The Journal of Energy Markets (27–46) Volume 2/Number 3, Fall 2009

Cointegration between gas and power spot prices

Cyriel de Jong
Kyos Energy Consulting/Erasmus University Rotterdam,
Lange Herenstraat 38 zwart, 2011 LJ Haarlem, the Netherlands;
email: dejong@kyos.com

Stefan Schneider
EON Energy Trading, Holzstrasse 6, 40221 Düsseldorf, Germany;
email: stefan.schneider@eon.com

In this paper we show how cointegration can be applied to capture the joint dynamics of multiple energy spot prices. For an example system we study the Title Transfer Facility, the Zeebrugge gas spot market and the National Balancing Point gas spot market, and, additionally, the Amsterdam Power Exchange power spot market, since these markets are strongly connected in terms of physical transportation and generation of power from gas. We develop a cointegrating multi-market model framework that is able to plausibly connect different single-market spot-price models. This is achieved by considering the mean-reverting spot-forward price spreads instead of spot prices only. Our analysis shows that the gas prices are strongly cointegrated, with a specific connection pattern for the markets, whereas cointegration of gas and power prices is at long-term forward price levels only.

1 INTRODUCTION

Exposures in energy markets often involve a number of commodities. They are often directly attributable to spread positions rather than to outright positions. For example, a company trading between different countries faces geographical spread exposures, a company with gas-fired power generation assets faces spark spread exposures, a company with gas storage capacity faces time spread exposures, and so on. Yet most research focuses on the price dynamics of single commodities in single markets. For example, the first two papers in the Fall 2008 issue of this journal focused on Nord Pool electricity (Borak and Weron (2008)) and New York Mercantile Exchange (NYMEX) natural gas (Spargoli and Zagaglia (2008)). Most “energy” papers actually focus on electricity alone: examples are Longstaff and Wang (2004), Koekebakker and Ollmar (2005) and Lucia and Schwartz (2002), who investigate the dynamics of electricity prices on the Pennsylvania–New Jersey–Maryland (PJM) and Nord Pool electricity markets.

There is, nevertheless, a range of literature that tries to find common factors in forward prices of more than one commodity. Westgaard *et al* (2008) apply Kalman filtering to prices of propane, butane and naphtha; Lin and Tamvakis (2001) study spillover effects from one energy futures market to the other. One of the earliest studies on the joint dynamics of multiple commodity prices is presented in Pindyck (1999, 2001). Pindyck considers the long-run dynamics and finds evidence for cycles, though at a very low mean-reversion speed.

The motivation for writing this paper is to develop a simulation model for gas and power spot prices on different markets that can be used for valuation and risk management in a day-to-day environment in energy trading. The model should produce realistic dynamics and communicate well with any forward curve simulation model. More specifically, the aim of this paper is to show how cointegration can be applied to capture the joint dynamics of multiple spot prices. We use a formulation in which each spot series is mean reverting to prompt month forward prices (which is similar to the formulation in Blanco *et al* (2002)), thereby also capturing the linkage between spot and forward markets.

2 COINTEGRATION STUDIES IN ENERGY MARKETS

Cointegration is quite a general concept and has been applied to energy markets by several researchers, though mainly to forward prices. For example, Brunetti and Gilbert (2000) and Figuerola-Ferretti and Gonzalo (2008) study cointegration between NYMEX and the International Petroleum Exchange's crude oil futures prices; Gjolberg and Johnsen (1999) study cointegration between crude and refined oil prices. An even wider range of commodities is the subject of the cointegration analysis of Serletis and Herbert (1999): they look at cointegration between different energy commodities in the North American market and find predictable patterns in the natural gas–fuel oil forward spread, but not in forward spreads involving electricity prices. In the same issue of *Energy Economics* as the Serletis and Herbert paper, De Vany and Walls (1999) provide one of the few papers on cointegration in spot markets. Using data from the period 1994–96, the authors find cointegration patterns between 11 regional power spot prices in the US. They find strong and direct relationships between the markets, especially in off-peak periods, and conclude that the transportation capacities lead to efficient and stable power prices.

The papers listed above demonstrate the different sources for cointegration: it primarily stems from transportation linkages between different markets (power, crude oil) or from processing and substitution linkages between different commodities (power produced from natural gas, refined products produced from crude oil, natural gas as a substitute for fuel oil). This study addresses both types of linkages and includes the linkage between forward and spot prices. We provide a framework for capturing the dynamics of the same commodity (natural gas or power) on multiple markets, the dynamics between natural gas and power spot prices, and the dynamics between spot and month-ahead forward prices. At the same time, we do not model the joint dynamics of power and natural gas forward prices explicitly, but treat those as exogenous.

3 A DESCRIPTION OF THE WESTERN EUROPEAN GAS AND POWER MARKETS

The continental European natural gas market is in the middle of a process of liberalization. A lot of progress is being made: trading hubs are developing, new regulations aim for more even and fair access to gas assets and the member states are opening up their home markets further and further. The various wholesale gas markets in Europe are nevertheless at widely differing stages of development. They range from the UK's National Balancing Point (NBP), which is highly liquid and efficient, to markets such as the Austrian CEGH or the Spanish CDG, where liquidity is barely established.

The most widely developed continental European gas markets are not far from the UK. For a long time, the only direct connection between the UK and continental Europe has been through a pipeline called the Interconnector, which lands at Zeebrugge in Belgium, a location at which pipelines from Norway also hit the continent. At this Zeebrugge location another trading hub, with rather good liquidity, has developed. A third important trading hub, called the Title Transfer Facility (TTF), has been implemented in the Netherlands. This hub has a virtual nature, bringing together all the gas in the Dutch high-calorific grid.

Natural gas is primarily used for heating, industrial processes and power production. An important driver of the fluctuation in gas demand comes from gas-fired power plants. Gas-fired generation acts as a swing or peak power supply in many European countries, in combination with hydropower. European countries have different shares of gas-fired generation. Our empirical analysis focuses on the Dutch power market, which has a gas-fired generation share of almost 60%. Its spot market, the Amsterdam Power Exchange (APX), is not the largest in Europe, but it is probably the most dependent on what happens in the gas market. If we find clear relationships in the Dutch market, then they may be translated to other markets. This connection is most likely to be found with the nearby countries: Germany (consuming 22% of western European electricity), Belgium and France ("market-coupling" takes place between the Netherlands, Belgium and France).

