

Gas Storage Valuation

under Multifactor Lévy Processes

Mark Cummins*

Greg Kiely[†]

Bernard Murphy[‡]

Abstract

A practical problem for energy companies is instituting a consistent framework across its supply and trading activities to deliver on all-important *P&L* and *at-Risk* reporting requirements. With a focus on storage assets and wider natural gas market exposures, we present a gas storage valuation methodology, which uniquely uses a flexible multifactor Lévy process setting that allows for consistent valuation

**Corresponding Author:* DCU Business School, Dublin City University, Dublin 9, Ireland. Email: Mark.Cummins@ul.ie. Tel: +353 1 700 8827.

[†]Gazprom Marketing & Trading Limited, 20 Triton Street, London, United Kingdom NW1 3BF. Email: Greg.Kiely@gazprom-mt.com. *The author's views are his own and do not necessarily reflect those of Gazprom Marketing & Trading Limited or any of its affiliates.*

[‡]Kemmy Business School, University of Limerick, Limerick, Ireland. Email: Bernard.Murphy@ul.ie.

and risk management reporting across a general derivative book. Our approach is capable of replicating the complex covariance structure of the natural gas forward curve and capturing time spread volatility, a key driver of extrinsic storage value, while being simultaneously capable of accurately calibrating to market traded options. We begin by extending a single factor Mean Reverting Variance Gamma process to an arbitrary number of dimensions and, by way of specific examples, show how the traditional Principal Component Analysis based view of gas forward curve dynamics can be incorporated into a primarily market based valuation. We develop in the process an innovative implied moments based calibration technique, which allows for efficient calibration of general multifactor forward curve models to delivery period options common in energy and commodity markets. Furthermore, to accommodate the forward curve and traded options market consistency, we propose an appropriate joint market based calibration and historical estimation methodology. Through a formal model specification analysis, we provide evidence that the multifactor Lévy models we propose provide a better joint fit to NBP natural gas options-forward market data, relative to comparative benchmark models. Finally, we develop a novel multidimensional fast Fourier transform based storage valuation algorithm and provide empirical evidence that the multifactor Lévy model suite is better specified to more accurately capture extrinsic value.

Keywords: gas storage valuation; multifactor Lévy processes; Mean Reverting Variance Gamma processes; implied moments calibration; fast Fourier transform.

forthcoming in the Journal of Banking and Finance

1 Introduction

As noted by Cummins et al. (2017), storage has become an increasingly prominent feature of the European natural gas markets, allowing market practitioners an effective means to manage the risk of supply distribution, to smooth out seasonal supply-demand imbalances, and to increase overall market liquidity. Natural gas storage contracts can refer to both capacity in physical storage units or “virtual” storage capacity, typically sold as a simplified tranche of physical storage. While typically of high materiality, storage assets are but one component of the overall natural gas supply and trading activities of energy companies and so should be valued within a consistent framework, to deliver on the all-important *P&L* and *at-Risk* reporting requirements of practitioners. While storage value is driven by time spread volatility rather than outright volatility, such that consistency with the options market might be argued to be a secondary consideration, 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, similar to storage assets. 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. In this paper, we therefore present

a class of finite dimensional Markov models of the entire forward curve where the underlying state variables are driven by Lévy processes, with the explicit aim of accurately capturing time spread volatility, which drives the storage value, whilst maintaining consistency with the vanilla options market.

Given a liquid forward market, traders have the ability to lock in a base or *intrinsic value* to the storage asset by locking in prices for future traded volumes, while inherent *extrinsic value* can be extracted through dynamically trading in the underlying forward and options market. Methods like *intrinsic basket of spread options*, *rolling intrinsic* and *rolling basket of spread options* capture in alternative ways some of the flexibility of the storage, but they fail to capture fully the real optionality. While there is no consensus modelling approach to the problem of gas storage valuation, the main approaches can be broadly categorised as (i) spot based and (ii) forward curve based. Spot based optimisation methods include those of Manoliu (2004), Boogert and De Jong (2008), Chen and Forsyth (2009) and Felix and Weber (2012). Despite the real option based merits of these approaches, storage traders may eschew spot based models as they neither reflect the available range of forward contracts, nor the multifactor structure of the forward curve dynamics. We therefore position our multifactor Lévy model development work to align with the forward based modelling literature referenced below.¹

¹While we cite a relevant selection of spot and forward curve based storage valuation literature, a wider literature exists, a good deal of which relies on the use of industry heuristics, such as futures based static or rolling intrinsic and options based basket of spreads trading strategies. However, recent and noteworthy research has leveraged insights and results from the extensive operations research literature (Carmona and Ludkovski 2010;

On the use of Lévy models in the literature, one of the more notable applications of single factor Lévy driven Ornstein–Uhlenbeck processes in the energy markets is given by Deng (2000). Benth et al. (2007) model electricity prices using a mean reverting Gamma process for modelling positive price spikes, while Luciano (2009) utilises time-changed Brownian motions with Lévy subordinators in the modeling of the spark spread. Attempts at incorporating multifactor forward curve models in storage valuation include Boogert and De Jong (2011), who propose modeling the spot price using a three-factor Ornstein-Uhlenbeck process, capturing spot price volatility, forward curve volatility and winter-summer spread volatility. Bjerksund et al. (2008) propose a six factor forward curve model, which aims to replicate the historical dynamics of the UK gas market to a high degree of accuracy, utilising though a rolling intrinsic valuation method, presumably, to overcome the high dimensionality of the problem. 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 optimise the storage value. Warin (2012) investigates the valuation and hedging of storage

Lai et al., 2010; Nadarajah et al., 2015). Lai et al. (2010), among others, have recognised the complexity of handling the trade-off between the multifactor model dimensionality, which is a stylised fact of natural gas spot and forward dynamics, and the intractability of stochastic dynamic programming (SDP) solution methods. Lai et al. (2010) have dealt with the problem by benchmarking these practice based heuristics, which involve the use of either deterministic dynamic programming or linear programming methods, against optimal control solutions generated by approximate stochastic dynamic programming (ADP) methods. Interestingly, Lai et al. (2010) find that the heuristic methods used in practice can perform reasonably well against their ADP benchmark valuation, that is, provided they are re-optimised over time and state.

assets under a two factor forward curve model, with the first factor chosen to represent the short-term volatility term structure and the second, the long-term volatility of the gas market. Our study is most closely aligned though to the single factor mean reverting Lévy work of Cummins et al. (2017), where we extend this framework to a multifactor setting, providing both modelling and storage valuation solutions. In so doing, we address some shortcomings in this literature. Most notable is that the extant models are inadequately specified to capture well the excess kurtosis evident in natural gas market returns, while all models are either estimated using historical data only or calibrated using options data only.

We extend forward curve modelling methods in such a way that offers dual benefits. First is the manner in which we cast our storage valuation exercise within a flexible multifactor Lévy process setting and, in contrast to much of the existing literature, demonstrate how one can construct models with the capability of capturing excess kurtosis and time spread volatility, whilst maintaining consistency with the vanilla options market. Specifically, we propose a multifactor Mean Reverting Variance Gamma (MRVG) modelling framework that is capable of being *simultaneously* forward curve consistent and calibrated to market traded options and, in a similar manner to Andersen (2010), we specify each factor to represent a latent principal component of the underlying forward curve covariance matrix. While our use of Lévy-driven processes is novel, it is consistent with previous literature (Deng, 2000; Benth et al., 2007; Bjerk Sund et al., 2008; Luciano, 2009; Warin, 2012;

Parsons, 2013). The multifactor nature of the Lévy model setting is capable of replicating the complex covariance structure of the natural gas forward curve, a key driver of extrinsic storage value. In allowing for accurate calibration to market traded options, we present the first class of multifactor storage models developed with the explicit intention of providing a link between the underlying model dynamics and the options market. To that end, we further develop in this study an innovative implied moments based calibration technique, which allows for efficient calibration of general multifactor forward curve models. We closely follow Guillaume and Schoutens (2013), who provide market based, and thus model free, valuation formulae for the square, cubic, and quartic “derivative” contracts, which can then be used to derive measures of market implied variance, skewness, and kurtosis. This approach relies upon the existence of a liquid options market on the underlying which is being modelled. While 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 in this methodology involves estimating the implied moments of the monthly delivery forward price in order to calibrate an 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 utilising this approach in model calibration, as we demonstrate, is that the moments of the forward

price can be easily derived for any reasonable forward curve model regardless of dimensionality. Further to this, to accommodate the forward curve and traded options market consistency of our Lévy models, we develop an appropriate joint market calibration-estimation approach. The former utilises our market implied moments procedure, while the latter is based on fitting the model derived volatility function to the historically estimated volatility function returned through Principle Component Analysis. Given the novelty of the proposed Lévy model suite, we conduct a formal model specification analysis exercise following an approach similar to that of Bakshi et al. (1997), whereby the daily performance of the models is evaluated using an extensive database of NBP forward and options data. For comparative purposes, we benchmark the Lévy model suite performance against the Mean Reverting Jump Diffusion (MRJD) model of Deng (2000) and the single factor MRVG model of Cummins et al. (2017), which are the closest related models available in the literature. We find that the multifactor MRVG models developed provide a better joint fit to the NBP natural gas options-forward data.

The second of the dual benefits that our work offers is our exploitation of the power of integral transform based approaches to option pricing (Heston, 1993; Duffie et al., 2000; Lewis, 2001) and hence the design a computationally efficient fast Fourier transform (FFT) based pricing methodology for gas storage valuation, drawing on the example of the financial markets literature (Carr and Madan, 1999; Andricopoulos et al., 2003; Chourdakis, 2004; O’Sullivan, 2005). Utilising the convolution method of Lord et al. (2007), combined

with the extended Fourier time-stepping method of Jaimungal and Surkov (2011), the storage value is derived by repeatedly evaluating the continuation value at each point in the state space through numerical integration. The only restriction on the associated spot price model is that it is Markov and all one needs is the conditional characteristic function (CCF) of the transition density. We proceed to derive the analytic CCFs for the multifactor Lévy-driven model suite considered. Kjaer (2008) notes that, in general, the characteristic functions of mean reverting Lévy-driven models are not available in analytical form and need to be calculated numerically, hence reducing the appeal of computationally efficient FFT methods. Thus, the MRJD model proposed by Deng (2000) is to date a special case of this model class, which has an analytical solution for the CCF. We therefore add to the availability of Lévy-driven models in the energy space with analytic CCFs, which are more likely to be adopted in practice for the computational efficiency they offer. Showcasing our proposed Lévy model set and applying our multidimensional FFT storage valuation algorithm, we use the fitted models obtained from the formal model specification analysis to appraise the performance of the models in valuing an assumed storage contract with specified physical constraints. We find the MRVG model suite offers greater flexibility in capturing extrinsic storage value relative to the benchmark MRJD model

The remainder of the paper is organised as follows. Section 2 sets out the development of the Lévy driven forward curve model suite. Section 3 presents the joint calibration-estimation approach taken, based respectively

on market implied moments calibration and historical volatility function estimation. Given the novelty of the multifactor MRVG models proposed, we conduct a formal model specification analysis of the competing models in Section 4 to ascertain the improvement in capturing natural gas market dynamics, informed by futures and options markets. Section 5 presents the innovative multidimensional FFT based valuation algorithm developed to price early exercise claims, such as storage contracts. Section 5 further presents the results of an extensive storage valuation exercise, leveraging the model specification analysis of Section 4, emphasising the merits of our proposed Lévy-driven forward curve models in accurately capturing extrinsic value in particular. Section 6 provides concluding remarks.

2 Lévy Storage Model Development

A major development within the natural gas storage valuation literature in recent years has been the need to accurately capture the covariance structure of the commodity forward curve. The price return dependency between contracts will fully determine the extrinsic value accruing to the owner of a storage asset. We therefore extend the forward curve modelling literature by means of casting our storage valuation exercise within a flexible multifactor Lévy process setting, presenting a class of finite dimensional Markov models of the entire forward curve where the underlying state variables are driven by Lévy processes, and which are sufficiently specified to accurately capture

time spread volatility, whilst maintaining consistency with the vanilla options market. This contrasts to the recent storage valuation literature, see Boogert and De Jong (2011) and Bjerksund et al. (2008). The business case for pricing these risks within a consistent framework has been set out in the opening of Section 1.

