

Schlüter, Stephan; Davison, Matt

Working Paper

Pricing an European gas storage facility using a continuous-time spot price model with GARCH diffusion

IWQW Discussion Papers, No. 02/2010

Provided in Cooperation with:

Friedrich-Alexander University Erlangen-Nuremberg, Institute for Economics

Suggested Citation: Schlüter, Stephan; Davison, Matt (2010) : Pricing an European gas storage facility using a continuous-time spot price model with GARCH diffusion, IWQW Discussion Papers, No. 02/2010, Friedrich-Alexander-Universität Erlangen-Nürnberg, Institut für Wirtschaftspolitik und Quantitative Wirtschaftsforschung (IWQW), Nürnberg

This Version is available at:

<https://hdl.handle.net/10419/30184>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IWQW

Institut für Wirtschaftspolitik und Quantitative
Wirtschaftsforschung

Diskussionspapier
Discussion Papers

No. 2/2010

**Pricing an European Gas Storage Facility using a
Continuous-Time Spot Price Model with GARCH
Diffusion**

Stephan Schlüter
University of Erlangen-Nuremberg
Matt Davison
University of Western Ontario, Canada

ISSN 1867-6707

Pricing an European Gas Storage Facility using a Continuous-Time Spot Price Model with GARCH Diffusion

STEPHAN SCHLÜTER, MATT DAVISON

*Department of Statistics and Econometrics, University of Erlangen-Nuremberg,
Lange Gasse 20, Nuremberg, D-90403*

*Department of Applied Mathematics, and of Statistical & Actuarial Sciences, The
University of Western Ontario, London, Ontario, N6A 5B7, Canada*

Email: stephan.schlueter@wiso.uni-erlangen.de, mdavison@uwo.ca

**This work was done while visiting The University of Western Ontario
under support of DAAD, MITACS & NSERC.**

Abstract

In this article we present both a theoretical framework and a solved example for pricing an European gas storage facility and computing the optimal strategy for its operation. As a representative price index we choose the Dutch TTF day-ahead gas price. We present statistical evidence that the volatility of this index is time-varying, so we introduce a new continuous-time model by incorporating GARCH diffusion into an Ornstein-Uhlenbeck process. Based on this price process we use dynamic programming methods to derive partial differential equations for pricing a storage facility.

As an example we apply our methodology to a storage site located in Epe at the German-Dutch border. In this context we investigate the effects of multiple contract types, and perform a sensitivity analysis for all model parameters. We obtain a value surface displaying the properties of a financial straddle. Both volatility and mean reversion influence the facility value – but only around the long-run mean of the gas price. The terminal condition, which includes information about the contract provisions, is of importance if it contains e.g. penalty terms for low inventory levels. Otherwise its influence is diminishing for increasing lease periods.

JEL-Classification: C31, C61.

Keywords: TTF gas price; GARCH diffusion; natural gas storage; dynamic computing.

1 Introduction

Despite the recent economic downturn, European natural gas markets recorded a growth in traded volume by a significant 57 % in 2008 (IEA, 2009). They are, in contrast to their North American counterparts, not yet fully liberalized, but deregulation is in progress and several important steps to open the markets have been taken. Diverse gas trading hubs have been established across Europe. The most important are the British National Balancing Point (founded in 1996), the Belgian hub in Zeebrugge (founded in 2000) and the Dutch Title Transfer Facility TTF (founded in 2003). In Germany, the total traded volume is about equal to that traded on the TTF. The German market is still divided into different geographic areas (see IEA, 2009), and so the TTF day-ahead index is commonly used as a reference for short-term prices. The Dutch exchange recently established itself as the main continental trading hub (with physical delivery within the Dutch natural gas grid), and comprises both spot and future markets. Liquidity has improved substantially since 2003, especially on the day-ahead market. In particular, Figure 2 displays the recent absence of price spikes. In addition, the volume traded on the TTF has more than doubled from 2007 to 2008, as shown in Figure 1 and by IEA (2009).

Independent of liquidity issues, there is significant fluctuation in both long- and short-term gas prices. Gas is still widely used for heating purposes, so many gas products show a distinct summer-winter seasonality. Gas is also increasingly used for power-production (see Kiesel & Herkner, 2008). Although the fuel is expensive, gas turbines are very flexible and are therefore mainly switched on to cover the high demand of peak hours. This links short-term gas prices to the highly volatile electricity market and increases their fluctuation. In contrast to electricity, natural gas can be stored and adequate storage facilities can be utilized to reduce volatility. In Germany, gas is stored in depleted oil & gas reservoirs or salt caverns. In areas where this is not possible, aquifers are developed as expensive but adequate substitutes. Storing gas in liquid state (LNG gas) is also becoming an issue as technology advances (see e.g. IEA, 2009).

Storage facilities play an important financial, social and political role in Europe. They reduce the seasonal gas price pattern and guarantee the supply of gas for winter heating (see Figure 3). They also serve as strategic reserves for events like the recent dispute between Russia and the Ukraine about the price for natural gas, during which the amount of gas delivered to Germany was significantly reduced (IEA, 2009).

The way a gas storage facility can be used strongly depends on its operational characteristics. Each type requires a certain amount of base gas that remains unretrievable in the facility and is required to maintain a certain level of pressure. More


important are the injection/withdrawal and cycling rates. Both the injection and the withdrawal rates depend on the pressure inside the storage which in turn depends on the fill level. The cycling rate quantifies how many times a year gas can be withdrawn or injected. Depleted oil reservoirs have a low injection/withdrawal rate. They are operated on a single annual cycle and used mainly for damping seasonal patterns (for more details, see Thompson et al., 2009). Salt caverns allow higher injection/withdrawal rates and can be operated multi cycle. This makes them adequate for hedging spot market transactions or using them in a bundle with gas turbines. With the increase of gas-fired power-plants the demand for these kinds of storage is rising and capacities have been extended over the last years (see IEA, 2009).