Whereas Germany may be the dominating western European power market, it is not "the leader" in the natural gas market. The German gas market is still developing

FIGURE 1 Price history of the spot markets APX, TTF, Zeebrugge and NBP as examined in this paper.

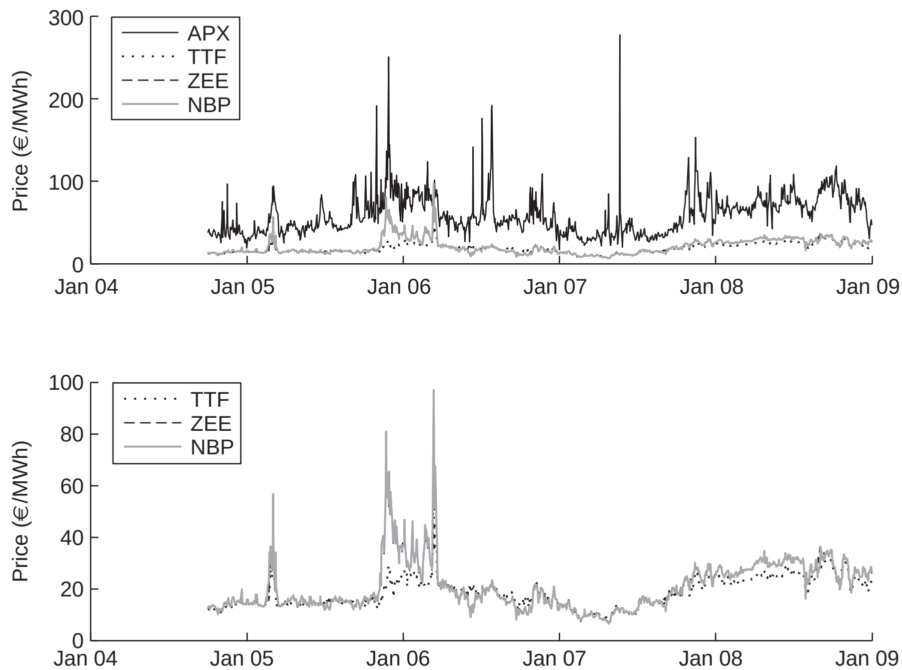

NBP prices have been converted to €/MWh.

and the history of available market data is too short for an empirical study. The UK (NBP), Belgian (the Zeebrugge gas spot market, or ZEE) and Dutch (TTF) natural gas markets are more developed than the German one; data from those markets has been used for this analysis. We therefore study their dynamics in conjunction with power price data from the Dutch power market APX, see Figure 1 for a plot of historical price data of these markets.

4 ENERGY SPOT-PRICE CHARACTERISTICS: CORRELATION AND COINTEGRATION

Price levels and dynamics on different gas markets may vary, but they generally exhibit various forms of seasonality, mean reversion, time-varying volatility, correlation and cointegration between markets.

The concepts of correlation and cointegration deserve some extra explanation. In financial and energy markets, time series are often assumed to be correlated in returns.

This is a useful concept and is especially applicable for analysis with a short-term horizon, such as short-term risk and hedging calculations. However, even a strong correlation of close to 1 does not ensure that prices of different time series stay together over longer horizons. In energy markets, such economically driven longer-term fundamentals do exist.

For circumstances where economic fundamentals eventually enforce a specific relationship between two or more price series, the concept of cointegration becomes very useful. Important work on this concept has been performed by Clive Granger and Robert Engle. The two professors shared the 2003 Nobel Memorial Prize for their innovation, first described in 1987.

The idea behind cointegration is that individual price series may be non-stationary, but one or more weighted combinations of the series are stationary. In many cases, stationarity can be achieved by first-order differencing, or by other mathematical transformations such as seasonal adjustments. A typical application in financial market analysis is therefore to model returns or price differences rather than price levels. Yet this approach of differencing non-stationary economic series into stationary series has been criticized for throwing out and ignoring valuable long-run equilibrium information (Engle and Granger (1987)). Intuitively, cointegration among a set of variables implies that there exist fundamental economic forces that make the variables move together stochastically over time (Niemi (2003); Urbain (1993)). These movements in variables are related in a predictable way to the discrepancy between observed and equilibrium states.

The enforcement of the cointegration relationship is easiest to understand in the so-called error-correction mechanism of two price series X and Y. Suppose that X and Y are both random walks, so integrated of order 1, $I(1)$. If there is an equilibrium relationship between X and Y equal to:

$$X(t) = a_0 + a_1 Y(t) \tag{1}$$

this means there is a linear combination of X and Y that is stationary or integrated of order 0, $I(0)$. The actual dynamics may consequently be described as:

$$X(t) = X(t-1) - \kappa(X(t-1) - a_0 + a_1 Y(t-1)) + u_t \tag{2}$$

with u_t being normally distributed and independent and $\kappa > 0$ being the mean-reversion rate.

The formulation shows that cointegration may be interpreted as mean reversion in a weighted combination of two (or more) variables. This is how we implemented the concept, in combination with a mean-reverting spot-to-forward price model.

5 MEAN-REVERTING SPOT-PRICE MODELS FOR NATURAL GAS

In energy markets, the concept of stationarity is often directly studied in the context of mean reversion. Mean reversion is most often defined in terms of a so-called Ornstein–Uhlenbeck process, meaning that the model is specified in log prices ($x = \ln P$) not absolute prices (P). Assuming zero drift, it is specified as follows:

$$x_t - x_{t-1} = \alpha\left(\mu - x_{t-1} - \frac{\sigma^2}{2\alpha}\right) + \sigma\varepsilon_t \quad \text{with } \varepsilon_t \sim N(0,1) \tag{3}$$

The parameter α is the daily mean-reversion rate, the parameter σ is the daily standard deviation of returns and μ is the mean-reversion level.

