Expected vs. Observed Natural Gas Storage Usage: Limits to Intertemporal Arbitrage

Anne Neumann[1] and Georg Zachmann[2]

Paper submitted to the Annual Congress of the Verein für Socialpolitik

September 8th – 11th, 2009 in Magdeburg

Abstract

The aim of this paper is to investigate if storage operators holding capacities in less competitive markets differ in their use of natural gas capacities from companies active in a competitive market. Storage is a function of spot and forward prices and costs, subject to technical characteristics of storage facilities and operators maximize profits. Our main result of the empirical analysis is that perfect arbitrage theory fails to explain storage operation in Germany. Possible explanations encompass not only technical limits to storage, but also strategic behavior of market players active in natural gas storage.

Keywords: commodity markets, natural gas, storage

JEL-Codes: L13, G14, C53

[1] Correspondent author. DIW Berlin, Mohrenstr. 58. D-10117 Berlin. aneumann@diw.de.

[2] Larsen Paris-Sud (Paris XI), 27 avenue Lombart, 92260 Fontenay aux Roses.

1 Introduction

Germany is one of Europe's largest natural gas importers and consumers. Given falling domestic reserves, natural gas storage therefore plays an increasingly important role. However, in regulatory terms Germany holds the "red lantern" in Europe, with very little institutional reform progress and a largely non-competitive natural gas sector. Subsequently, storage capacities are inefficiently used and the signals for new storage investment are distorted. In this paper we analyze the structure of natural gas storage in Germany and apply a simple econometric model to see if a particular storage site is efficiently used. The paper starts with theoretical considerations on the theory of storage and provides some technical details of storing natural gas. We then introduce the natural gas storage activities in Germany (Section 3). It is dominated by depleted gas and oil fields, but aquifers and salt caverns also play a significant role. The inefficient access to existing storage sites of the incumbents has prompted new market entrants to invest massively into new sites. In Section 4 we develop a model to evaluate the usage strategy of the observed use of storage with the "perfect arbitrage" solution. By comparing the optimal benchmarking behavior withobserved data, we can infer if the storage market works competitively. In Section 5 we apply the model to real-time data of a large storage site, Dötlingen owned by BEB. We find that injection and withdrawal decisions are not based on the profit maximizing behavior of a small player in a liquid market.

2 Empirical Models and Theoretical Foundations

In general, storage transfers a commodity from one period to the next, including the related costs due to intertemporal arbitrage. Consumers holding inventories receive an income stream referred to as convenience becoming due in times of production and/or supply shocks. Therefore, theory implies difference of spot and forward prices of a commodity at a level given by storage and interest costs (for storing) less convenience yields. Moreover, marginal convenience declines with increasing aggregate level of storage following a convex shape (Fama and French, 1987). The convex shape of the convenience yield implies a modest impact of changes in stock level on marginal costs of storage. Therefore, variations in spot prices are directly related to the benefit of holding inventory and inversely related to the correlation between spot and forward prices. Storage serves to balance short term differences in demand and supply. Entrepreneurial decision criterions for the use of storage are essentially described by: "Store until the expected gain on the last unit put into store just matches the current loss from buying - or not selling it - now" (Williams and Wright, 1991, p. 25). Storage facilities therefore induce arbitraging potential in functioning markets. Traders consider storage as an option derived as the sum of intrinsic and extrinsic values. In other words, the value derived from forward quotations and volatility of spot prices. Wright and Williams (1982) show that storage in a model where production is stochastic and both production and storage are performed by competitive profit-maximizers is favourable for consumers.

Deaton and Laroque (1996) investigate commodity prices for harvest assuming existence of speculators and competitive storage. Defining risk-neutral and profit-maximizing stockholders implies the nexus of spot prices over time periods. The authors show that the effect of storage on prices is only modest, but stronger on the mean and variance of the following period. Wright and Williams (1989) argue that backwardation[3] reflect a risk premium that drives futures prices down. Moreover, they argue that a negative price for storage is a positive difference between full carrying cost and expected rate of change of the spot price.