In this new scenario of liquid short-term gas markets the storage facility becomes a more active element. Instead of just buffering seasonality, it can also be used as a physical hedge for options on the gas or electricity price. Such strategies offer more profit than merely transacting on the forward market, but are far more complex and more risky. These arguments drive the need for powerful methods to determine the optimal strategy for operating storage facilities and their expected storage value. In the recent years several attempts were made to answer this question (see e.g. Thompson et al. 2009, Boogert and de Jong, 2006, or Bjerksund, 2008), however still certain questions remain unanswered. Existing work does not capture gas prices with dynamic volatility which is, however, a key feature of TTF spot prices. Within this paper we intend to close this gap and find in Section 2 that using a GARCH diffusion for volatility is adequate. In addition, various storage level arrangements have different terminal provisions, and existing works ignore these distinctions. We also adress this point in the current work.

For pricing the storage facility itself, real options theory suggests computing the facility value as the expected value over all (discounted) future cash flows (see Dixit & Pindyck, 1994). Monte Carlo simulations or the binomial tree method can be applied as numerical tools, but both have some shortcomings, which leaves us with partial differential equations (PDEs). A pricing algorithm which respects dynamic volatility is developed. We do also analyze the effects of other price parameters on the storage value and especially see a significant influence of the size of mean reversion. Our algorithm describes the scenario of a limited-time storage lease. Costs of withdrawal/injection as well as lease period and terminal condition are core elements of any lease contract. To show the effects of contract terminal conditions, we compute the sensitivity of lease values to some common contract specifications.

The paper is structured as follows: The basis for every storage pricing algorithm is an adequate spot-price model. In Section 2 we suggest such a model for recent TTF data. Subsequently an overview of different storage pricing methods is given

Figure 1: The Volume at TTF



The solid (dashed) line represents the traded (physical) volume (Gas Transport Services, 2010).


in Section 3. In Section 4, we apply our method to the storage facility in Epe. The focus of our analysis lies on evaluating the influence of dynamic volatility and of different contract features on the storage value. A brief summary concludes the paper.

2 The TTF Day Ahead Gas Price

Recent literature suggests different models for a short-term gas price process. Thompson et al. (2009) propose using a jump-diffusion process. In theory they allow for dynamic volatility, however, eventually, when deriving the gas storage value, they require the volatility to be constant. Boogert & de Jong (2006) propose a mean reverting diffusion with constant volatility. Chen & Forsyth (2010) do the same and add seasonality to the long-term equilibrium price. Bjerksund et al. (2008) criticize the simplicity of the mentioned models and construct a complex model based on a principal component analysis. Cartea & Williams (2008) refine the basic bivariate model for commodity prices of Schwartz & Smith (2000) by adding a seasonal component. This approach splits the log-prices into a long-run Brownian motion with drift and a short-term mean-reverting term. The individual processes are correlated.

Based on a TTF day-ahead time series data set (provided by the APX Group) we investigate whether one of the above models fits the data. In Figure 2 we clearly see a structural break in about 2006 after which large spikes disappear. Therefore we

Figure 2: The TTF Day-Ahead Price



consider a regime-switching model in the form of a jump component added to the process as inadequate. The approach of Bjerksund et al. (2008) is too complicated for most practical applications and – in contrast to Cartea & Williams (2008) who analyzed the UK day-ahead gas price index – we cannot find a significant seasonal pattern in the day-ahead prices. Moreover, the time series is too short to effectively realize the advantages of their two-factor model.


So, as we excluded almost all mentioned approaches, we decide to stick to the world of stochastic models and begin with fitting a simple mean reverting Ornstein-Uhlenbeck process to the data (similar to Thompson et al., 2009). We find a mean reversion and looking at the identified residuals we see that volatility is not constant over time. One way to achieve this for discrete time series is to use generalized autoregressive conditional heteroskedastic (GARCH) models in which the volatility follows an autoregressive process (Bollerslev 1986 and Engle 1982). Under weak conditions, Drost & Werker (1996) derive a continuous time limit of GARCH. We make use of their result to formulate this model:

$$\begin{aligned} dP_t &= \lambda(\mu - P_t)dt + \sigma_t dW_t^{(1)}, \\ d\sigma_t^2 &= \theta(\omega - \sigma_t^2)dt + \sqrt{2\eta\theta}\sigma_t^2 dW_t^{(2)}, \end{aligned} \tag{2.1}$$

where $\lambda \in \mathbb{R}$ denotes the long-term rate of mean reversion and $\mu \in \mathbb{R}$ denotes the long-term mean. The stochastic components $W^{(1)}, W^{(2)}$ are independent standard Brownian motions, $\omega > 0, \theta > 0$ and $\eta \in (0, 1)$. Using maximum likelihood estimation we find the following parameterisation to best fit the TTF day-ahead prices starting from Sept 2nd 2006 and ending Sept 9th 2009:

$$dP_t = 0.01(17.76 - P_t)dt + \sigma_t dW_t^{(1)}, \quad d\sigma_t^2 = 0.01(5.26 - \sigma_t^2)dt + 0.31\sigma_t^2 dW_t^{(2)}.$$

Figure 3: The Monthly German Gas Storage Balance



Source: FMET (2010).

We see that both the price itself and the volatility have a quite small mean reversion. The long-term mean lies around 18 EUR/MWh with a long-term volatility of little more than 2 EUR/MWh.

For computing the storage price, a risk neutral version of (2.1) is required. The martingale measure equivalent to $(W^{(1)}, W^{(2)})$ is again a Brownian motion, but we must adjust the long-run mean (of both price and volatility) by constant risk premia. A common method to quantify these premia involves estimating forward prices. The TTF market, while more liquid than others on the European continent, nonetheless provides insufficient data for this estimation. Guo (1996) and Stein & Stein (1991) solve this problem by simply assuming the premia to be zero. We address this issue in Section 4.3 where we study the effect of reducing both long-run means on the storage value.

3 Pricing a Gas Storage Facility

The operational strategy of a storage is to a large part determined by its injection and withdrawal rates. The more flexible a storage the bigger the variety of possible strategies. In Germany so far mainly forward-based strategies, which aim to benefit from the seasonal pattern, were applied – as shown by the monthly aggregated storage balance depicted in Figure 3. This is, however, suboptimal for storages with high injection/withdrawal rates. At the TTF forward market the tradable months or quarters are highly correlated and volatility is relatively low. The spreads between different months are therefore relatively small, as are the profits made by establishing static hedges between forwards of different maturities. Profits may be increased by dynamic hedging on the spot market, where a larger volatility is observed.