The rate of mean reversion, α, is an important statistic for various applications. First, mean reversion ensures that the distribution of prices over longer horizons is narrower than in a process of Brownian motion. Second, mean reversion introduces a form of predictability to spot prices: for example, when prices are above their mean, we may expect prices to fall in the following days, on average, which is useful information for a storage operator (Boogert and De Jong (2008)).

The previously defined mean-reversion model has only one source of randomness and is therefore a one-factor model. It has been used by several researchers (De Vany and Walls (1999); Eydeland and Wolyniec (2003)) to conclude that spot prices in power and gas markets are non-stationary. More realistically, the mean level should itself be modeled as a stochastic factor. We therefore incorporate forward curve information in the mean level. Blanco *et al* (2002) propose using the natural logarithm of the front-month forward price as the mean level. This leads to the following basic formulation of a gas spot-price model for market i, with log spot price $\ln S_{i,t}$ and log month-ahead price $\ln M_{i,t}$:

$$\ln S_{i,t} = \ln S_{i,t-1} + \alpha_i(\ln M_{i,t-1} - \ln S_{i,t-1}) + u_{i,t} \tag{4}$$

$$u_{i,t} \sim N(0, \sigma_{i,t}^2), \qquad \sigma_{i,t}^2 = k_1 + k_2 u_{i,t-1}^2 + k_3\sigma_{i,t-1}^2 \tag{5}$$

The random element $u_{i,t}$ is assumed to be normally distributed with GARCH(1, 1) volatility. It is correlated across markets i but not between time periods t. For simplicity we do not assume volatility spillover effects, which could be captured in a multivariate GARCH.

In Figure 2 on the facing page we show the joint dynamics of log spot and log M1 (we use simply M in place of M1 in the formulas in this paper) prices on the APX,[1] the TTF, the Zeebrugge (ZEE) and the NBP markets from October 2004 to December 2008.

[1] Note that the power forward prices are actually from the European Energy Derivatives Exchange, the Dutch power forward market, since APX is a spot market only. However, to keep the notation simple we denote both the Dutch spot and the forward as APX.

FIGURE 2 Historical log spot and log M1 prices (bold lines).

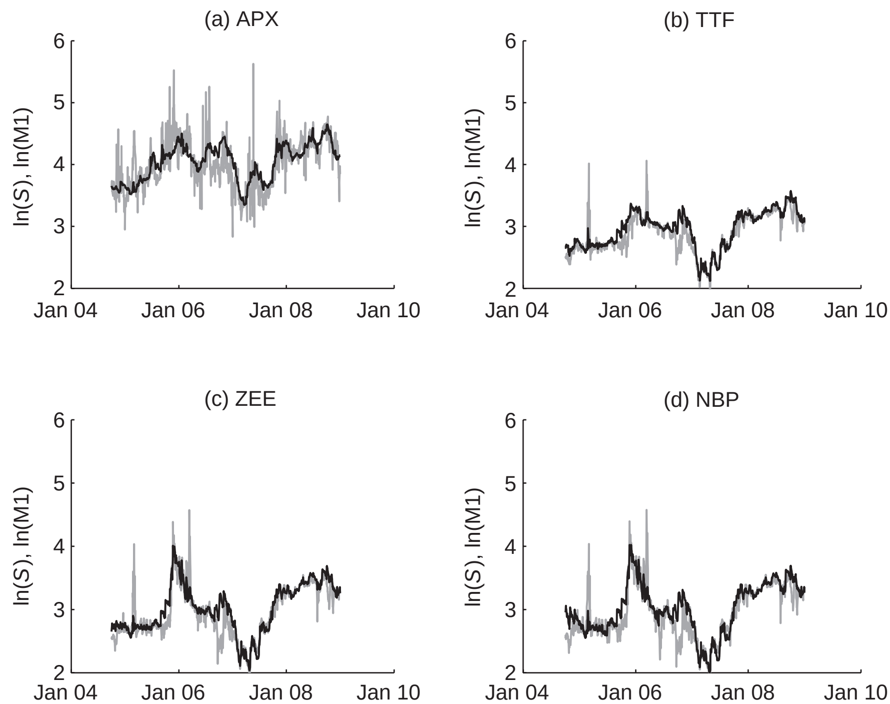

6 MEAN-REVERTING AND REGIME-SWITCHING SPOT-PRICE MODELS FOR POWER

So far we have ignored the non-normal nature of the residuals of a mean-reverting price model. Incorporation of jumps and spikes is quite important, however, in particular for power spot prices. In De Jong and Huisman (2003) and De Jong (2006), an empirical analysis of different power spot-price models is provided. The proposed solution in most markets is to use regime-switching models: more specifically, the "independent spike model". In this paper we work with a model variant of De Jong (2006) containing three potential regimes: a stable mean-reverting regime, a low-price regime capturing down spikes and a high-price regime capturing up spikes. It takes into account that spot prices in the Western European power markets may sometimes briefly and strongly deviate from the usual price levels but that these deviations are typically short lived and largely independent of previous and subsequent price levels. The model is formulated in terms of the deseasonalized natural logarithm of price levels x_t, with a time-varying mean level μ_t defined later on.

Mean-reverting regime M:

$$\mathrm{d}x_t^M = \alpha(\mu_t - x_{t-1}^M) + \sigma\varepsilon_t \tag{6}$$

Spike regimes:

$$x_t^S = \mu_t + \sum_{i=1}^{n_t+1} Z_{t,i}$$

High spike regime H:

$$Z_{t,i} \sim N(\mu^H, \sigma^H), \quad n_t \sim \mathrm{POI}(\lambda^H), \quad \mu^H > 0$$

Low spike regime L:

$$Z_{t,i} \sim N(\mu^L, \sigma^L), \quad n_t \sim \mathrm{POI}(\lambda^L), \quad \mu^L < 0$$

Markov transition matrix:

$$\Pi = \begin{bmatrix} 1 - \pi^{MH} - \pi^{ML} & \pi^{MH} & \pi^{ML} \\ \pi^{HM} & 1 - \pi^{HM} & 0 \\ \pi^{LM} & 0 & 1 - \pi^{LM} \end{bmatrix}$$

An additional modification was, however, necessary. This is because the month-ahead price not only reflects the expected price level in the stable regime but is in fact a probability-weighted expectation of price levels in all three regimes. For example, when the high-price regime is especially likely to occur, then the mean level in the stable mean-reverting regime should be lower than the month-ahead price.