Markets for natural gas have been of interest for an application of storage theory. This is mainly due to the peculiarities of energy sources as compared to wheat or coffee: natural gas storage is limited by technical factors influencing operability of facilities induced by geological and technical characteristics, and strong seasonality. However, the existence of a number of spot markets (with futures and options traded at) for natural gas and the intertwining of former regionally segmented markets in the US resulted in applications of storage theory. Susmel and Thompson (1997) provide empirical evidence demonstrating that an increase in price volatility was followed by investment in additional storage facilities. The increase of the variance inherent in spot prices (due to changing market structure and institutional framework) theoretically results in an increased use of storage (an increase in volatility increases the marginal benefit of holding inventory). An application to the Californian market for natural gas is provided by Uria and Williams (2007) arguing that injection decisions rather than resulting stock levels respond to price differences ("despite official seasons, regulatory requirements, and operational rigidites"). Using daily flow data the authors show that injection in Californian facilities increases slightly with a strengthening intertemporal spread on NYMEX. Serletis and Shahmoradi (2006) test the theoretical prediction that when inventory is high, large inventory responses to shocks imply roughly equal changes in spot and futures prices, whereas when inventory is low, smaller inventory responses to shocks imply larger changes in spot prices than in futures prices. Their tests on North American spot and futures natural gas prices confirm these predictions of the theory of

[3] Backwardation refers to a situation in which a commodity's future price for future delivery is below the price for immediate delivery.

storage. Wei and Zhu (2006), Dencerler, Khokher, and Simin (2005) and Khan, Khoker and Simin (2005) model risk premiums and the dependence of futures prices on inventory levels with a focus on mean-reverting behaviour for natural gas among other US commodities. Chaton et al. (2008) develop a model of seasonal natural gas demand taking into account the exhaustibility of the resource as well as supply and demand shocks. In a competitive setting the effect of policy instruments, i.e. tariffs or price caps, are investigated and applied to the North American market.

The technology of underground natural gas storage differs in the physical and economic characteristics of the sites. Deliverability rate, porosity, permeability, retention and capability of a site are the main physical of each storage type. To make operation of a storage site financially viable site preparation, maintenance costs, deliverability rates or cycling capacity are the main features. Key for profitable site operation is capacity and deliverability rate. The more natural gas injected or withdrawn the higher the economics of scale. Flexibility and therefore the ability to react to short-term price signals require reasonable deliverability.

Depleted gas and oil fields (DGF, DOF) can be converted to storage while making use of existing wells, gathering systems, and pipeline connections. Natural aquifers are suitable for storage if the water bearing sedimentary rock formation is overlaid with an impermeable cap rock. Whereas aquifers are similar to depleted gas fields in their geology they require more base (cushion) gas and greater monitoring of withdrawal and injection performance. Deliverability of the site can be enhanced if there is an active water drive. The highest withdrawal and injection rates relative to their working gas capacity are provided by salt caverns. Moreover, base gas requirements are relatively low. Constructing salt cavern storage facilities in salt dome formations is more costly than depleted field. But the ability to perform several withdrawal and injection cycles each year reduces the per-unit cost of natural gas injected and withdrawn.

The fundamental characteristics of an underground storage facility distinguish between the characteristic of a facility (i.e. capacity), and the characteristic of the natural gas within the facility (i.e. inventory level). Total natural gas *storage capacity* is the maximum volume of natural gas that can be stored in an underground storage facility at a particular time. *Base gas* (or cushion gas) is the volume of natural gas intended as permanent inventory in a storage reservoir to maintain adequate pressure and deliverability rates throughout the withdrawal season. *Working gas capacity* is the volume in the reservoir above the level of base gas and is available to the storage operator. *Deliverability* is a measure of the amount of natural gas that can be delivered (withdrawn) from a storage facility on a daily basis (often referred to also as deliverability rate, withdrawal rate, or withdrawal capacity. Deliverability varies and depends on factors such as the amount of natural gas in the reservoir, the pressure within the reservoir, compression capability available to the reservoir, the configuration and capabilities of surface facilities associated with the reservoir, and other factors. It is highest when the reservoir is full and declines as working gas is withdrawn. *Injection capacity* (or rate) is the complement of the deliverability and is the amount of natural gas that can be injected into a storage facility on a daily basis. It is inversely related to the total amount of natural gas in storage (EIA, 2004).

Depending on the type of storage, investment costs, lead times and operating costs differ. There are no exact figures on natural gas storage sites available, but Grewe (2005) provides some good estimates which are presented in Table 1.1.

Table 1: Storage costs

	DGF/DOF	Aquifer	Salt cavern
Specific investment costs per m^3 working gas [Euro/m^3]	0,18-0,33	0,38-0,40	0,54
Specific investment costs in withdrawal rate [Euro/m^3]	11,4-22,7	26,5-34,8	13,6
Total costs per cycled m3 working gas p.a. [Euro cent/m^3][1]	5,86	6,73	9,81
Total costs per (m^3/day) withdrawal rate p.a. [Euro cent/(m^3/day)/a]2	3,82	5,87	1,99

[1] Capital costs plus fix and variable operating costs.

