

Gas storage valuation and hedging. A quantification of the model risk.

Patrick Henaff (1), Ismail Laachir (2) and Francesco Russo (3)

December 12th 2013

Abstract.

This paper focuses on the valuation and hedging of gas storage facilities, using a spot-based valuation framework coupled with a financial hedging strategy implemented with futures contracts. The first novelty consist in proposing a model that unifies the dynamics of the futures curve and the spot price, which accounts for the main stylized facts of the US natural gas market, such as seasonality and presence of price spikes. The second aspect of the paper is related to the quantification of model uncertainty related to the spot dynamics.

Key words and phrases. Energy markets; commodities; natural gas storage; model uncertainty.

JEL Classification: C4; C5; C6; C8; G11; G13; G17.

MSC Classification 2010: 91G10; 91G60; 91G70; 91G80.

(1) Patrick Henaff, IAE Paris, Université Paris I - Panthéon Sorbonne, France

(2) Ismail Laachir, Université de Bretagne-Sud and ENSTA ParisTech, France

(3) Francesco Russo, ENSTA ParisTech, Unité de Mathématiques appliquées, France.

1 Introduction

Natural gas storage units are used to reconcile the variable seasonal demand for gas with the more constant rate of natural gas production. These gas storage facilities are mainly owned by distribution companies which use them for system supply regulation, and to reduce the risk of shortages. In fact, due to regulation laws, local distribution companies are obliged to own storage units, to ensure their gas supply and to be able to meet any sudden increase in demand or any disruption in the pipeline transportation system.

In practice there are several techniques for gas storage valuation: The two most popular are the classical *intrinsic valuation* based on physical futures contracts, which is equal to the optimal discounted value of calendar spreads, and the *extrinsic valuation*, which uses spot trading strategies. Traditionally the demand for natural gas is seasonal, with peaks during winter, and lows during summer. This motivates the first valuation methodology, which exploits the predictable seasonal shape of the natural gas futures curve. Following this strategy, the storage manager observes the initial futures curve, at the beginning of the storage contract, decides to buy/sell multiple futures contracts, and consequently receives/delivers natural gas at their expiration. In order to determine the optimal futures positions, a linear optimization problem has to be solved, with constraints imposed by the physical and financial conditions of the storage contract (see Annexe B which refers to Eydeland and Krzysztof [2002]). We emphasize that storage manager keeps the optimal futures positions for the whole storage contract duration, so this strategy does not take advantage of possible profitable movements of the futures curve.

This static methodology is extended by Gray and Khandelwal [2004] to the *rolling intrinsic valuation*, to take advantage of the changing dynamics of the futures curve. In that case, optimal futures positions are chosen at the beginning of the storage contract, but when the futures curve moves away from its initial shape, the new optimal futures positions is recalculated and if managers find it more profitable, the portfolio is rebalanced. Hence, the evaluation of the rolling intrinsic strategy requires a model for the dynamics of the futures curve.

These two futures-based approaches capture the predictable seasonal pattern of natural gas prices: In particular they lead to buy cheap summer futures and sell expensive winter futures, and the storage value obtained greatly depends on the seasonal spread between cold and warm periods of the year.

The intrinsic valuation methodology has been popular in the storage industry, especially during periods when seasonal patterns are very pronounced. However, during the last years, the seasonal spreads have been reduced, which puts into perspective the futures-based methodology. In fact, the *2011 State of the Markets* report, by the US Federal Energy Regulatory Commission (FERC) FERC [2012], noticed the following: “We have also seen a decline in the seasonal difference between winter and summer natural gas prices. Falling seasonal spreads reflect increased production and storage capacity, as well as greater year-round use of natural gas by power generators. This decline has developed over the past several years and we expect the trend to continue.” The narrowing winter/summer spread, mentioned earlier, is mainly due to two factors that put a downward pressure on winter gas prices and an upward pressure on summer prices. The first factor is the recent surge of the non conventional shale gas supply, with geographical locations that are closer to gas consumption areas. The

main result of this new abundant source of gas is a downward pressure on winter prices. The second factor is related to power consumption by cooling systems during summer periods and the growing use of natural gas as a fuel for electricity generation, instead of other sources of power. This puts an upward pressure on summer gas prices.

The combination of these two factors has the logical consequence of narrowing the seasonal spreads between winter and summer prices, diminishing the intrinsic value of gas storage units. The use of futures contracts, exclusively as instruments to monetize the value of the gas storage, is no longer sufficient: in particular it sometimes fails to recover the operating expenses. This motivates the interest in valuation strategies based on the gas spot prices (instead of futures). This so-called *extrinsic valuation* can still take advantage of the remaining seasonality and more importantly allows the monetization of the high volatility of natural gas spot prices.

The first contribution of this paper to the topic is the proposal of a new modeling framework, unifying the natural gas futures curve and the spot price, taking into account the stylized facts that are essential in the gas storage valuation problem i.e., seasonality and spikes. The second aspect of the paper is related to the quantification of model uncertainty related to the spot dynamics.

In fact, the main result of Section 7 is the significant sensitivity of gas storage value with respect to the specification and estimation of the spot model. This result puts into perspective the extensive literature on gas storage valuation, and calls for a more careful study of the model risk inherent to our problem. We believe that it is crucial to turn more attention to the choice of the spot-futures modeling framework, rather than concentrate all the effort on the specification of an optimal trading strategy.

This paper is organized as follows. In Section 2 we describe the important stylized facts related to natural gas markets. In Section 3 we describe the characteristics of the gas storage unit and valuation using an optimal spot strategy, and the futures-based hedging methodology. Section 4 is devoted to a review of the modeling approaches in the gas storage literature. In Section 5 we introduce the modeling framework combining futures and spot dynamics and in Section 6 we perform several numerical tests. In Section 7 we introduce two natural model risk measures to quantify the sensitivity of a class of models with respect to the parameters; those risk measures are computed in several test cases.

2 Natural gas stylized facts

In this section, we highlight important stylized facts about natural gas markets that influence the value of a storage unit. These properties are related to the demand and use of natural gas. In fact, the demand for natural gas for heating in cold periods of the year produces a

seasonal behaviour for prices during winter periods, while unpredictable changes in weather can cause sudden shifts in gas prices. These facts are the two main sources of value for a gas storage unit, which allows for the exploitation of seasonality and sudden variations in demand.

As for all other commodities, the prices of natural gas (NG) are influenced by their geographical location. In our study, we will be interested in the United States market, specifically in a storage location near Henry Hub (Louisiana), which justifies the use of gas daily spot prices and the Nymex natural gas futures as hedge instruments. Essentially, it is possible to buy natural gas using spot or futures contracts. In the spot market the prices are settled every day for a delivery on the day after. In the futures market the prices of 72 monthly futures contracts are available on every business day, but there are only about 24 or fewer liquid contracts.

In what follows, S_t will denote the spot price of natural gas at date t , and $(F(t, T_i))_i$ represent the futures contracts prices at t , for a set $\{T_i\}$ of maturities. We consider monthly spaced maturities, so every futures contract is related to a month of a year. Also, we denote by P_t the price of prompt contract, i.e. the futures contract with the closest maturity to current time t . Natural gas prices are quoted in U.S. dollars per million British thermal units (MMBtu).

As mentioned above, the first main feature of natural gas prices is constituted by the presence of a seasonality component: we plot the NG futures curve for several dates in Figure 1. As noted above, we remark repeated rises in prices that occur during every winter, which are clearly due to the demand for heating during cold periods of the year. In addition to this traditional seasonal feature, the use of natural gas for electricity generation has created a second smaller increase during the summer period, because of the cooling systems during warm periods. These expected patterns in natural gas prices are the first source of value of a storage unit, since it is possible to buy summer futures contracts and store the gas delivered during summer, and sell more valuable winter futures contracts, which implies gas withdrawal during cold periods. This is the basic idea behind the so-called "intrinsic strategy," which is based only on futures contracts and exploits the calendar spreads in the futures curve (see Appendix B for more details about the intrinsic value of a storage unit.)

The second important aspect of natural gas prices, is the presence of sudden moves due to unexpected rise in demand (caused in general by an unpredicted weather change), technical problems in the supply chain, or a poor anticipation of the global stored gas available in the market. Because of delivery constraints (one day) of spot contracts, such sudden changes are almost instantaneously reflected in the spot dynamics, giving rise to large shifts in prices, rapidly absorbed by the storage capacities available in the market. These large and quickly absorbed jumps, commonly called spikes, can be viewed in Figure 2, which shows many sud-

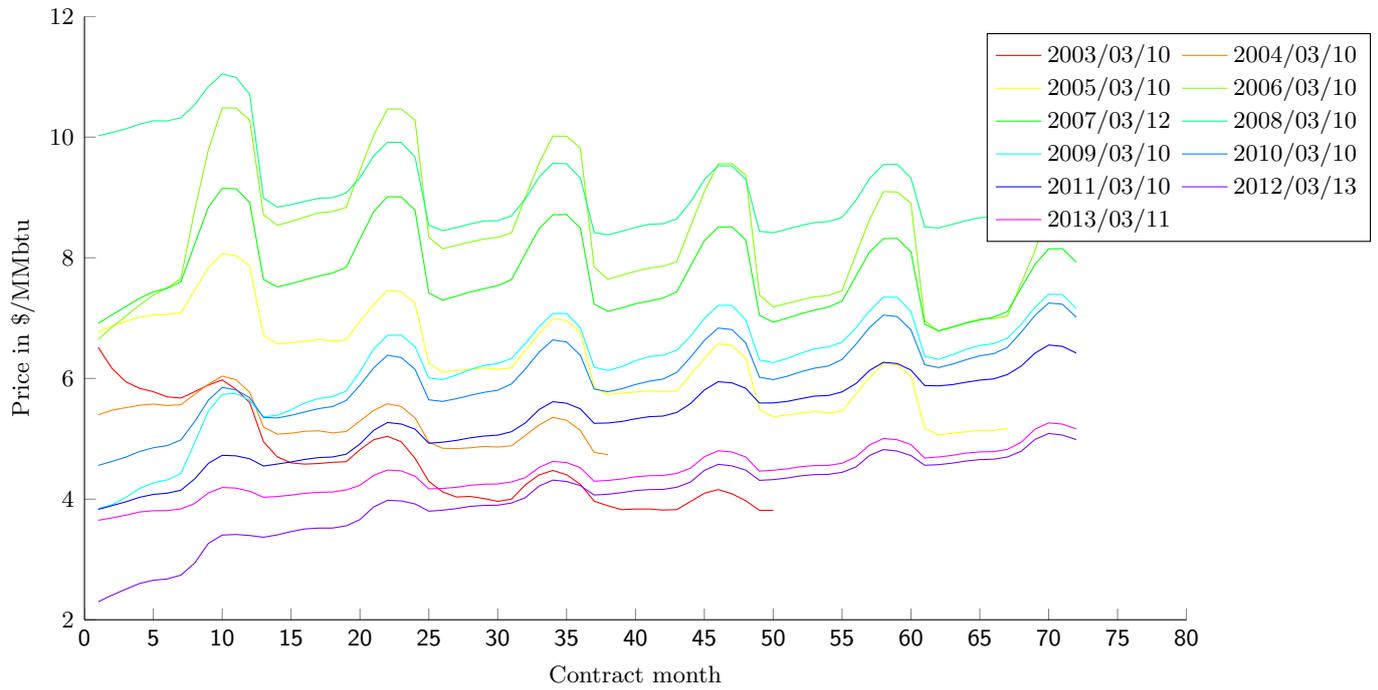


Figure 1: Futures curve for Nymex NG at different observation dates

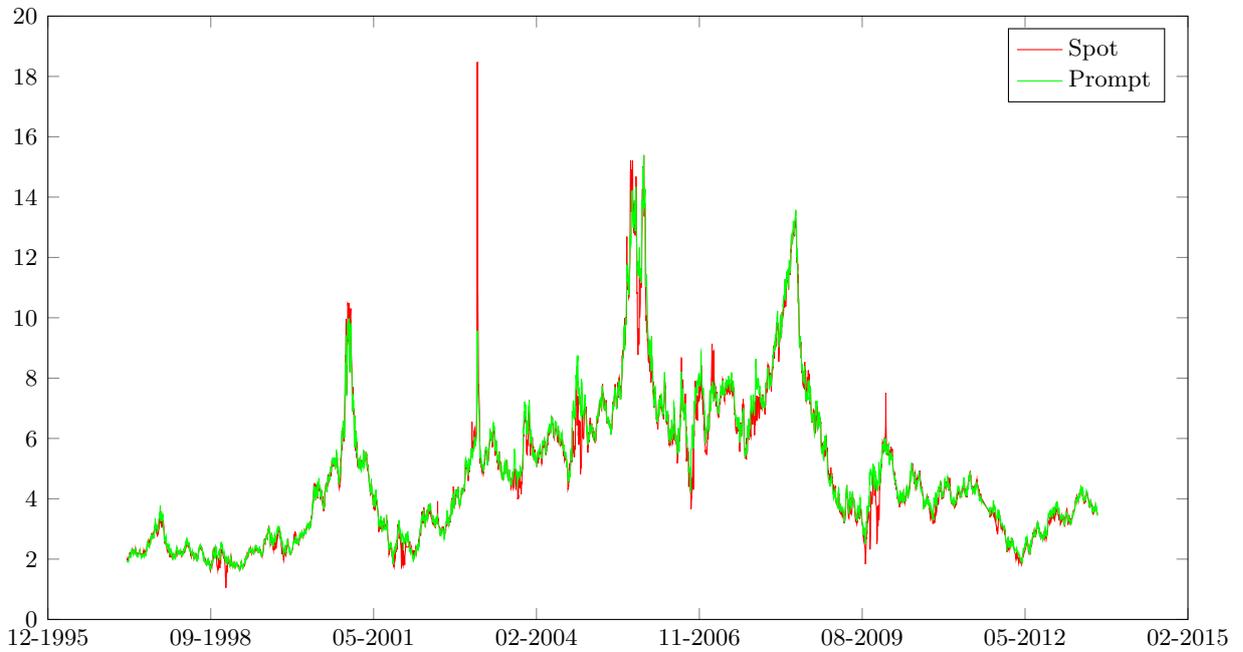


Figure 2: Spot and prompt historical prices

den dislocations between spot and prompt prices.¹ For example, we can notice a large spike in the spot price during late February 2003, when the natural gas price jumped by almost 78.00% and 54.26% in two successive days, then went down by -43.34% and -19.58% during the two following days. As noted by the US Federal Energy Regulatory Commission (FERC) by FERC [2003], this spike in gas price was due to “physical market conditions leading to low supply and high demand for a short time.” FERC [2003] also observed that “similar natural gas price spikes are possible when episodes of cold weather occur at times when storage inventories are limited.”

