UK Gas Markets: the Market Price of Risk and Applications to Multiple Interruptible Supply Contracts

Álvaro Cartea[*][†]

Commodities Finance Centre

Birkbeck College, University of London, UK

Email: a.cartea@bbk.ac.uk

Thomas Williams

Dresdner Kleinwort, London

Email: Thomas.Williams@dresdnerkleinwort.com

Forthcoming: Energy Economics

March 5, 2007

---

[*]Corresponding author a.cartea@bbk.ac.uk.

[†]The authors are grateful for comments from Andrew Shaw, Gareth Davies, Murray Hartley, two anonymous referees and from seminar participants at the SIAM meeting Boston 2006, University of Oslo, University of Ulm, Risk Measurement and Management Summer School (Rome), Energy Forum (London 2005). Cartea is thankful for the hospitality and generosity provided by the Finance Group at the Saïd Business School, Oxford, where part of this research was undertaken.

Abstract

We employ the Schwartz and Smith (2000) model to explore the dynamics of the UK gas markets. We discuss in detail the short-term and long-term market prices of risk borne by the market players and how deviations from expected cyclical storage affect the short-term market price of risk. Finally, we illustrate an application of the model by pricing interruptible supply contracts that are currently traded in the UK.

Keywords: Interruptible supply contracts, gas markets, commodities, market price of short-term and long-term risk, multi-exercise Bermudan options, convenience yield. JEL Classification: G12, C61.

1   Introduction

Over the last twenty years the UK natural gas market has undergone dramatic changes. Starting with the 1982 Oil and Gas Act, the British Government passed a succession of laws designed to bring competition to the transmission and distribution of natural gas, areas previously monopolised by the publicly owned British Gas. With the 1995 Gas Act, the groundwork was laid for the introduction of full retail competition in the natural gas industry, creating licensing schemes for companies to engage in the transport and supply of gas. This was followed in 1996 by the Network Code, a legal framework for the relationship between the operator of the pipeline system (the now privatised British Gas Transco) and shippers, those using the pipeline system to transport gas.

Network Code was designed to provide a set of market-based mechanisms to ensure the optimal operation of the UK gas pipeline system by Transco. Successfully and safely running a gas pipeline system is a complex task. The system operator has to carefully monitor and control the system intake (gas injected into the pipeline system by producers) the system off-take (gas withdrawn from the system by end-users) and the physical transportation of the gas around the whole national network. This is done to maintain an equilibrium between instantaneous supply and demand for natural gas at the various Local Distribution Zones throughout the country, and maintain system pressure and quality.

2

A shipper running an imbalance, either injecting or withdrawing more gas than it is contracted to, imposes a cost on other users of the system. To discourage the externality caused by these imbalances, the Network Code allows Transco to impose severe balancing penalties on shippers found to be breaching a certain tolerance level. The resulting price is, often significantly, above the market price for gas when the shipper is short gas and below the market price when the shipper is long.

This balancing mechanism requires the system operator to be provided with an up-to-date market price of natural gas on which to base the balancing price. Before deregulation, gas had overwhelmingly been sold to end-users on long-term contracts with terms agreed upon in October of a given year. However, under the new system, smaller independent end-user suppliers entered the market often purchasing 100% take-or-pay contracts from producers.[1] This created a demand for short-term contracts to allow new entrants to meet their balancing needs, selling their surpluses back to the market. As a result, a highly liquid spot market for gas developed. On January 31st, 1997 standardised gas futures contracts were launched on London's International Petroleum Exchange for delivery via a virtual system hub, the National Balancing Point (NBP). This allowed for a system-wide trading point and a national spot market needed for the purpose of balancing.

The development of the UK gas markets exposes participants to different types of risks. One way in which market participants may manage their exposure to price and volume fluctuations is by buying or selling instruments written on gas. One of the most common and important types of these contracts has been the interruptible supply contract, which gives the gas supplier the right to cease supplying his customers with gas for a finite number of days throughout the life of the contract.

The contribution of this article is twofold. First, we employ Schwartz and Smith's model (Schwartz and Smith 2000) to explore the dynamics of the UK natural gas industry to determine what economic factors influence spot and forward prices. Given the idiosyncrasies of the storage facilities in the UK gas markets, for example the constraints on inflow and outflow, we argue that the relationship between the short-term market price of risk and

---

[1]Take-or-pay means the buyer of gas commits to buying a set annual contract quantity for which he or she is obliged to pay, but if all the gas is not required there is no obligation to take it. This is sometimes referred to as a 'buyer's option' agreement as the buyer has the option to take the gas or not.

storage is, to a large extent, determined by deviations from *expected seasonal storage levels* rather than absolute levels. Second, we price interruptible supply contracts using actual contracts traded in the UK as a basis (E.ON Energy 2005).

The rest of this article is structured as follows. Section 2 discusses the UK gas market and motivates the choice of model. Section 3 proposes a spot model driven by short-term and long-term shocks, plus a seasonal component, under both the physical and risk-neutral measure. Section 4 discusses the estimation of the parameters for the model. Section 5 prices standard UK interruptible supply gas contracts. Finally, section 6 concludes and discusses further work.

2   Spot Prices, Forward Curves, Risk Premia and Convenience Yield

In this section we discuss the UK spot and forward data and the connection between the two. Although earlier data is available we will use spot data from March 2003 to Jan 2006. This is because in 1998, the Interconnector, a large pipeline connecting the UK gas entry beach Bacton to the Belgian port of Zeebrugge, came online, creating a link between the UK and Continental gas markets. Once these markets had adjusted to the new setup, UK prices became heavily determined by the factors that determined the European gas price such as the oil market. This structural change means that the pre-2000 data will no longer be relevant to the current market.

In Figure 1 we can see the UK NBP Day-Ahead price from March 2003 onwards. The path can be seen to include a long-term upward drift, a seasonal component (high in the winters and low in the summers) and random shocks throughout. The peaks in spot prices coincide with the coldest periods of each year's winter, usually occurring in January but occurring during March for 2005. They point to the fact that the seasonality in spot gas prices is driven heavily by weather conditions, especially in the winter. In the winter months, the colder weather increases the demand for gas heating from households and businesses, as well as producing adverse conditions for production and supply from the gas
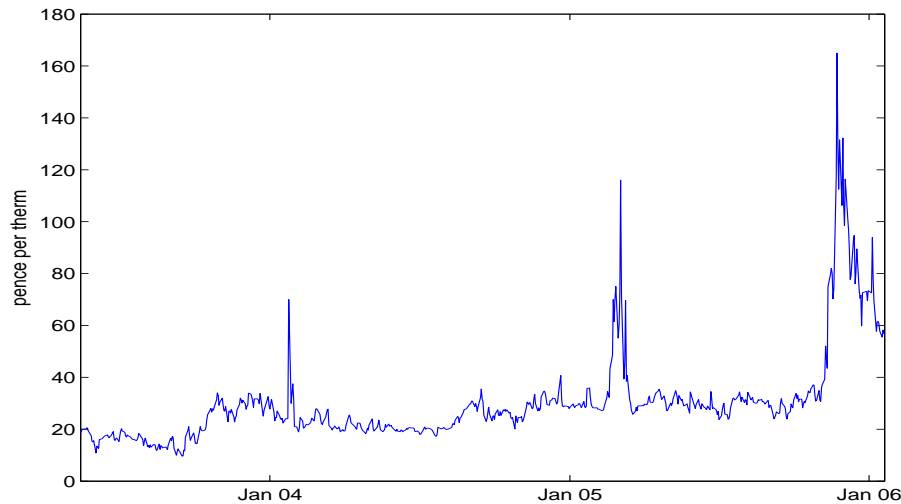
4

Figure 1: Weekday Day-Ahead NBP gas prices: March 2003-Jan 2006

fields in the North Sea. These combine to cause tight supply and demand conditions, which is reflected in higher winter spot prices.

