Modelling the Volatility of Spot Electricity Prices

Nektaria V. Karakatsani and Derek W. Bunn[1]

Department of Decision Sciences
London Business School

March 2004

Abstract

This paper presents a structural approach to model stochastic volatility in spot electricity prices. The peculiarities of electricity imply a complex structure, present both in price levels and volatility, which although critical for market and risk assessments, is neglected in stylised models and remains non trivial to model. In this methodology, prices are first detached from systematic components, such as economic fundamentals, risk measures, strategic and market design effects. Then, four *alternative* approaches are presented, where residual volatility is attributed to: i) the non-linear impacts of fundamentals, i.e. *GLS heteroscedasticity,* ii) the asymmetric volatility responses to lagged price shocks, i.e. a regression + *TGARCH* structure, iii) the *evolution* of the underlying price model due to market adaptation, i.e. *time-varying* regression effects and iv) the alteration of price structure during temporal market irregularities, i.e. *regime-switching* regression dynamics. Each alternative is motivated by different aspects of agent behaviour, but all derive stochastic volatility assuming a non-linear, *structural* specification for either the price formulation process (iii, iv) or the random shocks (i,ii). Implementation of this modelling to the UK market reveals strategic behaviour in agent reactions to shocks, with significant intra-day variation, and suggests that volatility inferences are sensitive to the assumed price model. For instance, GARCH effects diminish after adjusting for the time-varying price structure.

JEL: G14, L94
Keywords: Electricity, Spot Price Volatility, Heteroscedasticity, Time-varying parameter regression, Regime-switching, Market power

1. Introduction

Following the worldwide trend of restructuring public utilities during the 1990s, electricity has emerged as an actively traded commodity in spot, forward and derivatives markets. The most mature markets are those of the UK and Scandinavia, which started their operation at the beginning of the 1990s, followed shortly by Australasia and several S. American countries, and towards the end of the decade, by

[1] London Business School, Sussex Place, Regent's Park, London, NW1 4SA, UK.
Email: nkarakatsani@lbs.ac.uk, dbunn@london.edu
The authors gratefully acknowledge financial support from ESRC and the National Grid Company.

Spain, Germany, the Netherlands and some US States. Electricity prices have developed salient and general characteristics[3], most notably that of spot volatility, orders of magnitude higher than in financial assets and other commodities. Induced[4] by physical constraints and perhaps by generators' strategic behaviour, erratic volatility dynamics are present not only at the high-frequency level of trading periods but also the aggregated daily level, as illustrated for the UK in Figure 1. This volatility poses complications for hedging and security of supply, but may suggest profitable strategies for those agents who understand and can anticipate[5] its complexity. Although the proper specification remains challenging, stochastic volatility models are fundamental for trading, production scheduling, derivatives pricing, capacity investments and generation asset evaluation. Furthermore, volatility models, if linked to economic fundamentals or strategic effects, can elucidate agent reactions to shocks, and thus, reveal aspects of market performance of interest to regulators.

In the research literature, stylised stochastic models, inspired from financial markets and adapted to electricity, replicate some of the *statistical* price peculiarities but still, disregard the stochastic nature of volatility or do not clarify its *causalities* and *structural* properties. Thus, price models tend to introduce a jump component for spikes but, in order to facilitate analytical derivatives formulae, adopt the unrealistic assumption of *constant* volatility for the regular price process (Johnson and Barz, 1999; Lucia and Schwartz, 2002). Although this simplification is corrected in Deng (2000), where volatility is specified as a stochastic, mean-reverting process, jump-diffusion models do not disentangle the effects of mean-reversion and jump reversal (Huisman and Mahieu, 2001) implying a mis-specification of volatility. An alternative class of models to jump-diffusion, regime-switching models (Ethier and Mount, 1998; Huisman and Mahieu, 2001), postulate that volatility alternates stochastically, according to a Markovian process, between distinct values, estimated with probabilistic inference. In this stylised framework, volatility tends to be mis-specified due to implicit restrictions on the price process, such as stationarity under both regimes, instant reversion to normal levels after an episode or constant transition probabilities. In addition, the Markovian assumption could be restrictive given the nature of the occurrence of spikes, which could more explicitly be signified by a market variable, such as expected capacity surplus over predicted demand.

Conventional forms of time-varying volatility, such as conditional heteroscedasticity models (GARCH), although intuitively appealing, derive erroneous

[3] These include mean-reversion to a long-run level, multi-scale seasonality (intra-day, weekly, seasonal), calendar effects, erratic extreme behaviour with fast-reverting spikes as opposed to "smooth" regime-switching,, non-normality manifested as positive skewness and leptokurtosis, unstable correlations with fuel prices due to the alternation of marginal plant technologies.

[4] Due to the instantaneous nature of the commodity, spot volatility cannot be smoothed with economic inventory but remains exposed to real-time uncertainties, technical or strategic, such as plant outages, interconnector failures or demand shocks. The effect of non-storability is amplified by the limited demand elasticity to price in the short-term and oligopolistic market structures. In this setting, portfolio generators have the ability to behave strategically and induce volatility in the spot market in order to create incentives for forward contracting and increase risk premia. Another potential source of volatility is the presence of multiple, parallel markets for electricity trading, with limited volume and inadequate hedging instruments to link electricity with fuel prices.

[5] This prospective becomes more appealing given the currently introduced volatility swaps, which pose the challenge of correct pricing.

results for electricity prices (Duffie et. al., 1998), which is attributed to the presence of extreme values. These complications however, could be the result of a mis-specified price process and are indeed reduced in the presence of a richer price specification, as in Escribano et al. (2001), where the price model specifies mean-reversion, jump-diffusion and seasonality in the deterministic component and jump intensity. Applying GARCH modelling, Knittel and Roberts (2001) document asymmetric responses of volatility to positive and negative shocks and an inverse to financial assets leverage effect, but no explanation is suggested for this interesting idiosyncrasy.

Albeit insightful for medium-term price simulations and derivatives pricing, stylised stochastic models, present limitations when adopted for market assessments or short-term trading. Retaining too high a level of analysis, the fundamental sources of spot volatility are not addressed and hence, our understanding of agent reactions to shocks and of the way information is processed remains limited. As the links of volatility to market fundamentals are ignored, a significant component of uncertainty is retained, which, if further modelled, could dramatically reduce short-term risk exposure. Stylised models disregard the stage of model validation and do not allow for market or agent specificities, such as the integration of private expectations into the econometric specification. This synthesis is however critical for short-term predictions and, given the information asymmetries prevailing in electricity markets, translates to an obvious strategic advantage. Although an empirical approach that models volatility responses to fundamentals has not emerged yet, empirical evidence alludes to the presence of a complex underlying structure. Duffie et al. (1998) suggest that stochastic volatility models should account for[6] price levels, trading volume and spreads between spot and forward prices, whereas Knittel and Roberts (2001) generally emphasise the need for structural price modelling.

In this paper, we have adopted an econometric analysis of UK half-hourly spot prices, from June 2001 to April 2002, along with several fundamental and strategic variables, in order to model sources of stochastic volatility and clarify agent reactions to shocks. The extent to which market prices, after the introduction of the New Electricity Trading Arrangements (NETA) in 2001, were cost reflective, responded to risk measures, or manifested some forms of strategic pricing was discussed in Karakatsani and Bunn (2004). Despite the low price era, and its fundamental explanation, the daily price dynamics revealed that prices were particularly sensitive to margin variations and capacity withdrawal seemed plausible, throughout the day and most intensely around the evening demand peak. Still, market power was not exercised to its full potential. Extending this research on market efficiency, this paper specifies regression models for price evolution, of static, time-varying or regime-switching specification, and focuses on the stochastic and structural properties of the residual volatility that relates to them. This approach perceives volatility as an intrinsic, unobservable measure of price movement intensity, linked to the flow of financial activity and information, and simultaneously a reflection of incomplete modelling.

[6] The first element can be represented with GARCH modelling, the second seems plausible and consistent with evidence in other energy markets, whereas the last requires some clarification, as forward bias is more an implication rather than a causal effect of volatility.

After disentangling spot prices from systematic components, such as economic fundamentals, strategic, risk-reflecting and market design effects, four alternative sources of uncertainty are suggested to explain volatility behaviour; i) the impacts of market fundamentals on unconditional residual variance, ii) the asymmetric influences of lagged shocks on conditional residual variance, iii) the adaptation of price structure due to learning and iv) the occurrence of abnormal market states because of market power abuse. Each modelling approach provides an alternative interpretation for the non-linear nature of volatility and focuses on distinct statistical properties of the underlying process. The first two formulations impose complexity on the volatility specification, whereas the last two on the price formulation process. All reveal aspects of market performance and agent behaviour and have the potential to reduce short-term risk compared to stylised models, as they allow agents to include their predictions for market fundamentals. All approaches are plausible depending on market conditions, whereas hybrids, albeit theoretically appealing, entail convergence complications. As half-hourly trading periods are differentiated w.r.t. technical, economic and strategic characteristics, they are modelled with separate price specifications. This allows appealing intra-day properties of volatility structure to emerge, which are obscured in aggregated analyses.

In terms of methodology, the paper raises the issue of how sensitive to price specification volatility modelling is. The results suggest that, in the presence of multiple complexities in the price process, an adequate representation of price structure is crucial for the correct specification of volatility. Specifically, the apparent GARCH effects are concealed or exaggerated depending on the price model and seem to *diminish*, when the adaptive price structure or unconditional heteroscedasticity are accounted for. In a stylised framework, this indicates that modelling the dynamic structure of the price process (e.g. an autoregressive model with *time-varying* parameters) may describe volatility dynamics better than a price model with complex variance structure (e.g. AR + asymmetric GARCH).

The paper is structured as follows. Section 2 introduces a rich regression model for spot electricity prices, based on which parsimonious specifications are derived, distinct for each load period. In section 3, the *non-linear responses* of unconditional residual variance to economic and strategic impacts are modelled followed by the autoregressive structure of conditional residual variance. In section 4, time-varying volatility is linked with the evolving structural impacts on prices, as market adapts to endogenous or exogenous changes. Section 5, investigates the reaction of volatility to temporal market abnormalities. Section 6 concludes the paper.

2. Structural Price Specification

Volatility analysis is especially pertinent to the reformed UK spot market under the 2001 New Electricity Trading Arrangements (NETA), which induced new risks for market participants, some of which are unhedgeable. Among the regulator's deliberate intents was to increase the volatility in short-term markets, in order to reward flexible capacity and encourage forward contracts. Thus, we would expect a rather complex volatility process in the spot market, where agent behaviour would reflect both the ten year's mature experience of wholesale power trading, and the new learning associated with the market mechanism change.

2.1 Data Set

The market regime preceding NETA, the "Pool", was a compulsory spot market with uniform pricing, where dispatch was derived day-ahead with a cost-minimisation algorithm on offers submitted by generators. As prices were specified on a day-ahead basis, they were insensitive to real-time uncertainties and hence, not particularly volatile. Forward contracts allowed the diversification and hedging of spot price risk. NETA introduced a market institution with bilateral trading and discriminatory prices in a sequence of voluntary, un-administered markets with contract horizons up to several years ahead. More than 95% of electricity is currently traded forward in over-the-counter markets. For adjustments of contractual positions close to real time, Power Exchanges (PXs) have emerged. These operate on a day-ahead basis up to Gate Closure, which was initially defined as 3½ hours before real time and reduced to one hour in July 2002. At Gate Closure, participants notify their final physical positions (FPN) to the System Operator. After this point and in order to retain system stability, the System Operator administers a balancing mechanism (BM), where generators and suppliers submit bids and offers to deviate from their declared positions at Gate Closure. Imbalances, i.e. deviations between notified and ex-post metered positions of firms are penalised with dual discriminatory pricing. This renders unhedgable risk with energy deficiency in the BM much more costly than energy surplus.

In this bilateral environment, the structural properties of short-term markets were expected to influence prices in the preceding bilateral markets and seem the appropriate initial point for market analyses (Sweeting, 2000). The UKPX, analysed here, is the day-ahead spot market that provides half-hourly spot indicators, perceived as a replacement of the Pool Purchase Price. The functions of UKPX include physical delivery, adjustment of contractual positions, forward contracts and derivatives linked with the spot index.

The data consists of half-hourly values for the UKPX spot prices and the explanatory variables of their variation. For a given day, load Period 1 is defined as 23.00-23.30 (prior to the day), period 2 as 23.30-0.00, period 3 as 0.00-0.30 and so on up to period 48 (22.30-23.00). The sampling period was specified as 6th June, 2001-1st April, 2002 including 300 days. Considering 10 months of data for each load period was sufficient to derive reliable estimates and induced sufficient stationarity in the demand and margin series. Although the reforms were implemented in 27th March, 2001, the first 2 months of NETA were disregarded due to the pronounced market instability and the data quality issues that emerged. These mainly involved the occurrence of 'artificial' price spikes due to mistakes and numerical deficiencies of the price algorithm. June 2001 was suggested by industrial analysts as an appropriate initial point for representative analyses.