The appearance of the spikes is correlated to the spot and prompt prices spread. Indeed, while the prompt contract is a good proxy for the spot price, it does not suffer from sudden shocks of the same amplitude as spot prices, because of time-to-maturity factor. The spikes provoke ‘unusual’ gaps between the two contracts. In our study, we detect spikes by identifying the outliers from the time series (x_t) of the spread between the spot price S_t and the prompt price P_t given by $x_t := \frac{S_t - P_t}{P_t}$; we separate the study of positive and negative spikes, since they reflect two different market conditions. In fact, positive spikes are often caused by unpredicted weather changes, such as a cold front or a heat wave. On the other hand, negative spikes are generally due to a poor anticipation of market-wide gas storage levels. In Figure 3 we plot the number of occurrences of negative and positive spikes during each month. We remark that the repartition of the spikes is clearly dependent on their sign. In fact, most of the positive spikes happen during the winter months of January and February and the summer month of June, which can be explained by the occurrence of an unpredicted cold front or heat wave. On the other hand, the negative spikes appear during the months of pre-heating periods, i.e. September, October and November. One plausible explanation is given by Mastrangelo [2007], which states: “October is the last month of the refill season. There may be increased competition from storage facilities looking to meet end-of-season refill goals as well as increased anticipation regarding the upcoming heating season.”

In order to take into account the stylized facts mentioned above, first, our futures model incorporates seasonality in the futures curve, and second, the spot model describes the existence of spikes and takes into account the correlation between spot and futures prices, through the prompt contract. As far as our knowledge is concerned, these two facts were not taken into account in the literature related to gas storage valuation; we believe that they constitute the main sources of storage value. Indeed, storage managers can exploit the winter/summer seasonality of gas prices, buying during summer and selling during winter. The presence of spikes can be monetized (if the injection/withdrawal rates are high enough), which is possible for some high deliverability storage units (see Appendix A.)

Before discussing modeling issues, we start by specifying the gas storage unit valuation

¹We use a 1997-2013 historical data of spot and prompt price, published by the U.S. Energy Information Administration. cf http://www.eia.gov/dnav/ng/ng_pri_fut_s1_d.htm

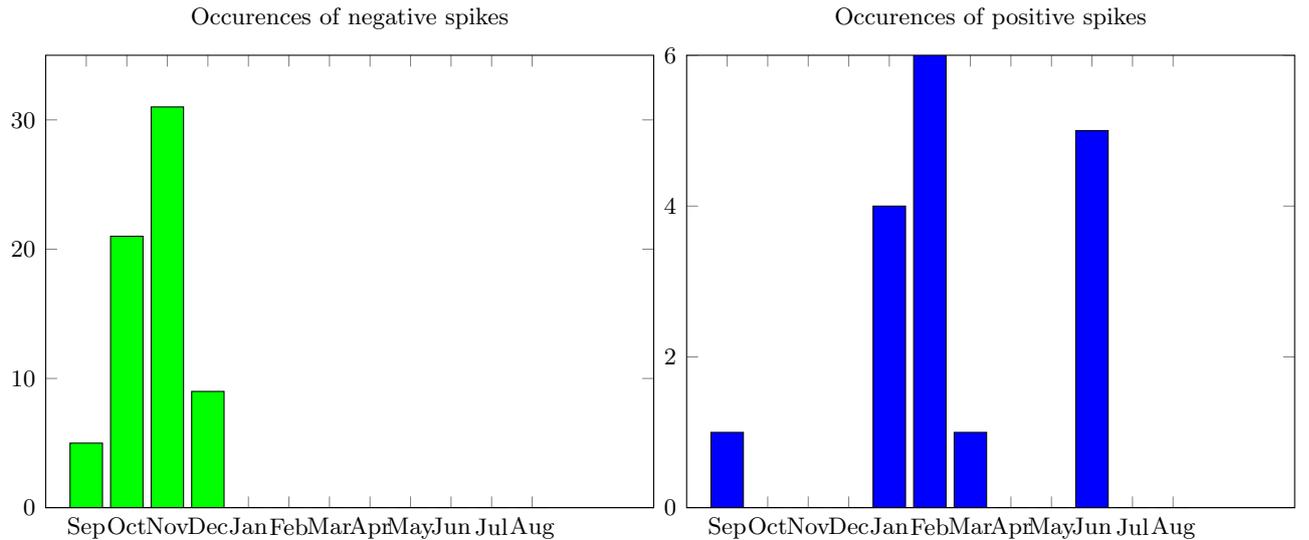


Figure 3: Occurrences of spikes

and hedging problem, and recall numerical simulation algorithms.

3 Valuation and hedging of a gas storage utility

The problem of valuing gas storage units has been discussed from many angles in the literature, yielding different approaches and numerical methods. Leasing a gas storage unit is equivalent to paying for the right, but not the obligation, to inject or withdraw gas from the unit. Hence the goal of the owner is to optimize the use of the gas storage facility, by injecting or withdrawing gas from the unit and, at the same time, trading gas on the spot and/or futures market. All these decisions have to be made under many operational constraints, such as maximal and minimal volume of the storage, and limited injection and withdrawal rates. This forces the resolution of a constrained stochastic control problem.

The gas spot price is modeled by a process, denoted by S . We suppose that this process is given as a function of a Markov process X in term of which the optimal control problem will be expressed. For example, in the framework (4.12), used by Boogert and De Jong [2008], the spot process is a Markov process, so we take obviously $X = S$.

Warin [2012] takes the futures curve as the underlying process, such that

$$F(t, T) = F(0, T) \exp \left[-\frac{1}{2} V(t, T) + \sum_{i=1}^n e^{-a_i(T-t)} W_t^i \right], \quad (3.1)$$

with $W_t^i = \int_0^t \sigma_i(u) e^{-a_i(T-u)} dZ_u^i$, Z^i being standard Wiener processes, and $V(t, T) = \text{var} \left(\sum_{i=1}^n e^{-a_i(T-t)} W_t^i \right)$. Indeed, the Markov process X can be chosen to be equal to the random sources $\bar{X} = (W^1, \dots, W^n)$.

In fact the futures prices modeled in (3.1) are martingales. In particular one supposes that the underlying probability is risk-neutral probability, and not necessarily unique.

Remark 1. *For simplicity, in the rest of this article, we suppose the discount interest rate to vanish, and consider the problem specification in a time discrete setting.*

In the next part of the paper we present the specification of the gas storage valuation problem, using the notations of Warin [2012].

3.1 Gas storage specification

We consider a gas storage facility with technical constraints (either physical or regulatory) on the volume of stored gas, V_{min} and V_{max} i.e. at each time, the volume level of the stored gas V should verify $V_{min} \leq V \leq V_{max}$.

We suppose a discrete set of dates $t_i = i\Delta t$ for $i = 0, \dots, n-1$ with $\Delta t = T/n$. At each date t_i and starting from a volume V_{t_i} , the user has the possibility to make one of three decisions: either inject gas at rate of a_{inj} , or withdraw gas at rate of a_{with} or take no action. We denote by u_i the decision at time t_i , and write $u_i = inj$ (resp. $with$, no) if the decision is injecting gas at rate a_{inj} (resp. withdrawing gas at rate a_{with} , no action).

If the user decides to follow a strategy $(u_i)_{i=0\dots n-1}$, then the volumes of gas in the storage $(V_i)_i$ are given by the iteration

$$V_0 = v, \tag{3.2}$$

$$V_{t_{i+1}}(u) = \begin{cases} \min(V_{t_i}(u) + a_{inj}\Delta t, V_{max}) & \text{if } u_i = inj \\ \max(V_{t_i}(u) - a_{with}\Delta t, V_{min}) & \text{if } u_i = with \\ V_{t_i}(u) & \text{if } u_i = no, \end{cases} \tag{3.3}$$

for $i = 0, \dots, n-2$. The generated cash flow -selling, in the case of gas withdrawal, or buying, in the case of gas injection- is given by

$$\phi_{u_i}(S_{t_i}) := S_{t_i}(V_{t_{i+1}}(u) - V_{t_i}(u)).$$

In general the maximum and minimum injection and withdrawal speeds (a_{inj} and a_{with}) are functions of the amount of gas in storage. However, without loss of generality, we assume for simplicity that these bounds are constant. In Table 1 we summarize the possible decisions and their consequences on gas volume and generated cash flow. We recall that a_{inj} indicates the injection rate per time unit Δt and a_{with} the withdrawal rate per time unit Δt . Consequently, the wealth generated by following a strategy u is given by

$$\text{Wealth}_{\text{spot}}(u) = \sum_{i=0}^{n-1} \phi_{u_i}(S_{t_i}). \tag{3.4}$$

Decision u	Next volume	Cash flow
Injection: $u_i = inj$	$V_{t_{i+1}}(u) = \min(V_{max}, V_{t_i}(u) + a_{inj}\Delta t)$	$\phi_{inj} = S_{t_i}(V_{t_i}(u) - V_{t_{i+1}}(u))$
Withdrawal: $u_i = with$	$V_{t_{i+1}}(u) = \max(V_{min}, V_{t_i}(u) - a_{with}\Delta t)$	$\phi_{with} = S_{t_i}(V_{t_i}(u) - V_{t_{i+1}}(u))$
No Action: $u_i = no$	$V_{t_{i+1}}(u) = V_{t_i}(u)$	$\phi_{no} = 0$

Table 1: Possible decisions

Finally, we are interested in the expectation of this cumulative cash flows, which we denote by J . More precisely we set

$$\begin{aligned}
J(t_0, x_0, v_0; u) &:= \mathbb{E}[\text{Wealth}_{\text{spot}}] \\
&= \mathbb{E}\left[\sum_{i=0}^{n-1} \phi_{u_i}(S_{t_i})\right],
\end{aligned} \tag{3.5}$$

J is a function that depends on the initial time t_0 , the value of the Markov process $X_0 = x_0$, the initial volume in the storage v_0 and the strategy u .

The goal of the storage operator is to find a strategy u maximizing the expected cumulative cash flows. We denote this optimal value by J^* . So the problem to solve is the following:

$$\begin{aligned}
J^*(t_0, x_0, v_0) &= \max_{(u_i)_{i=0\dots n-1}} J(t_0, x_0, v_0; u) = \max_{(u_i)_{i=0\dots n-1}} \mathbb{E}\left[\sum_{i=0}^{n-1} \phi_{u_i}(S_{t_i})\right] \\
&= J(t_0, x_0, v_0; u^*).
\end{aligned} \tag{3.6}$$

A priori, the underlying probability is the historical probability measure, at least if the futures do not intervene in the spot model. This quantity constitutes an objective for the manager. However, it is not a “fair” price in the sense of “absence of arbitrage,” since the spot is not traded as a financial asset; in that case the price would be an expectation with respect to a risk-neutral probability. On the other hand, J^* constitutes a price indicator; practitioners trade gas storage units at a proportion of this price.

In our proposed framework (see Section 5) the spot and the futures are jointly modeled on a product space $\Omega = (\Omega_s, \Omega_f)$ equipped with a probability \mathbb{Q} . The futures are first directly described as martingales on Ω_f with respect to their corresponding risk-neutral probability \mathbb{P}^* and they are extended trivially to Ω . Formally \mathbb{Q} is defined by $\mathbb{Q}(d\omega_s, d\omega_f) = Q^{\omega_f}(d\omega_s)P^*(d\omega_f)$, where $Q^{\omega_f}(d\omega_s)$ is a random probability kernel (historically considered), describing the random behavior of S for each realization ω_f of the futures asset. The expectation of the optimal cumulative cash flows with respect to Q will be then a price indicator, compatible with classic financial principles, as far as futures assets are concerned. Indeed, we will also estimate the volatility parameters for the diffusion describing the futures assets F using historical data, that is under some historical probability \mathbb{P} and not \mathbb{P}^* as we would

need. However, the probability \mathbb{P}^* is equivalent to \mathbb{P} and this justifies the coherence of the estimation.

3.2 Dynamic programming equation

From Table 1, we recall that $V_{t_{i+1}}(u)$ only depends on $V_{t_i}(u)$ and u_i . To emphasize this fact, if $V_{t_i} = v$, we also express $V_{t_{i+1}}(u)$ by $\widehat{V}_{u_i}(v)$.

At time t , for $X_t = x$ and with current volume level v , the (optimal) value for gas storage will be of course denoted by $J^*(t, x, v)$. The dynamic programming principle implies

$$J^*(t_i, x, v) = \max_{u_i \in \{inj, no, with\}} \{ \phi_{u_i} + \mathbb{E}[J^*(t_{i+1}, X_{t_{i+1}}, \widehat{V}_{u_i}(v)) | X_{t_i} = x, V_{t_i} = v] \}. \quad (3.7)$$

The classic way to solve this problem numerically is to use Monte Carlo simulations, combined with the Longstaff and Schwartz [2001] algorithm, which approximates the above conditional expectation, using a regression technique. This backward algorithm yields an estimate of the optimal strategy. As noted by Boogert and De Jong [2008], the main difficulty comes from the fact that the value function also depends on volume level, which in turn depends on the optimal strategy.

