the different approaches to calibration and parameter estimation risk measurement outlined above. We will begin with a non-technical review of the pricing models to be tested and relate them to the information we wish to incorporate into a storage pricing model, namely

the volatility surface and the scaled covariance matrix of forward curve returns ⁶. As discussed previously, the choice of models at our disposal is limited due to the fact there is very little in the literature which approaches the storage valuation problem with a view to maintaining consistency with the volatility smile. Thus although price models developed for storage valuation have the ability to capture the forward curve time-spread dynamics well, for example Parsons (2013) and Boogert and De Jong (2011), they are generally driven by diffusion processes and thus fail to replicate the volatility smile observable in gas option markets. Conversely, traditional price models which are capable of replicating the smile, Heston⁷ or SABR⁸ for example, will yield perfectly correlated forward curve returns and therefore little or no extrinsic value to the storage.

We will proceed by first constructing push forward densities for the storage values associated with each model specification and show how a trading desk can use this information to construct calibration and parameter estimation risk adjusted bid-offer levels capable of providing a buffer against p&l risk owing to model calibration/estimation. Finally, we will discuss how the AVaR model uncertainty measure can be used by a risk or model validation team to rank the individual model specifications by focusing on a practical example of a proposed model being tested against an incumbent.

The first of our test models is the Mean Reverting Variance Gamma (MRVG) model detailed in Section 2.1. Recall, the log-spot price model is specified by the following dynamics.

$$dx(t) = \left(\frac{\partial f(0,t)}{\partial t} - \kappa_j (\exp(-\alpha t)) + \alpha f(0,t) - \alpha \int_0^t \kappa_j (\exp(-\alpha(t-s))) ds - \alpha x(t) \right) dt + dX(t) \quad (4.2.1)$$

where $f(0,t)$ is the initial log forward price for time t and $dX(t)$ is a driftless Variance-Gamma process with κ_j its cumulant.

As demonstrated in Chapter 2, this model can be calibrated directly to market option prices and thus the measure distribution R can be obtained using the approach of Bannör and Scherer (2013b). The model is specified such that the instantaneous variance of different maturities along the forward curve shows exponential decay as time to maturity increases. This property, known commonly as the Samuelson Effect Serletis (1992), is a stylized feature of many commodity markets and particularly natural gas. In the case of constant σ , the exponential decay rate, α , fully controls the structure of the covariance amongst different points on the curve. This results in non-parallel shifts in the forward curve due to changes in the underlying stochastic driver. It is this last element which is crucial for storage valuation as it results in the operator exercising their optionality to switch planned injection and withdrawals thus creating extrinsic

⁶The scaled covariance matrix is defined as the covariance of returns matrix relative to the volatility of the shortest maturing contract.

⁷Heston (1993)

⁸Hagan et al. (2002)

value. In the spot price representation of the model, this parameter is more commonly referred to as the mean reversion rate as it controls the speed at which the spot price reverts to the expected forward level. The ν parameter controls the implied volatility smile attenuation and ensures that the model is consistent with the initial volatility surface

The second model is the MRJD detailed in Section 2.1. The model, which was first specified by Deng (2000) and later used in swing contract valuations by Kjaer (2008), is equivalent to Equation 4.2.1. The only difference is in the choice of stochastic driver, $dX(t)$, which in this case is a jump-diffusion process with Compound Poisson jump process driven by a double Exponential distribution. Thus, as in the MRVG model, the presence of the parameter α is the primary driver of the storage value. The parameters relating to the jump-diffusion are σ , the diffusion volatility, λ , the jump arrival rate, and μ which controls the size of both positive and negative jumps. The jump diffusion specification is what allows us to replicate the volatility smile using this model, in a similar manner to the MRVG model.

The final model, the MRVG-3x presented in Section 3.1, is a multifactor extension of the first. The inclusion of the second factor allows more flexibility in modelling the covariance structure of the forward curve and therefore should produce storage values more representative of the underlying dynamics. Recall that the dynamics of the log spot price under the MRVG-3x model are given by

$$dx(t) = \frac{\partial f(0,t)}{\partial t} + dy^{(1)}(t) + dy^{(2)}(t)$$

where

$$\begin{aligned} dy^{(1)}(t, T) &= \left(-\kappa_{j(1)} (b \exp(-\alpha(T-t)) \sigma) \right) dt + b \exp(-\alpha(T-t)) \sigma dX(t) \\ dy^{(2)}(t, T) &= -\frac{1}{2} (\exp(-\varepsilon(T-t)) c_1 \sigma)^2 dt + (\exp(-\varepsilon(T-t)) c_1) \sigma dW(t) \end{aligned} \quad (4.2.2)$$

and $dW(t)$ is a standard Brownian Motion.

The model is composed of two factors. The first, which accounts for the majority of the forward curve variability is a Mean Reverting Variance Gamma process. As with the MRVG model, the main parameters for this factor can be calibrated to the options market. The parameter b represents the proportion of total variance attributed to the first factor and would need to be estimated from the historical data. The second factor is specified such that it approximates the typical shape of the sensitivity, which we refer to as the volatility function, of the forward curve to the second principal component of the forward curve returns covariance matrix. An example of this is given below by Figure 4.2.1




Figure 4.2.1: MRVG-3x 2nd Factor Volatility Function

The fit between the model and historically estimated second volatility function for the MRVG-3x model.

As outlined in Chapter 3, the parameters relating to the second factor can be estimated directly from the eigenvector values. The parameter c_1 determines the value of the function as it approaches zero time to maturity and ε controls the decay of the function as maturity increases. The shape of the volatility function, which is determined by these two parameters, will have a direct impact on the covariance of different maturities along the forward curve. A sharply decaying curve will decrease the covariance between prompt forward prices and the back of the curve, which will lead to greater time-spread variance and thus higher storage value. This model, as it encompasses the first, will generally attribute more value to a storage asset through this additional decorrelation of the forward curve returns. One obvious concern, which we hope to address, is the extent to which this additional value should be discounted on the basis of the parameter estimation risk inherent in the model estimation.

We will again focus on the storage deal valued in previous chapters, a simple *20in/20out* deal. This means the asset takes 20 days to inject to full capacity from empty (In) and 20 days to empty once full (Out). The deal commences immediately on the options quote date and lasts 1 year. As before, the option data was sourced from Bloomberg with a quote date of 19th December 2012. The data set used to estimate the historical covariance matrix covers the 3 years ending 19th December 2012, where each day in the sample contains the day-ahead and month ahead quotes spanning 11 months.

The procedure for estimating calibration risk outlined by Bannör and Scherer (2013b) begins by discretizing the parameter space and then evaluating the pricing error to benchmark instruments for each parameter combination in the space. Thus they reduce the multidimensional parameter constellation to

a vector of error terms. Given that a large majority of these errors would fall outside of what would be deemed a reasonable calibration, the authors propose discarding parameter combinations which return an error greater than some cut-off point. In their numerical examples the authors use a cut-off for the error term, given by the root mean square error, of 2% above the minimum error term. Obviously the choice of cut-off will depend upon the average bid-offer percentage spread observed in the options market. For this example, we have chosen a cut-off of 3% in order to reflect the relatively lower liquidity of natural gas options.

There are two main sources of risk associated with a given parameter combination which yields a RMSE within this cut-off level.

1. Risk due to the option pricing model, emanating from the assumed dynamics of the model, and any numerical error inherent in the pricing algorithm. Further, as outlined in Cont and Tankov (2003) the calibration error landscape in exponential Lévy models may have flat regions in which the error is insensitive to variations in the parameters, thus the behaviour of the objective function minimization algorithm within these areas creates an additional source of risk. Combined these risks are characterized by an acceptance that the true minimum lies within a narrow region containing the calibrated values. Whilst the variability in parameters may be low within this region the potential effect on the pricing of complex derivatives means the risk can still be quite substantial. In collectively identifying these risks we will follow the convention of Deryabin (2012) who refers to the risk due to small variations in the parameters as “local” calibration risk.
2. Risk due to the presence of multiple local minima. Generally speaking, the calibration problem is ill-posed and the solution may have a strong dependency on the starting point of the search algorithm. This problem is made worse as one widens the acceptable cut-off error, which may lead to valid solutions at many distant points in the parameter space. Cont and Tankov (2003) demonstrate the dependency of the Merton jump-diffusion model on the choice of initial parameter values when calibrating to equity options and conclude that the number of local minima depends less on the number of model parameters relative to option prices and more on the convexity of the objective function. Obviously, this issue is more severe than the first as it could potentially destroy all confidence in the calibrated parameters and lead to a very high level of model risk. We will refer to this risk as “multiple minima” calibration risk. Although some authors suggest overcoming this issue through regularization of the objective function, Avellaneda et al. (1997); Crépey (2003), we tend to agree with Cont and Ben Hamida (2005) that the presence of multiple minima contains information on the level of model risk which may be lost through regularization.

For our purposes, one could in theory, search over a dense and wide ranging discretization of the parameter space in order to identify regions containing local minima. However, there is a computational demand when proceeding in this manner which necessitates the need for approximation. This is particularly true

when calibrating to non-standard option contracts, as is the case here for monthly swaptions contracts. We therefore suggest, first searching for local minima by selecting the initial parameter values in ones search function from a wide and relatively sparse discretization of the parameter space. Once each of these local minima have been identified one can construct a region encompassing each in order to identify the local calibration risk.

For each of the models specified above the calibration to our options data using the simplex method of Lagarias et al. (1998) did produce multiple local minima dependent upon the choice of initial parameters. However, in each case the local minimum either yielded a RMSE far in excess of our acceptable cut-off or returned parameter values similar to our optimal values and thus could be considered “local” calibration risk as defined above. Therefore we will focus only on a reduced parameters space encompassing our initial parameter estimates.

Example 1: Market Based Calibration Risk

In this first example we will focus on the MRVG and MRJD models as they require estimation from market data alone. Following the example of Bannör and Scherer (2013b), we choose $h_m()$ to be the normal transformation given by Equation 4.1.3. For the scaling parameter, λ , we have used the sample variance of the calibrated errors centered on the minimum value and chosen c such that the density values sum to unity. For the MRVG model, we see from Equation 4.2.1 the parameter set of interest is given by $\theta_m = \{\alpha, \sigma, \nu\}$. The initial calibrated parameters are given in Table 4.2.1. The associated root mean square error (RMSE) is 1.07%.

	α	σ	ν	RMSE
MRVG	0.2162	0.201	0.2560	1.07%

Table 4.2.1: MRVG Model Parameters.

The calibrated model parameters for the MRVG model. ‘RMSE’ refers to the root mean square error between the model and market option prices.

To investigate the local risk of our initial calibration, we select α values ranging from 70% to 130% in steps of 5% of the calibrated α . For the σ and ν values we use a much tighter range of 90% to 110% in steps of 1.5%, which reflects the greater sensitivity of the RMSE to changes in these parameters. In total, this yielded 2548 distinct parameter combinations. We next iterate over this parameter space and discard any parameter combinations which exceed the minimum RMSE by the 3% limit. The resulting parameter space consisted of 807 parameter combinations. Using the error terms associated with these parameters we derive an error density as per Equation 4.1.3 and then evaluate the storage value at each point in

this reduced parameter space. The resulting calibration risk induced storage value density is displayed graphically in Figure 4.2.2⁹ and the valuation results are given in Table 4.2.2. As is evident from the graph, the density appears almost uniform over the acceptable calibration range. The apparent presence of columns of points in Figure 4.2.2 correspond to storage values at each α point in the parameter space. This behaviour can be rationalized by observing that changes in α have little effect on the calibration error relative to changes in σ and ν due to the sensitivity of the implied volatility smile to changes in model volatility and kurtosis. This means that a much wider range of α values fall within the acceptable parameter space. Further, due to the sensitivity of storage values to the model covariance structure, the choice of α will impact the value much more than both σ and ν .




Figure 4.2.2: MRVG Storage Value Push Forward Density

The empirical calibration risk induced storage value probability density for the MRVG model. The red line corresponds to the storage value with the smallest calibration error.

	Value
Modus	11.2087
Expected Value	11.2116
Coefficient of Variation	0.38%
Skewness	-0.034
Minimum RMSE	1.05%

Table 4.2.2: MRVG Storage Value Push Forward Distribution

Summary statistics of the calibration risk induced storage value distribution under the MRVG model.

⁹Values were obtained using the FFT-m based algorithm detailed in Chapter 2 with a grid size of 2¹¹.

For the MRJD model, the parameter set is $\theta_m = \{\alpha, \sigma, \lambda, \mu\}$. The initial calibrated parameters are given below in Table 4.2.3, the RMSE is slightly higher than the MRVG model at 1.10%.

	α	σ	λ	μ	RMSE
MRJD	0.2099	0.0334	8.7966	0.047	1.10%

Table 4.2.3: MRJD Model Parameters

The calibrated model parameters for the MRJD model. 'RMSE' refers to the root mean square error between the model and market option prices.

We have again selected the α range to be 70% to 130% in steps of 5% of the calibrated α . The values for σ, λ and μ were selected to range from 90% to 110% in steps of 4% of the initial calibrated values, again this tighter range is reflective of the greater sensitivity of the RMSE to changes in these parameters. In total this gave 2808 distinct parameter combinations, of which 795 were within the cut-off threshold. The resulting push forward storage density is displayed in Figure 4.2.3 and summary statistics are given in Table 4.2.4.

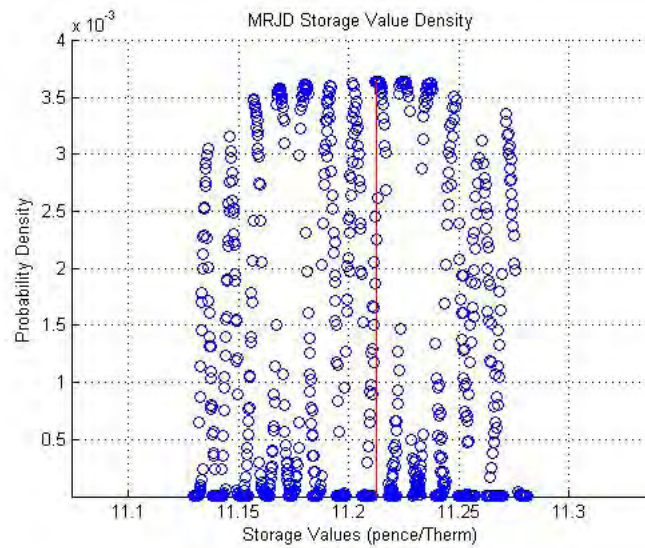


Figure 4.2.3: MRJD Storage Value Push Forward Density

The empirical calibration risk induced storage value probability density for the MRJD model. The red line corresponds to the storage value with the smallest calibration error.

	Value
Modus	11.2127
Expected Value	11.2039
Coefficient of Variation	0.36%
Skewness	0.021
Minimum RMSE	1.08%

Table 4.2.4: MRJD Storage Value Push Forward Distribution

Summary statistics of the calibration risk induced storage value distribution under the MRJD model.

As we can see visually, and also from the summary statistics, there is not much difference between the push forward densities under each model. The MRVG model returns a higher expected value than the MRJD model but also has a slightly higher level of variability. It would appear on the basis of this analysis that both models carry a roughly equivalent level of calibration risk. Later, we will use these results to construct bid-offer levels which may shed some additional light on the relative riskiness of these models.

Example 2: Joint Calibration and Parameter Estimation Risk

As discussed above the MRVG-3x model requires calibration to market data and also parameter estimation from historical data. We see from Equation 4.2.2 that the full parameter set in this case is $\theta = \{\alpha, \sigma, \nu, b, c_1, \varepsilon\}$. The market calibrated parameter set is $\theta_m = \{\alpha, \sigma, \nu\}$. Recall that these parameters determine the shape of the volatility term structure of the first factor of the model. The historically estimated parameters, $\theta_h = \{b, c_1, \varepsilon\}$, determine the shape of the volatility term structure of the second factor of the model. The initial market calibrated and historically estimated model parameters are given in Table 4.2.5.

	α	σ	ν	b	c_1	ε	RMSE
MRVG-3x	0.1148	0.2518	0.1675	0.7511	0.6254	12.734	1.86%

Table 4.2.5: MRVG-3x Model Parameters

The calibrated and historically estimated model parameters for the MRVG-3x model. 'RMSE' refers to the root mean square error between the model and market option prices.

Recall from Equation 4.1.4 that the calibration risk induced density for this model shall be decomposed into the conditional market calibration error density $h_m(\varepsilon(\theta) | \theta_h)$ and the parameter estimation risk density corresponding to the historically estimated parameters $p(\theta_h)$. We shall proceed by first estimating the former. We have chosen the same parameter space discretization as in the MRVG case, that is α

values ranging from 70% to 130% in steps of 5% of the calibrated α and σ, ν values ranging from 90% to 110% of their initial calibrated values in steps of 1.5%. This returned 2548 parameter combinations, of which 838 were within the threshold of 3% of the minimum calibration error of 1.35%¹⁰. The push forward density is displayed graphically in Figure 4.2.4, and summary statistics are given in Table 4.2.6.

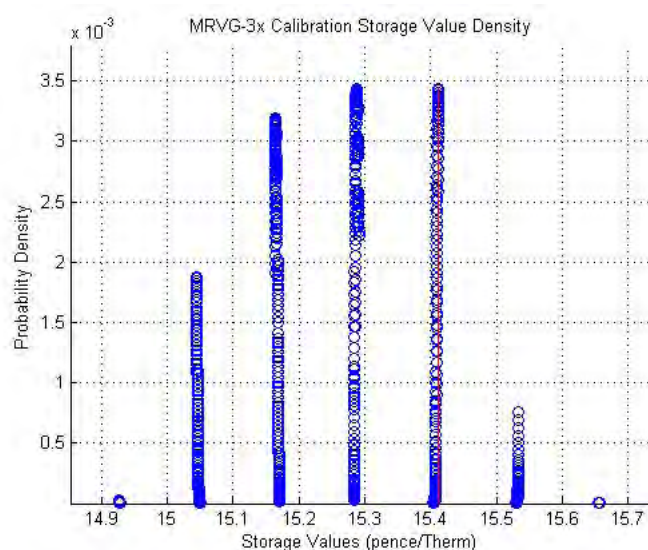


Figure 4.2.4: MRVG-3x Calibration Error Storage Value Density

The empirical calibration risk induced storage value probability density for the MRVG-3x model. This density is conditional upon the historically estimated parameters. The red line corresponds to the storage value with the smallest calibration error.