The half-hourly price series[7] displayed the typical empirical features of spot electricity markets, i.e. large volatility, positive skewness and excess kurtosis. The daily average prices, average daily price profile, inter- and intra-day volatilities, displayed in Figure 1, reveal the complexity of spot price dynamics, despite the

[7] Pilipovic (1998) showed that the logarithms of daily average electricity prices are normally distributed. In this high-frequency study, the emphasis was instead on half-hourly prices and thus the log transformation was not adopted. Despite its variance stabilizing properties, wherever applied, it did not alter the inferences and complicated the interpretation of the regression coefficients.

relatively short sampling period. In Figure 2, half-hourly prices for selected periods illustrate the rich intra-day variation in price evolution and the additional modelling complexity that this poses as well as the potential for misleading inferences, when diurnal patterns are smoothed with averaging. Stationarity tests (Augmented Dickey-Fuller and Phillips-Perron tests) for prices in each load period, after adjusting for serial correlation, rejected the presence of a unit root at the 5% significance level. The rejection of the random walk hypothesis motivated the detailed exploration of price structure.

2.2 Influential Variables

The modelling approach involves the formulation of a regression model for the evolution of spot prices, with static, time-varying or regime-switching specification, and the study or explicit modelling of the residual variance that emerges. A non-trivial issue is to define or quantify the factors reflected upon prices, such as economic fundamentals, plant constraints, strategic effects, risk perceptions, trading inefficiencies, learning, forward trading and market design implications.

A proper specification[8] of *Demand* is crucial, both in itself as a fundamental driver of daily price variation, and in order to formulate a well-specified background from which to estimate properly other, perhaps more subtle, influences on price. In the context of electricity markets, demand can be perceived as an exogenous variable to price, because of the absence of demand elasticity in the short-term. In our application, the non-linear demand effect was identified as a *quadratic* polynomial. To reflect the timing of the spot market and in order to avoid endogeneity, *Demand* was defined as an expectation, the 12 p.m. day-ahead forecast conducted by the National Grid. Due to their high correlation, the coexistence of the two demand terms would lead to an ill-conditioned matrix. To resolve the collinearity, the demand polynomial was decomposed into two orthonormal functions.

Furthermore, since the existence of a balancing mechanism in NETA could induce the backward migration of some pricing of plant dynamics into the preceding PX trading, the *Slope* and *Curvature*[9] of demand were also considered. More specifically, *Demand Slope,* the rate of change in demand, could be particularly influential and represent the periods when the more flexible plant is able to achieve higher prices. *Demand Variation*, due to temporal, weather and consumption patterns, imposes difficulties in load prediction and plant scheduling and eventually implies balancing costs. In addition, unanticipated demand paths influence agents' risk attitudes and, given NETA's asymmetric penalty for energy Imbalances, could encourage suppliers' over-contracting. However, the notion of demand uncertainty is not obvious to quantify. Relevant measures include the unexpected demand derived from a predictive model or the historic volatility of demand. Here, it is assumed that market participants update their perceptions about demand fluctuations considering the

---

[8] Demand appears as a state-variable in equilibrium stochastic models (e.g. Eydeland and Geman, 1998), a critical variable in threshold autoregressive models (Stevenson, 2001), a causal factor in neural networks and linear regression (e.g. Vucetic et al., 1999). The last specification involves a third order demand polynomial.

[9] Demand changes continuously but for simplicity, *Demand Slope* (*Curvature*) was approximated by the rate of change of demand (demand slope) in successive half-hour periods, i.e. by the first (second) differences of half-hourly demand measurements.

sequence of the 7 most recent demand values for each period. This time-horizon was selected for meteorological reasons and ensures always the presence of all weekdays and a weekend in a trader's database. *Demand Volatility* was then, defined as the coefficient of variation, i.e. (standard deviation/mean) in a weekly moving window. Due to large demand fluctuations across the year, this standardisation was essential in order to avoid misleading inferences. *Demand Forecast Error*[10] by NGC is a cause for over or under-contracting and thus, imbalance. It was defined as Actual Demand minus the 12pm day-ahead Forecast.

Market information and possible strategic effects were reflected in the following variables:

i) *Margin* is a measure of excess generation capacity, defined as the aggregated maximum possible output (the final notification at "gate closure") minus the day-ahead demand forecast from the National Grid.
ii) *Expected Imbalance,* is defined as Indicated Generation minus Predicted Demand at Gate Closure.
iii) *Scarcity* is derived from the *Ratio=Margin/Demand* (Visudhiphan and Illic, 2000) as: $\max\{Lower\ Quartile\ of\ Ratio - Ratio, 0\}$, where the lower quartile is calculated from the historic distribution of Ratio in each load period. This variable is intended to capture the steep impact of capacity surplus on price after a threshold.

Historic market conditions were captured by spot *price* for the same load period on the previous day and week as well as *daily average* price on the previous day. The latter created the desired link between by-period bidding and signals from the entire day. *Price Volatility*, an index of instability and risk, was defined similarly with demand volatility, as the coefficient of variation of prices in the preceding week. Finally, *Spread* was included, defined as the difference between the two balancing prices for insufficient and excess capacity[11], and representing *unhedgable risk* in the balancing mechanism. As a measure of risk exposure, it could impact on forward premia and be manifested in spot prices. Although the value of spread on the day is derived after the spot price, *Spread* on the previous day could still signal relevant risks and influence bidding, particularly if there was a tendency in the market (due to speculators, regulator or grid activities) towards smoothing imbalance risk. Alternatively, one could define the differences *PX-SSP* and *SBP-PX*, two measures of arbitrage that reflect the value or cost of flexibility close to real-time. These indices are essential for contract evaluations and informative about the relative attractiveness of spot and balancing markets.

Each load period displays a rather distinct price profile reflecting the daily variation of: demand, operational constraints and costs implied by different plant technologies, market depth and potential for market power. In order to control for these dissimilarities, the modelling was implemented *separately for each load period*. This was also inspired by the extensive research on demand forecasting which has

---

[10] The prediction error could, in principle, be correlated with demand and incur multi-collinearity, but this did not occur in our application.
[11] In the balancing mechanism, *Spread* is the difference between the System Buy Price (SBP) and the System Sell Price (SSP).

generally favoured this multi-model approach for accurately forecasting daily demand (Bunn, 2000). The different demand profiles for weekends would be expected to induce systematic elements in the evolution of demand, margin and hence prices, and also a shift in the morning and night peaks, compared to weekdays. However, separate modelling of weekdays, Saturdays and Sundays suggested that the above issues did not alter the inferences, possibly due to a well-specified inclusion of demand and margin in our model. *Seasonality,* however, was important as a proxy for the yearly pattern of fuel prices and approximated with a sinusoidal function peaking in winter.

The linear regression model[12] for spot prices is specified as:

$$P_{jt} = X'_{jt}\beta_j + \varepsilon_{jt}\,,\ \varepsilon_{jt} \sim N(0,\sigma_j)$$

where, P_{jt} denotes the spot price on day t and load period j, $t = 1,2,..,T$ and $j=1,2,...,48$, β_j a 16x1 vector of parameters, X_{jt} a 16x1 vector of exogenous explanatory variables, defined after preliminary analysis as: X_{jt} = *(1, P* $_{j(t-1)}$*, P* $_{j(t-7)}$*, Average P* $_{t-1}$*, Spread* $_{j(t-1)}$*, Price Volatility* $_{j(t-1)}$*, Demand* $_{jt}$ *(Linear, Quadratic Term), Demand Slope* $_{jt}$*, Demand Curvature* $_{jt}$*, Demand Volatility* $_{j(t-1)}$*, Margin* $_{jt}$*, Margin* $_{j(t-1)}$*, Scarcity* $_{jt}$*, Time, Seasonal Component* $_{jt}$ $)'$ and ε_{jt} a random and serially uncorrelated error term.

This rich model revealed significant intra-day variation of the structural effects on prices (Karakatsani and Bunn, 2004). Therefore, in the volatility modelling that follows, distinct parsimonious models are selected for each trading period. These are not uniquely defined, as several specifications may be equally plausible for the data depending on the adopted optimality criterion and variable selection procedure (e.g. forward addition, backward elimination, two-direction stepwise, best subset selection). In addition, each volatility modelling approach poses different convergence complications even for the same load period. The formulations presented in the illustrative examples are the more robust for the particular setting.

3. Heteroscedasticity Modelling

After detaching prices from fundamental structure with a parsimonious model, questions are raised about the structural properties of *residual variance*, $\mathrm{Var}(\varepsilon_t)$. This is an ex-post measure of price risk that reflects the model's uncertainty around fitted prices. Alternatively, it can be interpreted as the predictive uncertainty of a trader who formulates his price expectations with the regression model and knows

---

[12] In the above model with stochastic regressors, standard assumptions include:

i) $\{P_{jt}, X_{jt}\}$ is jointly stationary and ergodic, which precludes trending regressors.

ii) The regressors are predetermined, which excludes endogenous regressors but allows lagged dependent variables.

Under these assumptions, which were deemed to be plausible for the data, the OLS estimates are consistent and asymptotically normally distributed even under i.i.d., non- normal errors.

apriori[13] the model specification and the correct values of exogenous variables. In trading practice, an agent would apply the price model substituting the exogenous variables by his own predictions. If there is, in addition, an explicit relationship that links residual variance to market fundamentals and past shocks, the trader could derive also an estimate for his implicit price risk.

Systematic components in residual variance could arise due to inadequacies of the price model or heteroscedasticity, conditional or unconditional, of the random shocks. Although it is not feasible to disentangle the effects of price model mis-specification and heteroscedasticity, it is possible to assess their compounded implications on residual variance and quantify its responses to market fundamentals and past shocks.

3.1 Unconditional Heteroscedasticity

To clarify the *non-linear responses* of residual variance to strategic and economic effects, Generalised Least Squares (GLS) modelling is adopted. This explicit modelling of unconditional price heteroscedasticity allows a formal exploration of volatility hypotheses and indicates how agents react to uncertainty under different market conditions. Correcting the restrictive assumption of homogeneity implicit in the OLS price model, GLS estimation should induce more reliable inferences regarding the significance, sign and magnitude of structural effects on price levels.

The GLS regression model with non-spherical uncorrelated disturbances is specified as:

$$P_{jt} = X'_{jt}\beta_j + \varepsilon_{jt},$$

$$\varepsilon_j \sim N_n(0,\Sigma_j),\quad Var(\varepsilon_{jt}) = g(v_{jt})\ \ \Sigma_j = diag\{g(v_{jt})\} \neq I_n$$

where v denotes the covariate driving the variance, $g(v)$ the variance function, $\Sigma = E(\varepsilon\varepsilon')$ the error covariance matrix. A critical issue is the selection of the function g, which should replicate non-linearities similar to those emerging due to non-convex marginal costs and strategic behaviour. Two functions seem appealing:

i) The power formulation, defined as $g(v) = (\alpha + |v|^b)^2$, where *a, b* denote unknown parameters. The quadratic form allows for a rich volatility structure, not necessarily monotonic, which includes the linear dependence as a special case.

ii) The exponential formulation, defined as $g(v) = e^{tv}$, where t is an unknown parameter, replicates monotonic, kinked effects on volatility.

GLS modelling of spot prices across several trading periods suggested that residual variance presents an heterogeneous structure within the day. Still, clusters of periods with similar responses seemed to emerge according to their position in the demand curve and particularly the degree of demand stability or adjustment. This intra-day variation is illustrated with the peak load period 25 and the morning period 15, which represents a transitory stage for the demand curve when flexible plants start operating. The price models were parsimonious specifications selected with stepwise procedures from the rich model in section 2.2 and the profiles of residual volatility were explored

[13] The former assumption is plausible if the model equation is stable over time and derived from past data. The latter applies to traders with accurate information about future market fundamentals, which in our application reduces to the variable Capacity Margin.

with the power specification, which proved here to be more robust to the selection of initial values, less sensitive to outliers and easier to converge than the exponential model. The estimates of the parameter α were effectively zero in all the presented examples, which revealed monotonic impacts of the considered covariates on volatility. The significance and signs of the variables in the price model were quite robust across variance specification, which suggested that an OLS model is a valid approximation of price dynamics, even if GLS effects are ignored.

The results[14] for period 25 are summarised in Table 1. Residual volatility responded positively to price signals, displaying a linear dependence on *Lagged Price* and a parabolic, less steep than quadratic, link to *Expected Price.* These positive effects suggested an inverse to the "leverage effect" documented in financial markets, a peculiarity of electricity prices to be discussed later. The increase[15] in volatility as margin declines indicated the diversity of strategies and possibly arbitrariness of bidding under relative scarcity. Residual volatility and demand were almost inversely proportional and this hyperbolic link could be attributed to several modes of agent behaviour. The augmented uncertainty during low loads might signal the decline of prices to unexpected levels, below their fundamental value, possibly due to over-supply. In contrast, high demand could motivate similar price expectations and perhaps easier collusion, which translated to more predictable prices in the context of a price model that addresses strategic effects.

Alternatively, the demand effect could relate to the conservative bidding of flexible plants, which achieve excessive profits during the morning and evening peaks and simply wish to retain their operation in intermediate periods. If demand is low on a specific day, these stations' reward declines even further, which possibly creates incentives for more aggressive and unpredictable bids between peaks. Another conjecture is suppliers' over-contracting under high demand, when exposure to imbalance prices is particularly penal, which induces more activity in the spot market and more representative prices. As opposed to the previous remarks, empirical evidence suggests the increase of price volatility with demand. This contradiction is only superficial, as the volatility measure analysed here is not the statistical price variance but the residual uncertainty *after* fundamental structure has been subtracted.