To circumvent this difficulty, Boogert and De Jong [2008] suggest discretizing the volume into a finite grid, $v_l = V_{min} + l\delta$, $l = 0, \dots, L = (V_{max} - V_{min})/\delta$ where L is the number of volume subintervals. However, the fact that the time grid is discrete and that at each time the storage unit manager has only three possible actions implies that the number of attainable volumes for any strategy is finite. In fact, at each time t_i , the set $\mathcal{V}(i)$ of possible volumes is given by

$$\mathcal{V}(i) = \{V_i = v_0 + ka_{inj}\Delta t + la_{with}\Delta t, \text{ such that } V_{min} \leq V_i \leq V_{max} \text{ and } k, l \in \mathbb{N}, k + l \leq i\};$$

consequently, it is possible to solve the dynamic programming equation (3.7) for all volumes in $\mathcal{V}(i)$. The only motivation to use a restricted volume grid would be the reduction of computation time.

We then get the following equation at time t_i , for each path simulation X^m , where $m = 1, \dots, M$, M being the total number of realization paths, and each volume level $v_l \in \mathcal{V}(i)$:

$$J^*(t_i, X_{t_i}^m, v_l) = \max_{u_i \in \{inj, no, with\}} \{ \phi_{u_i} + \mathbb{E}[J^*(t_{i+1}, X_{t_{i+1}}, \widehat{V}_{u_i}(v_l)) | X_{t_i} = X_{t_i}^m, V_{t_i} = v_l] \}. \quad (3.8)$$

The conditional expectation above is estimated using Longstaff-Schwartz regression algorithm, at each volume grid point v_l .

This will give us an estimation of the optimal strategy, denoted u^* , and the initial value of the gas storage unit $J^*(t_0, X_{t_0}, v_0) = J(t_0, X_{t_0}, v_0; u^*)$. The numerical resolution of problem (3.7) is done in two phases. The first stage consists of estimating the optimal strategy u^* , by

performing the backward iterations of equation (3.8), using regression techniques to estimate the conditional expectations, along a set of simulated paths. The second phase consists of estimating the value function J^* through the forward iterations of (3.8), along a new set of simulated paths, where we apply the estimated optimal strategy, as given by the backward algorithm.

One important remark about problem (3.7) is that the maximization is carried out for the expected wealth generated by the spot-trading strategy. Consequently, following the corresponding estimated optimal strategy on a single path will not ensure that the manager will recover the initial storage value J^* . There will certainly be a discrepancy between the realized cumulative cash flows on a given path and the expected value J^* . Hence, it is crucial for the storage manager to reduce the variance of the cumulative cash flows, which is a random variable. As we will explain in the next section, this will be realized by conducting a financial hedging strategy, based on futures contracts on natural gas.

3.3 Financial hedging strategy

After estimating the optimal gas strategy, the storage unit manager will follow these optimal decisions on the sample path revealed by the market. But one should keep in mind that, if one follows this optimal strategy u^* , the cumulative wealth is only the realization of a random variable whose expectation equals the initial price J^* of the gas storage unit. This motivates the interest in hedging strategies enabling better tracking of the storage value. This can be done by combining optimal gas strategies and additional financial trades, so that the expectation of the related cumulative wealth generated by both physical and financial operations is still J^* , but its variance (or some other risk criterion) is reduced. Analogously to Bjerksund et al. [2011], who treats the intrinsic value case, this additional financial hedging strategy plays a similar role to the control variate in the variance reduction of Monte Carlo simulations, as it preserves the expected cumulative cash flows and reduces its variance. In general to reduce the variance of a Monte Carlo estimator of a r.v. Y , one adds to it a mean zero control variate, which is highly (negatively) correlated to Y . Since futures contracts are the most liquid assets in the natural gas market, and are strongly correlated to the spot price, they form an ideal hedging instrument. In fact, although a futures contract price $F(t, T)$ does not converge to the spot price, when the time to maturity $T - t$ goes to zero, the correlation between the prompt contract (for example) and the spot price is very high, and often the two contracts move in the same direction. In practice, in the market, one has access to a set of futures standard contracts, with specified maturities $\{T_j\}_{1 \leq j \leq m}$. The basic idea of a financial hedging strategy is to add to the physical spot trading, a strategy of buying and selling, at a trading date t_i , a quantity $\Delta(t_i, T_j)$ of futures contracts $F(., T_j)$ for $1 \leq j \leq m$. Logically, those quantities will depend on the spot and futures prices, S and $\{F(., T_j)\}_{1 \leq j \leq m}$, but also on the current volume level.

If the gas storage manager follows such a hedging strategy, in addition to the spot physical trading, then the cumulative cash flows of those two combined strategies is equal to

$$\text{Wealth}_{\text{spot+futures}} = \sum_{i=0}^{n-1} \phi_{u_i^*}(S_{t_i}) + \sum_{i=0}^{n-1} \sum_{j=1}^m \Delta(t_i, T_j)(F(t_{i+1}, T_j) - F(t_i, T_j)). \quad (3.9)$$

Because the futures contract $F(\cdot, T_j)$ stops trading after its expiration date T_j , we use the convention $\Delta(t, T_j) = 0$, for $t \geq T_j$.

Since the futures price process is a martingale under the risk neutral probability, we have $\mathbb{E}_{t_i}[F(t_{i+1}, T_j)] = F(t_i, T_j)$. Hence, the expectation of this hedging strategy is null i.e.

$$\mathbb{E}\left[\sum_{i=0}^{n-1} \sum_{j=1}^m \Delta(t_i, T_j)(F(t_{i+1}, T_j) - F(t_i, T_j))\right] = 0.$$

Consequently, following the optimal spot strategy in parallel with a futures hedging portfolio gives the same cash flows in expectation, but very likely with lower variance.

$$\mathbb{E}[\text{Wealth}_{\text{spot+futures}}] = \mathbb{E}[\text{Wealth}_{\text{spot}}] = \mathbb{E}\left[\sum_{i=0}^{n-1} \phi_{u_i^*}(S_{t_i})\right] = J^*.$$

The specification of such a hedging strategy will of course depend on the nature of the relation between the spot price and the futures curve.

A heuristic strategy that is widely used in the industry is to take the quantity $\Delta_1(t_i, T_j)$ of futures $F(\cdot, T_j)$ to be equal to the conditional expectation of volume to be exercised during the delivery period of the futures contract, conditional on the information at t_i . More precisely, the heuristic delta is equal to the t_i -conditional expectation

$$\Delta_1(t_i, T_j) = \mathbb{E}_{t_i}\left[\sum_{T_{j-1} \leq t_l < T_j} V_{l+1}(u^*) - V_l(u^*)\right]. \quad (3.10)$$

We also propose a modification of this heuristic delta, where we use the concept of tangent process (Warin [2012]). If we assume that the prompt converges towards the spot, then we can write

$$\begin{aligned} \text{Wealth}_{\text{spot}} &= \sum_{i=0}^{n-1} (V_{i+1}(u^*) - V_i(u^*)) S_{t_i} \\ &\simeq \sum_{i=0}^{n-1} (V_{i+1}(u^*) - V_i(u^*)) P_{t_i} \\ &= \sum_j \sum_{T_{j-1} \leq t_l < T_j} (V_{l+1}(u^*) - V_l(u^*)) F(t_l, T_j). \end{aligned}$$

So, another heuristic delta Δ_2 can be defined, using the concept of tangent process

$$\Delta_2(t_i, T_j) = \mathbb{E}_{t_i}\left[\sum_{T_{j-1} \leq t_l < T_j} (V_{l+1}(u^*) - V_l(u^*)) \frac{F(t_l, T_j)}{F(t_i, T_j)}\right] \quad (3.11)$$

We emphasize that the definition of these two hedging strategies is based on heuristic reasoning. Therefore, the hedging will not be perfect and a residual risk will still remain.

As we will see in the numerical experiments, this financial hedging strategy allows for a significant reduction of the cash flows uncertainty of the spot trading strategy. In fact, the variance of the spot trading is quite reduced while conducting a futures hedging strategy, and an out-of-sample test applied over a price history of 10 years shows better wealth tracking for the hedging strategy.

In the next section we will present several approaches for spot and futures prices modeling, and study the consequences of a model choice.

4 Literature on price processes

Generally, the problem of gas storage unit valuation has been studied from the angle of numerical resolution, and not much interest has been paid to the modeling itself and its effects on the final output of the numerical scheme. In fact, we encountered two modeling approaches in the literature.

The first approach consists of modeling only the spot model, with classical mean-reverting models, as proposed by Boogert and De Jong [2008], with no spike features. This approach, while realistic, does not take into account the dependence of the problem with respect to the futures curve and its dynamics, and does not offer the possibility of a hedging strategy based on futures. The second approach is based on the modeling of the futures curve by multi-factor log-normal dynamics, and makes the assumption of considering the spot price equal to the limit of futures prices with time to maturity going to zero. We note, however, that this assumption does not conform to the commodities market. On the other hand, it enables the definition of a delta hedging strategy based on available futures contracts on the market. We give more details about these two approaches in what follows.

A very common framework consists of modeling the spot price as a mean-reverting process. For instance, Boogert and De Jong [2008] developed a Monte Carlo method for storage valuation, using the Least Square Monte Carlo method, as proposed by Longstaff and Schwartz [2001] for American options. They consider one factor model for the spot price, which is calibrated to the initial futures curve. The price process S is given by

$$\frac{dS_t}{S_t} = \kappa[\mu(t) - \log(S_t)]dt + \sigma dW_t, \quad (4.12)$$

where W is a standard Brownian motion, μ is a time-dependent parameter, calibrated to the initial futures curve $(F(0, T))_{T \geq 0}$, provided by the market; the mean reversion parameter κ

and the volatility σ are two positive constants. As pointed out by Bjerksund et al. [2011], this framework has several drawbacks with respect to the goal of capturing the value of the gas storage. In fact, the calibration of the time-varying function $\mu(t)$ is, as expected, quite unstable and gives unrealistic sensitivity of the spot dynamics, and hence of the gas storage value with respect to the initial futures curve. Also, more importantly, this spot modeling does not take into account the futures market and the possibility of trading strategies on futures contracts, while it is well known that spot price is strongly correlated with the evolution of the futures curve, especially the short-term contracts (prompt contract). The futures curve is only used as an initial input to calibrate the parameter μ , but no dynamics for the futures curve is assumed. Indeed, modeling the futures curve is important in order to formulate hedging strategies based on futures contracts.

To take into account the correlation between futures curve and spot, it seems reasonable to introduce models that combine the spot and prompt price in the same dynamics. Finally, 4.12 is too poor to describe stylized facts about the spot price, as seasonality and the presence of spikes, which are the main sources of the gas storage value. An enhancement of this model is proposed by Parsons [2013], who considers the following two-factor mean-reverting model:

$$\frac{dS_t}{S_t} = a[\mu(t) + \log(L_t) - \log(S_t)]dt + \sigma_{S,t}dW_t, \quad (4.13)$$

$$\frac{dL_t}{L_t} = b[\log(\mathcal{L}) - \log(L_t)]dt + \sigma_{L,t}dZ_t, \quad (4.14)$$

where S is the spot price described by mean-reversion dynamics whose long-run mean has a stochastic component L and deterministic value \mathcal{L} .

While this model is more realistic than the one factor model, it still suffers from the instability of the deterministic function μ , and still does not include the possibility of spikes in the spot price. The author defines the futures contract price as the expectation of the spot price at maturity date T . We emphasize that this definition implies that natural gas is delivered at the futures expiration T ; however in reality the delivery period is spread over a whole calendar month.

A second way to model the spot process is to consider it as the limit of the futures contract price as time to maturity goes to zero; in particular we have $S_t = \lim_{T \downarrow t} F(t, T)$. This approach was adopted by Warin [2012]; the author considers a n -factor log-normal dynamics for the futures curve:

$$\frac{dF(t, T)}{F(t, T)} = \sum_{i=1}^n \sigma_i(t) e^{-a_i(T-t)} dZ_t^i, \quad (4.15)$$

and by continuity the spot is given by $S_t = F(t, t)$. In this framework, the author presented a similar algorithm to Boogert and De Jong [2008] to estimate the optimal strategy for the

spot; moreover he gives formulae for the sensitivities of storage value with respect to futures contracts, which enable a hedging strategy to be set up, based on futures, in parallel to the spot optimal trading strategy, which reduces the uncertainty of the realized cash flows. This futures-based hedging strategy presents a big advantage, compared to the first approach, since it increases the manager's chances of recovering the storage value and consequently the price paid to rent the storage unit.

In conclusion, we have decided to jointly model the futures curve and the spot prices. Indeed, we will formulate a multi-factor model for the futures curve, that includes seasonality of natural gas futures prices, and a dynamics of spot price that is correlated to futures curve, more precisely to its short end. We will also incorporate the presence of spikes in the spot prices.

5 Our modeling framework

In Section 2 we discussed the main stylized facts of natural gas prices, which are seasonality and spikes. We believe that the incorporation of these two features is essential in order to monetize these two sources of value. Also, we emphasize that it is crucial to use a modeling framework that combines spot and futures curve dynamics, and takes into account the existence of a basis between spot and prompt prices.

In Section 5.1, we will introduce a two-factors model for the futures curve, with a seasonal component for instantaneous volatility. This parsimonious model has easy-to-interpret parameters and an efficient calibration procedure using futures curve historical data.

In Section 5.2, we discuss spot price modeling: we consider two models, with a clear relation to the prompt contract. We also include spikes by means of a fast-reverting jump process, similar to a model by Hambly et al. [2009], which was applied to the electricity market.

5.1 Modeling the futures curve

The first models for energy futures curves $F(t, T)$ were obtained through conditional expectations of S_T with respect to the current information at time t , where S is the spot price process, which is indeed modeled, also taking into account possibly stochastic quantities such as the convenience yield and the interest rate. This approach has several drawbacks such as the difficulty of observing or estimating those quantities and the problem of fitting the initial curve $F(0, T)$.

Hence a second stream of models was proposed to directly describe the futures curve, using multi-factor log-normal dynamics. For instance, Clewlow and Strickland [1999b] proposes a one-factor model for the futures curve; this was then extended by Clewlow and Strickland

[1999a] to a multi-factor setting. A two-factor version of this model can be expressed as

$$\frac{dF(t, T)}{F(t, T)} = e^{-\lambda(T-t)} \sigma_{ST} dW_t^S + \sigma_{LT} dW_t^L,$$

where λ , σ_{ST} and σ_{LT} are positive constants, and W^S and W^L are two correlated Brownian motions. This model has the advantage of exactly fitting the initial futures curve, and the dependence of the volatility on the maturity parameter, i.e. it is of term-structure type; however it does not take into account the essential seasonality feature. Note that this model is an adaptation of the well-known Gabillon [1991] model, originally proposed for spot prices. Our framework slightly modifies previous models, adding a seasonality component and introducing parameters that have an economical significance.

