

Implications of a Regime-Switching Model on Natural Gas Storage Valuation and Optimal Operation

Zhuliang Chen* Peter A. Forsyth †

September 13, 2007

Abstract

In this paper, we propose a one-factor regime-switching model for the risk adjusted natural gas spot price and study the implications of the model on the valuation and optimal operation of natural gas storage facilities. We calibrate the model parameters to both market futures and options on futures. Calibration results indicate that the regime-switching model is a better fit to market data compared to a one-factor mean-reverting model similar to those used by other authors to value gas storage. We extend a semi-Lagrangian timestepping scheme from Chen and Forsyth (2007) to solve the gas storage pricing problem, essentially a stochastic control problem, and conduct a convergence analysis of the scheme. Numerical results also indicate that the regime-switching model can generate operational strategies for gas storage facilities that reflect the existence of multiple regimes in the market as well as the regime shifts due to various exogenous events.

Keywords: Gas storage, regime switching, stochastic control

JEL Classification: G12, G13

1 Introduction

As noted in Pilopović (1998), natural gas prices exhibit seasonality dynamics due to fluctuations in demand. As such, natural gas storage facilities are constructed to provide a cushion for such fluctuations by releasing natural gas in storage in seasons with high demand.

The valuation of gas storage facilities is characterized as a stochastic control problem resulting in Hamilton-Jacobi-Bellman (HJB) equations. Recently, several authors Chen and Forsyth (2007); Ahn et al. (2002); Thompson et al. (2003); Ware (2005); Ware and Li (2006); Li (2006); Carmona and Ludkovski (2005); Barrera-Esteve et al. (2006) have discussed the no-arbitrage value of natural gas storage facilities (or, equivalently, the value of gas storage contracts for leasing the storage facilities). Their work has almost exclusively assumed that natural gas spot prices follow one-factor mean-reverting processes. However, as demonstrated in Schwartz (1997); Jaillet et al. (2004)

*David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1 e-mail: z4chen@uwaterloo.ca

†David R. Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1 e-mail: paforsyt@uwaterloo.ca

and again this paper, one-factor mean-reverting models do not seem to be able to capture the dynamics of typical gas forward curves. Consequently, we need to resort to other more complex stochastic models for natural gas prices in order to more accurately price gas storage contracts.

A number of multi-factor models for the natural gas spot price are suggested in (Xu, 2004; Manoliu and Tompaidis, 2002; Jaillet et al., 2004; Cartea and Williams, 2006). More general multi-factor models for commodity spot prices are developed in (Schwartz, 1997; Cortazar and Schwartz, 2003). Nevertheless, it is computationally expensive to apply the two-factor and three-factor commodity spot price models in (Schwartz, 1997; Cortazar and Schwartz, 2003) to price complex commodity derivatives such as the gas storage contracts, although they seem able to fit the market futures prices.

Consequently, we will focus on one-factor regime-switching models for natural gas spot prices. Initially proposed in Hamilton (1990), a regime-switching model consists of several regimes; within each regime the gas price follows a distinct stochastic process. The price process can randomly shift between these regimes due to various reasons, such as changes of weather conditions, alterations of demand and supply, and surprise events such as political instability. Regime-switching models have been used in several areas. For example, Hardy (2001) develops a regime-switching model for equities. Gray (1996) and Bansal and Zhou (2002) use regime-switching processes to model term structures of interest rates. Various regime-switching models are calibrated to electricity spot prices in (Deng, 2000; Huisman and Mahieu, 2003; de Jong and Huisman, 2002; de Jong, 2005; Schindlmayr, 2005).

Our goal in this paper is twofold. First, we develop a simple one-factor regime-switching model that outperforms a typical one-factor mean-reverting model in terms of capturing the dynamics of the gas forward curves. Second, we show that we can price gas storage contracts under the proposed regime-switching model and furthermore, we demonstrate that the regime-switching model can produce operational strategies for the storage facility that reflect the existence of multiple regimes in the market (if we believe this is true) as well as the regime shifts due to various exogenous events. This results in a richer set of operational strategies compared to one-factor mean-reverting models.

Our one-factor regime-switching model consists of two regimes. By adjusting parameter values, the deseasoned process in each regime is either a mean-reverting process or a geometric Brownian motion (GBM) like process with a positive/negative drift. Hence this produces several variations of the model. We are interested in three variations: an MRMR variation where the processes in both regimes are mean-reverting; an MRGBM variation where the process in one regime is mean-reverting while that in the other regime follows a GBM with a positive drift; a GBMGBM variation where the process in one regime follows a GBM with a positive drift while that in the other regime follows a GBM with a negative drift. The MRMR variation is able to generate a behavior similar to the two-factor mean-reverting model given in Xu (2004), while the GBMGBM variation can mimic the behavior of the two-factor model in Schwartz (1997). As a comparison, we also consider a one-factor mean-reverting model, denoted by MR, which is similar to the models proposed in (Pilopović, 1998; Xu, 2004).

We emphasize here that our purpose is not to develop the *ultimate* model for gas prices. Our objective is to propose a model which can (at least qualitatively) fit forward curves, yet is not too computationally expensive to use in gas storage valuation.

The main results of this paper are summarized as follows:

- Among the three gas price models that we examine, the calibration results show that the MRMR and MRGBM variations of the regime-switching model are capable of fitting the

market gas forward curves more accurately than the MR model. The GBMGBM variation does not appear to be consistent with market data.

- We calibrate model parameters to both market futures and options. We use a two phase calibration process. Under these models, a subset of the parameters can be obtained by calibration to the forward curves. The remaining parameters can be determined by calibration to options. As a result, the computational requirements of the calibration process are reduced compared to more general models. Note that since we are interested in valuation and operation of gas storage, we calibrate to futures and options prices, which gives us the \mathbb{Q} measure parameters directly. This is, of course, distinct from the econometric approach of examining spot price time series, which would generate \mathbb{P} measure parameters.
- We use the calibrated regime-switching model to price the value of a natural gas storage facility, resulting in coupled HJB equations. We extend the fully implicit, semi-Lagrangian timestepping scheme from Chen and Forsyth (2007) to solve the nonlinear pricing equations. Provided a strong comparison result holds, we prove the convergence of the scheme to the viscosity solution of the equations using the results in (Barles and Souganidis, 1991; Barles, 1997). The numerical results demonstrate that the scheme converges to the solution at a first order rate.
- We study the implications of the proposed models on the optimal operational strategies for gas storage facilities. Our observations indicate that the regime-switching model, in contrast to one-factor mean-reverting models, is able to produce distinct controls for each regime that are consistent with the dynamics of gas spot prices in that regime. This effect is, of course, absent in the standard MR model. However, whether to pick the MRMR or MRGBM variation depends on one's perspectives about gas price dynamics. It is difficult to distinguish between these two models if we simply calibrate to market data.

2 Natural Gas Spot Price Models

In this section, we specify two one-factor models that we use to examine the dynamics of the natural gas spot price. Since we are interested in pricing derivatives on natural gas, we will consider directly the risk adjusted (or risk neutral) price processes with parameters given under the \mathbb{Q} measure.

2.1 One-factor mean-reverting model (MR model)

Let P denote the natural gas spot price. In the MR model, the gas spot price follows a mean-reverting stochastic process with the seasonality effect represented in the drift term. The risk adjusted gas spot price is modeled by a stochastic differential equation (SDE) given by

$$dP = \alpha(K_0 - P)dt + \sigma PdZ + S(t)Pdt, \quad (2.1)$$

$$S(t) = \beta_A \sin(2\pi(t - t_0 + C_A(t_0))) + \beta_{SA} \sin(4\pi(t - t_0 + C_{SA}(t_0))), \quad (2.2)$$

where

α $\alpha > 0$ is the mean-reversion rate,

K_0 $K_0 > 0$ is the long-term equilibrium price,

$\sigma > 0$ is the volatility,

dZ is an increment of the standard Gauss-Wiener process,

$S(t)$ is a time-dependent term so that $S(t)Pdt$ is the price change at time t contributed by the seasonality effect. Note that multiplying $S(t)$ with P guarantees the price of natural gas always stays positive,

β_A is the annual seasonality parameter,

t_0 is a reference time satisfying $t_0 < t$.

$C_A(t_0)$ is the annual seasonality centering parameter for t_0 . We define

$$C_A(t_0) = A_0 + D(t_0), \quad (2.3)$$

where A_0 is a constant time adjustment parameter obtained through calibration; $D(t_0)$ is the distance between the reference time t_0 and the first date in January in the year of t_0 . Thus, by calibrating the value of A_0 , we are able to determine the evolution of the annual seasonality effect over time.

β_{SA} is the semiannual seasonality parameter,

$C_{SA}(t_0)$ is the semiannual seasonality centering parameter for t_0 . Similar to the definition of $C_A(t_0)$, we define

$$C_{SA}(t_0) = SA_0 + D(t_0), \quad (2.4)$$

where the constant time adjustment parameter SA_0 is obtained from a calibration process.

This simple model is considered by several authors (Pilopović, 1998; Xu, 2004), although the seasonality feature is handled in a slightly different manner.

Remark 2.1 (Effect of the seasonality term on gas price dynamics). *We can rewrite equation (2.1) as*

$$dP = \alpha K_0 dt + (S(t) - \alpha)Pdt + \sigma PdZ. \quad (2.5)$$

Since $-(|\beta_A| + |\beta_{SA}|) \leq S(t) \leq |\beta_A| + |\beta_{SA}|$ according to equation (2.2), if

$$|\beta_A| + |\beta_{SA}| > \alpha, \quad (2.6)$$

then there exists certain periods of time within which $S(t) - \alpha > 0$. In this case, if P is large and $(S(t) - \alpha)Pdt \gg \alpha K_0 dt$ in equation (2.5), then the process (2.1) becomes a GBM process with positive drift rate due to the strong seasonality effect. At other times, the process is mean-reverting. Note that the deseasoned process (i.e., setting $S(t) = 0$ in SDE (2.1)) is a mean-reverting process. As indicated in our calibration results in Section 3.1.4, condition (2.6) is typically satisfied by the calibrated parameters.

2.2 Regime-switching model

In order to capture the gas price dynamics more accurately than a one-factor model, Schwartz (1997); Xu (2004) propose different two-factor models for the natural gas spot price. In this subsection, we present a one-factor regime-switching model that is able to exhibit behavior similar to the models introduced in Schwartz (1997); Xu (2004) without introducing an additional stochastic factor.

Roughly speaking, our model consists of two regimes; each regime corresponds to a distinct stochastic process (with the same stochastic factor). At any time, the natural gas spot price follows one of these two processes. However, the price process can jump to another regime with some finite probability.