Furthermore, to value derivatives, such as interruptible contracts, we have to be able to model the spot price process under the risk-neutral measure. This naturally requires understanding the market's attitude to risk, as well as the value it places on it. In the absence of a complete market this will have to be estimated by observing the quoted price of derivatives where risk will already have been accounted for. The simplest and most liquidly traded natural gas contingent claims are forward contracts.

Forwards curves in contango, ie positively sloped in T, are associated with times when supplies are plentiful. We can see in Figure 2 that in June the forward curves are in contango going into the winter quarters. In the winter quarters the curve then shifts into backwardation, ie negatively sloped in T, going into the summer. The market can be seen to place a premium on ownership of gas in the winter relative to ownership in the coming summer, ownership coming from production or from gas held in storage. The marked seasonality in the forward curve prompts the question from where does this premium come? Looking beyond the seasonal fluctuations we can also see that the overall slope of the curve is backwardated, a fact that might point to the effect of long-term risk exposure on prices. We will address the matter of long-term risk first.
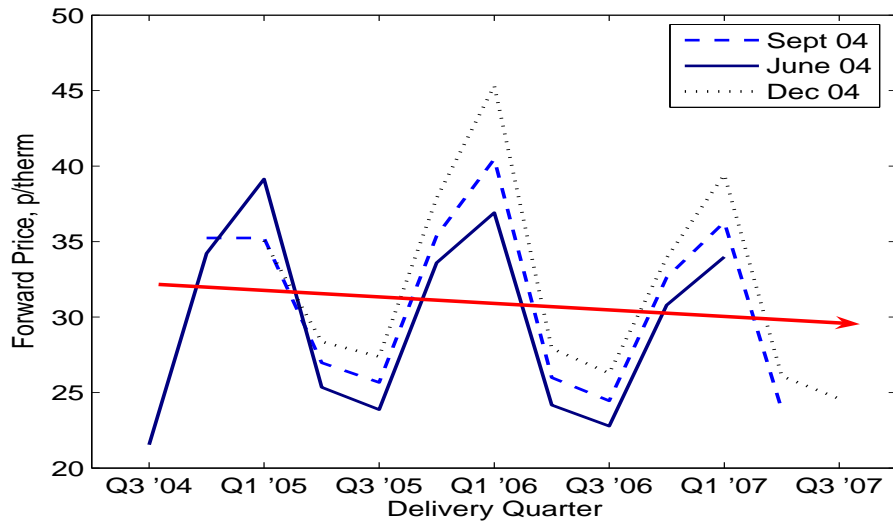
5

Figure 2: Quarter forward curves late-2004

The theory of normal backwardation, originally postulated by Keynes (1930), moved away from traditional backwardation and contango by shifting the emphasis away from whether forward prices are above or below the current spot price, S_t, and on to how forward prices relate to *expected* future spot prices, under the physical measure $\mathbb{P}$. He investigated this relationship by linking spot forecasts to forward prices through a forward risk premium, $\pi_t = \mathbb{E}_t^P[S_T] - F(t,T)$, where $\mathbb{E}_t^P$ is the expectation operator with respect to the physical measure with information up until time t and $F(t,T)$ is the price of the forward at time t with delivery T. When this risk premium is positive then forward prices are below expected spot prices and the forward curve is said to be normally backwardated. If the premium is instead negative, placing forward prices above expected spot prices, then the curve is said to be in normal contango.[2]

Observing the steady backwardated trend in the forward curve we can theorise as to what state the UK gas market is in over the medium and long-term, ie a year or longer into the future. For example, if it is assumed that market conditions are generally the same from year to year, and that expected spot prices will stay at broadly the same levels across time, then it would seem that since the forward curve is actually backwardated, the risk premium

---

[2]In the literature the forward risk premium is also defined as $F(t,T) - \mathbb{E}_t^P[S_T]$.

must be positive and must increase in magnitude with the length or expiry of the forward contract, a situation consistent with normal backwardation.

Normal backwardation such as the one that seems to occur in the UK gas market can be explained through long-run market structure and hedging demands. As a result of the large sunk costs inherent to energy production due to exploration, production and processing, gas producers have effectively "purchased" their supplies for a long period in advance. For example Gazprom's recently developed Yuzhno-Russkoye field in Siberia by itself holds enough gas to supply the entire UK market for fourteen years, see Cahill and Gismatullin (2005). It is natural, therefore, that producers would wish to sell very long-dated forwards to reduce their exposure to adverse changes in the equilibrium gas price over this time. Wholesale gas consumers, on the other hand, do not have such extreme sunk costs to cover and therefore only require shorter-term hedges, purchasing annual or biannual contracts. Therefore, although the hedging demand for customers with short-term positions decreases as the maturity of the forward increases, the hedging demand of producers with long-dated gas exposures does not. This means that these hedges have to be provided by speculators who demand a risk premium as compensation for supplying what is essentially insurance for producers.

Although the risk of long-run changes in price can be examined and evaluated through a Keynesian forward risk premium, short-term price risk must be approached differently. This is due to the fact that natural gas storage, practically unavailable over the long-run, can be exploited in the short-term. Traditionally the effect of storage has been explained through the concept of convenience yield of storage, which allows the application of more traditional arbitrage arguments to forward commodity pricing.

We have to question whether storability and a deterministic convenience yield are suitable assumptions for the UK gas market. One difficulty comes from the fact that the rates at which gas can be injected into and withdrawn from storage systems are limited. Often during the main winter cold snap of the year, withdrawals from UK gas storage sites have been at maximum outflow, whilst the spot price has continued to increase dramatically due to the price inelasticity of short-term gas supply and demand. Parties with gas in storage were prevented from taking full advantage of this price increase due to limitations in the

7

withdrawal capability of the system. The arbitrage opportunities of storage therefore are not as clear or effective as standard theory would suggest.

Further evidence against the assumptions of convenience yield and storage comes from the fact that given a constant convenience yield, forward price volatility must be equal to spot price volatility. This contradicts a well-known and observed property of commodities futures prices called the *Samuelson effect* which states that forward price volatility will decrease as the time to maturity of the futures contract increases. More recent papers such as Schwartz (1997) have built the Samuelson effect into spot models by modelling the convenience yield as a stochastic process in itself.