The general modelling framework we utilise in constructing our Lévy models mirrors and generalises that of Andersen (2010). Each model can be viewed as a generalisation of the single factor MRVG model of Cummins et al. (2017) and collectively they represent a unique family of models designed to reflect the rich dynamics of the forward market while allowing for accurate calibration to the option implied volatility surface. In order to price storage assets, and other derivative contracts, we furthermore derive the forward curve consistent characteristic functions associated with our proposed forward curve models, which allows us to later exploit computationally efficient transform based pricing methods within a dynamic programming setting. Spot-based optimal control solutions generated by dynamic programming methods, such as ours, will always result in initial storage valuations *at least* as high as forward or option based trading strategies, such as rolling-intrinsic or even the dynamic basket of spread options approach, given a consistent market model. In order to be internally consistent with forward-based strategies, the corresponding spot price model must have a rich multifactor specification and be capable of efficient calibration to both the market forward curve and volatility term structures. Our methodology being premised on

a multifactor forward curve modelling approach achieves this, and allows traders to avoid locking in or financially contracting the storage asset in the forward market, in order to keep alive the possibility of exercising into more profitable forward spreads in the future. Furthermore, our combined calibration-estimation approach enables the efficient estimation, and isolation of the corresponding risk exposure, of the principal component volatility function, which is the main driver of twist movements in the forward curve, and hence allows traders to monetise the real optionality embedded in the gas storage contract in an equivalent fashion to that implemented in a spot price only based optimisation tool.

2.1 Single Factor Lévy Model

To begin, we provide an introduction to the Lévy framework and the single factor MRVG model of Cummins et al. (2017), which we extend to higher dimensions. Price discontinuities and mean reversion are accepted realities of energy and commodity markets (Nomikos and Andriosopoulos, 2012; Maslyuka et al., 2013). In the context of storage contracts, spikes, both positive and negative, 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. Mean reversion has the effect of increasing the spread variance between different points on the forward curve depending upon their relative maturities, implying greater time spread variability and hence extrinsic value of the storage contract. With these stylised market fea-

tures in mind, Cummins et al. (2017) develop a single factor MRVG model for the purposes of storage valuation.

A Lévy process is a right-continuous stochastic process possessing stationary and independent increments characterised 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 is the diffusion volatility, and $\mu(dj)$ is a Lévy measure, satisfying $\int_{\mathbb{R} \setminus 0} \min(j^2, 1) \mu(dj) < \infty$, which measures the relative frequency of different jump sizes j . A common method for constructing Lévy processes is via subordination of simpler processes such as Brownian motion. Intuitively, the subordinating process controls the stochastic passage of time in the subordinated process and can be used to induce heavy tails. As such, a subordinator can be viewed as a ‘building-block’ Lévy process, with a corresponding Lévy triplet in which the Brownian part does not exist. The process has a non-negative drift and the process is non-decreasing in the sense that the Lévy measure which describes how jumps, and hence heavy tails, occur is zero on the negative half-line, i.e. positive-only increments with the additional $\mathbb{R} \setminus \{0\}$ requirement that the corresponding Lévy density, if it exists, must have zero mass at the origin. We refer the interested reader to Sato (2001) and Barndorff-Nielsen et al. (2012) for a thorough introduction to the theory of Lévy processes, while for commodity market applications, see Li and Linetsky (2014).

For our purposes, we are interested in utilising Lévy driven Ornstein-Uhlenbeck processes in modelling the natural gas spot price. We define $x(t)$

as the log spot price, governed by

$$dx(t) = [\alpha(\omega - x(t))] dt + dJ(t) \quad (1)$$

where $J(t)$ is a pure jump process defined such that

$$J(t) = \int_0^t \int_{\mathbb{R} \setminus 0} j(v(dj, dc) - \mu(j) dj) dc$$

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. Assume a Variance Gamma process $\{X(t), (\theta, \sigma, v)\}$, which is an infinite activity,² finite variation pure jump process constructed as a time-changed Brownian Motion with activity rate driven by a Gamma process, such that $X(t) = \theta\Gamma(t; 1, v) + \sigma W(\Gamma(t; 1, v))$. It follows that :

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

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 v . The skewness of the distribution is controlled by θ . However, the implied volatility smile is generally symmetric in log-strike and as such

²A 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.

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 parameterised by (α, σ, v) , where α is the mean reversion rate, σ is the process volatility, and v controls the variance of the jump magnitudes.

The solution to Eq. (1) is then

$$\begin{aligned} x(t) &= x(s) \exp(-\alpha(t-s)) + \omega(1 - \exp(-\alpha(t-s))) \\ &\quad + \int_s^t \mathbb{R}_{\setminus 0} j \exp(-\alpha(t-c)) (v(dj, dc) - \mu(j) djdc) \end{aligned}$$

It follows that the CCF of the log spot price $x(t)$ driven by the MRVG process, $\Phi_{x(t)}(z; x(s), s, t)$, is

$$\begin{aligned} &\exp(izx(s) \exp(-\alpha(t-s)) + iz\omega(1 - \exp(-\alpha(t-s)))) \times \\ &\quad \exp\left(\int_s^t \varphi(z \exp(-\alpha(t-c))) dc\right) \end{aligned}$$

where $\varphi(z)$ is the characteristic exponent of the Variance Gamma process (with $\theta = 0$)

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

Solving the integral in the exponent gives

$$\begin{aligned} &\exp(izx(s) \exp(-\alpha(t-s))) \quad (3) \\ &\quad \times \exp(iz\omega(1 - \exp(-\alpha(t-s)))) \exp(A(z, s, t)) \end{aligned}$$

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 Eq. (3), 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 $|w| < \sqrt{\frac{2}{\sigma^2 v}}$.

When constructing a market model it is obviously crucial that the model is capable of reproducing the prices of liquid traded instruments, particularly those that are likely to form part of a hedging portfolio. At the very least, the model should be consistent with the observed market forward curve. To this end, an extended forward curve consistent single factor MRVG model is developed by Cummins et al. (2017). It is readily shown that process for the log spot price becomes:

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

which is based off the assumption of a general Lévy driven model of the log forward curve $f(t, T)$, and where $\bar{v} \equiv v(dj, ds) - \mu(dj)$ is the compensated jump measure and $\kappa_j(\cdot)$ is the cumulant function of the jump process. The

CCF of the log-price, $\Phi_{x(t)}(z; x(s), s, t)$, can then be shown to be

$$\begin{aligned} & \exp\left(izx(s)\exp(-\alpha(t-s)) + iz \int_s^t \omega(c)\exp(-\alpha(t-c))dc\right) \\ & \times \exp\left(\int_s^t \varphi(iz\exp(-\alpha(t-c)))dc\right) \end{aligned}$$

where $\omega(t)$ is introduced to denote the time-dependent drift of the process $dx(t)$. Cummins et al. (2017) proceed to show how Fourier methods can be used to price early exercise claims and consider in particular the problem of valuing storage contracts.

2.2 Multifactor Lévy Model

While the single factor Lévy model of Cummins et al. (2017) is effective and efficient in its approach to storage valuation, it has limitations in its one dimensional construction. In this section, we extend the work of Cummins et al. (2017) and present a general multifactor Lévy model that offers greater flexibility in capturing the dynamics of the natural gas markets. We begin by specifying the dynamics of the log-forward price $f(t, T)$:

$$\begin{aligned} df(t, T) = & \left(-\frac{1}{2}\beta(t, T)\beta(t, T)' - \kappa_{\vec{j}}(\gamma(t, T)) \right) dt \\ & + \beta(t, T)d\vec{W}(t) + \int_{\mathbb{R}^{\kappa \setminus 0}} \gamma(t, T)\vec{j}\vec{v}(dj; dt) \end{aligned} \quad (4)$$

where $\vec{W}(t)$ is a Brownian Motion of dimension $(K \times 1)$, $J(\vec{t}) \equiv \int_{\mathbb{R}^K \setminus 0} \vec{j} \vec{v}(dj; dt)$ is a pure jump process of dimension $(K \times 1)$ with Poisson random measure \vec{v} , and $\kappa_{\vec{j}}(\cdot)$ is the cumulant function of the jump process. The functions $\beta(t, T)$ and $\gamma(t, T)$, of dimension $(1 \times K)$, are respectively the sensitivity of the log-return of the T -maturity forward price to the diffusion and pure-jump stochastic drivers of forward curve returns. Given \mathcal{F}_t , the filtration defined on the probability space $(\Omega, \mathcal{F}, \mathbb{Q})$, if the \mathcal{F}_t -measurable continuous differentiable functions $\beta(t, T)$ and $\gamma(t, T)$ can be separated in terms of time and maturity, such that

$$\beta(t, T) = \varsigma(T) \eta(t) \tag{5}$$

$$\gamma(t, T) = \varsigma(T) \theta(t) \tag{6}$$

with $\eta(t)$ and $\theta(t)$ both mapping $\mathbb{R} \rightarrow \mathbb{R}^{K \times K}$ and $\varsigma(T)$ mapping $\mathbb{R} \rightarrow \mathbb{R}^K$, then the evolution of the entire forward curve can be represented via a Markov dynamical system of K state variables. In this case, we can write Eq. (4) as

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

where

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

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. While imposing dependence amongst Lévy processes is well covered in the literature (Cont and Tankov, 2004; Marfè, 2009; Luciano and Semeraro, 2010), 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 allows us to specify the effect of each state variable on the forward curve dynamics independently, which is in agreement with traditional Principal Component Analysis (PCA) of forward curve movements (Clewlow and Strickland, 1999, Carmona and Coulon, 2014).

For the purposes of valuing storage, we next need to derive the forward curve consistent dynamics of the log spot price process, $x(t) \equiv f(t, t)$,

$$x(t) = f(0, t) + \sum_{k=1}^K y^{(k)}(t) \tag{7}$$

where $y^{(k)}(t) \equiv y^{(k)}(t, t)$. This is crucial for the purposes of storage valuation

as it links the general forward curve model above with the implied spot price model used in the spot optimisation valuation approach. This pattern, of first introducing the forward curve model and then the implied spot price model, will therefore be repeated for each specific model introduced in this section. The spot factor evolution is in general given by the following, where we explicitly use the time-maturity separated terms of Eqs. (5)-(6):

$$\begin{aligned}
dy^{(k)}(t) &= \left[-\frac{1}{2} \eta^{(k)}(t)^2 \varsigma^{(k)}(t)^2 - \varsigma^{(k)'}(t) \int_0^t \eta^{(k)}(s)^2 \varsigma^{(k)}(t) ds \right. \\
&\quad \left. - \kappa_{j^{(k)}}(\theta^{(k)}(t) \varsigma^{(k)}(t)) - \varsigma^{(k)'}(t) \int_0^t \theta^{(k)}(s) \kappa_{j^{(k)}}'(\theta^{(k)}(s) \varsigma^{(k)}(t)) ds \right. \\
&\quad \left. + (y^{(k)}(t) + \int_0^t \kappa_{j^{(k)}}(\theta^{(k)}(s) \varsigma^{(k)}(t)) ds) \frac{\varsigma^{(k)'}(t)}{\varsigma^{(k)}(t)} \right] dt \\
&\quad + \eta(t) \varsigma(t) dW^{(k)}(t) + \mathbb{R}/_0 \theta^{(k)}(t) \varsigma^{(k)}(t) j \bar{v}^{(k)}(dj, dt) \\
&= \left(\omega^{(k)}(t) + \frac{\varsigma^{(k)'}(t)}{\varsigma^{(k)}(t)} y^{(k)}(t) \right) dt + \eta(t) \varsigma(t) dW^{(k)}(t) \\
&\quad + \mathbb{R}/_0 \theta^{(k)}(t) \varsigma^{(k)}(t) j \bar{v}^{(k)}(dj, dt)
\end{aligned}$$

where $\omega^{(k)}(t)$ is introduced to capture the collected terms of the drift. This calculation of the drift adjustment, $\omega^{(k)}(t)$, allows one to define the characteristic function of the implied spot price model and thus utilise Fourier transform based methods.