More precisely, we need to derive $A = M_t / E[S_t \mid S_t \text{ is no spike}]$, assuming that the month-ahead price equals M, which is the unconditional expected spot-price level. To simplify the adjustment, we first assume that the probability of each regime is time independent and therefore equal to the unconditional probability for each regime. We call these probabilities π^L and π^H. Second, we approximate the Poisson distribution of the log prices in the spike regimes with a normal distribution. It then holds that:

$$A = \pi^M + \pi^L \exp(\mu^L + \tfrac{1}{2} Y^L \sigma^{L^2}) + \pi^H \exp(\mu^H + \tfrac{1}{2} Y^H \sigma^{H^2}) \tag{7}$$

where Y^L and Y^H are the expected number of jumps in regimes L and H, respectively, $\exp(\mu^L + \frac{1}{2} Y^L \sigma^{L^2}) - 1$ is the expected percentage deviation of the price in regime L

from the price in the stable regime M and $\exp(\mu^H + \frac{1}{2}Y^H {\sigma^H}^2) - 1$ is the expected percentage deviation of the price in regime H from the price in the stable regime M.

7 A MULTI-MARKET COINTEGRATED SPOT-PRICE MODEL

Having formulated the individual market models we now move on to the multi-market gas-power model. It is capable of generating daily simulations of day-ahead spot prices of multiple power and gas markets. Each individual spot price is connected to the spot and forward prices of one or more other markets.

Note that the model focuses on spot prices and assumes that the month-ahead price simulations are exogenously generated. Of course, the performance of the spot-price simulation model partly depends on how the relationships between products are defined in the forward curve simulation model. We advocate the use of cointegration models for forward prices; we have implemented such cointegrated forward price models for daily use in our own companies. However, in this study we stick to historical forward price dynamics and focus on the spot–forward spread in order to avoid mixing up results from two separate models.

As described earlier, we include correlation in the residuals $(u_{i,t})$ between different markets. This is an elementary approach to connect prices. However, as explained before, this is not expected to tie prices together sufficiently in the long run. A likely behavior is that prices of TTF and Zeebrugge, for example, diverge too often if the simulation model only contains this correlation. In the history of both markets such a departure has occurred only once, whereas in all other time periods the TTF–Zeebrugge spread has seldom exceeded 10%. This is visible in Figure 3 on the next page and Figure 4 on page 37.

The connection between the eight time series of spot and month-ahead price levels has been tested using the Johansen (1988) cointegration test. The results indicate that all variables are cointegrated with each other at the 95% confidence level. One interpretation is that the connection between the series goes beyond the connection between spot and M1 prices in the same market. We incorporate this in a mathematical framework that stays close to the original single-market formulation.

To achieve this, we observe that mean reversion of spot-to-M1 price levels is almost equal to mean reversion in the spot–M1 spread to zero. The only difference is that the spread on the left-hand side of the second equation is lagged by one day:

$$\begin{aligned} &\ln S_{i,t} = \ln S_{i,t-1} + \alpha_i(\ln M_{i,t-1} - \ln S_{i,t-1}) + u_{i,t} \\ &\qquad \iff \quad \ln S_{i,t} - \ln M_{i,t-1} = (1-\alpha_i)(\ln S_{i,t-1} - \ln M_{i,t-1}) + u_{i,t} \end{aligned} \tag{8}$$

If this spot–M1 spread at time t for market i is additionally correlated to the same spread at time t in markets $j \in J_i$, then we capture cointegration between different

FIGURE 3 Historical log spot prices.

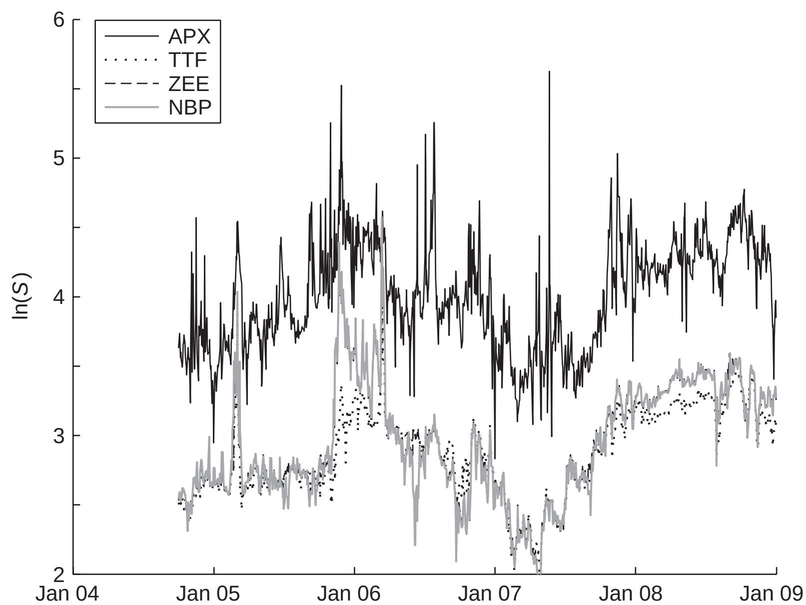

markets on the level of the spot–M1 spreads:

$$\begin{aligned}\ln S_{i,t} - \ln M_{i,t-1} = (1-\alpha_i)(\ln S_{i,t-1} - \ln M_{i,t-1})\\ + \sum_{j \in J_i} (1-\alpha_{ij})(\ln S_{j,t-1} - \ln M_{j,t-1}) + u_{i,t} \quad (9)\end{aligned}$$

where the residuals $u_{i,t}$ are normally distributed and are correlated to the other residuals. This correlation in residuals captures the short-term relationships between markets, whereas the parameters $\alpha_{i,j}$ capture the long-term relationships between markets. In the context of this article we incorporate the regime switches in the power market by simply assuming that the spike states are totally disconnected from the dynamics in any other market. More sophistication is possible though, in particular by letting the spike probabilities and levels across different power markets be correlated and by introducing correlation between spikes and gas spot-price levels.