Finally, instead of a structural volatility equation, serial correlation was assumed for the innovation terms ε_t in the form of an AR(1) process, where $\text{cov}(\varepsilon_t, \varepsilon_{t-1}) = \rho, \forall t = 1,...,n.$ This specification seemed plausible but autocorrelation diminished when a structural model was assumed for variance, which suggests that AR effects are simply a surrogate for omitted factors. The adequacy of a variance

---

[14] The parameters in the variance function are estimated with an iterative procedure by maximising the marginal likelihood of the residuals from the least-squares price model. The regression coefficients are subsequently estimated by maximum likelihood assuming that the variance structure is known, as proposed in McCullagh and Nelder (1989). The derived $\hat{\beta}_{GLS}$ is simply the weighted-least-squares estimator: $\hat{\beta}_{GLS} = (X'\hat{\Sigma}^{-1}X)^{-1}X'\hat{\Sigma}^{-1}P$

with estimated covariance matrix: $\hat{V}ar(\hat{\beta}_{GLS}) = (X'\hat{\Sigma}^{-1}X)^{-1}$.

[15] The same signs of the coefficients associated with demand and margin are consistent, as the linear correlation between the two variables is low for this period (-0.29).

specification can be assessed with the value of the Log Likelihood function and the Residual Standard Error (RSS).

Table 1. PX Price, Period 25. Variance Structure from GLS Modelling.

Price Model	P_t ~ *Constant* + P_{t-1} + MP_{t-1} + SBP_{t-1} + *Price Volatility* + *Demand*$_t$ + *Demand Curvature*$_t$ + *Margin* $_t$ + ε_t		
Volatility Model	Var (ε_t) = g (v) = v^{2b}		
Covariate v	Coefficient (2b)	Log Likelihood	RSE
Expected Price	1.62	-725.89	0.23
Demand $_t$	-0.88	-730.22	28.16
Margin $_t$	-0.44	-729.84	22.12
P_{t-1}	1.06	-726.80	0.54
Autocorrelation (ρ)	0.26	-736.32	2.77

Inferences about the variance structure in period 15 are summarised in Table 2. The effects of Margin and Demand remained negative, as in period 25, but their magnitudes were considerably different, with the former portraying a limited impact of the order of ¼ root and the latter a more dramatic, almost inverse quadratic effect. One interpretation is that flexible plants, which start their operation at this stage and require a premium for the implicit risk, behave in a more predictable manner as demand increases, possibly because cost implications are more dramatic or collusion easier. Finally, Expected and Lagged Price displayed negative effects on residual volatility, which implies more uncertainty around low prices, the opposite condition of the peak period 25. This dissimilarity in the response of volatility to price signals possibly arises from the different position of the two periods in the demand curve. Firm behaviour seems to be dominated by profit maximisation in period 15 vs. operating constraints in period 25.

Table 2. PX Price, Period 15. Variance Structure from GLS Modelling.

Price Model	P_t ~*Constant*+P_{t-1} + P_{t-7} + *Demand* $_t$+ *Demand Slope* $_t$+ *Margin* $_t$ + *Margin* $_{t-1}$ + ε_t		
Volatility Model	Var (ε_t) = g (v) = v^{2b}		
Covariate v	Coefficient (2b)	Log Likelihood	RSE
Expected Price	- 0.62	-708.80	5.42
Demand $_t$	- 1.8	-705.87	27
Margin $_t$	-0.28	-709.59	9.89
P_{t-1}	-0.32	-709.02	3.66
Autocorrelation (ρ)	0.52	-698.27	2.73

3.2 Conditional Heteroscedasticity

The previous modelling suggested that unconditional heteroscedasticity is a plausible explanation for the erratic volatility dynamics. The volatility responses to fundamentals tended to be non-linear and reversed sign within the day reflecting dynamic operating constraints and firm strategies or inadequacies of the price model. It should be emphasised that, after accounting for the impact of fundamentals on volatility, GLS residuals presented no significant GARCH structure. This implied that the autoregressive volatility structure, observed in practice, is eliminated, if a fundamental explanation for variance is postulated. Although a surrogate for omitted factors, conditional heteroscedasticity is still an appealing price property for trading and consistent with the realised paths of electricity prices, often characterised by periods of high instability followed by periods of relative tranquillity. This volatility clustering implies some predictability, which could be enhanced when accounting for asymmetries and non-linearities in the response of volatility to news. The occurrence of jumps however, prohibits the applicability of GARCH models. Duffie et al. (1998) conclude that erroneous results, such as integrated volatility processes, are usually derived due to the bias introduced by extreme prices.

Such undesirable effects are constrained however, if regression-GARCH modelling is adopted, as GARCH effects are then explored in regression disturbances and not pure prices. Detached from systematic components, which may present extreme values for certain values of the covariates, regression residuals are smoother, even during spikes, and allow conventional volatility modelling. The regression-GARCH approach introduces implicitly a distinction between *unexpected* shocks outside the model boundary, which reflect news and create volatility, and *extreme* values of influential variables, possibly anticipated and persistent for a time period, which induce high prices.

The regression model with GARCH (1,1) normal errors is defined as:

$$P_t = X_t'\beta + \varepsilon_t$$
$$\varepsilon_t = \sqrt{h_t}u_t \ , u_t \sim N(0,1)$$
$$h_t = a_o + a_1\varepsilon_{t-1}^2 + a_2 h_{t-1}, \ \ a_o > 0, a_1, a_2 \geq 0.$$

where P_t is the spot price in a specific load period (the subscript j is omitted for simplicity), ε_t an i.i.d. serially uncorrelated innovation process, stationary under the condition $a_1 + a_2 < 1$, with conditional variance $h_t = Var(\varepsilon_t | I_{t-1}) = E(\varepsilon_t^2 | \ I_{t-1})$, a time-varying, positive and measurable function of I_{t-1}, the information set at time t-1. In order to account for the leptokurtosis of electricity prices, a standardized Student's t distribution is assumed for u_t.

Having defined the Regression-GARCH model, the volatility properties of spot electricity prices are explored with different specifications for the conditional mean and conditional variance. The equations, displayed in Table 3, vary w.r.t. adequacy of price description and complexity. The naïve price model I, conventional in financial practice, implies that prices follow a leptokurtic distribution fluctuating randomly around a long-run mean value. To reflect mean-reversion, model II imposes an AutoRegressive (AR) process with optimal number of lags w.r.t. to the AIC criterion. Equation III postulates a regression model intended to capture systematic price structure. Among the variance equations, IV implies a *symmetric* GARCH (1,1) structure for the residuals of the price model, i.e. assumes that positive and negative disturbances have the same impact on h_t. Model (V) imposes the Threshold GARCH (1,1) structure, introduced by Zakoian (1994), which allows for asymmetric effects of positive and negative lagged shocks on the conditional standard deviation. Positive news ($\varepsilon_{t-1} > 0$) have an effect of α_1 on $\sqrt{h_t}$, whereas negative news have an effect of $\alpha_1 + \gamma$. The price and variance equations are estimated jointly, as the former involves lagged values of the response. The likelihood function[16] is maximized via dual quasi-Newton. The starting values for the regression parameters are obtained from OLS estimation or Yule-Walker equations when autoregressive parameters are present.

Table 3. Price and Variance Specifications in GARCH Modelling.

Price Model	Variance Model
$P_t = c + \varepsilon_t$ (I)	$h_t = a_0 + \alpha_1 \varepsilon_{t-1}^2 + a_2 h_{t-1}$ (IV)
$P_t = c + L(P_t) + \varepsilon_t$ (II)	$\sqrt{h_t} = a_0 + \alpha_1 \varepsilon_{t-1}^2 + \gamma\, S_{t-1} \varepsilon_{t-1}^2 + a_2 \sqrt{h_{t-1}}$ (V)
$P_t = X_t \beta + \varepsilon_t$ (III)	

In the above, L denotes a distributed lag polynomial and S_{t-1} the indicator function: $S_{t-1} = \begin{cases} 1, & if\ \varepsilon_{t-1} < 0 \\ 0, & otherwise \end{cases}$.

Implementation of the modelling to UK spot prices showed that the specifications postulated for price do affect the identification of volatility dynamics and may even suggest conflicting conclusions. As an illustration, Tables 4- 5 present GARCH results for two peak periods with particularly volatile dynamics. For period 25, the optimal AR(6) model did not converge and was replaced by an AR(4). GARCH effects, revealed under other price specifications, proved insignificant after removing the structural component of prices. For period 35, the regression price model implied significant volatility asymmetry, which was however insignificant under AR models. To facilitate convergence, the optimal AR (24) was substituted by the robust AR(1) specification.

The results could be summarised as follows:

i) Within the Regression + GARCH model, the complications documented by Duffie et al. (1998) diminish. This could be attributed to the fact that high prices frequently emerge from the same structural model as regular prices but for extreme values of the covariates. Whereas the regression model could anticipate to some extent the abnormal price level, a non-structural model could erroneously perceive high prices as large shocks and thus, bias the estimation. More specifically, two types of results emerged. For some load periods, when a regression price specification was

---

[16] The Log likelihood function under the t error distribution is:

$$\log L = \sum_{t=1}^{T} \log(\Gamma(\frac{v+1}{2})) - \log(\Gamma(\frac{v}{2})) - \frac{1}{2}\log((v-2)h_t)$$ where Γ is the gamma function and v denotes the degrees of freedom in the conditional t distribution, an additional parameter to be estimated.

assumed, a strong autoregressive structure was detected in the variance. However, GARCH effects seemed insignificant or the model failed to converge, when prices were instead described with an AR process. In other cases, the opposite result was observed. The persistence of extreme but anticipated market conditions (e.g. high demand), which induced a series of high prices, was interpreted as persistence of random shocks by simplistic GARCH models. The above observations suggest the sensitivity of volatility inferences to the price model.

ii) Model I was entirely inadequate and led to erroneous results, such as explosive variance. Imposing an optimal AR model (II) was often sufficient to avoid unreasonable estimates, possibly because autocorrelation had a strong presence in the new market due to learning and trading inefficiencies. As it was expected, residual uncertainty was reduced significantly[17] when price levels were modelled with a structural formulation. This was reflected both in long-term (asymptotic variance) and short-term (impact of lagged shocks on conditional variance) measures of uncertainty. This could be partially attributed to the fact that realised values of exogenous variables, such as Margin, were used in the price model. Still, simulations of expected Margin suggested that this property of regression-GARCH modelling is retained in forecasting settings. This finding is crucial for day-ahead trading, as forward prices convey little information about intra-day spot price fluctuations, especially when reported only as aggregated indices (peak/baseload).

iii) In several trading periods, conditional variance seemed to respond asymmetrically to positive and negative past shocks. The negativity of the parameter γ suggested that the inverse of the leverage effect identified in financial assets applies to electricity prices. This means that positive disturbances have a stronger impact on volatility than negative ones, as documented in Knittels and Robert (2000) for California prices. One possible explanation is that in the case of stocks, low prices increase the leverage exposition of the firm and have a direct impact on the perceptions of stockholders; in contrast, in the case of electricity, prices higher than expected may attract regulatory intervention or new entry, both undesirable outcomes for generators and sufficient to create temporal uncertainty. An alternative interpretation suggests that as prices exceed their anticipated levels, often due to market power abuse, there is a lack of consensus about the true value of the commodity and a tendency towards more aggressive bidding. Speculators attracted by potential profits add to this dispersion of expectations, which causes more intense price fluctuations than during lower prices.

iv) When a covariate was introduced into the GARCH equation, such as Demand or Margin, the coefficients were insignificant or convergence was infeasible.

---

[17] Even when they involved more variables, AR+GARCH models indicated more volatile innovation processes than Regression +GARCH models. The more the extreme prices predicted by the regression model, the more dramatic the deviations in the results were.

Table 4. GARCH Modelling for PX Price, Period 25.

Price Equation	Model (I)	Model (II)	Model (III)
Variance Equation	$h_t^2 = a_0 + \alpha_1 \varepsilon_{t-1}^2 + a_2 h_{t-1}^2$ (IV)		
α_o	3.01 (0.03)	2.46 (0.01)	3.87 (0.002)
α_1	0.39 (0.003)	0.29 (0.007)	0.28 (0.015)
α_2	0.46 **(0.006)**	0.46 **(0.002)**	0.17 **(0.2)**
t-distribution (df)	23.56	6.21	10.75
Asymptotic St.Deviation	**4.60**	**3.20**	**2.63**

Table 5. GARCH Modelling for PX Price, Period 35.

Price Equation	Model (I)	Model (II)	Model (III)
Variance Equation	$\sqrt{h_t} = a_0 + \alpha_1 \varepsilon_{t-1}^2 + \gamma\, S_{t-1} \varepsilon_{t-1}^2 + a_2 \sqrt{h_{t-1}}$ (V)		
α_o	6.76 (0)	0.51 (0.06)	1.22 (0.01)
α_1	1.25 (0)	0.26 (0.02)	0.17 (0.02)
α_2	0.17 (0)	0.82 (0)	0.92 (0)
γ	**-0.86** (0.03)	**-0.18** (0.10)	**-0.51** (0.007)
t-distribution (df)	5.27	3.02	3.42
Asymptotic St. Deviation	**24.56**	**7.53**	**2.12**

4. Evolution of Price Structure

4. 1 Motivation

This section intends to assess whether the evident time-varying spot volatility reflects an evolving price structure, result of a highly repeated auction and continuous agent adjustments to changes in market structure and rules. A dynamic fundamental price analysis is thus proposed to reveal the direction towards which prices evolve, such as more cost-reflecting, risk-related or strategically abused levels. In order to follow the complex process of market adaptation, the structural price models assumed in previous sections are re-specified here with time-varying parameters.