[2] Capital costs plus fix operating costs.

Source: Following Grewe (2005).

3 Market Based Use of Storage Capacities – A Model

The overview of German storage facilities and the corresponding operator in the previous section reveals a significant share of incumbents in the market for natural gas storage. In this section we test the hypothesis that the usage strategy observed at Dötlingen (a large depleted gas field operated by BEB) is not closely related to perfect or liquid market mechanisms. To evaluate the usage strategy of the facility actual storage decisions have to be compared with some benchmark. Therefore, we proceed in three steps: First, we define the storage optimization strategies. Second, we calculate the behavior given the defined strategy (benchmark). Finally, we compare the benchmark behavior with the observed strategy.[4]

[4] This section draws on work where we compare storage operation in the UK and Germany (Zachmann and Neumann, mimeo). The basic idea is that a competitive market such as the UK will use natural gas storage according to the theory of storage.

The benchmark that we want to compare with the observed storage decisions is the "perfect market" strategy. It is characterized by full price taking behavior of the storage customer. Therefore, the profit function can be written as $\Pi = \sum \triangle v_t p_t - c(\triangle v_t p_t, V_t)$ where $\triangle v_t$ is the storage decision, p_t the price and $c(\triangle v_t p_t, V_t)$ are associated cost at time t.

In a second step, we calculate the storage customer's strategy. This is done by maximizing its profit with respect to stochastic prices, a non-linear cost function and non-linear constraints. Before presenting the algorithm the core components of the profit optimization are introduced.

A storage facility is essentially characterized by three factors: the injection rate, the withdrawal rate and the working gas volume (maximum less minimum volume). We consider the maximum and minimum observed storage level as best proxy for the real upper and lower constraints. This approach has the advantage that not only the purely technical constraints are included but also non-technical obligations e.g., such as strategic reserves in case of bad weather, are incorporated.

Maximum injection and withdrawal rates are more difficult to deduce as those generally depend on the storage level. If, for example, a storage facility is close to its capacity limit it is technically more diffcult to inject natural gas and if almost empty, withdrawal rates decline. Taking this behavior into account we estimate the corresponding relationship using observed data. Therefore, we first extract the maximum injection and withdrawal speed for each storage level. Then we estimate the relationships between maximum injection rate and storage level, and between maximum withdrawal rate and storage level using a polynominal (see Figure 2).

Figure 1: Observed injections/withdrawals (green) and corresponding estimated maximum injection (red)/withdrawal (blue) rates

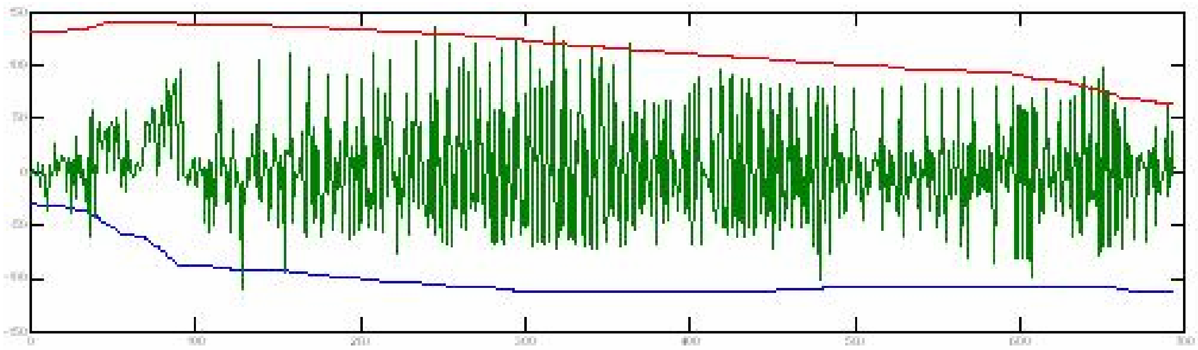

The cost function consists of four components: fuel cost, injection cost, withdrawal cost and storage cost. *Fuel cost* (fc) is a symmetric percentage (φ) of injections/withdrawals used for injection/withdrawal. *Used fuel* is valued at current prices and the fuel cost component is written as $fc_t = \varphi \triangle v_t p_t$ ($\triangle v_t$ is the storage decision). *Injection/withdrawal costs* (ic/wc) are additional cost depending only on injected/withdrawn volumes: $ic_t = \mu_i \triangle v_t$ if $\triangle v_t > 0$ and $ic_t = \mu_w \triangle v_t$ if $\triangle v_t < 0$. Finally, *storage cost* is the cost for holding gas in store: $sc = \varsigma V_t$ The assumptions for the four cost components are taken from Simmons (2000) and presented in Table 2.