	Value
Modus	15.4117
Expected Value	15.2558
Coefficient of Variation	0.69%
Skewness	0.0033
Minimum RMSE	1.35%

Table 4.2.6: MRVG-3x Conditional Storage Value Push Forward Distribution

Summary statistics of the conditional calibration risk induced storage value distribution under the MRVG-3x model.

¹⁰This RMSE value is much less than that of the initial calibration and was found when iterating through the parameter space. This is due to the use of the Moment Matching technique in calibrating the parameters which is less accurate than calibrating directly to option prices.

Comparing with the MRVG and MRJD cases we see that the value is much higher but also the coefficient of variation is almost double that of the single factor models and thus the confidence one would have in the calibrated value is lower. From the plot of the push forward density we observe a clustering of storage values at distinct levels. These points relate to different σ values in the discretized parameter space. Recall that in the MRVG and MRJD models there was an apparent clustering at distinct α values due to the sensitivity of the storage value to changes in this parameter. For the MRVG-3x model however, the value will be primarily driven by the relatively high estimated ε value, which will act to decorrelate the short and far maturities on the forward curve. From the model specification given in Equation 4.2.2 we can see that the model sensitivity to ε will be proportional to σ which gives rationale for the storage value density behaviour observable in Figure 4.2.4.

We now proceed to deriving the parameter estimation risk density associated with the historically estimated parameters. The link between the eigenvalue decomposition of the forward returns covariance matrix and the factor volatility as a function of relative maturity is detailed in Section C.1 of the Appendix. Essentially, the factor volatility for a given maturity is defined as the square root of the factor variance, given by the eigenvalue, multiplied by the related eigenvector entry for that maturity. We can derive a sampling error covariance matrix for each of these volatility parameters, which in turn can be used to derive a sample covariance matrix for our model parameters. In order to derive the parameter estimation risk induced storage value density according to Equation 4.1.5, we must first estimate the Jacobian of our storage value with respect to our model parameters θ_h . We do so by numerical differentiation, the results of which are displayed in Table 4.2.7.

Parameter	Delta
b	0.0039
c_1	8.8059
ε	0.1402

Table 4.2.7: MRVG-3x Parameter Deltas

Partial derivatives of the storage value with respect to the historically estimated model parameters (Delta). Derivatives were numerically estimated using finite differences.

The c_1 parameter which controls the percentage of spot variance accounted for by the second factor has the greatest impact on the storage value, followed by the ε parameter which is ultimately responsible for the extrinsic value attributable to the second factor. Intuitively, this can be understood as the ε parameter reducing the auto-correlation of the second factor and c_1 being responsible for passing this decorrelation through into the forward curve returns. The b parameter has little impact on the storage value which is due to the low levels of extrinsic value attributable to the first factor. The resulting parameter estimation

risk induced storage value density is shown in Figure 4.2.5.

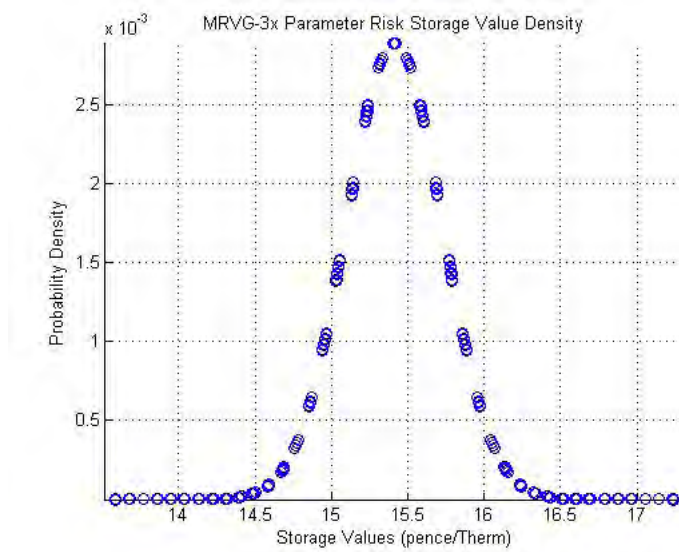


Figure 4.2.5: MRVG-3x Parameter Estimation Risk Storage Value Density

The historically estimated parameter estimation risk induced storage value density. The density is approximately Gaussian with variance given by Equation 4.1.5.

In order to construct the combined calibration and parameter estimation risk induced storage value density we make an assumption of independence between the calibration error and the values of the parameters estimated from historical data. This is primarily due to the computational effort involved in constructing the conditional calibration risk induced storage value density, which would grow exponentially if the density was a function of a second equally dense parameter space. This assumption reduces Equation 4.1.4 to

$$h(\varepsilon(\theta)) = h_m(\varepsilon(\theta))p(\theta_h)$$

The summary statistics of this combined density are given in Table 4.2.8. As we can see the variability of the value has increased dramatically due to the inclusion of the risk associated with the historically estimated parameters. The density is characterized by lower asymmetry reflected in the less negative skewness number.

	Value
Modus	15.4117
Expected Value	15.2528
Coefficient of Variation	2.17%
Skewness	0.0001

Table 4.2.8: MRVG-3x Calibration and Parameter Estimation Risk Push Forward Distribution.

Summary statistics of the joint parameter estimation and calibration risk induced storage value distribution under the MRVG-3x model.

Example 3: Deriving Risk Adjusted Price Levels

Now that we have evaluated the push forward densities associated with each of the models we can discuss the actionable analysis that can be carried out using this information. We begin with the choice of appropriate bid-offer levels attributable to parameter estimation and calibration risk. As discussed in the introduction, the main motivation for this is to allow traders to reserve some cash upfront to cover a potential loss due to the calibration of the pricing model being used. To inform this decision, we suggest first constructing the cumulative value distribution (CDF) and then choosing percentiles appropriate to one's risk tolerance as bid-offer levels. This approach can be seen as an informal alternative to that of Bannör and Scherer (2013b), where the authors use the risk functional $\Gamma(X)$ to derive bid and offer levels, $[-\Gamma(-X), \Gamma(X)]$.

Figure 4.2.6 displays the CDF associated with the MRVG model. Taking the 10th and 90th percentiles as bid and offer gives us price levels of 11.1507 – 11.2697, which equates to a spread of 1.06% of the mean value. In comparison, using the Average Value at Risk (AVaR) risk functional of Bannör and Scherer (2013b) naturally yields wider bid-offer levels for the 10th and 90th percentiles at 11.1449 – 11.2803. The MRJD cumulative distribution (Figure 4.2.7) gives similar bid and offer levels of 11.1473 – 11.2611, which is a spread of 1.02% of the mean value. The MRVG levels are marginally higher and wider than those of the MRJD model which is consistent with the results of Example 1. The spread given by the AVaR risk functional for the MRJD model is again wider at 11.1395 – 11.2688. Finally, for the MRVG-3x model, displayed in Figure 4.2.8, the bid-offer levels are at 14.8266 – 15.6772 or 5.58% of the mean. Thus, although the MRVG-3x model attributes more extrinsic value to the storage deal, there is a greater level of uncertainty as to the true value and a higher risk of suffering p&l variations attributable to the choice of model parameters. The bid-offer spread given by the AVaR induced risk functional in this case is 14.6753 – 15.8364.

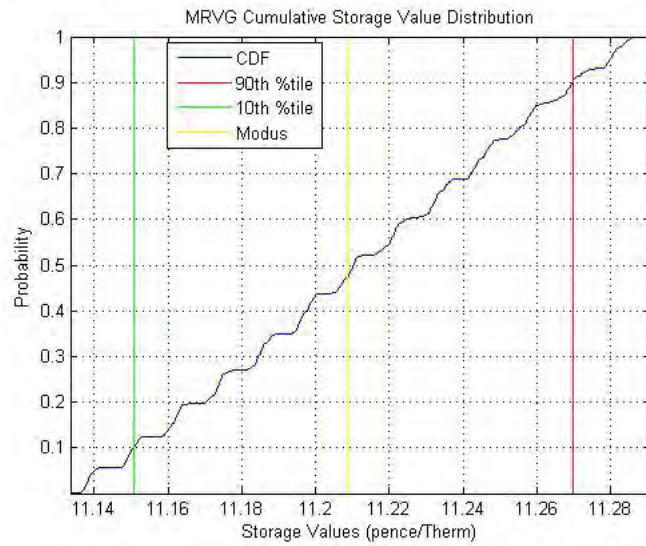


Figure 4.2.6: MRVG Storage Value Cumulative Distribution

The cumulative probability distribution of the calibration risk induced storage value density under the MRVG model.

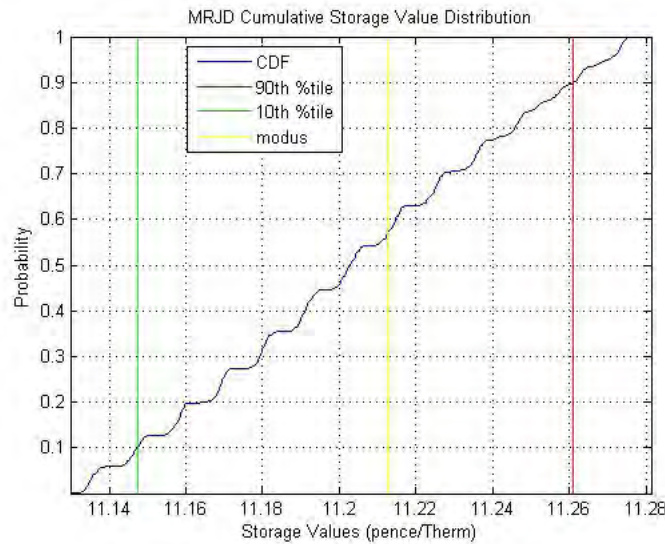


Figure 4.2.7: MRJD Storage Value Cumulative Distribution

The cumulative probability distribution of the calibration risk induced storage value density under the MRJD model.

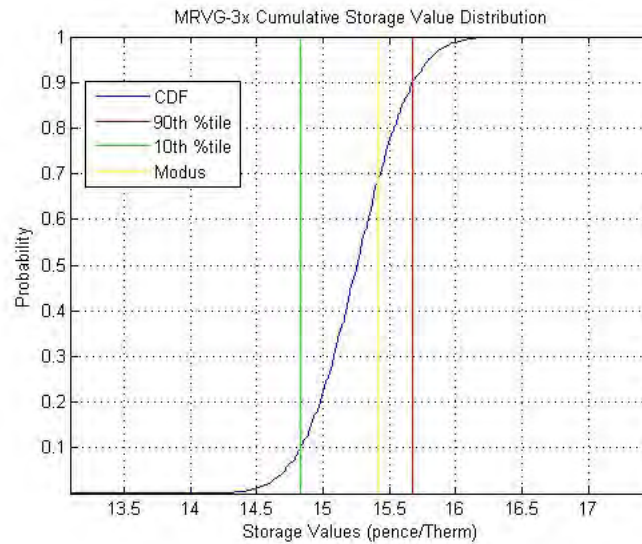


Figure 4.2.8: MRVG-3x Storage Value Cumulative Distribution

The cumulative probability distribution of the joint parameter estimation and calibration risk induced storage value density under the MRVG-3x model.

Example 4: Model Comparison

We now move on to addressing the issue of authorizing the use of one model over the other, which would typically fall on a model validation team within the risk control function. The approach we outline is informal in the sense that it focuses only on the relative levels of parameter risk and should be viewed as a compliment to a more rigorous model specification analysis. To frame the example, let us assume the the MRVG model is the incumbent and the trading desk have requested the use of the MRVG-3x model for valuation and p&l reporting. Once a rigorous theoretical examination of the assumptions and inputs underpinning the model has been carried out there is then an obvious need for a quantitative methodology capable of appropriately ranking the models with respect to their parameter risk. Comparing the valuations returned by both models in the context of the model risk they carry and for the products most relevant to the company is the main goal of such an analysis. Here, we will utilize the Average Value at Risk induced risk functional $R^*AVaR_\alpha(X)$, given by Bannör and Scherer (2013b), in ranking the storage parameter risk under both models:

Average Value at Risk (AVaR) is a coherent risk measure, defined for a given percentile α as

$$\begin{aligned} AVaR_\alpha(X) &= \frac{1}{\alpha} \inf_{z \in \mathbb{R}} (E_{\mathbb{Q}}[z - X]^+ - \alpha z) \\ &= \frac{1}{\alpha} \int_0^\alpha VaR_b(X) db \end{aligned}$$

where $VaR_b(X)$ is the $b\%$ value at risk for a position X . The AVaR at 100% therefore corresponds to the expected value of a position under a given probability measure. We refer interested readers to Föllmer and Knispel (2011) for a comprehensive overview of the AVaR risk measure. Bannör and Scherer (2013b) utilize AVaR as the underlying coherent risk measure in the risk functional, $R * AVaR_\alpha(X)$, defined as

$$R * AVaR_\alpha(X) = AVaR_\alpha(\mathbb{Q} \rightarrow E_{\mathbb{Q}}[X])$$

where R is the distribution of the measures \mathbb{Q} .

In the case where one model is being proposed as a potential replacement for an incumbent and they rank equally in all other potential model comparison metrics, a naive approach would be to choose the model which returns the highest expected value on a given asset, ignoring the model risk implicit in the valuations. A conservative parameter risk based rule for switching would be to require the new model to yield a bid price, $-R * AVaR_\alpha(-X)$, at a chosen percentile to be greater than the offer price, $R * AVaR_\alpha(X)$, of the incumbent model. Not only does this guarantee that the model risk adjusted price levels in the new model exceed those of the incumbent, it also returns a minimum cash value associated with switching i.e. the bid minus the offer. Figure 4.2.9 and Figure 4.2.10 display the risk adjusted bid and offer values, associated with the MRVG and MRVG-3x models respectively, at varying percentile levels. In this case, the MRVG-3x bid values dominate the MRVG offers at all percentiles. This would imply that the MRVG-3x model is an acceptable replacement for the MRVG.

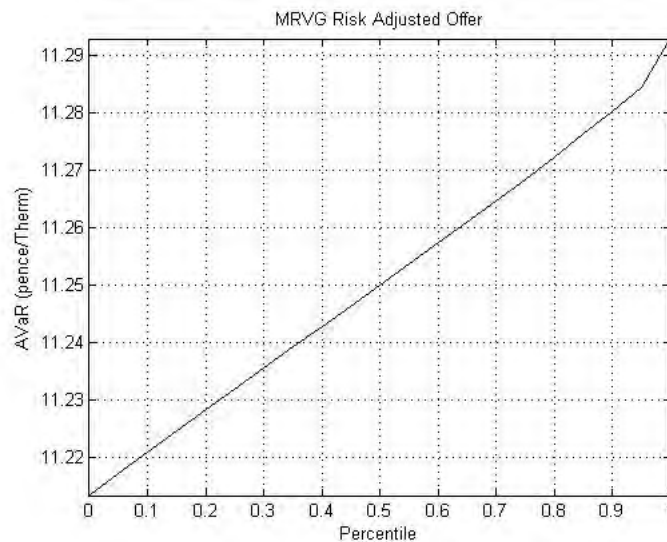


Figure 4.2.9: MRVG Storage Risk Adjusted Offers

The risk adjusted storage value offer level under the MRVG model using the Average Value at Risk model risk functional for a range of percentile levels.

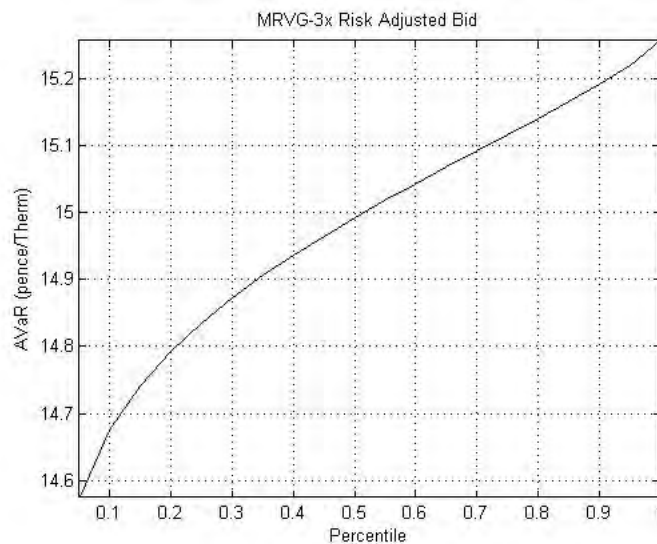


Figure 4.2.10: MRVG-3x Storage Risk Adjusted Bids

The risk adjusted storage value bid level under the MRVG-3x model using the Average Value at Risk model risk functional for a range of percentile levels.

In conclusion, it is clear from the above examples that a range of detailed analyses can be carried out under our proposed joint calibration and parameter estimation risk framework. Further, as demonstrated above, the chief benefit of this approach is the ability to compare models with varying reliance on market and historical information in a consistent manner. While there is always a degree of subjectivity in drawing conclusions from analyses of this type, for example in the percentile levels chosen for bids and offers, the framework presented would allow practitioners to make more informed decisions when evaluating the risk associated with asset valuations owing to both model choice and estimation. In summary, the main contributions to the literature have been

- A detailed analysis of the calibration risk inherent in real asset valuations extending previous research in this area which focused solely on the historical parameter estimation risk.
- Building on previous work by Bannör and Scherer (2013b), we presented a model risk methodology which incorporates both market calibrated and historically estimated model parameters thus allowing for a unified estimation of the model risk associated with pricing models which rely on both sources of information for parameter estimation.