The set J_i contains the markets j to which the prices in market i are "connected". It is intuitive to assume, for example, that the spot–M1 spread of the TTF not only mean reverts to 0 but also to the same spread in the APX power market. So, when power spot prices are temporarily above their (expected) month-ahead level, then gas spot prices are probably also above their expected month-ahead level. This formulation has the

FIGURE 4 Historical log M1 prices.

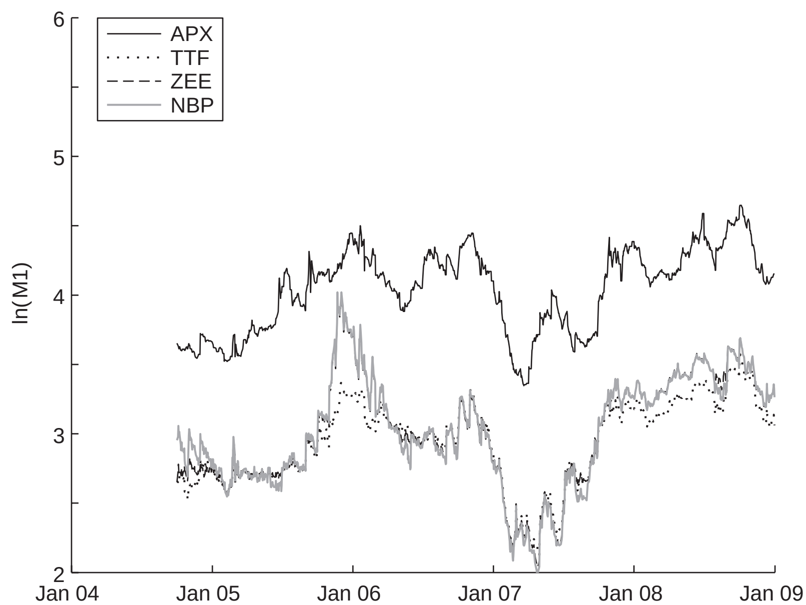

additional advantage that we can keep on treating the M1 prices as exogenous: given M1 prices, which may be cointegrated between markets as well, and given simulated power spot prices, we can express the gas spot prices. Likewise, the Zeebrugge and NBP gas spot–M1 spreads can be modeled as mean reverting to 0 and to the same spread as in the other markets. Additional gas markets could be added, connecting either the gas or power prices or a combination of these.

The primary motivation for our model specification lies in investigating the fundamental structures between spot and month-ahead prices as well as between markets. In a more standard covariance or cointegration specification, where each series is treated similarly, we would not be able to capture those fundamentals properly.

8 EMPIRICAL RESULTS

The model has been calibrated using historical data from the APX (power), the TTF, Zeebrugge and the NBP (gas) markets for the period from October 2004 to December 2008, for working days only. NBP prices have been converted to €/MWh to get the same unit for all price series. We applied a combination of least-squares regression and maximum-likelihood estimation. Maximum-likelihood estimation was needed

TABLE 1 GARCH (1,1) parameters.

	APX	TTF	ZEE	NBP
k_1	—	0.0003	0.0006	0.0007
k_2	—	0.3168	0.1386	0.2970
k_3	—	0.6138	0.7128	0.5544
Volatility	0.1952	0.0693	0.0635	0.0709

for estimation of the GARCH$(1,1)$ variance structure in the gas markets and the regime-switching specification in the power market.

The period from November 2005 to March 2006 deserves special attention. It was marked by a very tight supply situation in the UK gas market and a technical problem with the largest UK gas storage facility (named Rough) due to a fire. In the data it shows up as a very volatile period with very high prices on the NBP and Zeebrugge gas markets, and a relative disconnection from the TTF market. We decided to leave this period out of our analysis because we believe it was a (temporary) fundamental break requiring an analysis of its own. This is the only situation in which we removed outliers.

Using this data set of almost 1,000 days, we first estimate a simple mean-reverting specification for each market individually. All four spot-price series are clearly mean reverting to front month M1 with mean-reversion rates around 10% for the gas markets and over 33% for the APX power market. The residuals of the gas mean-reverting models exhibit large swings over time, which are captured in the GARCH$(1,1)$ specification by a rather high parameter for the previous day's squared residual (the "response parameter" k_2), see Table 1. A simple mean-reverting model such as GARCH is certainly not sufficient to capture the dynamics in the power spot market. We therefore estimate the regime-switching model with three independent regimes, see the results in Table 2 on the facing page (see de Jong (2007) for a detailed description of the maximum-likelihood method).

The mean-reversion parameter is lower than it is without the spike regimes. This indicates that mean reversion is usually more limited and is only occasionally quite strong: when prices move from the spike regime back to the normal regime. The switch probabilities indicate that there is an approximately equal probability (4%) of moving out of the normal regime to any of the spike regimes. Prices tend to stay somewhat longer in the high-price regime than in the low-price regime.

We now move on to the model involving the four spot markets and the four month-ahead markets as specified by Equation (9). The assumed relationship is captured as follows and is not related to causality:

TABLE 2 Parameter estimates for the regime-switching model APX.

	Time-series parameters		Switch probabilities	
Normal regime N	α	0.2230	From N to H	0.0426
	μ	3.5092	From H to N	0.2292
	σ	0.0945	From N to L	0.0374
			From L to N	0.3704
High-spike regime H	μH	0.2118		
	σH	0.1585		
	λH	0.6778		
Low-spike regime L	μL	0.1756		
	σL	0.1138		
	λL	0.6889		

TABLE 3 Cointegration dependencies.