Table 2: Assumptions on cost components

Fuel used at each injection/withdrawal (φ)	1%
Cost associated to each injection (μ_i)	0,02 USD/MMBtu
Cost associated to each withdrawal (μ_w)	0,02 USD/MMBtu
Cost for holding natural gas in store (ς)	0,40 USD/MMBtu

Source: Simmons (2000).

To optimize its day-to-day injection/withdrawal decision a storage customer needs to have some knowledge on future price developments. Futures and forward prices should represent the best guess of future spot price development that can be represented by the so-called price forward curve (PFC). This PFC is calculated based on current futures prices. While weekly or monthly futures are traded near to spot month, seasonal or annual futures are traded for longer time horizons. Thus, the PFC is calculated by smoothing and adding seasonalities (see Figure 2).

Nevertheless it is clear to all market participants that future spot prices will generally deviate from the PFC. Therefore we assume natural gas spot prices to be stochastic in the short run while reverting to the corresponding PFC in the long run. The related parameters (mean reversion speed, volatility) are estimated using real data from the Dutch Title Transfer Facility (TTF).[5]

[5] In an application to the UK market Ha¤ et al. (2008) find a non-linear effect of storage on the relation of spot and futures prices.

Figure 2: Three price forward curves for TTF

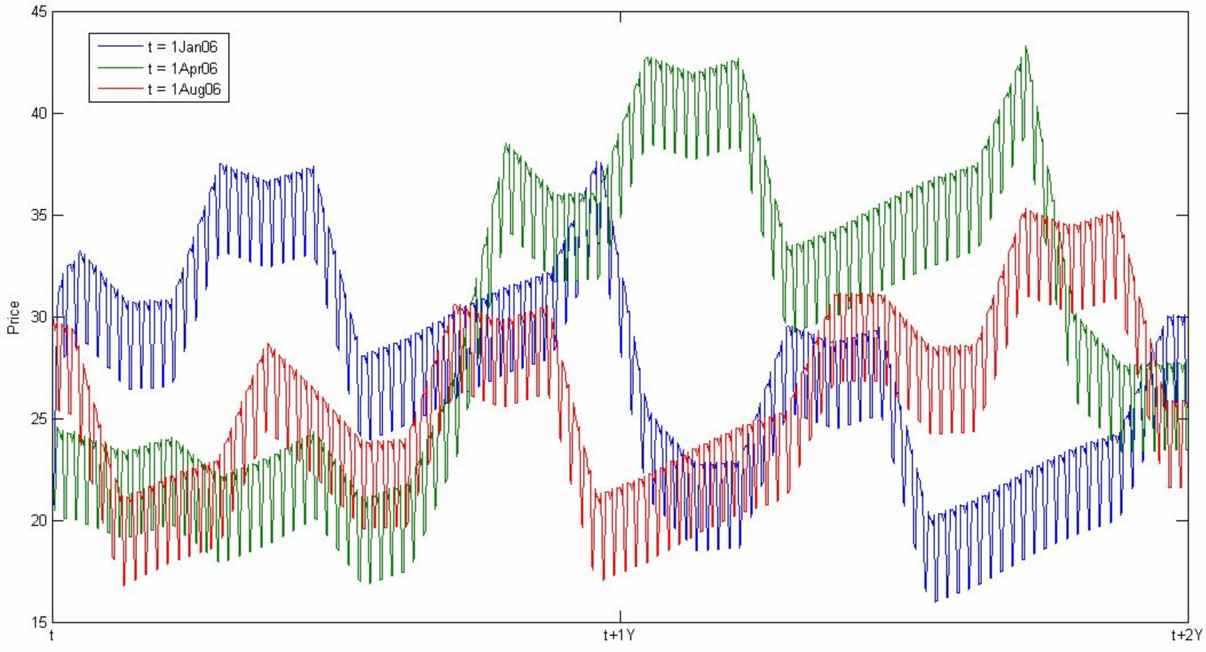

Different optimization algorithms for maximizing the profits from natural gas storage usage have been proposed in the literature. Generally two approaches can be distinguished. While solving a Bellman equation provides a closed form solution given certain price generating functions, Monte Carlo simulations are very flexible with respect to constraints and price models but have no analytic solution.