An apriori assessment of the market orientation is not obvious. The half-hourly interactions among market participants and between them and the system operator induce *learning,* which reveals profitable strategies but simultaneously motivates more efficient reactions by the system operator. Under the new uncertainties, induced by NETA, *individual* actors or groups (generators vs. suppliers, speculators, generation technologies, NGC) update their utility functions, risk aversion parameters and strategies. It is however questionable whether the adjustments of several parties with conflicting interests are counteracted or move towards a similar direction, which attitude dominates and how the orientation of the entire market evolves as a result of these dynamic interactions. Assessing price convergence without a fundamental analysis might be misleading in the presence of seasonality, contradictory or unstable market signals, and multiple exogenous or endogenous shocks affecting trading.

4.2 Time-Varying Parameter Regression Model

In order to follow the evolution of structural effects on prices and hence, clarify the time-heterogeneity of volatility, a Time-Varying Parameter (TVP) regression model is specified:

$$P_{jt} = X'_{jt}\beta_{jt} + \varepsilon_{jt} \qquad \text{Measurement equation}$$

$$\beta_{jt} = \beta_{j(t-1)} + v_{jt} \qquad \text{Transition equation}$$

where, $\varepsilon_{jt} \sim i.i.d.N(0,\sigma^2_{\varepsilon_j})$, $v_{jt} = (v_{j1t}, v_{j2t},..., v_{jkt})'$, $v_{jt} \sim N_k(0,\Sigma_j)$, $E(\varepsilon_{jt}v_{jt}) = 0$ and $\Sigma_j = diag\{\sigma^2_{v_{jk}}\}$.

In the above state-space formulation, the regression coefficients are not unknown constants but latent, stochastic variables that follow random walks. This specification was plausible for several load periods, as indicated by stability[18] tests (typically at significance levels of 5-10%). This result is intuitive as many shocks during the sampling period had a permanent or cumulative rather than diminishing effect in market adaptation. Such shocks include rule modifications, regulator's announcements about the length of the Balancing Market and a new rule for price calculation, policies regarding renewables, mergers and acquisitions in the electricity industry in addition to the Enron collapse.

The state-space model was estimated with discrete Kalman Filter. The filter recursions indicate how a rational economic agent would revise his estimates of the model parameters in a Bayesian fashion within an environment of uncertainty, as new information becomes available. In our context, the recursions indicate how the market as a whole evolves attaching varying importance to fundamentals. The Kalman Filter Algorithm for the TVP model is described in Appendix I.

It should be emphasised that TVP models are differentiated from the ARCH class of models implemented in 3.2 with respect to the features of uncertainty they intend to capture (Kim and Nelson, 2001). In the latter type of models, changing uncertainty about the future is focused on the conditional heteroscedasticity in regression disturbances. In time-varying regression however, an agent's uncertainty about the future arises partially from future random terms. It also reflects uncertainty about current parameter values and the model's ability to link the present to the future. The uncertainty about current regression coefficients results in the changing conditional variance of price. This decomposition of uncertainty is captured in the equation for the variance of the conditional forecast error: $H_{jt|t-1} = X_{jt|t-1}P_{jt|t-1}X'_{jt|t-1} + \sigma^2_{\varepsilon_j}$, where $P_{jt|t-1}$ represents the degree of uncertainty associated with an inference on β_{jt} conditional on information up to time t-1.

[18] Two types of stability tests were performed; the homogeneity test (Brown, Durbin, Evans, 1995) against the alternative hypothesis of unstable regression coefficients and the Engle and Watson (1985) against the alternative hypothesis of random walk coefficients

4.3 Dynamics of Price Structure

Implementation of the TVP modelling to NETA elucidated several aspects of the market evolution process. Although clusters of periods with similar patterns of structural evolution emerged, intra-day variation was still considerable. This indicated that the fragmentation of trading across periods allowed the persistence of different structural trends within the day and delayed market convergence. Even when the variances of the effects were not statistically significant, the dynamic estimation procedure uncovered subtle details of the adaptation process, which were neglected when a static regression model was assumed. In general, adjustment effects remained strong up to December 2001 and were still present one year after the introduction of NETA. The analysis indicated a gradual shift of market orientation towards more sophisticated trading and possibly more cost-based prices with greater responsiveness to perceived risks. In contrast to the erratic dynamics of demand effects, the decreasing impact of margin was evident and signified a trend towards less strategic bidding, at least in the conventional sense. Some evidence of capacity withholding, implied by the negative effect of Lagged Margin, also diminished in January 2002. The *decline* of strategic impacts and autocorrelation indicated that the initially prevailing inefficiencies in the reformed market were progressively being eliminated to a large extent.

Typical stochastic patterns of the evolving price structure are illustrated with periods 25 and 35. The price models selected were again the more robust for the specific setting. Table 6 reports model formulae and parameter estimates for the illustrative models. Figures 3-4 depict Kalman Filter the dynamic regression coefficients conditional on information up to time t-1. The recursions illustrate how price sensitivities to various factors were revised during the sampling period. It is apparent that a significant proportion of price variation is due to the evolution of regression coefficients in the price equation.

In specific, the impact on spot price of price signals from the previous day and week decreased over time for various load periods. This aspect of bidding behaviour, displayed in Figures 3a-b and 4a-c, indicated increasing limitations in price forecasting with autoregressive models. One plausible interpretation is that bidding became progressively more *sophisticated* or more based on private data rather than historical prices. This was consistent with the market tendency towards vertical integration and within-firm trading for risk management. The alternative conjecture that the market became more efficient and gradually cancelled price autocorrelation seems quite unrealistic given the persistence of relative illiquidity. In contrast with past prices, the role of historic PX volatility increased dramatically for some peak periods. As Figure 3g illustrates, in period 25 the volatility impact on price, reflecting forward premia, tripled in March 2002 compared to July 2001. Even if the initial estimates were sensitive to the selected prior, the increasing trend was still obvious. This implied that the hedging of risk via the day-ahead market became more expensive over time in response to the unpredictable SBP and primarily, credit risks. The Enron collapse and the first signs of generators' bankruptcies created an insecure trading environment, where risks were converted to more expensive pricing. An alternative plausible view is that generators tended to exploit progressively suppliers' exposure to penal imbalance prices, at least occasionally. A final component manifested in this dynamic behaviour is the regulatory risk induced with NETA and demystified over time with the discussions about a new pricing scheme and the reduction of Gate Closure to one hour.

The response of prices to Margin evolved in a fairly uniform fashion across periods. As Figures 3g and 4g illustrate, the partially strategic effects of Margin *declined* systematically, despite the possibly favourable winter conditions and reached in March 2002 approximately 60% and 18% of their initial values in periods 25 and 35 respectively. This decline of strategic elements in day-ahead pricing could imply that the market is converging gradually to a competitive state given overcapacity and fragmented suppliers' more active role, generators' market structure or their intention to eliminate inefficiencies in the fear of eminent reforms. Standardisation of coefficients further revealed that for several periods margin *ceased* to be the most influential factor of spot prices after December 2002. In this respect, the market gradually reached more *cost-based* outcomes. The pattern displayed in Figure 4h arose in several periods and clarified the enigmatic effect of Lagged Margin on prices. Although negative in the beginning, as expected, the coefficient reversed sign during the winter, possibly manifesting capacity withholding. This non-competitive element of pricing was inverted in January 2002.

The evolution of the demand effect was particularly volatile without a consistent pattern across periods. The evident intra-day and annual instability reflected the absence of central cost-minimising dispatch and the implied mixture of different technologies, the alternation of marginal capacity across the year, the path of fuel costs and the variability in operational costs such as start-up, which under NETA would be internalised in bids conditional on daily scheduling. In periods of peak demand, such as 35, the impacts of the quadratic demand term increased dramatically in response to weather conditions and declined abruptly after January 2002 (Figures 20, 21). In shoulder periods however, such as 25, patterns were different. (Figures 3e, 3f). Both the linear and quadratic components of demand were mostly negative and extremely volatile but diminished or converged to zero. Although periodic cycles could be distinguished, a seasonal interpretation was not apparent. The most plausible speculation is that as spot prices were collapsing, flexible generators were becoming more reluctant to compromise with low prices for shoulder periods in order to retain their presence in the evening peak. As a result, demand was becoming less influential on bidding.

As discussed in section 3.2, strong ARCH effects were detected when price models with fixed coefficients were assumed. Serial correlation was not detected however in the squared forecast errors of the TVP specifications after adjusting them for the conditional TVP heteroscedasticity ($\mathrm{H}_{t/t-1}^{1/2} n_{t/t-1}$). This implies that the existence of GARCH effects could be due to the varying structural components of price levels. Figures 5-6 illustrate conditional standard deviation for periods 25 and 35, estimated from the previous TVP regression models and the corresponding Regression-GARCH models. Regarding the latter, asymmetric effects of positive and negative shocks, in the form of TGARCH (1,1), were significant for period 35 but not 25, where GARCH (1,1) was sufficient. It should be noted that the TVP estimates are apriori volatility expectations, based on information up to t-1, whereas the GARCH estimates are derived ex-post. Still, within the TVP framework, uncertainty is significantly reduced. TVP models seem to uncover much more subtleties of the volatility process, particularly in unstable time periods. Day 1 is slightly different in the two classes of models, as TVP required more observations for filter initialisation.

Table 6. PX Price, Periods 25 and 35. Parameter Estimates for the Time-Varying Regression Model.

Period 25		**Period 35**	
σ_{ε}^{2}	0.75	σ_{ε}^{2}	2.4
Variable	σ_{v_k}	Variable	σ_{v_k}
Intercept	0.20	Intercept	0.03
P_{t-1}	0.012*	P_{t-1}	0.12*
P_{t-7}	0.0026	P_{t-7}	0.19*
Spot Volatility	0.0075*	MP_{t-1}	0.008
Demand- Linear	0.27*	Demand- Linear	0.54*
Demand-Quadratic	0.011	Demand-Quadratic	0.63*
$Margin_t$	0.00038*	$Margin_t$	0.000032*
Demand Volatility	0.00036	$Margin_{t-1}$	0.000012*
Demand Curvature	0.00009	Demand Curvature	0.0005

The asterisk denotes a variance significantly different than zero at the 95% level.

5. Discontinuities in Price Structure

5.1 Motivation

The previous analysis suggested strategic manipulation of capacity to a certain extent throughout the day, but most intensely around the evening demand peak. This reveals some kind of selective agent behaviour, which induces discontinuous volatility, and may also appear within a single trading period. To uncover the stochastic dynamics of such selective behaviour, the magnitude of any strategic effects and their implications on volatility, regression regime-switching is adopted at the high frequency level. The more heuristic approach of simply analysing extreme prices would be insufficient, as price effects should be assessed after controlling for various fundamentals, i.e. in the context of a structural model, rather than purely from their levels.

One of the peculiarities of spot electricity prices is that they exhibit regular patterns disrupted by recurrent but aperiodic, fast-reverting "spikes", which induce severe financial risks. These extreme prices signify temporal market *irregularities*, such as unexpected weather/demand, technical shocks (e.g. plant/interconnector failures, contingencies in transmission networks), strategic behaviour, trading inefficiencies (e.g. illiquidity), cross-commodity leakages (e.g. fuel price explosion) or accumulation of credit risks. Each of the above causalities is linked with different attributes of the commodity (e.g. non-storability, demand inelasticity to price) or the market design and structure[19].

[19] Specifying the source of irregularity in each case would require an analysis of the values of the covariates in every abnormal case. Such a detailed evaluation exceeds the purposes of this study.

The modelling question explored in this section is whether the discontinuous price structure, arising from the temporal irregularities discussed above, could be sufficiently captured with a few *structural regimes* with distinct volatilities. In the context of electricity prices, regime-switching has been adopted to replicate the erratic market alternations between “normal” and “abnormal” equilibrium states of supply and demand. Existing models refer to daily average prices and often assume an autoregressive process under both regimes (Ethier and Mount, 1998; Deng, 2000), which introduces estimation bias, as it does not disentangle mean-reversion from spike reversal and incorrectly imposes stationarity on the irregular price process. A continuous-time process that corrects this mis-specification is derived by Kholodnyi (2000), where self-reversing non-Markovian spikes are added to a Markovian regular price process. The mis-specification is also eliminated in discrete time by (Huisman and Mahieu, 2001) with the assumption of three regimes, a regular state of mean-reverting price, a jump regime that creates the spike and a jump reversal regime that ensures with certainty price reversion to their previous normal level. This regime-transition structure is however restrictive, as it does not allow for consecutive spikes. This constraint is relaxed by de Jong and Huisman (2002), where a stable mean-reverting regime is proposed and an *independent* spike regime of log-normal prices, which implies closed-form solutions for option pricing.

In contrast to these stylised multiple-regime models, regime-switching is adopted here within a regression model and at a high-frequency level, separately for each trading period. This structural specification allows the replication of more realistic price paths for intra-day trading and primarily, addresses issues of market performance and agent conduct. In principle, the proposed modelling presents the following properties:

i) It provides a more adequate description and potentially short-term prediction of price dynamics. In particular, it resolves the limitations of a static model by deriving properties such as: skewness and leptokurtosis (due to the mixture of price distributions) and discontinuous shifts in price levels and volatility.

ii) It clarifies the structural profiles of dissimilar and recurrent pricing regimes. If systematic structure is identified in extreme prices, this means that their magnitude is not arbitrary but reflects a recurrent generators’ reaction to market abnormalities, whenever these arise and irrespectively of their exact causality. Depending on the plausible interpretations of the latent market state (which relates to market specificities and model formulation), the modelling could indicate how often and for how long on average: a) Short-term strategic effects persist given the fear for regulatory intervention and new entry, b) A temporary shock in demand or supply inflates prices, c) The market remains illiquid before agents get attracted by price levels and induce activity.

iii) It reveals the stochastic dynamics of regime-switching based on fundamentals rather than solely price levels and thus, allows a more accurate evaluation of the risks induced by price spikes.