Chapter 5

Conclusion

The primary goal of this thesis has been to derive a derivative valuation and risk framework which is capable of consistently pricing the risk exposure of storage contracts and vanilla monthly swaptions on natural gas. To that end, we aimed to derive a family of pricing models capable of reflecting both the statistical properties of the forward curve dynamics whilst also allowing for accurate calibration to the options market. Further, we looked to develop a range of valuation, calibration and model analysis tools applicable to general natural gas derivative structures and as such forming a basis for a potential derivative pricing and risk management library. Lastly, to complement the model development work, we wished to provide a framework for analyzing and quantifying the model risk associated with general energy derivative models. We will now discuss in detail the results of our work, some limitations that exist, and the potential for further research in this area.

We began by introducing the general Lévy driven forward curve consistent spot price model to be utilized in the context of single factor storage valuations. We presented the general conditional characteristic function for this model which would allow us to utilize the Fourier based pricing methods developed later. This general form allowed us to separately specify the stochastic drivers underlying the spot price model. We demonstrated how one such model commonly referenced in the literature, namely, the Mean Reverting Jump Diffusion model, could easily be incorporated within our general specification. We went on to present one of the major contributions of this thesis, the derivation of the conditional characteristic function of the Mean Reverting Variance Gamma process. We derived the first four moments of the process and demonstrated, with reference to the related Mean Reverting Diffusion model, its flexibility in determining the level of model kurtosis and hence the attenuation of the resulting implied volatility smile.

In order to calibrate these models to market swaption contracts we extended previous work in pricing swaptions under diffusion processes to general Lévy driven single factor forward curve models. To test whether the Mean Reverting Variance Gamma (MRVG) process was capable of meeting our modelling requirements we calibrated it, along with the Mean Reverting Diffusion (MRD) and Mean Reverting

Jump Diffusion (MRJD) models, to an option surface on NBP monthly forward contracts. The results showed a significant improvement in replicating the implied volatility smile over the more commonly used diffusion model. The MRVG and MRJD models returned almost identical “fits” to the implied volatility surface. Given the greater parsimony of the MRVG model over the MRJD, which in practical terms means having less parameters to hedge, we concluded that this gives sufficient justification for its usage as the main stochastic driver within our modelling framework.

There remains a significant potential for further research in relation to calibrating this model to market. As discussed within the body of the thesis, to aid in the mathematical derivations we assumed time-independent parameters. A straightforward extension would be to allow the parameters of the Lévy driving process to be piece-wise flat functions of time. In practical terms, this would primarily impact the evaluation of the characteristic function for a given time to maturity. However, this can be trivially solved by breaking any integrals of the process over time into the non-overlapping regions over which the parameters are constant. This would allow for far more flexibility in matching the term structure of implied volatility. Further, it may free the mean reversion rate parameter to be estimated by reference to the historical structure of relative maturity forward curve returns in a manner similar to the multifactor framework. Of even more interest, is the possibility of estimating this parameter from the prices of more liquid storage products, effectively linking the parameter to the level of extrinsic value observable in the storage market. Another interesting area of potential future research is in obtaining an exact rather than approximate fit to the options market. To that end, we have noted with interest the recent work of Carr and Nadtochiy (2014) on the Local Variance Gamma model, the results of which may, in theory, be applied to the MRVG model.

We finished the theoretical component of the single factor modelling framework with the derivation of the Fast Fourier Transform (FFT) based storage valuation algorithm, following the approach of Lord et al. (2007). We then presented an extension to this method, (FFT-m), which removes the need for a dampening parameter and which in theory, as discussed in the algorithm error analysis given in the Appendix, could lead to an increased rate of valuation convergence. We tested both algorithms in relation to an hypothetical storage deal by first benchmarking to the “forest of trees” approach using the MRD model. We then presented convergence results for both methods under each of the previously calibrated models. As expected, the FFT-m version of the algorithm displayed a greater rate of convergence for each of the price models considered. In terms of future research, there remains much room for improvement in the computational performance of the algorithm, in particular, the interpolation of the value function at each time-step. Further, the need for an efficient method for the calculation of model Greeks under this valuation framework remains a potentially rich area for future research.

In Chapter 3, we extended the single factor framework to a multifactor setting, in keeping with recent developments within the storage valuation literature. We presented a family of multifactor models, belonging to the Cheyette model class, which were forward curve consistent, options market consistent, and

reflective of the statistical dynamics of the NBP gas forward curve returns. For each model, we derived the factor implied forward curve which was then used to construct a general multidimensional swaption algorithm for Markovian multifactor Lévy processes. Further, to aid in our analysis of these models, we derived the distribution moments of the stochastic factors and also the covariance generating function related to each model specification. Due to concerns surrounding the efficiency of calibrating directly to market option prices, we extended previous work utilizing market implied moments to incorporate general instantaneous forward curve models. This innovative approach allows one to calibrate to options on forwards of varying contract delivery periods, further, the computational effort is broadly independent of the number of factors in the underlying model.

In order to price storage using this family of models, we next proceeded to the derivation of the multi-dimensional extension of the single factor valuation algorithm, which as noted previously, could be used for general multi-asset valuations in addition to the multifactor storage algorithm presented. For example, one could price a path dependent option on multiple underlyings, such as an American Basket Option or an Oil Indexed Take-or-Pay contract. The main difficulty here would be in the non-trivial derivation of the joint characteristic function for the underlying stochastic processes, and, in a non Gaussian setting, the difficulties in modelling the co-dependency between the underlyings. We finished the chapter with the presentation and analysis of a storage valuation under two of our calibrated multifactor models. We demonstrated how one can utilize the factor implied forward curve methods derived previously to both analyze and rationalize the dependence of the resulting storage value on the model specification and calibrated parameter values. One issue that became apparent when valuing the storage deal was the computational burden of the algorithm, therefore we see the need for optimization of the algorithm as a key area for future research and a necessary pre-requisite to this method being deemed acceptable by practitioners.

As mentioned previously, the requirement that the forward curve model remain Markovian in the underlying factors places a constraint on the structure of the factor volatility functions. It is our belief that, from a practitioners standpoint, not having the flexibility to freely determine the nature of the forward curve volatility function, $\sigma(t, T)$, could be a potential deficiency of this modelling approach. To address this, we feel that further research into approximating the dynamics of a non-Markovian forward curve model within a low dimensional state space is required. Of course, the use of a general form for the forward curve volatility would induce a path dependent drift in the implied spot price model. Therefore, we would envisage future work in this area placing greater emphasis on accounting for this path dependency rather than the current approach where it is removed through approximation of the forward curve dynamics.

In Chapter 4 we focused on the problem of analyzing and quantifying the level of model risk inherent in the modelling framework presented. We began with a thorough review of the concepts of model uncertainty and parameter risk, and the methodologies used to incorporate these risks within financial and

physical asset valuations. We identified the lack of a consistent method for unifying both calibration and parameter estimation risk within a single risk metric as a key shortcoming within the model risk literature. Given the dependency of our modelling framework on both market based and historical information, we built upon previous work in this area in deriving a general parameter risk methodology which combines calibration and parameter estimation risk into a single metric.

Focusing on the specific case of energy forward curve models, we demonstrated how dimensionally reduced Markovian models could be incorporated within this framework. We ended with the presentation of a number of real world case studies which utilized the models derived in the earlier chapters. This served the purpose of demonstrating the flexibility of our proposed approach and also the breadth of detailed and actionable analysis which can be carried out. One aspect of the empirical examples presented that we feel would benefit from further research is the addition of other structured derivative contracts. Products which are exposed to both outright and spread volatility, such as take-or-pay swing, would provide an informative additional insight into the parameter risk of the model specifications presented.

In conclusion, we believe that the market consistent modelling approach presented in this thesis, whilst not without limitations, does provide a robust framework for effectively managing the volatility risk, and accurately quantifying the model risk, inherent in a natural gas derivatives business.

Bibliography

- Andersen, L., 2010. Markov models for commodity futures: theory and practice. *Quantitative Finance* 10 (8), 831–854.
- Andricopoulos, A., Widdicks, M., Duck, P., Newton, D., 2003. Universal option valuation using quadrature methods. *Journal of Financial Economics* 67 (3), 447–471.
- Artzner, P., Delbaen, F., Eber, J., Heath, D., 1999. Coherent measures of risk. *Mathematical finance* 9 (3), 203–228.
- Avellaneda, M., Friedman, C., Holmes, R., Samperi, D., 1997. Calibrating volatility surfaces via relative-entropy minimization. *Applied Mathematical Finance* 4 (1), 37–64.
- Bakshi, G., Kapadia, N., Madan, D., 2003. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies* 16 (1), 101–143.
- Bakshi, G., Madan, D., 2000. Spanning and derivative-security valuation. *Journal of Financial Economics* 55 (2), 205–238.
- Bannor, K., Kiesel, R., Nazarova, A., Scherer, M., 2013a. Model risk and power plant valuation. Available at SSRN 2259615.
- Bannör, K., Scherer, M., 2013b. Capturing parameter risk with convex risk measures. *European Actuarial Journal* 3 (1), 97–132.
- Barndorff-Nielsen, O., Benth, F., Veraart, A., 2010. Modelling energy spot prices by lévy semistationary processes. *CREATES Research Paper* 18.
- Barndorff-Nielsen, O., Mikosch, T., Resnick, S., 2012. *Lévy processes: theory and applications*. Springer Science & Business Media.
- Barndorff-Nielsen, O., Shephard, N., 2001. Non-gaussian ornstein–uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63 (2), 167–241.

- Benth, F., Kallsen, J., Meyer-Brandis, T., 2007. A non-gaussian ornstein–uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance* 14 (2), 153–169.
- Benth, F. E., Di Nunno, G., Khedher, A., Schmeck, M. D., 2015. Pricing of spread options on a bivariate jump market and stability to model risk. *Applied Mathematical Finance* 22 (1), 28–62.
- Beyna, I., Wystup, U., 2011. Characteristic functions in the cheyette interest rate model. Tech. rep., CPQF Working Paper Series.
- Bjerksund, P., Stensland, G., Vagstad, F., 2008. Gas storage valuation: Price modelling v. optimization methods. *Optimization Methods* (October 17, 2008). NHH Dept. of Finance & Management Science Discussion Paper (2008/20).
- Black, F., 1976. The pricing of commodity contracts. *Journal of financial economics* 3 (1), 167–179.
- Boogert, A., De Jong, C., 2008. Gas storage valuation using a monte carlo method. *The Journal of Derivatives* 15 (3), 81–98.
- Boogert, A., De Jong, C., 2011. Gas storage valuation using a multifactor price process. *The Journal of Energy Markets* 4 (4), 29–52.
- Breslin, J., Clewlow, L., Elbert, T., Kwok, C., Strickland, C., 2008. Gas storage: Overview and static valuation. Tech. rep.
- Carmona, R., Coulon, M., 2014. A survey of commodity markets and structural models for electricity prices. In: *Quantitative Energy Finance*. Springer, pp. 41–83.
- Carmona, R., Ludkovski, M., 2010. Valuation of energy storage: An optimal switching approach. *Quantitative Finance* 10 (4), 359–374.
- Carr, P., Madan, D., 1999. Option valuation using the fast fourier transform. *Journal of Computational Finance* 2 (4), 61–73.
- Carr, P., Madan, D., 2009. Saddlepoint methods for option pricing. *Journal of Computational Finance* 13 (1), 49.
- Carr, P., Nadtochiy, S., 2014. Local variance gamma and explicit calibration to option prices. *Mathematical Finance*.
- Černý, A., Kyriakou, I., 2011. An improved convolution algorithm for discretely sampled asian options. *Quantitative Finance* 11 (3), 381–389.
- Chen, Z., Forsyth, P., 2009. A semi-lagrangian approach for natural gas storage valuation and optimal operation. *SIAM Journal on Scientific Computing* 30 (1).

- Cheyette, O., 2001. Markov representation of the heath-jarrow-morton model. Available at SSRN 6073.
- Chourdakis, K., 2004. Option pricing using the fractional fft. *Journal of Computational Finance* 8 (2), 1–18.
- Clelow, L., Strickland, C., 1999a. A multi-factor model for energy derivatives. Tech. rep.
- Clelow, L., Strickland, C., 1999b. Valuing energy options in a one factor model fitted to forward prices. Quantitative Finance Research Group, Working paper, University of Technology, Sydney, Australia.
- Cont, R., 2006. Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical finance* 16 (3), 519–547.
- Cont, R., Ben Hamida, S., 2005. Recovering volatility from option prices by evolutionary optimization. *Journal of Computational Finance* 8 (4).
- Cont, R., Tankov, P., 2003. Calibration of jump-diffusion option pricing models: a robust non-parametric approach. *Journal of Computational Finance* 7, 1–49.
- Cont, R., Tankov, P., 2004. *Financial modelling with jump processes*. Vol. 2. Chapman & Hall.
- Crépey, S., 2003. Calibration of the local volatility in a trinomial tree using tikhonov regularization. *Inverse Problems* 19 (1), 91–127.
- De Jong, F., 2000. Time series and cross-section information in affine term-structure models. *Journal of Business & Economic Statistics* 18 (3), 300–314.
- Dempster, M. A. H., Hong, G., 2002. Spread option valuation and the fast fourier transform. In: *Mathematical Finance Bachelier Congress 2000*. pp. 203–220.
- Dempster, M. A. H., Tang, K., 2011. Estimating exponential affine models with correlated measurement errors: Applications to fixed income and commodities. *Journal of Banking & Finance* 35 (3), 639–652.
- Deng, S., 2000. *Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes*. Citeseer.
- Deryabin, M., 2012. On bounds for model calibration uncertainty. *Journal of Risk Model Validation*, Forthcoming.
- Ding, D., 2010. An accurate and stable fft-based method for pricing options under exp-lévy processes.
- Driessen, J., Klaassen, P., Melenberg, B., 2003. The performance of multi-factor term structure models for pricing and hedging caps and swaptions. *Journal of Financial and Quantitative Analysis* 38 (03), 635–672.

- Driscoll, T., Fornberg, B., 2001. A padé-based algorithm for overcoming the gibbs phenomenon. *Numerical Algorithms* 26 (1), 77–92.
- Duan, J., Simonato, J., 1999. Estimating and testing exponential-affine term structure models by kalman filter. *Review of Quantitative Finance and Accounting* 13 (2), 111–135.
- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68 (6), 1343–1376.
- Eberlein, E., Kluge, W., 2007. Calibration of lévy term structure models. In: *Advances in mathematical finance*. Springer, pp. 147–172.
- Eberlein, E., Raible, S., 1999. Term structure models driven by general lévy processes. *Mathematical Finance* 9 (1), 31–53.
- EFET, 2009. Gas storage in europe - adding security through flexibility. Tech. rep.
- Ellsberg, D., 1961. Risk, ambiguity, and the savage axioms. *The quarterly journal of economics*, 643–669.
- Epps, T., 1993. Characteristic functions and their empirical counterparts: geometrical interpretations and applications to statistical inference. *American Statistician*, 33–38.
- Epstein, L., 1999. A definition of uncertainty aversion. *The Review of Economic Studies* 66 (3), 579–608.
- Fang, F., Oosterlee, C., 2008. A novel pricing method for european options based on fourier-cosine series expansions. *SIAM Journal on Scientific Computing* 31 (2), 826–848.
- Fang, F., Oosterlee, C., 2009. Pricing early-exercise and discrete barrier options by fourier-cosine series expansions. *Numerische Mathematik* 114 (1), 27–62.
- Felix, B., Weber, C., 2012. Gas storage valuation applying numerically constructed recombining trees. *European Journal of Operational Research* 216 (1), 178–187.
- Figlewski, S., 1998. *Derivatives risks, old and new*.
- Filipović, D., Tappe, S., 2008. Existence of lévy term structure models. *Finance and Stochastics* 12 (1), 83–115.
- Flannery, B., Press, W., Teukolsky, S., Vetterling, W., 1992. *Numerical recipes in c*. Press Syndicate of the University of Cambridge, New York.
- Föllmer, H., Knispel, T., 2011. Entropic risk measures: Coherence vs. convexity, model ambiguity and robust large deviations. *Stochastics and Dynamics* 11 (02n03), 333–351.

- Föllmer, H., Schied, A., 2002. Convex measures of risk and trading constraints. *Finance and stochastics* 6 (4), 429–447.
- Gapeev, P., Küchler, U., 2006. On markovian short rates in term structure models driven by jump-diffusion processes. *Statistics & Decisions* 24 (2), 255–271.
- Gibson, R., Lhabitant, F., Pistre, N., Talay, D., 1998. Interest rate model risk: an overview. *Journal of Risk* 1, 37–62.
- Guillaume, F., Schoutens, W., 2013. A moment matching market implied calibration. *Quantitative Finance* 13 (9), 1359–1373.
- Gupta, A., Reisinger, C., Whitley, A., 2010. Model uncertainty and its impact on derivative pricing. *Rethinking Risk Management and Reporting: Uncertainty, Bayesian Analysis and Expert Judgement*. Risk Books.
- Hagan, P., Kumar, D., Lesniewski, A., Woodward, D., 2002. Managing smile risk. *The Best of Wilmott*, 249.
- Haslip, G., Kaishev, V., 2014. Lookback option pricing using the fourier transform b-spline method. *Quantitative Finance* 14 (5), 789–803.
- Heath, D., Jarrow, R., 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*, 77–105.
- Heather, P., 2012. Continental european gas hubs: are they fit for purpose. *Oxford Institute for Energy Studies*. Working paper 63.
- Henaff, P., Laachir, I., Russo, F., 2013. Gas storage valuation and hedging. a quantification of the model risk. *arXiv preprint arXiv:1312.3789*.
- Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of financial studies* 6 (2), 327–343.
- Ho, T., Lee, S., 1986. Term structure movements and pricing interest rate contingent claims. *Journal of Finance*, 1011–1029.
- Holland, A., Walsh, C., 2013. An equilibrium analysis of third-party access to natural gas storage. *The Journal of Energy Markets* 6 (2), 3.
- Homescu, C., 2011. Implied volatility surface: Construction methodologies and characteristics. *arXiv preprint arXiv:1107.1834*.