Dincerler, Khokher, and Simin (2004) state that although a great deal of what drives the convenience yield is still undetermined, it is generally agreed that inventory levels have a strong impact with the marginal convenience yield declining as a function of storage. What is not necessarily agreed upon though, is how levels of inventory affect the price of convenience yield risk. Brennan (1958) suggests that speculators will become wary of holding stocks as the overall level of inventories increase, perhaps wary of being crowded out of arbitrage opportunities as described earlier. To account for this, Brennan suggests speculators will price an increasing risk-adjustment factor into the cost of carrying inventory. Others such as Ribeiro and Hodges (2004) have suggested that during times of comfortable supply, when injections into storage increase, speculators will require lower premia in their expected returns, so risk will be priced more cheaply as inventory increases. We, on the other hand, argue that it is deviations from expected seasonal storage that producers take into account. Producers must pay particular attention to the profile of inflow/outflow of gas from storage. Storage facilities tend to be at their peak approaching the winter season (first couple of weeks of December) and are normally depleted towards the end of April. Hence, producers bear the risk of finding themselves out of line from where seasonal storage levels need to be in order to maximise expected profits from storing gas.

3 The Long-Term / Short-Term Model

Working commodity spot price models, using more than one factor, truly began with Gibson and Schwartz (1990) who introduced a mean-reverting stochastic convenience yield as a second cause of uncertainty in the determination of prices. The spot price itself was modelled as a geometric Brownian motion (GBM). The model was then solved numerically and was shown to be capable of displaying the desired Samuelson effect, with futures contract volatility decreasing as maturity increased.

Schwartz (1997) continued his work with this model proposing

$$dX_t = \left(\mu - \delta_t - \frac{1}{2}\sigma_1^2\right) dt + \sigma_1 dW_1, \tag{1}$$

$$d\delta_t = \kappa(\alpha - \delta_t)dt + \sigma_2 dW_2, \tag{2}$$

with W_1 and W_2 correlated Wiener processes. Here X_t modelled the spot price with drift μ and volatility σ_1, whilst δ_t modelled the stochastic convenience yield, an Orstein-Uhlenbeck process with mean reversion rate, κ, mean reversion level α and volatility σ_2.

A shift away from direct convenience yield modelling occurred when Schwartz and Smith (2000) devised a two-factor model in which log-spot prices were described as the sum of two state variables, a mean-reverting short-term variation component χ_t, and a long-term equilibrium price component ξ_t, modelled as a GBM. They proved that this model was mathematically equivalent to the two-factor convenience model, (1) and (2), with the short-term deviations being related to the convenience yield. This two-factor, long-term/short-term (LT/ST) model, proved to have distinct advantages over the earlier model as it replaced the fairly opaque concept of convenience yield with the simpler idea of short-term deviations from the long-run trend price. The two-factors were related by the correlation between their driving processes, and therefore the model became more 'orthogonal' than the model in (1) where the level of convenience yield δ_t directly affects the evolution of X_t.

We employ the LT/ST model with an added deterministic seasonality function $g(t)$, decomposing spot prices into three components, $\ln S_t = g(t) + \chi_t + \xi_t$. The stochastic components evolve according to the following SDEs:

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi dW_\chi, \tag{3}$$

$$d\xi_t = \mu_\xi dt + \sigma_\xi dW_\xi, \tag{4}$$

with the two driving processes dW_χ and dW_ξ correlated with $dW_\chi dW_\xi = \rho_{\chi\xi} dt$. This model has also been applied by Lucia and Schwartz (2002) to the Scandinavian electricity market.

A report by economic consultancy Global Insight into the UK forward gas market identifies the key real-world factors driving UK gas spot prices (Global Insight 2005). This gives us an indication of what the three components of our model may represent. The report claims that long-run equilibrium ξ_t, is driven not by long-run gas demand or long-run marginal cost of gas but by crude oil prices. This feature is imported from the Continental market through the Interconnector pipeline because European gas prices are index-linked to European oil prices. The seasonality seems broadly to be a reflection of British weather patterns, prices rising as temperatures fall with the onset of winter. Finally, short-term variations χ_t are caused by unusual weather patterns: a prolonged frost for example; unexpected production and transportation problems, such as a gas field experiencing technical difficulties; or rumours of either of these reaching traders and speculators.

To price derivatives, interruptible contracts for example, we need to be able to choose a risk-neutral martingale measure $\mathbb{Q}$, equivalent to the physical measure under which we have already defined the LT/ST model, to model the risk-neutral dynamics of spot prices. In line with most of the commodities literature (see for example Schwartz (1997), Cartea and Figueroa (2005), Benth and Saltyte-Benth (2006)), we introduce two parameters to represent the market prices of risk for short-term deviations and long-term equilibrium price changes, λ_χ and λ_ξ respectively. Hence,

$$d\chi_t = (-\kappa\chi_t - \lambda_\chi)dt + \sigma_\chi dW_\chi^*, \tag{5}$$

$$d\xi_t = (\mu_\xi - \lambda_\xi)dt + \sigma_\xi dW_\xi^*, \tag{6}$$

where dW_χ^* and dW_ξ^* are the increments of Brownian motion under the $\mathbb{Q}$-measure with $dW_\chi^* dW_\xi^* = \rho_{\chi\xi} dt$.

As we have observed, producers with production schedules spanning far into the future bear the long-term forward price risk λ_ξ, of potential price changes damaging the value of these sunk commitments. The situation is reversed in the short-run with producers injecting or withdrawing gas from storage and consumers coming to market with price inelastic supply and demand. This shifts the risk exposure onto consumers whose inflexible consumption and balancing needs could leave them to bear the brunt of short-run price rises in the event of an unexpected reduction in supply. This second source of risk is modelled as λ_χ, the short-term market price of risk.

Above we raised the debate as to how inventories affect the convenience yield risk. The parameter λ_χ in the LT/ST model is directly proportional to the convenience yield risk in the two-factor Schwartz and Gibson model shown in equations (1) and (2). When the short-run deviations from the equilibrium price rises, inventories will often be depleted to take advantage of the higher price and vice-versa. We can therefore examine how changes in inventory affect convenience yield or short-term risk through this relationship by making the short-term risk become a linear function of χ_t:

$$\lambda_\chi = \alpha + \beta\chi_t. \tag{7}$$

The sign of β provides an insight into how the market reacts to changes in short-term prices. For example, if β is negative, χ_t and λ_χ go in opposite directions. We note that with the specification (7), the drift component of the risk-neutral χ_t process becomes

$$-(\kappa + \beta)\chi_t - \alpha,$$

hence we may write the risk-neutral process (5) as

$$d\chi_t = (-\kappa^*\chi_t - \alpha)dt + \sigma_\chi dW_\chi^*,$$

where $\kappa^* = \kappa + \beta$. It is straightforward to see the effects β will have on the risk-neutral mean reversion rate. For example, a positive β implies a higher mean reversion rate. In this risk-neutral world a positive χ_t causes a higher risk premium to be demanded, increasing the magnitude of the negative drift and causing mean reversion to occur more quickly. A negative β means that the risk-neutral world has slower mean reversion than under the physical measure, with the risk-averse customers exposed to short-term increases in prices

11

acting as if those increases will last longer than expected under the physical measure. Intuitively, we would expect that $\beta < 0$ since risk-averse agents will tend, as in standard pricing theory, to give more weight to 'unwanted' outcomes, like a price spike, and less weight to 'favourable' scenarios.

