



UvA-DARE (Digital Academic Repository)

Asset Pricing in Incomplete Markets: Valuing Gas Storage Capacity

Zhao, L.; van Wijnbergen, S.

DOI:

[10.2139/ssrn.2653647](https://doi.org/10.2139/ssrn.2653647)

[Link to publication](#)

Citation for published version (APA):

Zhao, L., & van Wijnbergen, S. (2015). *Asset Pricing in Incomplete Markets: Valuing Gas Storage Capacity*. (Tinbergen Institute Discussion Paper; No. 15-104/VI/DSF95). Amsterdam: Tinbergen Institute.
<https://doi.org/10.2139/ssrn.2653647>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.



Duisenberg school of finance – Tinbergen Institute Discussion Paper

TI 15–104 / VI / DSF 95

Asset Pricing in Incomplete Markets: Valuing Gas Storage Capacity

Lin Zhao

Sweder van Wijnbergen

Faculty of Economics and Business, University of Amsterdam, and Tinbergen Institute.

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Asset Pricing in Incomplete Markets: Valuing Gas Storage Capacity

August 28, 2015

Lin Zhao

University of Amsterdam and Tinbergen Institute

and

Sweder van Wijnbergen

University of Amsterdam and Tinbergen Institute

Abstract

We investigate the relationship between the gas spot market and the price of gas storage capacity. Contrary to the common belief, the auction prices for gas storage are mostly affected by the volatility of current market prices rather than by the winter-summer price differences. This paper provides a numerical solution for pricing storage capacity, by taking investor's activities through the spot market and storage service into account. A bivariate Generalized Autoregressive Score (GAS) model is employed for modeling the dynamics of the day-ahead and month-ahead spot market prices, as well as the time-varying volatilities and correlations. Under an incomplete market setting, our model is able to approximate the realized auction prices. Moreover, one interesting implication is that the implied average risk aversion of investor for a storage contract increases with the volatility of the spot market. This is an intuitive result because storage capacity can serve as an effective hedging product for the spot market, and the demand for this product is high when the market becomes risky: more risk averse investors are participating in the auctions. Moreover, a sensitivity analysis on different injection/withdrawal rates is also included, and particularly, contracts with higher capacity rates are priced at a higher level.

Key words: stochastic volatility, Generalised Autoregressive Score modeling, incomplete markets, real options, utility indifference pricing, gas storage, capacity constraints

JEL codes: C61, C63, G12, G13, Q41

1 Introduction

Gas storage allows the investors to inject a certain amount of gas when gas demand and spot prices are low (e.g. in summer) and to withdraw when market gas demand and spot prices are high (e.g. in winter). By purchasing the storage capacity, the investors gain the right to choose both the timing and the quantity of gas injections/withdrawals in accordance with the anticipated or realized demand changes. Therefore, theoretically, the value of the storage capacity consists of both intrinsic and extrinsic values: the former is intuitive, the price difference between winter and summer, and the latter values the flexibility of the storage opportunity. Many studies have been conducted in order to correctly model the two components. However, there has been a mismatch between the theory and practice of the capacity pricing. This paper tries to bridge that gap by on the one hand taking the theoretical problems seriously into account, but on the other hand modeling the problem with sufficient complexity to do justice to the real world structure of the contracts traded. In addition to constructing a satisfactory statistical model to solve the capacity pricing problem, we will in the process also shed light on investors' behavior by comparing the modeling results with actual auction outcomes. The aim of this paper is to value storage capacity and to investigate its interaction with the Dutch gas spot market, namely the Title Transfer Facility (TTF) market, by taking the complex option structure into consideration.

In the Netherlands, storage capacity is traditionally traded through semi-annual auctions. All eligible capacity buyers submit their bids including the price for one Standard Bundled Unit (SBU) and the amount of SBUs they require. After rounds of auctions, successful bidders will be offered the demanded storage correspondingly at their bidding prices. During the contract period, natural gas is then delivered at TTF market. Therefore, in essence, storage capacities serve as a bundle of options on the gas spot market, because it offers the capacity buyer the right to ask for gas delivery or storage at a predetermined cost.

Only European-style fixed maturity options on the TTF market are publicly traded. The lack of flexibility and low liquidity of that market are the main reason for the popularity of storage capacity contracts. Since the storage capacity contract is essentially an option-like contract on gas, its pricing can be determined through option pricing techniques. However, the classical pricing methods can not readily be applied. First, the standard option pricing assumptions (complete markets allowing for arbitrage-free price determination) do not hold due to the complicated dynamics of underlying gas price processes. Second, this real option problem has a more complex option structure than typically considered in the academic literature: the option to withdraw or to

inject gas can be exercised many times on the amount requested by the option holder; but the amount of withdrawal or injection is constrained by the decisions in earlier periods. Third, the Dutch TTF market was only initiated in 2003 and its corresponding option market has not become active until 2011. Thus the illiquid and thinly traded futures and options market leads to the failure of market completeness, which is a fundamental assumption required by the classical option pricing theory. As a result, option pricing methods based on no arbitrage conditions (risk neutral pricing) lose their applicability and we need an effective approach to deal with this issue.

Also, the method for valuing and estimating extrinsic values used in most of the existing literature appears to be too simplified, not taking the investors' active hedging behavior over spot market into consideration. For instance, as one can see from historical data, if an investor observes a stable TTF market for a long period, she is apparently less interested in bidding for storage capacity, or in other words, she values capacity less. This characteristic feature of capacity pricing is missing from most of current literature and industry applications. Examples are the relatively regular failures of capacity auctions held by GasTerra in 2011 (both March 28, 2011 and November 29, 2011) and in early 2013. Further studies therefore are needed to investigate the investors' bidding and pricing behavior.

Another issue concerns the impact of volatility on the valuation of the storage capacity, because of the option feature of storage capacity contracts. So far, only scenario analysis has been performed and results are usually given separately for high or low but always constant volatilities. In this paper, we allow for a more complicated model structure by taking into account the time-varying dynamics of volatilities. We apply both multivariate GARCH models and GAS (Generalized autoregressive score) models, so as to capture the changing of volatility of gas spot prices. Introducing non-trivial volatility risk factors once again leads to a violation of the assumptions needed for the applicability of risk neutral pricing methods, so we use a full fledged stochastic dynamic optimization approach to embed optimal investor behavior into the pricing analysis. We use Least Square Monte Carlo (Longstaff and Schwartz [2001]) to reduce the dimensionality problem endemic in Dynamic Programming problems down to manageable proportions, but in a different manner than Boogert and De Jong [2008] have proposed. Details will be discussed in Section 5. Furthermore, the invalidity of preference-free pricing in our setting makes it necessary to explicitly consider the impact of the degree of risk aversion of investors (capacity buyers, in our case) on valuation. This has the added benefit of allowing us to

infer the implied risk aversion of capacity buyers by observing realized auction prices.

Our methodology is able to produce valuations for storage capacity that approximate the realized auction prices quite well. Our findings suggest that the investor assigns a higher value to the storage capacity when facing a riskier market. This can be explained by the fact that in a more volatile market, there is more demand for an effective hedging product. The revealed risk aversion of the investor is far from zero, which again suggests the importance of parametrizing the risk aversion of individual investor. Also, we compare the values of several possible storage contracts with various injection/withdrawal capacity constraints. In accordance with intuition, we find that storage with higher injection/withdrawal rate (faster storage) results in higher contract values.

This paper is organized as follows. Section 2 reviews studies that have been conducted on pricing storage capacity and on GAS models. An extensive analysis of the spot market and historical auction prices can be found in Section 3 and Section 4. Section 5 discusses the setup of the model for one unit of storage (one SBU). Section 6 deals with the statistical models for gas prices, while Section 7 discusses the option pricing implications. Section 8 concludes.