Beyond that it is possible – if physical properties allow – to obtain even more profit from reacting to short-term price fluctuations (e.g. in form of options trading) and using the storage as a physical hedge. This strategy, however, is more risky and far more complex than a future-based one. On the other hand – and this holds especially for emerging markets like the gas market – if there is no liquid forward market, it might be impossible to operate the optimal forward-based strategy.

For these reasons we decide to price a storage based on spot-market operations. For this purpose we provide a brief introduction into different valuation methods. Subsequently we derive our valuation method and address the numerical techniques which arise.

3.1 A Review of Methodologies

There are three different methods of gas storage valuation: Monte-Carlo simulation, binomial trees and PDE techniques. The advantage of the first is that the price process and the optimal strategy are separated. This makes it easier to test and use different price models, especially when applying other stochastic factors than a Brownian motion or Poisson jumps. However, simulating prices leads to instabilities in the storage value, which makes it hard to compute the effect of instantly injecting/withdrawing gas on the future value. Therefore it is quite difficult to find a stable optimal control strategy with Monte-Carlo simulation. Advanced methods like Least Squares Monte Carlo, applied e.g. by Boogert and de Jong (2006), solve this problem, but this technique again evokes further questions. One is the problem of choosing a set of basis functions for the regression. This is difficult for higher dimensions and is a source for further approximation errors.

The binomial tree method can be considered as a numerical translation of forward differencing, which is a technique to solve PDEs (see e.g. Hull, 1999). Its main problem is, that it requires large computational resources. This demand is dramatically rising when increasing the spatial dimension. Besides that it is hard to incorporate flux limiters, which are necessary to cope with the arising hyperbolic equations.

This leaves us with the PDE technique. If implemented adequately these relatively fast standard techniques are able to handle hyperbolic equations. Furthermore, when starting from the Bellman equation, we have a necessary and sufficient condition for an optimal solution. However, as the derivation involves Ito's Lemma, we are restricted to Brownian motion based processes (possibly including Poisson jumps). In the subsection below we will use this method to derive a gas storage pricing scheme based on the price model found in Section 2.

3.2 Gas Storage Pricing via Dynamic Computing

We now generalize the approach of Thompson et al. (2009) to dynamic volatility and derive the required PDEs in order to price a gas storage. We first define the parameters used in the sequel: Let P be the current natural gas price, T the terminal time and I the current amount of stored (working) gas with boundaries I_{min}, I_{max} . The central variable is the control $c(I, P)$, which is the amount of gas currently released ($c > 0$) resp. injected ($c < 0$). The boundaries of $c(I, P)$ vary with I because of the relation between I and the internal reservoir pressure (see B), and are denoted by $c_{min}(I)$ and $c_{max}(I)$. Eventually we introduce $a(I, c)$ which serves as a parameter to collect the different costs occurring while operating the storage, i.e. costs for injection/withdrawal or costs due to leakages. Constant costs such as the general storage capacity and connection lease may be ignored as they do not influence the optimization process. Yet this definition still offers various options. Thompson et al. (2009) used

$$a(I, c) = \begin{cases} 0 & \text{for } c \geq 0. \\ K & \text{for } c < 0, \end{cases} \quad (3.1)$$

with $K > 0$. We can also design the costs of injection as a linear function with a factor $h > 0$ and

$$a(I, c) = \begin{cases} 0 & \text{for } c \geq 0, \\ -h \cdot c & \text{for } c < 0. \end{cases} \quad (3.2)$$

Both the injection and the withdrawal of gas causes costs, but it is a common design in practice to charge all costs when storing in (see e.g. RWE, 2010).

Having chosen a definition for each variable, the current value of the gas storage is the discounted cash flow of all future activities until time T . The cash flow in every time step dt with an optimal control c is simply $(c - a(c, I))Pdt$. Given a risk neutral measure $E[\cdot]$ for both the price and the volatility process and a risk free interest rate $\rho \in \mathbb{R}$ we can write

$$V(P, \sigma^2, t = 0, I) = \max_{c(P, I)} E \left[\int_0^T e^{-\rho\tau} (c - a(I, c)) P d\tau \right] \quad (3.3)$$

$$\text{s.t. } c_{min}(I) \leq c \leq c_{max}(I).$$

The interest rate ρ may also fluctuate in time provided it is uncorrelated with the price process. For the price process P we use (2.1), and the change of I with

time is simply the negative control c in time t , i.e. $dI = -c \cdot dt$. Now having introduced all variables we can start deriving the PDEs. We use methods from dynamic programming and follow the steps proposed by Thompson et al. (2009). We set out by identifying the storage value at time t as

$$V(P, \sigma^2, t, I) = \max_{c(P, I, t)} E \left[\int_t^T e^{-\rho(\tau-t)} (c - a(I, c)) P_\tau d\tau \right]. \quad (3.4)$$

We now split the above integral by introducing the time step dt :

$$\begin{aligned} & \max_c E \left[\int_t^{t+dt} e^{-\rho(\tau-t)} (c - a(I, c)) P d\tau + \int_{t+dt}^T e^{-\rho(\tau-t)} (c - a(I, c)) P d\tau \right] \\ = & \max_c E \left[\int_t^{t+dt} e^{-\rho(\tau-t)} (c - a(I, c)) P d\tau + e^{-\rho dt} V(P + dP, \sigma^2 + d\sigma^2, I + dI, t + dt) \right]. \end{aligned}$$

Following this procedure, we arrive at the Bellman equation

$$V = \max_c E \left[(c - a(I, c)) P dt + e^{-\rho dt} V(P + dP, \sigma^2 + d\sigma^2, I + dI, t + dt) \right].$$

It is significant that we obtain a Bellman equation, as it can be shown that solving this equation is equivalent to finding an optimal solution to the original problem (Bellman, 1957). A first step towards a solution is to use an extended version of Ito's lemma. We get

$$\begin{aligned} V &= \max_c E \left[(c - a(I, c)) P dt + (1 - \rho dt) V + (1 - \rho dt) \left(V_t + \frac{1}{2} \sigma^2 V_{PP} \right. \right. \\ &+ \left. \frac{1}{2} 2\lambda\theta (\sigma^2)^2 V_{\sigma^2\sigma^2} + \eta(\mu - P) V_P + \theta(\omega - \sigma^2) V_{\sigma^2} - V_I dI \right) dt \\ &+ \left. (1 - \rho dt) \left(\eta(\mu - P) V_P dW^{(1)} + \theta(\omega - \sigma^2) V_{\sigma^2} dW^{(2)} \right) \right]. \end{aligned}$$