The switch between two regimes can be modeled by a two-state continuous-time Markov chain $m(t)$, taking two values 0 or 1. The value of $m(t)$ indicates the regime in which the risk adjusted gas spot price resides at time t . Let $\lambda^{0 \rightarrow 1} dt$ denote the probability of shifting from regime 0 to regime 1 over a small time interval dt , and let $\lambda^{1 \rightarrow 0} dt$ be the probability of switching from regime 1 to regime 0 over dt . Then $m(t)$ can be represented by

$$dm(t) = (1 - m(t-))dq^{0 \rightarrow 1} - m(t-)dq^{1 \rightarrow 0}, \quad (2.7)$$

where $t-$ is the time infinitesimally before t , and $q^{0 \rightarrow 1}$ and $q^{1 \rightarrow 0}$ are the independent Poisson processes with intensity $\lambda^{0 \rightarrow 1}$ and $\lambda^{1 \rightarrow 0}$, respectively.

In the regime-switching model, the risk adjusted natural gas spot price is modeled by an SDE given by

$$dP = \alpha^{m(t-)}(K_0^{m(t-)} - P)dt + \sigma^{m(t-)}PdZ + S^{m(t-)}(t)Pdt, \quad (2.8)$$

$$S^{m(t-)}(t) = \beta_A^{m(t-)} \sin(2\pi(t - t_0 + C_A(t_0))) + \beta_{SA}^{m(t-)} \sin(4\pi(t - t_0 + C_{SA}(t_0))). \quad (2.9)$$

As indicated in equations (2.8-2.9), within a regime $k \equiv m(t-)$ the gas spot price follows the process (2.1-2.2) with parameters $\alpha^k, K_0^k, S^k(t), \sigma^k$ (but the signs of α^k and K_0^k are not constrained). Meanwhile, the stochastic factors for the two regimes are perfectly correlated. Note that we assume that the centering parameters $C_A(t_0)$ and $C_{SA}(t_0)$, as given in equations (2.3) and (2.4), respectively, are identical for two regimes in order to reduce the number of calibrated parameters.

Remark 2.2 (Mean-reverting or GBM-like process). *From the model (2.8-2.9), the deseasoned spot price in regime $m(t-)$ can follow either a mean-reverting process or a GBM-like process by setting parameter values.*

If we choose $\alpha^{m(t-)} > 0$ and $K_0^{m(t-)} > 0$, then the deseasoned gas price (obtained from setting the seasonality term $S^{m(t-)}(t) = 0$ in SDE (2.8)) follows a mean-reverting process

$$dP = \alpha^{m(t-)}(K_0^{m(t-)} - P)dt + \sigma^{m(t-)}PdZ \quad (2.10)$$

with equilibrium level $K_0^{m(t-)}$ and mean-reversion rate $\alpha^{m(t-)}$.

If we set $K_0^{m(t-)} = 0$ in equation (2.8), then the deseasoned gas price SDE becomes

$$dP = -\alpha^{m(t-)}Pdt + \sigma^{m(t-)}PdZ. \quad (2.11)$$

This is a GBM-like process. Specifically, if the drift coefficient $-\alpha^{m(t-)} > 0$, then SDE (2.11) is a standard GBM process, i.e., gas price P will drift up at a rate $|\alpha^{m(t-)}|$ at time t ; if $-\alpha^{m(t-)} < 0$, then the gas price will drift down at a rate $|\alpha^{m(t-)}|$.

2.2.1 Variations of the regime-switching model

As indicated in Remark 2.2, the deseasoned spot price in each regime can follow either a mean-reverting process or a GBM-like process. Consequently, there exist many possible variations of the regime-switching model by choosing different combinations of the stochastic processes in two regimes. We are interested in the following three variations:

MRMR variation

The processes in both regimes are mean-reverting with different equilibrium levels, i.e., $K_0^k > 0$, $\alpha^k > 0$, $k \in \{0, 1\}$ in SDE (2.8). In this variation, the equilibrium level of the gas spot price switches between two constants, K_0^0, K_0^1 , which thus creates a sort of mean-reverting effect on the equilibrium level. This simulates the behavior of the equilibrium price in the two-factor model proposed by Xu (2004), where the gas spot price P follows a one-factor mean-reverting process and its equilibrium price evolves over time according to the other one-factor mean-reverting process.

MRGBM variation

The process in one regime is mean-reverting while the other regime is a GBM process with a positive drift, i.e., $K_0^0 > 0, K_0^1 = 0, \alpha^0 > 0, \alpha^1 < 0$ in SDE (2.8). The mean-reverting regime represents the normal price dynamics, and the GBM regime can be regarded as the sudden drifting up of the gas price driven by exogenous events.

GBMGBM variation

The processes in both regimes are GBM processes with a positive drift in one regime and a negative drift in the other, i.e., $K_0^0 = K_0^1 = 0, \alpha^0 < 0, \alpha^1 > 0$ in SDE (2.8). This simulates the behavior of the two-factor model in Schwartz (1997), where the risk adjusted commodity spot price process is modeled by a GBM-like process given by

$$dP = (r - \delta)Pdt + \sigma PdZ. \quad (2.12)$$

Here r is the constant riskless interest rate; δ is the instantaneous convenience yield, following an Ornstein-Uhlenbeck mean-reverting process. The drift coefficient $r - \delta$ can switch between positive and negative values during a time interval since the value of δ is stochastic and may change signs during the interval. Thus the gas price P will either drift up or drift down at any time depending on the sign of $r - \delta$. According to (2.11), the GBMGBM variation can produce a behavior similar to the SDE (2.12).

3 Empirical Calibration

In this section, we calibrate the models proposed in Section 2 to the market gas futures prices and options on futures.

3.1 Calibration to Futures

We illustrate the calibration of the regime-switching model (2.8-2.9) to the market futures prices; the calibration procedure of the MR model (2.1-2.2) is similar but simpler, and hence omitted here.

The calibration results indicate that the regime-switching model can better fit the market forward curves compared with the MR model.

3.1.1 Futures Price Valuation

Let $F^k(P, t, T)$ denote the natural gas futures price in regime k , $k \in \{0, 1\}$, at time t with delivery at T , while the gas spot price resides at P . Assuming the risk adjusted natural gas spot price follows the regime-switching model (2.8-2.9), we can write $F^k(P, t, T)$ as the risk neutral expectation of the spot price at T

$$F^k(p, t, T) = E^{\mathbb{Q}}[P(T) \mid P(t) = p, m(t) = k], \quad (3.1)$$

where $m(t)$ is the two-state Markov chain given in (2.7), representing the regime in which the risk adjusted gas spot price resides at t . From equation (3.1), F^k satisfies two PDEs that are coupled with each other given by

$$F_t^k + [\alpha^k(K_0^k - P) + S^k(t)P]F_P^k + \frac{1}{2}(\sigma^k)^2P^2F_{PP}^k + \lambda^{k \rightarrow (1-k)}(F^{1-k} - F^k) = 0, \quad k \in \{0, 1\} \quad (3.2)$$

with the boundary conditions

$$F^k(P, T, T) = P, \quad k \in \{0, 1\}. \quad (3.3)$$

The solution to PDEs (3.2) has the form

$$F^k(P, t, T) = a^k(t, T) + b^k(t, T)P, \quad (3.4)$$

where functions a, b are independent of P . Substituting equation (3.4) into equations (3.2-3.3) gives an ODE system

$$\begin{aligned} a_t^k + \lambda^{k \rightarrow (1-k)}(a^{1-k} - a^k) + \alpha^k K_0^k b^k &= 0 \\ b_t^k - [\alpha^k - S^k(t) + \lambda^{k \rightarrow (1-k)}]b^k + \lambda^{k \rightarrow (1-k)}b^{1-k} &= 0, \quad k \in \{0, 1\} \end{aligned} \quad (3.5)$$

subject to the boundary conditions

$$a^k(T, T) = 0; \quad b^k(T, T) = 1, \quad k \in \{0, 1\}. \quad (3.6)$$

Remark 3.1. For the regime-switching model, equation (3.4) and the ODE system (3.5-3.6) imply that the futures prices $F^k(t, T)$, $k \in \{0, 1\}$ at time t when the gas spot price is known are independent of the spot price volatilities σ^0, σ^1 . Similar observations indicate that the futures price is independent of spot price volatility for the MR model. Consequently, the volatility needs to be calibrated using financial instruments other than futures contracts; in this paper, we choose options on futures (see Section 3.2 for the detailed calibration procedure).

3.1.2 Data

The data used to test the models consist of monthly observed delivery prices of NYMEX Henry Hub natural gas futures contracts. The data are publicly available on the website <http://www.econstats.com>.

Our data set contains 51 observations in 51 months (one observation each month) during the period from February 2003 to July 2007. Each observation contains delivery prices for the first 14 contracts that correspond to the deliveries in the next 14 consecutive months starting from the month of the observation day.

In order to carry out the calibration, we need to input the gas spot prices. Although there exists a gas spot market in Henry Hub that trades the next day delivery contract, we cannot use the delivery price of the contract as the spot price because the delivery periods for the contracts in the spot market and futures market are different: the delivery lasts for only 24 hours for the former and normally over a whole month for the latter. However, we can regard the delivery price of the next month futures contract, each month on the last trading day of the contract, as the proxy for the gas spot price, since it corresponds to the delivery starting three days later and delivering over the next month¹. The same approach is used in Jaillet et al. (2004). Thus, our monthly observation is made on the last trading day of the next month delivery contract², where the delivery price for that contract is used as the market spot price during calibration and the delivery prices for the rest of 13 contracts from the observation are used as the market futures prices during calibration (this amounts to a total of 663 futures prices).

3.1.3 Calibration procedure

Let $\theta = \{\alpha^k, K_0^k, \beta_A^k, \beta_{SA}^k, A_0, SA_0, \lambda^{k \rightarrow (1-k)} \mid k \in \{0, 1\}\}$ denote the set of parameters we need to obtain through calibrating to the futures price data.

We use a least squares approach for the calibration. For each observation day t , we need to determine the regime in which the risk adjusted gas spot price lies. Following the approach in (Cortazar and Schwartz, 2003; Xu, 2004), we treat the regime number as a latent variable and reveal its value through calibration. Specifically, we perform the calibration by solving the following optimization problem:

$$\min_{\theta} \sum_t \sum_T \left(\hat{F}^{\hat{k}(t;\theta)}(P(t), t, T; \theta) - F(t, T) \right)^2, \text{ where} \quad (3.7)$$

$$\hat{k}(t; \theta) = \arg \min_{k \in \{0, 1\}} \sum_T \left(\hat{F}^k(P(t), t, T; \theta) - F(t, T) \right)^2,$$

where $F(t, T)$ is the market futures price on the observation day t with maturity T ; $\hat{F}^{\hat{k}(t;\theta)}(P(t), t, T; \theta)$ is the corresponding model implied futures price in regime $\hat{k}(t; \theta)$ calculated numerically from equations (3.4-3.6) using the market spot price $P(t)$ and the parameter set θ . In equation (3.7), the range of t consists of all the observation days in our sample data and that of T covers the thirteen consecutive delivery months starting two months after the month of t .