- Hull, J., Suo, W., 2002. A methodology for assessing model risk and its application to the implied volatility function model. *Journal of Financial and Quantitative Analysis* 37 (02), 297–318.
- Hureau, G., 2015. Gas storage in europe, recent developments and outlook to 2035. European Gas Conference, 27-29 January 2015, Vienna.
- Jackson, K., Jaimungal, S., Surkov, V., 2007. Fourier space time-stepping for option pricing with lévy models. *Journal of Computational Finance*, Vol. 12, No. 2, pp. 1-29, 2008.
- Jaimungal, S., Surkov, V., 2011. Levy based cross-commodity models and derivative valuation. To Appear: *SIAM Journal of Financial Mathematics*.
- Joshi, M., Yang, C., 2011. Fourier transforms, option pricing and controls. *Option Pricing and Controls* (October 9, 2011).
- Kahl, C., Lord, R., 2010. Fourier inversion methods in finance. Tech. rep., Working Paper, Commerzbank, and Cardano.
- Kerkhof, J., Melenberg, B., Schumacher, H., 2010. Model risk and capital reserves. *Journal of Banking & Finance* 34 (1), 267–279.
- Kerkhof, J., Schumacher, J., Melenberg, B., 2002. Model risk and regulatory capital. Available at SSRN 301531.
- Kjaer, M., 2008. Pricing of swing options in a mean reverting model with jumps. *Applied mathematical finance* 15 (5-6), 479–502.
- Knight, F., 1921. *Risk, uncertainty and profit*. Boston: Houghton Mifflin.
- Lagarias, J., Reeds, J., Wright, M., Wright, P., 1998. Convergence properties of the nelder–mead simplex method in low dimensions. *SIAM Journal on optimization* 9 (1), 112–147.
- Lai, G., Margot, F., Secomandi, N., 2010. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations research* 58 (3), 564–582.
- Le Fevre, C., 2013. *Gas storage in Great Britain*. Oxford Institute for Energy Studies.
- Lee, R., 2004. The moment formula for implied volatility at extreme strikes. *Mathematical Finance* 14 (3), 469–480.
- Leentvaar, C., Oosterlee, C., 2007. Multi-asset option pricing using a parallel fourier-based technique. Tech. rep., Delft University of Technology, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft Institute of Applied Mathematics.

- Lewis, A., 2001. A simple option formula for general jump-diffusion and other exponential lévy processes. Envision Financial Systems and OptionCity. net.
- Li, L., Linetsky, V., 2014. Time-changed ornstein–uhlenbeck processes and their applications in commodity derivative models. *Mathematical Finance* 24 (2), 289–330.
- Litterman, R., Scheinkman, J., 1991. Common factors affecting bond returns. *The Journal of Fixed Income* 1 (1), 54–61.
- Lord, R., Fang, F., Bervoets, F., Oosterlee, C., 2007. A fast and accurate fft-based method for pricing early-exercise options under lévy processes. Center for Mathematics and Computer Science (CWI).
- Lord, R., Kahl, C., 2007. Optimal fourier inversion in semi-analytical option pricing.
- Luciano, E., 2009. Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy. Vol. 445. Wiley.
- Luciano, E., Semeraro, P., 2010. Multivariate variance gamma and gaussian dependence: a study with copulas. In: *Mathematical and Statistical Methods for Actuarial Sciences and Finance*. Springer, pp. 193–203.
- Madan, D., Carr, P., Chang, E., 1998. The variance gamma process and option pricing. *European Finance Review* 2 (1), 79–105.
- Manoliu, M., 2004. Storage options valuation using multilevel trees and calendar spreads. *International Journal of Theoretical and Applied Finance* 7 (04), 425–464.
- Marfè, R., 2009. A multivariate two% factor variance gamma process for asset returns. Tech. rep., Working Paper, Swiss Finance Institute.
- Maslyuka, S., Rotarub, K., Dokumentovc, A., 2013. Price discontinuities in energy spot and futures prices. Tech. rep., Monash University, Department of Economics.
- Morini, M., 2011. *Understanding and Managing Model Risk: A practical guide for quants, traders and validators*. John Wiley & Sons.
- Nomikos, N., Andriosopoulos, K., 2012. Modelling energy spot prices: Empirical evidence from nymex. *Energy Economics* 34 (4), 1153–1169.
- O’Sullivan, C., 2005. Path dependant option pricing under lévy processes. In: *EFA 2005 Moscow Meetings Paper*.
- Pan, C., 1993. Gibbs phenomenon suppression and optimal windowing for attenuation and q measurements. Tech. rep.

- Parsons, C., 2013. Quantifying natural gas storage optionality: a two-factor tree model. *The Journal of Energy Markets* 6 (1).
- Pinho, C., Madaleno, M., 2011. Using exponential lévy models to study implied volatility patterns for electricity options.
- Ruijter, M., Oosterlee, C., 2012. Two-dimensional fourier cosine series expansion method for pricing financial options. *SIAM Journal on Scientific Computing* 34 (5), B642–B671.
- Sato, K., 2001. Basic results on lévy processes. In: *Lévy processes*. Springer, pp. 3–37.
- Schwartz, E., 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance* 52 (3), 923–973.
- Serletis, A., 1992. Maturity effects in energy futures. *Energy Economics* 14 (2), 150–157.
- Thompson, M., Davison, M., Rasmussen, H., 2009. Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics (NRL)* 56 (3), 226–238.
- Tichy, T., 2006. Modeling the electricity spot price at the czech and austrian markets by extended vg model. *Journal of Information, Control and Management Systems* 4 (2).
- Vasicek, O., 1977. An equilibrium characterization of the term structure. *Journal of financial economics* 5 (2), 177–188.
- Warin, X., 2012. Gas storage hedging. In: *Numerical Methods in Finance*. Springer, pp. 421–445.

Appendix

For ease of exposition, many of the technical details around the modelling innovations and contributions to the existing storage literature presented in the body of the thesis are given in the following Appendix. To aid the reader, we will now outline the main contributions detailed in the following sections.

- Section A.2 presents the derivation of the characteristic function for the Mean Reverting Variance Gamma (MRVG) process. This work contributes to the literature another example of a mean reverting Lévy process with both positive and negative jumps, in addition to the model of Deng (2000), where the characteristic function is available in closed form.
- Section A.3 presents the derivation of the first four central moments of the MRVG process, which were used in Chapter 2 to provide insight into the properties of this model.
- In Section A.5 we adapt previous work in the context of interest rates to provide a detailed derivation of the conditions under which a single factor forward curve model has a Markovian spot price representation.
- Section A.5 presents details of a swaption pricing algorithm which extends previous work by Clewlow and Strickland (1999b) to single factor Lévy driven forward curve models.
- In Section B.1 we derive the time-dependent drift of the implied spot price dynamics for each of the forward curve model specifications presented in Chapter 3.
- In Sections B.2 and B.3 we derive the characteristic functions of our family of multifactor models.
- Sections B.5 gives details of the factor implied forward price for each model, which is used in Section B.7 in deriving a general multifactor swaption pricing algorithm.
- In Section B.8 we present detailed derivations of the first four moments of generic forward contracts under each model specification. This work is a crucial pre-requisite to the use of the implied moments calibration technique detailed in Chapter 3.
- Finally, in Sections B.4 and B.6 we present the derivation of a number of utility methods which aid in analyzing the behaviour of the different models specified in Chapter 3. Section B.4 gives the

first four central moments of the log spot price under each model. In Section B.6 the instantaneous covariance function of the each model is presented, these results are a useful tool in analyzing the impact of the model parameters on the model covariance structure, which is of course a key driver of storage value.

Appendix A

Chapter 2

A.1 Parseval's Theorem

The statement of the theorem is as follows, given two complex functions $f, g \in \mathcal{L}^2$

$$\int_{\mathbb{R}} g(x) \overline{f(x)} dx = \frac{1}{2\pi} \int_{\mathbb{R}} \hat{g}(x) \overline{\hat{f}(u)} du$$

where \overline{f} denotes complex conjugation.

We will now review the contribution of Lewis (2001) who utilizes the above in deriving transform based pricing functions for general payoffs. Firstly, we make the usual assumptions of frictionless markets, the absence of arbitrage, and the existence of an equivalent market measure \mathbb{Q} , under which all expectations are taken unless otherwise stated. Then given a stochastic process X_t defined on the probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t, t \geq 0\}, \mathbb{Q})$ we define the price of some contingent claim $V(X_t)$ as

$$\begin{aligned} V(X_t) &= E[g(X) | \mathcal{F}_t] \\ &= \int_{-\infty}^{\infty} g(x) f(x) dx \end{aligned}$$

where $g(x)$ is some arbitrary payoff function.

Lewis (2001) applies a generalized version of Parseval's Theorem to the function $V(X_t)$ such that

$$\begin{aligned} V(X_t) &= \int_{-\infty}^{\infty} g(x) f(x) dx \\ &= \frac{1}{2\pi} \int_{iw-\infty}^{iw+\infty} \hat{g}(z) \hat{f}(-z) dz \end{aligned} \tag{A.1.1}$$

where the function $\hat{g}(z)$ is the general Fourier Transform of $g(x) : \mathbb{R} \rightarrow \mathbb{R}$, and $\hat{f}(z)$ is the general Fourier Transform of $f(x) : \mathbb{R} \rightarrow \mathbb{R}$ for $z = u + iw \in \mathbb{C}$. The choice of the constant w is determined such that the functions $\hat{g}(z)$ and $\hat{f}(-z)$ are finite for all $u \in \mathbb{R}$.

Proof:

Let

$$\begin{aligned} \int_{-\infty}^{\infty} h(x) \overline{k(x)} dx &= \int_{-\infty}^{\infty} g(x) \exp(wx) \overline{f(x) \exp(-wx)} dx \\ &= \int_{-\infty}^{\infty} g(x) \overline{f(x)} dx \end{aligned}$$

Then by Parseval's Theorem we have,

$$\begin{aligned} \int_{-\infty}^{\infty} h(x) \overline{k(x)} dx &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{h}(u) \overline{\hat{k}(u)} du \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{h}(u) \hat{k}(-u) du \end{aligned}$$

as f is real. Evaluating $\hat{h}(u)$ and $\hat{k}(-u)$ gives

$$\begin{aligned} \hat{h}(u) &= \int_{-\infty}^{\infty} g(x) \exp(wx) \exp(-iux) dx \\ &= \int_{-\infty}^{\infty} g(x) \exp(-i(u + iw)x) dx \\ &= \hat{g}(z) \end{aligned}$$

$$\begin{aligned} \hat{k}(-u) &= \int_{-\infty}^{\infty} f(x) \exp(-wx) \exp(iux) dx \\ &= \int_{-\infty}^{\infty} f(x) \exp(-i(-u - iw)x) dx \\ &= \hat{f}(-z) \end{aligned}$$

A.2 Derivation of the Characteristic Function of the Mean Reverting Variance Gamma Process

We can use Equation 2.1.3 to solve for the characteristic function, in which case the only non-trivial derivation is that of the exponential

$$\exp\left(\int_s^t \varphi_{vg}(z \exp(-\alpha(t-c))) dc\right)$$

Recall that the characteristic exponent of the Variance-Gamma Process (with $\theta = 0$) is given by

$$\varphi_{vg}(z) = -\frac{1}{v} \ln\left(1 + \left(\frac{\sigma^2 v}{2}\right) z^2\right)$$

and thus,

$$\varphi_{vg}(z \exp(-\alpha(t-c))) = -\frac{1}{v} \ln\left(1 + \left(\frac{\sigma^2 v}{2}\right) z^2 \exp(-2\alpha(t-c))\right)$$

Letting $x = -\left(\frac{\sigma^2 v}{2}\right) z^2 \exp(-2\alpha(t-c))$, we can write the above integral as

$$-\frac{1}{2v\alpha} \int_{l_0}^{l_1} \frac{\ln(1-x)}{x} dx \tag{A.2.1}$$

where

$$\begin{aligned} l_1 &= -\left(\frac{\sigma^2 v}{2}\right) z^2 \\ l_0 &= -\left(\frac{\sigma^2 v}{2}\right) z^2 \exp(-2\alpha(t-s)) \end{aligned}$$

which by definition is,

$$\frac{1}{2v\alpha} [Li_2(l_1) - Li_2(l_0)]$$

where $Li_2(z)$ is the dilogarithm function.

A.3 MRVG Distribution Moments

We can calculate the moments of the MRVG process directly from the moment generating function, which can be derived in manner in similar to the derivation of the characteristic function above.

$$m(u) = E\left[\exp\left(\int_s^t u \exp(-\alpha(t-c)) dX(c)\right)\right] = \exp\left(\int_s^t k_{vg}(u \exp(-\alpha(t-c))) dc\right) \tag{A.3.1}$$

where $k_{vg}(u)$ is the log cumulant function of the variance gamma process, and is given by

$$k_{vg}(u) = -\frac{1}{v} \ln \left(1 - \frac{\sigma^2 v}{2} u^2 \right)$$

We can then derive expressions for the moments through repeated differentiation of the function $m(u)$. To keep the calculations simple we will use the following shorthand, $m(u) = \exp(A(u))$ with $m^{(n)}(u)$, $A^{(n)}(u)$ denoting the n^{th} derivative.

$$\begin{aligned} A(u) &= \int_s^t k_{vg}(u \exp(-\alpha(t-c))) dc \\ A(0) &= 0 \end{aligned}$$

$$\begin{aligned} A^{(1)}(u) &= \frac{1}{v} \int_s^t \frac{\sigma^2 v u \exp(-2\alpha(t-c))}{1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))} dc \\ A^{(1)}(0) &= 0 \end{aligned}$$

$$\begin{aligned} A^{(2)}(u) &= \frac{1}{v} \int_s^t \left[\frac{\sigma^2 v \exp(-2\alpha(t-c))}{1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))} + \frac{\sigma^4 v^2 u^2 \exp(-4\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^2} \right] dc \\ A^{(2)}(0) &= \frac{\sigma^2}{2\alpha} (1 - \exp(-2\alpha(t-s))) \end{aligned}$$

$$\begin{aligned} A^{(3)}(u) &= \frac{1}{v} \int_s^t \frac{\sigma^4 v^2 u \exp(-4\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^2} + \frac{2\sigma^4 v^2 u \exp(-4\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^2} \dots \\ &\quad + \frac{2\sigma^6 v^3 u^3 \exp(-6\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^3} dc \\ A^{(3)}(0) &= 0 \end{aligned}$$

$$\begin{aligned}
A^{(4)}(u) &= \frac{1}{v} \int_s^t \frac{\sigma^4 v^2 \exp(-4\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^2} + \frac{2\sigma^6 v^3 u^2 \exp(-6\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^3} \dots \\
&+ \frac{2\sigma^4 v^2 \exp(-4\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^2} + \frac{4\sigma^6 v^3 u^2 \exp(-6\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^3} \dots \\
&+ \frac{6\sigma^6 v^3 u^2 \exp(-6\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^3} + \frac{6\sigma^8 v^4 u^4 \exp(-8\alpha(t-c))}{\left(1 - \frac{\sigma^2 v}{2} u^2 \exp(-2\alpha(t-c))\right)^4} dc \\
A^{(4)}(0) &= \frac{1}{v} \int_s^t 3\sigma^4 v^2 \exp(-4\alpha(t-c)) dc \\
&= \frac{3\sigma^4 v}{4\alpha} (1 - \exp(-4\alpha(t-s)))
\end{aligned}$$

Now, the 1st moment is given by

$$\begin{aligned}
m^{(1)}(u) &= m(u)A^{(1)}(u) \\
m^{(1)}(0) &= 0
\end{aligned}$$

the 2nd moment by

$$\begin{aligned}
m^{(2)}(u) &= m^{(1)}(u)A^{(1)}(u) + m(u)A^{(2)}(u) \\
m^{(2)}(0) &= \frac{\sigma^2}{2\alpha} (1 - \exp(-2\alpha(t-s)))
\end{aligned}$$

the 3rd moment by

$$\begin{aligned}
m^{(3)}(u) &= m^{(2)}(u)A^{(1)}(u) + m^{(1)}(u)A^{(2)}(u) \\
&+ m^{(1)}(u)A^{(2)}(u) + m(u)A^{(3)}(u) \\
m^{(3)}(0) &= 0
\end{aligned}$$

and finally, the 4th by

$$\begin{aligned}
m^{(4)}(u) &= m^{(3)}(u)A^{(1)}(u) + m^{(2)}(u)A^{(2)}(u) \\
&+ 2m^{(2)}(u)A^{(2)}(u) + 2m^{(1)}(u)A^{(3)}(u) \\
&+ m^{(1)}(u)A^{(3)}(u) + m(u)A^{(4)}(u) \\
m^{(4)}(0) &= \frac{3\sigma^4}{4\alpha^2} (1 - \exp(-2\alpha(t-s)))^2 + \frac{3\sigma^4 v}{4\alpha} (1 - \exp(-4\alpha(t-s)))
\end{aligned}$$

A.4 Markov Implied Spot Price Model

The dynamics of the log forward price $f(t, T)$ are,

$$df(t, T) = \left(-\frac{1}{2}\beta(t, T)^2 - \kappa_j(\gamma(t, T)) \right) dt + \beta(t, T) dW(t) + \int_{\mathbb{R}/0} \gamma(t, T) j\bar{v}(dj, ds)$$

where for brevity we have used the compensated jump measure, $\bar{v} = v(dj, ds) - \mu(dj)$. Integrating yields,

$$f(t, T) = f(0, T) + \int_0^t \left(-\frac{1}{2}\beta(s, T)^2 - \kappa_j(\gamma(s, T)) \right) ds + \int_0^t \beta(s, T) dW(s) + \int_0^t \int_{\mathbb{R}/0} \gamma(s, T) j\bar{v}(dj, ds)$$

Now, the log spot price $x(t)$ is equal to $f(t, t)$, thus

$$x(t) = f(0, t) + \int_0^t \left(-\frac{1}{2}\beta(s, t)^2 - \kappa_j(\gamma(s, t)) \right) ds + \int_0^t \beta(s, t) dW(s) + \int_0^t \int_{\mathbb{R}/0} \gamma(s, t) j\bar{v}(dj, ds) \quad (\text{A.4.1})$$