Now we eliminate every term that goes faster to zero than dt (except $-\rho dt V$) and thus

$$\begin{aligned} V &= \max_c E \left[(c - a(I, c)) P dt + (1 - \rho dt) V + \left(V_t + \frac{1}{2} \sigma^2 V_{PP} \right. \right. \\ &+ \left. \lambda\theta (\sigma^2)^2 V_{\sigma^2\sigma^2} + \eta(\mu - P) V_P + \theta(\omega - \sigma^2) V_{\sigma^2} - V_I dI \right) dt \\ &+ \left. \eta(\mu - P) V_P dW^{(1)} + \theta(\omega - \sigma^2) V_{\sigma^2} dW^{(2)} \right]. \end{aligned}$$

Taking the bivariate expectation, eliminating V on both sides and dividing by dt yields (as we deal with independent Brownian motion increments $dW^{(1)}, dW^{(2)}$)

$$\begin{aligned} \max_c \left[(c - a) P - c V_I - \rho V + V_t + 0.5 \sigma^2 V_{PP} + \right. \\ \left. \lambda\theta (\sigma^2)^2 V_{\sigma^2\sigma^2} + \eta(\mu - P) V_P + \theta(\omega - \sigma^2) V_{\sigma^2} \right] = 0. \quad (3.5) \end{aligned}$$

Only two of these terms depend on c . Thus, optimizing the equation above with respect to c means solving

$$\begin{aligned} \max_c \quad & [-cV_I + (c - a(I, c))P] \\ \text{s.t} \quad & c_{min}(I) \leq c \leq c_{max}(I). \end{aligned} \quad (3.6)$$

We can interpret (3.6) as follows: It is only reasonable to sell gas if the instantly payoff is larger than the effect of reducing the inventory on the future storage value. In this case, it is optimal to withdraw the maximum amount, i.e. $c_{max}(I)$. If not, it might be reasonable to inject gas, but only if the future payoff (reduced by the costs of injection) surpasses the instantly costs. Then the maximal possible injection rate is optimal. Otherwise, we remain idle. This is called a bang-bang control structure. Having found an optimal $c = c_{opt}$ we can rewrite (3.5) as follows:

$$\begin{aligned} & (c_{opt} - a)P - c_{opt}V_I - \rho V + V_t + 0.5\sigma^2V_{PP} \\ & + \lambda\theta(\sigma^2)^2V_{\sigma^2\sigma^2} + \eta(\mu - P)V_P + \theta(\omega - \sigma^2)V_{\sigma^2} = 0. \end{aligned} \quad (3.7)$$

So, iteratively, using PDE methods, we identify in every time step first the optimal control c_{opt} and subsequently the corresponding storage value V . Yet, in order to solve the PDE, we have to define boundary conditions. The terminal condition has a substantial influence on the storage value. From (3.4) follows straight forward

$$V(\cdot, \cdot, T, \cdot) = 0. \quad (3.8)$$

This condition has a strong implication on the control variable as it promotes gas withdrawal within the proximity of T at almost every price level. A less extreme condition would be

$$V(\cdot, \cdot, T, \cdot) = \mu \cdot I, \quad (3.9)$$

i.e. any inventory at time T will be reimbursed by the long-term average price. To implement this in our pricing algorithm we merely have to modify the PDE at time $T - 1$ and replace the value at time T by the above formula. Another scenario is a fixed inventory level I_{end} at time T . A positive difference is reimbursed by the current price and a negative difference will be billed include a prespecified penalty, i.e.

$$V(P, \cdot, T, I) = \begin{cases} P \cdot (I - I_{end}) & \text{if } I > I_{end}, \\ \nu \cdot P \cdot (I - I_{end}) & \text{if } I < I_{end}, \\ 0 & \text{if } I = I_{end}, \end{cases} \quad (3.10)$$

whereby ν denotes a penalty if $\nu > 1$. RWE, for example, demands twice the market price (see RWE, 2010).

For the price boundary condition we cite Thompson et al. (2009) and use the argument that if P is very large/small the optimal strategy remains the same, so the value varies linearly in P . In other words

$$V_{PP} \rightarrow 0 \text{ for large } P, \quad V_{PP} \rightarrow 0 \text{ for } P \rightarrow 0.$$

For the volatility boundaries we draw a comparison to option pricing. We can look on the storage (as we see in Section 4.2) as a straddle-like mixture of a call and a put option. Both call and put reach their maximum value for $\sigma^2 \rightarrow \infty$, which is independent of σ^2 . In case of $V_{\sigma^2} \rightarrow 0$ the price is almost deterministic, a small change in volatility has almost no effect on the value of our facility. So we set

$$V_{\sigma^2} \rightarrow 0 \text{ for } \sigma^2 \rightarrow \infty, \quad V_{\sigma^2} \rightarrow 0 \text{ for } \sigma^2 \rightarrow 0.$$

This, however, does not mean that for other times the price must also be deterministic.

Eventually we derive our boundaries for I from the storage characteristics itself. If the storage is full, i.e. $I = I_{max}$, we cannot inject gas and thus derivatives with respect to I have to be computed from storage levels smaller than or equal to I_{max} (and vice versa for $I = I_{min}$).

Now, having the boundaries, the storage value can be computed using a simple explicit finite difference scheme (see e.g. Hull & White, 1990). We start at terminal time T and work ourselves backward in time using (3.6) resp. (3.7). However, as (3.7) is hyperbolic in I , some numerical problems arise and we address them in the section below.

3.3 Numerical Issues

For solving the PDE's described in (3.6) resp. (3.7) we follow again Thompson et al. (2009) and use an explicit finite difference scheme, as well as their suggestions to cope with instabilities. For computing both V_P , V_{PP} and V_{σ^2} , $V_{\sigma^2\sigma^2}$ we use the central difference method. Boundary values are derivated using the above defined boundary conditions.