	APX	TTF	ZEE	NBP
Estimates				
APX	0.6586	0.0006	–0.0210	–0.0395
TTF	—	0.9043	0.6062	0.0240
ZEE	—	—	0.5151	0.5284
NBP	—	—	—	0.5413
T-ratio				
APX	27.6794	0.0652	–2.5250	–4.2734
TTF	—	55.0105	22.0589	0.5680
ZEE	—	—	23.5939	13.7080
NBP	—	—	—	25.8424

- the TTF market moves together with the APX market;
- the Zeebrugge market moves together with the TTF and APX markets; and
- the NBP market moves together with the Zeebrugge, TTF and APX markets.

The results of the parameter estimation are displayed in Table 3 and Table 4 on the next page. The anticipated strong link between the NBP and Zeebrugge markets is confirmed by the high cointegration parameter (0.5284, $t = 13.7$) and daily return correlation (61.15%). Also strongly connected are the TTF and Zeebrugge markets, with estimates of 0.6062 ($t = 22.0589$) and 42.33%. We conclude that the spot–M1

TABLE 4 Correlations between residuals.

	APX	TTF	ZEE	NBP
APX	100.00%	1.73%	–0.20%	1.08%
TTF	–1.73%	100.00%	42.33%	49.81%
ZEE	–0.20%	42.33%	100.00%	61.15%
NBP	1.08%	49.81%	61.15%	100.00%

spread of two “adjacent” gas markets (TTF–ZEE and ZEE–NBP) is explained by its own spread of the previous day and by the “adjacent” spread of the previous day.

Somewhat surprising at first sight is the low cointegration (0.0240, $t = 0.5680$) between TTF and NBP, at least in the chosen setting. To investigate this further we also estimated the model in different orders. When we look at how TTF is cointegrated with NBP without Zeebrugge in the equation, there is also a strong and significant cointegration between TTF and NBP, albeit a somewhat lower one (cointegration parameter $= 0.4660$, correlation of residuals $= 42.06\%$). The exact nature of this connection becomes even clearer when the places of NBP and ZEE are swapped in the dependency chain. The last equation then yields a high cointegration parameter of ZEE on TTF (0.3994), even higher than that on NBP (0.2702). Apparently, TTF can be said to be linked to both Zeebrugge and NBP, but more strongly to Zeebrugge, the closer market. A similar conclusion is that the data shows that the chain of connection runs from TTF to Zeebrugge to NBP, or vice versa.

Maybe less expected is that the spot–M1 spreads of the three gas markets are barely, or even negatively, correlated with the spot–M1 spread in the APX power market. This holds for the price relationships (cointegration) as well as for the daily return relationships (correlations). The negative estimates of the NBP market with the APX market could indicate some form of multi-collinearity: the NBP market is strongly connected to the Zeebrugge market and other potential relationships are better left out of the equation.

A plausible explanation for the lack of a positive relationship between power and gas markets in our model is that the relationship is already captured in the month-ahead forward markets. The additional movements of spot prices around forward prices in the power market are independent between power and gas. The fact that very few power plants trade in the spot market to handle shortages and surpluses in natural gas probably plays a role. Instead, most gas is sourced through longer-term contracts with oil-indexed pricing; a little gas is sourced in the gas forward markets, leaving almost nothing to be sourced in the spot market.

In order to verify the stability of the parameter estimates, we estimated parameters with a moving time window of one year. The time step taken is half a year (110 days), creating eight subsamples, each with about 50% overlap with a neighbouring subsample. Table 5 on the next page shows some results and we highlight three findings. First, the lack of cointegration between the power market APX and any of the gas markets is very consistent. Second, the cointegration between the gas markets, eg, between Zeebrugge and TTF, is strong in all subsamples. Third, there seems to be interplay between correlation and cointegration. In the subsamples with a relatively high cointegration parameter, the correlation of residuals is somewhat below average, and vice versa. Whether this is a coincidence or not is hard to judge based on only eight estimates with overlapping data.

Based on the estimated parameters we generated 250 Monte Carlo simulations of the four markets' cointegrated log spreads. We limit our attention to the period starting in April 2006: after the tight-supply period in the UK market. The month-ahead forward price scenarios are the actual historical prices over this period. Note that our model is independent of the nature of M1 price scenarios as it is only dealing with the difference between log spot and log M1 prices. With 250 scenarios and 718 trading days per market, the total sample size is 179,500 per market.

To better understand the meaning of the parameters, Figure 5 on page 43, Figure 6 on page 44 and Table 6 on page 44 display the comovements of the (log) spot–M1 spreads between the different markets. For example, the spot–M1 spreads of the TTF and Zeebrugge markets follow each other quite closely, whereas the TTF and APX spot–M1 spreads are visibly disconnected. A very similar pattern is observed in Figure 6 on page 44, displaying one representative scenario of the 250 simulations.

Whereas the connection between markets seems to be well captured, the distribution of an individual market's returns displays some discrepancies between history and simulations, see Table 6 on page 44.

Simulated spreads in the three gas markets have hardly any skewness, which is an artefact of the more or less symmetrical nature of the mean-reversion and GARCH model elements. In contrast, historical spreads exhibit skewness; especially the large positive (0.90) skewness for TTF is not taken up by the simulations. The different historical skewness levels of NBP and Zeebrugge, compared with TTF, make it questionable though whether positive skewness for TTF is actually expected for the future when the three markets will probably become even more similar.

The properties of (log) spread changes show a more consistent pattern between the various markets: skewness is strongly negative for all gas markets (-0.58 to -1.15) and kurtosis is rather high (13.21 to 22.65). The simulations have some left-skewness as well (-0.33 to -0.48) but less than in the historical data. Similarly, the simulations do not reach the high historical kurtosis levels. The simulations for the power market (APX) match the historical properties much better. The regime-switching model for power (APX) does allow for positive skewness and achieves a good fit with history in terms of both skewness and kurtosis.

TABLE 5 Moving window estimates of cointegration coefficients and of correlations.