2 Literature Review

2.1 Pricing Storage Capacity

De Jong [2015] discusses and compares four commonly used valuation approaches to valuing gas storage, namely, intrinsic, rolling intrinsic, basket of spreads, and spot trading approaches. These methods are not mutually exclusive, they all share certain similarities. Intrinsic and rolling intrinsic methods focus mainly on optimal trading in forward markets, with the latter allowing for recalculating optimal strategies over time and for adjusting the trading strategy as new information flows in. However, as indicated in De Jong [2015], applications of these two methods overlook the fact that the required extensive forward contracts are often not traded in the forward markets, resulting in an upwards bias in estimated storage values. The basket of spreads method can be seen as a simplified rolling intrinsic approach. The number of feasible spreads is limited due to the restrictions embedded in a storage capacity problem. This results in a possible downward bias of this approach, in contrast to the intrinsic and rolling intrinsic approaches. Finally, the spot trading method

includes the hedging behaviors of spot trading on a frequent basis. Secomandi and Seppi [2014] use a different classification based on relevant underlying assets, including spot price models (e.g. Boogert and De Jong [2008]), futures term structure models (e.g. Lai et al. [2010]), and equilibrium asset pricing models. This is a more general classification than the one given by De Jong [2015] since they included the equilibrium pricing models. The reason of the unpopularity of the equilibrium asset pricing models among practitioners lies in their low tractability. These classifications highlight the two main aspects of modeling and pricing storage capacity: the dynamics of underlying asset prices and the option pricing techniques used. We discuss the literature on both aspects separately.

Underlying processes Jaillet et al. [2004], Secomandi [2010] and Boogert and De Jong [2008] all assume that the underlying asset follows a mean-reversion process with constant volatility. Moreover, Thompson et al. [2009] add jumps to constant volatility models. Similarly, Carmona and Ludkovski [2008] study the optimal switching regimes between electricity and gas markets by employing an Ornstein-Uhlenbeck process with jumps. In addition, based on Parsons [2013], Henaff et al. [2013] studies the hedging problem with interactions between spot market, futures market and storage capacity, where a two-factor model is considered for the future prices and a GARCH model with jumps is used for spot prices. Deviating from the above models, Boogert and De Jong [2011] and Parsons [2013] decompose the long-term and short-term behavior of the observed price dynamics and describe them separately in a two-factor model.

Option pricing techniques Both Jaillet et al. [2004] and Secomandi [2010] apply trinomial lattice techniques to solve for the capacity prices. However, this method suffers from the curse of dimensionality and therefore requires heavy computations. Another popular technique is the Least Square Monte Carlo method where continuation values are approximated by functions of concurrent state variables, with their parameters estimated from earlier simulations. Boogert and De Jong [2008] modify the LSMC method in their valuation of storage capacity and test the basis functions and convergence for LSMC. They formulate the basis functions based on both storage space and gas prices, where grids of storage space are created for each period. This is an efficient adjustment for the use of LSMC but by using the grids without considering any restrictions on the injection/withdrawal rate or the gas-in-storage prior to certain periods, it is likely to lead to an upward bias, overestimating the value of the storage contracts. To solve the stochastic control problem, both Thompson

et al. [2009] and Carmona and Ludkovski [2008] also use (a modified) version of the LSMC method for its simplicity, flexibility and fast convergence properties. We have also settled on LSMC, also with modifications necessitated by the structure of our pricing problem; we elaborate on our specific approach below.

Incomplete market assumption Despite the extensive discussions on complex mathematical modeling, the existence and consequences of market incompleteness are effectively overlooked or trivialized by most of the literature by simply assuming zero risk premia or by setting the physical probabilities equal to the risk-neutral adjusted ones. For instance, Thompson et al. [2009] state that the embedded risks are not fully hedgeable but they simply use real probabilities to avoid the incomplete market issue; Parsons [2013] argues that neither side of the market has bargaining power, therefore risk premium is too small that the physical probabilities become equivalent to the risk-neutral probabilities. None of these simplifying assumptions is satisfactory as we will show below.

2.2 GAS Models

The importance of correctly modeling underlying asset prices is widely accepted in the literature. For example, by allowing simple trading strategy for managing the storage with fixed spot price thresholds for actions including injection, withdrawal, and doing nothing, Secomandi [2010] indicates that the dynamics of stochastic spot prices have major effects on the optimal decision policies. He also emphasizes the importance of correctly modeling the uncertainty embedded in the underlying processes and shows how the valuations differ for high and low volatility assumptions. However, most of earlier research ignores the complex volatility structure into consideration. We extend the existing applications and take one step away from GARCH models by employing and econometrically testing GAS models, in order to characterize some unique features of the time series.

GAS models have recently been developed by Creal et al. [2013], Harvey [2013]. Creal et al. [2013] first propose a model framework called Generalized autoregressive score (GAS) Model, which updates the dynamics based on the scaled score of the loglikelihood. This GAS model is able to fully exploit the information provided by the loglikelihood function. Koopman et al. [2015] investigate the forecasting performance of GAS models comparing to state space models. Their Monte Carlo study confirms that GAS models outperform many observation-driven models in terms of predictive accuracy; and they perform closely to correctly specified parameter-driven models in terms of forecasting errors. Creal et al. [2011] consider a multivariate GAS model

for times series, particularly with fat tails, where they develop a novel method to estimate the time-varying correlations and volatilities, a method we also use. Our estimation and simulation results similarly show that the multivariate t-GAS model slightly outperforms GARCH/DCC models in terms of loglikelihood values, and provides more accurate out-of-sample predictions. In this paper, we extend it to a bivariate t-GAS model, in order to further investigate the time-varying volatilities and time-varying correlations between two correlated time series.

3 Data Analysis

3.1 Statistical Descriptions

Initiated in 2003, TTF has quickly grown into one of the biggest gas markets in Europe in term of daily trading volumes. Figure 1 plots the daily spot prices for day-ahead, month-ahead, quarter-ahead¹, season-ahead², and year-ahead data covering the period from early 2005 to early 2015. The day-ahead price is defined as the price for a standard unit of natural gas (one megawatt hour) to be delivered the next working day. The month-ahead price is therefore the price that will be paid for delivering one megawatt hour through each working day of the next calendar month³.

¹Quarters: Q1 = January to March, Q2 = April to June, Q3 = July to September, Q4 = October to December.

²Seasons: Winter = October to March, Summer = April to September.

³Source: Continental Gas Snapshot Methodology by ICIS.