Here an optimal regime number $\hat{k}(t; \theta)$ is determined for each observation date t and for a given parameter set θ . Let $\hat{\theta}$ be the parameters returned by the calibration procedure. Then the regime $\hat{k}(t; \hat{\theta})$ is where the realized market gas price actually resides at time t , revealed from the calibration. Note that the time spent in a given regime, as determined by calibration, will depend on the \mathbb{P} measure transition densities, not the \mathbb{Q} measure densities.

Remark 3.2. *In (Schwartz, 1997; Cortazar and Schwartz, 2003), the commodity spot price serves as an unknown parameter, instead of an observed market value, and is estimated from the calibration process. This may help improve the calibration results.*

¹In NYMEX, the trading of the next month delivery contract each month terminates three business days prior to the first calendar day of the next month.

²Occasionally, the futures prices on that day are not available on the source website. If that is the case, we use the available price data on the day closest (usually within five days) to the last trading day in the month of the day.

In Cortazar and Schwartz (2003), the estimated spot price is assured to be consistent with the simultaneously calibrated model volatilities and correlations. Note that in our case the spot price is observed from the market and the latent variable $\hat{k}(t; \theta)$ does not have such a consistency requirement. This is because the transition densities under the \mathbb{P} and \mathbb{Q} measures are not the same.

3.1.4 Calibration results

As suggested in Section 2.2.1, we are interested in three variations of the regime-switching model. We will determine the best model through calibration, i.e., the model that optimally fits the market data. For this purpose, we set the initial parameter values so that the calibration procedure starts from each of the three variations in Section 2.2.1.

Our calibration results are sensitive to the starting values used in the optimization procedure. For example, if the initial estimates for the parameters has either the MRMR or MRGBM form, the calibrated parameters retain the same form. As we shall see below, good fits to the data can be obtained with either MRMR or MRGBM. However, if we use initial parameters consistent with GBMGBM, then the optimization procedure converges to the MRGBM parameters.

This indicates that the MRMR or MRGBM models are consistent with the market data, while the GBMGBM model does not appear to be consistent with market data.

The behavior of the volatility of the futures price as $T - t$ becomes large depends in general on the calibrated parameters. However, for both regimes in the MRMR model, the volatility of the futures prices declines as the maturity increases. For the MRGBM model, the volatility of the futures prices declines in the MR regime, but the behavior in the GBM regime is a complicated function of the calibrated parameters and the seasonality terms. It is worthwhile mentioning that the seasonality effects are very large for gas futures prices, and a much larger set of futures prices expiring at long dates would be needed to determine the long term volatility of the futures prices.

Table 3.1 presents the calibrated risk adjusted parameter values for the MR model (2.1-2.2) and for the MRMR and MRGBM variations. In our calibration procedure, we set a lower bound of $\beta_{SA}^k = 0$. As shown in Table 3.1, the semi-annual seasonality parameters β_{SA}^k for three models are zero, suggesting that a single trigonometric term can satisfyingly approximate the seasonality trend in the futures price data. Meanwhile, the table reveals a strong annual seasonality behavior under the risk neutral world: condition (2.6) is satisfied for the MR model and also for the processes in regime 0 of the MRMR and MRGBM variations. Consequently, Remark 2.1 implies that the corresponding gas price dynamics incorporating the seasonality effect are mean-reverting within certain periods of time and switch to (essentially) GBM with positive drift at other times. Regime 1 of the MRMR variation, nevertheless, always shows a mean-reverting effect and that of the MRGBM variation always follows a GBM with positive drift.

From Table 3.1, for the MRMR variation, the equilibrium level in regime 1 is considerably higher than in regime 0. As a result, regime 0 represents the relatively low price regime and regime 1 represents the relatively high price regime. Similarly, for the MRGBM variation, regime 0 can be regarded as the low price regime and regime 1 represents the regime where the gas price drifts up quickly (according to the value of α^1).

Comparing the calibrated parameter values in regime 0 of the MRMR variation with those in the MR model, we observe that these values are similar except for the equilibrium price: the former has $K_0^0 \approx 4.5$ while the latter has $K_0 \approx 8.7 > K_0^0$. The situation is reversed for regime 1 with $K_0 < K_0^1 \approx 11.7$. The above observation also holds for the MRGBM variation: $K_0^0 > K_0$ in regime 0 and the effective equilibrium price is greater than K_0 in regime 1 (we can imagine that

the GBM regime is equivalent to a mean-reverting regime with equilibrium level at $+\infty$).

Note that the risk neutral parameters in Table 3.1 are not necessarily consistent with their counterparts under the real world \mathbb{P} measure. To further illustrate this point, in Table 3.2 we give, for the MRGBM variation of the regime-switching model, the regimes $\hat{k}(t; \theta)$ where the realized gas spot price resides at various times in our sample, calibrated using the procedure in Section 3.1.3. We can observe that the duration of time spent in each regime, implied from Table 3.2, is inconsistent with the risk neutral regime shift intensities $\lambda^{0 \rightarrow 1}$, $\lambda^{1 \rightarrow 0}$ under the \mathbb{Q} measure in Table 3.1: the realized gas price stays in regime 1 for over 45% of the time, while the risk neutral gas price resides at regime 1 for only about 11% of the time. Regime 1 is a regime in which gas price drifts up quickly. The risk neutral price stays in this regime (on average) a much shorter time than the realized price. This observation is consistent with the common paradigm \mathbb{Q} is more pessimistic than \mathbb{P} , i.e. investors in gas are risk averse and price gas contracts with a pessimistic view of future gas prices.

Parameter	Description	MR	MRMR	MRGBM
		Estimate	Estimate	Estimate
α (α^0)	Mean-reversion rate (for regime 0)	0.406	0.430	0.435
K_0 (K_0^0)	Equilibrium price (for regime 0)	8.678	4.466	4.748
β_A (β_A^0)	Annual seasonality parameter (for regime 0)	0.527	0.600	0.550
β_{SA} (β_{SA}^0)	Semiannual seasonality parameter (for regime 0)	0	0	0
A_0	Annual seasonality time adjustment parameter	0.483	0.441	0.457
α^1	Mean-reversion rate for regime 1		1.033	-0.650
K_0^1	Equilibrium price for regime 1		11.709	0
β_A^1	Annual seasonality parameter for regime 1		0.571	0.555
β_{SA}^1	Semiannual seasonality parameter for regime 1		0	0
$\lambda^{0 \rightarrow 1}$	Intensity of the jump from regime 0 to regime 1		0.304	0.283
$\lambda^{1 \rightarrow 0}$	Intensity of the jump from regime 1 to regime 0		0.975	2.290

TABLE 3.1: *Estimated parameter values for the three models using 663 monthly observed futures price data from February 2003 to July 2007. The column MR represents the MR model. The columns MRMR and MRGBM represent the MRMR and MRGBM variation of the regime-switching model, respectively.*

Table 3.3 provides the dollar and the percentage mean absolute errors between the model and market prices for futures contracts with different delivery months across all observation days. The table illustrates that the MR model performs the worst in terms of both the dollar and the

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
2003	N/A	0	0	0	0	0	0	0	0	0	0	0
2004	0	0	0	0	0	0	0	1	1	0	0	1
2005	1	1	0	0	N/A	0	0	0	0	0	0	0
2006	1	1	1	1	1	N/A	N/A	1	1	1	1	1
2007	1	1	1	1	1	1	1	N/A	N/A	N/A	N/A	N/A

TABLE 3.2: Regimes where the realized market gas spot price resides at various times, where the spot price follows the MRGBM variation of the regime-switching model. The Table shows that 29 months correspond to regime 0 and 22 months correspond to regime 1. The N/A in the table corresponds to missing data.

percentage errors. On the other hand, the MRMR and MRGBM variations result in similar errors (with the difference of the overall errors less than 7%), while the MRMR variation outperforms the MRGBM variation for the contracts with relatively long maturities.

Note that these fits were obtained with eleven parameters fitting 663 data points. This fit may not be good enough for trading purposes. However, an *exact fit* can be obtained to any set of futures prices at a given time by adding a time dependent fitting function to the gas price process SDE. However, this fitting function would only have to account for the approximately 7% error obtained from the global calibration, hence would be relatively small. It seems that the overall forward curves for gas can be fit reasonably well with either the MRMR or MRGBM models.

Figure 3.1 illustrates the model implied futures prices and the market prices for the longest maturity contract, which corresponds to the largest calibration errors among all the contracts, in the sample across all observation days starting from February 2003. Figure 3.1a indicate that the MR model fits the market prices poorly in observation days close to February 2003. On the contrary, Figures 3.1b and 3.1c show that the MRMR and MRGBM variations of the regime-switching model can reasonably fit the data across all the observation days. Therefore, we conclude that among these models, the regime-switching models outperform the MR model in terms of fitting the market gas forward curves.

3.2 Calibration to Options on Futures


As stated in Remark 3.1, the spot price volatilities for the models in Section 2 need to be estimated using derivatives other than the futures contracts. In this section, we calibrate the volatility using market European call options on natural gas futures. We demonstrate the calibration procedure only for the regime-switching model (2.8-2.9). Similar but simpler procedures follow for the MR model.

3.3 Futures Option Valuation


Let $\bar{V}^k(F, t, T_v)$ denote the European call option value in regime k at time t with maturity at T_v , where F represents the price of the underlying futures contract at time t . Let K denote the strike price of the option. Let $F^k(t, T)$ represent the price of the underlying futures contract in regime k at time t with maturity at T , where T satisfies $T \geq T_v$.

Contract maturity	Mean absolute error					
	MR	MRMR	MRGBM	MR	MRMR	MRGBM
	In Dollars			In Percentage		
Month+2	0.278	0.240	0.248	3.98	3.57	3.64
Month+3	0.499	0.386	0.388	6.99	5.70	5.56
Month+4	0.645	0.487	0.470	8.96	7.01	6.57
Month+5	0.684	0.504	0.471	9.44	7.08	6.45
Month+6	0.741	0.486	0.502	10.13	6.60	6.60
Month+7	0.827	0.493	0.528	11.14	6.59	6.75
Month+8	0.872	0.492	0.548	11.86	6.66	7.11
Month+9	0.949	0.505	0.563	12.97	6.92	7.37
Month+10	1.011	0.557	0.574	13.93	7.61	7.53
Month+11	1.037	0.603	0.622	14.62	8.20	8.20
Month+12	1.075	0.580	0.640	15.60	8.21	8.53
Month+13	1.118	0.580	0.677	16.65	8.53	9.21
Month+14	1.152	0.585	0.698	17.68	8.97	9.74
Overall	0.838	0.500	0.533	11.84	7.05	7.17


TABLE 3.3: Mean absolute errors between the model and the market prices for the futures contracts with different delivery months, where the notation Month+k in the first column represents the kth month delivery after the month of the observation day. The errors are given both in dollars and in percentage. The column MR represents the MR model. The columns MRMR and MRGBM represent the MRMR and MRGBM variation of the regime-switching model, respectively.