To cope with a non-standard price function (reversion to moving mean) as well as nonlinearities in constraints and cost we follow Boogert and de Jong (2006) applying a Least Square Monte Carlo approach to natural gas storage contracts. Since identifying the optimal storage strategy is comparable to locating the exercise date of American options, Boogert and de Jong (2006) apply an option valuation algorithm proposed by Longstaff and Schwartz (2001). The general idea of the concept is to optimize storage usage decisions backwards in time using a discrete (daily) time grid, a discrete volume grid and n simulated price paths. The volume grid stretches from minimum to maximum storage level at equal distance volume steps: Vol_{min} : VolStep : Vol_{max}. These volume steps are defined to approximately represent a tenth of the daily decision spectrum (i.e., the difference of maximum injections and maximum withdrawals). Thus, at each day and volume combination, around ten different decisions are possible.

Time-values for a discrete set of allowed strategies are compared at each decision making point. Consequently, we first define a termination date and the payoff function at this date. We set the termination date $T = t_0 + 365$ (that is one year after the start date), and the payoff function at T is defined depending on the volume in storage at termination date (Vol_T). If the volume exceeds a desired level (Vol^*) the payoff is zero. We assume that a storage customer has to pay a punishment of double the time-value of the missing volume if the critical level Vol^* is undercut. This yields:

$$Payoff_T^i(p_T^i, Vol_T) = \begin{cases} 0 & Vol_T \geq Vol^* \\ -2(Vol^* - Vol_T)\,p_T^i & Vol_T < Vol^* \end{cases} \tag{1}$$

Therefore, the value of the storage contract $Value_T = Payoff_T(p_T, Vol_T)$ depends on the volume stored (Vol_T) and the price (p_T) at the termination date. The value is calculated for all simulated price paths and all allowed discrete volume levels. Departing from last day's storage values we then move back a day and calculate the optimal decisions for all price paths and allowed volumes. We define the value in *T-1* according to the current payoff of the optimal injection/withdrawal decision (*optDecision*), the discounted future value resulting from the volume after the optimal decision as well as the cost of the optimal decision:

$$\begin{aligned} Value_{T-1}^i\left(Vol_{T-1}, optDecision_{T-1}^i, p_{T-1}^i\right) = {} & p_{T-1}^i \times optDecision_{T-1}^i \\ & + discount\_rate \times Payoff_T^i(p_T^i, Vol_{T-1} + optDecision_{T-1}^i) \\ & - Cost\left(optDecision_{T-1}^i\right) \end{aligned} \tag{2}$$

where δ is the discount factor.

The optimal decision given each allowed (discrete) volume level is derived by maximizing the current storage value $Value^i_t(Vol_t; optDecision^i_t; p^i_t)$ with respect to the allowed discrete steps of $optDecision^i_t$. The volume level and the price path are the sole driver of the storage value and the $optDecision^i_t$ can thus be determined. According to equation 2 the value for each price-path volume-level combination is calculated. Similarly, the corresponding storage values for all points in time ($t = 1: T-1$) can be determined:

$$\begin{aligned} Value_t^i\left(Vol_t, optDecision_t^i, p_t^i\right) = {} & p_t^i \times optDecision_t^i \\ & + discount\_rate \times Value_{t+1}^i(p_{t+1}^i, Vol_{t+1}) \\ & - Cost\left(optDecision_t^i\right) \end{aligned} \tag{3}$$

where δ is the discount factor.

To address the fact that storage customers can not know the exact spot price development in advance, *n* price paths are simulated according to the PFC and the estimated stochastic behavior. Furthermore, it has to be taken into account that the future storage value (that is needed to calculate the current optimal decision) is unknown to a storage customer as it depends on the future price development. Therefore, the methodology proposed by Boogert and de Jong (2006) implies a "Least Square Step" where, based on available information (current prices), the future storage value is estimated. The idea of this "Least Square Step" is to mimic the storage customers belief on the future storage value by regressing the future values ($Value^i_{t+1}(p^i_{t+1}, Vol^i_{t+1})$) at each volume level and each price path on the current prices of this path p^i_{t+1}. The forecast of the future value is then given by $\hat{V}^i_{t+1} = \hat{\beta} p^i_t$.