Table 1: Statistical descriptions of TTF daily spot prices and daily returns

(a) Statistical descriptions of daily spot prices

Statistics		Day	Month	Quarter	Season	Year
Sample size		2616	2616	2616	2616	2616
Mean		20.1748	20.8477	22.0978	23.3084	23.9975
Std.Dev.		5.7739	5.6945	6.1462	6.1121	4.7082
Correlation	Day	1.0000	0.8920	0.7670	0.6048	0.7035
	Month		1.0000	0.9054	0.6676	0.7885
	Quarter			1.0000	0.7657	0.8357
	Season				1.0000	0.9034
	Year					1.0000

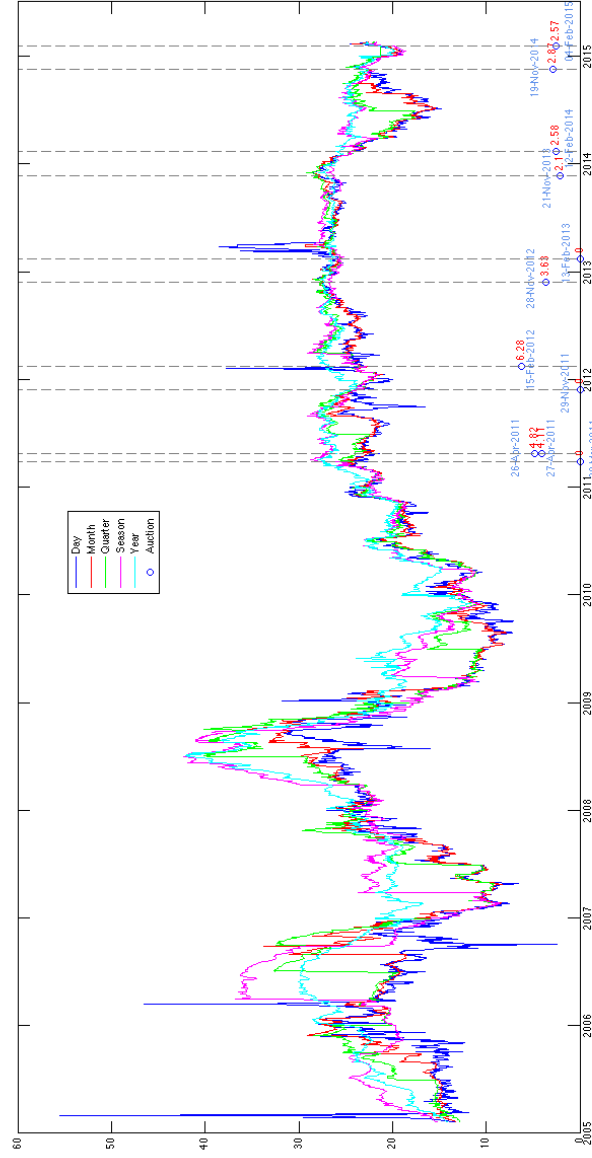
(b) Statistical descriptions on returns on daily spot prices

Statistics		Day	Month	Quarter	Season	Year
Sample size		2615	2615	2615	2615	2615
Mean ($\times 10^{-3}$)		0.1990	0.2036	0.2081	0.1364	0.1807
Std.Dev.		0.0694	0.0297	0.0294	0.0282	0.0197
Correlation	Day	1.0000	0.2661	0.2018	0.1610	0.1449
	Month		1.0000	0.3574	0.2541	0.3580
	Quarter			1.0000	0.2695	0.3063
	Season				1.0000	0.3501
	Year					1.0000

Statistical descriptions of daily spot prices and daily returns are shown in Table 1a and 1b respectively. It can be seen that the average spot price is increasing in the maturity of the contract. The correlations of the spot prices between a year-ahead contract and other contracts are shown in the last column of Table 1a, and the correlation coefficient increases as the differences of the delivery periods are smaller. This suggests the existence of price persistence in the data. From Table 1b, we observe that the standard deviation is declining with the maturity of the contract. Thus the data exhibit mean reversion.

Our study focuses the dynamics between Day-ahead and Month-ahead return data. To further investigate the features of the time series, we plot the spot return time series in Figure 2a and the spot squared return time series in Figure 2b. In both figures, clustering of volatility is clearly visible. Therefore we continue our analysis with GARCH and GAS models in the later sections. Furthermore, note that the time series presents certain features of heavy tails, therefore robust GAS model is expected to incorporate these characteristics.

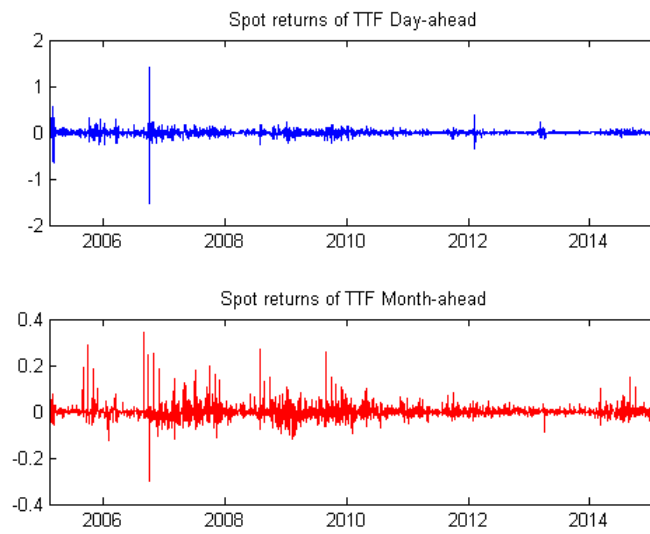
Figure 1: TTF daily spot prices (euro/MWh) and historical auction prices (euro/SBU)



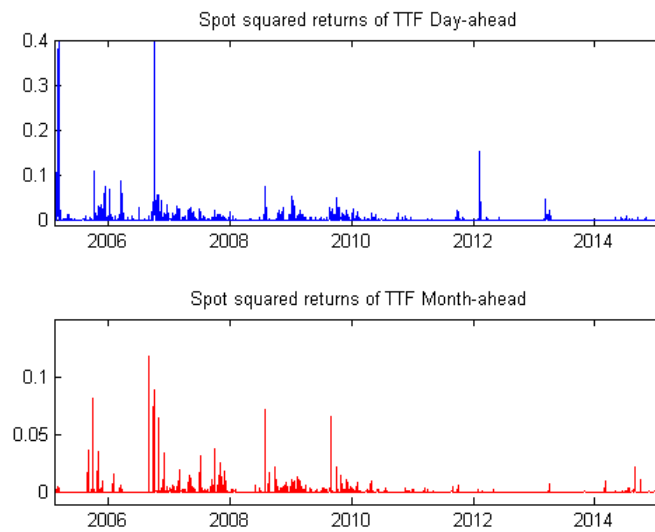
The vertical dashed lines denote auctions and the dots label the auction prices achieved through auctions.

Figure 2: Spot Returns and Spot Squared Returns of TTF Day-ahead and Month-ahead Data

(a) Spot Returns of TTF Day-ahead and Month-ahead



(b) Spot Squared Returns of TTF Day-ahead and Month-ahead



3.2 Seasonality

It is widely accepted that seasonality plays an essential role in energy spot markets due to the high sensitivity of gas demand and supply to temperature. Figure 3 gives a quick view on the seasonality in spot prices and returns for both day-ahead and month-ahead time series, including the mean and the standard deviation for each month.

For day-ahead and month-ahead data, the average highest (lowest) spot price over a year happens in December and October (August and April); while the most (least) volatile periods are March and September (January and January) respectively. The level and volatility of spot price for day-ahead is relatively high (low) during the winter (summer) period, which can be explained by the high (low) demand for gas in cold (warm) days. However, the price level of month-ahead time series is high and volatile in the autumn (September and October) before entering the winter and the price stays relatively high until spring. It can be explained that the gas providers are preparing for the high demand in the winter and start buying gas forwards (e.g. month-ahead contracts); while when spring comes, the gas providers feel less pressured to buy products on TTF market and the month-ahead price falls.