		Start and end dates							
		Oct 04 –Aug 05	**Mar 05 –Jun 06**	**Aug 05 –Nov 06**	**Jun 06 –Apr 07**	**Nov 06 –Sep 07**	**Apr 07 –Feb 08**	**Sep 07 –Jul 08**	**Feb 08 –Dec 08**
TTF									
Coefficients	APX	0.03	–0.04	0.02	0.03	–0.02	–0.03	–0.01	0.00
	TTF	0.83	0.78	0.93	0.91	0.92	0.82	0.85	0.96
Correlation	APX	–2.17%	–6.68%	–14.09%	–2.48%	1.98%	8.44%	15.28%	3.96%
	ZEE	33.40%	32.71%	28.18%	38.05%	52.12%	62.96%	49.55%	37.80%
	NBP	44.31%	48.76%	40.47%	41.71%	50.46%	58.63%	53.71%	39.16%
Zeebrugge									
Coefficients	APX	0.00	0.01	0.02	0.00	–0.05	–0.05	0.00	0.00
	TTF	0.79	0.70	0.45	0.47	0.45	0.20	0.30	0.91
	ZEE	0.33	0.40	0.65	0.64	0.65	0.71	0.56	0.25
Correlation	APX	–7.49%	–2.70%	–4.95%	2.59%	1.46%	5.49%	12.30%	–0.74%
	TTF	33.40%	32.71%	28.18%	38.05%	52.12%	62.96%	49.55%	37.80%
	NBP	52.73%	49.48%	57.30%	63.99%	71.83%	75.26%	70.48%	59.97%

FIGURE 5 Historical log spreads in a pairwise comparison.

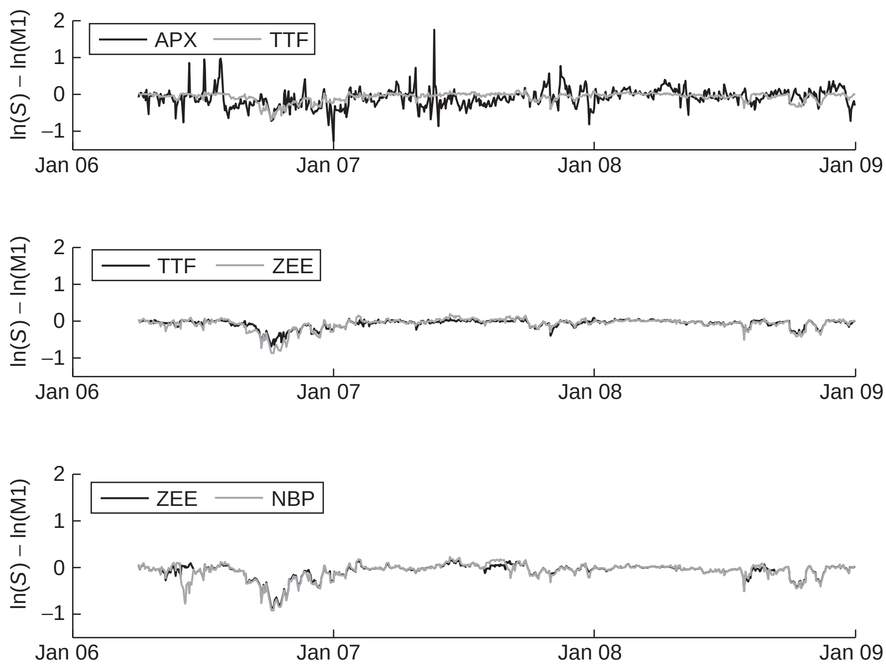

We conclude that it is worthwhile exploring a more advanced model of the gas market residuals, potentially one that involves regime switches. The exact market dynamics that create the non-normal behavior of the residuals is still to be investigated. One possible explanation is the regular occurrence of maintenance to pipelines.

9 CONCLUDING REMARKS

Our aim was to show how cointegration can be applied to capture the joint dynamics of multiple energy spot prices. We found clear indications of cointegration and a specific connection pattern between the gas spot prices of the TTF, Zeebrugge and NBP markets. For the connection between the gas spot markets and the power spot market APX, while it turned out that cointegration was also found, it was on the "forward time scale". We ascribe the day-to-day comovement of the gas spot prices to the strong physical connection of the markets through pipelines, whereas the looser connection between gas and power prices is presumably due to gas-fired power plants generally procuring on a longer-term basis.

FIGURE 6 Simulated log spreads in a pairwise comparison.

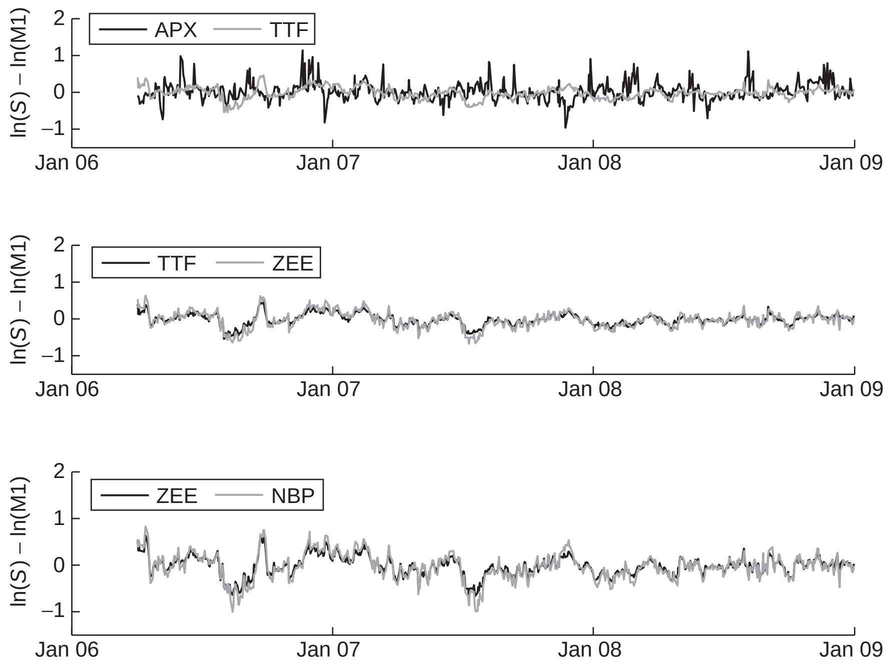

TABLE 6 Distribution statistics of historical and simulated data.