Applying the Fundamental Theorem of Calculus we can expand each of the integrals in the above expression such that¹

$$\begin{aligned} x(t) &= f(0, t) - \int_0^t \left(\frac{1}{2}\beta(s, s)^2 + \kappa_j(\gamma(s, s)) \right) ds - \int_0^t \int_s^t \beta(s, u) \frac{\partial \beta(s, u)}{\partial u} duds \\ &\quad - \int_0^t \int_s^t \kappa_j'(\gamma(s, u)) \frac{\partial \gamma(s, u)}{\partial u} duds + \int_0^t \beta(s, s) dW(s) + \int_0^t \int_s^t \frac{\partial \beta(s, u)}{\partial u} dudW(s) \\ &\quad + \int_0^t \int_{\mathbb{R}/0} \gamma(s, s) j\bar{v}(dj, ds) + \int_0^t \int_s^t \int_{\mathbb{R}/0} \frac{\partial \gamma(s, u)}{\partial u} jdu\bar{v}(dj, ds) \end{aligned} \quad (\text{A.4.2})$$

Now under the assumption that $\beta(s, u)$ and $\gamma(s, u)$ are separable functions and noting that

$$f(0, t) = f(0, 0) + \int_0^t \frac{\partial f(0, s)}{\partial s} ds$$

we can apply Fubini's Theorem and rewrite Equation A.4.2 as

$$\begin{aligned} x(t) &= x(0) + \int_0^t \frac{\partial f(0, s)}{\partial s} ds - \frac{1}{2} \int_0^t \eta(s)^2 \zeta(s)^2 ds - \int_0^t \kappa_j(\theta(s) \zeta(s)) ds \\ &\quad - \int_0^t \zeta'(u) \zeta(u) \int_0^u \eta(s)^2 ds du - \int_0^t \zeta'(u) \int_0^u \theta(s) \kappa_j'(\theta(s) \zeta(u)) ds du \\ &\quad + \int_0^t \eta(s) \zeta(s) dW(s) + \int_0^t \zeta'(u) \int_0^u \eta(s) dW(s) du \\ &\quad + \int_0^t \int_{\mathbb{R}/0} \theta(s) \zeta(s) j\bar{v}(dj, ds) + \int_0^t \zeta'(u) \int_0^u \int_{\mathbb{R}/0} \theta(s) j\bar{v}(dj, ds) du \end{aligned}$$

¹Here, $f'()$ represents the derivative of $f()$.

which in differential form is

$$\begin{aligned}
dx(t) = & \left(\frac{\partial f(0,t)}{\partial t} - \frac{1}{2}\eta(t)^2\zeta(t)^2 - \kappa_j(\theta(t)\zeta(t)) - \zeta'(t) \int_0^t \eta(s)^2\zeta(s) ds \right. \\
& \left. - \zeta'(t) \int_0^t \theta(s) \kappa'_j(\theta(s)\zeta(s)) ds + \zeta'(t) \int_0^t \eta(s) dW(s) + \zeta'(t) \int_0^t \int_{\mathbb{R}/0} \theta(s) j\bar{\nu}(dj, ds) \right) dt \\
& + \eta(t)\zeta(t) dW(t) + \int_{\mathbb{R}/0} \theta(t)\zeta(t) j\bar{\nu}(dj, ds)
\end{aligned} \tag{A.4.3}$$

Now, noting from Equation A.4.1 that

$$\begin{aligned}
x(t) \frac{\zeta'(t)}{\zeta(t)} = & \frac{\zeta'(t)}{\zeta(t)} f(0,t) - \frac{\zeta'(t)}{2} \int_0^t \eta(s)^2\zeta(s) ds - \frac{\zeta'(t)}{\zeta(t)} \int_0^t \kappa_j(\theta(s)\zeta(s)) ds \\
& + \zeta'(t) \int_0^t \eta(s) dW(s) + \zeta'(t) \int_0^t \int_{\mathbb{R}/0} \theta(s) j\bar{\nu}(dj, ds)
\end{aligned}$$

and inserting the above into Equation A.4.3 gives,

$$\begin{aligned}
dx(t) = & \left(\frac{\partial f(0,t)}{\partial t} - \frac{1}{2}\eta(t)^2\zeta(t)^2 - \frac{\zeta'(t)}{2} \int_0^t \eta(s)^2\zeta(s) ds - \kappa_j(\theta(t)\zeta(t)) \right. \\
& \left. - \zeta'(t) \int_0^t \theta(s) \kappa'_j(\theta(s)\zeta(s)) ds + \left(x(t) - f(0,t) + \int_0^t \kappa_j(\theta(s)\zeta(s)) ds \right) \frac{\zeta'(t)}{\zeta(t)} \right) dt \\
& + \eta(t)\zeta(t) dW(t) + \int_{\mathbb{R}/0} \theta(t)\zeta(t) j\bar{\nu}(dj, ds)
\end{aligned} \tag{A.4.4}$$

and thus $x(t)$ is Markov.

Clearly an obvious choice for the functions $\beta(t, T)$ and $\gamma(t, T)$ are $\sigma \exp(-\alpha(T-t))$ and $\exp(-\alpha(T-t))$ respectively, in which case we can rewrite Equation A.4.4 as

$$\begin{aligned}
dx(t) = & \left(\frac{\partial f(0,t)}{\partial t} - \frac{1}{2}\sigma^2 + \frac{1}{4}\sigma^2(1 - \exp(-2\alpha t)) - \kappa_j(\exp(-\alpha t)) + \right. \\
& \left. \alpha f(0,t) - \alpha \int_0^t \kappa_j(\exp(-\alpha(t-s))) ds - \alpha x(t) \right) dt + \sigma dW(t) + \int_{\mathbb{R}/0} j\bar{\nu}(dj, ds)
\end{aligned}$$

or more concisely as,

$$dx(t) = (\omega(t) - \alpha x(t))dt + dy(t) \tag{A.4.5}$$

which is of course of the form given by Equation 2.1.1 with time-varying mean log price level,

$$\omega(t) = \frac{\partial f(0,t)}{\partial t} - \frac{1}{2}\sigma^2 + \frac{1}{4}\sigma^2(1 - \exp(-2\alpha t)) - \kappa_j(\exp(-\alpha t)) + \alpha f(0,t) - \alpha \int_0^t \kappa_j(\exp(-\alpha(t-s))) ds$$

Equation A.4.5 has the following solution

$$x(t) = x(0) \exp(-\alpha t) + \int_0^t \omega(s) \exp(-\alpha(t-s)) ds + \int_0^t \exp(-\alpha(t-s)) dy(t)$$

or more generally,

$$x(t) = x(s) \exp(-\alpha(t-s)) + \int_s^t \omega(c) \exp(-\alpha(t-c)) dc + \int_s^t \exp(-\alpha(t-c)) dy(c)$$

Note that we can partially evaluate the integrated drift term to give

$$\begin{aligned} x(t) &= x(s) \exp(-\alpha(t-s)) + f(0,t) - \exp(-\alpha(t-s)) f(0,s) - \frac{\sigma^2}{2\alpha} (1 - \exp(-\alpha(t-s))) \\ &\quad + \frac{\sigma^2}{4\alpha} [\exp(-\alpha t) (\exp(-\alpha t) - \exp(-\alpha s) + \exp(\alpha t) - \exp(\alpha s))] \\ &\quad - \int_s^t (\kappa_j(\exp(-\alpha c)) + \alpha \varphi_j(-i, 0, c)) \exp(-\alpha(t-c)) dc + \int_s^t \exp(-\alpha(t-c)) dy(c) \end{aligned}$$

²The conditional characteristic function is then given by equating Equation A.4.5 with Equation 2.1.3 which gives

$$\Phi_{x(t)}(z) = \exp\left(izx(s) \exp(-\alpha(t-s)) + iz \int_s^t \omega(c) \exp(-\alpha(t-c)) dc + \int_s^t \varphi(iz \exp(-\alpha(t-c)), c, t) dc\right)$$

A useful feature of the above parametrization of the functions $\beta(t, T)$ and $\gamma(t, T)$ is that it allows us to define the forward price $f(t, T)$ as a function of the spot price $x(t)$ as follows

$$\begin{aligned} f(t, T) &= f(0, T) + \exp(-\alpha(T-t)) (x(t) - f(0, t)) \\ &\quad - \frac{\sigma^2}{4\alpha} \exp(-\alpha T) ((\exp(2\alpha t) - 1) (\exp(-\alpha T) - \exp(-\alpha t))) \\ &\quad - \int_0^t \kappa_j(\exp(-\alpha(T-s))) ds + \int_0^t \kappa_j(\exp(-\alpha(t-s))) \exp(-\alpha(T-t)) ds \\ &= f(0, T) + \exp(-\alpha(T-t)) (x(t) - f(0, t)) \\ &\quad - \frac{\sigma^2}{4\alpha} \exp(-\alpha T) ((\exp(2\alpha t) - 1) (\exp(-\alpha T) - \exp(-\alpha t))) \\ &\quad - \int_0^t \kappa_j(\exp(-\alpha(T-s))) ds + \varphi(-i, 0, t) \exp(-\alpha(T-t)) \end{aligned} \tag{A.4.6}$$

This last result is used later to construct the swaption pricing algorithm.

² $\varphi(-i, s, t)$ is the characteristic exponent function of the increments of the mean reverting jump process over the time period $(t-s)$.

A.5 Swaption Pricing Algorithm

The payoff for a call option on a forward delivering between times T_1 and T_2 is given by

$$\begin{aligned} V(F, K, t, T_1, T_2) &= \max [F(t, T_1, T_2) - K, 0] \\ &= \max \left[\frac{1}{N} \sum_{j=0}^{N-1} F(t, T_1 + j\Delta t) - K, 0 \right] \end{aligned}$$

where

$F(t, T)$ is the forward price observed at time t and delivering at time T

Δt is the delivery frequency of the commodity underlying the forward contract

N is the number of deliveries of the underlying forward contract.

Using Equation A.4.6 we can write the payoff function as

$$V(x(t), K, t, T_1, T_2) = \max \left[\frac{1}{N} \sum_{j=0}^{N-1} \exp(x(t) \exp(-\alpha(T_1 + j\Delta t - t)) + g(t, T_1 + j\Delta t)) - \exp(k), 0 \right] \quad (\text{A.5.1})$$

where

$$\begin{aligned} g(t, T_1 + j\Delta t) &= f(0, T_1 + j\Delta t) - \exp(-\alpha(T_1 + j\Delta t - t)) f(0, t) \\ &\quad - \frac{\sigma^2}{4\alpha} \exp(-\alpha(T_1 + j\Delta t)) ((\exp(2\alpha t) - 1) (\exp(-\alpha(T_1 + j\Delta t)) - \exp(-\alpha t))) \\ &\quad - \int_0^t \kappa_j (\exp(-\alpha(T_1 + j\Delta t - s))) ds + \varphi(-i, 0, t) \exp(-\alpha(T_1 + j\Delta t - t)) \end{aligned}$$

Following the example of Lewis (2001), we will price the option using Parseval's Theorem which will require the General Fourier Transform of the payoff function,

$$\tilde{V}(z) = \int_{-\infty}^{\infty} \max \left[0, \frac{1}{N} \sum_{j=0}^{N-1} \exp(x (\exp(-\alpha(T_1 + j\Delta t - t)) - iz) + g(t, T_1 + j\Delta t)) - \exp(-izx + k) \right] dx \quad (\text{A.5.2})$$

As outlined in Clewlow and Strickland (1999b) because the payoff is monotonic in x we can write the above as

$$\tilde{V}(z) = \int_{x^*}^{\infty} \frac{1}{N} \sum_{j=0}^{N-1} \exp(x (\exp(-\alpha(T_1 + j\Delta t - t)) - iz) + g(t, T_1 + j\Delta t)) - \exp(-izx + k) dx$$

where x^* solves

$$\frac{1}{N} \sum_{j=0}^{N-1} \exp(x(t) \exp(-\alpha(T_1 + j\Delta t - t)) + g(t, T_1 + j\Delta t)) - \exp(k) = 0$$

Thus,

$$\tilde{V}(z) = \frac{1}{N} \sum_{j=0}^{N-1} \frac{\exp(x^* (\exp(-\alpha(T_1 + j\Delta t - t)) - iz) + g(t, T_1 + j\Delta t))}{iz - \exp(-\alpha(T_1 + j\Delta t - t))} - \frac{\exp(-izx^* + k)}{iz} \quad (\text{A.5.3})$$

for all $z = u + iw, u, w \in \mathbb{R}$ and $w < -\exp(-\alpha(T_1 + j\Delta t - t)) \forall j$. The call option price is then given by,

$$C(F, K, t, T_1, T_2) = \frac{1}{2\pi} \int_{iw-\infty}^{iw+\infty} \tilde{V}(z) \Phi(z; x(0), 0, t) dz$$

where $\Phi(z; x(0), 0, t)$ is the conditional characteristic function of the underlying spot price model.

A.6 Valuation Algorithm Error Analysis

The main sources of error from using the valuation methodology outlined in Chapter 2 stem from the truncation and discretization of the integrals in Equation 2.3.4, and Equation 2.3.6 and the interpolation of the value function in order to account for mean reversion. Beginning with the latter, in our implementation we have chosen to use a cubic spline as our interpolation method, thus we are assuming the following functional form for the value function between successive points in the discretized log gas price domain Flannery et al. (1992),

$$\begin{aligned} V(x) &= AV(x_n) + BV(x_{n+1}) + CV''(x_n) + DV''(x_{n+1}) \\ A &= \frac{x_{n+1} - x}{x_{n+1} - x_n} \\ B &= \frac{x - x_n}{x_{n+1} - x_n} \\ C &= \frac{1}{6} (A^3 - A) (x_{n+1} - x_n)^2 \\ D &= \frac{1}{6} (B^3 - B) (x_{n+1} - x_n)^2 \end{aligned}$$

for $x \in [x_n, x_{n+1}]$. The error in this case will be $\mathcal{O}(dx^4)$.

A thorough examination of the truncation and discretization error is provided by Lord et al. (2007), who determine error bounds by focusing on the continuation value,

$$C(x) = \int_{-\infty}^{\infty} V(x) f(x) dx$$

The first approximation is to replace $V()$ by its Fourier Series expansion on $[x_0, x_N]$, such that

$$\begin{aligned} \int_{-\infty}^{\infty} V(x) f(x) dx &\simeq \int_{-\infty}^{\infty} \sum_{u=-\infty}^{\infty} \hat{v}_k \exp(2\pi i u x / N) f(x) dx \\ &= \sum_{u=-\infty}^{\infty} \hat{v}_u \exp(2\pi i u x / N) \Phi\left(\frac{2\pi u}{N}\right) \end{aligned}$$

where \hat{v}_u are the coefficients of the Fourier Series expansion of V given by

$$\hat{v}_u = \frac{1}{N} \int_{-N/2}^{N/2} V(x) \exp(-2\pi i u x / N) dx$$

The magnitude of this first error will depend primarily on the rate of decay of the density function near x_0, x_N .

The second approximation is given by the approximation of the infinite sum in the above by a finite sum,

$$\sum_{u=-\infty}^{\infty} \hat{v}_u \exp(2\pi i u x / N) \Phi\left(\frac{2\pi u}{L}\right) \simeq \sum_{u=-N/2}^{N/2} \hat{v}_u \exp(2\pi i u x / N) \Phi\left(\frac{2\pi u}{L}\right)$$

The remaining error is determined by the choice of the trapezoidal rule to approximate the integral. The authors show that for characteristic functions with power decay the error will be

$$\mathcal{O}\left(N^{-\min\{\beta_3, \beta_3 + \beta_2 - 1, \beta_2 + \beta_1 - 1\}}\right)$$

where

- β_3 is the error constant of the quadrature rule used to approximate the integral, such that its leading error is $\mathcal{O}\left(N^{-\beta_3}\right)$
- β_2 is determined by the rate of decay of the characteristic function, such that $|\Phi(u)| \leq |u|^{-\beta_2}$
- β_1 is determined by the rate of decay of the coefficients \hat{v}_u , that is, $|\hat{v}_u| \leq \frac{\eta_1(N)}{|u|^{\beta_1}}$ where $\eta_1()$ is a bounding constant.

As pointed out by the authors, for value functions that are C^1 on $\left[-\frac{N}{2}, \frac{N}{2}\right]$ but whose periodic extension is discontinuous, $\beta_1 = 1$. It is reasonable to assume that the storage value as a function of the log gas price, $V(x)$, will be non-periodic as $x \rightarrow \pm\infty$. This provides the motivation for developing the FFT-m methodology described in Section 2.3.

A.7 Initial Storage Value Function

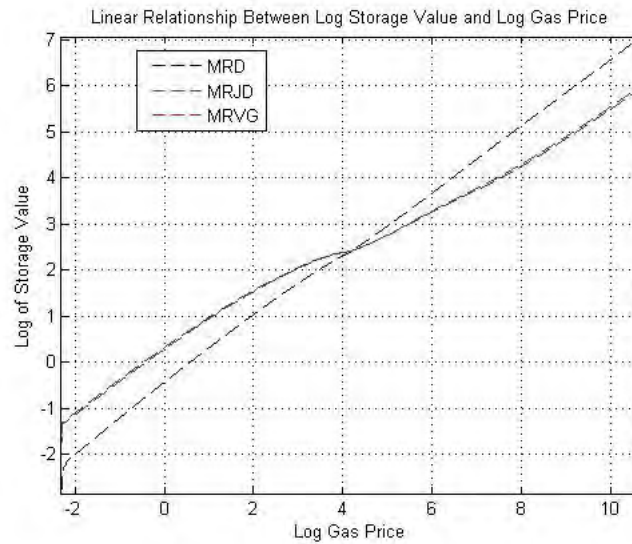


Figure A.7.1: Log Storage Value Function

The above graph shows the linear relationship between the natural logarithm of the storage value and the log gas price for each of the models calibrated in Chapter 2. The exponential dampening function used in the FFT algorithm reduces the difference in storage values at the endpoints of the gas price domain thereby reducing the error attributable to the Gibbs Phenomenon. From our numerical testing, an exponential coefficient of -1 , which implies log storage values which are linear in log-gas price with a slope of approximately 1, is sufficient in minimizing the impact of the Gibbs Phenomenon.