A first order upwind differencing scheme (see e.g. Seydel, 2009) is applied to derive V with respect to I in $I = I_{min}$ resp. $I = I_{max}$. However, for other inventory levels, this method is inadequate as it is inaccurate for all but very small spatial step sizes. Therefore, Thompson et al. (2009) suggest to use a method that shows the total variation diminishing (TVD) property introduced by Harten (1983). This

concept limits diffusion in the solution space and therefore guarantees finite derivatives. A method having the TVD property is the function minmod (see A) and we use it in the sequel to compute V_I for $I \in (I_{min}, I_{max})$.

This procedure guarantees a stable solution for the storage value V , however not for the control as c_{opt} itself depends on V_I in every time step. For this problem, we can use the trick of Thompson et al. (2009): by deriving (3.7) again with respect to I we get a PDE for V_I , which is, as we have to deal with c_I , a nonlinear maximization problem. Now applying the same numerical tools as in (3.7) to this new equation we obtain a stable V_I and by inserting this in (3.6) a stable control c , which can then be used in (3.7).

Eventually we want to address the question of choosing an adequate step size for the spatial grid. After discretizing the PDE we have for each time t an equation for the value V as a linear-combination of values in $t + 1$ at different price, inventory and volatility levels. To obtain numerical stability and convergence to a solution, we have to demand that these coefficients vary within $[0, 1)$. The two crucial relations within our PDE system are


$$\begin{aligned} 0 &\leq \frac{dt}{1 + \rho dt} \left(\frac{1}{dt} - \frac{\sigma_s^2}{(dP)^2} - \frac{2\lambda\theta(\sigma_s^2)^2}{d\sigma_s^2 d\sigma_{s+1}^2} \right) < 1, \\ 0 &\leq \frac{dt}{1 + \rho dt} \left(\frac{\lambda\theta(\sigma_s^2)dt}{d\sigma_s^2 + d\sigma_{s+1}^2} - \frac{\theta(\omega - \sigma_s^2)dt}{d\sigma_s^2 + d\sigma_{s+1}^2} \right) < 1. \end{aligned} \quad (3.11)$$

Thereby dP ($d\sigma_s^2$) represents the step size of the price (volatility) and $(\lambda, \theta, \omega)$ are parameters from (2.1). We consider a price step of $dP = 0.5$ adequate and therefore the above equation reduces to a relation between $dt, d\sigma^2$ and σ^2 . Decreasing the volatility step size and increasing the maximal volatility requires a substantial reduction of the time step size. The volatility domain is limited by the time step size on the right side and by zero on the left side. Using this knowledge, we can identify all admissible pairs $(\sigma^2, d\sigma^2)$. In Figure 4 these are all pairs that lie between both lines.

4 A Realistic Scenario

Within this section we pick a gas storage facility located in Epe at the German-Dutch border and operated by the Trianel Gasspeicher Epe GmbH & Co. KG (Trianel). The gas is stored in salt caverns which guarantee a high injection and withdrawal rate. Moreover, the densely populated Ruhrgebiet is little more than 100km distant and physical access to the German gas pipeline network is guaranteed. The proximity of the TTF hub justifies use of the TTF index as the reference gas

Figure 4: The Admissible Volatility Step Size



price and allows the assumption of liquid trading in the gas spot market. Using this scenario we test our storage pricing algorithm. The focus lies thereby on the influence of the parameters – especially of volatility.

4.1 The Storage Situation

The salt caverns located in Epe are Europe’s largest gas storage facility and different operators use the caverns. Trianel itself operates several caverns with a total size of 314 million m^3 (from now on Mm^3) of natural gas (LBEG, 2010). This amount is divided into into 237M m^3 working gas and 77M m^3 cushion gas, which has to be left in the cavern to supply base pressure and so cannot be retrieved (LBEG, 2009). The maximum injection rate is 3.6M m^3/day and the maximum withdrawal rate is 7.2M m^3/day . The type of gas stored in the caverns is H-gas with an approximate caloric value of 11,5 kWh/m^3 (Trianel, 2009). H-gas stands for “high gas” consisting of 87 - 99 % methane – contrary to L-gas (low gas) which consists of 80-87 % methane (see e.g. Cerbe, 2004).

Based on the above information we derive formulas for c_{min}, c_{max} (see B):

$$c_{min} = -36.3612\sqrt{\frac{1}{I+77} - \frac{1}{314}}, \quad c_{max} = 0.4677\sqrt{I}.$$

For the risk free interest rate we use the interest rate for short term German government bonds, which is currently 1% p.a. (Bundesrepublik Deutschland Finanzagentur GmbH, 2010). We price an annual contract and use weekday TTF day-ahead prices, i.e. one year is approximately 250 days.

4.2 Storage Pricing with Dynamic Volatility

For pricing the storage we use the parameters fitted to (2.1) in Section 2. These were estimated using the observed (not risk adjusted) data. The effects of a risk-adjustment are discussed in Section 4.3. Moreover we set K in (3.1) as $K = 0.1$ and as terminal condition we use in the first instance $V(\cdot, \cdot, T, \cdot) = \mu I$. With these parameter values, we compute the storage value in time t by first solving

$$\begin{aligned} & \max_{c_{min}(I) \leq c \leq c_{max}(I)} [-cV_I + (c - a(c))11500P] \\ \text{s.t.} & \\ & c = \begin{cases} c_{max} & \text{for } 11500P > V_I, \\ c_{min} & \text{for } 11500P < V_I \wedge (c_{min} - K)11500P > c_{min}V_I, \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

and then using this optimal c in

$$\begin{aligned} (c - a(c))P - cV_I + V_t + \frac{\sigma^2}{2}V_{PP} + 0.0663(\sigma^2)^2 V_{\sigma^2\sigma^2} \\ + 0.0142(17.78 - P)V_P + 0.0131(6.2551 - \sigma^2)V_{\sigma^2} = \rho V. \end{aligned}$$

From our algorithm we obtain two kinds of information: The optimal control c (Figure 5) and the corresponding expected maximal storage value (Figure 6) for different price and storage levels. In the control variable's plot we can clearly see three regions. If the price is above the discounted long-term mean, then it is reasonable to withdraw and sell the gas. If the price is below this border then - as we have positive operational costs - there is an area where the operational costs are higher than the benefit of injecting. As the costs are constant in this scenario, they increase relatively to the decreasing c_{min} . This leads to the concave boundary between idleness and injecting. Eventually there is an area where the price is so low that the benefit of injection surpasses its costs.