	Historical data				Simulated data			
	APX	**TTF**	**ZEE**	**NBP**	**APX**	**TTF**	**ZEE**	**NBP**
Log spread								
Mean	–0.04	–0.06	–0.07	–0.08	0.00	0.00	–0.01	–0.01
Standard deviation	0.26	0.13	0.17	0.19	0.24	0.16	0.23	0.30
Skewness	0.76	0.90	–0.11	–0.25	0.70	–0.02	–0.05	–0.07
Kurtosis	7.56	19.62	13.30	9.01	6.33	3.93	3.58	3.47
Log spread changes								
Mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Standard deviation	0.22	0.07	0.07	0.08	0.23	0.07	0.11	0.14
Skewness	–0.18	–1.15	–0.91	–0.58	–0.09	–0.33	–0.42	–0.48
Kurtosis	13.92	22.65	18.70	13.21	9.11	5.17	4.88	5.04

With our specification of a cointegrating model we were able to separate spot from forward effects and capture the strong connections of the gas markets. By defining a spot mean-reversion level on the basis of month-ahead forward prices we plausibly captured the comovement of the gas spot prices as cointegration of their spot–M1 spreads. The cointegration model is clearly structured and is flexible enough to be employed as an application framework: different spot and forward model types and a variable number of markets can easily be combined to meet the changing day-to-day requirements of energy trading.

A number of potential topics for further research follow from our study. The empirical results show that the GARCH$(1, 1)$ volatility model does not fully capture the complexity of the historical gas market distributions with significant skewness and kurtosis. A second interesting topic for further research is the linkage between spikes of power spot prices on various markets. It is not obvious how one would correlate regime probabilities, for example. Finally, the observed comovement of month-ahead prices, in particular between gas and power prices, calls for an appropriate cointegrated forward price model. The (increasing) interconnection of energy markets motivated us to develop a separate cointegration framework for forward markets as well. It includes related commodities, such as emission allowances, coal, oil and oil distillates.

The historical record of about 10 years of energy trading in Europe shows that the markets are characterized by mutual interactions in price dynamics. The joint dynamics stem from various sources, ranging from elementary physical connections to weather effects and economic conditions. With the many challenges that still need to be addressed to capture those dynamics in adequate price models, we hope this paper provides an interesting approach to capturing joint spot and forward price dynamics.

REFERENCES

Blanco, C., Soronow, D., and Stefiszyn, P. (2002). Multi-factor models of the forward price curve (II). *Commodities Now* September, 80–83.

Boogert, A., and de Jong, C. (2008). Gas storage valuation using a Monte Carlo method. *Journal of Derivatives* **15**(3), 81–98.

Borak, S., and Weron, R. (2008). A semi-parametric factor model for electricity forward curve dynamics. *The Journal of Energy Markets* **1**(3), 3–16.

Brunetti, C., and Gilbert, C. L. (2000). Bivariate FIGARCH and fractional cointegration. *Journal of Empirical Finance* **7**, 509–530.

Bunn, D. W. (ed). (2004). *Modelling Prices in Competitive Electricity Markets*. John Wiley & Sons.

De Jong, C. (2006). The nature of power spikes: a regime-switch approach. *Studies in Nonlinear Dynamics and Econometrics* **10**(3), Article 3.

De Jong, C., and Huisman, R. (2003). Option pricing for power prices with spikes. *Energy Power Risk Management* **7**(11), 12–16.

De Vany, A. S., and Walls, W. D. (1999). Cointegration analysis of spot electricity prices: insights on transmission efficiency in the western US. *Energy Economics* **21**, 435–448.

Engle, R., and Granger, C. (1987). Cointegration and error correction: representation, estimation and testing. *Econometrica* **55**, 251–276.

Eydeland, A., and Wolyniec, K. (2003). *Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging*. John Wiley & Sons.

Figuerola-Ferretti, I., and Gonzalo, J. (2008). Modelling and measuring price discovery on the NYMEX and IPE crude oil markets. Working Paper, Universidad Carlos III de Madrid.

Gjølberg, O., and Johnsen, T. (1999). Risk management in the oil industry: can information on long-run equilibrium prices be utilized? *Energy Economics* **21**, 517–527.

Johansen, S. (1988). Statistical analysis of cointegrating vectors. *Journal of Economic Dynamics and Control* **12**, 231–254.

Koekebakker, S., and Ollmar, F. (2005). Forward curve dynamics in the Nordic electricity market. *Managerial Finance* **31**, 73–94.

Lin, S. X., and Tamvakis, M. (2001). Spillover effects in energy futures markets. *Energy Economics* **23**, 43–56.

Longstaff, F. A., and Wang, A. W. (2004). Electricity forward prices: a high-frequency empirical analysis. *Journal of Finance* **59**(4), 1,877–1,900.

Lucia, J., and Schwartz, E. S. (2002). Electricity prices and power derivatives: evidence from the Nordic power exchange. *Review of Derivatives Research* **5**(1), 5–50.

Niemi, J. (2003). Cointegration and error correction modelling of agricultural commodity trade: the case of ASEAN agricultural exports to the EU. PhD dissertation, MTT Agrifood Research Finland.

Pindyck, R. (1999). The long-run evolution of energy prices. *Energy Journal* **20**(2), 1–27.

Pindyck, R. (2001). The dynamics of commodity spot and futures markets: a primer. *Energy Journal* **22**(3), 1–29.

Serletis, A., and Herbert, J. (1999). The message in North American energy prices. *Energy Economics* **21**(5), 471–483.

Spargoli, F., and Zagaglia, P. (2008). The comovements along the forward of natural gas futures: a structural view. *The Journal of Energy Markets* **1**(3), 17–35.

Urbain, J. P. (1993). *Exogeneity in Error Correction Models*. Springer.

Weron, R. (2006). *Modeling and Forecasting Electricity Loads and Prices: a Statistical Approach*. John Wiley & Sons.

Westgaard, S., Faria, E., and Fleten, S. E. (2008). Price dynamics of natural gas components: empirical evidence. *The Journal of Energy Markets* **1**(3), 37–68.