Another interesting property of the market prices of risk, beside the effect of inventories, is their comparative size. The worst-case scenario for a producer concerned about long-term equilibrium changes will be the price falling to zero but consumers exposed to short-run risk face potentially infinite price rises. As we can see in Figure 1 huge spot price rises over short periods of time are not unheard of in the gas market, as shown by the three price spikes visible in the data. These occurred in the winters of 2004, 2005 and early 2006 during times of simultaneous high demand, due to cold weather, and supply problems. We would therefore expect the average magnitude of λ_χ to be larger than λ_ξ. This effect could be exacerbated by there being a partial hedge against long-term price changes, namely the oil market due to the oil-index linked nature of European gas.

3.1   Forward Contract Valuation

We now have the full specification of our risk-neutral spot-price process, $\ln S_t = g(t) + \chi_t + \xi_t$, where

$$d\chi_t \quad = \quad (-\kappa^* \chi_t - \alpha)dt + \sigma_\chi dW_\chi^*, \tag{8}$$

$$d\xi_t \quad = \quad \mu_\xi^* dt + \sigma_\xi dW_\xi^*, \tag{9}$$

and $\kappa^* = \kappa + \beta$, $\mu_\xi^* = \mu_\xi - \lambda_\xi$.

This formulation allows us to value forward contracts $F(t,T)$ on the gas price by taking the expectation of the future spot-price under the risk-neutral measure. The forward price for delivery at maturity T, set at time t is

$$F(t,T) \quad = \quad \exp\left(g(T) + e^{-\kappa^*(T-t)}\chi_t + \xi_t + A(T-t) \right), \tag{10}$$

$$A(T-t) \quad = \quad \mu_\xi^*(T-t) - \left(1 - e^{-\kappa(T-t)} \right) \frac{\alpha}{\kappa^*}$$

$$+ \frac{1}{2}\left(\left(1 - e^{-2\kappa^*(T-t)} \right) \frac{\sigma_\chi^2}{2\kappa^*} + \sigma_\xi^2(T-t) + 2\left(1 - e^{-\kappa^*(T-t)} \right) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa^*} \right).$$

12

4 Estimation of parameters

In this section we discuss the calibration of the parameters, $\kappa, \alpha, \beta, \mu_\xi, \lambda_\xi, \sigma_\chi, \sigma_\xi$ and the correlation between the driving Wiener processes, ρ, to recorded UK market data. We must also approximate the other part of the spot price evolution, the annual seasonality, $g(t)$.

4.1 Contract Selection

We use data from the IPE Natural Gas Futures data published in various Heren Reports (Heren 2001-2006), a daily gas market newsletter. For our futures price data, we have taken two and a half years (Aug 2003 - Jan 2006) of daily *Heren Index* prices for the Month +1, Month +2, ..., Month +6 contracts, rolling the contracts over as one month ends and another begins. As there is no true spot market for gas we will use the Day-Ahead contract, the shortest maturity traded contract, as a proxy for the spot over the same period.[3]

4.2 Seasonality and Risk-Neutral Parameters

When using a relatively small data set, as we are, it is advisable to estimate as few parameters as possible using the Kalman filter, see Harvey (1989). One way in which it is possible to reduce the number of parameters we have to estimate is to remove the seasonality from the data before we apply the filter as in Cartea and Figueroa (2005). At this point we can proceed as usual and remove the seasonality present in spot prices from the data or we can obtain this seasonal component from forward data; we have chosen to use the latter. From equation (10) we see that for long-dated forwards the expression of the log-forward, as a consequence of mean reversion, is given by

$$\ln F(t,T) \sim g(T) + \xi_t + \mu_\xi^*(T-t) + \frac{1}{2}\sigma_\xi^2(T-t) - \frac{\alpha}{\kappa^*} + \frac{\sigma_\chi^2}{4\kappa^*} + \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa^*}. \tag{11}$$

---

[3]In other commodities markets, such as electricity, the day-ahead price is also used as the spot price Escribano, Peña, and Villaplana (2005), Villaplana (2006), Benth, Ekeland, Hauge, and Nielsen (2003), Benth and Koekebakker (2005).
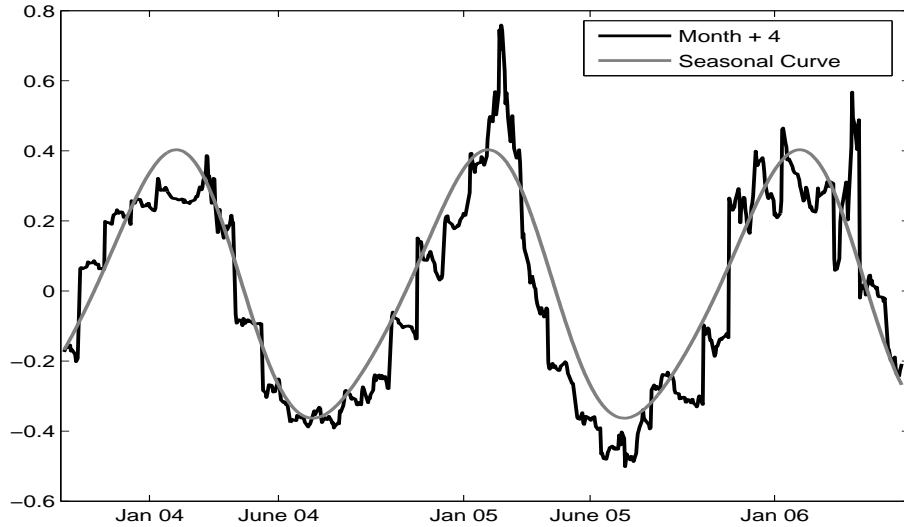
Figure 3: Fitted seasonal curve

Therefore, to isolate the seasonal component $g(T)$ we first detrend the log-prices of a long-dated forward (Month+4) and then fit a second-order Fourier series as in Cartea and Figueroa (2005).[4] We note that using the seasonal component $g(T)$ estimated from spot prices is also desirable, but only as long as current and expected market conditions do not differ from those in the past. On the other hand, forward contracts reflect market expectations of what this seasonal component will be and care must be taken not to include the risk-premium in the estimation of seasonality; a situation we have avoided by using detrended log-forwards with long maturity.

Despite removing the seasonality there are still a number of parameters to estimate via the Kalman filter: $\kappa, \alpha, \beta, \mu_\xi, \lambda_\xi, \sigma_\chi, \sigma_\xi$ and ρ. Running the Kalman filter over the futures data will not be sufficient to give us estimates of every parameter as the filter is unable to distinguish between the relative sizes of λ_ξ and μ_ξ or between κ and β. The filter can only pick up on their combined effects, μ_ξ^* and κ^* respectively.

To try and differentiate between the two we will first run the Kalman filter over the futures series to obtain the $\mathbb{Q}$-measure/risk-neutral parameters. Following that, we will run the filter again over just the spot data to obtain the $\mathbb{P}$-measure/real-world parameters. When running over just the spot data we will allow κ and μ_ξ to change and hold constant

---

[4]We could have used the M+5 or M+6 forwards but these are not as liquidly traded as M+4 forwards.

14

the remaining parameters, σ_χ, σ_ξ and ρ, as they are not affected by the change of measure. The differences between κ^* and κ, μ_ξ^* and μ_ξ will give general estimates of β and λ_ξ.