We will call it Seasonal Gabillon two-factor model. It is formulated as

$$\frac{dF(t, T)}{F(t, T)} = e^{-\lambda(T-t)} \phi(t) \sigma_S dW_t^S + (1 - e^{-\lambda(T-t)}) \sigma_L dW_t^L, \quad (5.16)$$

where W^S and W^L are two correlated Brownian motions, with $d\langle W^L, W^S \rangle_t = \rho dt$. The letters L and S stand respectively for Long term and Short term; λ , σ_S and σ_L are positive constants. The function $\phi(t) = 1 + \mu_1 \cos(2\pi(t - t_1)) + \mu_2 \cos(4\pi(t - t_2))$ weights instantaneous volatility with a periodic behaviour. It takes into account the winter seasonal peaks (resp. the secondary summer peak) by taking for example t_1 equal to January (resp. with t_2 equal to August). The coefficients μ_1 and μ_2 quantify the winter and summer seasonality contribution in the volatility: we expect the winter parameter μ_1 to be often larger, in absolute value, than the summer parameter μ_2 .

Our model constitutes an efficient framework, whose parameters are economically meaningful. Indeed, the parameters σ_L and σ_S play the role of a 'long-term', respectively 'short-term' volatility. Note that even if the model is expressed with a continuous set of maturities, in the real world we only have access to a finite number of maturities, for example, monthly spaced futures contracts.

In the next section we give more details about the meaning of each parameter and their estimation, using historical data of futures prices.

5.1.1 Model estimation

Many of the model parameters are almost observable, if we have sufficient historical data of futures curves at hand. In fact, σ_S and σ_L could be approximated by the volatility of short and long-dated continuous futures contracts, and ρ by their empirical correlation.

In fact, for $T \rightarrow \infty$, we can formally write $\frac{dF(t, T)}{F(t, T)} \simeq \sigma_L dW_t^L$, so a good approximation

for the long-term volatility is

$$\sigma_L^2 \simeq \frac{1}{m-1} \sum_{i=1}^m \left(\frac{z_{t_i}^L}{\sqrt{\Delta t_i}} - \bar{\mu}^L \right)^2,$$

where z_t^L is the log-return of a constant maturity long-dated contract, four years for example, and $\bar{\mu}^L = \frac{1}{m} \sum_{i=1}^m \frac{z_{t_i}^L}{\sqrt{\Delta t_i}}$.

On the other hand, for small times to maturity, i.e. $T-t \rightarrow 0$, we can ignore the long-term noise effect, and write $\frac{dF(t,T)}{F(t,T)} \simeq \sigma_S dW_t^S$, so that a good proxy for the spot volatility is the volatility of the rolling prompt contract, i.e. the contract with the nearest maturity

$$\sigma_S^2 \simeq \frac{1}{m-1} \sum_{i=1}^m \left(\frac{z_{t_i}^P}{\sqrt{\Delta t_i}} - \bar{\mu}^P \right)^2,$$

where z_t^P is the log-return of a prompt futures contracts and $\bar{\mu}^P = \frac{1}{m} \sum_{i=1}^m \frac{z_{t_i}^P}{\sqrt{\Delta t_i}}$.

We can also give an initial estimate for the correlation parameter ρ as

$$\rho \simeq \frac{1}{m-1} \frac{\sum_{i=1}^m \left(\frac{z_{t_i}^P}{\sqrt{\Delta t_i}} - \bar{\mu}^P \right) \left(\frac{z_{t_i}^L}{\sqrt{\Delta t_i}} - \bar{\mu}^L \right)}{\sigma_S \sigma_L}.$$

These rough estimates could be used directly, or as input parameters for a more rigorous statistical estimation procedure. For example, we can use the maximum likelihood method. For that, suppose we have a time series over dates t_1, \dots, t_m of futures prices maturing at T_1, \dots, T_n . We denote $z_t, t = t_i, i \in \{0, \dots, t_{m-1}\}$ the vector of price returns, Δt being the corresponding step $t_{i+1} - t_i$ and θ is the model parameters vector $\theta = (\lambda, \mu_1, \mu_2, \sigma_S, \sigma_L, \rho)$, we have

$$z_t = \begin{pmatrix} \frac{\Delta F(t, T_1)}{F(t, T_1)} \\ \cdot \\ \cdot \\ \cdot \\ \frac{\Delta F(t, T_n)}{F(t, T_n)} \end{pmatrix}, \quad H_t = \sqrt{\Delta t} \begin{pmatrix} e^{-\lambda(T_1-t)} \phi(t) \sigma_S, & (1 - e^{-\lambda(T_1-t)}) \sigma_L \\ \cdot & \cdot \\ \cdot & \cdot \\ \cdot & \cdot \\ e^{-\lambda(T_n-t)} \phi(t) \sigma_S, & (1 - e^{-\lambda(T_n-t)}) \sigma_L \end{pmatrix},$$

where $\Delta F(t, T_1) = F(t + \Delta t, T_1) - F(t, T_1)$. Then an Euler discretization of the SDE (5.16) gives the equation

$$z_t = H_t x_t, \quad t \in \{t_1, \dots, t_m\},$$

where (x_{t_i}) are independents Gaussian 2-d vectors such that

$$x_{t_i} \sim \mathcal{N}(0, \Sigma), \quad 1 \leq i \leq m,$$

where

$$\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}.$$

The likelihood maximization could then be written as the minimization of the function

$$L(x_{t_1}, x_{t_2}, \dots, x_{t_m} | \theta) = \frac{1}{m} \sum_{i=1}^m \log(\det(\Sigma)) + x_{t_i}^T \Sigma^{-1} x_{t_i},$$

and the $x_t, t \in \{t_1, \dots, t_m\}$ are given by $z_t = H_t x_t$, i.e.

$$x_t = (H_t^T H_t)^{-1} H_t^T z_t.$$

So, the model estimation procedure is equivalent to the following minimization problem

$$\begin{cases} \min L(x_{t_1}, x_{t_2}, \dots, x_{t_m} | \theta) = \log(\det(\Sigma)) + \frac{1}{m} \sum_{i=1}^m x_{t_i}^T \Sigma^{-1} x_{t_i} \\ \theta = (\lambda, \mu_1, \mu_2, \sigma_S, \sigma_L, \rho). \end{cases} \quad (5.17)$$

To illustrate, we apply this estimation procedure, to historical data based on a 1997-2007 futures curves history. As mentioned, the estimation problem (5.17) is solved using an optimization algorithm, with the rough estimates of σ_S , σ_L and ρ as initial point for the algorithm. We report in Table 2 the estimated parameters of the futures curve model.

Parameter	Value	Confidence interval
σ_S	0.4580	[0.4462,0.4698]
σ_L	0.1655	[0.1617,0.1694]
λ	0.7896	[0.7518,0.8274]
μ_1	0.0246	[-0.0015,0.0507]
μ_2	0.0038	[-0.0218,0.0294]
ρ	0.4113	[0.3737,0.4488]

Table 2: Estimated parameters using 1997-2007 futures curves history.

As expected, the short-term volatility is larger than the long-term volatility, which is a common feature in energy futures curve dynamics, and the winter contribution μ_1 in the seasonality component is larger than summer contribution μ_2 .

5.2 Modeling spot price

The main characteristic of the modeling framework introduced by Warin [2012] is the assumption that the futures price maturing at T , $F(t, T)$, converges to the spot price S_t when the time to maturity is close to zero; as we mentioned, this hypothesis is not realistic, for the simple reason that a futures contract $F(\cdot, T)$ delivery does not take place at the fixed expiration day T but spreads out over a period of one month. Besides, the delivery point of

futures and spot might be different. We also note that the spot is subject to additional noise compared to futures contracts, as unpredicted weather changes or technical incidents have a larger impact on the spot than on the futures itself.

All these considerations suggest the spot should be considered as a separate, but correlated, stochastic process when in general

$$S_t \neq \lim_{T \rightarrow t} F(t, T).$$

A model in that sense was proposed by Gray and Palamarchuk [2010], where the logarithm of the spot is a mean reverting process, whose mean-reversion level is a stochastic process equal to the prompt price. For a family of maturities $(T_i)_i$, the futures contract $F(t, T_i)$ is a log-normal process fulfilling

$$\frac{dF(t, T_i)}{F(t, T_i)} = \sigma(t, T_i) dW_t$$

and the spot price S_t evolves along

$$d \log(S_t) = (\theta_t + a \log(P_t) - a \log(S_t)) dt + \sigma_t^S dB_t, \quad (5.18)$$

where B and W are two correlated Brownian motions, and for the current date t , P_t denotes the prompt price, i.e.

$$P_t = F(t, T_i) \quad \text{for} \quad T_{i-1} \leq t < T_i.$$

In our opinion it is crucial to incorporate futures curve dynamics into the modeling of the spot prices, for instance a dynamics relating the spot and prompt futures price. Indeed, as shown by the historical paths of spot and prompt prices in Figure 2, the two processes are closely related. In fact they seem to move very often in the same direction, with some occasional dislocations of spot and prompt prices.

In what follows we will study two spot models, in relation to our futures curve model. These will be stated in discrete time.

5.2.1 Spot model 1

Our first spot model is similar to (5.18), which was introduced by Gray and Palamarchuk [2010]. Its dynamics, based on the spot log-return $y_t = \log(S_t/S_{t-1})$, is given by

$$\log(S_t/S_{t-1}) = a_1 + a_2 \log(P_{t-1}/S_{t-1}) + a_3 \log(P_t/P_{t-1}) + \epsilon_t \quad (5.19)$$

where (ϵ_t) is a Garch(p, q) process and again P is the prompt price.

Recall that a Garch(p, q) process ϵ verifies an autoregressive moving-average equation for the error variance

$$\begin{aligned} \epsilon_t &= \sigma_t z_t, \text{ where} \\ \sigma_t^2 &= \kappa + \sum_{i=1}^p \gamma_i \sigma_{t-i}^2 + \sum_{i=1}^q \alpha_i \epsilon_{t-i}^2 \end{aligned} \quad (5.20)$$

and z a white noise.

This model intends to capture both the heteroscedasticity of the natural gas spot price and the correlation between spot price and prompt futures price. In fact, similarly to (5.18), we remark here that in our dynamics (5.19), the spot price is mean reverting around a stochastic level equal to prompt price. On the other hand, the prompt log return is a supplementary explanatory variable of the spot log return. We recall that our futures model (5.16) incorporates seasonality in the futures curve dynamics; this implies that the spot dynamics itself, by means of the prompt price, has an implied seasonality component. This allows us to avoid the addition of some seasonal function into the spot dynamics (5.19).

5.2.2 Spot model 2

The second spot model we propose is based on the series of spot-prompt spread $y_t = \frac{S_t - P_t}{P_t}$. In fact, we model the spot-prompt spread, through the so-called front-back spread as regression variable. In other words, we write

$$\frac{S_t - P_t}{P_t} = a_1 + a_2 \frac{S_{t-1} - P_{t-1}}{P_{t-1}} + a_3 \frac{P_{t-1} - B_{t-1}}{B_{t-1}} + \epsilon_t, \quad (5.21)$$

where B_t is the price of the second nearby futures (also known as the back contract) and ϵ is a Garch(p, q) process.

(5.21) has the advantage of directly handling the spread between spot and prompt prices, which is probably a good indicator of the decisions to be made in gas storage management. In fact, a large positive spread value would possibly induce the decision to withdraw gas, while the inverse would motivate a gas injection. Also, as we pointed out in the introduction, the narrowing of the seasonal spread in the futures curve during last years has diminished the intrinsic value of gas storage units. Consequently, almost all the storage value is now concentrated in the extrinsic value, which is heavily dependent on the spot-prompt spread.

5.2.3 Spikes modeling

In Section 2, we showed that natural gas prices have two special characteristics: seasonality and presence of spikes. The first feature (seasonality), is included in the spot dynamics through the prompt (and the second nearby) futures contract. In fact, the futures curve dynamics (5.16) already has a seasonal component, so we have chosen not to add a supplementary seasonal part in the spot dynamics. On the other hand, spikes are included in the spot model via a jump process. These large and rapidly absorbed jumps are an essential feature of the spot, since they can be source of value for gas storage and they can be monetized if injection/withdrawal rates are high enough.

Indeed, we describe the spikes as a fast mean-reverting (jump) process, already introduced by Hambly et al. [2009], in the framework of the electricity market. These authors proposed a spot model for the power price that incorporates the presence of spikes via a process Y , being the solution of the equation

$$dY_t = -\beta Y_t dt + dZ_t, \quad Y_0 = 0, \quad (5.22)$$

where Z is a compound Poisson process of the type $Z_t = \sum_{i=1}^{N_t} J_i$, (N_t) is a Poisson process with intensity λ and $(J_i)_{i \in \mathbb{N}}$ is a family of independent identically distributed (iid) variables representing the jump size. Furthermore (N_t) and (J_i) are supposed to be mutually independent. The process Y can be written explicitly as

$$Y_t = Y_0 e^{-\beta t} + \sum_{i=1}^{N_t} e^{-\beta(t-\tau_i)} J_i. \quad (5.23)$$

We recall that the spot model is directly expressed as a discrete time process, indexed on the grid (t_i) introduced in Section 3. For that reason Y will be restricted to the same time grid.

Choosing a high value for the mean-reversion parameter β forces the jump process Y to revert very quickly to zero after the jump times τ_i , which constitutes a desired feature for natural gas spikes. In fact, the jumps in natural gas spot prices are rapidly absorbed, thanks to the storage capacities available in the market.