Appendix B

Chapter 3

B.1 Forward Curve Consistent State Variable Dynamics

Beginning with the three factor model, from Equation 3.1.8 the spot dynamics of each state variable are given by¹

$$\begin{aligned}y^{(1)}(t,t) &= -\int_0^t (\kappa_j (\exp(-\alpha(t-c)) b\sigma(c))) dc + \int_0^t \exp(-\alpha(t-c)) \sigma(c) dX(c) \\y^{(2)}(t,t) &= -\int_0^t \frac{1}{2} ((\exp(-\varepsilon(t-c)) (c_1 - c_2) + c_2) \sigma(c))^2 dc \\&\quad + \int_0^t (\exp(-\varepsilon(t-c)) (c_1 - c_2)) \sigma(c) dW(c) \\y^{(3)}(t,t) &= \int_0^t c_2 \sigma(c) dW(c)\end{aligned}\tag{B.1.1}$$

Following the argument given in the single factor derivation (Section A.4) we can see that the dynamics for the first state variable are given by

$$\begin{aligned}y^{(1)}(t) &= -\int_0^t \kappa_j (b\sigma(c)) dc + \int_0^t \alpha \int_0^u \exp(-\alpha(u-c)) b\sigma(c) \kappa_j' (\exp(-\alpha(u-c)) b\sigma(c)) dc du \\&\quad - \int_0^t \alpha \int_0^u \exp(-\alpha(u-c)) b\sigma(c) dX(c) du + \int_0^t b\sigma(c) dX(c)\end{aligned}$$

¹Note that $y^{(i)}(0, T) = 0 \forall i, T$

From Equations B.1.1 we have

$$\begin{aligned} -\alpha y^{(1)}(u) &= \alpha \int_0^u (-\kappa_j(\exp(-\alpha(u-c))b\sigma(c))) dc \\ &\quad -\alpha \int_0^u \exp(-\alpha(u-c))b\sigma(c)dX(c) \\ -\alpha \int_0^u \exp(-\alpha(u-c))b\sigma(c)dX(c) &= -\alpha y^{(1)}(u) - \alpha \int_0^u (\kappa_j(\exp(-\alpha(u-c))b\sigma(c))) dc \end{aligned}$$

and so,

$$\begin{aligned} y^{(1)}(t) &= -\int_0^t \kappa_j(b\sigma(c)) dc + \int_0^t \alpha \int_0^u \exp(-\alpha(u-c))b\sigma(c)\kappa_j'(\exp(-\alpha(u-c))b\sigma(c)) dc du \\ &\quad -\alpha \int_0^t \left(y^{(1)}(u) + \int_0^u (\kappa_j(\exp(-\alpha(u-c))b\sigma(c))) dc \right) du + \int_0^t b\sigma(c)dX(c) \end{aligned}$$

If we set $\sigma(c) = \sigma$ then

$$\alpha \int_0^u \exp(-\alpha(u-c))b\sigma\kappa_j'(\exp(-\alpha(u-c))b\sigma) dc = \kappa_j(b\sigma) - \kappa_j(\exp(-\alpha(u))b\sigma)$$

inserting into the above and differentiating gives us,

$$\begin{aligned} dy^{(1)}(t) &= \left(-\kappa_j(\exp(-\alpha(t))b\sigma) - \alpha y^{(1)}(t) - \alpha \int_0^t (\kappa_j(\exp(-\alpha(t-c))b\sigma)) dc \right) + b\sigma dX(t) \\ &= \left(\omega^{(1)}(t) - \alpha y^{(1)}(t) \right) dt + b\sigma dX(t) \end{aligned}$$

Similarly for $y^{(2)}(t)$ we have,

$$\begin{aligned} y^{(2)}(t) &= -\int_0^t \frac{1}{2}(c_1\sigma(c))^2 dc \\ &\quad + \varepsilon \int_0^t \int_0^u \exp(-\varepsilon(u-c))(c_1-c_2)\sigma(c)((\exp(-\varepsilon(u-c))(c_1-c_2)+c_2)\sigma(c)) dc du \\ &\quad + \int_0^t (c_1-c_2)\sigma(c)dW(c) - \varepsilon \int_0^t \int_0^u \exp(-\varepsilon(u-c))(c_1-c_2)\sigma(c)dW(c) du \end{aligned}$$

again from B.1.1 we have,

$$\begin{aligned} -\varepsilon y^{(2)}(u) &= \varepsilon \int_0^u \frac{1}{2}((\exp(-\varepsilon(u-c))(c_1-c_2)+c_2)\sigma(c))^2 dc \\ &\quad -\varepsilon \int_0^u (\exp(-\varepsilon(u-c))(c_1-c_2))\sigma(c)dW(c) \end{aligned}$$

and so,

$$\begin{aligned}
y^{(2)}(t) &= -\int_0^t \frac{1}{2} (c_1 \sigma(c))^2 dc \\
&\quad + \varepsilon \int_0^t \int_0^u \exp(-2\varepsilon(u-c)) (c_1 - c_2)^2 \sigma^2(c) + \exp(-\varepsilon(u-c)) (c_1 - c_2) c_2 \sigma^2(c) dc du \\
&\quad - \frac{\varepsilon}{2} \int_0^t \int_0^u ((\exp(-\varepsilon(u-c)) (c_1 - c_2) + c_2) \sigma(c))^2 dc du - \int_0^t \varepsilon y^{(2)}(u) du \\
&\quad + \int_0^t (c_1 - c_2) \sigma(c) dW(c) \\
&= -\int_0^t \frac{1}{2} (c_1 \sigma(c))^2 dc + \frac{\varepsilon}{2} \int_0^t \int_0^u \exp(-2\varepsilon(u-c)) (c_1 - c_2)^2 \sigma^2(c) dc du \\
&\quad - \frac{\varepsilon}{2} \int_0^t \int_0^u c_2^2 \sigma^2(c) dc du - \int_0^t \varepsilon y^{(2)}(u) du + \int_0^t (c_1 - c_2) \sigma(c) dW(c)
\end{aligned}$$

Setting $\sigma(c) = \sigma$ as before and differentiating gives,

$$\begin{aligned}
dy^{(2)}(t) &= \left(-\frac{1}{2} (c_1 \sigma)^2 + \frac{1}{4} (1 - \exp(-2\varepsilon t)) (c_1 - c_2)^2 \sigma^2 \right. \\
&\quad \left. - \frac{\varepsilon}{2} c_2^2 \sigma^2 t - \varepsilon y^{(2)}(t) \right) dt + (c_1 - c_2) \sigma dW(t) \\
&= \left(\omega^{(2)}(t) - \varepsilon y^{(2)}(t) \right) dt + \sigma(t) dW(t)
\end{aligned}$$

Finally, as the dynamics of $y^{(3)}(t, T)$ are independent of T we have

$$dy^{(3)}(t) = c_2 \sigma dW(t)$$

For the two factor model (MRVG-2) the spot price dynamics are given by

$$\begin{aligned}
y^{(1,2)}(t, t) &= -\int_0^t (\kappa_j (\exp(-\alpha(t-c)) b \sigma(c))) dc + \int_0^t \exp(-\alpha(t-c)) \sigma(c) dX(c) \\
&\quad - \int_0^t \frac{1}{2} ((\exp(-\alpha(t-c)) (c_1 - c_2) + c_2) \sigma(c))^2 dc \\
&\quad + \int_0^t (\exp(-\alpha(t-c)) (c_1 - c_2)) \sigma(c) dW(c) \\
y^{(3)}(t, t) &= \int_0^t c_2 \sigma(c) dW(c) \tag{B.1.2}
\end{aligned}$$

Thus, following the same logic as above we have,

$$\begin{aligned}
y^{(1)}(t) &= - \int_0^t \kappa_j(b\sigma(c)) dc + \int_0^t \alpha \int_0^u \exp(-\alpha(u-c)) b\sigma(c) \kappa_j'(\exp(-\alpha(u-c)) b\sigma(c)) dc du \\
&\quad - \int_0^t \alpha \int_0^u \exp(-\alpha(u-c)) b\sigma(c) dX(c) du + \int_0^t b\sigma(c) dX(c) \\
&\quad - \int_0^t \frac{1}{2} (c_1 \sigma(c))^2 dc \\
&\quad + \alpha \int_0^t \int_0^u \exp(-\alpha(u-c)) (c_1 - c_2) \sigma(c) ((\exp(-\alpha(u-c)) (c_1 - c_2) + c_2) \sigma(c)) dc du \\
&\quad + \int_0^t (c_1 - c_2) \sigma(c) dW(c) - \alpha \int_0^t \int_0^u \exp(-\alpha(u-c)) (c_1 - c_2) \sigma(c) dW(c) du
\end{aligned}$$

From Equations B.1.2 we have,

$$\begin{aligned}
-\alpha y^{(1)}(u) &= \alpha \int_0^u (\kappa_j(\exp(-\alpha(u-c)) b\sigma(c))) dc \\
&\quad + \alpha \int_0^u \frac{1}{2} ((\exp(-\alpha(u-c)) (c_1 - c_2) + c_2) \sigma(c))^2 dc \\
&\quad - \alpha \int_0^u \exp(-\alpha(u-c)) \sigma(c) dX(c) \\
&\quad - \alpha \int_0^u (\exp(-\alpha(u-c)) (c_1 - c_2)) \sigma(c) dW(c)
\end{aligned}$$

and so,

$$\begin{aligned}
y^{(1)}(t) &= - \int_0^t \kappa_j(b\sigma(c)) dc + \int_0^t \alpha \int_0^u \exp(-\alpha(u-c)) b\sigma(c) \kappa_j'(\exp(-\alpha(u-c)) b\sigma(c)) dc du \\
&\quad - \int_0^t \frac{1}{2} (c_1 \sigma(c))^2 dc \\
&\quad + \alpha \int_0^t \int_0^u \exp(-\alpha(u-c)) (c_1 - c_2) \sigma(c) ((\exp(-\alpha(u-c)) (c_1 - c_2) + c_2) \sigma(c)) dc du \\
&\quad - \alpha \int_0^t \left(y^{(1)}(u) + \int_0^u (\kappa_j(\exp(-\alpha(u-c)) b\sigma(c))) \right. \\
&\quad \left. + \int_0^u \frac{1}{2} ((\exp(-\alpha(u-c)) (c_1 - c_2) + c_2) \sigma(c))^2 dc \right) \\
&\quad + \int_0^t b\sigma(c) dX(c) + \int_0^t (c_1 - c_2) \sigma(c) dW(c)
\end{aligned}$$

Setting $\sigma(c) = \sigma$ we can partially evaluate the integral terms and differentiate with respect to t to yield,

$$\begin{aligned} dy^{(1)}(t) &= \left(-\kappa_j (\exp(-\alpha(t)) b\sigma) - \alpha \int_0^t (\kappa_j (\exp(-\alpha(t-c)) b\sigma)) dc \right. \\ &\quad \left. - \frac{1}{2} (c_1 \sigma)^2 + \frac{1}{4} (1 - \exp(-2\alpha t)) (c_1 - c_2)^2 \sigma^2 - \frac{\alpha}{2} c_2^2 \sigma^2 t - \alpha y^{(1)}(t) \right) dt \\ &\quad + b\sigma dX(t) + (c_1 - c_2) \sigma dW(t) \\ &= \left(\omega^{(1)}(t) - \alpha y^{(1)}(t) \right) dt + b\sigma dX(t) + (c_1 - c_2) \sigma dW(t) \end{aligned}$$

and as before,

$$dy^{(2)}(t) = c_2 \sigma dW(t)$$

Finally, for the abridged two factor model (MRVG-2x) we have

$$\begin{aligned} y^{(1)}(t,t) &= - \int_0^t (\kappa_j (\exp(-\alpha(t-c)) b\sigma(c))) dc + \int_0^t \exp(-\alpha(t-c)) \sigma(c) dX(c) \\ y^{(2)}(t,t) &= - \int_0^t \frac{1}{2} (\sigma_L(c))^2 dc + \int_0^t \sigma_L(c) dW(c) \end{aligned} \quad (\text{B.1.3})$$

and so constraining σ as before and setting $\sigma_L(c) = \sigma_L$ we have

$$\begin{aligned} dy^{(1)}(t) &= \left(-\kappa_j (\exp(-\alpha(t)) b\sigma) - \alpha y^{(1)}(t) - \alpha \int_0^t (\kappa_j (\exp(-\alpha(t-c)) b\sigma)) dc \right) + b\sigma dX(t) \\ &= \left(\omega^{(1)}(t) - \alpha y^{(1)}(t) \right) dt + b\sigma dX(t) \end{aligned}$$

$$dy^{(2)}(t) = -\frac{1}{2} (\sigma_L)^2 dt + \sigma_L dW(t)$$

B.2 Derivation of Characteristic Function of the MRVG-3 Model

The only non-trivial derivation needed in Equation 3.1.10 is that of the expectations

$$E^{(1)}(z^{(1)}) = E \left[\exp \left(\int_s^t iz^{(1)} (\exp(-\alpha(t-c))) b\sigma(c) dX(c) \right) \right] \quad (\text{B.2.1})$$

$$E^{(2)}(z^{(2)}, z^{(3)}) = E \left[\exp \left(\int_s^t i \left((z^{(2)} \exp(-\varepsilon(t-s)) (c_1 - c_2) + z^{(3)} c_2) \sigma(c) \right) dW(c) \right) \right] \quad (\text{B.2.2})$$

To derive a closed-form solution we first set $\sigma(t) = \sigma$ as a time dependent volatility curve would generally require numerical integration. The first expectation is then given by the characteristic function of the

MRVG model (Section A.2)

$$E^{(1)}\left(z^{(1)}\right) = \exp\left(\frac{1}{2v\alpha} [Li_2(l_1) - Li_2(l_0)]\right)$$

where

$$\begin{aligned} l_1 &= -\left(\frac{(b\sigma)^2 v}{2}\right) z^{(1)2} \\ l_0 &= -\left(\frac{(b\sigma)^2 v}{2}\right) z^{(1)2} \exp(-2\alpha(t-s)) \end{aligned}$$

In a similar manner, we can solve the second expectation by recalling that the characteristic exponent of the Normal Distribution is given by

$$\varphi_W(z) = -\frac{1}{2}z^2$$

and so we need to solve the the integral term in

$$\begin{aligned} E^{(2)}\left(z^{(2)}, z^{(3)}\right) &= \exp\left(\int_s^t \varphi_W\left(\left(z^{(2)} \exp(-\varepsilon(t-c))(c_1 - c_2) + z^{(3)}c_2\right) \sigma\right) dc\right) \\ &= \exp\left(-\frac{\sigma^2}{2} \int_s^t \left(z^{(2)} \exp(-\varepsilon(t-c))(c_1 - c_2) + z^{(3)}c_2\right)^2 dc\right) \end{aligned}$$

Squaring the integrand gives us

$$\left(z^{(2)}\right)^2 \exp(-2\varepsilon(t-c))(c_1 - c_2)^2 + 2z^{(2)} \exp(-\varepsilon(t-c))(c_1 - c_2)z^{(3)}c_2 + \left(z^{(3)}\right)^2 c_2^2$$

Integrating gives

$$\frac{1}{2\varepsilon} \left(z^{(2)}\right)^2 (c_1 - c_2)^2 (1 - \exp(-2\varepsilon(t-s))) + \frac{2}{\varepsilon} z^{(2)} (1 - \exp(-\varepsilon(t-s)))(c_1 - c_2)z^{(3)}c_2 + \left(z^{(3)}\right)^2 c_2^2 (t-s)$$

and so

$$\begin{aligned} E^{(2)}\left(z^{(2)}, z^{(3)}\right) &= \exp\left[\left(-\frac{\sigma^2}{2}\right) \left(\frac{1}{2\varepsilon} \left(z^{(2)}\right)^2 (c_1 - c_2)^2 (1 - \exp(-2\varepsilon(t-s)))\right.\right. \\ &\quad \left.\left. + \frac{2}{\varepsilon} z^{(2)} (1 - \exp(-\varepsilon(t-s)))(c_1 - c_2)z^{(3)}c_2 + \left(z^{(3)}\right)^2 c_2^2 (t-s)\right)\right] \end{aligned}$$

Now, we know from the single factor case that we need to restrict the function $E^{(1)}\left(z^{(1)}\right)$ by ensuring that $Li_2(l_0), Li_2(l_1) \in \mathbb{C} \setminus (1, \infty)$. This implies that

$$|w| < \sqrt{\frac{2}{(b\sigma)^2 v}}$$

The function $E^{(2)}(z^{(2)}, z^{(3)})$ is defined over the whole of \mathbb{C}^2 and thus $\Phi_{\vec{y}(t)}(\vec{z}; s, t)$ is defined on $\mathbb{C}^2 \cup S_y$ where S_y is a strip in the complex plane such that $S_y := \left\{ z = u + iw; w \in \left(-\sqrt{\frac{2}{(b\sigma)^2 v}}, \sqrt{\frac{2}{(b\sigma)^2 v}} \right) \right\}$. The characteristic function for the MRVG-3x model is given by setting $c_2 = 0$.