5.2 Regime - Switching Regression Model

The regression model with first-order Markov regime-switching is specified as:

$$P_t = X_t' \beta_{S_t} + \varepsilon_t$$
$$\text{where, } \varepsilon_t \sim N(0, \sigma^2_{S_t}),\ \Pr(S_t = i | S_{t-1} = j) = p_{ij}, \forall i, j \in S$$

P_t denotes the spot price in a given load period on day *t*, S_t the latent regime, $S = \{1,2,...,n\}$ the set of possible states X_t a kx1 vector of exogenous explanatory variables at t, β_{S_t} a kx1 vector of regression coefficients, $\sigma^2_{S_t}$ the error variance in regime S_t and p_{ij} the transition probability between states *i* and *j*.

The above model proposed by Hamilton (1990), assumes that the market at each time point is in one of n possible states, indexed by an unobservable discrete variable S_t, which alternates stochastically according to a first-order Markovian process. S_t is as an additional *endogenous* variable in the model and exogenous to the included fundamentals. Each market regime is characterised by a distinct regression price model, i.e. the model parameters are a function of the prevailing state at each time point. In other words, S_t indicates which price model applies at t. A regime shift occurs whenever the underlying market framework changes. The changes are not restricted to shifts in the intercepts or the impacts of the variables but could also be shifts in the residual variances, the non-explained part of the volatility. In this sense, the model is a method to capture stochastic volatility. It should be emphasised that prices are *not apriori* classified into distinct regimes. Both their categorisation (i.e. the identification of the latent state at each point) and the equation of their magnitude are *endogenously* derived with probabilistic inference. The estimation algorithm[20] is outlined in Appendix II.

Having described the structural model with Markov regime-switching, some remarks follow on our application.

i) The hidden market state could be assigned several interpretations depending on market specificities and the variables included in the structural price models. In our case, the irregular regime could reflect temporal market efficiencies, both in the economic and trading sense, such as illiquidity, extreme movements in fuel prices and abusive strategic behaviour of individual generators, not captured by the aggregated strategic variables.

ii) The model suggests that the extreme variation in electricity prices emerges due to the alternation of market regimes, which differ in the nature of *significant*

---

[20] A desirable feature of the model would be to disentangle the mean-reverting component anticipated in the regular price process from the jump-reversal process. Adopting an estimation procedure similar to the one implemented by de Jong and Mahieu (2002) would separate the autocorrelation within the normal regime from the jump towards a spike as well as the reversion after that. The results presented here are derived with conventional probabilistic inference followed in earlier papers. The contamination of the normal process from spikes was not expected to be an issue given the rich structure of the model and the focus on fundamentals rather than autocorrelation dynamics.

influential factors or the *magnitudes* of their effects. For instance, economic fundaments could dictate pricing in normal states, whereas strategic effects prevail in abnormal states. This speculation would not exclude spikes arising from fundamentals such as demand/supply shocks, as frequently in electricity markets, actual capacity scarcity is only moderate but agents react strategically and cause a price explosion.

iii) The serial correlation in the stochastic evolution of the latent state, implied by the Markov assumption, is appealing in the presence of a regulator and speculators, market rules and plant constraints, which pose a limit to the persistence of extreme prices and imply a consistent probabilistic structure for S_t. Particularly in NETA, extreme deviations from regular prices could not persist for long given the overcapacity. It could be argued that the *constant transition probabilities* assumed for S_t are restrictive for an adapted market and could instead be time-varying, i.e. linked to a covariate. However, static transition probabilities were specified for reasons of statistical validity, as the price model postulated a rich structure with implicit time effects and S_t should be exogenous to the included fundamentals.

5.3 Modelling Results

The introduction of regime-switching to the regression model captured more precisely the complex stochastic behaviour of high-frequency electricity prices. The implicit assumption was that dissimilar pricing profiles emerged due to changes in the influential factors or the magnitude of their impacts, in response to diverse market conditions. Indeed, the analysis revealed significant structural components in pricing under *both* market states. For the normal regime, the structural insights achieved with the one-regime regression model were retained to a large extent and more adequate representations were derived, as spikes were endogenously classified to a different regime. Evidence for *systematic* structure was also identified in the irregular regime, where *strategic* factors proved to be significant and much more influential, whereas cost-fundamentals limited compared to normal market conditions. The fact that causalities were uncovered behind the level of irregular prices exceeded the insights of the stylised literature and has significant implications for price forecasting. The conventional assumption of regime-switching within an AR model, instead of a regression one, was inadequate, in most load periods, to replicate the involved price dynamics arising from the interface of financial and physical trading.

The regimes derived with *probabilistic inference* were differentiated by price level, ‘low’ and ‘high’, although this criterion was not constraining. For instance, a moderate price was occasionally assigned to the irregular regime, if its value was perceived peculiar given the underlying market conditions. Similarly, relatively high values were assigned to the regular regime, whenever they seemed to arise naturally from the regular price process for extreme values of the covariates. Allowing for two regimes was often sufficient to capture non-modelled shifts in the market environment. Occasionally, a particularly small set of extremely high prices was classified into the irregular regime causing unreliable estimation. The specification of three regimes resolved this issue resulting to two categories of irregular prices: a set generated from the same structural equation and a minority of spikes that did not comply with the former, at least in the context of the model. The relative frequency of these two irregular regimes could be viewed as a proxy for ‘partially explainable’ versus ‘unanticipated’ spikes within the model boundary. An anticipated feature of the emerged regimes is that they *did not persist* for long periods. As suggested by the smoothed state probabilities, price tranquility was interrupted by fast-reversing spikes. The high frequency of the irregular state was consistent with the interpretation that S_t represented temporal, but highly influential and recurrent changes in the trading environment. This aspect does not characterise the switching dynamics of the less volatile, economic time-series, where regimes tend to be longer reflecting fundamental changes in the macro-economic environment.

The multiple-regime modelling is illustrated with typical examples derived from the spot market for trading periods 25 and 35. Reduced regression specifications were selected in order to facilitate the convergence of the algorithm and ensure reliable estimates in the infrequent irregular category. Despite their simplicity, the regression models retained primary elements of price structure. For the spot price, these involved: first-order autocorrelation as an indication of trading heuristics/inefficiency, historic spot volatility as a risk measure, predicted demand as an economic factor, margin and lagged margin as potentially strategic effects. The dissimilarity between normal and irregular regimes was clearly manifested in the intercepts terms. Beyond this divergence, statistically significant and appealing price structure was detected in extreme states. All effects were significant in at least one mode, but the magnitudes of the coefficients displayed great variation. Tables 7 and 9 illustrate the varying significance and magnitude of coefficients across regimes for the illustrative periods. Tables 8 and 10 summarise transition dynamics. Adding regime-switching to the static regression models, increased the R^2 from 62% and 79% to 74% and 92% respectively for the illustrative periods. The latter was over-estimated to some extent due to the small size of the third regime. Figure 7 displays actual versus fitted values and Figure 8 relative residuals (residuals over actual prices) across the sampling period.

The following discussion focuses on the structural insights into the extreme market state, which were qualitatively consistent across periods. This suggested that players implemented strategies repeatedly or reacted to market irregularities in a similar way in order to induce profits. In specific, for period 25 the margin effect under irregularities was three times more intense than normally, which reveals a substantial escalation of scarcity rents. This particular sensitivity of abnormal prices to margin could partially confound the lower values of margin on irregular days, but as this did not seem to be the case, it could be attributed to a large extent to strategic bidding. Analogously, the negative linear component of demand was inflated compared to the normal regime. This indicated that under unconventional conditions, flexible generators were more inclined to compromise on prices in order to retain the possibility of operating in the evening peak, an option that was becoming more profitable under a spike regime. As opposed to the normal regime, historic price volatility was not significant under the extreme conditions. This suggested that the spikes were not induced by risk-aversion and rational trading behaviour, at least as measured in the model. The autocorrelation parameter for the normal regime was magnified compared to the one-regime regression model of section 2, possibly reflecting a surrogate for the disregarded variables in this simplified model. Neither the ergodic probability of the irregular regime (0.31) nor its probability of persistence (0.34) was minor. Its expected duration was greater than one day suggesting that the assumption of instantly reverting spikes, frequently employed in the industry, is quite restrictive.

For the particularly volatile period 35, three latent states were assumed leading to the emergence of a normal price regime (denoted as II) and two irregular ones (I, III), the last of which captured the most extreme prices. The transition dynamics of regime III revealed persistence to the same state with probability 0.09, shift to the smoother regime I with probability 0.80 and abrupt reversion to normal levels (II) with probability 0.04. This switching pattern suggests a two-stage spike reversal to normal prices. In the irregular regimes, the effects of both Margin and Lagged Margin were inflated, implying abusive bidding and capacity withholding. The impact of the latter in regime I was double and in regime III ten-fold the regular value under regime II. The impact of demand was retained significant in regime I but with a sign reversal, which could signal strategic behaviour even in periods of low demand.

The systematic exercise of strategic behaviour of generators, when temporal irregularities arise, implies discontinuities in the volatility process. Indeed, volatility is multiplied by a factor of 2 or 3 during abnormalities, as shown below.

Table 7. PX Price, Period 25. Price Structure with a Two-Regime Regression Model.

Variable	Irregular Regime (I)	Regular Regime(II)
Intercept	36.82*[21]	11.47*
Lagged PX	-0.10	0.44*
PX Volatility	5.55	10.87*
Demand – Linear	-33.57*	-8.89*
Demand – Quadratic	-0.42	1.12
Margin	-0.0014*	-0.0005*
σ	1.12	0.42
R^2	0.74	

Table 8. PX Price, Period 25. Transition and Ergodic Regime Probabilities.

Regime	I	II
I	0.34	0.30
II	0.66	0.70
Ergodic Probabilities	0.31	0.69

Table 9. PX Price, Period 35. Price Structure with a Three-Regime Regression Model.

Variable	Irregular Regime (I)	Regular Regime (II)	Irregular Regime (III)
Intercept	40.72*	7.30*	-18.10*
Lagged PX	0.04	0.61*	1.4
Demand – Linear	-63.6*	16.09*	28.90
Demand – Quadratic	115.3*	21.03*	15.82
Margin	-0.0028*	-0.0007*	-0.0025*
Lagged Margin	0.001*	0.0005*	0.005*
σ	1.97	0.82	2.8

21 Asterisk denotes significance at the 95% confidence level.

R^2	0.92

Table 10. PX Price, Period 35. Transition and Ergodic Regime Probabilities.

Regime	I	II	III
I	0.33	0.07	0.80
II	0.57	0.90	0.13
III	0.10	0.03	0.09
Ergodic Probabilities	0.14	0.82	0.04

6. Conclusions

Spot electricity prices exhibit unusual volatility, orders of magnitude higher than financial assets and other commodities. Because spot price changes have a number of structural components, conventional price volatility models derive erroneous results (e.g. GARCH), whereas specifications that perceive variance as constant or a simple stochastic process (jump-diffusion, autoregressive regime-switching) facilitate analytical solutions in derivatives evaluation, but are inadequate to evaluate accurately short-term price dynamics and risks. This paper models explicitly the stochastic behaviour of volatility that relates to an underlying structural price specification, provides causal explanations for its emergence as well as a methodology to reduce short-term trading risk. The volatility measure modelled has an appealing interpretation; that of *residual variance* after detaching prices from fundamental structure. This represents the descriptive, ex-post, uncertainty of a rational expectations price model and is also a proxy for the predictive uncertainty of an agent who derives his price expectations from the econometric model and his prior beliefs about the influential factors of price variation. To clarify the nature and dynamics of volatility, four modelling approaches are proposed. Each captures distinct causal or stochastic aspects of the price and volatility process, such as heteroscedasticity, conditional or unconditional, of random shocks, evolution of the underlying price structure and discontinuities in price structure due to temporal market irregularities. Some results and comments are summarised below.

i) Each modelling approach motivates a convolution of fundamental, strategic and behavioural interpretations of volatility, which seem consistent with plant dynamics. Each is plausible depending on market conditions and agents' information, whereas a hybrid entails computational complexities. This multiplicity of plausible models is consistent with the view that electricity prices are only one of multiple equilibria and suggests that agents may not converge quickly to the same expectations model. The inevitable heterogeneity of expectations possibly induces additional volatility to that implied by non-storability and strategic behaviour.

ii) Volatility structure displayed significant intra-day variation. This indicates that agent responses to information arrivals are a function of plant operating constraints and strategic potential, both varying within the day.

iii) GLS modelling revealed non-linear responses of residual variance to strategic and economic impacts, which partially explains the erratic price movements.

iv) The convergence limitations of GARCH models reported in the literature were resolved when a structural model was assumed for prices and a leptokurtic distribution for the errors. GARCH effects (including the parameter of asymmetry) are concealed or exaggerated depending on the adequacy of the price model. This sensitivity can lead to mis-perceptions of the underlying price risk and mis-pricing of financial instruments. The analysis further indicated an inverse to the leverage effect typical in financial markets, which could relate to regulatory presence.

iv) Time-varying regression modelling revealed the dynamic structure of prices, as agents continuously update their strategies. The concept of market adaptation provides an adequate explanation for stochastic volatility, as GARCH effects diminish after adjusting prices for TVP heterogeneity. Dynamic price modelling seems more appropriate for highly evolving or unstable markets, such as those emerging after reforms or regulatory interventions.

vi) A structural price model with multiple, stochastically alternating, regimes clarified price uncertainty under temporal market abnormalities, when strategic behaviour dominates pricing. The quantification of volatility under extreme conditions is appealing for Value-At-Risk calculations.

vii) The applicability of these structural volatility specifications are contingent upon agents integrating their expectations about market fundamentals into a formal model which, given the information asymmetries in electricity markets, may well be quite diverse. The models could, however, be reduced to a stylised formulation, if an autoregressive model is assumed as the rational expectations price model.