On the value surface in Figure 6 we see a kink around the discounted long-term mean. The value is highest for full inventory and maximum prices and lowest for high prices and an empty storage. Having low prices and an empty storage gives the operator the chance to buy low and eventually sell for a higher price. To reveal further the option character of a gas storage we plot cross-sectional slices in Figure 7. For low inventory levels we see the classical shape of a put option and for high inventory levels the shape of a call option. Between these extremes, a straddle-like shape appears. In this case, the storage operator speculates on extreme prices. We

Figure 5: Optimal Control Strategy


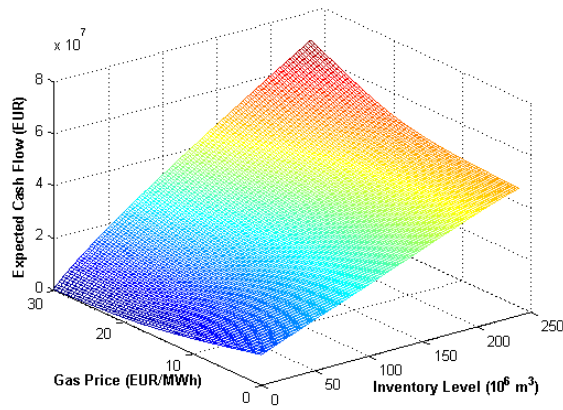



Figure 6: Expected Cash Flow



Volatility $\sigma^2 = 3.7$ at $t = 0$.

Figure 7: Cross-Sectional Slices of the Value Surface



The dashed line shows $\sigma^2 = 0.3$ at $t = 0$, the solid (dotted) line $\sigma^2 = 1.25$ (2.5).

notice that both control and value surface are similar to the results of Thompson et al. (2009).

From the results we can also see that volatility matters (see Figure 7). increasing the initial volatility yields a different value (but identical control) surface. The effect, however, restricts itself to the described borders at the control surface, i.e. between withdrawing and doing nothing resp. doing nothing and injecting. The latter effect is of minor size (and vanishes with increasing storage lease period). This result is reasonable, as higher volatility means more fluctuation around the critical boundary between injections and withdrawals (or doing nothing) and thus more chances to increase the storage value by trading. For the extreme regions, however, a stronger price fluctuation has no influence on the control. This corresponds with the fact that the kink around the (discounted) long-term mean becomes smoother with larger volatility (see Figure 7). The extreme price regions, however, show no difference.

The same results can be seen when comparing the storage value of our dynamic-volatility price model with the constant volatility model used in Thompson et al. (2009). Dynamic volatility introduces more oscillation and creates more value along the borders between injection, remaining idle and withdrawal (on the control surface).

4.3 Further Factors Influencing the Storage Value

Not only volatility or inventory level, but also other parameters – such as the cost function a – have a substantial influence on the storage value V . We now have a brief look on each of these factors as summarized below:

The interest rate


Increasing the interest rate means decreasing the slope of the storage value in time. For large interest rates, this slope can be negative, i.e. the value of the storage is decreasing with increasing time of operating the storage. This results seems paradoxical, but stems from the fact that future cash flows must be discounted, i.e. the larger ρ the smaller the current value of the future cash flow. However, as ρ is currently quite small, it is of minor importance.

The price process parameters

Changing the long-term mean μ shifts the border between withdrawal and remaining idle along the price axis on the control surface and therefore the location of the kink on the value surface. The kink is moved down along the price axis and the overall storage value is lowered as well. From this analysis we can easily derive the effect of introducing a risk premium in the price process.

The mean reversion λ has no influence on the control, but a significant influence

Figure 8: Cross-Sectional Slices of the Value Surface for Different Levels of Mean Reversion



The solid (dashed) line shows the value surface for $\lambda = 0.014$ (0.1).

on the value surface. The influence can be seen in Figure 8. The higher λ , the less volatile (so the more predictable) is the price process. The storage loses some of its option character and the value surface becomes flatter. In general (except for full inventories and high prices) it can be said that the higher the mean reversion the lower the storage value.

Modifying the volatility parameters, again, has no influence on the control, but a significant impact on the value surface. Increasing ω , the long-run mean of volatility, means more fluctuation, i.e. increases the storage value – however only within a certain area around the (discounted) long-term mean (as explained above). Introducing a volatility risk premium reduces ω , i.e. the fluctuation, and therefore eventually the storage value.

Increasing the volatility's mean reversion θ pushes the storage value around the

(discounted) long-term mean, but has no effect for extremely low or high prices – which is reasonable, as volatility is there of minor importance anyway. Modifying η , which controls the amplitude of the volatility dynamics results in changing not the control surface, but only the value surface which increases with η . The value surface is far less sensitive to η than to other parameters, most notably θ .

The cost function a

Increasing a constant a implies the enlargement of the area where remaining idle is optimal. The value surface is decreasing with growing K . Introducing a dynamic cost function a of the shape (3.2) creates a linear boundary between injecting and doing nothing independent of c_{min} at a price level of $(1 - h)/(1 + h)e^{-Tdt\rho}\mu$, where h is as defined in (3.2).


Terminal condition and terminal time

The terminal time – in combination with the terminal condition – influences the shape of the control surface and the value surface, if it is below the minimum time \mathcal{T} required to either fill or empty the storage given any inventory level. In Scenario (3.8) the optimal control is to withdraw always if $T < \mathcal{T}$. Because the price processes are mean reverting, increasing T and price increases the probability of a fall in spot price, and the attractions of the strategy of selling current inventory and replacing it later. This area on the control surface grows until – if T surpasses \mathcal{T} – the control surface is equal to scenario (3.9) which is independent of T .

With increasing T the effect of the terminal condition on the value surface is diminishing, as we can see in Figure 9, which shows two important results. In the left hand figure we see that the conditions (3.8) and (3.9) converge, and that the influence of the initial inventory level is decreasing in time. The right hand figure proves that the dynamics of the facility are much more important than details of e.g. lease period or initial fill level. This result can be observed even for short lease periods (7 days, i.e. 56 iterations, are enough). Higher price volatility of spot-markets is one argument in favour of an option-based operational storage strategy, but Figure 9 indicates that simply the fact of being more active than in a forward-based strategy is another major reason for yielding higher benefits from our facility.