We emphasize that the two models (5.19) and (5.21) alone do not take into account the possibility of sudden spikes in the spot price. The process Y will be indeed incorporated into the dynamics in (5.19) and (5.21), by multiplying the spot process by the process $\exp(Y_t)$, i.e.

$$\tilde{S}_t = \exp(Y_t) S_t.$$

As we noted in Section 2, the natural gas spikes are clearly distinguished by their signs. In fact, positive spikes, due to unpredicted weather changes, occur exclusively during the winter and summer months. On the other hand, negative spikes, generally caused by poor anticipation of the storage capacities of the market, happen mostly during shoulder months like October and November. This motivates a separate modeling for these two categories of spikes. We will consider two processes Y^+ and Y^- for positive and negative spikes, each one verifying a slightly modified version of the equation (5.23), as follows:

$$Y_t^+ = \sum_{i=1}^{N_t} e^{-\beta(t-\tau_i)} J_i \mathbb{1}_{\tau_i \in I^+}, \quad (5.24)$$

where I^+ (resp. I^-) represents the positive (resp. negative) spikes occurring period, i.e. winter and summer (resp. shoulder months), as we observed in Section 2.

Consequently, the spot process that we consider for our gas storage valuation is

$$\tilde{S}_t = \exp(Y_t^+ + Y_t^-)S_t. \quad (5.25)$$

Finally (5.25) will have all the desired properties: it includes seasonality by relating the futures curve to the spot dynamics and it allows the presence of positive and negative spikes, each one generated by a separate jump process Y^+ and Y^- .

5.2.4 Model estimation

Similarly to the model for futures, we use historical data for spot and futures prices. The parameters estimation for the two spot dynamics (5.19), (5.21) proposed above is based on regression techniques and the classic estimation procedure for Garch processes. Similarly to Hambly et al. [2009], we also use the likelihood method to estimate spike process parameters, after filtering the underlying time series to extract the jumps. Note that the coefficient β is heuristically fixed.

Regression parameters	Value
a_1	-0.0044
a_2	0.2622
a_3	0.4467
Garch(1,1) parameters	Value
κ	1.6928e-005
γ_1	0.8764
α_1	0.1138
Spike process Y^+	
β	300
λ	0.8331
Jump Law	$\mathcal{N}(0.2579, 0.3910)$
Spike process Y^-	
β	300
λ	2.9488
Jump Law	$\mathcal{N}(-0.7624, 0.6402)$

Table 3: Spot model 1 parameters using 1997-2007 data

An analysis of the spot and futures historical data shows that a Garch(p, q) process of order $p = 1$ and $q = 1$ is sufficient, using higher order being of insignificant impact. As mentioned before, we use a large value for the spike reversion parameter β .

To illustrate, the estimation procedure of the parameters of the spot model 1, using a Garch(1,1) process and a 1997-2007 history of spot and futures prices, yields the parameters in Table 3.

6 Numerical results

In this section we use our futures-spot modeling to value a storage contract, and compare it with the intrinsic valuation method. For this we will concentrate on fast and slow storages, which constitute two realistic cases. A fast gas storage has high injection/withdrawal rates, so that it can be filled in general within a month, but it has low gas capacity: salt caverns are a common example of high deliverability storage units. Slow gas storages are in general large depleted oil/gas fields, or aquifers, so they have very large gas capacities, but they suffer from low injection/withdrawal rates (see Appendix A for more details).

We will consider a fast and a slow storage unit with the characteristics described in Table 4, where for simplicity, all the quantities are expressed in 10^6 MMBtu², while the storage values are expressed in \$ million. This means that fast storage takes 25 days to fill, and almost 17 days to empty, while slow storage needs 125 days to be completely filled and 83 days to be completely emptied. We also consider null injection/withdrawal costs.

	Fast storage	Slow storage
Total capacity	100	100
Injection rate	4 per day	0.8 per day
Withdrawal rate	6 per day	1.2 per day
Initial gas volume	0	0
Final gas volume	0	0
Lease duration	1 year	1 year

Table 4: Gas storage characteristics (fast and slow units)

The experiments were run using the Matlab software, with 5000 simulations for the Monte Carlo method, and we used independent paths for the backward and forward phases, in the Longstaff & Schwartz algorithm (see Section 3.2). First, we simulate a set of spot and futures paths, then we apply the dynamic programming algorithm (3.8) to estimate the optimal spot strategy; in parallel we evaluate the hedging strategy, based on futures contracts, according either to (3.10) or (3.11). We then re-simulate a new set of spot and futures paths, independent from the paths used in the preceding backward phase, and we apply the estimated optimal spot strategy, combined with the futures hedging strategy, to the new trajectories.

²This energy unit can be naturally converted into a volume, under standard conditions for temperature and pressure.

We store the cumulative cash flows $\text{Wealth}_{\text{spot}+\text{futures}}(u^*)$ resulting from these physical and financial operations for each sample path, and we compute the empirical mean and standard deviations of those cash flows. The mean of the cumulative wealth gives an estimate of the extrinsic value J^* of the gas storage, given in (3.7), while the standard deviation is an indicator of the dispersion of the realized cash flows around the extrinsic value. We emphasize that the empirical mean estimates the cash flow generated by the optimal strategy, while the empirical standard deviation gives an indicator of the variance reduction obtained through the financial hedging strategy. A lower standard deviation means that the manager will face less uncertainty on a single realization of spot and futures prices. Numerical results will confirm that the hedging strategy indeed allows for a significant variance reduction of the cumulative cash flows. An example of the outputs of the above valuation procedure is represented in Figure 4 (fast storage) and 5 (slow storage) by samples of simulated spot trajectories, with the corresponding optimal gas volumes in the unit, for a contract starting in April 2007. Note that different colors correspond to different simulated spot trajectories.

Note also that the analysis described above depends on the choice of the model, because the backward and the forward phases are executed on the sample paths generated by the model itself. In order to make the comparison less model-dependent, we calculate the cumulative cash flows of the estimated optimal strategy, based on spot and futures *historical paths*. For this reason, we will consider a series of spot and futures curve data from 2003 to 2012, and split it into periods of one year: the storage lease contracts specified in Table 4 start in April each year, for a one-year period. We run the optimal strategy obtained in the backward phase (for the corresponding storage duration) on the spot and futures historical paths for the related period. This constitutes a real case test for the optimal strategy and corroborates the relevance of the spot modeling, since it provides the profit that would have been accumulated by the storage manager in a realized path. Figures 6 and 7 represent the historical spot path realized during the contract period (for both slow and fast units) from April 2007 to April 2008, and the natural gas volumes resulting from the optimal strategy computed on simulated paths (see Figures 4 and 5 for examples of these simulated paths).

We summarize the results of the valuation algorithm for each period in Tables 5 and 6 in fast and slow storage cases, when the spot paths are generated according to the spot model 2. proposed in (5.21). The tables report, for each period, the intrinsic value (IV) (see (O_t) in Appendix B) and the estimate of the extrinsic value (EV) on simulated paths and historical paths realized during the current period. The last two columns show the standard deviation of the simulated cash flows under the optimal strategy.

We expect that the extrinsic spot-based strategy will give a larger value than the intrinsic physical futures-based strategy, while our financial hedging strategy is supposed to reduce the uncertainty of gas storage cash flows. For example, the fast storage contract starting

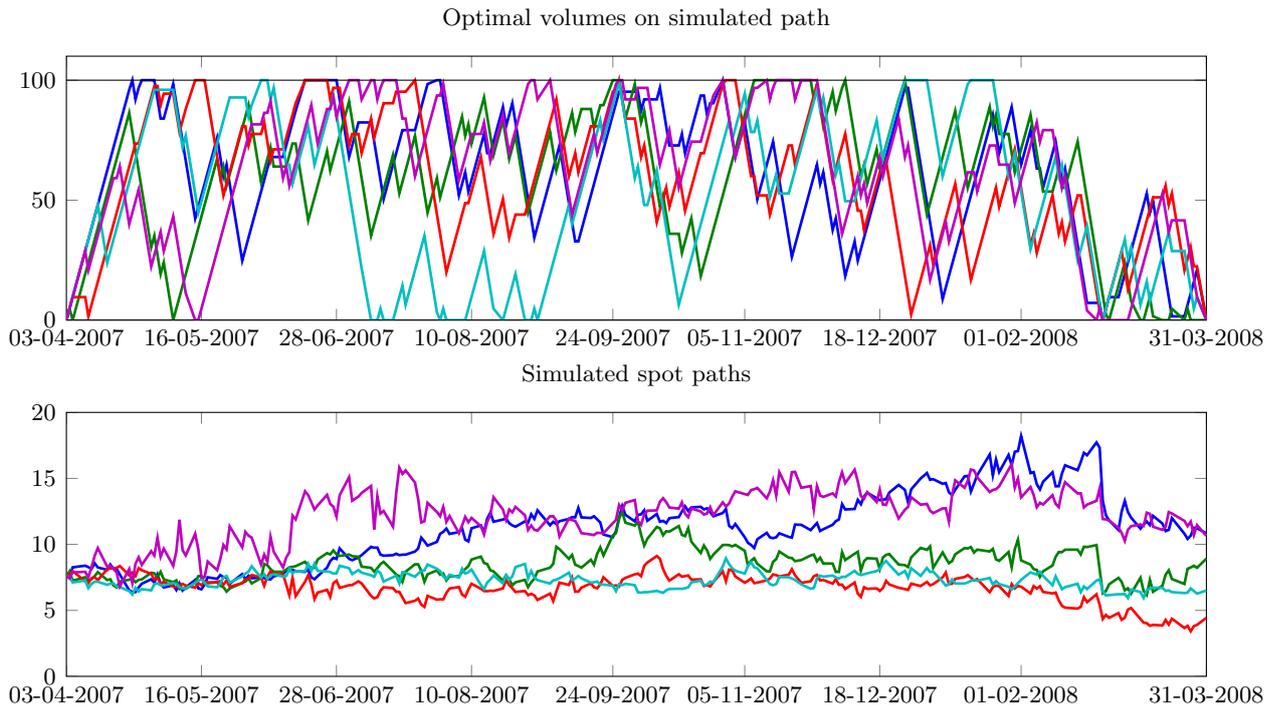


Figure 4: Simulated spot paths and optimal volumes (fast storage)

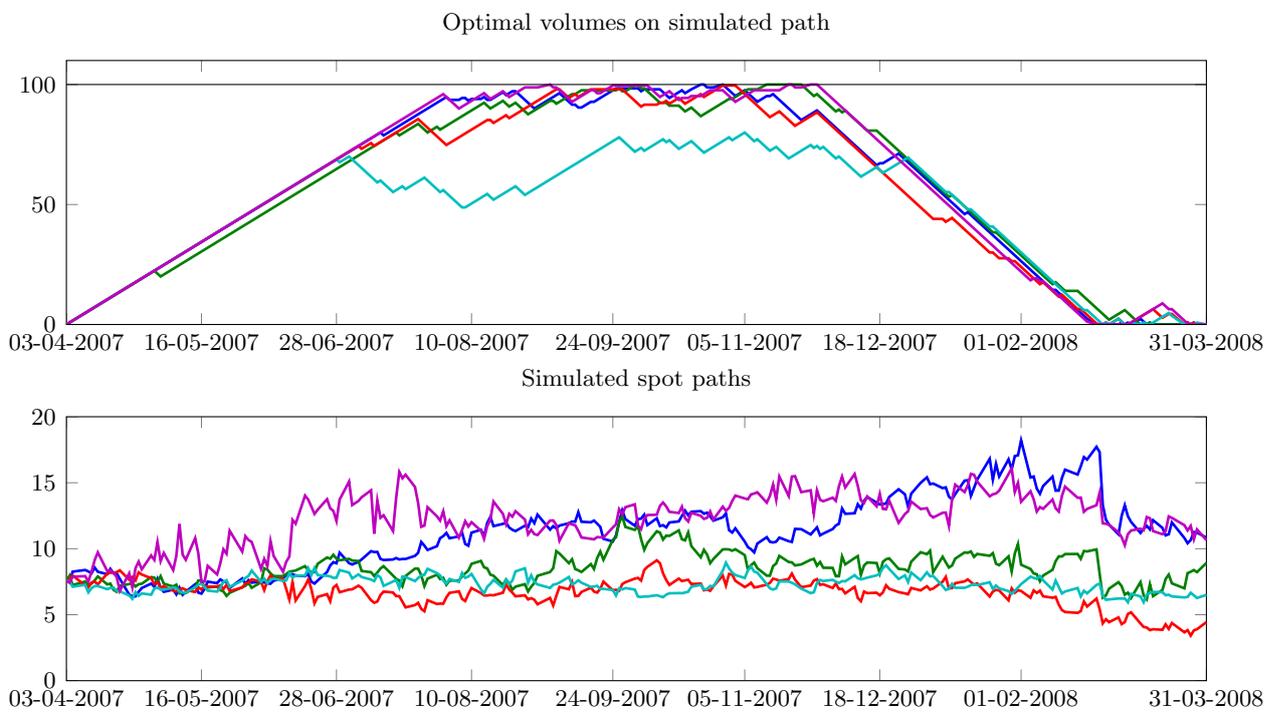


Figure 5: Simulated spot paths and optimal volumes (slow storage)

in April 2007 has an intrinsic value of $\$222.9689 \cdot 10^6$ while the spot-based strategy gives an extrinsic value of $\$697.0003 \cdot 10^6$. As expected, the extrinsic strategy allows better financial exploitation of the rights (without obligation) of injection/withdrawal natural gas compared to the conservative intrinsic strategy. In other words, the extrinsic strategy allows better extraction of the optionality of storage. We also note that the hedging strategy yields a significant empirical variance reduction of the cumulative cash flows from $\$340.2193 \cdot 10^6$ to $\$190.8546 \cdot 10^6$. On the other hand, the intrinsic value of slow storage is equal to $\$195.5517 \cdot 10^6$, while the spot-based strategy captures a larger optionality value of $\$251.0064 \cdot 10^6$. Similarly to fast storage, the financial hedging strategy allows an important variance reduction from $\$232.7825 \cdot 10^6$ to $\$28.0414 \cdot 10^6$.