The algorithm defined in equation 3 can now be iterated backwards from $t = T-1$ to $t = 1$. The number of allowed storage decisions should be at least equal to three (injection at maximum capacity, withdrawal at maximum capacity, no operation) while a finer grid would allow for more precise results. As the quantity of storage levels is proportional to the resolution of the grid, the number of allowed levels is substantial. Furthermore, a high number of price paths is desirable to obtain reliable results. And finally, the observation period should at least contain one full cycle (365 days). Consequently, computation time becomes an issue and has to be carefully balanced with precision.

To make the benchmark strategy comparable with the observed strategy it is necessary that the benchmark strategy algorithm departs from the information set that was available to the actual storage customers. Therefore, the "ex post optimal strategy" (which implies perfect foresight of future prices) is a misleading benchmark for the observed strategy. Given imperfect price foresight (i.e. price simulation), the optimization has to be rerun for every point in time to assure that the price information are updated. Thus, the optimal strategy under imperfect price foresight is calculated in 365 subsequent rolling windows, each of which containing an updated 365 day price forecast ($t_{0,1}$=1:365; $t_{0,2}$=2:366,…, $t_{0,365}$=365:730).

Optimizing the strategy given imperfect foresight, the question arises which starting and end volume to assume for each run. Essentially, two "ex ante optimal strategies" are available: one, departing from the past optimal decisions and one departing from past observed decisions. Starting each optimization from the observed level would on the one hand assure that each days information set is most accurately reflected. But on the other hand it implies that the cumulated decisions might surpass the technical constraints. Starting from the past optimal value, by contrast, assures that the cumulated decisions can be compared to the observed volumes. Since our analysis focuses on day-to-day injection/withdrawal decisions we optimize according to the observed initial volumes. Furthermore, final volumes are deduced from the observed data. As the optimization at the last day requires the end volume

365 days later we require 730 days of observed storage volumes to obtain the optimal strategy for 365 days.

The Dötlingen storage site is operated by a significant market player: BEB is partly owned by ExxonMobil which is an important natural gas trader in Europe. Flow data for Dötlingen are available from October 2005 onwards. Table 3 summarizes the main characteristics of the storage site.

Table 3: Key characteristics of Dötlingen storage site

Operator	BEB
Storage Type	Depleted gas field
Used working capacity in GWh (technical max.)	8.847 (17.899)
Max. injection in GWh/day (technical max.)	109 (217)
Max. withdrawal in GWh/day (technical max.)	135 (217)
Available data	Daily aggregated injections and withdrawals

Source: Operator's website.

To understand the observed usage of the natural gas storage facilities we use prices at the Dutch natural gas exchange (TTF) which can be considered as reference and are shown for the time period under consideration in Figure 3.