B.3 Derivation of the Characteristic Function of MRVG-2 Model

The derivation for the MRVG-2 model follows that of the three factor model where

$$\begin{aligned} E^{(1)}(z^{(1)}) &= E \left[\exp \left(\int_s^t i z^{(1)} (\exp(-\alpha(t-c))) b\sigma dX(c) \right) \right] \\ &= \exp \left(\frac{1}{2v\alpha} [Li_2(l_1) - Li_2(l_0)] \right) \end{aligned} \quad (\text{B.3.1})$$

with

$$\begin{aligned} l_1 &= - \left(\frac{(b\sigma)^2 v}{2} \right) z^{(1)2} \\ l_0 &= - \left(\frac{(b\sigma)^2 v}{2} \right) z^{(1)2} \exp(-2\alpha(t-s)) \end{aligned}$$

$$\begin{aligned} E^{(2)}(z^{(1)}, z^{(2)}) &= E \left[\exp \left(\int_s^t i \left(z^{(1)} \sigma \exp(-\alpha(t-c)) (c_1 - c_2) + z^{(2)} \sigma c_2 \right) dW(c) \right) \right] \\ &= \exp \left[\left(-\frac{\sigma^2}{2} \right) \left(\frac{1}{2\alpha} (z^{(1)})^2 (c_1 - c_2)^2 (1 - \exp(-2\alpha(t-s))) \right. \right. \\ &\quad \left. \left. + \frac{2}{\alpha} z^{(1)} (1 - \exp(-\alpha(t-s))) (c_1 - c_2) z^{(2)} c_2 + (z^{(2)})^2 c_2^2 (t-s) \right) \right] \end{aligned} \quad (\text{B.3.2})$$

Repeating the analysis carried out for the MRVG-3 model, we see that $\Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(s), s, t)$ is also defined on $\mathbb{C} \cup S_y$. For the abridged MRVG-2x model, the second expectation becomes

$$\begin{aligned} E^{(2)}(z^{(2)}) &= E \left[\exp \left(\int_s^t i z^{(2)} \sigma_L dW(c) \right) \right] \\ &= \exp \left(-\frac{(z^{(2)} \sigma_L)^2}{2} \right) \end{aligned}$$

B.4 Distribution Moments

We can derive the moments for the innovations on the MRVG-3 process with constant volatility through the moment generating function $m(u)$ in a manner identical to the MRVG model derivations in Section A.3. Further we can utilize equations Equation B.2.1 and Equation B.2.2 to aid in our analysis. First we define, the random variable

$$\begin{aligned} n_t &= \int_0^t (\exp(-\alpha(\tau))) b \sigma(\tau) dX^{(1)}(\tau) \\ &\quad + \int_0^t (\exp(-\varepsilon(\tau)) (c_1 - c_2) + c_2) \sigma dW^{(2)}(\tau) \end{aligned}$$

and therefore, we need

$$\begin{aligned} m(u) &= E[\exp(un_t)] \\ &= E^{(1)}(\exp(iu)) E^{(2)}(\exp(iu), \exp(iu)) \\ &= M^{(1)}(u) M^{(2)}(u) \end{aligned}$$

Since we wish to calculate the first 4 moments, we will need to derive the first 4 derivatives of $m(u)$. To clarify the notation used, we continue to identify a factor by the superscript (k) and for the remainder of this section, we will denote the n^{th} derivative of a function with the superscript n . Thus,

$$\begin{aligned} m^1(u) &= M^{(1)1}(u) M^{(2)}(u) + M^{(1)}(u) M^{(2)1}(u) \\ m^2(u) &= M^{(1)2}(u) M^{(2)}(u) + 2M^{(1)1}(u) M^{(2)1}(u) \\ &\quad + M^{(1)}(u) M^{(2)2}(u) \\ m^3(u) &= M^{(1)3}(u) M^{(2)}(u) \\ &\quad + 3M^{(1)2}(u) M^{(2)1}(u) + 3M^{(1)1}(u) M^{(2)2}(u) \\ &\quad + M^{(1)}(u) M^{(2)3}(u) \\ m^4(u) &= M^{(1)4}(u) M^{(2)}(u) \\ &\quad + 4M^{(1)3}(u) M^{(2)1}(u) + 6M^{(1)2}(u) M^{(2)2}(u) \\ &\quad + 4M^{(1)1}(u) M^{(2)3}(u) + M^{(1)}(u) M^{(2)4}(u) \end{aligned}$$

From the MRVG derivation given in A.3, we know the first 4 moments of

$$\int_0^t (\exp(-\alpha(\tau))) b \sigma dX^{(1)}(\tau)$$

therefore we simply need to derive the corresponding values of the diffusion process

$$\int_0^t (\exp(-\varepsilon(\tau))(c_1 - c_2) + c_2) \sigma dW^{(2)}(\tau)$$

The first moment and third moments of the diffusion are of course equal to 0, as are the first and third moments of the MRVG. Therefore,

$$E[n_t] = E[n_t^3] = 0$$

The second moment of the MRVG process we know is given by

$$\frac{(b\sigma)^2}{2\alpha} (1 - \exp(-2\alpha t))$$

therefore

$$m^2(0) = \frac{(b\sigma)^2}{2\alpha} (1 - \exp(-2\alpha t)) + M^{(2)2}(0)$$

Now,

$$\begin{aligned} M^{(2)}(u) &= \exp\left((u)^2 \left(\frac{\sigma^2}{2}\right) \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 (1 - \exp(-2\varepsilon t))\right.\right. \\ &\quad \left.\left. + \frac{2}{\varepsilon} (1 - \exp(-\varepsilon t)) (c_1 - c_2) c_2 + c_2^2 t\right)\right) \\ &= \exp\left((u)^2 \lambda\right) \end{aligned}$$

and so,

$$\begin{aligned} M^{(2)2}(u) &= 2\lambda M^{(2)}(u) + (2u\lambda)^2 M^{(2)}(u) \\ M^{(2)2}(0) &= (\sigma^2) \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 (1 - \exp(-2\varepsilon t))\right. \\ &\quad \left. + \frac{2}{\varepsilon} (1 - \exp(-\varepsilon t)) (c_1 - c_2) c_2 + c_2^2 t\right) \end{aligned}$$

The second moment of the innovations are thus

$$\begin{aligned} E[n_t^2] &= \frac{(b\sigma)^2}{2\alpha} (1 - \exp(-2\alpha t)) \\ &\quad + (\sigma^2) \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 (1 - \exp(-2\varepsilon t))\right. \\ &\quad \left. + \frac{2}{\varepsilon} (1 - \exp(-\varepsilon t)) (c_1 - c_2) c_2 + c_2^2 t\right) \end{aligned}$$

for large t the variance of the process is therefore equal to

$$\frac{(b\sigma)^2}{2\alpha} + (\sigma^2) \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 + \frac{2}{\varepsilon} (c_1 - c_2) c_2 + c_2^2 t \right)$$

The fourth moment of the MRVG-3 process is given by

$$\begin{aligned} m^4(0) &= M^{(1)4}(0) + M^{(2)4}(0) \\ &\quad + 6M^{(1)2}(0)M^{(2)2}(0) \end{aligned}$$

We know from the MRVG moment derivation that,

$$M^{(1)4}(0) = \frac{3(b\sigma)^4}{4\alpha^2} (1 - \exp(-2\alpha t))^2 + \frac{3(b\sigma)^4 v}{4\alpha} (1 - \exp(-4\alpha t))$$

Now,

$$\begin{aligned} M^{(2)4}(0) &= 12\lambda^2 \\ &= 12 \left(\frac{\sigma^2}{2} \right)^2 \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 (1 - \exp(-2\varepsilon t)) \right. \\ &\quad \left. + \frac{2}{\varepsilon} (1 - \exp(-\varepsilon t)) (c_1 - c_2) c_2 + c_2^2 t \right)^2 \end{aligned}$$

Therefore we have,

$$\begin{aligned} E[n_t^4] &= \frac{3(b\sigma)^4}{4\alpha^2} (1 - \exp(-2\alpha t))^2 + \frac{3(b\sigma)^4 v}{4\alpha} (1 - \exp(-4\alpha t)) \\ &\quad + 12 \left(\frac{\sigma^2}{2} \right)^2 \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 (1 - \exp(-2\varepsilon t)) \right. \\ &\quad \left. + \frac{2}{\varepsilon} (1 - \exp(-\varepsilon t)) (c_1 - c_2) c_2 + c_2^2 t \right)^2 \\ &\quad + 6 \frac{(b\sigma)^2}{2\alpha} (1 - \exp(-2\alpha t)) (\sigma^2) \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 (1 - \exp(-2\varepsilon t)) \right. \\ &\quad \left. + \frac{2}{\varepsilon} (1 - \exp(-\varepsilon t)) (c_1 - c_2) c_2 + c_2^2 t \right) \end{aligned}$$

and for large t we have

$$\begin{aligned} E [n_t^4] &= \frac{3(b\sigma)^4}{4\alpha^2} + \frac{3(b\sigma)^4 v}{4\alpha} + 12 \left(\frac{\sigma^2}{2} \right)^2 \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 \right. \\ &\quad \left. + \frac{2}{\varepsilon} (c_1 - c_2) c_2 + c_2^2 t \right)^2 + 6 \frac{(\sigma^2)^2}{2\alpha} (\sigma^2) \left(\frac{1}{2\varepsilon} (c_1 - c_2)^2 \right. \\ &\quad \left. + \frac{2}{\varepsilon} (c_1 - c_2) c_2 + c_2^2 t \right) \end{aligned}$$

Analogous formulae can easily be derived for the MRVG-3x model by simply setting $c_2 = 0$. The equivalent formulae for the MRVG-2 model are given by setting $\varepsilon = \alpha$ so that

$$\begin{aligned} E [n_t^2] &= \frac{(b\sigma)^2}{2\alpha} (1 - \exp(-2\alpha t)) \\ &\quad + (\sigma^2) \left(\frac{1}{2\alpha} (c_1 - c_2)^2 (1 - \exp(-2\alpha t)) \right. \\ &\quad \left. + \frac{2}{\alpha} (1 - \exp(-\alpha t)) (c_1 - c_2) c_2 + c_2^2 t \right) \end{aligned}$$

and

$$\begin{aligned} E [n_t^4] &= \frac{3(b\sigma)^4}{4\alpha^2} (1 - \exp(-2\alpha t))^2 + \frac{3(b\sigma)^4 v}{4\alpha} (1 - \exp(-4\alpha t)) \\ &\quad + 12 \left(\frac{\sigma^2}{2} \right)^2 \left(\frac{1}{2\alpha} (c_1 - c_2)^2 (1 - \exp(-2\alpha t)) \right. \\ &\quad \left. + \frac{2}{\alpha} (1 - \exp(-\alpha t)) (c_1 - c_2) c_2 + c_2^2 t \right)^2 \\ &\quad + 6 \frac{(b\sigma)^2}{2\alpha} (1 - \exp(-2\alpha t)) (\sigma^2) \left(\frac{1}{2\alpha} (c_1 - c_2)^2 (1 - \exp(-2\alpha t)) \right. \\ &\quad \left. + \frac{2}{\alpha} (1 - \exp(-\alpha t)) (c_1 - c_2) c_2 + c_2^2 t \right) \end{aligned}$$

Finally, the moments for the MRVG-2x model are given by

$$E [n_t^2] = \frac{\sigma^2}{2\alpha} (1 - \exp(-2\alpha t)) + \sigma_L^2 t$$

and

$$E [n_t^4] = \frac{3\sigma^4}{4\alpha^2} (1 - \exp(-2\alpha t))^2 + \frac{3\sigma^4\nu}{4\alpha} (1 - \exp(-4\alpha t)) + 3\sigma_L^4 t^2$$

B.5 Log Forward Price

Similar to the single factor case, we can derive the log forward price for maturity T , at any given future time t , as a deterministic function of the state variables $y^{(i)}(t)$ for $i = 1, 2, 3$.

Recall from Equation 3.1.8 that for the MRVG-3 we have,

$$y^{(1)}(t, T) = y^{(1)}(s, T) - \int_s^t (\kappa_j (b \exp(-\alpha(T-c)) \sigma(c))) dc + \int_s^t b \exp(-\alpha(T-c)) \sigma(c) dX(c)$$

and,

$$y^{(1)}(t, t) = y^{(1)}(s, t) - \int_s^t (\kappa_j (b \exp(-\alpha(t-c)) \sigma(c))) dc + \int_s^t b \exp(-\alpha(t-c)) \sigma(c) dX(c)$$

therefore,

$$\begin{aligned} y^{(1)}(t, T) &= y^{(1)}(s, T) - \int_s^t (\kappa_j (b \exp(-\alpha(T-c)) \sigma(c))) dc \\ &\quad + \exp(-\alpha(T-t)) \left(y^{(1)}(t, t) - y^{(1)}(s, t) + \int_s^t (\kappa_j (b \exp(-\alpha(t-c)) \sigma(c))) dc \right) \end{aligned}$$

Setting $\sigma(c) = \sigma$ we have,

$$\begin{aligned} y^{(1)}(t, T) &= y^{(1)}(s, T) - \int_s^t (\kappa_j (b \exp(-\alpha(T-c)) \sigma)) dc \\ &\quad + \exp(-\alpha(T-t)) \left(y^{(1)}(t, t) - y^{(1)}(s, t) \right) \\ &\quad + \exp(-\alpha(T-t)) \varphi(-i, s, t) \end{aligned}$$

where φ is the characteristic exponent of the MRVG model.

Similarly we have,

$$\begin{aligned} y^{(2)}(t, T) &= y^{(2)}(s, T) - \int_s^t \frac{1}{2} ((\exp(-\varepsilon(T-c))(c_1 - c_2) + c_2) \sigma(c))^2 dc \\ &\quad + \int_s^t (\exp(-\varepsilon(T-c))(c_1 - c_2)) \sigma(c) dW(c) \end{aligned}$$

$$\begin{aligned} y^{(2)}(t, t) &= y^{(2)}(s, t) - \int_0^t \frac{1}{2} ((\exp(-\varepsilon(t-c))(c_1 - c_2) + c_2) \sigma(c))^2 dc \\ &\quad + \int_s^t (\exp(-\varepsilon(t-c))(c_1 - c_2)) \sigma(c) dW(c) \end{aligned}$$

$$\begin{aligned} y^{(2)}(t, T) &= y^{(2)}(s, T) - \int_s^t \frac{1}{2} ((\exp(-\varepsilon(T-c))(c_1 - c_2) + c_2) \sigma(c))^2 dc \\ &\quad + \exp(-\varepsilon(T-t)) (y^{(2)}(t, t) - y^{(2)}(s, t)) \\ &\quad + \exp(-\varepsilon(T-t)) \left(\int_s^t \frac{1}{2} ((\exp(-\varepsilon(t-c))(c_1 - c_2) + c_2) \sigma(c))^2 dc \right) \end{aligned}$$

Again setting $\sigma(c) = \sigma$ gives,

$$\begin{aligned} y^{(2)}(t, T) &= y^{(2)}(s, T) + \exp(-\varepsilon(T-t)) (y^{(1)}(t, t) - y^{(1)}(s, t)) \\ &\quad - \frac{\sigma^2}{4\varepsilon} (\exp(-2\varepsilon(T-t)) - \exp(-2\varepsilon(T-s))) (c_1 - c_2)^2 \\ &\quad - \frac{\sigma^2}{\varepsilon} (\exp(-\varepsilon(T-t)) - \exp(-\varepsilon(T-s))) (c_1 - c_2) c_2 \\ &\quad + \frac{\sigma^2}{4\varepsilon} (1 - \exp(-2\varepsilon(t-s))) (c_1 - c_2)^2 \exp(-\varepsilon(T-t)) \\ &\quad + \frac{\sigma^2}{\varepsilon} (1 - \exp(-\varepsilon(t-s))) (c_1 - c_2) c_2 \exp(-\varepsilon(T-t)) \\ &\quad - \frac{\sigma^2}{2} c_2^2 (t-s) (1 - \exp(-\varepsilon(T-t))) \\ &= y^{(2)}(s, T) + \exp(-\varepsilon(T-t)) (y^{(1)}(t, t) - y^{(1)}(s, t)) \\ &\quad - \frac{\sigma^2}{4\varepsilon} \exp(-2\varepsilon(T-t)) (c_1 - c_2)^2 \times \\ &\quad ((1 - \exp(-2\varepsilon(t-s))) (1 - \exp(\varepsilon(T-t)))) \\ &\quad - \frac{\sigma^2}{2} c_2^2 (t-s) (1 - \exp(-\varepsilon(T-t))) \end{aligned}$$

and finally,

$$\begin{aligned} y^{(3)}(t, T) &= y^{(3)}(s, T) + \int_s^t c_2 \sigma(c) dW(c) \\ &= y^{(3)}(s, T) + y^{(3)}(t, t) - y^{(3)}(s, t) \end{aligned}$$

For the MRVG-2 model we therefore have

$$\begin{aligned} y^{(1)}(t, T) &= y^{(1)}(s, T) - \int_s^t (\kappa_j (b \exp(-\alpha(T-c)) \sigma)) dc \\ &\quad + \exp(-\alpha(T-t)) (y^{(1)}(t, t) - y^{(1)}(s, t)) \\ &\quad + \exp(-\alpha(T-t)) \varphi(-i, s, t) \\ &\quad - \frac{\sigma^2}{4\alpha} \exp(-2\alpha(T-t)) (c_1 - c_2)^2 \times \\ &\quad ((1 - \exp(\alpha(T-t))) (1 - \exp(-2\alpha(t-s)))) \\ &\quad - \frac{\sigma^2}{2} c_2^2 (t-s) (1 - \exp(-\alpha(T-t))) \end{aligned}$$

$$\begin{aligned} y^{(2)}(t, T) &= y^{(2)}(s, T) + \int_s^t c_2 \sigma(c) dW(c) \\ &= y^{(2)}(s, T) + y^{(2)}(t, t) - y^{(2)}(s, t) \end{aligned}$$

For the MRVG-2x model we have

$$\begin{aligned} y^{(1)}(t, T) &= y^{(1)}(s, T) - \int_s^t (\kappa_j (b \exp(-\alpha(T-c)) \sigma)) dc \\ &\quad + \exp(-\alpha(T-t)) (y^{(1)}(t, t) - y^{(1)}(s, t)) \\ &\quad + \exp(-\alpha(T-t)) \varphi(-i, s, t) \end{aligned}$$

$$\begin{aligned} y^{(2)}(t, T) &= y^{(2)}(s, T) - \frac{1}{2} \int_s^t \sigma_L^2(c) dc + \int_s^t \sigma_L(c) dW(c) \\ &= y^{(2)}(s, T) + y^{(2)}(t, t) - y^{(2)}(s, t) \end{aligned}$$

Finally, from Equation 3.1.3 we have,

$$\begin{aligned} df(t, T) &= \sum_{k=1}^K dy^{(k)}(t, T) \\ f(t, T) &= f(s, T) + \sum_{k=1}^K y^{(k)}(t, T) - y^{(k)}(s, T) \end{aligned}$$

Thus the current values of the state variables along with the initial log forward price are all that are required to derive the current value of the log forward price.