Previous observations about the 2007 contract remain valid for the other test periods; indeed the intrinsic futures strategy is always out-performed by the extrinsic spot-based strategy, in both simulated and historical paths. The historical backtesting over the period 2003-2012 shows that the extrinsic strategy allows for better extraction of storage unit optionality, with a ratio of extrinsic value to intrinsic value as high as 500% for a fast storage unit. This performance of the extrinsic strategy is less significant in the case of slow storage unit, with a ratio up to 100%. This is due to limitations in the deliverability of slow storage, since the storage manager is not able to profit completely from the gas price volatility and cannot respond rapidly to favorable price movements. On the other hand, hedging with financial instruments provides a significant reduction of the cumulative cash flows uncertainty. In fact, the last two columns of Tables 5 and 6 show a standard deviation reduction factor of up to 10, with better performance for slow storage. This gives the storage manager more insurance to recover a large percentage of the value of the storage contract.

Remark 2. *1. In Section 3.3, we presented two heuristic hedging strategies, (3.10) and (3.11), based on financial futures contracts. The numerical tests that we have conducted show that the hedging strategy defined by (3.11) gives better results in the variance reduction of the simulated cash flows under the optimal strategy; in addition, in the historical backtesting, (3.11) renders a better cumulative wealth performance than (3.10). We emphasize that we have only reported about the better performing hedging strategy (3.11).*

2. We also note that the historical intrinsic value of the gas storage attains a peak in 2006, and shows a clear decreasing effect afterwards. This can be intuitively explained observing the futures curve samples in Figure 1: in 2006, the seasonal spreads were very pronounced, while they were quite small in 2011.

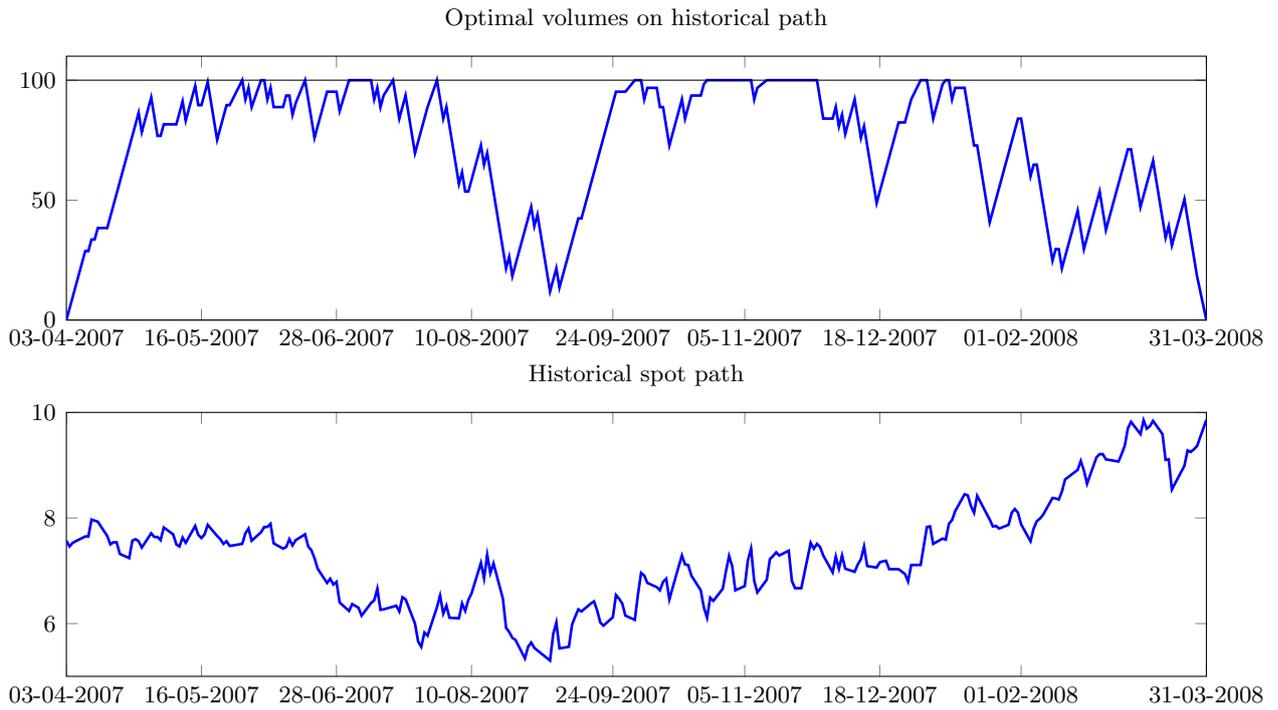


Figure 6: Historical spot path and optimal volumes (fast storage)

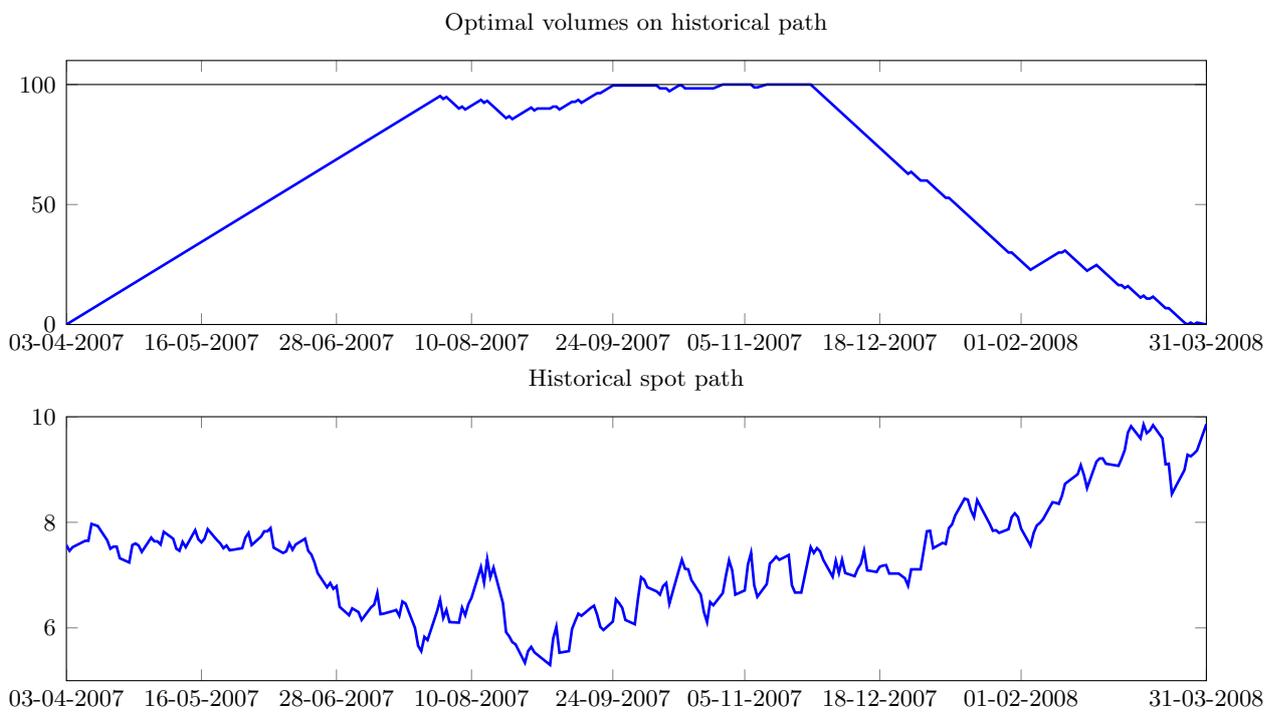


Figure 7: Historical spot path and optimal volumes (slow storage)

	Simulated paths test		Historical path test		Standard deviation	
Starting Date	IV	EV	IV	EV	Without hedge	With hedge
2003-Apr	39.9542	337.7276	42.6441	184.1178	189.5820	119.3606
2004-Apr	63.0335	395.6198	63.6000	347.2736	213.1796	126.3763
2005-Apr	115.2008	592.0854	112.0473	528.6510	306.3232	179.4792
2006-Apr	371.1724	860.9714	416.3992	616.2357	390.0864	194.9693
2007-Apr	222.9689	697.0003	241.8000	399.7347	340.2193	190.8546
2008-Apr	119.5200	674.6745	129.6000	427.9652	359.7650	210.6817
2009-Apr	204.6539	459.5531	205.9000	302.6958	203.2847	100.9753
2010-Apr	144.1958	420.2250	153.7000	259.1776	202.1802	112.6989
2011-Apr	86.5488	352.7785	92.9000	134.7794	190.0749	102.9312
2012-Apr	125.8968	272.5376	130.2000	215.9591	118.0606	55.4645

Table 5: Fast gas storage valuation (under spot model 2 (5.21))

	Simulated paths test		Historical paths test		Standard deviation	
Starting Date	IV	EV	IV	EV	Without hedge	With hedge
2003-Apr	24.6556	67.5795	26.0563	16.7890	83.9382	18.5218
2004-Apr	45.2183	91.1053	44.6833	53.8389	119.8064	20.1098
2005-Apr	93.6136	157.9486	92.2304	146.3097	189.6219	28.0376
2006-Apr	333.1988	386.4656	333.1972	356.0564	282.2993	29.6749
2007-Apr	195.5517	251.0064	195.3024	221.4466	232.7825	28.0414
2008-Apr	96.8824	169.6740	98.5936	141.5038	216.1477	32.6633
2009-Apr	180.5010	206.6439	180.4980	210.8117	148.3145	14.6222
2010-Apr	122.4013	152.9924	122.3784	128.1330	140.5389	16.5936
2011-Apr	68.5264	104.2509	68.1356	72.3083	118.2672	18.8294
2012-Apr	107.4493	122.5703	107.3928	110.0214	86.2897	8.4167

Table 6: Slow gas storage valuation (under spot model 2 (5.21))

We conclude from the numerical results presented above that the joint modeling of the natural gas spot price and futures curve is a pertinent framework for the gas storage valuation and hedging problem. It allows the unit manager to better exploit storage optionality by monetizing the spot price volatility and seasonality. Indeed, the historical backtesting shows that the extrinsic value under this modeling always outperforms the classical intrinsic value, even in the case of slow storage. A joint model for the futures curve with its own risk factors is a more realistic framework for spot and futures markets, since it takes into account the seasonality of the futures curve and the non-convergence of the futures price to the spot price, an unrealistic hypothesis that is often made in the literature. This also allows for a more relevant hedging strategy based on futures contracts, and better tracking of the extrinsic value of gas storage in real market conditions.

7 Model risk

As we showed in the introduction, seasonal spreads have become narrower these last years, which leads to a concentration of almost all the value of gas storage in the extrinsic part, based on spot trading. Hence it is very important to look closely into the spot modeling and its effect on storage valuation and hedging. We believe that the uncertainty of storage value is due more to the uncertainty of the spot modeling than the futures modeling, since only spot evolution determines the optimal strategy even though the futures contract prices intervene in spot modeling, see (5.19) and (5.21). Indeed their main contribution is devoted to variance reduction. In Section 7.1 we study the effects of the modeling hypotheses, and sensitivity with respect to the model parameters; in Section 7.2, we define a model risk measure to quantify this dependence, as proposed by Cont [2006]. Before this, in Section 7.1, we compare the performance of the two spot models proposed in Section 5.2, using historical data.

7.1 Spot modeling

In Section 5.2, we proposed two discrete models for the spot price dynamics. The first model, defined in (5.19), is a discrete version of a mean-reverting model, with a stochastic mean-reversion level equal to the prompt price. The second model, in (5.21), handles directly the spread between spot and prompt prices, which could be a decisive quantity in the optimal strategy. In fact, the unit manager will probably tend to buy and store gas if the spot-prompt spread is negative and withdraw and sell gas in the opposite case. Since the seasonality of gas prices has been getting weaker in recent years, the principal source of value for the storage unit is the spot-prompt spread rather than the winter-summer spreads, so the second model (5.21) is expected to give good results for storage valuation and hedging.

We run the valuation procedure explained in Section 6 for the two spot models, during the testing periods between 2003 and 2012, and we describe the performance of both models through historical spot paths: in particular, we report in Figures 8 and 9 the cumulative cash flows using the optimal spot strategy for historical spot trajectories. In the fast storage case, Figure 8 shows that spot-prompt spread model 2 yields slightly better results than spot model 1 in all the test cases, except for the year 2004. In the slow storage case, see Figure 9, the two spot models give comparable results for all periods. In the fast storage case, other tests show that in spot model 2, the cumulative cash flows generate a lower standard deviation than in spot model 1, which is in agreement with its better performance.