(A) MR



(B) MRMR



(C) MRGBM

FIGURE 3.1: Comparison between the model and the market futures prices for the contract with the longest maturity (for the delivery after 14 months) in the sample across all observation days starting from February 2003. The x-axis represents the number of days between the observation day and the starting date. The model implied prices are computed using the calibrated parameters in Table 3.1. MR represents the MR model. MRMR and MRGBM represent the MRMR and MRGBM variation of the regime-switching model, respectively.

In NYMEX, the trading of a European option ends on the business day immediately preceding the expiration of the underlying futures contract. As a result, we can assume $T_v = T$, and we will thus use T as the maturity for both an option and its underlying futures contract.

We can write $\bar{V}^k(F, t, T)$ in the form of the risk neutral expectation as

$$\begin{aligned}\bar{V}^k(f, t, T) &= e^{-r(T-t)} E^{\mathbb{Q}} \left[(F^{m(T)}(T, T) - K) \mathbf{1}_{F^{m(T)}(T, T) \geq K} \mid F^k(t, T) = f, m(t) = k \right] \\ &= e^{-r(T-t)} E^{\mathbb{Q}} \left[(P(T) - K) \mathbf{1}_{P(T) \geq K} \mid a^k(t, T) + b^k(t, T)P(t) = f, m(t) = k \right],\end{aligned}\quad (3.8)$$

where $\mathbf{1}_{x \geq y}$ is an indicator function that returns 1 if $x \geq y$, or 0 if $x < y$; the second equality above uses the fact that $F^{m(T)}(T, T) = P(T)$ at maturity T as well as the relation (3.4) between futures price F and spot price P at time t assuming the risk neutral gas spot price follows the regime-switching model (2.8-2.9). Let $V^k(P, t, T)$ represent a synthetic European call option on spot price P at time t , in regime k with maturity T . Then we can write $V^k(P, t, T)$ in the form of the risk neutral expectation as

$$V^k(p, t, T) = e^{-r(T-t)} E^{\mathbb{Q}} \left[(P(T) - K) \mathbf{1}_{P(T) \geq K} \mid P(t) = p, m(t) = k \right]. \quad (3.9)$$

Comparing equations (3.8) and (3.9), we have

$$V^k \left(\frac{f - a^k(t, T)}{b^k(t, T)}, t, T \right) = \bar{V}^k(f, t, T), \quad (3.10)$$

where a^k and b^k are computed from the ODE system (3.5-3.6). As a result, we can compute $\bar{V}^k(F, t, T)$ using equation (3.10) as long as we are able to solve for $V^k(P, t, T)$.

Let r denote the constant riskless interest rate. Assuming that the spot price process follows SDE (2.8-2.9) and using the risk neutral expectation formulation (3.9), we find that the synthetic option value V^k satisfies the coupled PDEs

$$\begin{aligned}V_t^k + [\alpha^k(K_0^k - P) + S^k(t)P]V_P^k + \frac{1}{2}(\sigma^k)^2 P^2 V_{PP}^k - rV^k + \\ \lambda^{k \rightarrow (1-k)}(V^{1-k} - V^k) = 0, \quad k \in \{0, 1\}\end{aligned}\quad (3.11)$$

subject to the boundary conditions

$$V^k(P, T, T) = \max[P - K, 0], \quad k \in \{0, 1\}, \quad (3.12)$$

We will solve equations (3.11-3.12) numerically in the computational domain $P \in [0, P_{\max}]$ with $P_{\max} \gg K$. For this purpose, we need to impose boundary conditions at the computational boundary $P = 0$ and $P = P_{\max}$. At the $P = 0$ boundary, taking the limit as $P \rightarrow 0$, equations (3.11) become

$$V_t^k + \alpha^k K_0^k V_P^k - rV^k + \lambda^{k \rightarrow (1-k)}(V^{1-k} - V^k) = 0, \quad k \in \{0, 1\}. \quad (3.13)$$

Since $\alpha^k K_0^k \geq 0$ for all variations of the regime-switching model we consider (see Section 2.2.1), the characteristics are outgoing in the P direction at $P = 0$. As a result, we can solve equations (3.13) without requiring additional boundary conditions, as we do not need information from outside the

computational domain. At the $P = P_{\max}$ boundary, we make the assumption that $V^k(P_{\max}, t, T) \rightarrow$ payoff. In other words, we impose the Dirichlet boundary condition

$$V^k(P_{\max}, t, T) = P_{\max} - K, \quad k \in \{0, 1\}. \quad (3.14)$$

The error introduced by this approximation will be small if P_{\max} is sufficiently large.

During calibration, we will solve equations (3.11-3.14) using a standard fully implicit finite different scheme that is stable, monotone and consistent and hence the scheme converges to the unique solution of equations (3.11-3.14). We skip the details here.

3.4 Calibration procedure and results

We choose as input the values of twelve European call options from NYMEX on $t = 6/26/2007$ with different strike prices. These options have the same underlying futures contract, which expires in August 2007, denoted by T_1 . The futures price is 7.002 \$/mmBtu at time t . The strike prices with respect to the twelve options range from 6.5 to 7.5 \$/mmBtu, that is, we pick both slightly in the money and slightly out of the money options³. We assume that the annual riskless interest rate is $r = 5\%$.

The first step of the calibration is to determine the regime in which the underlying futures contract lies; the approach is given in Section 3.1.3. After determining the optimal regime, denoted by \hat{k} , for the futures contract at time t , we use a least squares approach to calibrate the volatility σ^k by solving

$$\min_{\sigma^0, \sigma^1} \sum_K \left(\bar{V}^{\hat{k}}(F(t, T_1), t, T_1; \theta, K, \sigma^0, \sigma^1) - C(t, T_1; K) \right)^2, \quad (3.15)$$

where $C(t, T_1; K)$ represents the value of a market call option on futures at time t with maturity at T_1 and strike price at K ; $\bar{V}^{\hat{k}}$ is the corresponding model implied option value in regime \hat{k} using the market futures price $F(t, T_1)$, the parameter set $\theta = \{\alpha^k, K_0^k, \beta_A^k, \beta_{SA}^k, A_0, SA_0, \lambda^{k \rightarrow (1-k)} \mid k \in \{0, 1\}\}$ and the volatility pair σ^0, σ^1 . The values of $\bar{V}^{\hat{k}}$ are computed by first solving PDEs (3.11-3.14) numerically to obtain values of V^k at time t and at discrete grid nodes in the P direction and then linearly interpolating these discrete values using equation (3.10). We choose the mesh size so that the error in $\bar{V}^{\hat{k}}$ is less than 10^{-2} .

Table 3.4 gives the calibration results and mean absolute errors for the MR model and the MRMR and MRGBM variations of the regime-switching model.

4 Pricing Natural Gas Storage Contracts under the Regime-Switching Model

In previous sections we have showed that the regime-switching model (2.8-2.9) outperforms the one-factor mean-reverting model (2.1-2.2) in capturing the dynamics of natural gas futures prices. In this section, we apply the calibrated model to price the value of a natural gas storage facility (or, equivalently, a gas storage contract for leasing the storage facility). After giving the pricing

³The data set we choose is relatively small. Nevertheless, as an illustration in our simple constant parameter setting, it is sufficient to estimate the volatilities for two regimes. One can add more market data into calibration, such as American options. One can also imagine assuming a volatility surface $\sigma^k = \sigma^k(P, t)$ in model (2.8-2.9) and calibrating the surface using futures options with different maturities and different strike prices.

Model	Volatility		Mean absolute error	
	σ^0	σ^1	In Cents	In Percentage
MR	0.428		1.75	4.27
MRMR	0.406	0.453	0.15	0.59
MRGBM	0.394	0.416	0.15	0.60

TABLE 3.4: *Calibrated volatilities and mean absolute errors for the futures options. The errors are given both in cents and in percentage. The column MR represents the MR model. The columns MRMR and MRGBM represent the MRMR and MRGBM variation of the regime-switching model, respectively.*

equations for gas storage contracts, we discuss the corresponding boundary conditions. We then introduce the numerical scheme for solving the pricing equations. Finally, we conduct numerical convergence tests and investigate the optimal operational strategies on storage facilities implied from these gas spot price models. Readers can refer to Chen and Forsyth (2007) and the references therein for a detailed description of the gas storage valuation problem.

4.1 The PDEs for natural gas storage facilities

The natural gas storage valuation problem is characterized as a stochastic control problem that results in Hamilton-Jacobi-Bellman (HJB) equations. In Chen and Forsyth (2007), a gas storage HJB equation is formulated assuming the underlying risk adjusted gas spot price follows a one-factor mean-reverting process. Prior to presenting the gas storage equations in terms of the regime-switching model (2.8-2.9), we first introduce the following notation.

Let T denote the maturity of the gas storage contract. Let $V^k(P, I, \tau) = V(P, I, \tau, k)$ denote the value of a natural gas storage facility in regime k when the gas price resides at P , the working gas inventory lies at I and the current time to maturity of the contract is $\tau = T - t$. The inventory I can be any value lying within the domain $[0, I_{\max}]$. Let c be a control variable that represents the rate of producing gas from or injecting gas into the gas storage, where $c > 0$ represents production, $c < 0$ represents injection, and $c = 0$ represents doing nothing. Let $c_{\max}(I) \geq 0$ and $c_{\min}(I) \leq 0$ denote the maximum gas production rate and maximum gas injection rate as a function of I , respectively. Following Thompson et al. (2003), we assume

$$c_{\max}(I) = k_1 \sqrt{I}, \quad (4.1)$$

$$c_{\min}(I) = -k_2 \sqrt{\frac{1}{I + k_3} - \frac{1}{k_4}}, \quad (4.2)$$

where k_1, k_2, k_3 and k_4 are positive constants, and k_2, k_3, k_4 satisfy the constraint $c_{\min}(I_{\max}) = 0$ which means that no gas injection is possible if the gas storage is full. Let $a(c)$ be the rate of gas loss incurred by the operation with rate c . We assume $a(c) = 0$ if $c \geq 0$, and $a(c) = k_5$ if $c < 0$, where k_5 is a positive constant. Based on the arguments in Chen and Forsyth (2007), we require that the control c satisfies $c \in C(I)$ where

$$C(I) = [c_{\min}(I), -k_5] \cup [0, c_{\max}(I)]. \quad (4.3)$$

Let r be the riskless interest rate.