Table 2: Seasonality Tests on Day-ahead and Month-ahead Spot Prices/Returns

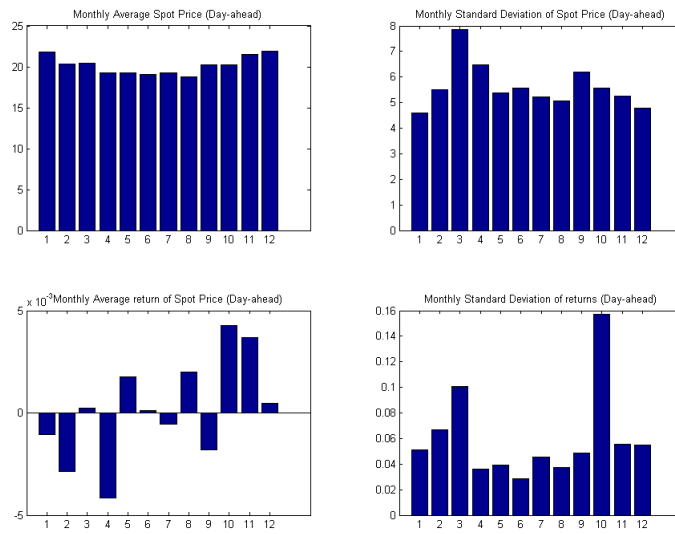
Month	Day-ahead		Month-ahead	
	Spot Price	Price Return	Spot Price	Price Return
December Level	21.8998***	0.0005	22.9422***	-0.0029
Jan	-0.0572	-0.0015	-0.9877*	0.0001
Feb	-1.5373***	-0.0033	-3.2412***	-0.0010
Mar	-1.4445***	-0.0002	-3.6333***	0.0039
Apr	-2.6608***	-0.0046	-3.8597***	0.0022
May	-2.6758***	0.0013	-3.4407***	0.0041
Jun	-2.8741***	-0.0003	-3.7248***	0.0028
Jul	-2.6251***	-0.0010	-3.3896***	0.0042
Aug	-3.1393***	0.0015	-3.1206***	0.0057**
Sep	-1.6454***	-0.0022	-0.5338	0.0099***
Oct	-1.6704***	0.0038	0.4623	0.0036
Nov	-0.3887	0.0032	0.2715	0.0016

***, **, * represent 1%, 5%, 10% significance level respectively.

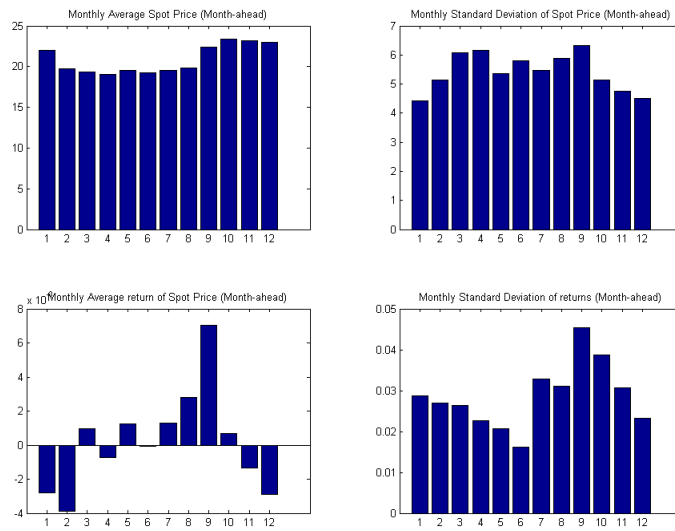
Unfortunately the spot price time series contains a unit root, therefore we instead focus on the return series

Figure 3: Seasonality

(a) Day-ahead



(b) Month-ahead



of day-ahead and month-ahead data, which are stationary according to the results of Phillips-Perron tests.

Table 2 illustrates the results of formal tests for seasonality in spot prices and returns for both day-ahead and month-ahead data. Seasonality effect is presented significantly for both spot day-ahead and month-ahead prices. Despite the test rejects the significance effects of seasonality in the daily price returns for day-ahead, seasonality still exists in the return series of month-ahead data: significantly higher average spot returns are found in August and September comparing to the average return in December. We have to take this into account in our modeling and a simple solution would be described as follows. First eliminate the seasonality effect and model the adjusted time series with standard GARCH or GAS models. Second, add back the seasonality effect to the simulated future data series.

4 Auction Analysis

The storage capacity auctions conducted by GasTerra are carried out over at least two rounds and at most five rounds. For each round, bidders have to submit their bids including both the price for one Standard Bundled Unit (euro per SBU) and the amount of SBUs required. After the first round, a start price for the second round will be announced, which serves as a minimum bidding price for the second round. Given this information, bidders may adjust their bidding prices in the next round, provided that only upward adjustment is allowed. Those who decide not to raise their bids above the minimum price of the subsequent round will be considered out of market. Unfortunately, we only have the final weighted average price paid by successful bidders, but not the auction data for each round of auctions in detail.

The first gas storage auction in the Netherlands was held by GasTerra on March 28, 2011, followed by ten auctions since then. Both Figure 1 and Table 3 present the historical auction prices for GasTerra Storage Service, obtained from ICE-ENDEX. More specifically, Table 3 shows the weighted average price of each auction held, where a price of zero means an auction failure. In addition, the last column of Table 3 gives the information on the highest bid even if the auction failed. The auction price is defined as euros per SBU.

Table 3: Historical Auction Prices From GasTerra Storage Service (Data source: ICE-ENDEX)

Auction Date	Auction Price (euro/SBU)	Highest bid if failed (euro/SBU)
28-Mar-2011	0	2.96
26-Apr-2011	4.82	
27-Apr-2011	4.11	
29-Nov-2011	0	3.01
15-Feb-2012	6.28	
28-Nov-2012	3.63	
13-Feb-2013	0	1.47
21-Nov-2013	2.10	
12-Feb-2014	2.58	
19-Nov-2014	2.87	
04-Feb-2015	2.57	

A natural question that follows is how the auction results are related to the market. One widely accepted belief is that the bidders consider the winter-summer difference as the main elements of the storage prices. In order to understand whether the statement is true or false, we compute the correlations between auction prices and the gas market daily spot prices/returns in Table 4, including the average daily spot price/return and standard deviation of daily price/return in the past 1/3/6 month(s), as well as the winter-summer price/return difference in the previous year and in the next year. One surprising observation is that, opposite to the popular opinion mentioned before, the correlations between auction prices and the winter-summer price/return difference are relatively low, for both cases of winter-summer price difference of the previous year and the following year.

In all the figures in Table 4, the standard deviations of the spot returns in the past 1/3/6 month(s) are the most highly and positively correlated to the auction prices (with highest bid prices considered given auction failure). Since observed volatility (standard deviation) is the standard benchmark for investors to assess risks, the result implies that when the observed market risk is high, the storage buyers assign high value to the storage capacity. This is reasonable since storage capacity is an effective tool to hedge market risk, given that the corresponding market on futures and options is poorly developed. This point is missing from most of the literature.

Moreover, the average daily price in the past 1/3/6 month(s) is negatively correlated to the auction price

Table 4: Correlation between auction prices and gas market daily spot prices/returns

Correlation between auction price and gas market daily spot price		Auction price (when failed, 0)	Auction price (when failed, the highest bid)
Winter-Summer price difference	in the previous year	-0.0062	-0.0687
	in the next year	0.1691	-0.1713
Average daily spot price	in the past 1 month	-0.1217	-0.1981
	in the past 3 months	-0.1925	-0.3192
	in the past 6 months	-0.1573	-0.3227
	in the past 1 month	0.5769	0.7365
Std. Dev. of daily spot price	in the past 3 months	0.3224	0.5125
	in the past 6 months	0.2110	0.2092
Correlation between auction price and gas market daily return		Auction price (when failed, 0)	Auction price (when failed, the highest bid)
Winter-Summer return difference	in the previous year	0.2082	0.2349
	in the next year	0.0191	0.3228
Average daily return	in the past 1 month	-0.0663	-0.0440
	in the past 3 months	-0.0750	-0.1016
	in the past 6 months	0.2588	0.2862
Std. Dev. of daily return	in the past 1 month	0.6040	0.7989
	in the past 3 months	0.4663	0.7692
	in the past 6 months	0.5184	0.7629

with failed prices considered. In other words, if the average price in the past few months is high, the storage buyers tend to bid less for the future storage; and vice versa. One possible explanation is that a low market price implies a high market supply, thus the potential storage investors have easy access to the gas products and do not need storage capacity for insurance reasons. Note that if the price of a failed auction is given as zero, the correlation between average daily price and auction price is still low and negative.