To determine an appropriate specification for the functions $\beta^{(k)}(t, T)$ and $\gamma^{(k)}(t, T)$, standard PCA shows that the first two principal components describe the usual “shift” and “twist” features of energy forward curve dynamics respectively, with a high degree of explanatory power. Including the

third principal component, which describes the usual “bend” feature, is possible but its contribution to total variation in the forward curve is found to be marginal.³ The PCA based decomposition of the forward curve returns naturally implies a model of the log forward price equal to the sum of the independent factors, as given by Eq. (7). 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. Experimentally, we find that the first volatility function, which accounts for 90% of the total variance (see Section 4), resembles the negative exponential volatility function associated with the traditional single factor model. Thus, we define $\gamma^{(1)}(t, T) \equiv b \exp(-\alpha(T - t))$, where b is a scaling parameter and α controls the decay of instantaneous volatility with respect to relative maturity. As in the single factor case, with the objective being to capture this primary source of forward curve variation, whilst also allowing one to accurately calibrate to the market for monthly options, the background stochastic driver of the model we chose is again the Variance Gamma process parameterised by Madan et al. (1998) as a time-changed Brownian Motion with the parameter triplet (θ, φ, v) . θ and φ are again the drift and volatility of the Brownian

³Furthermore, there are a number of practical considerations that require a trade off to be made. Most notably, including this third principal component within one’s price model, increases the computational effort needed for pricing complex derivatives, like gas storage contracts, considerably given the higher dimensionality. For this reason, we judiciously select to proceed with modelling just the first two principal components, recognising that the second “twist” component in particular is a key driver of storage value, capturing the movement of relative maturity spreads.

Motion respectively and v is the jump variance of the subordinating Gamma process. Again, for our purposes we explicitly set $\theta = 0$ and $\varphi = 1$, such that we have the process $\{dX(t); (0, 1, v)\}$.

The shape of the second volatility function, which accounts for a further 4.81% of the curve variability (see Section 4), 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 volatility function can be approximated as negative exponential tending to a negative asymptote (Andersen, 2010; Cheyette, 2001). We therefore set $\beta^{(2)}(t, T) \equiv \exp(-\epsilon(T-t))(c_1 - c_2) + c_2$, with $\epsilon > 0$, c_1 the sensitivity of the spot maturity ($T \rightarrow t$) to the second principal component, and c_2 the asymptotic sensitivity of the forward curve returns to the second principal component as $T \rightarrow \infty$. We have chosen to model this secondary source of variation using a diffusion process: $dW(t)$. This will enable us to approximate the market smile for vanilla options whilst also enforcing an inter-maturity covariance structure on the model.

With these volatility functions specified, it is possible for us to now set out the suite of multifactor MRVG models that we consider in our study. The overarching multifactor MRVG model is given by the following system:

$$dy^{(1)}(t, T) = (-\kappa_{j^{(1)}}(\gamma^{(1)}(t, T)\sigma(t))dt + \gamma^{(1)}(t, T)\sigma(t)dX(t) \quad (8)$$

$$dy^{(2)}(t, T) = -\frac{1}{2}(\beta^{(2)}(t, T)\sigma(t))^2dt + \beta^{(2)}(t, T)\sigma(t)dW(t) \quad (9)$$

with $E [dX_t dW_t] = 0$. 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 Eqs. (5)-(6).

The following sections outline the specific MRVG models we consider, while we refer to the pros and cons of each.

MRVG-3 Model Specification

We specify an equivalent three-dimensional SDE system representation of the overarching MRVG model of the log-forward price defined by Eqs. (8)-(9), which we label MRVG-3. This is a theoretical model specification that due to substantial computational constraints is impractical to implement but which we reduce in dimensionality in the next section to a model specification that may be practically implemented but which retains much of the attractiveness of the MRVG-3 model. We discuss the MRVG-3 model here to highlight these attractive modelling features. In order to ensure the model will be Markovian, we must first re-define the second volatility function, $\beta^{(2)}(t, T) \equiv \exp(-\epsilon(T-t))(c_1 - c_2)$, and add a third volatility function $\beta^{(3)}(t, T) \equiv c_2$. The dynamics of the MRVG-3 model are then given by Eqs. (8)-(9) augmented with a third factor, $dy^{(3)}(t, T) = \beta^{(3)}(t, T) \sigma(t) dW(t)$.

The associated spot price dynamics are then given by

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

where the functions $\omega^{(i)}(t), i = 1, 2$, are given in Appendix A. The three factors are required in order to obtain a Markovian log spot price representation. Note that if the percentage of spot variance explained by the two principal components is given by s then we have $b = \sqrt{s - c_1^2}$. The CCF that would be used for pricing is given in Appendix B. The main benefit of this model is that the factor specification allows us to match the typical shape of the returns' sensitivity to the first two principal components, i.e. the “shift” and “twist” features, to a high degree of accuracy. The model parameters may be estimated using a joint calibration-estimation procedure on current market option prices and historical price returns. The main drawback though is the impractical computational burden incurred when valuing storage. Indeed, the implementation of the FFT storage valuation algorithm devised in Section 5.1 would in effect be impossible in this three dimensional setting as the memory requirement would exceed that of any hardware available currently.⁴ We therefore proceed to the next model specification, which

⁴A two-dimensional FFT is implemented as part of the storage valuation algorithm (see Section 5.1) for the two-factor MRVG model specifications to be presented in the forthcoming sections. For a single implementation of the storage valuation algorithm under a given two-factor MRVG model, we find experimentally that the peak memory requirement

allows a critical reduction in dimensionality.

MRVG-3x Model Specification

We can reduce the dimensionality of the MRVG-3 model by explicitly setting $c_2 = 0$, in which case we use the label MRVG-3x.⁵ The model retains many of the benefits of the MRVG-3 model, in particular the ability to match the typical shape of the returns' sensitivity to the first two principal components to a high degree of accuracy and the Markovian nature of the log spot price process. What is different is that the model approximates the second volatility function as a negative exponential tending to zero rather than the negative volatility asymptote under the MRVG-3 model. This is a misrepresentation of the full nature of the second volatility function, which leads to a decrease in the variability of the spread between short and long maturity forward prices compared to the MRVG-3 model. A positive shock to this factor

is in the order of $\sim 600\text{MB}$ using a grid size of 256×256 for the FFT. The memory usage is recorded by monitoring the memory display provided by Microsoft Windows' Task Manager, observed over a number of single instance runs of the storage valuation algorithm. For the MRVG-3 model, we would need to move to a three-dimensional FFT, which if we use a similarly sized third dimension, would require a grid size of $256 \times 256 \times 256$. This would scale the peak memory requirement to $\sim 600\text{MB} \times 256 = \sim 153\text{GB}$ of memory, which is beyond the maximum memory capacity of currently available hardware. The largest memory availability on commercial hardware is, for instance, 128GB. The 256×256 grid size used in this experiment is much coarser than the grid sizes used here for the valuations; for instance, under the MRVG-2x and MRVG-3x model specifications, we set the grid sized to be 512×512 and 256×2048 respectively, in order to increase the accuracy of the storage valuations. Hence, the $\sim 153\text{GB}$ memory requirement calculated here is likely a conservative lower bound of what would be required.

⁵A note on the naming convention used here. Although the model clearly has two-factors, we have chosen "MRVG-3x" so as to identify the model as a derivative of the "MRVG-3" model with one parameter constrained. We also use this labelling to distinguish it from the alternative two factor MRVG-2, which we define later.

would typically lead to a positive return on the one-day maturity contract and a negative return on the, for example, 200-day maturity contract, thus widening the price spread between the two. The MRVG-3x model however, would imply an approximately zero return on the 200 day maturity contract thereby underestimating the price spread variability. However, the reduction in dimensionality is justified by the significant decrease in the computational burden when using the model to value complex derivative assets, such as in the natural gas storage case here.

MRVG-2 Model Specification

As an alternative to the MRVG-3/MRVG-3x modelling approach, it is possible to simplify the overarching multifactor MRVG model by means of setting $\epsilon = \alpha$ in the general setting, so that the decay of both the first and second volatility functions are equal. The model simplifies to a two factor system, which we label MRVG-2. The associated spot price dynamics are then

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

which are Markovian in nature as desired. The function $\omega^{(1)}(t)$ is given in Appendix A and the associated CCF in Appendix B. As with the MRVG-3x model, the two factor dimensionality of the MRVG-2 model is computation-

ally efficient relative to the three factor MRVG-3 model, at the expense of constraining both factors to share a single mean reversion rate. This latter feature constrains the MRVG-2 model in its approximation of the returns' sensitivity to the first two principal components relative to the MRVG-3x model.

MRVG-2x Model Specification

We further simplify the MRVG-2 model. For this, we constrain the parameters b , c_1 and c_2 such that $b = 1$, $c_1 = c_2 = \frac{\sigma_L(t)}{\sigma(t)}$, where $\sigma_L(t)$ is defined as the long-term volatility of the spot price. This is a model we label MRVG-2x. Unlike all models specified thus far, this two factor model is sufficiently simplified that it can be calibrated to the options market alone and as such should be viewed as a closely related extension of the single factor MRVG model of Cummins et al. (2017).

3 Joint Model Calibration-Estimation

The main advantage of the model suite developed in the previous section is the ability to accurately capture time spread volatility through replicating as much as possible the complex covariance structure of the natural gas forward curve, while maintaining consistency with the vanilla options market. To accommodate such forward curve consistency and calibration to market traded options, we require an appropriate joint model calibration-estimation

approach. We first discuss calibrating the forward curve models to the market for natural gas options, and subsequently outline the historical estimation approach adopted for the MRVG-2 and MRVG-3x multifactor models. For calibration purposes, we focus on monthly delivery period options, although the proposed calibration methodology can easily be applied to swaptions of varying delivery periods. As such, this approach can be seen as a generalisation of the methodology of Guillaume and Schoutens (2013), where the authors focus on single day delivery options. The proposed approach is appealing in our context as the associated computational effort is independent of the number of dimensions in the underlying model, facilitating our multifactor setting.

We propose a variation on the implied moments technique of Guillaume and Schoutens (2013). Following the formulation of Bakshi and Madan (2000), the expectation of any twice differentiable payoff function on a price F_T can be expressed as

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

where $C(F_t, \cdot, T)$ and $P(F_t, \cdot, T)$ respectively denote call and put prices. Defining the payoff function as $v(F_T) = \ln\left(\frac{F_T}{F_t}\right)^n$ and setting $\kappa = K_0$, the first listed strike below the current forward price F_t , the formula for the n^{th}

market moment is then given by Guillaume and Schoutens (2013) as

$$\begin{aligned}
E \left[\ln \left(\frac{F_T}{F_t} \right)^n \right] &= \ln \left(\frac{K_0}{F_t} \right)^n + n \ln \left(\frac{K_0}{F_t} \right)^{n-1} \left(\frac{F_t}{K_0} - 1 \right) \\
&+ \int_{K_0}^{\infty} \frac{n}{K^2} \left[(n-1) \ln \left(\frac{K}{F_t} \right)^{n-2} - \ln \left(\frac{K}{F_t} \right)^{n-1} \right] \times \\
&C(F_t, K, T) dK \\
&+ \int_0^{K_0} \frac{n}{K^2} \left[(n-1) \ln \left(\frac{K}{F_t} \right)^{n-2} - \ln \left(\frac{K}{F_t} \right)^{n-1} \right] \times \\
&P(F_t, K, T) dK \tag{10}
\end{aligned}$$

Discretising using a trapezoidal rule gives

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

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

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

and ΔK_i is the option spacing defined as

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

Although it would be desirable to derive the implied moments using only market quoted options, in practice there may be valid reasons for interpolating a volatility surface from the existing quotes or a subset thereof; for example, liquidity and arbitrage concerns. Furthermore, 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. In our implementation, we take as input the quoted NBP monthly volatility surface on a given date and begin by interpolating/extrapolating the volatility surface over a set of equally spaced strikes. We chose linear interpolation of implied variance to fill in the missing points within the original surface and linear extrapolation for points outside the original surface.