Figure 3: Observed storage level and prices

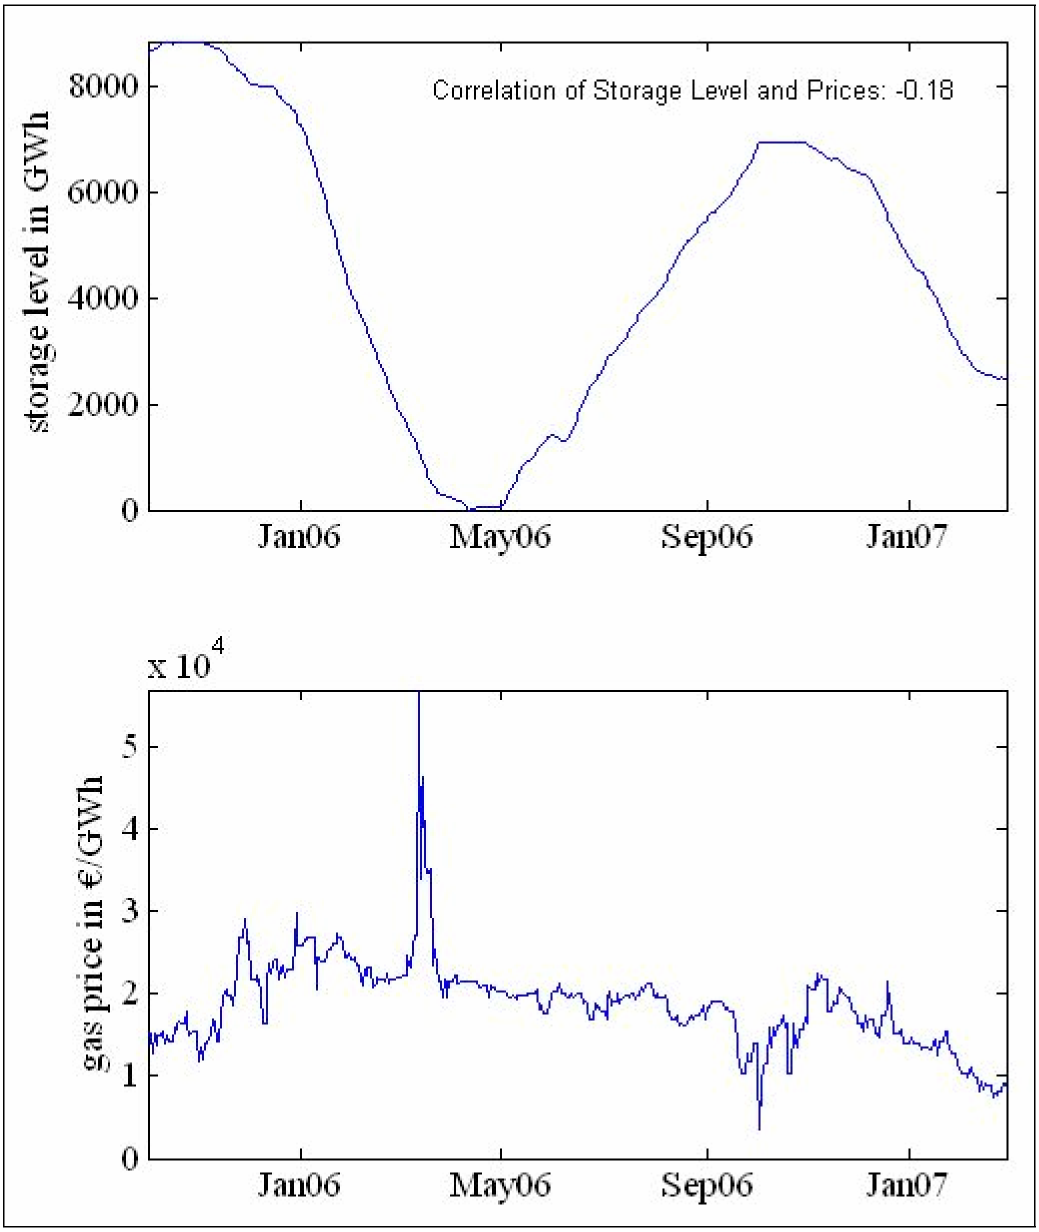

Applying the described algorithm provides the desired benchmark that can be compared to the observed data. The correlations of the observed storage flows with the corresponding benchmark (i.e., optimal decisions under imperfect price foresight) for the Dötlingen storage facility is 8%. This is a rather low degree of correlation when taking into account that the benchmark strategy is correlated with the perfect price foresight strategy at 21%. Even by international standards the explanatory power of the benchmark strategy for the observed flows is low since the equivalent for the Rough storage facility (UK) is three times bigger (23%) (Zachmann and Neumann, 2008). The low correlation of the observed flows and the benchmark indicate that natural gas injection and withdrawal decisions at Dötlingen are not based on the profit-maximizing behavior of a small player in a liquid market.

Table 4: Correlations

Correlation of observed flows with benchmark case	0,08
For comparison	
Correlation of benchmark case flows with perfect foresight decisions	0,21
Correlation of observed flows with benchmark case at Rough (UK)	0,23

There are several explanations for this finding

1. **Insufficiency of TTF**: Low liquidity of the spot and futures market as well as transport cost and potential congestion between TTF and Dötlingen might cause prices at TTF not to be a good proxy for the true (but unknown) natural gas price at the Dötlingen entry/exit point.
2. **Technical considerations:** Storage operators could take into account additional technical constraints not considered in the benchmark strategy. For example, storage can serve to provide short-run balance of supply and demand, regulate the pipeline pressure, level injections/withdrawals in the system (e.g., LNG tanker arrivals).
3. **Strategic reserve:** Storage facilities might serve as a physical hedge and therefore be operated more smoothly than implied by pure arbitrage considerations.
4. **Cost of storage operations:** Optimization of the storage facility may not be based on the variable cost of storage operations. Storage contracts usually allocate a share of variable costs to flat-rate components.
5. **Exercise of market power:** Storage operators or customers can withhold natural gas in periods of low price elasticity. This creates potentials for strategic behavior for a sufficiently big player.