In brief, both the low correlation between winter-summer difference and the storage capacity prices, and the high correlation between the market volatility and the storage capacity prices have led to the invalidation of the commonly accepted belief. Therefore, we need to develop a new model to accommodate and explain the counter-intuitive observations.

5 Decision Strategy

This section presents the model for pricing storage capacity. A key feature of our approach involves parameterizing key factors affecting the investors' valuation for the storage capacity. The capacity storage service provided by GasTerra is an example of slow storage⁴, constrained by both the storage space and the injection/withdrawal capacity. We assume there is a 2-day delay of the execution after the decision is made, and no further action can be undertaken before the execution of the previous action is finished. An indirect tradeoff of withdrawal or injection is that the larger amount is ordered, the longer periods are locked from new action being enforced. Therefore, in order to effectively take advantage of the storage capacity, we assume the capacity holder makes a decision with monthly frequency. Thus the capacity holder is able to continue with their injection/withdrawal in the following period or revise their initial long-term plan based on new market information.

5.1 Value Functions

Consider the pricing of a standard contract of one SBU, with the maturity of one year⁵ (i.e. $T = 1$). Define the maximum capacity or storage space C^{max} , minimum space C^{min} and the current usage level C_t , therefore

⁴A fast storage does not have the injection/withdrawal capacity constraint, which is therefore simpler to deal with regarding modeling and optimization.

⁵We consider a one-year contract offered by GasTerra covering the period from 1 April up to 31 March (incl.) the following year.

we have $C_t \in [C^{min}, C^{max}]$. For a GasTerra product, $C^{min} = 0\text{kWh}$ and $C^{max} = 1440\text{kWh}$. After observing the current storage usage level C_t and the current market gas price P_t^D (day-ahead) and P_t^M (month-ahead) at time t ($t \in [0, T]$), the capacity holder makes a decision of injection (A^I), or withdrawal (A^W), or doing nothing (A^N). Define π as the action set $\pi(C_t, P_t^D, P_t^M) = \{A_t^I, A_t^W, A_t^N\}$.

Here, the injection capacity is $r^I = 0.3333 \text{ kWh/h}$ and withdrawal capacity is $r^W = 1.0 \text{ kWh/h}$. Note that for an SBU offered by GasTerra, it takes 180 days to inject an empty capacity space until its maximum capacity and it takes at least 60 days (approximately 69 days) to deplete a full storage.

Costs for injection and withdrawal are $c^I = 0.00042 \text{ euros/kWh}$ and $c^W = 0.00003 \text{ euros/kWh}$ respectively, and it is never optimal to withdraw and to inject at the same time due to the positive costs related to both actions. The injection factor b^I is constant and equals 1.0, and the withdrawal factor b^W is fixed in the contract and defined as (also see Figure 10)

$$\begin{aligned} b_t^W &= b_t^W(C_t) = \min\left(1, \frac{1}{2160}C_t + 0.6\right) \\ &= \begin{cases} 1 & \text{if } 864\text{kWh} \leq C_t \leq C^{max} \\ \frac{1}{2160}C_t + 0.6 & \text{if } C^{min} \leq C_t < 864\text{kWh} \end{cases} \end{aligned}$$

The immediate payoff function at time t for injection can be written as

$$v_t^I = -(\min(P_t^D, P_t^M) + c^I) A_t^I;$$

and similarly, the withdrawal payoff function at t is (the gas extracted can be sold one day after the completion of withdrawal)

$$v_t^W = (\max(P_t^D, P_t^M) - c^W) A_t^W.$$

Naturally, the payoff function for A^N is simply

$$v_t^N = 0.$$

At the final date T , the action injection is not allowed, therefore the payoff function is

$$v_T^W = (\max(P_T^D, P_T^M) - c^{WT}) C_T$$

where a fine c^{WT} ($> c^W$) has to be paid out for any gas in storage C_T left at the final contract date.

The optimal value function at time t is $V_t(P_t^D, P_t^M, C_t)$. Given a discount rate β , the dynamic program for $t \in [0, T)$ can be formulated as,

$$V_t(P_t^D, P_t^M, C_t) = \max_{A_t^j \in \pi(C_t, P_t^D, P_t^W)} w_t(A_t^j, P_t^D, P_t^M, C_t)$$

$$w_t(A_t^j, P_t^D, P_t^M, C_t) = v_t^j + \beta \mathbb{E}^Q [V_{t+1}(P_t^D, P_t^M, C_t + A_t^j)]$$

where $j \in \{I, W, N\}$.

5.2 Utility Function

As we have discussed before, preference free pricing does not exist in incomplete markets, so we have to parametrize the risk aversion of the investor to deal with the market incompleteness. Since the auction price is only given as the weighted average for one SBU, it is reasonable to assume a CRRA (constant relative risk aversion) utility function. As a result, we assume the investor has a utility function as below:

$$u(x) = \frac{x^{1-\alpha} - 1}{1-\alpha}$$

where $\alpha (> 0)$ is her risk aversion parameter. Note that when $\alpha \Rightarrow 1$, this utility function approaches a log utility in the limit and can be replaced by logarithmic utility for $\alpha = 1$.

5.3 Modified LSMC Method

Boogert and De Jong [2008] adopt LSMC for capacity pricing by using grids of gas-in-storage volumes to form basis functions. However, as we have explained in the previous section, the gas-in-storage is highly dependent on previous actions and the grids may not be achievable, which may results in an overestimation of prices due to the ignorance of constraints. Thus, instead of assigning grids of storage levels for each period, we modify

LSMC in the following manner. First we compute the optimized storage strategy for each simulated day-ahead and month-ahead pair. Second, similar to LSMC, start from the final date T and compute the final cash flows. Third, regress the final cash flows on basis functions formulated from simulated day-ahead and month-ahead prices as well as gas-in-storage and injection/withdrawal amount at time $T - \Delta t$; and the fitted value from this regression is the continuation value. Fourth, compare the continuation value with the value of a project when no further action would be taken afterward, and work backward.

6 Estimation Results

We model the clustering volatility via both GARCH and GAS models, based on both normality and student t assumptions. The results are reported in Appendix 9.2. It is evident from both Table 6 and Table 7 that both the day-ahead and month-ahead data series are heavy-tailed, thus Student t assumption is preferred over a Gaussian one. When comparing t-GARCH models with t-GAS models, the latter yields a higher loglikelihood level, which implies a better fit of the data. Take the results for day-ahead data for example. From Table 6, we conclude that the t-GAS model fits the data better than t-GARCH one, with rejecting a loglikelihood ratio test $LR = 2 \times (-2144 + 2164) = 40 > \chi^2(1)$. Furthermore, the estimated degrees of freedom ν is 3.7238 with the t-GAS model and is significant with 1% level, which proves the heavy-tailed feature in the day-ahead return series. Further discussions in model selections and result comparisons between GARCH and GAS models are presented in Appendix 9.2.