In the context of natural gas markets, where option prices are generally most liquid on monthly delivery period contracts, 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

$$E[F(t, T_1, T_2)^n] = E\left[\left(\sum_{j=0}^{N-1} \frac{1}{N} F(t, T_1 + j\Delta t)\right)^n\right] \quad (11)$$

where t is the expiry date of the monthly option delivering over N periods between times T_1 and T_2 . The right hand side can be implied from the options market using a simplified version of Eq. (10), which once scaled by the n^{th} power of the underlying forward price and discretised as before becomes:

$$\begin{aligned} \frac{1}{F_0^n} E[F^n] &= \frac{K_0^n}{F_0^n} + n \frac{K_0^{n-1}}{F_0^n} (F_t - K_0) + \\ &\quad \frac{1}{F_0^n} \sum_{i=1}^M \Delta K_i n(n-1) K_i^{n-2} Q(K_i) \end{aligned} \quad (12)$$

For each of the MRVG forward curve models developed in Section 2, we assume that the spot volatility parameter is independent of time, i.e. $\sigma(t) = \sigma$, when deriving the implied spot price dynamics rather than assuming a particular time dependent functional form. This is done primarily to minimise the need for numerical integration when evaluating the CCF of a given model, which could add an additional source of error to our storage valuation algorithm. 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 2nd through to the 4th moments. The general formulae are presented here, while the exact formulae for the MRVG model suite are presented in Appendix C:

- second moment M_2

$$\frac{1}{F(0, T_1, T_2)^2} E \left[\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} F(0, T_1 + i\Delta t) F(0, T_1 + j\Delta t) \times \exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t)) | y(\vec{0}) \right];$$

- third moment M_3

$$\frac{1}{F(0, T_1, T_2)^3} E \left[\frac{1}{N^3} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} F(0, T_1 + i\Delta t) F(0, T_1 + j\Delta t) F(0, T_1 + k\Delta t) \times \exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t) + y(t, T_1 + k\Delta t)) | y(0) \right];$$

- fourth moment M_4

$$\frac{1}{F(0, T_1, T_2)^4} \times$$

$$E \left[\frac{1}{N^4} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} F(0, T_1 + i\Delta t) F(0, T_1 + j\Delta t) F(0, T_1 + k\Delta t) F(0, T_1 + l\Delta t) \times \right.$$

$$\left. \exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t) + y(t, T_1 + k\Delta t) + y(t, T_1 + l\Delta t)) |y(\vec{0}) \right]$$

Unlike the approach specified in Guillaume and Schoutens (2013) for simple payoffs, we 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. To obtain optimal calibrated parameters for a given model, we use the simplex algorithm in order to minimize the objective function in question; namely, the system equating the market and model moments:

$$Variance^{Market}(\cdot) = Variance^{Model}(\cdot)$$

$$Skewness^{Market}(\cdot) = Skewness^{Model}(\cdot)$$

$$Kurtosis^{Market}(\cdot) = Kurtosis^{Model}(\cdot)$$

In the case of the single factor MRVG-2x model, all parameters are calibrated as per the procedure just described. However, given the lack of a liquid market for time spread options, we need a methodology which combines historical and market data in order to estimate the parameters of the MRVG-2 and MRVG-3x models. Central to this approach is the idea outlined in Section 2 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. The historical estimation approach can be enumerated as follows: (i) construct the covariance matrix of historical forward curve relative maturity returns; (ii) scale the returns by the historical spot price volatility; (iii) perform an eigenvalue decomposition on the covariance matrix to retrieve the eigenvectors or “volatility functions”; (iv) estimate the relevant model parameters by minimizing the sum of squared differences between the model assumed and historically estimated volatility functions.

4 Model Specification Analysis

To evaluate the novel multifactor MRVG model suite proposed, we first conduct a formal model specification analysis, following closely the pricing performance approach of Bakshi et al. (1997). We collect from Bloomberg daily calendar maturity NBP natural gas futures and delivery month options data over the sample period 5th January 2015 – 24th December 2015. The options data, which comprises a sample of 253 trade dates, is used to conduct daily implied moment calibrations of the model suite and evaluate pricing performance. The overnight index swap (OIS) curve is used for derivatives discounting purposes, with its preferred use over LIBOR recommended by, for example, Hull and White (2012). The twelfth and sixth month expiry options

contracts are used in our daily calibrations, spanning the following moneyness levels, defined relative to the underlying futures price: 110%, 105%, 102.5%, 100%, 97.5%, 95%, and 90%. For historical estimation purposes, constant maturity NBP natural gas futures quotes, derived from Bloomberg off the calendar maturity NBP natural gas futures quotes, are used for the PCA based volatility function fitting. The constant maturity futures curve extends from month one (M1) through to month twelve (M12), sampled over the period 3rd January 2012 – 31st December 2014. For our model specification analysis, we fix this period and estimate the historical parameters of the multifactor MRVG-2 and MRVG-3x models once, holding these parameters fixed throughout the daily calibration period. For robustness, however, we performed the model specification analysis again but extended recursively the historical constant maturity futures data window on a daily basis, re-estimating the historical parameter estimates. This robustness check confirmed the findings we report below.

The daily calibration exercise is conducted for each of the MRVG models. For comparative purposes, we consider two appropriate benchmark models from the existing literature; namely, the single factor MRVG model of Cummins et al. (2017) and the single factor MRJD model of Deng (2000). The former model is of course the basis of the extended multifactor Lévy work conducted here and so emerges as a natural benchmark. The latter model is driven by a standard Brownian Motion diffusion component and a compound Poisson jump component described by a symmetric double exponential

distribution. The presence of the mean reversion parameter is the primary driver of the storage value in the MRJD model, controlling the forward curve covariance structure. The jump diffusion specification is what allows us to replicate the volatility smile. While not previously applied to the problem of gas storage, we choose the MRJD model as a second benchmark as it is the only other Lévy driven mean-reverting model in the literature with known characteristic function and so represents a natural comparator for our multifactor Lévy models.

Following the general approach of Bakshi et al. (1997), in-sample pricing performance is evaluated using a mean square error (MSE) criterion between market observed option prices and model derived option prices. Each of the MRVG models is fitted to the options data using the modified implied moment matching calibration technique set out in Section 3 and an estimated MSE is returned on each day. A constrained optimisation routine is utilised, with details of the constraints provided in Table 1. While the in-sample pricing performance analysis reveals important insights into the ability of the models to fit the options data and allows for informal ranking of models based on MSE, there is an inherent bias. As noted by Bakshi et al. (1997), a proposed model may outperform another in terms of in-sample fit simply by virtue of having more structural parameters than the other. However, this feature may lead to a penalisation or over-fitting when tested out-of-sample, hence resulting in poorer out-of-sample performance of this model relative to the other. We therefore perform an out-of-sample pricing performance

analysis to check for such bias. Christoffersen and Jacobs (2004) provide evidence that aligning the loss function used for the estimation and evaluation stages in the testing of alternative option pricing models is of critical importance. Incorrect judgments may be made about the out-of-sample performance of competing models in the case where loss function specifications used for estimation and evaluation are inconsistent. For our out-of-sample analysis, we therefore use the MSE criterion to evaluate the MRVG model specifications, whereby the MSE is calculated on any given date t using the model parameter estimates at time $(t - 1)$ as inputs, while all other inputs are observed at time t .

Table 1 summarises the results of the model specification analysis. Calibrated parameters are designated by (c) and historically estimated parameters are designated by (h). For the calibrated parameters, the averages of the daily calibration estimates are reported. Average MSEs are also reported for both the in-sample and out-of-sample analysis, where it can be seen that there is no evidence of out-of-sample penalisation based on differences in the number of structural parameters between models. The average MSEs confirm, at least informally, that the two factor MRVG-3x model provides the best joint fit to the NBP natural gas options-forward data, followed by the MRVG-2x model. Both models outperform the benchmark MRVG and MRJD models. Most notably though, the MRVG-2 model fails to outperform even the benchmark MRJD model, showing an inability to match the quoted options data sufficiently and reasonably. To understand this poor perfor-

mance, we note that the MRVG-2 model is the only specification where the model term structure of option volatility is entirely determined through historical parameter estimation, i.e. through the α parameter, which is basically used to match the term structure of implied volatility. It is unreasonable to expect therefore that this model would be able to replicate the term structure of volatility because of the use of historical data to estimate the parameters. This example demonstrates the complexity in the relationship between model specification and parameter estimation strategy and the potential pitfalls in the practical application of a model. This ultimately leads to an excessively high mean reversion α estimate for the MRVG-2 model. On this basis, we proceed to drop the MRVG-2 model from the analysis and utilise the MRVG-2x and MRVG-3x models in pricing the assumed storage deal set out in the next section.

When comparing the MRVG-2x and MRVG-3x models, it is notable that the estimated α for the latter model is considerably lower; $\sim 40\%$ that of the MRVG-2x estimate. In the case of the MRVG-3x model, the ϵ parameter serves to take some of the burden off the α parameter, with the two parameters working together to better capture the covariance structure of the underlying forward curve data - see the discussion to follow. This is supported by the descriptive statistics reported in Table 2 where it can be seen that across the upper percentiles of the α implied parameter series, the MRVG-2x estimates are higher, and often substantially higher, than the corresponding MRVG-3x estimates.

To complement the out-of-sample analysis above, and to consider further the penalisation of models on the basis of increased numbers of parameters, we move next to calculate AIC numbers for the model suite. The issue of course with option pricing model calibration is that the error distribution is unknown. To assign AIC numbers to the calibrated models therefore requires an assumed distribution to be imposed. We follow the model risk study of Detering and Packham (2016) and impose a normal distribution on the MSEs, allowing us to calculate model AIC numbers on each day as follows:

$$AIC = I [1 + \ln(2\pi) + \ln(MSE)] + 2(K + 1),$$

where I denotes the number of options used in the calculation of the MSE and K is the number of model parameters. Detering and Packham (2016) prove that in this setting, a (quasi-)maximum likelihood estimator (MLE) is equivalent to minimising the MSE corresponding to a particular model family, and as such, the MLE has the same properties as the MSE. In this context, the AIC provides a valid approach to model ranking. Table 1 presents the AIC results. Notably, while the out-of-sample pricing shows no evidence of penalisation (relative to the in-sample pricing) on the basis of differences in the numbers of structural parameters between models, the AIC differs somewhat in its penalisation. In particular, the MRVG-2x model is penalised on the basis of having one additional parameter relative to the MRVG model, with the AIC ranking the latter model ahead of the former model. The

AIC measure however aligns with the in-sample and out-of-sample MSEs in ranking the MRVG-3x model as the top model.

While the analysis so far gives informal insights into the relative ranking of the MRVG models, it is useful to consider the economic significance of the model fits reported. We augment the analysis with a number of additional layers of analysis. Table 1 reports (i) average mean absolute errors (MAEs) in price terms in order to give a monetary interpretation to the pricing performance and (ii) average MAEs in implied volatility terms in order to give a volatility interpretation to the pricing performance. For all of the MRVG models (excluding the MRVG-2 model), the absolute error between market and model option prices is in the order of 19p on average, while the implied volatility error is in the order of 1.4% on average. So the MRVG, MRVG-2x and MRVG-3x models appear to fit the options data comparably well. Recognising that model performance may of course vary across the moneyness-maturity dimensions, Table 3 expands on the MAE analysis. For the most part, the MRVG-3x model can be seen to provide the better fit across the two MAE measures, although there are areas of the moneyness-maturity space where the best performing model alternates to either the MRVG or MRVG-2x model.

The evidence thus far on model performance and ranking is of course informal and suggests that the MRVG-3x model provides the best fit to the options data, albeit that the MRVG models (with the exception of the MRVG-2 model) appear to perform comparably well. Additionally, there are

natural concerns over the normality assumption imposed on the MSEs in the AIC analysis – Jarque-Bera tests, for instance, confirm the MSE series to be non-normal under all model specifications. We therefore seek to formalise the model comparison. To this end, we follow the approach of Huang and Wu (2004) and implement t-testing on the in-sample MSEs to confirm whether or not there is a statistically significant difference in the MRVG models.⁶ The t-statistic is specified as follows on a pairwise model comparison basis:

$$\text{t-statistic} = \frac{M\bar{S}E^i - M\bar{S}E^j}{\text{stdev}(MSE_t^i - MSE_t^j) / \sqrt{T}},$$

where $M\bar{S}E^i$ and $M\bar{S}E^j$ are the mean MSEs for models i and j respectively, $\text{stdev}()$ is the standard deviation function, and $T = 253$ is the number of days in the estimation period. Table 4 presents the t-statistics for the pairwise MRVG model comparisons (excluding the MRVG-2 model), where it can be seen quite clearly that there is no evidence of a statistically significance difference between the MRVG model specifications, or indeed, for that matter, the benchmark MRJD. While this showcases the comparable ability of the alternative MRVG models to capture the smile dynamics in the NBP options markets, for the purposes of storage valuation, the ability of a model to represent the covariance structure of the underlying NBP forward curve is central to accurately determining the extrinsic value of a storage

⁶The implementation of t-testing on the out-of-sample MSEs follows the same procedure, where the results are found to be consistent with the analysis on in-sample MSEs presented here.

contract. We therefore need to consider how the MRVG models perform in this respect. To this end, we explicitly derive the model implied covariance structure under each of the MRVG model specifications, emphasising the role of the structural parameters. We then proceed to explore the ability of each specification to model the historical covariance structure. The evidence shows that, in general, the MRVG-3x model more accurately represents the co-movement observed in the forward curve than either the two factor MRVG-2x or single factor MRVG models and so can more accurately value extrinsic storage value. It is this feature of the MRVG-3x model in particular that offers practitioner appeal.