References

Brown, R.L., Durbin, J. and Evans, J.M., (1975) "Techniques for testing for the constancy of regression relationships over time (with discussion)", *Journal of the Royal Statistical Society* B **37** pp 149-192.

de Jong, C., and Huisman, R., 2002. Option Formulas for Mean-Reverting Power Prices With Spikes. Working Paper, Erasmus University.

Dempster, A.P, Laird, N.M., and Rubin, D.B., (1997) "Maximum Likelihood Estimation from Incomplete Data via the EM Algorithm" *Journal of Royal Statistical Society*, **39**, Series B, pp 1-38.

Deng, S., (2000) "Stochastic Models of Energy Commodity Prices and their applications: Mean Reversion with Jumps and Spikes" *working paper*, PWP-073, University of California Energy Institute.

Duffie, D., Gray, S., and Hoang, P., (1998) "Volatility in Energy Prices" *Managing Energy Price Risk*. Risk Publications.

Escribano, A., Peaea, J. and Villaplana, P., (2002) "Modelling Electricity Prices: International Evidence" *working paper*, Universidad Carlos III de Madrid.

Ethier, R., and Mount, T., (1998) "Estimating The Volatility Of Spot Prices In Restructured Electricity Markets And The Implications For Option Values" *working paper*, Cornell University"

Eydeland, A., and Gemand, H., (1998) "Pricing Power Derivatives", *Risk* **11** pp 71-73.

Hamilton, J., (1990) "Analysis of Time Series Subject to Changes in Regime, *Journal of Econometrics*" **45** pp 39-70.

Huisman, R., and Mahieu, R., (2001) "Regime Jumps In Power Prices" *working paper*, Energy & Power Risk Management.

Johnson and Barz, 1999. Selecting Stochastic Processes For Modelling Electricity Prices. *Energy Modelling and the Management of Uncertainty*, Risk Publications

Karakatsani, N.V., and Bunn, D.W., (2004) "Structural and Dynamic Properties of the Low British Electricity Wholesale Prices, 2001-2002", *working paper*, London Business School

Kholodnyi, V., (2001) "A Non-Markovian Process for Power Prices with Spikes for Valuation of European Contingent Claims on Power" *working paper*, TXU Energy Trading.

Kim, C.J., and Nelson, C.R., 1999. Space Models with Regime Switching. MIT Press.

Knittel, C.R., and Roberts, M., (2001) "An Empirical Examination of Deregulated Electricity Prices" *working paper*, PWP-087, University of California Energy Institute.

McCullagh, P., and Nelder, J. A., (1987) Generalised Linear Models. Chapman and Hall, London.

Pilipovic, D., (1998) *Energy Risk: Valuing and Managing Energy Derivatives* McGraw-Hill, New York

Stevenson, M., (2002) "Filtering and Forecasting Spot Electricity Prices in the Increasingly Deregulated Australian Electricity Market" *working paper*, University of Technology, Sydney.

Sweeting, A., (2000) "The Wholesale Market for Electricity in England and Wales: Recent Developments and Future Reforms" *working paper*, MIT.

Visudhiphan, P., and Ilic, M., (2000) "Dependence of generation market power on the demand/supply ratio: analysis and modelling" *IEEE Transactions on Power Systems* **2** pp 1115-1122

Vucetic S., Tomsovic, K., and Obradovic, Z., (2001) "Discovering Price-Load Relationships in California Electricity Market" *IEEE Transactions on Power Systems* **16** pp 280-286

Zakonian, J.M., (1994) "Threshold Heteroscedastic Models", *Journal of Economic Dynamics and Control* **18** pp 931-955.

Appendix I

The Kalman Filter Algorithm for the above TVP model consists of the following:

Step 1: *Prediction*

$$\beta_{t|t-1} = \beta_{t-1|t-1}$$

$$P_{t|t-1} = P_{t-1|t-1} + \Sigma \text{ , variance} - \text{covariance matrix of } \beta_{t|t-1}$$

$$n_{t|t-1} = P_t - X_t \beta_{t|t-1}, \textit{forecast error based on information up to } t-1$$

$$H_t = X_t P_{t|t-1} X' + \sigma_\varepsilon^2 \text{ ,conditional variance of forecast error } n_{t|t-1}$$

Step 2: *Updating*

$$\beta_{t|t} = \beta_{t|t-1} + K n_{t|t-1}$$

$$P_{t|t} = (I_k - K X_t) P_{t|t-1}$$

$$\textit{where, } K = P_{t|t-1} X_t H_t^{-1} \ (\textit{Kalman gain})$$

Step 3: *Smoothing*

$$\beta_{t|T} = \beta_{t|t} + P_t^* (\beta_{t+1|T} - \beta_{t|t})$$

$$P_{t|T} = P_{t|t} + P_t^* (P_{t+1|T} - P_{t+1|t}) P_t^*$$

$$\textit{where, } P_t^* = P_{t|t} P_{t+1|t}^{-1}$$

The estimate of β_t based on information up to time t, i.e. $\beta_{t|t}$, is defined as an optimal combination of the prior on β_t, i.e. $\beta_{t|t-1}$ and the forecast error $n_{t|t-1}$, the weight being the Kalman gain K. Hence, for starting values of $\beta_{o|-1}, P_{o|-1}$[22] and assuming that the values of σ_ε^2 and Σ are known[23], recursive estimates are derived for the mean and variance of the state vector β_t using observations on P_t and X_t

[22] As prior estimates of the parameters did not exist, the initialisation of the filter was not trivial. The coefficients derived from the static regression models were selected as initial values for the vectors β_j and a diffusion prior was assumed for the initial variances. Due to the non-linearities implicit both in the data and the optimisation problem, sensitivity of the solution to the initial values is by theory inevitable. Although the results were fairly robust to a wide range of values, they should still be interpreted with caution.

[23] In economics, the values of the above time-invariant parameters are seldom known. They can be estimated however, maximising the following log likelihood function, which is based on the forecasting errors and their conditional variance:

$$\log L = constant + \frac{1}{2} \sum_{t=1}^{T} \{\log |H_t| + n'_{t|t-1} H_t^{-1} n_{t|t-1}\}$$

Appendix II

Estimation of Regime-Switching Regression Model

As the underlying regime is unobserved, the distribution of price conditional on past information can be expressed as: $f(P_t|I_{t-1},\theta) = p^1_{t|t-1}\phi(\frac{P_t - X\beta_{S_t}}{\sigma_1}) + p^2_{t|t-1}\phi(\frac{P_t - X\beta_{S_t}}{\sigma_{21}})$,where $p^1_{t|t-1} = 1 - p^2_{t|t-1} =$

$\Pr(S_{t-1} = 1|I_{t-1};\theta) = p_{11}$, $\theta = (\beta_1, \beta_2, \sigma_1, \sigma_2, p_{11}, p_{22})$ and ϕ the standard normal p.d.f. The ex ante transition probability p_{ij} that state *i* will be followed by state *j*, *i,j*$\in s$, depends on the available information set. Expressions for $p^s_{t|t-1}$ are obtained by

Bayes' Rule: $p^1_{t|t-1} = p^1_{t-1|t-1}p_{11} + p^2_{t-1|t-1}p_{11}(1 - p_{22})$, $p^1_{t|t} = \dfrac{p^1_{t|t-1}\ \phi(\frac{P_t - X'_t\beta_1}{\sigma_1})}{f(P_t|I_{t-1})}$,

where $p^s_{t|t+\tau}$ is similarly defined. When τ is negative, $p^s_{t|t+\tau}$ represents a forecast of the probability that regime s will be realized in period τ. When τ is positive, $p^s_{t|t+\tau}$ represents a smoothed or updated inference of the probability that that regime s was in fact realized τ periods ago.

Although the state variable S_t is unobservable, it is possible from the estimated transition probabilities and regime-dependent models to derive a probabilistic inference about its latent value for each time point t in the sample. This inference is compounded in the filtered or smoothed state probability $p^f_s = P(S_t = s|I_T)$, expressing the probability of the market being in regime s on day t, conditional on I_T; the set that contains all the available information on prices and structural variables in the sample.

Parameter estimates could be derived with Maximum Likelihood Estimation (MLE) by constructing the log likelihood function $\sum \log(f(P_t|I_{t-1},\theta)$ and setting the scores to zero. Hamilton (1990) shows that the ML estimates can be obtained using an application of the Expectation Maximisation algorithm of Dempster, Laird and Rubin (1977). In the *expectation step* (E), the unobserved states S_t □are estimated by their smoothed probabilities $\hat{p}^s{}_{t|T}$. The conditional probabilities $\Pr(S_t|P,\theta^{l-1})$ are calculated with a filter and smoother using the estimated parameter vector θ^{l-1} of the last maximization step instead of the unknown true parameter vector. In the *maximization step* (M), an estimate of θ is derived as a solution of the FOCs of ML estimation, where the conditional regime probabilities $\Pr(S_t|P,\theta\ )$ are replaced by the smoothed probabilities $\hat{p}^s{}_{t|T}\ (\theta^{l-1})$ of the last expectation step. Thus, the dominant source of non-linearities in the FOCs is eliminated.

For the regime-switching structural model, the application of the EM algorithm results in the following set of recursive equations for iterations *l=1,2,…*.

$$\hat{\beta}_s^l = (X_s' X_s)^{-1} X_s' P_s$$

$$\hat{\sigma}_s^l = \frac{1}{\sum \hat{p}^s{}_{t|T}} (P_s - X_s' \beta_s)'(P_s - X_s' \beta_s)$$

$$P_{ss}^l = \frac{\sum \Pr(S_t = s, S_{t-1} = s; I_T, \theta^{l-1})}{\sum \Pr(S_t = s; I_T, \theta^{l-1})}$$

where $X_s = \sqrt{\hat{p}^s{}_{t|T}} \cdot X$, $P_s = \sqrt{\hat{p}^s{}_{t|T}} \cdot P$.

The recursion begins by selecting an initial parameter vector θ^0 computing $\hat{p}^s{}_{t|T}$, then computing θ^1 and so until convergence is achieved. The recursive nature of the estimation algorithm arises because of the need to construct $\hat{p}_{t|T}^s$, (which depends on θ^{l-1}) before computing θ^l. From the transition probabilities ergodic probabilities are derived, which refer to the steady-state i.e. long-term probability distribution of S_t. Expected (fitted) prices and variances of the estimators are calculated as: $E(P_t) = \sum_{s=1}^{2} E(P_t | S_t = s) \cdot p_s^f$,

$Var(P_t) = p_1^f \sigma_1^2 + p_2^f \sigma_2^2 + p_1^f p_2^f [E(P_t | S_t = 1) - E(P_t | S_t = 2)]^2$

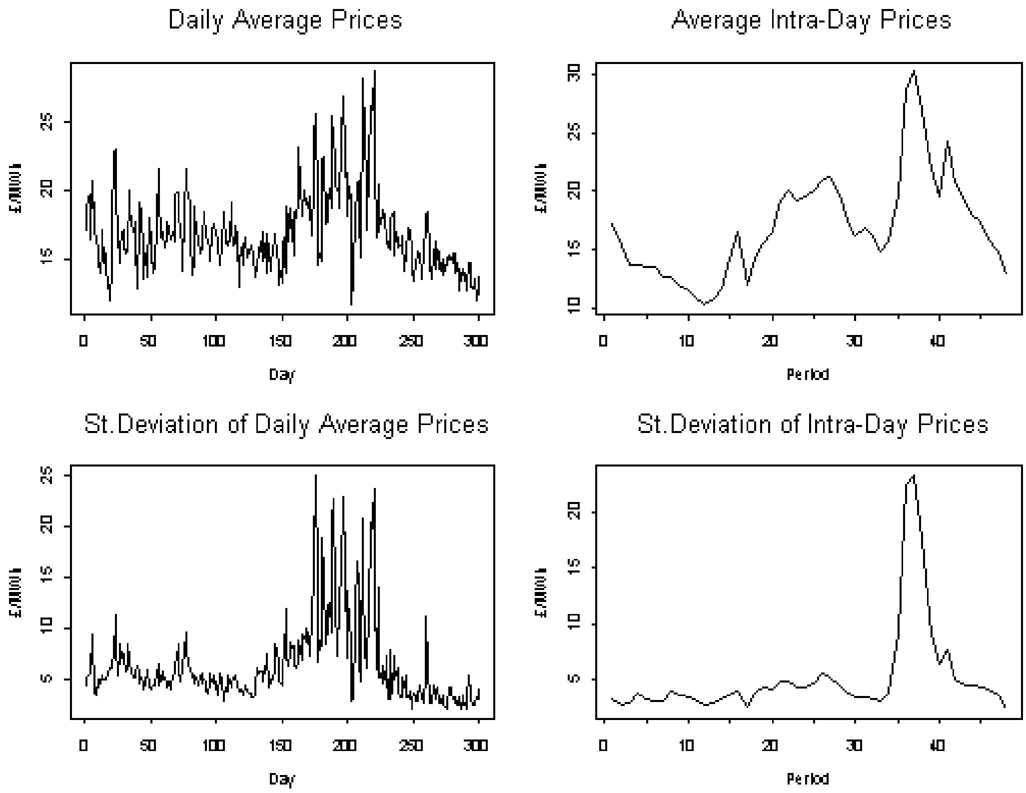

Figure 1. Inter-Day and Intra-Day Price and Volatility Profiles for Spot Prices during 6 June 2001 - 1 April 2002.