B.6 Forward Curve Covariance Function

Firstly, recall that the log dynamics of a forward price at maturity T under the MRVG-3 model are given by Equation 3.1.3 and Equation 3.1.8 as

$$df(t, T) = \sum_{k=1}^K dy^{(k)}(t, T)$$

$$\begin{aligned} dy^{(1)}(t, T) &= (-\kappa_j(\exp(-\alpha(T-t))b\sigma(t)))dt + \exp(-\alpha(T-t))b\sigma(t)dX(t) \\ dy^{(2)}(t, T) &= -\frac{1}{2}((\exp(-\varepsilon(T-t))(c_1 - c_2) + c_2)\sigma(t))^2 dt \\ &\quad + (\exp(-\varepsilon(T-t))(c_1 - c_2))\sigma(t)dW(t) \\ dy^{(3)}(t, T) &= c_2\sigma(t)dW(t) \end{aligned}$$

Now assuming constant σ , and taking two maturities T_s and T_l , the instantaneous covariance of the log forward price returns are given by

$$\begin{aligned} Cov[t; T_s, T_l] &= E \left[\left(\sum_{k=1}^K dy^{(k)}(t, T_s) - E \left[\sum_{k=1}^K dy^{(k)}(t, T_s) \right] \right) \left(\sum_{k=1}^K dy^{(k)}(t, T_l) - E \left[\sum_{k=1}^K dy^{(k)}(t, T_l) \right] \right) \right] \\ &= E [\exp(-\alpha(T_s - t))b\sigma dX(t) \exp(-\alpha(T_l - t))b\sigma dX(t)] \\ &\quad + E [((\exp(-\varepsilon(T_s - t))(c_1 - c_2)) + c_2)\sigma dW(t) \times \\ &\quad ((\exp(-\varepsilon(T_l - t))(c_1 - c_2)) + c_2)\sigma dW(t)] \\ &= \exp(-\alpha(T_l + T_s - 2t))b^2\sigma^2 dt + (\exp(-\varepsilon(T_l + T_s - 2t))(c_1 - c_2)^2)\sigma^2 dt \\ &\quad + c_2(c_1 - c_2)(\exp(-\varepsilon(T_l - t))\sigma^2 dt + \exp(-\varepsilon(T_s - t))) + c_2^2\sigma^2 dt \end{aligned}$$

Setting $c_2 = 0$ gives the MRVG-3x model equivalent

$$\begin{aligned} Cov[t; T_s, T_l] &= \exp(-\alpha(T_l + T_s - 2t))b^2\sigma^2 dt \\ &\quad + (\exp(-\varepsilon(T_l + T_s - 2t))(c_1)^2)\sigma^2 dt \end{aligned}$$

For the MRVG-2 model we have

$$\begin{aligned} Cov[t; T_s, T_l] &= \exp(-\alpha(T_l + T_s - 2t))b^2\sigma^2 dt + (\exp(-\alpha(T_l + T_s - 2t))(c_1 - c_2)^2)\sigma^2 dt \\ &\quad + c_2(c_1 - c_2)(\exp(-\alpha(T_l - t))\sigma^2 dt + \exp(-\alpha(T_s - t))) + c_2^2\sigma^2 dt \end{aligned}$$

and for the MRVG-2x model,

$$\text{Cov}[t; T_s, T_l] = \exp(-\alpha(T_l + T_s - 2t)) \sigma^2 dt + \sigma_L^2 dt$$

B.7 Multi-Factor Swaption Pricing Algorithm

Recall that the payoff for a call option is given by

$$\begin{aligned} V(F, K, t, T_1, T_2) &= \max[F(t, T_1, T_2) - K, 0] \\ &= \max \left[\frac{1}{N} \sum_{j=0}^{N-1} F(t, T_1 + j\Delta t) - K, 0 \right] \\ &= \max \left[\frac{1}{N} \sum_{j=0}^{N-1} \exp(f(0, T_1 + j\Delta t) + \vec{y}(t, T_1 + j\Delta t)) - K, 0 \right] \end{aligned}$$

where

$F(t, T)$ is the forward price observed at time t and delivering at time T .

$f(t, T)$ is the log of $F(t, T)$.

$\vec{y}(t, T)$ is a vector of underlying stochastic factors.

Δt is the delivery frequency of the commodity underlying the forward contract.

N is the number of deliveries of the underlying forward contract.

For clarity, we shall write the payoff as $V(\vec{y}(t), K, t, T_1, T_2)$. We can then derive the option prices numerically, following the general framework outlined in Section 3.3. The expected payoff² is given by

$$E \left[V(\vec{y}(t), K, t, T_1, T_2) | \vec{y}(0) \right] = \left(\frac{1}{2\pi} \right)^K \int_{\mathbb{C}} \dots \int_{\mathbb{C}} \tilde{v}(\vec{z}) \Phi(\vec{z}, \vec{y}(0), 0, t) dz^{(1)} dz^{(2)} \dots dz^{(K)}$$

where

$$\tilde{v}(\vec{z}) = \int_{\mathbb{R}^n} \exp(-i\vec{z}^T \cdot \vec{y}(t)) V(\vec{y}(t), K, t, T_1, T_2) d\vec{y}(t) \quad (\text{B.7.1})$$

The swaption payoff can be evaluated quickly by utilizing the formula for the log-forward prices, conditional upon the future values of the state variables, derived in Section B.5 of the Appendix. For example,

²For simplicity, we ignore the effect of discounting.

noting that $y(\vec{0}, T) = 0 \forall T$, for the MRVG-2x model we have

$$E \left[V \left(y(\vec{t}), K, t, T_1, T_2 \right) | y(\vec{0}) \right] = \left(\frac{1}{2\pi} \right)^2 \int_{\mathbb{C}} \int_{\mathbb{C}} \tilde{v} \left(z^{(1)}, z^{(2)} \right) \Phi \left(\vec{z}, y(\vec{0}), 0, t \right) dz^{(1)} dz^{(2)}$$

where

$$\tilde{v} \left(z^{(1)}, z^{(2)} \right) = \int_{\mathbb{R}} \int_{\mathbb{R}} \exp \left(-z^{(1)} y^{(1)}(t) - z^{(2)} y^{(2)}(t) \right) \max \left[\frac{1}{N} \sum_{j=0}^{N-1} \exp \left(f(0, T_1 + j\Delta t) + \exp \left(-\alpha (T_1 + j\Delta t - t) \right) y^{(1)}(t) + g(t, T_1 + j\Delta t) + y^{(2)}(t) \right) - \exp(k), 0 \right]$$

and $g(t, T + j\Delta t)$ is given by

$$g(t, T_1 + j\Delta t) = - \int_0^t \left(\kappa_j (b \exp(-\alpha (T_1 + j\Delta t - c)) \sigma) \right) dc + \exp(-\alpha (T_1 + j\Delta t - t)) \varphi(-i, 0, t)$$

B.8 Moment Formulae

We will begin by deriving the relevant expression for the most general model of interest, the MRVG-3 forward curve model. By placing the necessary restrictions on the model parameters we can then easily use these results to derive the moment formulae of the MRVG-2, MRVG-2x, and MRVG models. Beginning with the second moment, recall that we need to derive a solution for

$$\frac{1}{F(0, T_1, T_2)^2} E \left[\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} F(0, T_1 + i\Delta t) F(0, T_1 + j\Delta t) \exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t)) | y(\vec{0}) \right]$$

Taking the expectation operator inside the summation, the only difficulty is in evaluating the following

$$E \left[\exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t)) | y(0) \right]$$

For the MRVG-3 process the function $y(t, T_i)$ is given by

$$y(t, T_i) = y_1(t, T_i) + y_2(t, T_i) + y_3(t, T_i)$$

From Section B.5 of the Appendix we know that

$$\begin{aligned} y^{(1)}(t, T_i) &= y^{(1)}(s, T_i) - \int_s^t (\kappa_j (b \exp(-\alpha(T_i - c)) \sigma)) dc \\ &\quad + \exp(-\alpha(T_i - t)) \left(y^{(1)}(t, t) - y^{(1)}(s, t) \right) \\ &\quad + \exp(-\alpha(T_i - t)) \varphi(-i, s, t) \end{aligned}$$

where φ is the characteristic exponent of the MRVG model

$$\begin{aligned} y^{(2)}(t, T_i) &= y^{(2)}(s, T_i) + \exp(-\varepsilon(T_i - t)) \left(y^{(2)}(t, t) - y^{(2)}(s, t) \right) \\ &\quad - \frac{\sigma^2}{4\varepsilon} \exp(-2\varepsilon(T_i - t)) (c_1 - c_2)^2 \times \\ &\quad \left((1 - \exp(\varepsilon(T_i - t))) (1 - \exp(-2\varepsilon(t - s))) \right) \\ &\quad - \frac{\sigma^2}{2} c_2^2 (t - s) (1 - \exp(-\varepsilon(T_i - t))) \end{aligned}$$

and,

$$\begin{aligned} y^{(3)}(t, T_i) &= y^{(3)}(s, T_i) + \int_s^t c_2 \sigma(c) dW(c) \\ &= y^{(3)}(s, T_i) + y^{(3)}(t, t) - y^{(3)}(s, t) \end{aligned}$$

Setting $s = 0$, $y^{(k)}(0, T_i) = 0$ for all k , and using the shorthand $\sum_{x \in \{i, j\}} f_x = f_i + f_j$, we can write the exponent term in the expectation as

$$\begin{aligned} y(t, T_i) + y(t, T_j) &= y^{(1)}(t, t) \sum_{x \in \{i, j\}} \exp(-\alpha(T_x - t)) + y^{(2)}(t, t) \sum_{x \in \{i, j\}} \exp(-\varepsilon(T_x - t)) + 2y^{(3)}(t, t) \\ &\quad + \sum_{x \in \{i, j\}} - \int_0^t (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\ &\quad + \sum_{x \in \{i, j\}} - \frac{\sigma^2}{4\varepsilon} \exp(-2\varepsilon(T_x - t)) (c_1 - c_2)^2 \left((1 - \exp(\varepsilon(T_x - t))) (1 - \exp(-2\varepsilon t)) \right) \\ &\quad + \sum_{x \in \{i, j\}} - \frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\varepsilon(T_x - t))) \end{aligned}$$

Focusing on the non-deterministic part of the above and noticing that

$$E \left[\exp \left(y^{(1)}(t, t) \sum_{x \in \{i, j\}} \exp(-\alpha(T_x - t)) + y^{(2)}(t, t) \sum_{x \in \{i, j\}} \exp(-\varepsilon(T_x - t)) + 2y^{(3)}(t, t) \right) \right]$$

is equal to $\Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(0), 0, t)$ with

$$\vec{z} = \begin{pmatrix} -i \sum_{x \in \{i, j\}} \exp(-\alpha(T_x - t)) \\ -i \sum_{x \in \{i, j\}} \exp(-\varepsilon(T_x - t)) \\ -2i \end{pmatrix}$$

we obtain the second moment of the forward price for the MRVG-3 model as

$$\frac{1}{F(0, T_1, T_2)^2} \left[\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} F(0, T_i) F(0, T_j) \Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(0), 0, t) \exp(A) \right]$$

where $T_x = T_1 + x\Delta t$

$$\begin{aligned} A &= \sum_{x \in \{i, j\}} - \int_0^t (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\ &+ \sum_{x \in \{i, j\}} - \frac{\sigma^2}{4\varepsilon} \exp(-2\varepsilon(T_x - t)) (c_1 - c_2)^2 ((1 - \exp(\varepsilon(T_x - t))) (1 - \exp(-2\varepsilon t))) \\ &+ \sum_{x \in \{i, j\}} - \frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\varepsilon(T_x - t))) \end{aligned}$$

Similarly, the third moment is given by

$$\frac{1}{F(0, T_1, T_2)^3} \left[\frac{1}{N^3} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} F(0, T_i) F(0, T_j) F(0, T_k) \Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(0), 0, t) \exp(A) \right]$$

where

$$\vec{z} = \begin{pmatrix} -i \sum_{x \in \{i, j, k\}} \exp(-\alpha(T_x - t)) \\ -i \sum_{x \in \{i, j, k\}} \exp(-\varepsilon(T_x - t)) \\ -3i \end{pmatrix}$$

and

$$\begin{aligned} A &= \sum_{x \in \{i, j, k\}} - \int_0^t (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\ &+ \sum_{x \in \{i, j, k\}} - \frac{\sigma^2}{4\varepsilon} \exp(-2\varepsilon(T_x - t)) (c_1 - c_2)^2 ((1 - \exp(\varepsilon(T_x - t))) (1 - \exp(-2\varepsilon t))) \\ &+ \sum_{x \in \{i, j, k\}} - \frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\varepsilon(T_x - t))) \end{aligned}$$

and finally, the fourth moment

$$\frac{1}{F(0, T_1, T_2)^4} \left[\frac{1}{N^4} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} F(0, T_i) F(0, T_j) F(0, T_k) F(0, T_l) \Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(0), 0, t) \exp(A) \right]$$

where

$$\vec{z} = \begin{pmatrix} -i \sum_{x \in \{i, j, k, l\}} \exp(-\alpha(T_x - t)) \\ -i \sum_{x \in \{i, j, k, l\}} \exp(-\varepsilon(T_x - t)) \\ -4i \end{pmatrix}$$

and

$$\begin{aligned} A &= \sum_{x \in \{i, j, k, l\}} - \int_0^t (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\ &+ \sum_{x \in \{i, j, k, l\}} - \frac{\sigma^2}{4\varepsilon} \exp(-2\varepsilon(T_x - t)) (c_1 - c_2)^2 ((1 - \exp(\varepsilon(T_x - t))) (1 - \exp(-2\varepsilon t))) \\ &+ \sum_{x \in \{i, j, k, l\}} - \frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\varepsilon(T_x - t))) \end{aligned}$$

In a similar manner, we can now easily derive the moments for the MRVG-2 model by substituting the characteristic function for that model, removing the middle row of \vec{z} and redefining A as

$$\begin{aligned} A &= \sum_{x \in \{i, \dots\}} - \int_0^t (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\ &+ \sum_{x \in \{i, \dots\}} - \frac{\sigma^2}{4\varepsilon} \exp(-2\alpha(T_x - t)) (c_1 - c_2)^2 ((1 - \exp(\alpha(T_x - t))) (1 - \exp(-2\alpha t))) \\ &+ \sum_{x \in \{i, \dots\}} - \frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\alpha(T_x - t))) \end{aligned}$$

A similar reduction can be applied to the MRVG-2x and MRVG models such that

$$A = \sum_{x \in \{i, \dots\}} - \int_0^t (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t)$$

B.9 Valuation Convergence

	Grid Points: $N \times M \Rightarrow 2^N \times 2^M$	Value	Average Time (secs)
MRVG-3x	4×4	257.5192	6.30
	5×5	34.9837	12.51
	6×6	14.2158	37.74
	7×7	14.8167	195.49
	8×8	15.0857	928.35
	8×9	15.1390	2192.5
	9×9	15.1390	4190.8
	8×10	15.1469	4556.3
	8×11	15.1477	8688.4
MRVG-2x	4×4	52.2877	5.87
	5×5	15.7889	12.46
	6×6	12.3636	38.18
	7×7	11.9961	195.68
	8×8	11.9643	954.13
	9×9	11.9642	4192.10

Table B.9.1: Two Factor Convergence Results

Performance statistics across 5 runs for each grid size.

Appendix C

Chapter 4

C.1 Sample Parameter Covariance Matrix

Recall that the covariance matrix of log-forward returns is given by $C = WDW^{-1}$ where W represents the matrix of eigenvectors, \vec{w} , and D the diagonal matrix of eigenvalues, λ . We can therefore write C as function of the factor volatility functions such that $C = VV^T$ where

$$V = \begin{pmatrix} w_{1,1}\sqrt{\lambda_1} & \cdot & \cdot & \cdot & w_{1,M}\sqrt{\lambda_M} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ w_{M,1}\sqrt{\lambda_1} & \cdot & \cdot & \cdot & w_{M,M}\sqrt{\lambda_M} \end{pmatrix}$$

Define the log-forward returns as $df(\vec{t}, \tau) \in \mathbb{R}^M$ and let $p(df(\vec{t}, \tau); C)$ be the associated multivariate density. Assuming approximately normally distributed log-forward returns we have

$$\begin{aligned} p(df(\vec{t}, \tau); C) &= p(df(\vec{t}, \tau); V) \\ &= \left(\frac{1}{\sqrt{2\pi}}\right)^M \sqrt{|VV^T|} \exp\left(-\frac{1}{2} (df(\vec{t}, \tau) - \vec{\mu})^T VV^T (df(\vec{t}, \tau) - \vec{\mu})\right) \end{aligned}$$

with

$$\vec{\mu} = -\frac{1}{2} \begin{pmatrix} \sum_i V_{1,i}^2 \\ \cdot \\ \cdot \\ \cdot \\ \sum_i V_{M,i}^2 \end{pmatrix}$$

which we can use to estimate the inverse, Σ , of the sample Fisher Information Matrix.

We are now in a position to derive the sample covariance matrix relating to the parameters in our model. For a set of model parameters, $\theta \in \mathbb{R}^N$, we assume the existence of a function $g : \mathbb{R}^N \rightarrow \mathbb{R}^{M \times M}$ with inverse $g^{-1} \in C^1$ such that $g^{-1}(V) = \theta$. The function g^{-1} is in essence the calibration function used to match the shape of the volatility function of a specific model factor with the entries of V . The restriction that g^{-1} be differentiable is a pre-requisite for use of the delta method and may at first seem overly restrictive. However, the form of each factor's volatility curve is generally quite simple and can typically be expressed as a linear or log-linear function of the entries in V . As an example, in the case of the MRVG-3x model we could choose the following functional relationships between the model parameters and the covariance matrix

$$\begin{aligned} b &= V_{1,1} \\ c_1 &= V_{1,2} \\ \varepsilon &= -\ln\left(\frac{V_{2,2}}{V_{1,2}}\right) \frac{1}{t_2} \end{aligned}$$

The last equality is based on choosing an ε value given by the rate of decay of the volatility function between the day-ahead and month-ahead maturities.