We begin first with the derivation of the covariance structure function under the single factor MRVG model. 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 is the well-established Samuelson Effect (Serletis, 1992), a stylised feature of many commodity markets and particularly natural gas. For two log-forward prices, $f(t, T_1)$ and $f(t, T_2)$, the instantaneous covariance of returns under this model can be shown to be given by

$$\begin{aligned}
 & E [(df(t, T_1) - E[df(t, T_1)]) (df(t, T_2) - E[df(t, T_2)])] \\
 & = \exp(-\alpha(T_1 + T_2 - 2t)) \sigma^2 dt.
 \end{aligned}$$

Therefore, the mean reversion rate of the process, α , which captures the exponential decay of forward price volatility with respect to maturity, fully controls the structure of the forward curve covariance matrix. This results in non-parallel shifts in the forward curve due to changes in the underlying stochastic driver. It is this dynamic forward curve movement that is crucial for storage valuation as it motivates the operator to exercise their optionality to switch planned injection and withdrawals, thus creating extrinsic value. It should be noted that the variance of the jump magnitudes, ν , controls the implied volatility smile attenuation and ensures that the model is consistent with the initial volatility surface.

For the MRVG-2x model, the instantaneous covariance of returns is quite similar in form but also features the long-term volatility level. The instantaneous covariance of returns for the MRVG-2x model can be shown to be given by

$$E[(df(t, T_1) - E[df(t, T_1)])(df(t, T_2) - E[df(t, T_2)])] \\ = \exp(-\alpha(T_1 + T_2 - 2t))\sigma^2 dt + \sigma_L^2 dt.$$

where again, the mean reversion rate of the process, α , is the primary driver of extrinsic storage value.

In contrast, the two-factor MRVG-3x model offers much more flexibil-

ity. The first factor, which accounts for the majority of the forward curve variability, is an MRVG process and the main parameters for this factor can be calibrated to the options market. The parameters b and c_1 represent the proportion of total variance attributed to the first and second factors respectively and would need to be estimated from 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. The parameters relating to the second factor can be estimated directly from the eigenvector values. The parameter ϵ controls the decay of the volatility function as maturity increases. The slope of the volatility function 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. The instantaneous returns covariance for two log forward prices $f(t, T_1)$ and $f(t, T_2)$ in the case of the MRVG-3x model can be shown to be given by

$$\begin{aligned}
 & E[(df(t, T_1) - E[df(t, T_1)])(df(t, T_2) - E[df(t, T_2)])] \\
 & = \exp(-\alpha(T_1 + T_2 - 2t)) b^2 \sigma^2 dt + \exp(-\epsilon(T_1 + T_2 - 2t)) c_1^2 \sigma^2 dt.
 \end{aligned}$$

With $\epsilon \gg \alpha$ (as we estimate here), the MRVG-3x model, as it encompasses the one factor MRVG model, will attribute more value to a storage asset through this additional decorrelation of the forward curve returns. Appendix D provides full derivation details of the model implied covariance structures just presented.

With the model implied covariance structures now derived, we proceed to examine how well the MRVG model specifications fit the observed historical covariance structure. In calculating the model implied covariance matrices, we use the parameter estimates for the MRVG, MRVG-2x and MRVG-3x models as reported in Table 1. To showcase the respective fits to the historical covariance matrix, we present two heat maps in Figure 1; one (panel (a)) comparing the fit of the MRVG and MRVG-3x models to the historical covariance matrix, and the other (panel (b)) comparing the fit of the MRVG-2x and MRVG-3x models to the historical covariance matrix. To interpret the heat maps, negative values (black-green areas) show that the MRVG-3x model fits the historical covariance structure better than the MRVG (panel (a)) and MRVG-2x (panel (b)) models. The x- and y-axes label the twelve monthly forward maturities M1-M12. The construction of the heat maps proceeds as follows: calculate the historical covariance matrix, as used in the estimation of the MRVG-3x model; calculate the MRVG/MRVG-2x/MRVG-3x model implied covariance matrices (as per the derivations above) using the parameter estimates reported in Table 1, discretised over a single day period $\Delta t = 1/253$; calculate the absolute error

on an element-by-element basis between the historical covariance matrix and each of the model implied covariance matrices; for the heat map in panel (a), subtract the historical-MRVG absolute errors from the historical-MRVG-3x absolute errors on an element-by-element basis, such that a negative value indicates that the MRVG-3x model fits a given covariance matrix element better than the competing MRVG model; for the heat map in panel (b), subtract the historical-MRVG-2x absolute errors from the historical-MRVG3x absolute errors on an element-by-element basis, such that a negative value therefore indicates that the MRVG-3x model fits a given covariance matrix element better than the competing MRVG-2x model.

The MRVG-3x model clearly outperforms, almost entirely, the MRVG model in fitting the historical covariance matrix, except for a segment of the long dated maturities' co-movement. The MRVG-3x model also outperforms the MRVG-2x model for most of the medium-long dated maturities' co-movement, although the MRVG-2x model does seem to capture the covariance structure better for the short dated maturities' co-movement. In the round though, it can be concluded that the MRVG3x is better specified to represent the forward curve covariance structure and so is more flexible in accurately capturing extrinsic value. The MRVG-3x model specification therefore offers the most in this regard. To examine this further, the next section proceeds with a comprehensive storage valuation analysis across the MRVG model suite, first deriving a novel multidimensional Fourier based valuation algorithm.

MRVG Models				
Parameter	MRVG	MRVG-2x	MRVG-2	MRVG-3x
α	0.6365(c)	1.6164(c)	10.5668(h)	0.6459(c)
σ	0.2922(c)	0.2907(c)	0.8506(c)	0.3120(c)
ν	0.1130(c)	0.1935(c)	0.5993(c)	0.1103(c)
σ_L	-	0.1228(c)	-	-
b	-	$\equiv 1$	0.9367(h)	0.9367(h)
ϵ	-	$\equiv \alpha$	$\equiv \alpha$	10.2613(h)
c_1	-	$\equiv \frac{\sigma_L}{\sigma}$	0.2860(h)	0.2860(h)
c_2	-	$\equiv \frac{\sigma_L}{\sigma}$	-0.1022(h)	$\equiv 0$
Avg. MSE (in-samp)	0.90%	0.87%	16.14%	0.83%
Avg. MSE (out-samp)	0.90%	0.87%	16.16%	0.83%
Avg. MAE (price)	0.1959	0.1928	1.1276	0.1906
Avg. MAE (imp vol)	1.44%	1.39%	8.40%	1.38%
Avg. AIC	-11.2717	-9.5919	+12.5139	-11.6487
Benchmark MRJD				
Avg. MSE (in-samp)	1.24%		Avg. MSE (out-samp)	1.25%
Avg. MAE (price)	0.2287		Avg. MAE (imp vol)	1.65%
Avg. AIC	-8.8113			

The table presents the results of the model specification analysis on the multifactor Mean Reverting Variance Gamma (MRVG) model suite, as defined in Section 2. The market implied moments calibration procedure detailed in Section 3 is applied to estimate the parameters of the MRJD, MRVG and MRVG-2x models, using market options data only. The joint calibration-estimation procedure of Section 3, which incorporates calibration using the market implied moments procedure and historical estimation based on fitting the historical volatility functions, is applied to estimate the parameters of the MRVG-2 and MRVG-3x models. MSE is Mean Square Error, MAE is Mean Absolute Error and AIC is Akaike Information Criterion. MSEs are reported from the in-sample and out-of-sample analysis. MAEs are reported in price terms in order to give a monetary interpretation to the pricing performance and in implied volatility terms in order to give a volatility interpretation to the pricing performance. We use “(c)” to denote parameters that are calibrated, while we use “(h)” to denote parameters that are historically estimated. The calibrated parameters reported are averages of the parameters estimated from the daily calibrations conducted across the sample period, 5th January 2015 – 24th December 2015. We use “ \equiv ” to denote parameters that are constrained and set to the values described in Section 2. A constrained optimisation routine is used, where for consistency the same upper and lower bounds are used on all overlapping parameters between the MRVG models. For the parameter set $\{\alpha, \sigma, \nu, \sigma_L\}$, we set the upper bound constraints to be $\{4, 0.5, 1, 0.3\}$ and the lower bound constraints to be $\{0, 0, 0, 0\}$, with starting values $\{0.5, 0.25, 0.4, 0.1\}$.

Table 1: Model specification analysis on the MRVG model suite

Parameter	α			σ			ν			σ_L	
	MRVG-2x	MRVG-3x	MRVG-2x	MRVG-2x	MRVG-3x	MRVG-3x	MRVG-2x	MRVG-2x	MRVG-3x	MRVG-2x	MRVG-2x
Mean	1.6164	0.6459	0.2907	0.3120	0.1935	0.1103	0.4838	0.2686	0.2250	0.1727	0.1324
Max	3.9996	3.2801	0.5000	0.5000	1.0000	0.2990	0.2334	0.0552	0.1183	0.0668	
95 perc	3.8764	1.8370	0.4908	0.4448	0.6523	0.2990	0.2334	0.0552	0.1183	0.0668	
75 perc	2.6107	0.9144	0.3422	0.3463	0.2734	0.2334	0.2734	0.1638	0.0552	0.1324	
50 perc	1.3909	0.4573	0.2751	0.2971	0.1638	0.0552	0.1638	0.0552	0.1183	0.0668	
Std Dev	1.2712	0.5886	0.0980	0.0590	0.2141	0.1183	0.1183	0.1183	0.1183	0.1183	0.0668

Table 2: Descriptive statistics of implied parameter series

Moneyness	6 month maturity												1 year maturity											
	MAE (price)				MAE (imp vol)				MAE (price)				MAE (imp vol)											
	MRVG	MRVG2x	MRVG3x	MRVG	MRVG2x	MRVG3x	MRVG	MRVG2x	MRVG3x	MRVG	MRVG2x	MRVG3x	MRVG	MRVG2x	MRVG3x									
110%	0.0996	0.1005	0.0966	0.0095	0.0095	0.0092	0.1538	0.1556	0.1534	0.0093	0.0094	0.0093	0.0093	0.0094	0.0093									
105%	0.1583	0.1477	0.1472	0.0138	0.0128	0.0127	0.2273	0.2333	0.2296	0.0131	0.0135	0.0133	0.0131	0.0135	0.0133									
102.5%	0.1875	0.1742	0.1735	0.0160	0.0147	0.0146	0.2507	0.2586	0.2537	0.0143	0.0148	0.0145	0.0143	0.0148	0.0145									
100%	0.2054	0.1887	0.1897	0.0175	0.0158	0.0159	0.2510	0.2589	0.2531	0.0143	0.0148	0.0144	0.0143	0.0148	0.0144									
97.5%	0.2079	0.1907	0.1931	0.0182	0.0163	0.0166	0.2345	0.2398	0.2343	0.0135	0.0138	0.0134	0.0135	0.0138	0.0134									
95%	0.1960	0.1797	0.1832	0.0180	0.0162	0.0166	0.2173	0.2230	0.2156	0.0129	0.0132	0.0127	0.0129	0.0132	0.0127									
90%	0.1489	0.1404	0.1429	0.0165	0.0153	0.0156	0.2035	0.2077	0.2010	0.0141	0.0143	0.0138	0.0141	0.0143	0.0138									

The table presents Mean Absolute Error (MAE) results across moneyness-maturity dimensions for the MRVG, MRVG-2x and MRVG-3x models. MAEs are reported in price terms in order to give a monetary interpretation to the pricing performance and in implied volatility terms in order to give a volatility interpretation to the pricing performance. The maturity dimension spans the twelfth and sixth month expiry option maturities. The moneyness dimension spans the following levels, defined relative to the underlying futures price: 110%, 105%, 102.5%, 100%, 97.5%, 95%, and 90%.