Given that possible explanations (2), (3), (4) and (5) apply to both the UK and the German market (although potentially at a different degree) our results show that a large part of the observed difference (of correlations) is due to the absence of a liquid and transparent German natural gas market. Therefore, the importance of a short-term trading and the role of natural gas storage are interrelated.

4 Conclusions

In this paper we have argued that the role of natural gas storage and competitive usage of these facilities is gaining momentum in the process of market restructuring. Whereas most of the existing storage sites in Germany are owned by large companies active in long-distance pipeline transportation, investment in new capacities is also coming forward from market entrants. The development of a liquid trading point in the German pipeline system and the increasing interconnectivity with adjacent countries will further spur the development of new sites.

We have developed a model which uses real data for injection and withdrawal rates of a storage site in Germany, which is i) favorable located, ii) owned by a big player (BEB), and iii) has published utilization rates. Furthermore, we use natural gas futures prices from the Dutch Title Transfer Facility to compare a competitive usage of the storage site with observed behavior. The results show that the operation of this particular site in Germany does not follow the theory of storage. Even though there exist potential reasons to explain the discrepancy, our model also shows that German storage sites are not operated on a purely profit-maximizing behavior. This leads to the conclusion that the development of a competitive storage market ("merchant storage") is far from being completed. Hence, natural gas storage should be regulated and considered in a European context and the importance, both in economic and supply security terms, picked up by policy makers and other decision makers.

5 References

Boogert, Alexander and Cyriel de Jong (2006), "Gas Storage Valuation Using a Monte Carlo Method", BWPEF 0704, December 2006, Birkbeck, University of London.

Chaton, Corinne, Anna Creti, and Bertrand Villeneuve (2008), "Some Economics of Seasonal Gas Storage", Energy Policy, 36(11), 4235-4246.

Deaton, Angus and Guy Laroque (1996), "Competitive Storage and Commodity Price Dynamics", Journal of Political Economy, 104(5), 896-923.

Dencerler, C., Z. Khokher and T. Simin (2005), "An Empirical Analysis of Commodity Convenience Yields", Working Paper, June 2005, University of Ontario.

EIA (2004), "The Basics of Underground Natural Gas Storage", Washington.

Fama, Eugene F., and Kenneth R. French (1987), "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums and the Theory of Storage", Journal of Business, 60(1), 55-74.

Grewe, Joachim (2005), Auswirkungen der Liberalisierung auf die Erdgasspeicherung. Eine ökonomische Analyse für den deutschen Erdgasmarkt, Sonderpunkt Verlag Münster.

Khan, Sagib A., Zeigham Khoker and Timothy T. Simin (2005), "Inventories, Convenience Yields and Mean Reversion", Working Paper, October 2005, University of Western Ontario.

Longstaff, Francis A. and Eduardo S. Schwartz (2001), "Valuing American Options by Simulation: A Simple Least-Squares Approach", Review of Financial Studies, Oxford University Press for Society for Financial Studies, 14(1), 113-47.

Serletis, Apostolis, and Asghar Shahmoradi (2006), "Futures Trading and the Storage of North American Natural Gas", OPEC Review, 30(1), 19-26.

Simmons & Company International Energy Industry Research (2000), "Underground Natural Gas Storage", Houston, TX.

Susmel, Raul, and Andrew Thompson (1997), "Volatility, Storage and Convenience: Evidence from Natural Gas Markets", The Journal of Futures Markets, 17(1), 17-43.

Uria, Rocia, and Jeffrey Williams (2007), "The .Supply-of-Storage for Natural Gas in California", The Energy Journal, 28(3), 31-50.

Wei, Song Z.C. and Zhen Zhu (2006), "Commodity Convenience Yield and Risk Premium Determination: The Case of the US Natural Gas Market", Energy Economics, 28(4), 523-535.

Williams, Jeffrey C. and Brian D. Wright (1991): Storage and Commodity Markets, Cambridge University Press.

Wright, Brian D., and Jeffrey C. Williams (1982), "The Economic Role of Commodity Storage", The Economic Journal, 92(367), 596-614.

Wright, Brian D., and Jeffrey C. Williams (1989), "A Theory of Negative Prices for Storage", Journal of Futures Markets, 9(1), 1-13.

Zachmann, Georg and Anne Neumann (2008), "Storage Stories - Operation of Natural Gas Storage Facilities in Germany and the UK", Mimeo, TU Dresden.