Table 5 demonstrates the estimation results from a Bivariate t-GAS model. As can be seen, the estimated degree of freedom is 3.2935, which again confirms the fat-tails featured in the data.

Table 5: Estimation Results for the Bivariate t-GAS Model

Bivariate t-GAS	Seasonality Eliminated
ω (Variance)	1.3743***
	1.4263***
ω (Correlation)	0.2529***
A	0.0779***
	0.0304***
	0.0065***
B	0.9985***
	0.9980***
	0.9978***
ν	3.2935***
Loglikelihood	-2345.2
AIC	4710.4
BIC	4769.1

***, **, * represent 1%, 5%, 10% significance level respectively.

7 Results for Storage Capacity Prices

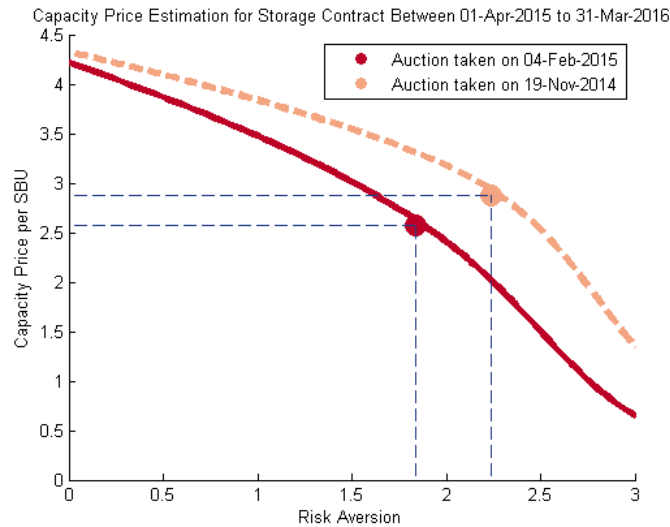
7.1 Comparison of two auctions for the same product

The hedging plan can be described as following: the investor compares the day-ahead and month-ahead price on a monthly basis, then buys and injects at the low price, and withdraws and sells at the high price. Note that injection and withdrawal are mutually exclusive actions, so the investor has to make a decision one way or the other at any given time.

We compare our calculations for two auctions held respectively on February 4, 2015 and November 19, 2014. Intuitively, the prices for the two auctions should be close to each other given that the auction target for both auctions is exactly the same contract. In fact, given other conditions the same, the price for the first auction should be expected to be slightly lower than the second since the latter is closer to the maturities of the options brought in by storage capacity. However, the realized auction prices not only differ from each other, with 2.87 euros/SBU for November 19, 2014, and 2.57 euros/SBU for February 4, 2015 respectively, but the second auction price is actually lower in spite of the shorter time to maturity.

Our modeling results for capacity prices are shown in Figure 4. The red dot presents the realized auction

Figure 4: Risk Aversion vs. Storage Capacity Price (November 19, 2014 and February 4, 2015)

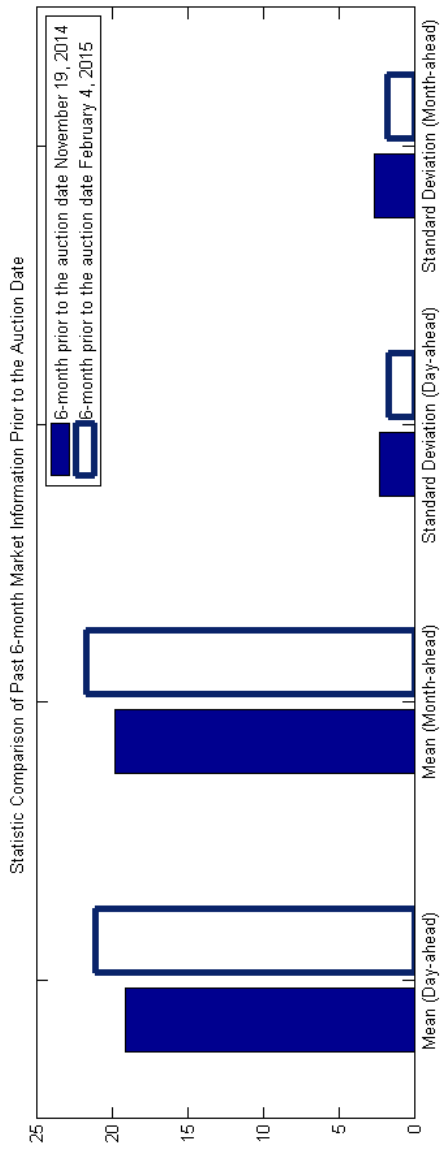


price on February 4, 2015 and the solid red line is the estimated price based on our model given different risk aversion parameters of *potential* storage capacity buyers; the orange dashed line and dot present estimated price and realized auction price for the auction held on November 19, 2014. Note that the contracts for both auctions are the same, i.e. one standard SBU contract covering the period of April 1, 2015 to March 31, 2016. For each calculation of the capacity prices, we use the market information until the date of the auction.

Several observations stand out. First *for given risk aversion*, bidders valued the contract less during the second auction than during the first; the solid line lies everywhere below the slotted line in 4. To see why, look at Figure 5: there we show that between the two auction dates, both the market price and the market volatility changed, the price went up somewhat but volatility went down. Table 4 shows that both price and volatility are positively correlated with the auction price, which makes sense given the call option characteristics of the contract. But the volatility effect apparently dominated, which can be seen from Figure 4: the estimated prices for November 19, 2014 dominate those for February 4, 2015, for all risk aversion levels. Higher volatility leads to higher option values, as is to be expected given the convexity of the payouts of the implied options embedded in the storage contract. If the storage capacity is an effective hedging contract, it is reasonable that the bidder is more keen on a hedging product when facing a more risky market.

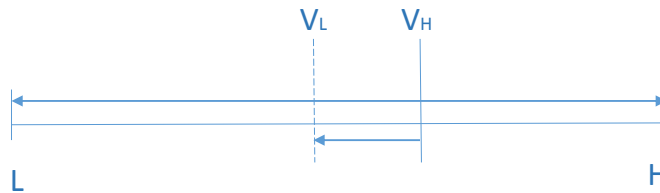
Second, the revealed risk aversion of the average capacity buyer on February 4, 2015 is lower than the one

Figure 5: Statistics Comparison of Past 6-month Market Information Before Auction Date (November 19, 2014 and February 4, 2015)



on November 19, 2014⁶. The implied risk aversion of the successful bidders is around 2.3 for the auction taking place on November 19, 2014, and 1.8 for the second auction on February 4, 2015. This can be understood from the incomplete markets structure of the problem. In a complete market setting risk aversion heterogeneity of bidders has no impact on the price (preference free pricing or risk neutral valuation) because additional risk will be compensated for at the market price of risk. However with markets incompleteness that neutrality breaks down and shifts in the composition of the pool of bidders can have an impact on the price. The lower value given risk aversion (the solid line lies below the slotted line in 4) has as a consequence that only less risk averse investors are drawn into the auction the second time around, with as a result a lower average implied degree of risk aversion. This can be further explained by Figure 6, which ranks all *potential* bidders in order of increasing risk aversion from left to right. As their value of storage capacity goes down from V_H to V_L , more risk averse bidders will be driven out of the market. As a result, the implied average risk aversion of bidders goes down.