Table 3: Mean Absolute Error analysis across moneyness-maturity dimensions

	MRJD	MRVG	MRVG2x	MRVG3x
MRJD	-	0.013	0.014	0.015
MRVG		-	0.003	0.008
MRVG2x			-	0.011

The table presents the t-statistics associated with the formal model comparison approach of Huang and Wu (2004). The t-statistics are specified as follows on a pairwise model comparison basis:

$$t\text{-statistic} = \frac{M\bar{S}E^i - M\bar{S}E^j}{\text{stdev} \left(MSE_t^i - MSE_t^j \right) / \sqrt{T}},$$

where $M\bar{S}E^i$ and $M\bar{S}E^j$ are the mean MSEs for models i and j respectively, $stdev()$ is the standard deviation function, and $T = 253$ is the number of days in the estimation period. The t-statistics are based off the in-sample MSEs.

Table 4: Model comparison t-testing

(a)
 << INSERT FIGURE 1 - PANEL (a) >>
 (b)
 << INSERT FIGURE 1 - PANEL (b) >>

The figure presents two heat maps; panel (a) comparing the fit of the MRVG and MRVG-3x model implied covariance matrices to the historical covariance matrix, and panel (b) comparing the fit of the MRVG-2x and MRVG-3x model implied covariance matrices to the historical covariance matrix. To interpret the heat maps, negative values (black-green areas) show that the MRVG-3x model fits the historical covariance structure better than the MRVG (panel (a)) and MRVG-2x (panel (b)) models. The x- and y-axes label the twelve monthly forward maturities M1-M12. The construction of the heat maps proceeds as follows: calculate the historical covariance matrix, as used in the estimation of the MRVG-3x model; calculate the MRVG/MRVG-2x/MRVG-3x model implied covariance matrices (as per the derivations described in Section 4 and detailed in Appendix D) using the average parameters reported in Table 1, discretised over a single day period $\Delta t = 1/253$; calculate the absolute error on an element-by-element basis between the historical covariance matrix and each of the model implied covariance matrices; for the heat map in panel (a), subtract the historical-MRVG absolute errors from the historical-MRVG-3x absolute errors on an element-by-element basis, such that a negative value indicates that the MRVG-3x model fits a given covariance matrix element better than the competing MRVG model; for the heat map in panel (b), subtract the historical-MRVG-2x absolute errors from the historical-MRVG3x absolute errors on an element-by-element basis, such that a negative value therefore indicates that the MRVG-3x model fits a given covariance matrix element better than the competing MRVG-2x model.

Figure 1: Heat maps for historical-model implied covariance structure fits

5 Storage Valuation

We proceed next to valuing an assumed storage deal with physical constraints under the estimated MRVG-2x and MRVG-3x models as outlined and fitted in the previous sections. A core contribution of this section is the development of a multidimensional algorithm, exploiting the power of the FFT for the purposes of storage valuation. The approach represents a generalisation of the single factor FFT based valuation algorithm of Cummins et al. (2017) 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.

5.1 Multidimensional Valuation Algorithm

We will begin by introducing the valuation problem with reference to the standard physical constraints and operating characteristics of the assumed 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: (i) the allowed injection/withdrawal nomination times over the valuation period belong to a discrete set $\{t_j\}$; and (ii) 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 one's 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

storage value is derived through backward stochastic dynamic programming, where the value function is given by

$$V(x_{t_{j-1}}(\vec{y}_{t_{j-1}}), I) = \sup_{I^*} \theta(x_{t_{j-1}}(\vec{y}_{t_{j-1}}), I^*; I) + E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}] \quad (13)$$

and where, in our multidimensional setting, $\vec{y}_t \in \mathbb{R}^K$. We therefore simply need to solve the expectation in Eq. (13) at each time-step and for each combination of (x_{t_j}, I^*) . Firstly, $E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}]$ may be written as

$$\mathbb{R} \dots \mathbb{R} V\left(f(0, t_j) + \sum_{k=1}^K y_{t_j}^{(k)}, I^*\right) \times f\left(y_{t_j}^{(1)}, y_{t_j}^{(2)}, \dots, y_{t_j}^{(K)} | y_{t_{j-1}}^{(1)}, y_{t_{j-1}}^{(2)}, \dots, y_{t_{j-1}}^{(K)}\right) dy_{t_j}^{(1)} dy_{t_j}^{(2)} \dots dy_{t_j}^{(K)}$$

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

$$\left(\frac{1}{2\pi}\right)^K \mathbb{C} \dots \mathbb{C} \tilde{v}(\vec{z}; I^*) \Phi_{\vec{y}_{t_j}}(\vec{z}; \vec{y}_{t_{j-1}}, t_{j-1}, t_j) dz^{(1)} dz^{(2)} \dots dz^{(K)} \quad (14)$$

where

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

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

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

The associated CCF 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 used to denote the element wise multiplication of the vector \vec{z} and $\vec{a} = \{a^{(1)}(t_{j-1}, t_j), a^{(2)}(t_{j-1}, t_j), \dots, a^{(K)}(t_{j-1}, t_j)\}$. Therefore, we can rewrite Eq.(14) to get the following form for $E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}]$:

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

Applying the substitution $\vec{z}' = \vec{a}\vec{z}$ gives us the following simplification for $E[V(x_{t_j}, I^*) | \vec{y}_{t_{j-1}}]$:

$$\begin{aligned} & \left(\prod_{k=1}^K \frac{1}{a^{(k)}(t_{j-1}, t_j)} \right) \left(\frac{1}{2\pi} \right)^K \int_{\mathbb{C}} \dots \int_{\mathbb{C}} \tilde{v} \left(\frac{\vec{z}'}{a}; I^* \right) \\ & \times \exp \left(i\vec{z}'^\top \cdot \vec{y}_{t_{j-1}} \right) \exp \left(\psi \left(\frac{\vec{z}'}{a}, t_{j-1}, t_j \right) \right) dz'^{(1)} dz'^{(2)} \dots dz'^{(K)} \end{aligned} \quad (15)$$

We then apply the scaling property of the Fourier transform to derive $\tilde{v} \left(\frac{\vec{z}'}{a}; I^* \right)$:

$$\left(\prod_{k=1}^K a^{(k)}(t_{j-1}, t_j) \right) \mathcal{F} \left[V \left(f(0, t_j) + \sum_{k=1}^K a^{(k)}(t_{j-1}, t_j) y_{t_j}^{(k)}, I^* \right) \right] \quad (16)$$

This lends itself to efficient evaluation using a multidimensional FFT algorithm. The next section discusses discretising the expectation given in Eq. (15) and recasting it in a form suitable for evaluation using the FFT.

Discretisation & Solution

To solve for Eq. (15) we first truncate and discretise the domains of \vec{y} and \vec{z} such that $y^{(k)} \in [y_0^{(k)}, \dots, y_{N^{(k)}-1}^{(k)}]$, 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

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

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 Eq. (16) then becomes

$$\left(\prod_{k=1}^K a^{(k)}(t_{j-1}, t_j) \right) \sum_{n^{(1)}=0}^{N^{(1)}-1} l_{n^{(1)}}^{(1)} \dots \sum_{n^{(K)}=0}^{N^{(K)}-1} l_{n^{(K)}}^{(K)} \exp \left(-iz'^{(1)} y_{n^{(1)}}^{(1)} \dots - iz'^{(K)} y_{n^{(K)}}^{(K)} \right) \times \\ V \left(f(0, t_j) + \sum_{k=1}^K a^{(k)}(t_{j-1}, t_j) y_{n^{(k)}}^{(k)}, I^* \right) \Delta y^{(K)} \dots \Delta y^{(1)}$$

where $l_{n^{(k)}}^{(k)} = \frac{1}{2}$ for $n^{(k)} = 0, N^{(k)} - 1$ and $l_{n^{(k)}}^{(k)} = 1$ elsewhere. Each term $\exp\left(-iz'^{(k)}y_{n^{(k)}}^{(k)}\right)$ can now be expanded as follows:

$$\begin{aligned} \exp\left(-iu_0^{(k)}n^{(k)}\Delta y^{(k)} - im^{(k)}\Delta u^{(k)}n^{(k)}\Delta y^{(k)} - iu_0^{(k)}y_0^{(k)}\right. \\ \left.- im^{(k)}\Delta u^{(k)}y_0^{(k)} + w^{(k)}n^{(k)}\Delta y^{(k)} + w^{(k)}y_0^{(k)}\right) \end{aligned}$$

Terms which are independent of $n^{(k)}$ can be factored out to give the following expression for $\tilde{v}\left(\frac{\vec{z}'}{a}; I^*\right)$:

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

Similarly, we discretise Eq. (15) to get the following expression for $E [V (x_{t_j}, I^*) | \vec{y}_{t_{j-1}}]$:

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

Expanding $\tilde{v}\left(\frac{\vec{z}}{a}; I^*\right)$ we finally get:

$$\begin{aligned} & \left(\prod_{k=1}^K \frac{\Delta u^{(k)} \Delta y^{(k)}}{2\pi} \right) \exp \left(\sum_{k=1}^K i u_0^{(k)} n^{(k)} \Delta y^{(k)} - w^{(k)} n^{(k)} \Delta y^{(k)} \right) \times \\ & \sum_{M^{(1)}=0}^{M^{(1)}-1} l_{m^{(1)}}^{(1)} \dots \sum_{m^{(K)}=0}^{M^{(K)}-1} l_{m^{(K)}}^{(K)} \exp \left(\sum_{k=1}^K i m^{(k)} \Delta u^{(k)} n^{(k)} \Delta y^{(k)} \right) \psi \left(\frac{\vec{z}}{a}, t_{j-1}, t_j \right) \times \\ & \sum_{n^{(1)}=0}^{N^{(1)}-1} l_{n^{(1)}}^{(1)} \dots \sum_{n^{(K)}=0}^{N^{(K)}-1} l_{n^{(K)}}^{(K)} \exp \left(\sum_{k=1}^K -i m^{(k)} \Delta u^{(k)} n^{(k)} \Delta y^{(k)} \right) \times \\ & \exp \left(\sum_{k=1}^K -i u_0^{(k)} n^{(k)} \Delta y^{(k)} + w^{(k)} n^{(k)} \Delta y^{(k)} \right) V \left(f(0, t_j) + \sum_{k=1}^K a^{(k)}(t_{j-1}, t_j) y_{n^{(k)}}^{(k)}, I^* \right) \end{aligned}$$

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

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

In implementing this Fourier based storage valuation algorithm, we exploit the FFT-m approach of Cummins et al. (2017), which is a modification that leads to an increase in valuation convergence. Note that for the MRVG-2x and MRVG-3x models, grid sizes of 512×512 and 256×2048 are used respectively. For the single factor MRVG and MRJD benchmark models, the number of grid points used in the FFT-m implementation is 4026.

5.2 Storage Valuation Example

Here we apply the multidimensional valuation algorithm derived in the previous section to an assumed storage contract with physical constraints. We value the storage contract using the daily estimated MRVG-2x and MRVG-3x models, determined in Section 4, while for comparative purposes we additionally value the same contract under the MRVG and MJRD benchmark models. In order to allow us focus on the relationship between model specification and storage value, we have chosen a straightforward, yet realistic storage contract. The deal parameters are as follows:

- Tenor: 1 year;
- Capacity: 29.3 GWh;
- Max Injection/Withdrawal: 1.465 GWh per day;
- Underlying Gas Price: NBP (National Balancing Point) in pence/therm.

We assume the same deal parameters on each day in our sample period of 5th January 2015 – 24th December 2015 but use the fitted model on each day for the valuation exercise. Therefore, on each day we are valuing a storage contract with a one year tenor from that date. Table 5 presents the average storage values for each model considered. The average intrinsic value calculated across the sample period is also reported so that we can appraise the ability of the models in capturing extrinsic value.