Based on the standard hedging arguments in the financial valuation literature (see, e.g., Thompson et al. (2003)), the value of a natural gas storage facility $V^k(P, I, \tau)$, assuming that the risk adjusted gas spot price follows the regime-switching model (2.8-2.9) is given by the coupled HJB equations

$$V_\tau^k = \frac{1}{2}(\sigma^k)^2 P^2 V_{PP}^k + [\alpha^k(K_0^k - P) + S^k(t)P] V_P^k + \max_{c \in C(I)} \{(c - a(c))P - (c + a(c))V_I^k\} - (r + \lambda^{k \rightarrow (1-k)})V^k + \lambda^{k \rightarrow (1-k)}V^{1-k}, \quad k \in \{0, 1\}. \quad (4.4)$$

4.2 Boundary conditions for the pricing PDEs

In order to completely specify the gas storage problem, we need to provide boundary conditions. As for the terminal boundary conditions, we use the following penalty payoff function given by Carmona and Ludkovski (2005):

$$V^k(P, I, \tau = 0) = \text{const.} \cdot P \cdot \min(I(t = T) - I(t = 0), 0), \quad k \in \{0, 1\}. \quad (4.5)$$

This has the meaning that a penalty will be charged if the gas inventory in storage when the facility is returned is less than the gas inventory at the inception of the contract.

The domain for the PDEs (4.4) is $P \times I \in [0, \infty] \times [0, I_{\max}]$. For computational purposes, we need to solve PDEs in a finite computational domain $[0, P_{\max}] \times [0, I_{\max}]$. As $I \rightarrow 0$ or $I \rightarrow I_{\max}$, according to the arguments in Chen and Forsyth (2007), no boundary conditions are needed since the characteristics are outgoing in the I direction.

Taking the limit of equations (4.4) as $P \rightarrow 0$, we obtain

$$V_\tau^k = \alpha^k K_0^k V_P^k + \max_{c \in C(I)} \{-(c + a(c))V_I^k\} - (r + \lambda^{k \rightarrow (1-k)})V^k + \lambda^{k \rightarrow (1-k)}V^{1-k}, \quad k \in \{0, 1\}. \quad (4.6)$$

Since $\alpha^k K_0^k \geq 0$ for all variations of the regime-switching model we consider (see Section 2.2.1), the characteristics are outgoing in the P direction and we can solve equations (4.6) without requiring additional boundary conditions.

As $P \rightarrow \infty$, we make the common assumption that $V_{PP}^k \rightarrow 0$ (see Windcliff et al. (2004)). We need to deal with one major issue in that the resulting boundary equations require information from outside the computational domain. To see the problem, assuming $V_{PP}^k \rightarrow 0$ as $P \rightarrow \infty$, then the pricing equations (4.4) become

$$V_\tau^k = [\alpha^k K_0^k + (S^k(t) - \alpha^k)P] V_P^k + \max_{c \in C(I)} \{(c - a(c))P - (c + a(c))V_I^k\} - (r + \lambda^{k \rightarrow (1-k)})V^k + \lambda^{k \rightarrow (1-k)}V^{1-k}, \quad k \in \{0, 1\}. \quad (4.7)$$

Using the calibrated parameter values from Table 3.1, we find that $S^0(t) - \alpha^0 > 0$ are positive for certain ranges of t . In this case, the characteristics of equations (4.7) are incoming in the P direction at $P \rightarrow \infty$ and consequently, a monotone discretization of the equation will require information from outside the computational domain.

This issue can be resolved using the following approximation. The assumption $V_{PP}^k \rightarrow 0$ as $P \rightarrow \infty$ implies that

$$V^k(P, I, \tau) \approx f^k(I, \tau)P + g^k(I, \tau), \quad (4.8)$$

where functions f^k and g^k are independent of P . If we assume that $f^k(I, \tau)P \gg g^k(I, \tau)$ as $P \rightarrow \infty$, we can further write

$$V^k \approx f^k(I, \tau)P. \quad (4.9)$$

Note that the approximation (4.8) is consistent with the payoff (4.5). Now the drift term in the boundary equation (4.7) can be written as

$$\begin{aligned} [\alpha^k K_0^k + (S^k(t) - \alpha^k)P]V_P^k &\approx (S^k(t) - \alpha^k)PV_P^k \\ &\approx (S^k(t) - \alpha^k)V^k, \end{aligned} \quad (4.10)$$

where the first approximation follows since $(S^k(t) - \alpha^k)P \gg \alpha^k K_0^k$ as $P \rightarrow \infty$ and the second approximation is due to equation (4.9). Substituting equation (4.10) into equation (4.7) results in

$$\begin{aligned} V_\tau^k &= \max_{c \in C(I)} \{(c - a(c))P - (c + a(c))V_I^k\} - (r + \alpha^k - S^k(t) + \lambda^{k \rightarrow (1-k)})V^k \\ &\quad + \lambda^{k \rightarrow (1-k)}V^{1-k}, \quad k \in \{0, 1\}. \end{aligned} \quad (4.11)$$

Since the drift term in equations (4.11) is zero, we are able to provide a monotone discretization for the equation without requiring information from outside the computational domain. (Refer to Section 4.3.2 for more details.)

4.3 Numerical solution

This section briefly introduces the fully implicit, semi-Lagrangian timestepping scheme for solving the gas storage PDEs (4.4) and the associated boundary equations (4.6) and (4.11). We also analyze the l_∞ stability and monotonicity properties of the scheme and show the convergence of the scheme to the viscosity solution of the equations (4.4), (4.6) and (4.11).

4.3.1 Discretization

Fully implicit and a Crank-Nicolson timestepping schemes using the semi-Lagrangian approach are developed in Chen and Forsyth (2007) to solve the gas storage equation for a one-factor mean-reverting model. We can easily extend the semi-Lagrangian discretizations to solve the gas storage equations (4.4) and boundary equations (4.6) and (4.11) in the regime-switching framework.

Prior to presenting the matrix form of semi-Lagrangian discretizations, we introduce the following notation. We use unequally spaced grids in the P and I directions for the PDE discretization, represented by $[P_0, P_1, \dots, P_{i_{\max}}]$ and $[I_0, I_1, \dots, I_{j_{\max}}]$, respectively. We use the discrete timesteps $0 = \tau^0 < \dots < \tau^N = T$ to discretize the PDEs. Let $U_{i,j,k}^n$ denote an approximation of the exact solution $V^k(P_i, I_j, \tau^n)$, where $k \in \{0, 1\}$. Let U^n denote a column vector that includes all values of $U_{i,j,k}^n$ with the index order arranged as $U^n = [U_{0,0,0}^n, \dots, U_{i_{\max},0,0}^n, \dots, U_{0,j_{\max},0}^n, \dots, U_{i_{\max},j_{\max},1}^n]'$. For future reference, assuming M is a square matrix, then we denote $[MU^n]_{ijk} = (MU^n)_{i,j,k}$, and denote $[MU^n]_{jk}$ as the vector $[(MU^n)_{0,j,k}, \dots, (MU^n)_{i_{\max},j,k}]'$, where the index of $(MU^n)_{i,j,k}$ in MU^n is the same as that of $U_{i,j,k}^n$ in U^n .

Let \mathcal{L} be a differential operator represented by

$$\mathcal{L}V^k = \frac{1}{2}(\sigma^k)^2 P^2 V_{PP}^k + [\alpha^k(K_0^k - P) + S^k(t)P]PV_P^k - (r + \lambda^{k \rightarrow (1-k)})V^k + \lambda^{k \rightarrow (1-k)}V^{1-k} \quad (4.12)$$

for $P \in [0, \infty)$. While $P \rightarrow \infty$, according to the boundary equations (4.11), we define operator \mathcal{L} to be

$$\mathcal{L}V^k = -(r + \alpha^k - S^k(t) + \lambda^{k \rightarrow (1-k)})V^k + \lambda^{k \rightarrow (1-k)}V^{1-k}. \quad (4.13)$$

The operator \mathcal{L} can be discretized using standard finite difference methods. Let $(\mathcal{L}_h V)_{i,j,k}^n$ denote the discrete value of the operator \mathcal{L} at node (P_i, I_j, τ^n, k) . Using central, forward, or backward differencing in the P, I directions, we can write

$$\begin{aligned} & (\mathcal{L}_h V)_{i,j,k}^n \\ &= \begin{cases} \gamma_{i,k}^n U_{i-1,j,k}^n + \beta_{i,k}^n U_{i+1,j,k}^n - (\gamma_{i,k}^n + \beta_{i,k}^n + r + \lambda^{k \rightarrow (1-k)})U_{i,j,k}^n + \lambda^{k \rightarrow (1-k)}U_{i,j,1-k}^n & i < i_{\max}, \\ -(r + \alpha^k - S^k(T - \tau^n) + \lambda^{k \rightarrow (1-k)})U_{i,j,k}^n + \lambda^{k \rightarrow (1-k)}U_{i,j,1-k}^n & i = i_{\max}, \end{cases} \end{aligned} \quad (4.14)$$

where constants $\gamma_{i,k}^n, \beta_{i,k}^n$ can be determined using a similar algorithm as that given in Chen and Forsyth (2007). The algorithm guarantees that $\gamma_{i,k}^n, \beta_{i,k}^n$ satisfy a positive coefficient condition:

$$\gamma_{i,k}^n \geq 0 \ ; \ \beta_{i,k}^n \geq 0 \quad i = 0, \dots, i_{\max} - 1 \ ; \ n = 1, \dots, N. \quad (4.15)$$

Following Chen and Forsyth (2007), we define a semi-Lagrangian trajectory $\mathcal{I}(\tau)$ that follows the ODE

$$\frac{d\mathcal{I}}{d\tau} = c + a(c). \quad (4.16)$$

According to (4.16), we can write the terms $V_\tau^k + (c + a(c))V_I^k$ in equations (4.4), (4.6) and (4.11) in the form of a Lagrangian directional derivative

$$\frac{DV^k}{D\tau} = \frac{\partial V^k}{\partial \tau} + \frac{\partial V^k}{\partial I} \frac{d\mathcal{I}}{d\tau}. \quad (4.17)$$

Then equations (4.4), (4.6) and (4.11) can be written together as

$$\min_{c \in \mathcal{C}(I)} \left\{ \frac{DV^k}{D\tau} - (c - a(c))P - \mathcal{L}V^k \right\} = 0, \quad k \in \{0, 1\}. \quad (4.18)$$

We can construct a semi-Lagrangian discretization of equations (4.18) by first integrating the equations along the semi-Lagrangian trajectory $\mathcal{I}(\tau)$ in (4.16) from $\tau = \tau^n$ to $\tau = \tau^{n+1}$, and then evaluating the resulting integrals using different numerical integration methods: for example, using the rectangular rule provides a fully implicit timestepping scheme, while using the trapezoidal rule gives a Crank-Nicolson timestepping scheme. The semi-Lagrangian approach assumes that the arrival point of the trajectory at $\tau = \tau^{n+1}$ resides at a grid node (P_i, I_j) , while the departure point of the trajectory at $\tau = \tau^n$ is determined by values of optimal control c along this path. We denote by $U_{i,j(i,n+1),k}^n$ the discrete solution value at the departure point of the semi-Lagrangian trajectory $\mathcal{I}(\tau^n)$. Note that the departure point $\mathcal{I}(\tau^n)$ will not necessarily coincide with a grid point in the I direction. Consequently, we need to interpolate the value $U_{i,j(i,n+1),k}^n$ using discrete grid values $U_{i,j,k}^n$, $i = 0, \dots, i_{\max}$, $j = 0, \dots, j_{\max}$. The arguments in Chen and Forsyth (2007) shows that using linear interpolation to compute $U_{i,j(i,n+1),k}^n$ is able to achieve the first-order global discretization error.