7.1.1 Effect of spikes modeling

The presence of spikes in natural gas prices is an essential feature of the dynamics of spot prices. In fact, as we noted in Section 2, these jumps are sudden dislocations of prices between

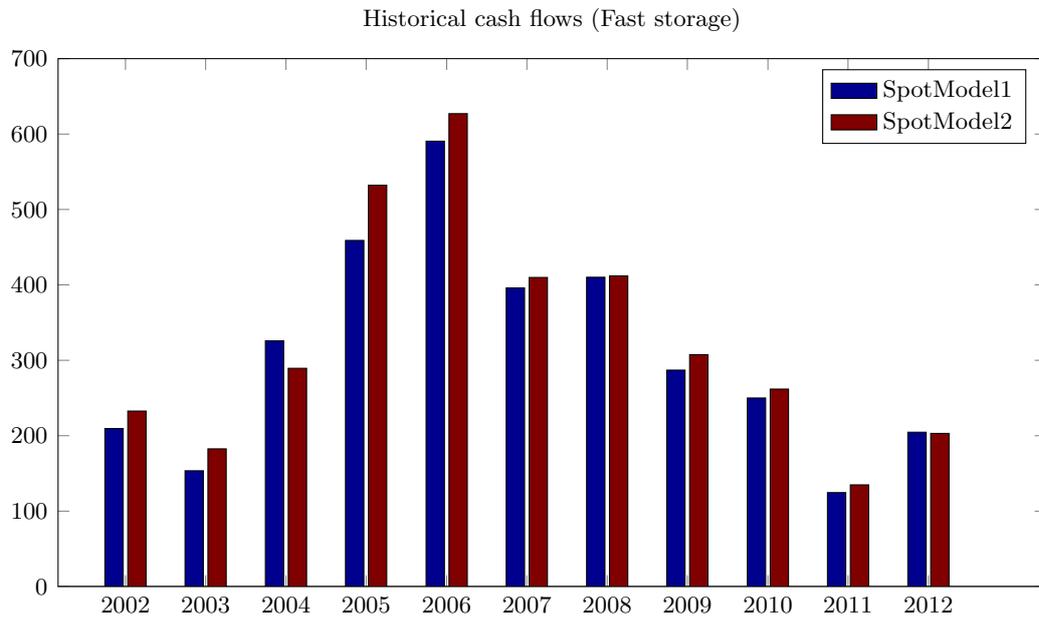


Figure 8: Historical cash flows for spot models 1 and 2 (fast storage)

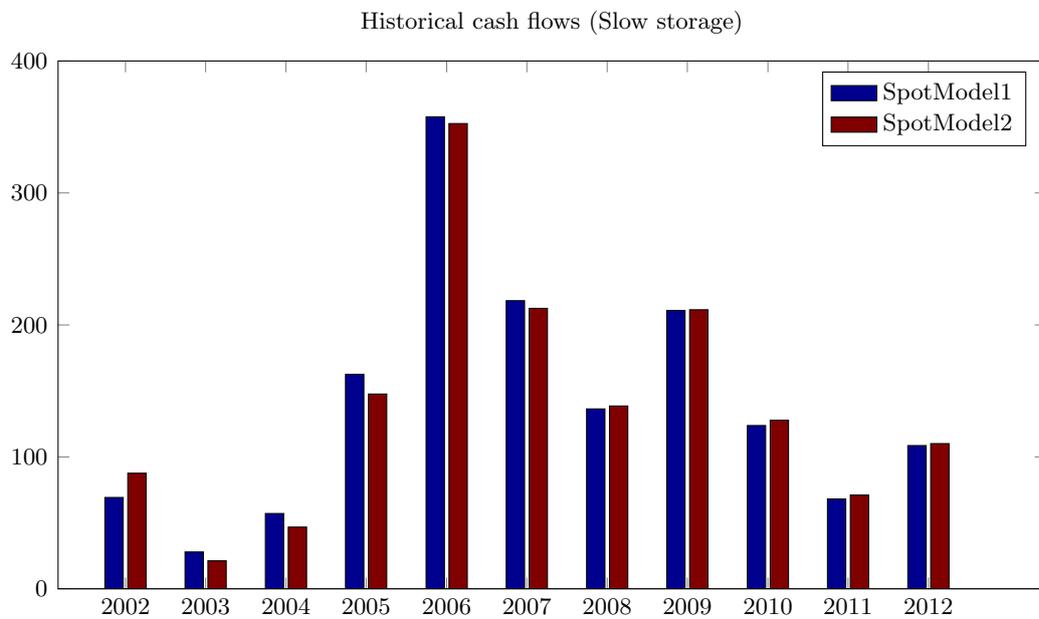


Figure 9: Historical cash flows for spot models 1 and 2 (slow storage)

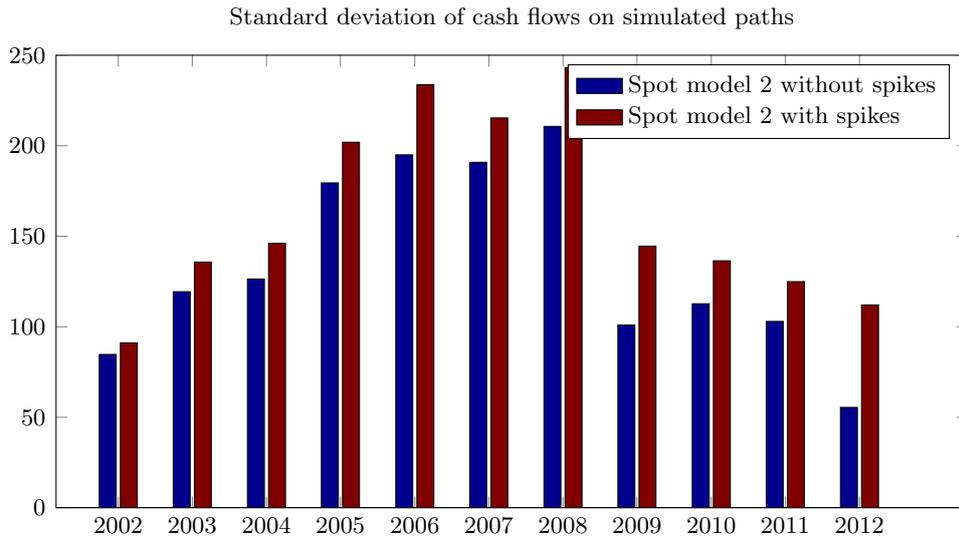


Figure 10: Standard deviation of cash flows on simulated paths

spot and prompt contracts, due to unpredicted weather changes, technical problems in the transport chain or simply to poor anticipation of the storage capacities of the market.

These spikes can be a source of value for the storage manager, since a large gap between spot and prompt prices can be monetized by buying gas during a negative spike, and selling gas during a positive spike. Since these are rapidly absorbed by the market, the value that can be captured from them strongly depends on storage characteristics. In fact, numerical tests show that with slow storage, the spike modeling has less effect on the extrinsic value, compared to fast storage. For this reason we only concentrate on the latter.

Figure 11 represents the expected cumulative cash flows, on simulated paths under the spot model (5.21), for fast storage. All the test periods show that modeling the spikes in the spot dynamics gives a larger extrinsic value for the storage unit, but at the same time it introduces a larger standard deviation for the cumulative cash flows, as illustrated in Figure 10. This foresees a more significant uncertainty for the cash flows on a single sample path, when the spikes are taken into account in the spot modeling.

A final test of the effect of the spikes modeling is performed on historical spot paths for each test period, and results are shown in Figure 12. In fact, according to this graphic, it seems that modeling the spikes does not make a relevant contribution. This accords with the fact that the models with spikes produce a large standard deviation. Finally this historical back testing does not show significant advantages of spike modeling.

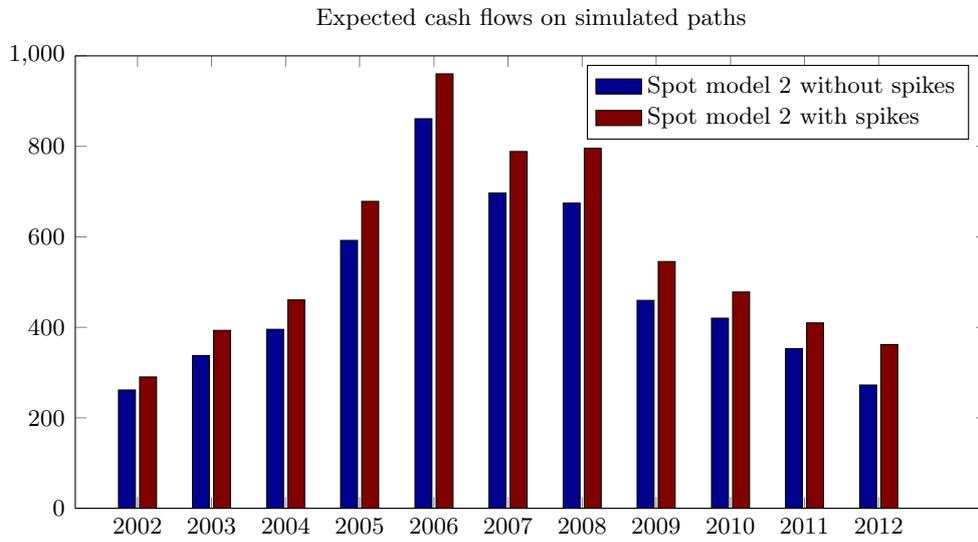


Figure 11: Expected cash flows on simulated paths

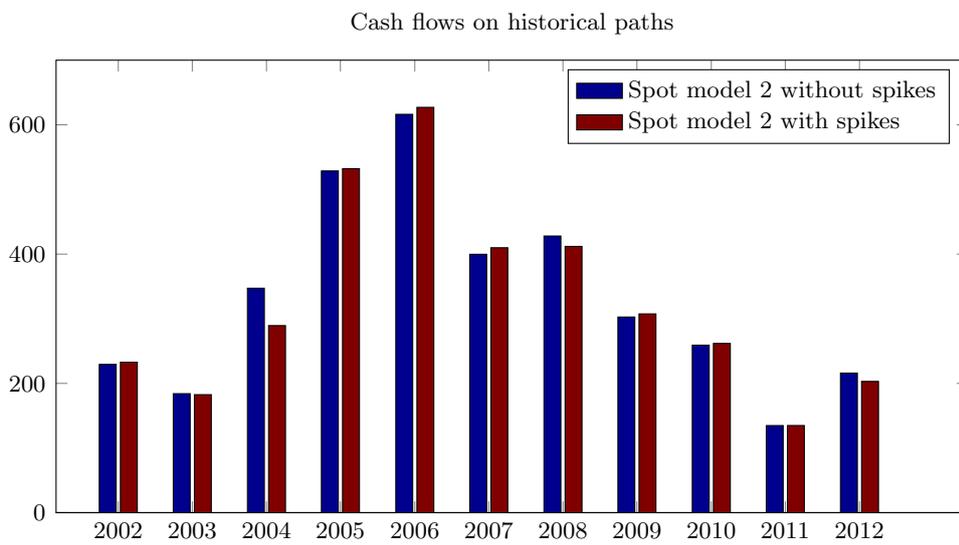


Figure 12: Cash flows on historical paths

7.2 Model risk measure

In order to quantify the modeling uncertainty, as anticipated, we will consider an approach introduced by Cont [2006], a method that was proposed for the uncertainty of stock models in view of pricing exotic derivative products. In that case, the market data are a set of vanilla option prices (or bid/ask intervals). Then the model uncertainty for an exotic payoff H , is quantified by computing the range of prices of this exotic product, using a set of risk neutral models Γ that calibrate the reference vanilla prices, i.e.

$$\pi(H) = \max_{Q \in \Gamma} \mathbb{E}^Q[H] - \min_{Q \in \Gamma} \mathbb{E}^Q[H]. \quad (7.26)$$

For our gas storage valuation problem, we will adapt this risk measure, by using as “calibration” data, the historical prices of the futures and spot contracts. The constraint of calibration on vanilla prices is replaced by the success of suitable statistical tests and closeness to the optimal likelihood objective function value of the model. Indeed, in our case, the family Γ consists of a set of spot models, (5.19) or (5.21), which pass statistical tests imposed by the modeling hypothesis for the noise (ϵ_t) to be of type Garch(1,1), and have a likelihood function value close to the optimal one found during the model estimation. This methodology for the generation of the models set Γ is partially similar to the one proposed in the case of multi-asset options by Dumont and Lunven [2006]. In this study, the authors calibrate a multi-assets model to single-asset vanilla options, then build the set Γ by perturbation of the correlation matrix. This yields a family of models that fit perfectly the reference vanilla options, but differ by their correlation matrix.

In our case of gas storage valuation, the statistical estimation of the spot model parameters, in (5.19) or (5.21), is realized by classical maximum likelihood methods. The estimation procedure is a maximization problem

$$\max_{\theta = \{a_1, a_2, a_3, \kappa, \gamma_1, \alpha_1\}} L(\theta),$$

where $L(\theta)$ is the likelihood function associated with the spot model (5.19) or (5.21). This maximization yields an optimal parameters vector $\theta^* = \{a_1^*, a_2^*, a_3^*, \kappa^*, \gamma_1^*, \alpha_1^*\}$, an optimal likelihood function value $L(\theta^*)$, and an empirical variance-covariance matrix Σ^* associated with model parameters, which specifies the confidence interval for the estimated parameters, up to a confidence level.

In order to generate the spot models family, we perturb the optimal parameters θ^* by adding a Gaussian noise with the specified covariance matrix Σ^* to θ^* . This yields a set of perturbed parameters $\{\theta_i\}_{i \in I}$, from which we retain only the perturbed models that satisfy two constraints: first, the inferred Garch white noise $z(\theta_i)$ in (5.20) has to pass a statistical test of normality³; second, the corresponding likelihood function value $L(\theta_i)$ has to be close

³We use a Kolmogorov-Smirnov test for the normality test of the inferred noise z .

to the optimal value $L(\theta^*)$, by a small factor ϵ i.e. $L(\theta_i) > (1 - \epsilon)L(\theta^*)$.

In our framework, Γ will be the set of $\theta_i, i \in I$, fulfilling the two conditions above.

After having constructed the models set Γ , we can now define the associated model risk. In our storage valuation problem, the analogue risk measure to (7.26) can be expressed using the value function $J^*(\theta)$ in (3.7), where we emphasize the dependence of this value function with respect to the spot model parameters θ , and we express the risk measure in relative terms. We set

$$\pi_1 = \frac{\max_{\theta_i \in \Gamma} J^*(\theta_i) - \min_{\theta_i \in \Gamma} J^*(\theta_i)}{J^*(\theta^*)}. \quad (7.27)$$

In this risk measure evaluation, each $J^*(\theta_i)$, is calculated using spot and futures paths simulated under the perturbed model θ_i .

Moreover, we propose a second model risk measure based on the performance on realized historical spot and futures paths. For this we define

$$\pi_2 = \frac{\max_{\theta_i \in \Gamma} \text{Wealth}_{\text{spot+futures}}(\theta_i) - \min_{\theta_i \in \Gamma} \text{Wealth}_{\text{spot+futures}}(\theta_i)}{\text{Wealth}_{\text{spot+futures}}(\theta^*)}, \quad (7.28)$$

where $\text{Wealth}_{\text{spot+futures}}$ represents the cumulative cash flows, computed on the historical path, as defined in (3.9).

The two risk measures π_1 and π_2 are computed for each of the test periods from 2003 to 2012, under the two spot models 1 and 2, using a set of 30 perturbed models. The results reported in Table 7 again show a better performance for spot model 2. In fact, this model seems to be less subject to model risk, since it gives a smaller range of prices, compared to spot model 1.