Model	Intrinsic Value	Full Value	Extrinsic Value	% Δ Extrinsic Value
MRJD	7.0638	7.7551	0.6913	-
MRVG	7.0638	7.7863	0.7225	4.5%
MRVG-2x	7.0638	9.5540	2.4902	260.2%
MRVG-3x	7.0638	9.0853	2.0215	192.4%

Table 5: Model Based Storage Valuation Results

The table presents the average storage valuation results for the assumed storage contract as set out in Section 5.2 under the two factor MRVG-2x and MRVG-3x models, along with storage valuation results for the single factor MRVG and MRJD benchmark models. The deal parameters are as follows: tenor of 1 year; capacity of 29.3 GWh; max injection/withdrawal of 1.465 GWh per day. All monetary values are expressed in pence/therm. We assume the same deal terms on each day in our sample period of 5th January 2015 – 24th December 2015 but use the fitted model on each day for the valuation exercise. Therefore, on each day we are valuing a storage contract with a one year tenor from that date. The average intrinsic value is calculated across the sample period in order to appraise the ability of the MRVG models in capturing extrinsic value. '% Δ Extrinsic Value' gives the increase in extrinsic value captured by the MRVG models relative to the extrinsic value captured by the benchmark MRJD model.

Notably, the results show that on average the MRVG-2x and MRVG-3x models assign considerably higher value to the storage contract than the benchmark models, at 9.5540 pence/therm and 9.0853 pence/therm respectively. These represent increases of approximately 260% and 192% respectively in extrinsic value over the MRJD model. The average storage value assigned by the MRVG-2x model is higher than the MRVG-3x case by some 0.47 pence/therm but it should be noted from Section 4 that MRVG-3x is established as being better specified to represent the forward curve covariance structure and so more flexible in accurately capturing extrinsic value. The results overall though highlight the sensitivity of the storage asset to the model implied time spread price variability. Recall that both models incorporate exponentially decaying positive or negative shifts in the forward curve, which add a certain degree of time spread price variability. However, the inclusion of the rapidly decaying second factor on the MRVG-3x acts to more accurately represent this variability and avoid the potential for resulting misvaluation.

Incorporating an accurate representation of the forward curve dynamics in conjunction with a market based calibration of the general level of curve variability is an important feature of modelling storage value. 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 value being bid/offered. Not only does the MRVG-3x model

meet this requirement, the parsimony of the model allows one to adjust one's price levels significantly by adjusting a single parameter, ϵ , without significantly impacting the ability of the model 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 utilising historical information may not be necessary in coming years. Even 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 covariance structure is controlled by two parameters, would be considerably more suited for this purpose. Indeed, this feature means the MRVG-3x model offers potential as a dynamic model outside of the natural gas markets, in energy and commodity markets where greater levels of market liquidity may be observed.

6 Conclusion

We extend the existing range of gas storage valuation methods uniquely within a flexible multifactor Lévy process setting. Specifically, we develop a family of multifactor Mean Reverting Variance Gamma (MRVG) models, which are forward curve consistent, while also being broadly consistent with the options market, and so more reflective of the statistical dynamics of the NBP gas forward curve returns we examine. We extend the single factor

MRVG process of Cummins et al. (2017), set out in Section 2.1, to an arbitrary number of dimensions and by way of specific examples show how the traditional Principal Component Analysis based view of gas forward curve dynamics can be incorporated into a primarily market based valuation. This methodological approach allows for consistent valuation and risk management reporting across an energy trading firm; a practitioner problem of real relevance for industry. We demonstrate how to construct models with the capability of accurately capturing the time spread volatility, whilst maintaining consistency with the vanilla options market. We present the first class of multifactor storage models developed with the explicit intention of providing a link between the model and the options market, developing in the process an innovative implied moments based calibration technique, which allows for efficient calibration of general multifactor forward curve models. Due to concerns surrounding the efficiency of calibrating directly to market option prices, we extend previous work utilising 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, where the computational effort is broadly independent of the number of factors in the underlying model. To accommodate forward curve and traded options market consistency, we go further and develop an appropriate joint market based calibration and historical estimation approach. A formal model specification analysis provides evidence that the two factor MRVG-2x and MRVG-3x models we propose, set out in Section 2.2,

provide a better fit to the natural gas options-futures markets relative to the benchmark MRVG model of Cummins et al. (2017) and the Mean Reverting Jump-Diffusion (MRJD) model of Deng (2000). Finally, we present empirical valuation results for a stylised storage deal under selected specifications from the multifactor Lévy model framework, which showcases the merits of the models. We develop a novel multidimensional fast Fourier transform (FFT) based early exercise claim valuation algorithm and showcase its application to valuing natural gas storage. We leverage the model fitting performed as part of the formal model specification analysis to investigate the performance of the MRVG model suite in terms of storage valuation, identifying that the two factor MRVG models generally capture higher levels, and more accurate levels, of extrinsic value. We discuss the dependence of the resulting storage values on the model specification and calibrated/estimated parameter values.

One issue when valuing the storage deal is the computational burden of the algorithm in higher dimensions, such as the overarching MRVG-3 model. As identified in Section 2.2, the memory requirement to implement the MRVG-3 model under a three-dimensional version of the FFT algorithm would exceed that limits of modern hard ware. We therefore see the need for optimisation of the algorithm as a key area for future research, with potential to explore cloud based parallel computing solutions.

Finally, whilst we have demonstrated the additional theoretical value of the storage asset under our enhanced suite of multifactor models, a full back-testing of the model "greeks" ability to protect this value through dynamic

hedging in both the forward and options market would be needed before drawing any firm conclusions on which model is best in practice. Such a study though is out of scope for this paper and left for future research. However, by way of laying a blue print for this work, the following is important to note. While there is no industry standard approach to model backtesting, we would see value in carrying out a model value attribution analysis where the ability of the first and higher order derivatives of the model inputs are used to explain the daily changes in the model value with respect to changes in these inputs. The exact steps in carrying out this type of analysis would be as follows: (1) select several time periods, each spanning at least one year, which represent different market regimes and ideally include a number of extreme short-term macro events, so as to cover low and high volatility environments; (2) for each day in the sample, calibrate/estimate the model parameters using all necessary market and historical information available as of that day; (3) on the first day of the time period being studied, run the storage valuation and record the value and model parameter "greeks"; (4) on the next business day in the time period, repeat step (3) and determine the ability of the model "greeks" to explain the change in storage value using a Taylor Series type analysis of the change in model value with respect to actual changes in the model inputs, the main metric to be recorded is the unattributed value change defined as the difference between actual and predicted model value; and (5) repeat step (4) for each day in the period and collect a sample of daily unattributed value changes, which can then be

analysed to infer the stability of the model under real world conditions. We would consider the model specification and parameter estimation to be satisfactory if the daily unattributed value change has a statistically significant zero sample mean and low sample variance. Two key challenges to be considered in future work are: the time required to evaluate the required number of model value derivatives with respect to the model parameters via numerical differentiation for each day in a sufficiently sized sample would not be practical; and the impact of numerical error on the robustness of these estimates for higher order "greeks" would likely invalidate any conclusion which one could draw from the analysis.

Acknowledgments

The authors would like to thank Prof. Michael Dempster (University of Cambridge) for reviewing an earlier draft of this paper and for very helpful comments provided. The authors would also like to thank the participants of the Energy and Commodity Finance Conference 2016, in particular the discussant Tomasso Pellegrino, for a range of valuable suggestions. The authors would also like to thank the Associate Editor and two anonymous reviewers for comments and recommendations that have led to a substantially improved study.

References

- Andersen, L., 2010. Markov models for commodity futures: theory and practice. *Quantitative Finance* 10, 831–854.
- Andricopoulos, A., Widdicks, M., Duck, P., Newton, D., 2003. Universal option valuation using quadrature methods. *Journal of Financial Economics* 67, 447–471.
- Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. *The Journal of finance* 52, 2003–2049.
- Bakshi, G., Madan, D., 2000. Spanning and derivative-security valuation. *Journal of Financial Economics* 55, 205–238.
- Barndorff-Nielsen, O., Mikosch, T., Resnick, S., 2012. Lévy processes: theory and applications. Springer Science & Business Media.
- 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, 153–169.
- 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 .
- Boogert, A., De Jong, C., 2008. Gas storage valuation using a monte carlo method. *The Journal of Derivatives* 15, 81–98.

- Boogert, A., De Jong, C., 2011. Gas storage valuation using a multifactor price process. *The Journal of Energy Markets* 4, 29–52.
- 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, 359–374.
- Carr, P., Madan, D., 1999. Option valuation using the fast fourier transform. *Journal of Computational Finance* 2, 61–73.
- Chen, Z., Forsyth, P., 2009. A semi-lagrangian approach for natural gas storage valuation and optimal operation. *SIAM Journal on Scientific Computing* 30.
- 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, 1–18.
- Christoffersen, P., Jacobs, K., 2004. The importance of the loss function in option valuation. *Journal of Financial Economics* 72, 291–318.
- Clelow, L., Strickland, C., 1999. A multi-factor model for energy derivatives. Technical Report.

- Cont, R., Tankov, P., 2004. Financial modelling with jump processes. volume 2. Chapman & Hall.
- Cummins, M., Kiely, G., Murphy, B., 2017. Gas storage valuation under lévy processes using the fast fourier transform. *Journal of Energy Markets* forthcoming.
- Deng, S., 2000. Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes. Citeseer.
- Detering, N., Packham, N., 2016. Model risk of contingent claims. *Quantitative Finance* 16, 1357–1374.
- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Felix, B., Weber, C., 2012. Gas storage valuation applying numerically constructed recombining trees. *European Journal of Operational Research* 216, 178–187.
- Guillaume, F., Schoutens, W., 2013. A moment matching market implied calibration. *Quantitative Finance* 13, 1359–1373.
- Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of financial studies* 6, 327–343.

- Huang, J.z., Wu, L., 2004. Specification analysis of option pricing models based on time-changed lévy processes. *The Journal of Finance* 59, 1405–1439.
- Hull, J.C., White, A., 2012. Libor vs. ois: The derivatives discounting dilemma. *Jornal of Investment Management* 11, 14–27.
- Jaimungal, S., Surkov, V., 2011. Levy based cross-commodity models and derivative valuation. To Appear: *SIAM Journal of Financial Mathematics* .
- Kiely, G., Murphy, B., Cummins, M., 2015. Gas storage valuation under levy processes using the fast fourier transform. Available at SSRN:2561684 .
- Kjaer, M., 2008. Pricing of swing options in a mean reverting model with jumps. *Applied mathematical finance* 15, 479–502.
- 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, 564–582.
- 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, 289–330.

- 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) .
- Luciano, E., 2009. Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy. volume 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, 79–105.
- Manoliu, M., 2004. Storage options valuation using multilevel trees and calendar spreads. *International Journal of Theoretical and Applied Finance* 7, 425–464.
- Marfè, R., 2009. A multivariate two% factor variance gamma process for asset returns. Technical Report. Working Paper, Swiss Finance Institute.
- Maslyuka, S., Rotarub, K., Dokumentovc, A., 2013. Price Discontinuities in Energy Spot and Futures Prices. Technical Report. Monash University, Department of Economics.
- Nadarajah, S., Margot, F., Secomandi, N., 2015. Relaxations of approximate linear programs for the real option management of commodity storage. *Management Science* 61, 3054–3076.

- Nomikos, N., Andriosopoulos, K., 2012. Modelling energy spot prices: Empirical evidence from nymex. *Energy Economics* 34, 1153–1169.
- O’Sullivan, C., 2005. Path dependant option pricing under lévy processes, in: EFA 2005 Moscow Meetings Paper.
- Parsons, C., 2013. Quantifying natural gas storage optionality: a two-factor tree model. *The Journal of Energy Markets* 6.
- Sato, K., 2001. Basic results on lévy processes, in: *Lévy processes*. Springer, pp. 3–37.
- Serletis, A., 1992. Maturity effects in energy futures. *Energy Economics* 14, 150–157.
- Warin, X., 2012. Gas storage hedging, in: *Numerical Methods in Finance*. Springer, pp. 421–445.