Figure 6: Rank All *Potential* Bidders in Order of Increasing Risk Aversion

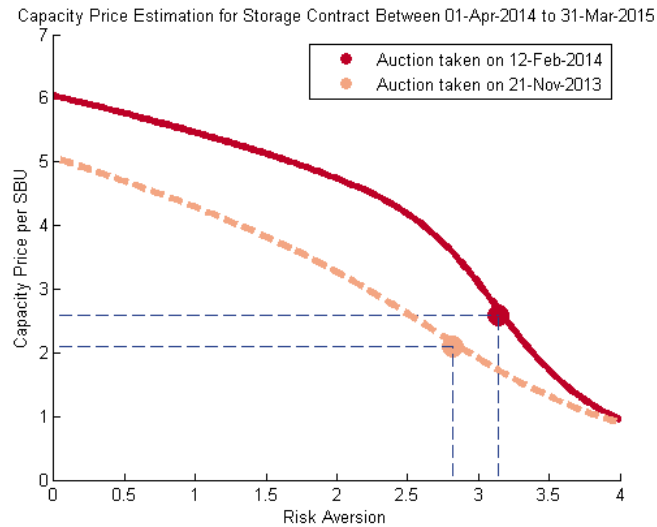


A third interesting observation is that the plot of capacity price as a function of risk aversion is declining faster in case of the auction taken on November 19, 2014 than the one taken on February 4, 2015. This suggests that investors with various risk aversion levels disagree more in valuing the capacity product when the market is more volatile (November 19, 2014).

⁶The auction design is such that every successful bidder pays his/her bidprice, and the reported auction price is the average price paid, so the implicit risk aversion is a (complicated) weighted average of the risk aversion parameter of all successful bidders.

7.2 Comparing two more auctions: November 21, 2013 and February 12, 2014

Figure 7: Risk Aversion vs. Storage Capacity Price (November 21, 2013 and February 12, 2014)



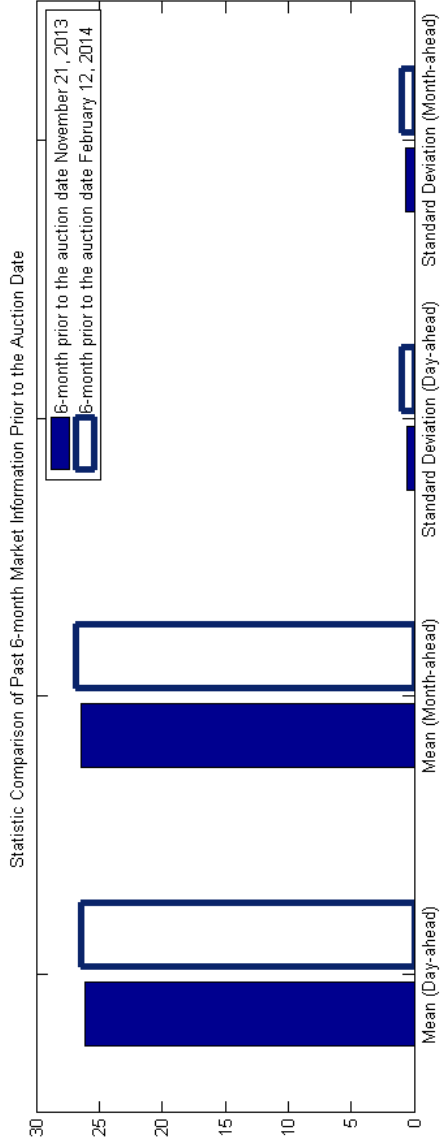
A similar line of reasoning explains the results stemming from comparing another set of auction results, those of November 21, 2013 and February 12, 2014, as shown in Figure 7. Note that the market risk is higher for the second auction in this case, see Figure 8. In this case, higher volatility causes the second auction schedule to lie above the first, so more investors are drawn into the auction and the average implied coefficient of risk aversion is higher the second time around.

We only compare these two pairs of auctions because at least one auction failed for the other pairs.

7.3 Injection/Withdrawal Rate vs. Capacity Price

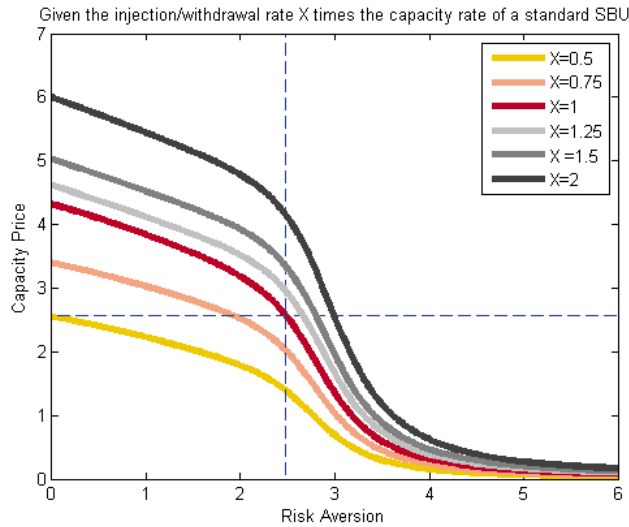
We have calculated the price for one standard SBU as specified in the GasTerra gas storage auction, which is a slow storage as explained in Section 5. In this section we do not assume a fixed rate, but we investigate how the injection/withdrawal capacity affects the value of the storage contract. Assume six different contracts with the injection/withdrawal rate equal to X times the original contract on November 19, 2014. The benchmark case is $X = 1$, which is the same contract as we considered before (the second auction in Section 6). If $X > 1$, it implies a faster storage and naturally results in a higher value of the contract since the capacity holder would

Figure 8: Statistics Comparison of Past 6-month Market Information Before Auction Dates (November 21, 2013 and February 12, 2014)



have more flexibility within the contract periods; and vice versa.

Figure 9: Capacity Price Given the Injection/Withdrawal Rate X times the Capacity Rate of a Standard SBU



This impact of capacity rates is demonstrated in Figure 9. For instance, when the capacity rate is only half as the benchmark ($X = 0.5$), the capacity price would drop to 1.38 euros per (nonstandard) SBU given the implied risk aversion ($\alpha = 2.48$), nearly half the price that the benchmark implies (2.57 euros/SBU); while when $X = 2$, the price increases to 4.14 euros per (nonstandard) SBU with $\alpha = 2.48$, about 1.61 times of the benchmark price. Therefore we conclude that the value of the contract is increasing with the capacity rate, however, the marginal effect of capacity rate is decreasing.

8 Conclusion

In this paper, we have developed a model for pricing gas storage capacity, a model that focuses on the interactions between the gas spot market, price volatility and gas storage capacity. We explicitly model the time-varying volatilities and correlations between TTF spot day-ahead and month-ahead returns. Due to the heavy-tailed feature of the data, we adopt the recently proposed GAS models. Given a feasible hedging strategy, we then adopt and modify the Least Square Monte Carlo method for pricing purposes. Several interesting findings are worth mentioning.

First, by comparing two auctions with the same target contract, we find that the investors are willing to pay higher prices for storage contracts when the market is more volatile, given the same risk aversion level. This stems from the insurance feature of a capacity contract, a feature that is missing from most of the existing literature. A second feature of our results is inextricably linked to the incomplete markets nature of our set up: the price is affected by the pool of bidders. A higher value given any degree of risk aversion draws in more investors, changing the implied average risk aversion of the successful pool of bidders. The implied risk aversion is in all auction cases considered not zero, invalidating the assumption of risk-neutrality typically made in the literature. The point to note is that the revealed risk aversion of the bidders is far from zero, contradicting the assumption of risk neutrality made in the literature on capacity pricing. Apparently market incompleteness is essential and cannot be overlooked. Third, our results also suggest that investors disagree more on the storage capacity prices when the market is riskier. Finally, a high injection/withdrawal capacity rate would increase the contract value, which shows that fast storage is more valuable than slow storage, given everything else.