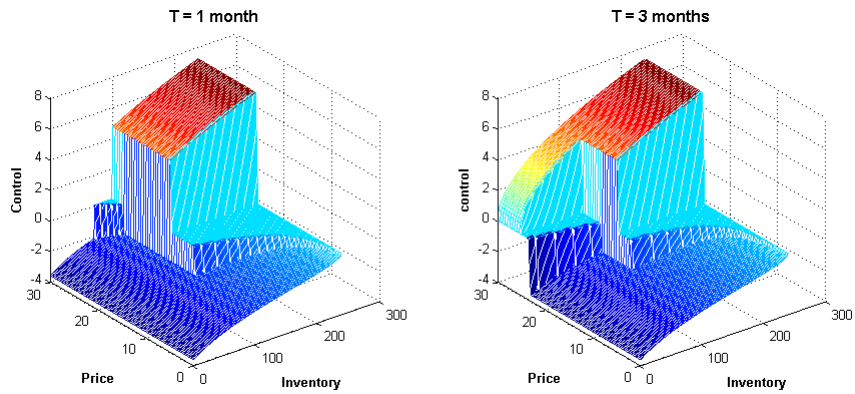
The case of (3.10) is a bit more complicated, especially when $T < \mathcal{T}$ (see Figure 10). Eventually, for the lease period of one year we have the following situation: it is optimal to inject, if $I < I_{end}$ and T is below the time needed to reach the fill level I_{end} . With increasing T and the higher the price the bigger the probability that prices fall and one can make profit by selling high and buying low. Eventually for $I < I_{end}$ we reach a structure similar to (3.9) – however with a long-run mean

Figure 9: Distance of Different Terminal Conditions for Increasing Terminal Time T



Terminal condition A: $V_A(T) = \mu * I$, terminal condition B: $V_B(T) = 0$.

Figure 10: The Effect of Introducing a Penalty Term in the Terminal Condition



of $e^{-\rho T dt} \nu \mu$ instead of μ . If $I \geq I_{end}$ and $T \geq \mathcal{T}$ we have a shape identical with (3.9). For $I \geq I_{end}$ the control surface is similar to (3.9). Below \mathcal{T} the area where $c_{opt} = 0$ is increasing to the right side (i.e. with lower prices) with decreasing T . The value surface reflects this by showing an abrupt jump in the area where the control differs from either (3.8) or (3.9). Moreover, below I_{end} the value surface is negative. For the value surface of (3.8) we refer to Thompson et al. (2009). There we see a monotonically increasing function with lowest value when price and inventory are zero.

5 Conclusions

In this article we analyze TTF spot prices and find their volatility to be quite variable. We show that existing concepts are not adequate to model this feature. Therefore we introduce a new continuous-time gas price model in which the volatility parameter follows a GARCH diffusion.

This result is subsequently incorporated in a PDE-based algorithm for pricing a gas storage facility. We discuss the design of the various storage parameters and give different versions for storage costs and terminal condition.

Using this pricing algorithm we perform a sensitivity analysis for all parameters. The larger the volatility at the beginning of the storage lease, the higher the storage value, although the effect is decreasing for extremely high and low prices. Volatility parameters have – as long as they enhance price oscillation – the same effect. A higher mean reversion reduces the price fluctuation and has a negative influence on the storage value (and vice versa). Knowledge about these effects is valuable, as one can easily estimate the effects of changing markets on the storage strategy and the facility value.

We also discuss the influence of different cost structures and terminal conditions. This does not only help to design storage lease contracts but also gives an important argument in favour of a dynamic storage operating strategy. We show that the dynamic features of a storage facility are much more important than details of e.g. lease period or lease time. Together with a higher spot-market volatility is this the major reason for preferring an option-based storage strategy to a forward-based one.

A The Function MINMOD

$$\text{minmod}(a, b) = \begin{cases} a & \text{if } |a| < |b| \ \& \ ab > 0, \\ b & \text{if } |b| < |a| \ \& \ ab > 0, \\ 0 & \text{if } ab \leq 0. \end{cases}$$

B Deriving the Control Boundaries

Thompson et al. (2009) conclude from fluid mechanics that the maximum outflow must be in a square root relation to the inventory level, i.e. $c_{max}(I) = K\sqrt{I}$, $K \in \mathbb{R}$. Our maximum deliverability is $7.2Mm^2/day$ and the maximum working gas is $237Mm^3$. Therefore $7.2 = K \cdot \sqrt{237} \rightarrow K = 0.4677$.

The injection of gas is more complicated as pump pressure has to be respected as well. Thompson et al. (2009) derive the relation

$$c_{min}(I) = -K_1\sqrt{1/(I + I_b) + K_2}, \quad K_1, K_2 \in \mathbb{R}.$$

To find K_1, K_2 we require two equations. In case of a full storage, no gas can be injected and therefore we can conclude $K_2 = -1/314$ and in case of an empty storage (i.e. maximal injectivity) we have $3.6 = K_1\sqrt{1/77 - 1/314} \rightarrow K_1 = -36.3612$.

References

- [1] Bellman R E. Dynamic programming. Princeton University Press: Princeton, NJ; 1957.
- [2] Boogert A, de Jong C. Gas storage valuation using a Monte Carlo method. Journal of Derivatives 2008; 15(3); 81–91.
- [3] Bollerslev T. Generalized autoregressive conditional heteroscedasticity. Journal of Econometrics 1986; 31; 307–327.
- [4] Bjerksund P, Stensland G, Vagstad F. Gas storage valuation: price modelling v. optimization methods. Instituted for Fretaksokonomi. Norwegian School of Economics and Business Administration. Working Paper 2008.
- [5] Bundesrepublik Deutschland Finanzagentur GmbH. Terms and conditions for bying German governmental bonds. URL: <http://www.deutsche-finanzagentur.de> [last access: Feb 10th, 2010].