The Kalman filter was applied to the forward data and the resulting maximum-likelihood (ML) risk-neutral parameters are recorded in Table 1.

	κ^*	α	σ_χ	μ_ξ^*	σ_ξ	ρ
UK NBP	10.18	1.29	1.38	0.15	0.24	-0.33
t-statistic	5.3839	2.6544	16.975	6.3131	4.6830	-1.6507

Table 1: Maximum-Likelihood Risk-Neutral Parameters and t-statistics

It is interesting to note that the correlation between the short and long-term shocks is negative, ie $\rho < 0$. This result may be expected since it implies that long-term movements do not necessarily affect, on average, short-term prices. In other words, a long-term shock that would move the spot price up is compensated by movements in the opposite direction via negative short-term shocks. Furthermore, in the UK gas market it is generally observed that short-term 'blips' in the spot market are immediately followed by a shift in the forward market in the opposite direction (Global Insight 2005), which again supports the finding that $\rho < 0$.

4.3   Real-World Parameters and the Market Prices of Risk

Re-running the Kalman filter over the spot data provides far less certainty with respect to the parameters, but we are able to obtain the parameter ranges shown in Table 2.

κ	β	μ_ξ	λ_ξ
$\approx$ 12 - 25	$\approx$(-15) - (-2)	$\approx$ 0.265 - 0.275	$\approx$ 0.07 - 0.08

Table 2: Parameter ranges under the physical measure

The spot data was sufficient to discern that β is negative but not to determine its magnitude.[5] Figure 4 shows the filter-predicted short-term market-price of risk, λ_χ with a range

---

[5]β was also found to be negative in an extension to Schwartz and Smith (2000) performed on crude oil but there too the data wasn't sufficient to estimate its size.
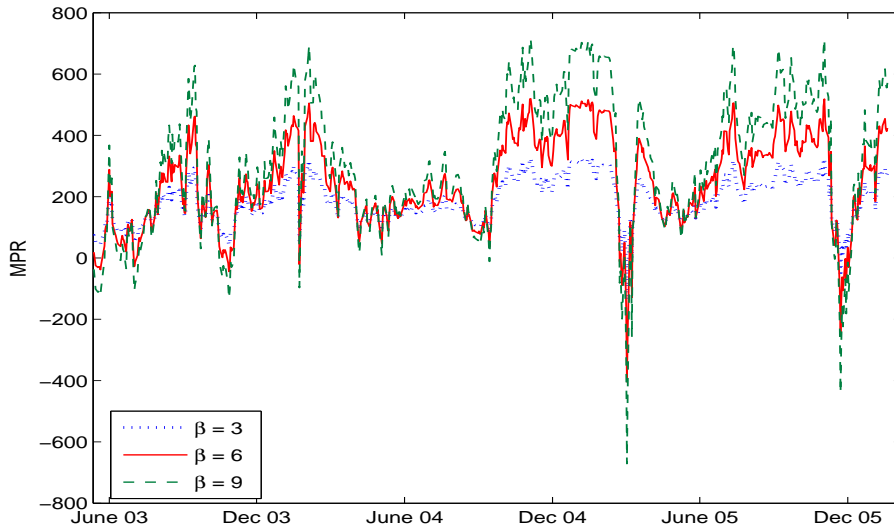
Figure 4: Predicted λ_χ with varying β

of possible βs. We can see the short-term price of risk is very high during the uncommonly mild weather that occurred from October to late-February during the winter of 2004/2005. At that time traders were reported to have believed that supplies in store were more than enough to comfortably withstand the rest of the winter (Heren 2001-2006). This supports our claim that the market demands a higher risk-premium when inventory levels are higher than seasonal storage plans would have predicted, owners of storage being aware that any further gas placed into storage might not actually be used. We remark that this is a similar finding to Brennan (1958) except that we have explained that risk-price adjustment using inventory levels *relative* to seasonal expectations, as opposed to *absolute* inventory levels.

This belief, that winter was essentially over,[6] meant that when the prolonged February/March 2005 cold-spell actually came (occurring at the same time as supply problems in the North Sea) prices reacted violently and the expected ability of storage to cope with demand was completely reversed. The British cold weather was mirrored in Europe, creating high continental demand and causing Interconnector imports to fall to almost nothing. This meant that "No longer able to depend on continental gas for swing volumes, the issue of how much is left in UK storage also became critical."[7] As we can see in Figure 4 the

---

[6]In early February some market commentators were actually mooting starting putting gas back into storage for the next winter.
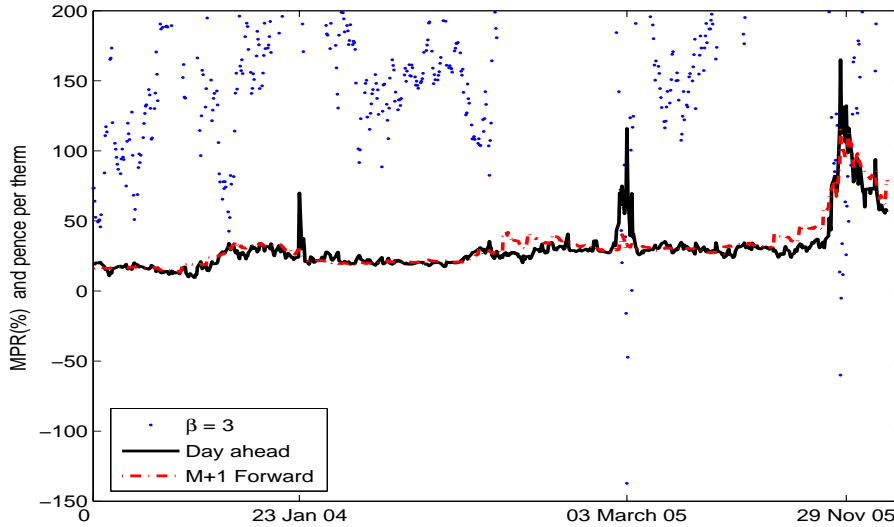
[7]*Heren Report*, March 4th 2005.

16

Figure 5: Predicted λ_χ with $\beta = 3$, Day ahead and M+1 forward prices

sudden, desperate need for storage supplies meant that the market price of risk fell dramatically, actually becoming negative at the height of the spike. This demonstrates the flip-side of our finding: that the market asks for a far lower (or even negative) price on risk during periods when inventory levels are below what seasonal market needs would require them to be. Moreover, Figure 5 shows short-term market price of risk for $\beta = 3$ with spot and M+1 forward prices. Note that during periods of positive short-term shocks like in March and Nov 2005, a large (in magnitude) negative short-term market price of risk (ie $\lambda_\chi < 0$) induces a relatively large positive drift in futures prices, see equation (10), which seems to be corroborated by an increase in the M+1 forward during those periods.

The LT forward risk premium λ_ξ, indicated a state of normal backwardation for UK gas forwards with μ^* being less than μ. It was, as postulated, a lot smaller in magnitude than λ_χ being only about 7%.