A Drift Corrections: MRVG Spot Dynamics

It can be readily shown from a simplification of the spot dynamics of the state variables, while setting $\sigma(t) = \sigma$, that the drift corrections for the MRVG-3 are

$$\begin{aligned}
w^{(1)}(t) &= \left(-\kappa_j (\exp(-\alpha(t)) b\sigma) - \alpha_0^t (\kappa_j (\exp(-\alpha(t-c)) b\sigma)) dc \right) \\
\omega^{(2)}(t) &= -\frac{1}{2} (c_1\sigma)^2 + \frac{1}{4} (1 - \exp(-2\epsilon t)) (c_1 - c_2)^2 \sigma^2 - \frac{\epsilon}{2} c_2^2 \sigma^2 t
\end{aligned}$$

Following similar derivations, it can be shown for the MRVG-2 model that the drift correction is given by

$$\begin{aligned}
w^{(1)}(t) &= \left(-\kappa_j (\exp(-\alpha(t)) b\sigma) - \alpha_0^t (\kappa_j (\exp(-\alpha(t-c)) b\sigma)) dc \right. \\
&\quad \left. -\frac{1}{2} (c_1\sigma)^2 + \frac{1}{4} (1 - \exp(-2\alpha t)) (c_1 - c_2)^2 \sigma^2 - \frac{\alpha}{2} c_2^2 \sigma^2 t \right)
\end{aligned}$$

Finally, for the abridged MRVG-2x model, we have

$$w^{(1)}(t) = \left(-\kappa_j (\exp(-\alpha(t)) \sigma) - \alpha_0^t (\kappa_j (\exp(-\alpha(t-c)) \sigma)) dc \right).$$

B Characteristic Functions

Beginning with the MRVG-3 model, we have

$$\begin{aligned}
\Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(s), s, t) &= \exp\left(iz^{(1)}y^{(1)}(s)\exp(-\alpha(t-s)) + iz^{(2)}y^{(2)}(s)\exp(-\epsilon(t-s))\right) \times \\
&\exp\left(izy^{(3)}(s) + iz^{(1)}\int_s^t \omega^{(1)}(c)\exp(-\alpha(t-c))dc\right) \times \\
&\exp\left(iz^{(2)}\int_s^t \omega^{(2)}(c)\exp(-\epsilon(t-c))dc\right) \times \\
&\exp\left(\int_s^t \varphi_{vg}(z\exp(-\alpha(t-c))b\sigma)\right) \times \\
&\exp\left[\left(-\frac{\sigma^2}{2}\right)\left(\frac{1}{2\epsilon}(z^{(2)})^2(c_1 - c_2)^2 \times\right.\right. \\
&\left.\left.(1 - \exp(-2\epsilon(t-s))) + \frac{2}{\epsilon}z^{(2)}(1 - \exp(-\epsilon(t-s))) \times\right.\right. \\
&\left.\left.(c_1 - c_2)z^{(3)}c_2 + (z^{(3)})^2c_2^2(t-s)\right)\right]
\end{aligned}$$

for $\vec{z} \in S_y \cup \mathbb{C}^2$, $S_y := \left\{z = u + iw; w \in \left(-\sqrt{\frac{2}{(b\sigma)^2v}}, \sqrt{\frac{2}{(b\sigma)^2v}}\right)\right\}$. φ_{vg} is the characteristic function of the Variance-Gamma process, and the solution to $\int_s^t \varphi_{vg}(z\exp(-\alpha(t-c))b\sigma)$ is given in Kiely et al. (2015).

The characteristic function reduced MRVG-2 model is given by

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

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

Finally, for the MRVG-2x we have

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

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

C 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 and MRVG-2x models. Beginning with the second moment, recall that we need to derive a solution for

$$\frac{1}{F(0, T_1, T_2)^2} \times E \left[\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} F(0, T_1 + i\Delta t) F(0, T_1 + j\Delta t) \exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t)) | y(\vec{0}) \right]$$

Taking the expectation operator inside the summation, the only difficulty is in evaluating the following

$$E[\exp(y(t, T_1 + i\Delta t) + y(t, T_1 + j\Delta t)) | y(0)] \quad (17)$$

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 the dynamics of the MRVG-3 model, we know that

$$\begin{aligned} y^{(1)}(t, T_i) &= y^{(1)}(s, T_i) - \int_s^t (\kappa_j (b \exp(-\alpha(T_i - c)) \sigma)) dc \\ &\quad + \exp(-\alpha(T_i - t)) (y^{(1)}(t, t) - y^{(1)}(s, t)) \\ &\quad + \exp(-\alpha(T_i - t)) \varphi(-i, s, t) \end{aligned}$$

where φ is the characteristic exponent of the single factor MRVG model,

$$\begin{aligned} y^{(2)}(t, T_i) &= y^{(2)}(s, T_i) + \exp(-\epsilon(T_i - t)) (y^{(2)}(t, t) - y^{(2)}(s, t)) \\ &\quad - \frac{\sigma^2}{4\epsilon} \exp(-2\epsilon(T_i - t)) (c_1 - c_2)^2 \times \\ &\quad ((1 - \exp(\epsilon(T_i - t))) (1 - \exp(-2\epsilon(t - s)))) \\ &\quad - \frac{\sigma^2}{2} c_2^2 (t - s) (1 - \exp(-\epsilon(T_i - t))) \end{aligned}$$

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 of Eq. (17), i.e. $y(t, T_i) + y(t, T_j)$ as

$$\begin{aligned} & y^{(1)}(t, t) \sum_{x \in \{i, j\}} \exp(-\alpha(T_x - t)) + y^{(2)}(t, t) \sum_{x \in \{i, j\}} \exp(-\epsilon(T_x - t)) + 2y^{(3)}(t, t) \\ & + \sum_{x \in \{i, j\}} -\int_0^t (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\ & + \sum_{x \in \{i, j\}} -\frac{\sigma^2}{4\epsilon} \exp(-2\epsilon(T_x - t)) (c_1 - c_2)^2 ((1 - \exp(\epsilon(T_x - t))) (1 - \exp(-2\epsilon t))) \\ & + \sum_{x \in \{i, j\}} -\frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\epsilon(T_x - t))) \end{aligned}$$

Focusing on the non-deterministic part of the above and noticing that

$$E \left[\exp \left(y^{(1)}(t, t) \sum_{x \in \{i, j\}} \exp(-\alpha(T_x - t)) + y^{(2)}(t, t) \sum_{x \in \{i, j\}} \exp(-\epsilon(T_x - t)) + 2y^{(3)}(t, t) \right) \right]$$

is equal to $\Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(0), 0, t)$ with

$$\vec{z} = \begin{pmatrix} -i \sum_{x \in \{i,j\}} \exp(-\alpha(T_x - t)) \\ -i \sum_{x \in \{i,j\}} \exp(-\epsilon(T_x - t)) \\ -2i \end{pmatrix}$$

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 $T_x = T_1 + x\Delta t$

$$\begin{aligned} A &= \sum_{x \in \{i,j\}} -\int_0^t (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\ &+ \sum_{x \in \{i,j\}} -\frac{\sigma^2}{4\epsilon} \exp(-2\epsilon(T_x - t)) (c_1 - c_2)^2 ((1 - \exp(\epsilon(T_x - t))) (1 - \exp(-2\epsilon t))) \\ &+ \sum_{x \in \{i,j\}} -\frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\epsilon(T_x - t))) \end{aligned}$$

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

$$\vec{z} = \begin{pmatrix} -i \sum_{x \in \{i,j,k\}} \exp(-\alpha(T_x - t)) \\ -i \sum_{x \in \{i,j,k\}} \exp(-\epsilon(T_x - t)) \\ -3i \end{pmatrix}$$

and

$$\begin{aligned} A = & \sum_{x \in \{i,j,k\}} -{}^t_0 \kappa_j (b \exp(-\alpha(T_x - c)) \sigma) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\ & + \sum_{x \in \{i,j,k\}} -\frac{\sigma^2}{4\epsilon} \exp(-2\epsilon(T_x - t)) (c_1 - c_2)^2 ((1 - \exp(\epsilon(T_x - t))) (1 - \exp(-2\epsilon t))) \\ & + \sum_{x \in \{i,j,k\}} -\frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\epsilon(T_x - t))) \end{aligned}$$

Finally, the fourth moment is given by

$$\frac{1}{F(0, T_1, T_2)^4} \times \left[\frac{1}{N^4} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} F(0, T_i) F(0, T_j) F(0, T_k) F(0, T_l) \Phi_{\vec{y}(t)}(\vec{z}; \vec{y}(0), 0, t) \exp(A) \right]$$

where

$$\vec{z} = \begin{pmatrix} -i \sum_{x \in \{i,j,k,l\}} \exp(-\alpha(T_x - t)) \\ -i \sum_{x \in \{i,j,k,l\}} \exp(-\epsilon(T_x - t)) \\ -4i \end{pmatrix}$$

and

$$\begin{aligned}
A &= \sum_{x \in \{i,j,k,l\}} -{}^t_0 (\kappa_j (b \exp(-\alpha(T_x - c)) \sigma)) dc + \exp(-\alpha(T_x - t)) \varphi(-i, 0, t) \\
&+ \sum_{x \in \{i,j,k,l\}} -\frac{\sigma^2}{4\epsilon} \exp(-2\epsilon(T_x - t)) (c_1 - c_2)^2 ((1 - \exp(\epsilon(T_x - t))) (1 - \exp(-2\epsilon t))) \\
&+ \sum_{x \in \{i,j,k,l\}} -\frac{\sigma^2}{2} c_2^2 t (1 - \exp(-\epsilon(T_x - t)))
\end{aligned}$$

D Model Implied Covariance Matrices

Firstly, recall that the log dynamics of a forward price at maturity T under the MRVG-3 model, as described in Section 2.2, are given by

$$df(t, T) = \sum_{k=1}^3 dy^{(k)}(t, T),$$

$$\begin{aligned}
dy^{(1)}(t, T) &= (-\kappa_j (\exp(-\alpha(T - t)) b \sigma(t))) dt + \exp(-\alpha(T - t)) b \sigma(t) dX(t) \\
dy^{(2)}(t, T) &= -\frac{1}{2} ((\exp(-\epsilon(T - t)) (c_1 - c_2) + c_2) \sigma(t))^2 dt \\
&\quad + (\exp(-\epsilon(T - t)) (c_1 - c_2)) \sigma(t) dW(t) \\
dy^{(3)}(t, T) &= c_2 \sigma(t) dW(t)
\end{aligned}$$

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

$$\begin{aligned}
Cov [t; T_s, T_l] &= E \left[\left(\sum_{k=1}^K dy^{(k)}(t, T_s) - E \left[\sum_{k=1}^K dy^{(k)}(t, T_s) \right] \right) \right. \\
&\quad \left. \times \left(\sum_{k=1}^K dy^{(k)}(t, T_l) - E \left[\sum_{k=1}^K dy^{(k)}(t, T_l) \right] \right) \right] \\
&= E [\exp(-\alpha(T_s - t)) b \sigma dX(t) \exp(-\alpha(T_l - t)) b \sigma dX(t)] \\
&\quad + E [((\exp(-\epsilon(T_s - t))(c_1 - c_2)) + c_2) \sigma dW(t) \times \\
&\quad ((\exp(-\epsilon(T_l - t))(c_1 - c_2)) + c_2) \sigma dW(t)] \\
&= \exp(-\alpha(T_l + T_s - 2t)) b^2 \sigma^2 dt + (\exp(-\epsilon(T_l + T_s - 2t))(c_1 - c_2)^2) \sigma^2 dt \\
&\quad + c_2(c_1 - c_2) (\exp(-\epsilon(T_l - t)) \sigma^2 dt + \exp(-\epsilon(T_s - t))) + c_2^2 \sigma^2 dt
\end{aligned}$$

Setting $c_2 = 0$ gives the MRVG-3x model equivalent

$$\begin{aligned}
Cov [t; T_s, T_l] &= \exp(-\alpha(T_l + T_s - 2t)) b^2 \sigma^2 dt \\
&\quad + \exp(-\epsilon(T_l + T_s - 2t)) c_1^2 \sigma^2 dt
\end{aligned}$$

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,

$$\text{Cov}[t; T_s, T_l] = \exp(-\alpha(T_l + T_s - 2t)) \sigma^2 dt + \sigma_L^2 dt$$