As demonstrated in Chen and Forsyth (2007), the first order fully implicit timestepping scheme is a better choice than the Crank-Nicolson timestepping scheme since the latter does not converge at a higher than first order rate and cannot guarantee convergence to the viscosity solution of the pricing HJB equation. As a result, in this paper we only consider the fully implicit timestepping scheme. In the above paragraphs we present the main idea on how to construct semi-Lagrangian discretizations, we then directly write out the matrix form of the fully implicit timestepping scheme in the following without showing the intermediate steps; readers can refer to Chen and Forsyth (2007) for more details. Let L^n denote a matrix such that

$$[L^n U^n]_{ijk} = (\mathcal{L}_h V)_{i,j,k}^n, \quad (4.19)$$

where the discrete value $(\mathcal{L}_h V)_{i,j,k}^n$ is given in equation (4.14).

We denote by Φ^{n+1} a Lagrange linear interpolation matrix such that

$$[\Phi^{n+1} U^n]_{ijk} = U_{i,j(i,n+1),k}^n + \text{interpolation error}. \quad (4.20)$$

Let P denote a column vector satisfying $[P]_i = P_i$. Let ζ_{jk}^n be a diagonal matrix containing control values that need to be determined. Let $a(\zeta_{jk}^n)$ denote a diagonal matrix with diagonal entries $[a(\zeta_{jk}^n)]_{ii} = a([\zeta_{jk}^n]_{ii})$. Let I be an identity matrix. Let $\Delta\tau^n = \tau^{n+1} - \tau^n$. Given the above notation, the fully implicit timestepping scheme for the gas storage pricing equations (4.4), (4.6) and (4.11) can be written as

$$\begin{aligned} [[I - \Delta\tau^n L^{n+1}] U^{n+1}]_{jk} &= [\Phi^{n+1} U^n]_{jk} + \Delta\tau^n (\zeta_{jk}^{n+1} - a(\zeta_{jk}^{n+1})) P, \\ \text{where } [\zeta_{jk}^{n+1}]_{ii} &= \arg \max_{[\zeta_{jk}^{n+1}]_{ii} \in C_{jk}^{n+1}} \left\{ [[\Phi^{n+1} U^n]_{jk} + \Delta\tau^n (\zeta_{jk}^{n+1} - a(\zeta_{jk}^{n+1})) P]_i \right\} \end{aligned} \quad (4.21)$$

for $j = 0, \dots, j_{\max}$ and $k = 0, 1$. Here $[\zeta_{jk}^{n+1}]_{ii}$ represents the optimal control in the admissible set C_{jk}^{n+1} (see Chen and Forsyth (2007) for the definition of the admissible set).

The discrete optimization problems in scheme (4.21) can be solved using both the *bang-bang* method and *no bang-bang* method as proposed in Chen and Forsyth (2007). As shown in Thompson et al. (2003), the exact solution of equations (4.4), (4.6), (4.11) has the property that the controls are of the bang-bang type, i.e. the optimal controls can take on only the values in a finite set which, in our case, consists of the maximum production rate $c = c_{\max}(I)$, maximum injection rate $c = c_{\min}(I)$ and zero rate $c = 0$. However, the solutions of the discrete optimization problems may allow controls which are not possible controls for the exact solution. The bang-bang method, roughly speaking, considers only those bang-bang controls as appearing in the exact solution, and the no bang-bang method actually maximizes the discrete problems. We know that both methods will converge to the correct solution with bang-bang type controls, for sufficiently small mesh spacing and timestep size. We will conduct numerical experiments using both the bang-bang and no bang-bang methods.

4.3.2 Properties of the numerical scheme

As shown in Pooley et al. (2003), in the case of nonlinear option pricing problem, seemingly reasonable discretizations that do not satisfy the sufficient convergence conditions for viscosity solutions may either never converge or converge to a non-viscosity solution. Provided a strong comparison

result for the PDE applies, (Barles and Souganidis, 1991; Barles, 1997) demonstrate that a numerical scheme will converge to the viscosity solution of the equation if it is l_∞ stable, pointwise consistent, and monotone.

The consistency of the semi-Lagrangian discretization (4.21) for the gas storage equations in the regime-switching setting directly follows from the consistency results in Chen and Forsyth (2007). In this section, we discuss the l_∞ stability and monotonicity properties of the scheme (4.21).

Let us define matrix

$$Q^{n+1} = I - \Delta\tau^n L^{n+1}. \quad (4.22)$$

First, we prove the following Lemma

Lemma 4.1 (*M-matrix property of Q^{n+1}*). *Assuming that scheme (4.21) satisfies the positive coefficient condition (4.15) and the timestep condition*

$$\Delta\tau^n < \frac{1}{\|\alpha\| + \|S\|}, \quad n = 0, 1, \dots, N, \quad (4.23)$$

where

$$\|\alpha\| = \max(|\alpha^0|, |\alpha^1|) \quad ; \quad \|S\| = \max_{t \in [0, T], k \in \{0, 1\}} |S^k(t)|, \quad (4.24)$$

then the matrix Q^{n+1} is an *M* matrix.

Proof. From equations (4.14), using conditions (4.15) and (4.23), we can verify that $-L^{n+1}$ has nonpositive offdiagonal elements and the sum of elements in each row in matrix Q^{n+1} is nonnegative. Note that condition (4.23) is needed to show that the above diagonal dominant property holds for the last row of the matrix. Hence Q^{n+1} is an *M*-matrix. \square

Remark 4.2 (Explanation of the timestep condition (4.23)). *Condition (4.23) is a mild timestep-ping constraint since $S^k(t)$ is bounded above by a relatively small constant. For example, using the parameter values from Table 3.1, condition (4.23) is equivalent to $\Delta\tau^n < 0.49$ and $\Delta\tau^n < 0.52$ for the MRMR and MRGBM variation of the regime-switching model, respectively. This indicates that a timestep of 0.4 year is sufficient to satisfy the condition.*

Let us define $\Delta\tau_{\max} = \max_n(\tau^{n+1} - \tau^n)$, $\Delta\tau_{\min} = \min_n(\tau^{n+1} - \tau^n)$. The following Lemma verifies the l_∞ stability of scheme (4.21).

Lemma 4.3 (*l_∞ stability*). *Let $D_\tau = (\Delta\tau_{\max})/(\Delta\tau_{\min})$, where $D_\tau > 0$ is assumed to be a $O(1)$ constant. Assuming that discretization (4.21) satisfies the conditions required for Lemma 4.1, then as $\Delta\tau_{\min} \rightarrow 0$, scheme (4.21) satisfies*

$$\|U^{n+1}\|_\infty \leq D_1 \|U^0\|_\infty + D_2, \quad (4.25)$$

where D_1, D_2 are bounded constants given by

$$D_1 = \exp((\|\alpha\| + \|S\|)D_\tau T) \quad ; \quad D_2 = \frac{1}{\|\alpha\| + \|S\|} (D_1 - 1) \cdot P_{\max} \cdot \max\{|C_{\max}(I_{\max})|, |C_{\min}(0)|\}. \quad (4.26)$$

Here $\|U^{n+1}\|_\infty = \max_{i,j,k} |U_{i,j,k}^{n+1}|$. Therefore, according to the definition of l_∞ stability given in Chen and Forsyth (2007), the scheme (4.21) is unconditionally l_∞ stable.

Proof. The proof directly follows from using Lemma 4.1 and applying the maximum principle to the discrete equation (4.21). We omit the details here. Readers can refer to (D’Halluin et al., 2005, Theorem 5.5) and Forsyth and Labahn (2007) for complete stability proofs of the semi-Lagrangian fully implicit scheme for American Asian options and the finite difference schemes for controlled HJB equations, respectively. \square

Next we discuss the monotonicity of scheme (4.21). We can write the discrete equations (4.21) at each node (P_i, I_j, k) as

$$\begin{aligned} & G_{i,j,k}^{n+1}(U_{i,j,k}^{n+1}, \{U_{l,j,k}^{n+1}\}_{l \neq i}, U_{i,j,1-k}^{n+1}, \{U_{i,j,k}^n\}) \\ & \equiv \min_{\{\zeta_{jk}^n\}_{i \in C_{jk}^{n+1}}} \left\{ [Q^{n+1}U^{n+1}]_{ijk} - [\Phi^{n+1}U^n]_{ijk} - \Delta\tau^n \left[(\zeta_{jk}^{n+1} - a(\zeta_{jk}^{n+1}))P \right]_i \right\} \\ & = 0, \end{aligned} \quad (4.27)$$

where $\{U_{l,j,k}^{n+1}\}_{l \neq i}$ is the set of values $U_{l,j,k}^{n+1}$, $l \neq i$, $l = 0, \dots, i_{\max}$, and $\{U_{i,j,k}^n\}$ is the set of values $U_{i,j,k}^n$, $i = 1, \dots, i_{\max}$, $j = 1, \dots, j_{\max}$.

Definition 4.4 (Monotonicity). *The scheme $G_{i,j,k}^{n+1}(U_{i,j,k}^{n+1}, \{U_{l,j,k}^{n+1}\}_{l \neq i}, U_{i,j,1-k}^{n+1}, \{U_{i,j,k}^n\})$ given in equation (4.27) is monotone if*

$$\begin{aligned} & G_{i,j,k}^{n+1}(U_{i,j,k}^{n+1}, \{U_{l,j,k}^{n+1} + \epsilon_{l,j,k}^{n+1}\}_{l \neq i}, U_{i,j,1-k}^{n+1} + \epsilon_{i,j,1-k}^{n+1}, \{U_{i,j,k}^n + \epsilon_{i,j,k}^n\}) \\ & \leq G_{i,j,k}^{n+1}(U_{i,j,k}^{n+1}, \{U_{l,j,k}^{n+1}\}_{l \neq i}, U_{i,j,1-k}^{n+1}, \{U_{i,j,k}^n\}); \quad \forall i, j, k, \quad \forall \epsilon_{i,j,k}^n, \epsilon_{i,j,k}^{n+1} \geq 0. \end{aligned} \quad (4.28)$$

The definition of monotonicity is equivalent to that introduced in Barles (1997).

Following the proof in (Barles and Jakobsen, 2007; Forsyth and Labahn, 2007) and using Lemma 4.1, we can obtain the monotonicity of scheme (4.21).