17

5 Model Application: Interruptible supply contracts

Now that the spot model has been posed and calibrated we are ready to price interruptible contracts and other contingent claims based on the spot price of gas[8]. Interruptible contracts[9] give the supplier a set number of rights, typically 45, to temporarily cease supplying gas to their customer for periods of a day at a time. The supplier can exercise these rights at their discretion, with exercise giving the potential payoff of

$$Z_t \quad = \quad \max\{S_t - K_t, 0\},$$

where K_t is the contract price for the t^{th} day.

An interruptible contract with one interruption can therefore be viewed as an American or Bermudan call option on the gas spot price, whilst an interruptible contract with many exercises becomes a multi-exercise Bermudan option. However, because of the impossibility of simultaneous exercise, only one exercise is ever active at once. Previously, attempts to value similar structures in energy markets, such as swing contracts, have focused on using methods such as trinomial trees (Jaillet, Ronn, and Tompaidis 2004), but with higher dimensional problems, such as our LT/ST model, make these methods very time-consuming, especially for derivatives with large numbers of exercises. More recently Monte Carlo methods have been applied to early-exercise problems, first by Longstaff and Schwartz (2001) for single-exercise problems, and then for the multi-exercise case, Meinshausen and Hambly (2004), Thanawalla (2005), Ibáñez (2004) and Ibáñez and Zapatero (2004). In this section we will discuss these methods and then apply them to the pricing of interruptible contracts.

5.1 Extended Least-Squares Monte Carlo Method

When deciding whether to exercise an American-style option before it's maturity date the option holder has to determine whether the current *exercise value*, Z_t, of the option is higher

---

[8]At the time of writing we are aware of no other studies into the pricing of gas interruptible contracts.

[9]All terms and conditions are taking from EON Energy's standard contract (E.ON Energy 2005)

than the *continuation value*, Q_t, of instead holding onto the option, where Q_t is the expected value of the option in the next period:

$$Q_t(x) = \begin{cases} \mathbb{E}\left[V_{t+1}\,\middle|\,X_t = x\right] & t < T \\ 0 & t = T \end{cases},$$

where T is the maturity-date of the option.

Knowing the continuation value of the option yields the value function of the option:

$$V_t(x) = \max\left\{Z_t(x), Q_t(x)\right\}, \tag{12}$$

and a natural optimal stopping-time for the option, τ:

$$\tau = \min\left\{t : Z_t > Q_t\right\}.$$

The central idea of the Least-Squares Monte Carlo (LSM) method of Longstaff and Schwartz is to approximate this continuation value function for an American option (and by proxy the value function and optimal-stopping rule of the option) using least-squares regression.

The LSM algorithm has proved robust and successful at pricing options, Moreno and Navas (2003). One important place it falls down though is that as it is driven by an approximation to the optimal stopping rule it can only provide a lower bound to the true value. However, for the LSM to be useful to price interruptible contracts it must be extended from the single-exercise case to the multi-exercise case. This extension has been proposed and implemented by a number of authors: Dörr (2003), Meinshausen and Hambly (2004), Thanawalla (2005) and we will refer to it as Extended Least-Squared Monte Carlo, (XLSM).[10] The XLSM algorithm differs from the LSM algorithm because instead of approximating the optimal stopping-time τ for one exercise we must approximate the optimal stopping-policy $\pi = \{\tau_n, \ldots, \tau_1\}$ for n separate exercises.

5.2 Approximating the Continuation Values

For the XLSM we approximate the optimal stopping rule by calculating not just one continuation value function at every time-point but n continuation value functions, $Q_t^n(x), \ldots, Q_t^1(x)$,

---

[10]We will use the notation and algorithms of Meinshausen and Hambly (2004).

one for every possible remaining exercise amount. The decision rule for the multi-exercise case becomes "exercise the m^{th} exercise if the marginal continuation value of the m^{th} exercise, $\Delta Q_t^m(x)$, is less than the exercise-payoff, Z_t", ie:

$$
\begin{aligned}
Z_t(x) \quad &> \quad \Delta Q_t^m(x), \\
&= \quad Q_t^m(x) - Q_t^{m-1}(x).
\end{aligned}
$$

By once again simulating a large number of sample paths we are ready to price options. Now the XLSM algorithm works very similarly to the LSM algorithm, starting at maturity, T, and iterating backwards. At time t we have continuation value functions for $m = 1, \ldots, n$ exercises for every time-point larger than t. With these we can calculate the optimal future cash-flows for each path, given each of the possible amounts of exercises $m = 1, \ldots, n$ that one could have at that time. We then use these to perform the least-squares regression, as in the LSM algorithm, to calculate $Q_t^n(x), \ldots, Q_t^1(x)$ respectively. Once we have iterated back to $t = 0$ we have an approximation to the marginal continuation value of each exercise and therefore we have an approximation to the optimal stopping policy. With these we are able to value multi-exercise claims.

5.3   XLSM Lower and Upper Bounds

Because of the numerical approximation error in the LSM and XLSM algorithms it desirable to construct both an upper and lower bound for the option value. With the approximated optimal stopping policy calculated in section 5.1 we can receive a lower bound for the option by simulating a number of price-paths and determining what their average payoff would have been under our sub-optimal, approximated stopping policy.

Calculating upper bounds in LSM, Rogers (2002), and XLSM, Meinshausen and Hambly (2004), is based around the idea that the value of American-style options can be expressed as the infimum over a family of expectations. For completeness we will reproduce the main theorem presented in Rogers (2002) for the case of one-exercise options. This provides the intuition behind the upper bound and we will then state the theory in the case of the multi-exercise problem as presented in (Meinshausen and Hambly 2004) and explore its practical implementation.

Theorem 1 *Consider the amount paid to the holder of an American option at exercise as an adapted process, $(Z_t)_{0 \leq t \leq T}$, with a finite time horizon, $T > 0$, defined on a filtered probability space, $(\Omega, \Im, (\Im)_{0 \leq t \leq T}, \mathbb{Q})$. Then the time-0 value, Y_0^* of the American option is given by:*

$$Y_0^* = \inf_{M \in H_0^1} \mathbb{E}\left[\sup_{0 \leq t \leq T} (Z_t - M_t)\right], \tag{13}$$

where H_0^1 is the space of martingales, M, such that $\sup_{0 \leq t \leq T} |M_t| \in L^1$ and $M_0 = 0$.

This theorem provides a natural method for calculating the upper bound. First, choose and construct a martingale, $M \in H_0^1$. And second, evaluate $\mathbb{E}\left[\sup_{0 \leq t \leq T} (Z_t - M_t)\right]$ using numerical simulation techniques.

The multi-exercise upper bound is generated in a similar way through calculating the upper bounds of the marginal values of each individual exercise, $\Delta V_0^{\uparrow,n}, \ldots, \Delta V^{\uparrow,1}$. The theory was introduced by Meinshausen and Hambly (2004), whose major result is as follows.