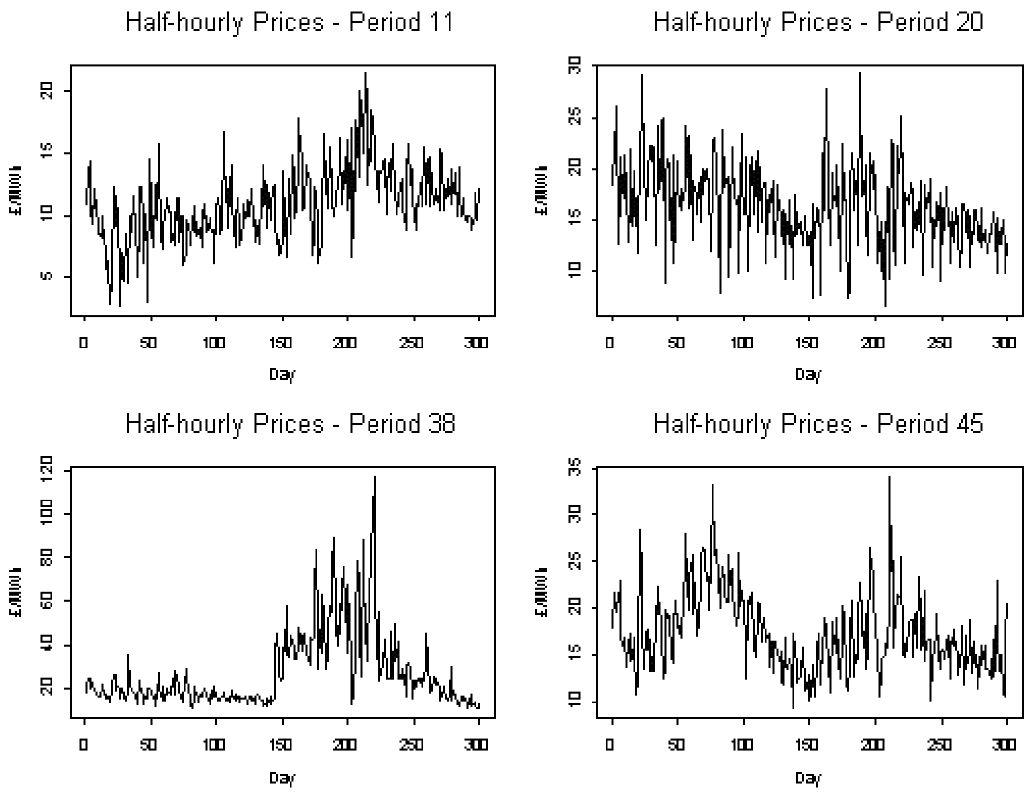

Figure 2. Half-Hourly Spot Prices for selected periods during 6 June 2001 – 1 April 2002.

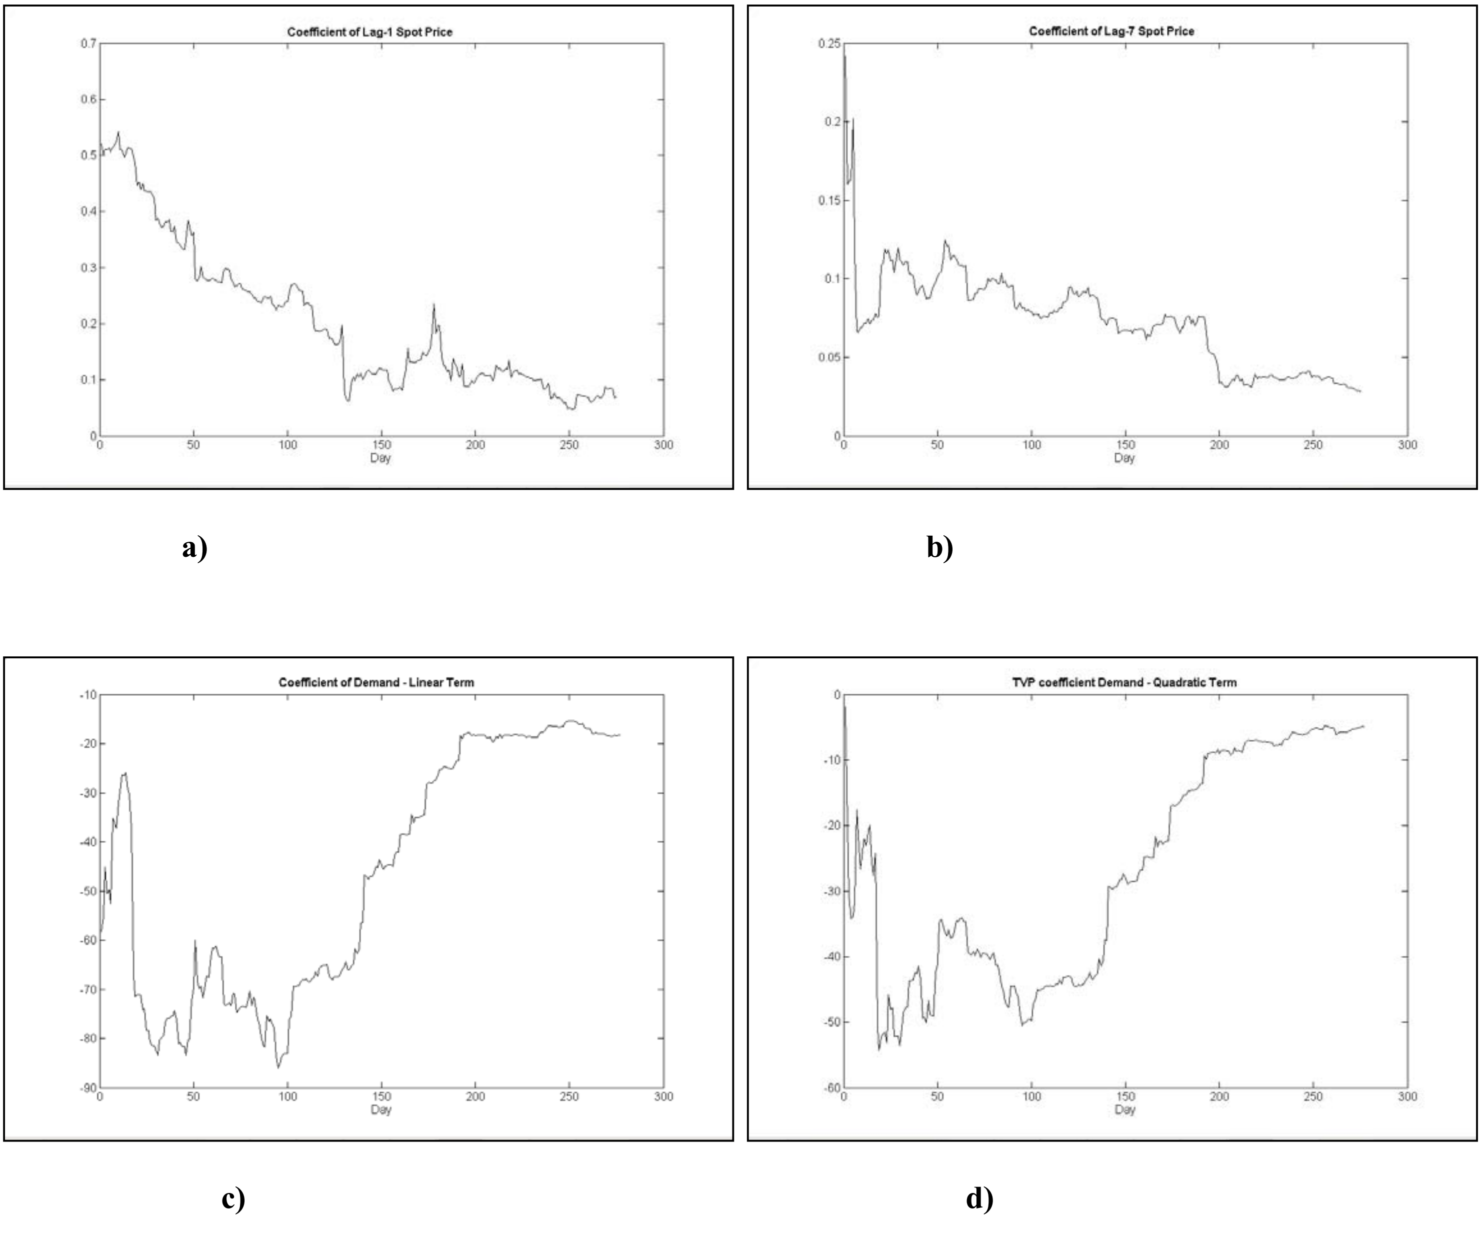

Figure 3. Time-Varying Regression Coefficients for Spot Prices in Period 25.

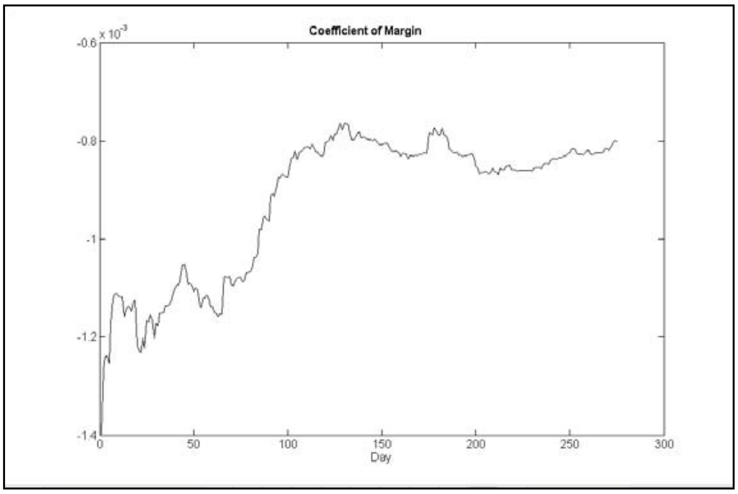

e)

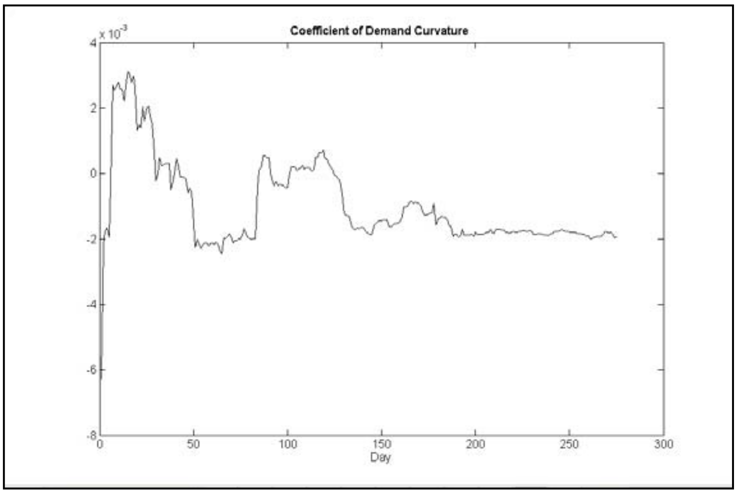

f)

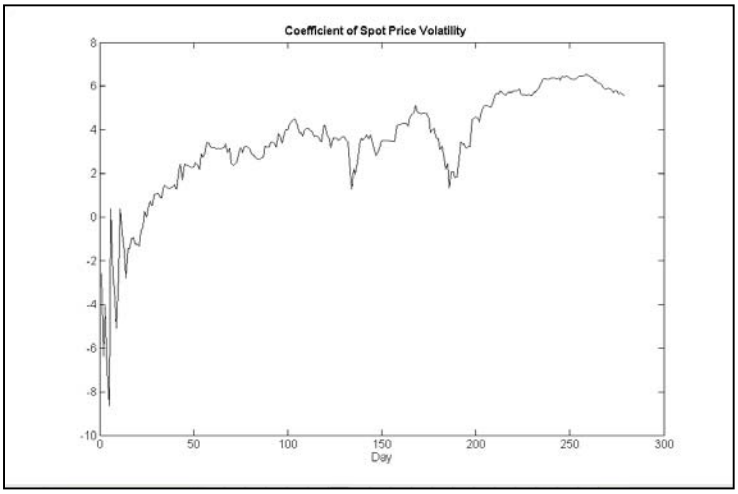

g)

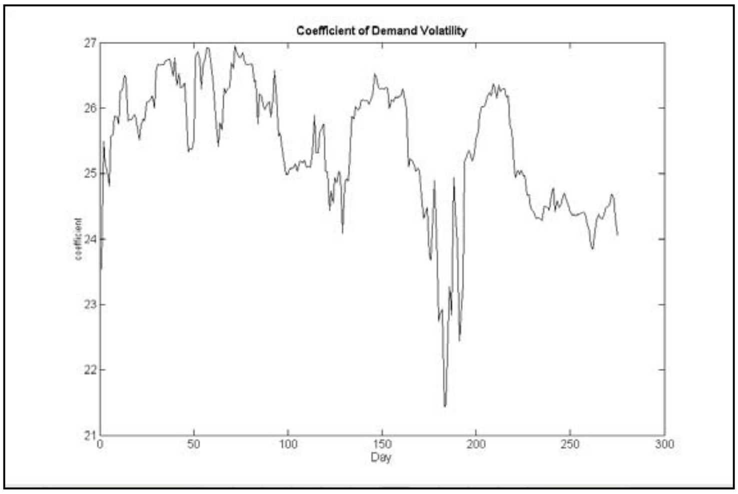

h)

Figure 3. Time-Varying Regression Coefficients for Spot Prices in Period 25.