Lemma 4.5 (Monotonicity). *If scheme (4.21) satisfies the conditions required for Lemma 4.1, then the discretization (4.27) is monotone according to Definition 4.4.*

Remark 4.6. *For different types of controlled HJB equations/inequalities, (Barles and Jakobsen, 2002; Jakobsen, 2003; Barles and Jakobsen, 2007) introduce various definitions of monotonicity in order to obtain convergence rate estimates for their numerical schemes. These definitions impose some other conditions in addition to condition (4.28). The discretization (4.27) may not satisfy these additional conditions since the coefficient of the V^k term in the boundary equations (4.11) may become negative at the node $i = i_{\max}$. However, as shown in Barles (1997), only condition (4.28) is needed in order to prove the convergence of a numerical scheme to the viscosity solution.*

From Lemmas 4.3 and 4.5, using the results in (Barles and Souganidis, 1991; Barles, 1997), we can obtain the following convergence result:

Theorem 4.7 (Convergence to the viscosity solution). *Assuming that discretization (4.21) satisfies all the conditions required for Lemmas 4.3 and 4.5, and assuming that a strong comparison result for equations (4.4), (4.6) and (4.11) holds (refer to (Chen and Forsyth, 2007; Barles and Burdeau, 1995; Barles and Rouy, 1998) for explanations of strong comparison results for controlled PDEs), then scheme (4.21) converges to the viscosity solution of gas storage equations (4.4), (4.6) and (4.11).*

Cont and Voltchkova (2005) demonstrate that a financially meaningful discretization should satisfy arbitrage inequalities, which means the inequality of contract payoffs is preserved in the inequalities of contract values. Using the M -matrix property of matrix Q^{n+1} (Lemma 4.1), we can show that scheme (4.21) satisfies the following arbitrage inequalities.

Theorem 4.8 (Discrete arbitrage inequalities). *Assuming that discretization (4.21) satisfies the conditions required for Lemma 4.1, then suppose U and W are two solutions of (4.21), with $U^n > W^n$, then $U^{n+1} > W^{n+1}$.*

Proof. The proof directly follows from the approach in (D’Halluin et al., 2005, Theorem 6.2) by using the M -matrix property in Lemma 4.1. \square

4.4 Numerical Experiments

Having introduced the fully implicit semi-Lagrangian discretization scheme in the previous section, this section conducts numerical experiments based on the proposed scheme. We use “dollars per million British thermal unit” (\$/mmBtu) and “million cubic feet” (MMcf) as the default units for gas spot price P and gas inventory I , respectively. The numerical experiments use the following nonlinear payoff function from Carmona and Ludkovski (2005):

$$V^k(P, I, \tau = 0) = -2P \max(1000 - I, 0), \quad k \in \{0, 1\}. \quad (4.29)$$

Equation (4.29) indicates that severe penalties are charged if the gas inventory is less than 1000 MMcf and no compensation is received when the inventory is above 1000 MMcf. Naturally, such a payoff structure will force the operator of a gas storage facility to maintain the gas inventory as close to 1000 MMcf as possible at maturity to avoid revenue loss.

We first carry out a convergence analysis, assuming that the risk adjusted natural gas spot price follows the MRMR variation of the regime-switching model (2.8-2.9) and taking model parameter values from Tables 3.1 and 3.4. Table 4.1 lists other input parameters for pricing the value of a gas storage contract.

The convergence results for two regimes at the node $(P, I) = (6, 1000)$ obtained from refining the mesh spacing and timestep size are shown in Table 4.2, where we use both the bang-bang and no bang-bang methods for solving the discrete optimization problem in scheme (4.21). The results indicate that the both methods achieve the first-order convergence. A similar observation is found for the MRGBM variation.

Parameter	Value	Parameter	Value
r	0.05	k_2	730000
T	3 years	k_3	500
I_{\max}	2000 MMcf	k_4	2500
k_1	2040.41	k_5	$1.7 \cdot 365$

TABLE 4.1: *Input parameters used to price the value of a gas storage contract, where I_{\max} is the maximum storage inventory; k_1, k_2, k_3, k_4, k_5 are parameters in equations (4.1-4.3). The values of $I_{\max}, k_1, k_2, k_3, k_4, k_5$ are taken from Thompson et al. (2003).*

Our next step is to investigate the optimal operational strategies implied from the gas spot price models. Figure 4.1a plots the optimal control surface of the MR model (2.1-2.2) as a function

P grid nodes	I grid nodes	No. of timesteps	Bang-bang method		No bang-bang method	
			Value	Ratio	Value	Ratio
Regime 0						
87	61	500	3902629	n.a.	4011939	n.a.
173	121	1000	3920017	n.a.	3978633	n.a.
345	241	2000	3932943	1.35	3962706	2.09
689	481	4000	3939135	2.09	3954988	2.06
1377	961	8000	3942166	2.04	3951148	2.01
Regime 1						
87	61	500	4836234	n.a.	4962610	n.a.
173	121	1000	4864240	n.a.	4929533	n.a.
345	241	2000	4881661	1.61	4913620	2.08
689	481	4000	4889176	2.32	4905978	2.08
1377	961	8000	4892715	2.12	4902130	1.99

TABLE 4.2: Value of a natural gas storage facility in two regimes at $P = 6$ \$/mmBtu and $I = 1000$ MMcf. Risk adjusted gas spot price follows MRMR variation of the regime-switching process (2.8-2.9) with model parameter values given in Tables 3.1 and 3.4. Convergence ratios are presented for the bang-bang and the no bang-bang methods in two regimes. The convergence ratio is defined as the ratio of successive changes in the solution, as the timestep and mesh size are reduced by a factor of two. Constant timesteps are used. Payoff function is given in (4.29). Other input parameters are given in Table 4.1. We assume that today is the January 1st of the year.

of forward time t and gas price P when $I = 1000$ MMcf. We can verify from the figure that the optimal controls are of the bang-bang type: the controls are either producing at the maximum rate $c = c_{\max} > 0$, or injecting at the maximum rate $c = c_{\min} < 0$, or doing nothing with $c = 0$.

Another observation from Figure 4.1a is that for a fixed time t , the control is price dependent; at the same time, the controls evolving over time follow a repeated seasonal pattern. Specifically, in winter months, it is optimal to produce when the price is sufficiently high, to inject when the price is relatively low, and to do nothing when the price lies in between. The higher the price, the longer the production period will be; the lower the price, the longer the injection period will be. Furthermore, the equilibrium level $P = 8.678$ approximately resides in the center of the three bang-bang control regions (i.e., the regions of injection, doing nothing and production). This is due to the mean-reverting behavior of the MR model during the winter period (see the discussions in Remark 2.1 and Section 3.1.4). In summer months, however, the optimal control is to inject or to do nothing, but never to withdraw. The gas price during this period essentially follows a GBM with a positive drift. As such, it is never optimal to withdraw since the price tends to drift up during this period due to the strong seasonality effect. From the discussions above, we conclude that the optimal controls are consistent with the gas price dynamics implied from the calibrated MR model.

We can also observe from Figure 4.1a that the controls converge to zero when $t \rightarrow T$ in order to avoid the revenue loss due to the payoff structure (4.29).

To see the seasonality effect on the control strategies more clearly, in Figure 4.1b we present the optimal control curve obtained by taking a slice of the control surface in Figure 4.1a at $P = 6$ \$/mmBtu along the t direction. The figure shows that it is optimal to produce between February

and March (i.e., in the cold season when the gas prices are relatively high), inject between July and October (i.e., in the warm season when the gas prices are relatively low) and do nothing in other seasons.

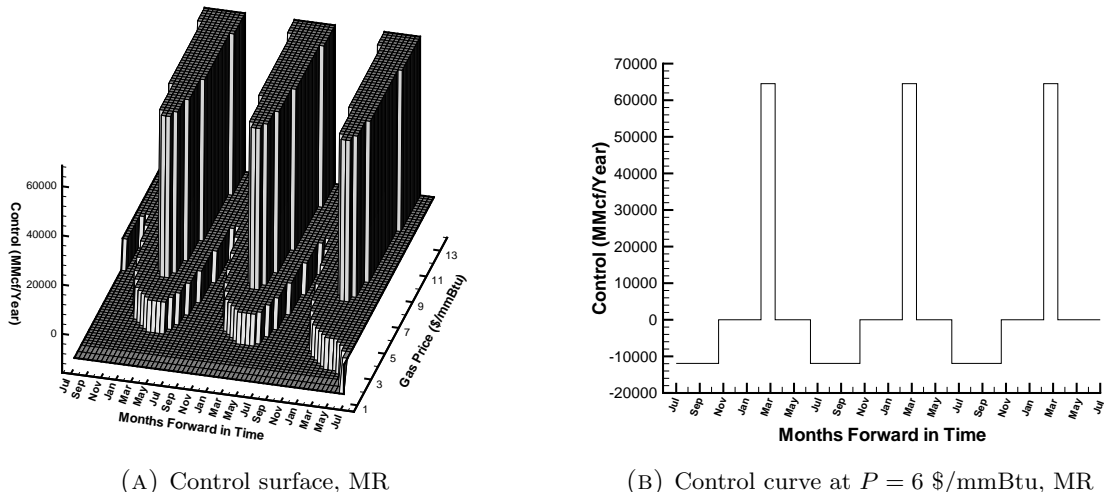




FIGURE 4.1: Optimal control surface for the MR model as a function of forward time t and gas spot price P as well as the corresponding control curve as a function of t when $P = 6$ \$/mmBtu, where the gas inventory resides at $I = 1000$ MMcf. Model parameter values are given in Tables 3.1 and 3.4. Fully implicit timestepping with the no bang-bang method and with constant timesteps is used. Other input parameters are given in Table 4.1. We assume that the starting time $t = 0$ corresponds to the July 1st of the year.

As a comparison, Figure 4.2 plots, for the MRMR variation of the regime-switching model, the optimal control surface as a function of t and P with $I = 1000$ MMcf and the corresponding control curve at $P = 6$ \$/mmBtu. Comparing three control surfaces in Figures 4.1a, 4.2a and 4.2c, we observe that they have similar seasonal patterns except that the three bang-bang control regions in each control surface align according to different gas price P , or more precisely, according to the equilibrium prices of the stochastic processes with respect to the three control surfaces. Consequently, the MRMR variation generates different control strategies for two regimes that are consistent with the calibrated processes in those regimes. Moreover, as a regime shift occurs due to an exogenous event, the seasonal control pattern will change accordingly. For example, if the regime shifts from 0 to 1 in March, then the control will switch from producing gas to injecting gas when $P = 6$ \$/mmBtu. As a result, the MRMR variation of the regime-switching model is able to generate controls that reflect the existence of multiple regimes (if we believe this is true) in the market as well as the regime shifts. Therefore, we regard the MRMR variation as a richer model, which has more complex optimal control strategies.