Theorem 2 *The marginal value $\Delta V_0^{*,n}$ is equal to:*

$$\Delta V_0^{*,n} = \inf_{\pi} \inf_{M \in H_0^1} \mathbb{E}\left[\max_{u \in (\daleth \setminus \{\tau_{n-1}, \ldots, \tau_1\})} (Z_u - M_u)\right], \tag{14}$$

where $\daleth = \{0, \ldots, T\}$ is the set of possible exercise dates and $0 \leq \tau_{n-1} < \ldots < \tau_1$ are stopping times and $M_t \in H_0^1$. Furthermore, the infimum is attained by the optimal policy of stopping times, π^*, and the optimal martingale M^* with $M_0^* = 0$, whose increment at time t is given by:

$$M_t^* - M_{t-1}^* = \Delta V_t^{*,m} - \mathbb{E}_{t-1}\left[\Delta V_t^{*,m}\right], \tag{15}$$

where m is the smallest natural such that $t > \tau_m$.[11]

This theorem has the advantage of identifying the optimal martingale. Given our approximated optimal stopping policy we generate an approximation of the martingale used to value the upper bound of the marginal value of the n^{th} exercise, $\Delta V_0^{\uparrow,n}$, along a certain path by means of the following algorithm:

---

[11]There is a typo here in the original paper corrected here after correspondence with the author.

1. First we must calculate the stopping-times, $\{\tau_{n-1}, \ldots, \tau_1\}$, along the path, $\{X_t\}_{t=1,\ldots,T}$, that would have occurred under our approximated stopping policy if we started with $n-1$ exercises.

2. Next we have to approximate the martingale increments in (15) along our path. We use the stopping-times generated, as well as defining $\tau_n = 0$, to choose the appropriate continuation value function to use to calculate $\Delta V_t^m(X_t)$:

Time	Continuation Value Function
$0 < t \leq \tau_{n-1}$	$Q_t^n(x)$
$\tau_{n-1} < t \leq \tau_{n-2}$	$Q_t^{n-1}(x)$
$\vdots$	$\vdots$
$\tau_m < t \leq \tau_{m-1}$	$Q_t^m(x)$

Table 3: Appropriate Continuation Functions Used For Upper Bounds

3. To generate the martingale increment $M_t - M_{t-1}$ equation (15) requires us to approximate two values: ΔV_t^m and $\mathbb{E}_{t-1}[\Delta V_t^m]$. The first of these is given by

$$
\begin{aligned}
\Delta V_t^m(X_t) &= V_t^m(X_t) - V_t^{m-1}(X_t), \\
\text{where} \quad V_t^m(x) &= \max\left\{ Z_t(x) + Q_t^{m-1}(x), Q_t^m(x) \right\},
\end{aligned}
$$

and the second is approximated by Monte Carlo simulation:

$$
\mathbb{E}_{t-1}[\Delta V_t^m] \approx \frac{1}{k} \sum_{i=1}^k \Delta V_t^m(X_t^i), \tag{16}
$$

where X_t^i, $i = 1, \ldots, k$ are independent one time-step evolutions of the path conditional from X_{t-1}.

4. Once the martingale is generated a sample of the upper bound for the n^{th} exercise is evaluated by taking

$$
\Delta V_0^{\uparrow, n} = \max_{u \in (\daleth \setminus \{\tau_{n-1}, \ldots, \tau_1\})} (Z_u - M_u),
$$

The upper bound of the exercise value is the sample mean of a number of such samples, each requiring its own martingale. Clearly this can become very computationally expense

22

for an option with a large number of exercises. Once $\Delta_0^{\uparrow,m}$ is calculated $\forall m = 1, \ldots, n$, then the overall upper bound is given by,

$$V_0^{\uparrow,n} = \sum_{m=1}^{n} \Delta V_0^{\uparrow,m}. \tag{17}$$

5.4 Implementation and Benchmarking

We implemented the XLSM algorithm in *MATLAB*. The implementation was designed to be reusable for different driving processes. This allowed the actual XLSM algorithm to be tested against the results presented in Meinshausen and Hambly (2004) for an AR(1) process with 1000 time-steps and a number of exercise amounts ranging from 1 to 100. The results for the lower bound of our implementation and the results presented in Meinshausen and Hambly (2004) were within 0.5% of one another whilst the upper bound estimates were within $1 - 2\%$ of each other, though our implementation reported a higher standard error for the upper bound. The lower bound was also tested, using GBM as the driving process, against American put option prices reporting prices within 0.5% of those calculated by finite-difference methods.

5.5 Contract Value and Interruptible Discount

The value of an interruptible contract can be expressed in one of two ways: either as a straight monetary value (per therm of gas) or as a gas price discounted relative to the cost of a firm supply of gas. Clearly these two values have a simple relationship: a one-pence per therm discount on a year-long gas contract is worth the present value of a cash-flow of 1p for every day of that year. Taking the annual interest rate as $r = 4.5\%$ this means that a penny discount is worth $\sum_{t=1}^{365} e^{-rt/365} \approx 357\text{p}$ at the start of the contract.[12] Given

---

[12]This is only a rough approximation and does not take into account the fact that no payments are made on an interruption day.
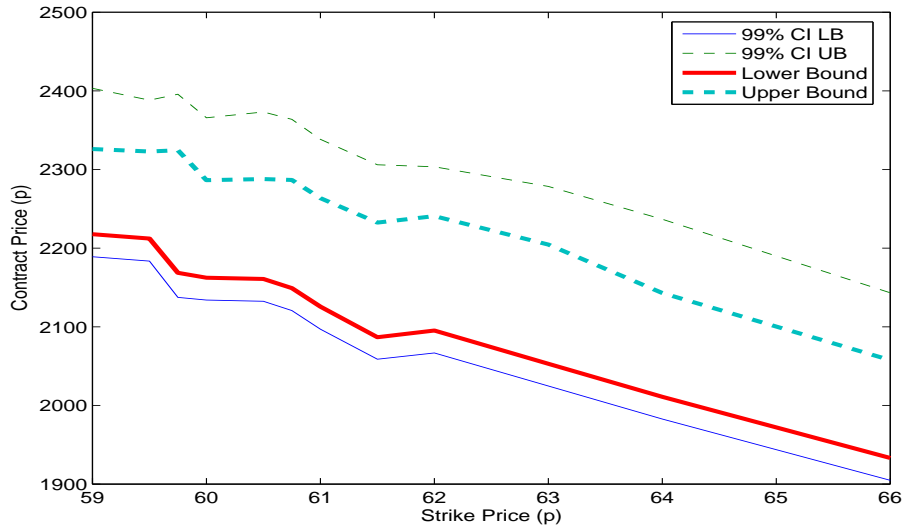
Figure 6: 99% Confidence interval for the price of a 45 interruption winter start contract.

a gas price for firm supply, K_F, and a set number of interruptions, n, we can calculate the corresponding price for interruptible gas, K_I, by finding the point where:

$$\underbrace{V^n(K_I)}_{\text{Interruption Value}} = \underbrace{357 \times (K_F - K_I)}_{\text{Discount Value}}.$$

5.5.1 Strike Price

The first contract variable we will alter is the contract strike price, K_I. Figure 6 shows a 99% confidence interval for the value of an interruptible contract with 45 exercises spanning from October 1st 2005 to September 30th 2006. The slope of the lower bound curve, whilst seemingly linear, is not as steep as might be expected. This is because during the typical life of a contract not every exercise available will actually be used.