Starting date	Risk measure π_1		Risk measure π_2	
	Spot model 1	Spot model 2	Spot model 1	Spot model 2
2003-Apr	51.33 %	44.8085 %	70.8465 %	39.3852 %
2004-Apr	25.4987 %	23.6942 %	26.5597 %	22.3195 %
2005-Apr	26.0388 %	27.0318 %	50.7306 %	38.352 %
2006-Apr	14.9666 %	15.9873 %	10.6853 %	6.6954 %
2007-Apr	93.8336 %	14.7645 %	29.4626 %	18.6143 %
2008-Apr	37.9839 %	13.8195 %	16.6811 %	8.6166 %
2009-Apr	20.7969 %	10.1216 %	15.1415 %	8.1936 %
2010-Apr	26.7845 %	12.8976 %	33.0669 %	7.5285 %
2011-Apr	25.9442 %	12.3857 %	35.8704 %	30.9282 %
2012-Apr	16.7783 %	9.1489 %	13.1014 %	7.1694 %

Table 7: Model risk measure for spot models 1 and 2.

One observation that follows clearly from Table 7 is that the range of prices induced by the model uncertainty and measured by π_1 and π_2 represents a large proportion of the storage value. This shows that the dependence of gas storage valuation on spot modeling is quite significant. While the literature has concentrated its efforts until now on the specification of an optimal valuation strategy, we believe that one should pay more attention to the choice of spot-futures modeling framework. A second comment that we can infer from Table 7 is that model 2 appears to be less sensitive to the change of parameters and is therefore more robust. Fortunately, this is in concordance with the better performance of spot model 2 already observed in Section 7.1. Table 7 shows that the spot-futures valuation framework is subject to a large model risk (average: 25%). For comparison, the model risk for a basket option has been evaluated to 3% (see Dumont and Lunven [2006]).

8 Conclusion

In this paper we consider the problem of gas storage valuation. After restating the main stylized facts of natural gas prices, specifically seasonality and spikes, we present a joint modeling framework for the futures curve and the spot, with two different spot models. Using a Monte Carlo simulation method, we estimate the extrinsic optimal spot strategy; for the purpose of variance reduction of the cumulative cash flow, we set up a financial hedging strategy. We also conduct back testing using historical data of futures and spot prices over a period of 10 years. This demonstrates the better performance of the extrinsic strategy compared to the classic intrinsic futures-based strategy. In fact, the spot strategy allows the manager to better track the value of gas storage, in real market conditions, even in the case of a slow storage unit.

In the final section, we study the model uncertainty and its effect on storage value, concentrating on the risk associated with spot modeling. After a quantitative comparison of the two spot models we proposed, we conclude that the model based on the spot-prompt spread performs better. In order to quantify the stability of these results with respect to model uncertainty, we define two model risk measures, inspired by the work of Cont [2006], but based on historical prices. Using those risk measures, we observe the great sensitivity of gas storage value with respect to the modeling assumptions. In fact the model uncertainty, as measured by the size of price range, represents a large proportion of storage value. This puts into perspective the concentration of effort in the literature on the specification of an optimal valuation strategy. In fact, much more attention should be probably devoted to the discussion of modeling assumptions.

Appendix A Different types of gas storage facilities

Natural gas storage units are underground facilities, so their characteristics depend essentially on the geological properties of the storage area. There are three types of gas storage units: depleted gas/oil fields, aquifers and salt caverns. The main characteristics that distinguish these gas storage units are their injection/withdrawal rates, the total capacity and the so-called cushion volume and working volume. The cushion volume is the quantity of gas that must remain in the storage unit to provide the required pressurization, and the working volume is the volume of gas that can be extracted. Using the notations in this article, the cushion volume corresponds to the minimum volume V_{min} , the total capacity corresponds to V_{max} , and the working volume is represented by the actual volume minus the cushion volume, i.e. $V_t - V_{min}$.

These characteristics distinguish two different types of gas storage: base-load and peak-load. Base-load units are used to meet seasonal demand (a more or less predictable phenomenon). In fact the demand for gas is highly concentrated in the winter season, so in order to ensure sufficient supply, gas is bought and stored in the summer season then withdrawn and sold in winter. The main characteristics of base-load units are their large volume capacity and low deliverability rates.

On the other hand, peak-load units are used to mitigate the risk of unpredictable increases in the gas demand, generally caused by weather changes or technical problems in the pipeline system. Hence, they have to be very reactive and have high deliverability rates, higher injection/withdrawal rates, and in general they contain less gas than base-load units. While the filling/withdrawal cycle duration of a base-load is in general one year, peak-loads can have a turn-over period of a few weeks.

The depleted gas/oil fields and aquifers are of the base-load type, while salt caverns are peak-load facility. We summarize here their main characteristics.

- **Depleted gas and oil fields.** They are the most commonly used underground storage sites because of their wide availability. Besides their large capacities, they benefit from the already available wells and injection/withdrawal equipments, pipelines etc. Their main drawbacks are low deliverability rate and the large *cushion* gas percentage (although part of this non-usable gas already exists in the geological formation). Therefore, these depleted fields naturally belong to the base-load category.
- **Aquifers.** These are underground, porous and permeable rock formations that act as natural water reservoirs.

They are flexible units with small volume, but more expensive than depleted fields since everything has to be built from scratch (wells, extraction equipments, pipelines, etc). In some instances, the installment of aquifer infrastructure can take four years, which is more than twice the time needed to transform depleted reservoirs into storage facilities.

On the other hand, aquifers require a greater percentage of *cushion* gas than depleted reservoirs: up to 80% of the total gas volume.

Like depleted fields, aquifers operate on a single annual cycle, so they still belong to the base-load category.

- **Salt cavern.** Salt caverns are the third common choice for gas storage. They are created by dissolving and extracting a certain amount of salt from the geological formation; this process then leaves a cavern that can be used for natural gas storage.

A salt cavern offers storage with high deliverability, with low *cushion* gas requirements (30% of *cushion* gas), but with lower capacities than depleted fields and aquifers. They cannot be used to meet base-load storage requirements, but they are well suited to rapid actions, which are distinctive features of the peak-load category.

Table 8, compiled by the Federal Energy Regulatory Commission (FERC),⁴ summarizes the three types of storage and their characteristics.

Type	Cushion to working gas ratio	Injection period (days)	Withdrawal period (days)
Aquifer	Cushion 50% to 80%	200 to 250	100 to 150
Depleted oil/gas reservoirs	Cushion 50%	200 to 250	100 to 150
Salt cavern	Cushion 20% to 30%	20 to 40	10 to 20

Table 8: Types of natural gas storage.

⁴ *Current State of and Issues Concerning Underground Natural Gas Storage*, Federal Energy Regulatory Commission (FERC), Staff Report, September 30, 2004. cf: <http://www.ferc.gov/EventCalendar/Files/20041020081349-final-gs-report.pdf>

Appendix B Futures-based valuation methods

Futures-based valuation is still very commonly used in natural gas storage management. This is mainly due to its simplicity and its low risk profile. It is based on trading natural gas futures with physical delivery, combined with gas injection/withdrawal. The idea is the following: at the beginning of the storage contract, the manager observes the initial futures curve, and decides to buy/sell multiple forward contracts, and consequently receives/delivers natural gas at their expiration, during a delivery period.⁵

In order to determine the optimal futures positions, a linear optimization problem has to be solved, with constraints imposed by the physical conditions of the storage unit, cf Eydeland and Krzysztof [2002]. Let us denote by $(F(t, T_j))_{j=1, \dots, N}$ the available futures contracts in the markets (the maturities T_j being generally spaced monthly), at date t , and $\alpha_j(t)$ the quantity of futures $F(t, T_j)$ bought/sold. Using the notations introduced in Section 3.1, at the initial date t_0 , the storage manager chooses futures positions $\alpha_j(t) \equiv \alpha_j(t_0)$ in order to maximize the profit, under the physical constraints fixed by the maximum V_{max} and the minimum V_{min} storage capacity and the injection a_{inj} and withdrawal a_{with} rates of the storage unit.

A natural optimization problem, at each time t , in the variables $\alpha_j(t)$, is the following:

$$\begin{aligned} IV(t) &:= \max_{(\alpha_j(t))_{j=1, \dots, N}} - \sum_j \alpha_j(t) F(t, T_j) \\ -a_{with} \Delta T_j &\leq \alpha_j(t) \leq a_{inj} \Delta T_j, \text{ for } j = 1, \dots, N \\ V_{min} &\leq V(t) + \sum_{j=1}^n \alpha_j(t) \leq V_{max}, \text{ for } n = 1, \dots, N, \end{aligned} \tag{O_t}$$

where $\Delta T_j = T_{j+1} - T_j$ is equal to the time length of the delivery period of the $F(t, T_j)$ futures and $V(t)$ is the gas volume in the storage at date t .

We emphasize that the storage manager, who follows the intrinsic value methodology, solves (O_t) only at time $t = t_0$. $IV(t_0)$ represents the optimal profit given by the maximization problem (O_t) ; it is generally called *intrinsic value*. Indeed the storage manager keeps the optimal futures positions $\alpha_j^*(t_0)$ for the whole storage contract duration, so this strategy does not take advantage of possible profitable movements of the futures curve.

This static methodology was extended by Gray and Khandelwal [2004] to the *rolling intrinsic valuation*, to take advantage of the changing dynamics of the futures curve. We consider a set of trading dates $t_0 < t_1 < \dots < t_{n-1} < t_n$, where t_n is the maturity of the

⁵For example, the Nymex NG futures have monthly spaced maturities, and the delivery period spreads out over the month following each maturity date.

storage contract, at which the manager can buy/sell physical futures. In the sequel, for $t = t_i, 1 \leq i \leq n - 1$, we set $\Delta t = t_{i+1} - t_i$.

At the beginning date t_0 of the storage contract, optimal futures positions are chosen, but after the time period Δt , one recalculates the new optimal futures positions: if the manager finds it more profitable, the portfolio is rebalanced. More precisely, suppose that at date t , the manager owns a futures portfolio $\alpha^*(t)$; then, at date $t + \Delta t$, the manager solves $O_{t+\Delta t}$, calculating an optimal portfolio $\alpha(t + \Delta t)^*$ and $IV(t + \Delta t)$. The profit/cost of rebalancing the futures portfolio from $\alpha(t)^*$ to $\alpha(t + \Delta t)^*$ would be equal to

$$C(t, \Delta t) := \sum_j [\alpha_j^*(t) - \alpha_j^*(t + \Delta t)] F(t + \Delta t, T_j),$$

so, the manager will switch positions only if this rebalancing profit/cost $C(t, \Delta t)$ is positive. We denote by $RI(t)$ the so-called rolling intrinsic value at time t . At time t_0 we set $RI(t) = IV(t)$; recursively we define $RI(t + \Delta t) := RI(t) + \max(C(t, \Delta t), 0)$; $RI(t)$ represents the cumulative profit generated by this enhanced strategy. Obviously, at each rebalancing date t , the rolling intrinsic value $RI(t)$ is always greater or equal than the intrinsic value $IV(t_0)$.

Remark 3. *The intrinsic and rolling intrinsic methodologies capture the predictable seasonal pattern of natural gas prices: they basically lead to buying cheap summer futures and selling expensive winter futures. Indeed the obtained storage value greatly depends on the seasonal spread between cold and warm periods of the year. The recent tightening of seasonal spreads (cf Section 2) implies that the (rolling) intrinsic value becomes less attractive for practitioners.*

ACKNOWLEDGEMENTS: The research of the second and third named authors was partially supported by the ANR Project MASTERIE 2010 BLAN-0121-01.

References

- P. Bjerksund, G. Stensland, and F. Vagstad. Gas storage valuation: Price modelling v. optimization methods. *The Energy Journal*, 32 (1):203–227, 2011.
- A. Boogert and C. De Jong. Gas storage valuation using a Monte Carlo method. *The Journal of Derivatives*, 15(3):81–98, 2008.
- L. Clewlow and C. Strickland. A multi-factor model for energy derivatives. *School of Finance and Economics, University of Technology, Sydney, Working Paper*, 1999a.
- L. Clewlow and C. Strickland. Valuing energy options in a one factor model fitted to energy prices. *Quantitative Finance Research Group, University of Technology, Sydney, Research*, 1999b.
- R. Cont. Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance*, 16(3):519–547, 2006.

- P. Dumont and C. Lunven. Risque de modèle et options multi sous-jacents. *Rapport de stage, Groupe de Travail ENSAE*, 2006.
- A. Eydeland and W. Krzysztof. *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. Wiley, 2002.
- FERC. Report of the natural gas price spike of february 2003. Technical report, Federal Energy Regulatory Commission, 2003. URL <http://www.ferc.gov/legal/maj-ord-reg/land-docs/The-Price-Spike-Report-07-23-03.pdf>.
- FERC. State of the markets, 2011. Technical report, Federal Energy Regulatory Commission, 2012. URL <https://www.ferc.gov/market-oversight/reports-analyses/st-mkt-ovr/som-rpt-2011.pdf>.
- J. Gabillon. The term structure of oil futures prices. *Oxford Institute of Energy Studies, Working Paper, No. M17*, 1991.
- F. Gray and K. Palamarchuk. Calibration of one and two-factor models for valuation of energy multi-asset derivative contracts. <http://arxiv.org>, 2010.
- J. Gray and P. Khandelwal. Towards a realistic gas storage model. *Commodities Now*, 2004.
- B. Hambly, S. Howison, and T. Kluge. Modelling spikes and pricing swing options in electricity markets. *Quantitative Finance*, 9(8):937–949, 2009.
- F.A. Longstaff and E.S. Schwartz. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1):113–47, 2001.
- E. Mastrangelo. An analysis of price volatility in natural gas markets. *U.S. Energy Information Administration*, 2007.
- C. Parsons. Quantifying natural gas storage optionality: a two-factor tree model. *Journal of Energy Markets*, 2013.
- X. Warin. Gas storage hedging. In René A. Carmona, Pierre Del Moral, Peng Hu, and Nadia Oudjane, editors, *Numerical Methods in Finance*, volume 12 of *Springer Proceedings in Mathematics*, pages 421–445. Springer Berlin Heidelberg, 2012. ISBN 978-3-642-25745-2. doi: 10.1007/978-3-642-25746-9_14. URL http://dx.doi.org/10.1007/978-3-642-25746-9_14.