In Figure 4.3, we plot the optimal control surface and the corresponding control curve for the MRGBM variation of the regime-switching model. The control strategies in regime 0 of the MRMR and MRGBM variations are similar. However, in regime 1 of the MRGBM variation, for all gas prices, it is never optimal to produce, even in the winter period (for fixed $I = 1000$). Again, this is in accordance with the GBM behavior of the gas price process in this regime. Therefore, similar to




(A) Control surface, Regime 0, MRMR



(B) Control curve at $P = 6$ \$/mmBtu, Regime 0, MRMR



(C) Control surface, Regime 1, MRMR



(D) Control curve at $P = 6$ \$/mmBtu, Regime 1, MRMR

FIGURE 4.2: Optimal control surface for the MRMR variation of the regime-switching model as a function of forward time t and gas spot price P as well as the corresponding control curve as a function of t when $P = 6$ \$/mmBtu, where the gas inventory resides at $I = 1000$ MMcf. Model parameter values are given in Tables 3.1 and 3.4. Fully implicit timestepping with the no bang-bang method and with constant timesteps is used. Other input parameters are given in Table 4.1. We assume that the starting time $t = 0$ corresponds to the July 1st of the year.

the MRMR variation, the MRGBM variation also produces regime specific control strategies that are consistent with the gas price dynamics in each regime. Consequently, the MRGBM variation can also produce very different optimal strategies compared to the MR model.

Finally, we note that from a calibration perspective, it is difficult to distinguish between the MRMR or MRGBM models. We would need other evidence to determine whether the gas price dynamics in the high price regime is mean reverting or GBM.

5 Conclusion

In this paper, we propose a one-factor regime-switching model for the risk adjusted natural gas spot price and study the implication of the model on the valuation and optimal operations of natural gas storage facilities.

By adjusting parameter values, the deseasoned process in each regime follows either a mean-reverting process or a geometric Brownian motion (GBM) like process with a positive/negative drift. This produces several variations of the basic model. We calibrate model parameters to both market futures prices and options on futures. The calibration to futures indicates that the MRMR variation (the processes in both regimes are mean-reverting) and the MRGBM variation (the process in one regime is mean-reverting and that in the other regime follows a GBM with a positive drift) of the regime-switching model are able to capture the dynamics of the market futures prices more accurately than a one-factor mean-reverting model. One-factor mean-reverting models have been the standard choices of other authors who have developed methods for valuing natural gas storage.

The valuation of a natural gas storage facility is characterized as a stochastic control problem, resulting in Hamilton-Jacobi-Bellman (HJB) equations. We extend the fully implicit, semi-Lagrangian discretizations in Chen and Forsyth (2007) to solve the pricing equations in the case of the regime-switching model. Provided a strong comparison result holds, we then prove the convergence of the scheme to the viscosity solution of the pricing equations. Numerical results show that our numerical scheme can achieve a first-order convergence rate.


We also observe from the numerical experiments that the optimal control strategies implied from the MRMR and MRGBM variations of the regime-switching model are able to produce distinct controls for each regime that are consistent with gas price dynamics in that regime. Therefore, the regime-switching model can reflect the existence of multiple regimes in the market as well as the regime shifts driven by various exogenous events, and hence it is a richer model for pricing the gas storage contracts than one-factor mean-reverting models. The regime-switching models result in qualitatively different optimal strategies compared to the simple mean-reverting model.

Of course, there are many ways that the regime-switching model can be improved to obtain better fits to the market data. Nevertheless, it appears that this comparatively simple model is a major improvement compared to the mean-reverting models used in previous work on gas storage.


In addition, solution of the HJB equation with the regime switching model is only about twice as computationally expensive as solving the HJB equation with the simple mean-reverting model. This is much less computationally expensive compared to a full two factor model.

References


Ahn, H., A. Danilova, and G. Swindle (2002). Storing arb. *Wilmott Magazine* 1.



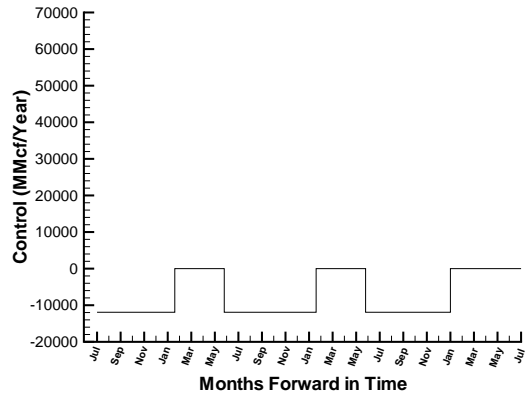
(A) Control surface, Regime 0, MRGBM



(B) Control curve at $P = 6$ \$/mmBtu, Regime 0, MRGBM



(C) Control surface, Regime 1, MRGBM



(D) Control curve at $P = 6$ \$/mmBtu, Regime 1, MRGBM

FIGURE 4.3: Optimal control surface for the MRGBM variation of the regime-switching model as a function of forward time t and gas spot price P as well as the corresponding control curve as a function of t when $P = 6$ \$/mmBtu, where the gas inventory resides at $I = 1000$ MMcf. Model parameter values are given in Tables 3.1 and 3.4. Fully implicit timestepping with the no bang-bang method and with constant timesteps is used. Other input parameters are given in Table 4.1. We assume that the starting time $t = 0$ corresponds to the July 1st of the year.

- Bansal, R. and H. Zhou (2002). Term structure of interest rates with regime shifts. *The Journal of Finance* 57(5), 1997–2043.
- Barles, G. (1997). Convergence of numerical schemes for degenerate parabolic equations arising in finance. In L. C. G. Rogers and D. Talay (Eds.), *Numerical methods in finance*, pp. 1–21. Cambridge University Press, Cambridge.
- Barles, G. and J. Burdeau (1995). The Dirichlet problem for semilinear second-order degenerate elliptic equations and applications to stochastic exit time control problems. *Communications in Partial Differential Equations* 20, 129–178.
- Barles, G. and E. Jakobsen (2007). Error bounds for monotone approximation schemes for parabolic Hamilton-Jacobi-Bellman equations. *Mathematics of Computation* 76, 1861–1893.
- Barles, G. and E. R. Jakobsen (2002). On the convergence rate of approximation schemes for Hamilton-Jacobi-Bellman equations. *M2AN Mathematical Modelling and Numerical Analysis* 36, 33–54.
- Barles, G. and E. Rouy (1998). A strong comparison result for the Bellman equation arising in stochastic exit time control problems and its applications. *Communications in Partial Differential Equations* 23, 1945–2033.
- Barles, G. and P. E. Souganidis (1991). Convergence of approximation schemes for fully nonlinear equations. *Asymptotic Analysis* 4, 271–283.
- Barrera-Esteve, C., F. Bergeret, C. Dossal, E. Gobet, A. Meziou, R. Munos, and D. Reboul-Salze (2006). Numerical methods for the pricing of swing options: a stochastic control approach. *Methodology and Computing in Applied Probability* 8(4), 517–540.
- Carmona, R. and M. Ludkovski (2005). Gas storage and supply guarantees: an optimal switching approach. Working paper, Princeton University.
- Cartea, A. and T. Williams (2006, August). UK gas markets: the market price of risk and applications to multiple interruptible supply contracts. Working Paper, Birkbeck, University of London.
- Chen, Z. and P. A. Forsyth (2007). A semi-Lagrangian approach for natural gas storage valuation and optimal operation. To appear in *SIAM Journal on Scientific Computing*.
- Cont, R. and E. Voltchkova (2005). Integro-differential equations for option prices in exponential Levy models. *Finance and Stochastics* 9(3), 299–325.
- Cortazar, G. and E. S. Schwartz (2003). Implementing a stochastic model for oil futures prices. *Energy Economics* 25, 215–238.
- de Jong, C. (2005). The nature of power spikes: a regime-switching approach. Working paper, Rotterdam School of Management at Erasmus University.
- de Jong, C. and R. Huisman (2002). Option formulas for mean-reverting power prices with spikes. Working paper, Rotterdam School of Management at Erasmus University.

- Deng, S. (2000). Pricing electricity derivatives under alternative stochastic spot price models. In *Proceedings of the 33rd Hawaii International Conference on System Sciences*.
- D'Halluin, Y., P. A. Forsyth, and G. Labahn (2005). A Semi-Lagrangian approach for American asian options under jump diffusion. *SIAM Journal on Scientific Computing* 27(1), 315–345.
- Forsyth, P. A. and G. Labahn (2007). Numerical methods for controlled Hamilton-Jacobi-Bellman PDEs in finance. to appear, *Journal of Computational Finance*.
- Gray, S. (1996). Modeling the conditional distribution of interest rates as a regime switching process. *Journal of Financial Economics* 42, 27–62.
- Hamilton, J. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics* 45, 39–79.
- Hardy, M. R. (2001). A regime-switching model of long-term stock returns. *North American Actuarial Journal* 5(2), 41–53.
- Huisman, R. and R. Mahieu (2003). Regime jumps in electricity prices. *Energy Economics* 25, 425–434.
- Jaillet, P., E. I. Ronn, and S. Tompaidis (2004, July). Valuation of commodity-based swing options. *Management Science* 50(7), 909–921.
- Jakobsen, E. R. (2003). On the rate of convergence of approximation schemes for Bellman equations associated with optimal stopping time problems. *Mathematical Models and Methods in Applied Science* 13, 613–644.
- Li, H. (2006, August). Adaptive wavelet collocation methods for option pricing PDEs. PhD Thesis, University of Calgary.
- Manoliu, M. and S. Tompaidis (2002). Energy futures prices: term structure models with kalman filter estimation. *Applied Mathematical Finance* 9(1), 21–43.
- Pilopović, D. (1998). *Energy Risk*. McGraw-Hill, New York.
- Pooley, D. M., P. A. Forsyth, and K. R. Vetzal (2003). Numerical convergence properties of option pricing PDEs with uncertain volatility. *IMA Journal of Numerical Analysis* 23, 241–267.
- Schindlmayr, G. (2005). A regime-switching model for electricity spot prices. Working paper, EnBW Trading GmbH.
- Schwartz, E. S. (1997). The stochastic behavior of commodity prices: implications for valuation and hedging. *The Journal of Finance* 52(3), 923–973.
- Thompson, M., M. Davison, and H. Rasmussen (2003, July). Natural gas storage valuation and optimization: a real options application. Working paper, University of Western Ontario.
- Ware, T. (2005, July). Swing options in a mean-reverting world. Presentation at: Stochastic Calculus and its applications to Quantitative Finance and Electrical Engineering, a conference in honour of Robert Elliott, Calgary.

- Ware, T. and H. Li (2006, June). Swing options with continuous exercise. Presentation at the Canadian Mathematical Society meeting, Calgary.
- Windcliff, H., P. A. Forsyth, and K. Vetzal (2004). Analysis of the stability of the linear boundary condition for the Black-Scholes equation. *Journal of Computational Finance* 8(1), 65–92.
- Xu, Z. (2004). Stochastic models for gas prices. Master’s thesis, University of Calgary.