Using these results we can see what discount we would offer the firm on the price of gas. On October 1st, 2005 the Gas-Year Ahead 05/06 contract was trading at 66.75p/per therm. Using the results in Figure 6 we can create a confidence interval for the appropriate interruptible gas price. As we can see in Table 4 the lower bound contract prices suggest setting $K_I = 60.75$p/therm, a 5.9625p discount. The upper bound prices suggest setting $K_I = 60.25$p/therm, a 6.4625p discount. Moreover, in the example shown here, for

Strike K_I	60	60.25	60.5	60.75
Contract CI	[2164.35, 2298.45]	[2162.2, 2286.45]	[2160.8, 2287.8]	[2149.25, 2286.75]
Discount CI	[6.067, 6.443]	[6.061, **6.410**]	[6.057, 6.413]	[**6.025**, 6.410]
Discount	6.7125	**6.4625**	6.2125	**5.9625**

Table 4: Appropriate discount against firm price of 66.75 p/therm. All figures are in pence.

strike prices between 60.25 and 60.75 around 42.5 interruptions are used. As expected, the inclusion of further interruptions increases the price of the interruptible contract but the marginal increase in the contract value is decreasing in a linear fashion with respect to strike for ranges $K_I \in [40, 100]$.

6 Conclusions and Further Work

In this article we make two main contributions. Firstly, we explored the short-term market price of risk in the gas market. In our analysis we explain changes in the risk-price, previously related to the absolute level of inventory in the literature Brennan (1958), in terms of the inventory level relative to the expected seasonal storage patterns. Motivating this, is the intuition that in the UK gas industry supplies are put into storage over the summer and withdrawn over the winter in a pre-planned yearly cycle designed to help the market cope with the tight winter supply/demand situation. What matters therefore is not whether supplies are being put into or taken out of storage but how these injections and withdrawals compare with what the market expectes them to be. In our work we tie these unexpected deviations, modelled by χ_t, to the short-term market price of risk, $\lambda_\chi = \alpha + \beta \chi_t$.

Running a Kalman filter over a data set of UK forward and spot gas prices we find β to be negative. This implies that in the UK natural gas market, the short-term market price of risk is higher when storage inventories are being depleted more slowly than the seasonal storage plans suggest they should be. Essentially, if there is more gas in storage than the market expects it will require a higher risk premium to be paid, for the duration of a period of tight supply, in order to convince a market participant to put further gas into storage than would have been demanded if there was a shortage in inventories relative to seasonal need.

The second contribution comes through the use of the multi-exercise Extended Least Squares Monte Carlo algorithm to value interruptible gas contracts. Given a price per therm for the firm supply of gas we are able to give upper and lower bounds for the appropriate per therm discount that a consumer should demand as compensation for entering into a standard interruptible contract. The upper and lower bounds are found to be within 0.6p of one another, a difference of less than 1% relative to our suggested price for interruptible gas.

Both of these contributions open opportunities for further work. A larger time series of forward data and a calibration method that accounted for the continuous delivery of gas forwards, such as a particle filter, would allow a better calibration of the LT/ST model to market data. Moreover, it is straightforward to see that a more realistic model should include jumps in the short-term deviations process that will in turn affect the price of interruptible contracts. Although the incorporation of jumps seems a natural way to model gas price dynamics, the application of the standard Kalman filter would not be possible. Hence, the richness of the insights provided by the understanding of the short-term market price of risk would have been more difficult to obtain if a non-Gaussian model had been chosen.

References

Benth, F.E., L. Ekeland, R. Hauge, and B.N. Nielsen, 2003, A note on arbitrage-free pricing of forward contracts in energy markets, *Applied Mathematical Finance* 10, 325–336.

Benth, F.E., and S. Koekebakker, 2005, Stochastic modeling of financial electricity contracts, Working paper.

Benth, F.E., and J. Saltyte-Benth, 2006, Analytical Approximation for the Price Dynamics of Spark Spread Options, Forthcoming: Studies in Nonlinear Dynamics and Econometrics.

Brennan, M., 1958, The Supply of Storage, *American Economic Review* 48, 50–72.

Cahill, T., and E. Gismatullin, 2005, Russia's Power Play, Bloomberg.

Cartea, Á., and M. G. Figueroa, 2005, Pricing in Electricity Markets: a Mean Reverting Jump Diffusion Model with Seasonality, *Applied Mathematical Finance* 12, 313–335.

Dincerler, C., Z. Khokher, and T. Simin, 2004, The Convenience Yield and Risk Premia of Storage, Working paper.

Dörr, U., 2003, aluation of Swing Options and Examination of Exercise Strategies by Monte Carlo Techniques, Msc thesis University of Oxford.

E.ON Energy, 2005, *Interruptible Gas For Business: Terms and Conditions of Supply*.

Escribano, Á., J.I. Peña, and P. Villaplana, 2005, Modelling electricity prices: International evidence, Working paper.

Gibson, R., and E. Schwartz, 1990, Stochastic Convenience Yield and the Pricing of Oil Contingent Claims, *The Journal of Finance* 45, 959–976.

Global Insight, 2005, Ensuring Effective and Efficient Forward Gas Markets: A Report to the DTI, Working paper, Global Insight.

Harvey, A.C., 1989, *Forcasting, Structural Time Series Models and the Kalman Filter*. (Cambridge).

Heren, 2001-2006, The Heren Energy, Working paper, Heren.

Ibáñez, A., 2004, Valuation by Simulation of Contingent Claims with Multiple Early Exercise Opportunities, *Mahtematical Finance* 14, 223–248.

Ibáñez, A., and F. Zapatero, 2004, Monte Carlo Valuation of American Options through Computation of the Optimal Exercise Frontier, *Journal of Financial and Quantitative Analysis* 39.

Jaillet, P., E. Ronn, and S. Tompaidis, 2004, Valuation of Commodity-Based Swing Options, *Management Science* 50.

Keynes, J.M., 1930, *A Treatise on Money* vol. 2. (Palgrave Macmillan).

Longstaff, F., and E. Schwartz, 2001, Valuing American Options by Simulation: A Simple Least-Squares Approach, *Review of Financial Studies* 14, 113–147.

Lucia, J., and E. Schwartz, 2002, Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange, *Review of Derivatives Research* 5, 5–50.

Meinshausen, N., and B. Hambly, 2004, Monte Carlo Methods for the Valuation of Multiple Exercise Options, *Mathematical Finance* 14, 557–583.

Moreno, M., and J. Navas, 2003, On the Robustness of Least-Squares Monte Carlo for Pricing American Derivatives, *Review of Derivative Research* 6, 107–128.

Ribeiro, D., and S. Hodges, 2004, Equilibrium Model for Commodity Prices: Competitive and Monopolistic Markets, Working paper.

Rogers, L.C.G., 2002, Monte Carlo Valuation of American Options, *Mathematical Finance* 12, 271–286.

Schwartz, E., 1997, The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging, *The Journal of Finance* 52, 923–973.

Schwartz, E., and J. Smith, 2000, Short-Term Variations and Long-Term Dynamics in Commodity Prices, *Management Science* 46, 893–911.

Thanawalla, R., 2005, Valuation of Swing Options using an Extended Least Squares Monte Carlo Algorithm, Working paper.

Villaplana, P., 2006, Valuation of electricity forward contracts: the role of demand and capacity, Working paper.