- [6] Cartea A, Williams T. UK gas markets: the market price of risk and applications to multiple interruptible supply contracts. *Energy Economics* 2008; 30(3); 829–846.
- [7] Cerbe G: *Grundlagen der Gastechnik Gasbeschaffung, Gasverteilung, Gasverwendung*, 6th ed. Hanser: Munich; 2004.
- [8] Dixit A K, Pindyck R S. *Investment under uncertainty*. Princeton University Press: Princeton, NJ; 1994.
- [9] Drost F C, Werker B J M: Closing the GARCH gap: continuous time GARCH modelling. *Journal of Econometrics* 1996; 74; 31–57.
- [10] Engle R F. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 1982; 50(4); 987–1007.
- [11] FMET (Federal Ministry of Economics & Technology). URL: <http://www.bmwi.de/English/Navigation/root.html> [last access: Feb 10th, 2010].
- [12] Forsyth P, Chen Z. Implications of a regime-switching model on natural gas storage valuation and optimal operation. Forthcoming in *Quantitative Finance* 2010.
- [13] Gas Transport Services. Development of TTF liquidity. URL: <http://www.gastransportservices.nl/en/transportinformation/ttf-volume-development> [last access: Feb 10th, 2010].
- [14] Harten A. High resolution schemes for hyperbolic conservation laws. *Journal of Computational Physics* 1983; 49; 357–393.
- [15] Hull J C, White A. Valuing derivative securities using the explicit finite difference method. *Journal of Financial and Quantitative Analysis* 1990; 25(1); 87–100.
- [16] Hull J C. *Options, futures, & other derivatives*, 4th ed. Prentice Hall: New Jersey; 1999.
- [17] International Energy Agency (IEA). *Natural gas market review 2009*. Maulde & Renou: Paris; 2009.
- [18] Kiesel F, Herkner T: *Energiemarkt Deutschland, Zahlen und Fakten zur Gas- und Stromversorgung*. VWEW Energieverlag GmbH: Frankfurt am Main; 2008.

- [19] Landesamt für Bergbau, Energie und Geologie (LBEG). Current capacities of gas storage facilities in the federal state of Niedersachsen. URL: <http://www.lbeg.niedersachsen.de>. [last access: Dec 10th, 2010].
- [20] RWE. Exemplary contract of a gas storage lease. URL: <http://www.rwe.com/web/cms/de/307376/rwe-rheinland-westfalen-netz/gasspeicher/speicherservices/speicherzugang-epe/download/> [last access: Dec 10th, 2010].
- [21] Schwartz E, Smith J E: Short-term variations and long-term dynamics in commodity prices. *Management Science* 2000; 46(7); 893–911.
- [22] Sedlacek R: Untertage-Gasspeicherung in Deutschland. *Erdöl Erdgas Kohle* 2007; 123(11); 422–432.
- [23] Seydel R: *Tools for computational finance*. Springer: Berlin; 2009.
- [24] Thompson M, Davison M, Rasmussen H: Natural gas storage valuation and optimization: a real options application. *Naval Research Logistics* 2009; 56(3); 226–238.
- [25] Trianel Gasspeicher Epe GmbH & Co. KG (Trianel). URL: <http://www.trianel-gasspeicher.com/index.php?id=364> [last access: Dec 10th, 2010].

Diskussionspapiere 2010 Discussion Papers 2010

- 01/2010 **Mosthaf, Alexander, Schnabel, Claus and Stephani, Jens:** Low-wage careers: Are there dead-end firms and dead-end jobs?

Diskussionspapiere 2009 Discussion Papers 2009

- 01/2009 **Addison, John T. and Claus Schnabel:** Worker Directors: A German Product that Didn't Export?
- 02/2009 **Uhde, André and Ulrich Heimeshoff:** Consolidation in banking and financial stability in Europe: Empirical evidence
- 03/2009 **Gu, Yiquan and Tobias Wenzel:** Product Variety, Price Elasticity of Demand and Fixed Cost in Spatial Models
- 04/2009 **Schlüter, Stephan:** A Two-Factor Model for Electricity Prices with Dynamic Volatility
- 05/2009 **Schlüter, Stephan and Fischer, Matthias:** A Tail Quantile Approximation Formula for the Student t and the Symmetric Generalized Hyperbolic Distribution
- 06/2009 **Ardelean, Vlad:** The impacts of outliers on different estimators for GARCH processes: an empirical study
- 07/2009 **Herrmann, Klaus:** Non-Extensivity versus Informative Moments for Financial Models - A Unifying Framework and Empirical Results
- 08/2009 **Herr, Annika:** Product differentiation and welfare in a mixed duopoly with regulated prices: The case of a public and a private hospital
- 09/2009 **Dewenter, Ralf, Haucap, Justus and Wenzel, Tobias:** Indirect Network Effects with Two Salop Circles: The Example of the Music Industry
- 10/2009 **Stuehmeier, Torben and Wenzel, Tobias:** Getting Beer During Commercials: Adverse Effects of Ad-Avoidance

- 11/2009 **Klein, Ingo, Köck, Christian and Tinkl, Fabian:** Spatial-serial dependency in multivariate GARCH models and dynamic copulas: A simulation study
- 12/2009 **Schlüter, Stephan:** Constructing a Quasilinear Moving Average Using the Scaling Function
- 13/2009 **Blien, Uwe, Dauth, Wolfgang, Schank, Thorsten and Schnabel, Claus:** The institutional context of an “empirical law”: The wage curve under different regimes of collective bargaining
- 14/2009 **Mosthaf, Alexander, Schank, Thorsten and Schnabel, Claus:** Low-wage employment versus unemployment: Which one provides better prospects for women?

Diskussionspapiere 2008 Discussion Papers 2008

- 01/2008 **Grimm, Veronika and Gregor Zoettl:** Strategic Capacity Choice under Uncertainty: The Impact of Market Structure on Investment and Welfare
- 02/2008 **Grimm, Veronika and Gregor Zoettl:** Production under Uncertainty: A Characterization of Welfare Enhancing and Optimal Price Caps
- 03/2008 **Engelmann, Dirk and Veronika Grimm:** Mechanisms for Efficient Voting with Private Information about Preferences
- 04/2008 **Schnabel, Claus and Joachim Wagner:** The Aging of the Unions in West Germany, 1980-2006
- 05/2008 **Wenzel, Tobias:** On the Incentives to Form Strategic Coalitions in ATM Markets
- 06/2008 **Herrmann, Klaus:** Models for Time-varying Moments Using Maximum Entropy Applied to a Generalized Measure of Volatility
- 07/2008 **Klein, Ingo and Michael Grottko:** On J.M. Keynes' “The Principal Averages and the Laws of Error which Lead to Them” - Refinement and Generalisation