References

- Alexander Boogert and Cyriel De Jong. Gas storage valuation using a monte carlo method. *The journal of derivatives*, 15(3):81–98, 2008.
- Alexander Boogert and Cyriel De Jong. Gas storage valuation using a multifactor price process. *The Journal of Energy Markets*, 4(4):29–52, 2011.
- René Carmona and Michael Ludkovski. Pricing asset scheduling flexibility using optimal switching. *Applied Mathematical Finance*, 15(5-6):405–447, 2008.
- Drew Creal, Siem Jan Koopman, and André Lucas. A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Business & Economic Statistics*, 29(4), 2011.
- Drew Creal, Siem Jan Koopman, and André Lucas. Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, 28(5):777–795, 2013.
- Cyriel De Jong. Gas storage valuation and optimization. *Journal of Natural Gas Science and Engineering*, 24(0):365 – 378, 2015. ISSN 1875-5100.

- Andrew C. Harvey. *Dynamic Models for Volatility and Heavy Tails*. Cambridge University Press, 2013.
- Patrick Henaff, Ismail Laachir, and Francesco Russo. Gas storage valuation and hedging. a quantification of the model risk. *Working paper, arXiv:1312.3789*, 2013.
- Patrick Jaillet, Ehud I Ronn, and Stathis Tompaidis. Valuation of commodity-based swing options. *Management science*, 50(7):909–921, 2004.
- Siem Jan Koopman, Andre Lucas, and Marcel Scharth. Predicting time-varying parameters with parameter-driven and observation-driven models. *Review of Economics and Statistics*, 2015.
- Guoming Lai, François Margot, and Nicola Secomandi. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations research*, 58(3):564–582, 2010.
- Francis A Longstaff and Eduardo S Schwartz. Valuing american options by simulation: a simple least-squares approach. *Review of Financial studies*, 14(1):113–147, 2001.
- Cliff Parsons. Quantifying natural gas storage optionality: a two-factor tree model. *Journal of Energy Markets*, 6(1):95–124, 2013.
- Nicola Secomandi. Optimal commodity trading with a capacitated storage asset. *Management Science*, 56(3): 449–467, 2010.
- Nicola Secomandi and Duane J Seppi. Real options and merchant operations of energy and other commodities. *Foundations and Trends® in Technology, Information and Operations Management*, 6:3–4, 2014.
- Matt Thompson, Matt Davison, and Henning Rasmussen. Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics (NRL)*, 56(3):226–238, 2009.

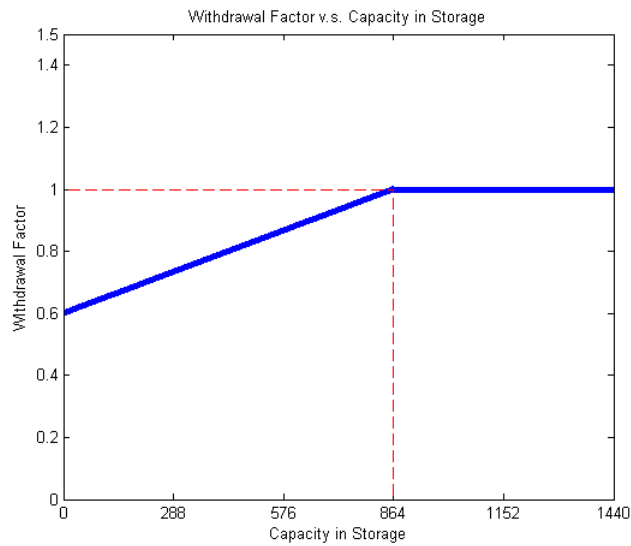
9 Appendix

9.1 A Standard GasTerra Contract

An SBU for a standard GasTerra Product (source: iceendex.com) is specified as follows.

1. Available storage space: 1440 kWh
2. Injection cost: EUR 0.00042 per kWh
3. Injection factor: 1.0
4. Withdrawal capacity: 1.0 kWh/h
5. Withdrawal cost: EUR 0.00003 per kWh
6. Withdrawal factor: if the gas-in-storage is between 0 kWh and 864 kWh, it increases linearly from 0.6 to 1; if the gas-in-storage is between 864 kWh and 1440 kWh, it equals 1. See Figure 10.
7. A Late Storage Fee has to be paid out for any gas remaining in Gas-in-Storage at the end of the contract Period.

Figure 10: Withdrawal Factor, storage service agreement 2015-2016, ICE-ENDEX



9.2 GARCH vs. GAS

Table 6: Estimate Results for Day-ahead Daily Logarithm Return Data

	Gaussian GARCH	Gaussian GAS	t-GARCH	t-GAS
omega	12.3954	-0.8901***	5.8559	-0.8943***
A	0.2285***	0.1493***	0.1568***	0.1754***
B	0.9992***	0.9244***	0.9985***	0.9546***
ν			3.5305	3.7238
Log-likelihood	-2399.465	-13657.141	-2164.468	-2144.376

*, **, *** represent a significance level of 10%, 5%, and 1% respectively.

Table 7: Estimate Results for Month-ahead Daily Logarithm Return Data

	Gaussian GARCH	Gaussian GAS	t-GARCH	t-GAS
omega	0.3937***	-0.9501***	0.4414	-1.7865***
A	0.0803***	0.0293***	0.0780***	0.0851***
B	0.9934***	0.9829***	0.9973***	0.9865***
ν			2.6344	2.6297
Log-likelihood	-1103.944	-140115.742	-405.425	-402.205

*, **, *** represent a significance level of 10%, 5%, and 1% respectively.

Moreover, as shown in Figure 1 and Table 1, the high correlation between the day-ahead and month-ahead time series could not be simply overlooked. Therefore, we extend the univariate t-GAS models into a bivariate t-GAS model. Let y_{1t} and y_{2t} be the logarithm daily returns of day-ahead and month-ahead respectively, and we have

$$\mathbf{y}_t = \epsilon_t \text{ or } \begin{pmatrix} y_{1t} \\ y_{2t} \end{pmatrix} = \begin{pmatrix} \epsilon_{1t} \\ \epsilon_{2t} \end{pmatrix}$$

where ϵ_t follows a bivariate student t distribution, whose covariance matrix is

$$\Sigma_t = \begin{pmatrix} \sigma_{1t}^2 & \rho_t \sigma_{1t} \sigma_{2t} \\ \rho_t \sigma_{1t} \sigma_{2t} & \sigma_{2t}^2 \end{pmatrix}$$

and degrees of freedom is ν . Following Creal et al. [2011], we choose the factor

$$f_t = \begin{pmatrix} \sigma_{1t}^2 \\ \sigma_{2t}^2 \\ \text{vech}(Q_t) \end{pmatrix},$$

where

$$R_t = \Delta_t^{-1} Q_t \Delta_t^{-1} = \begin{pmatrix} 1 & \rho_t \\ \rho_t & 1 \end{pmatrix}$$

is the correlation matrix and Q_t is symmetric and positive definite. To guarantee the properties of the correlations matrix, (e.g. symmetric, positive definite, and $|\rho_t| < 1$), the elements of the diagonal matrix Δ_t is

chosen as the square root of the diagonal elements of Q_t . The factor follows an autoregressive process as

$$f_{t+1} = \omega + A_1 s_t + B_1 f_t$$

where the scaled score function s_t is defined as

$$s_t = S_t \nabla_t \text{ and } \nabla_t = \frac{\partial \log p(\mathbf{y}_t | f_t, \mathcal{F}_{t-1}; \theta)}{\partial f_t}$$

with S_t being the inverse of the information matrix.