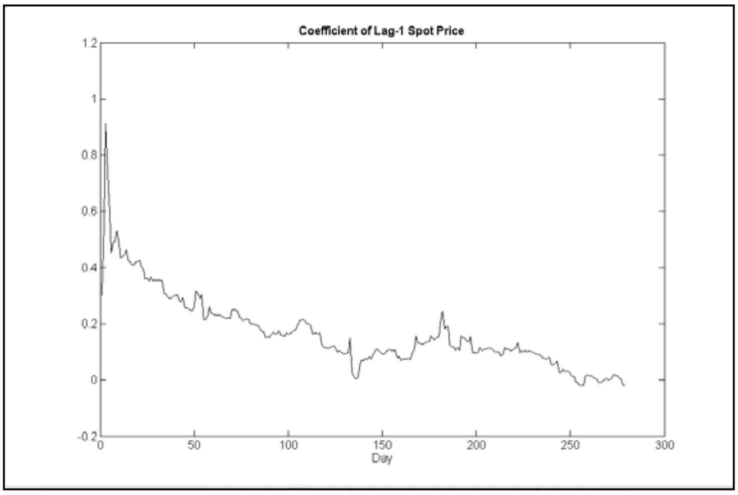

a)

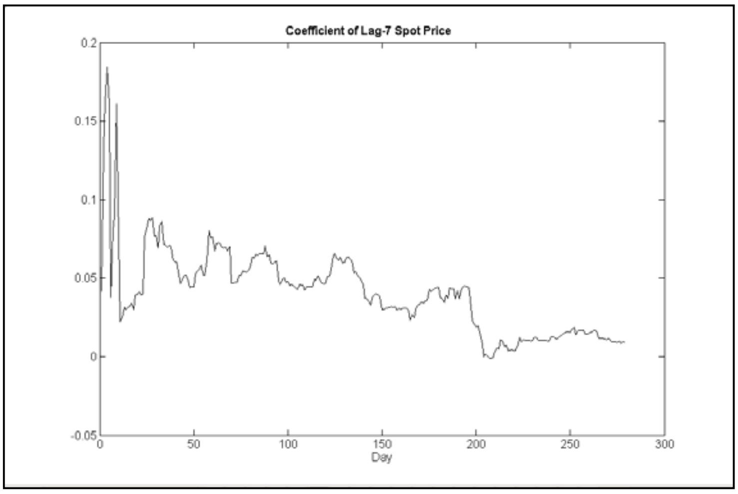

b)

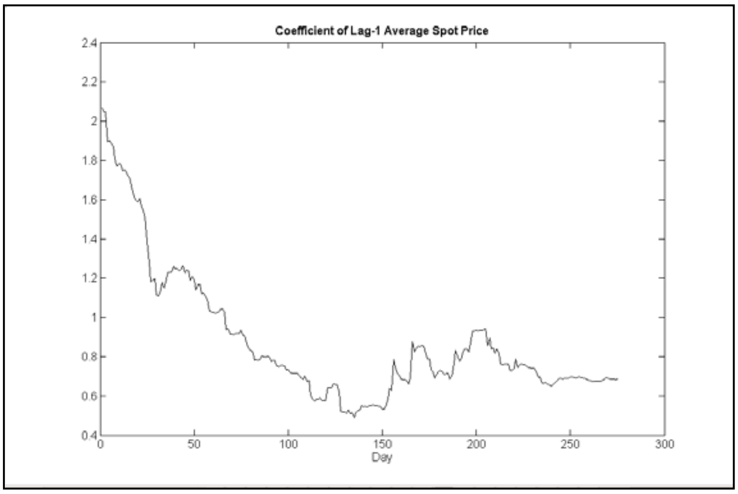

c)

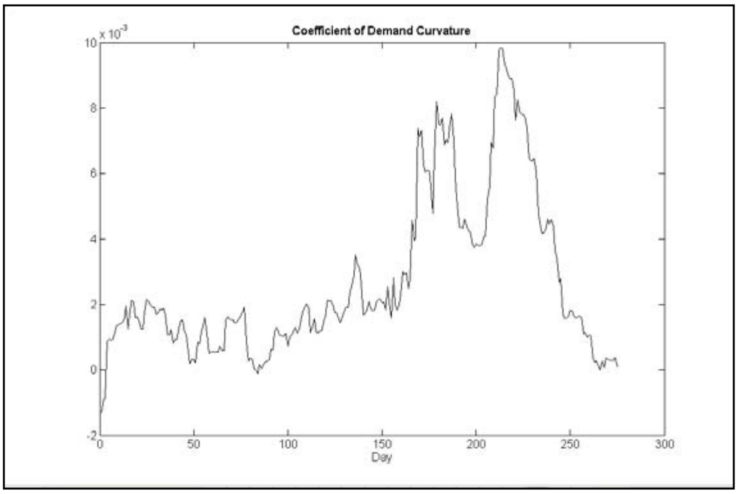

d)

Figure 4. Time-Varying Regression Coefficients for Spot Prices in Period 35.

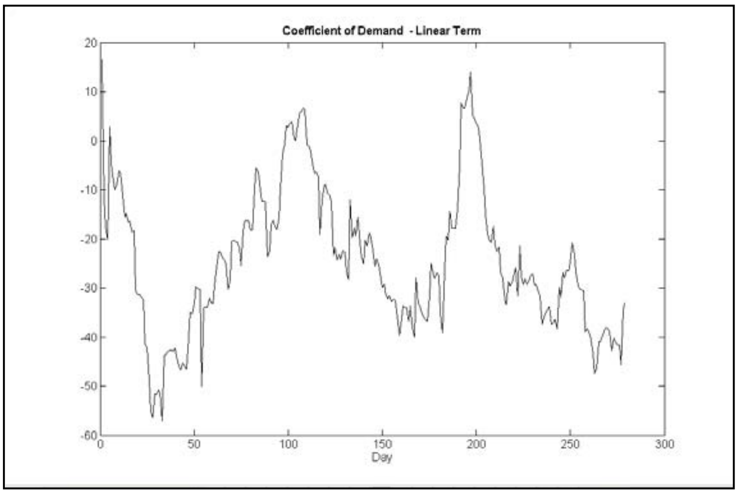

e)

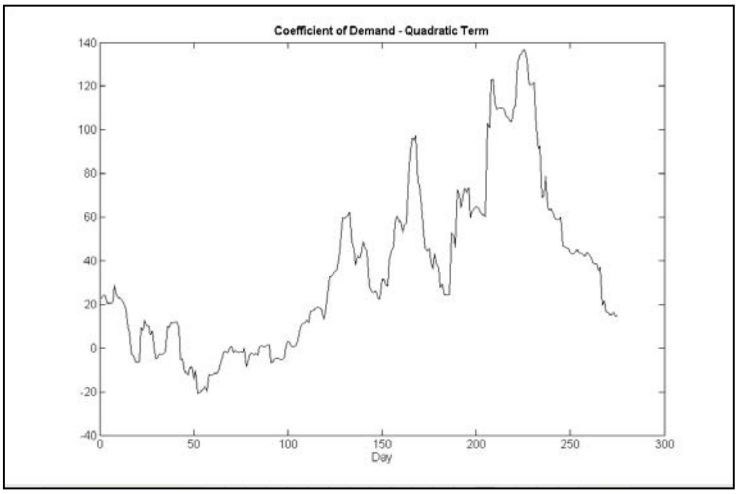

f)

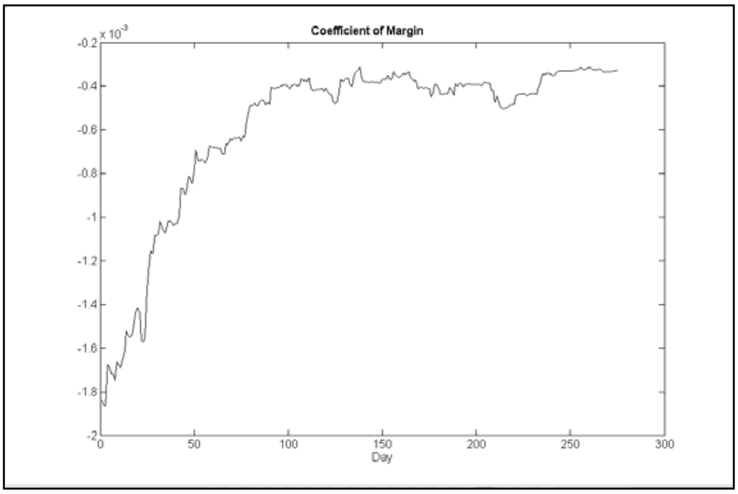

g)

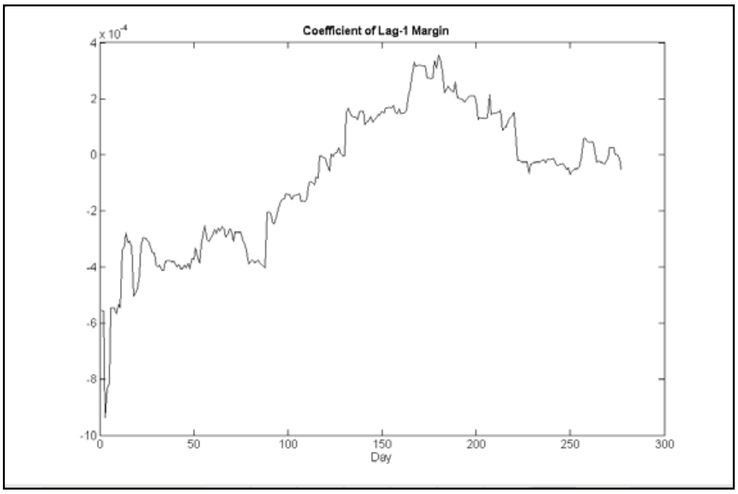

h)

Figure 4. Time-Varying Regression Coefficients for Spot Prices in Period 35.

Period 25 **Period 35**

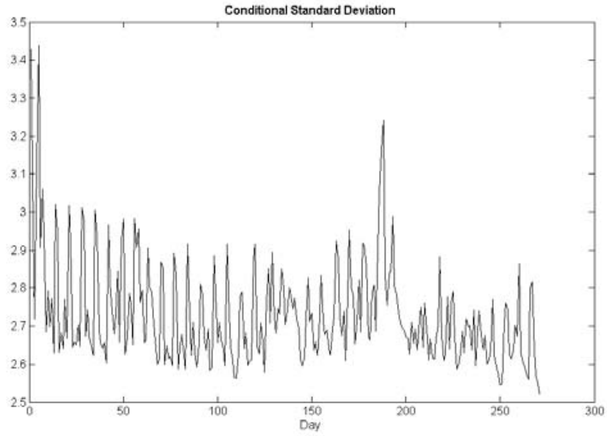

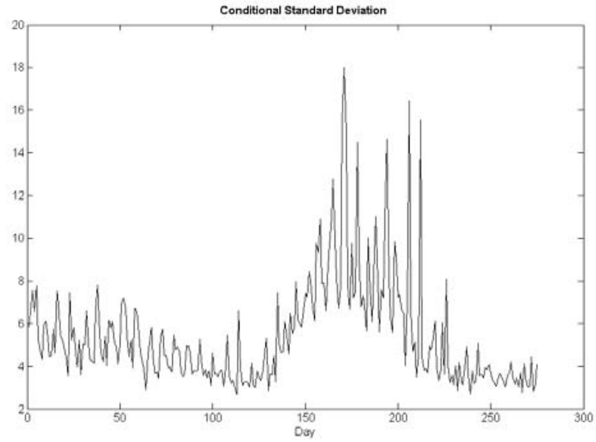

Figure 5. Conditional Standard Deviation with TVP modelling for Spot Prices in Periods 25 and 35.

Period 25 **Period 35**

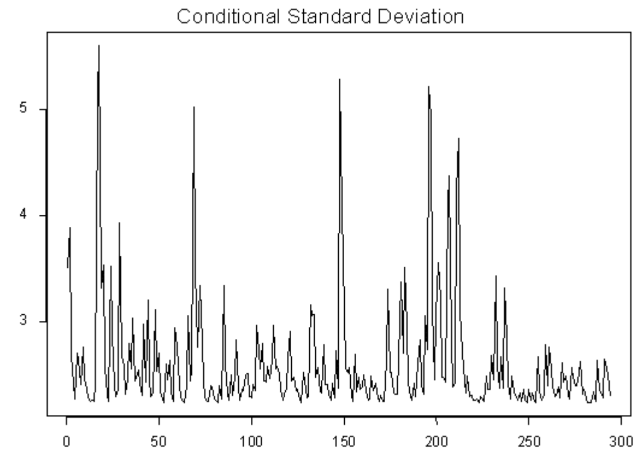

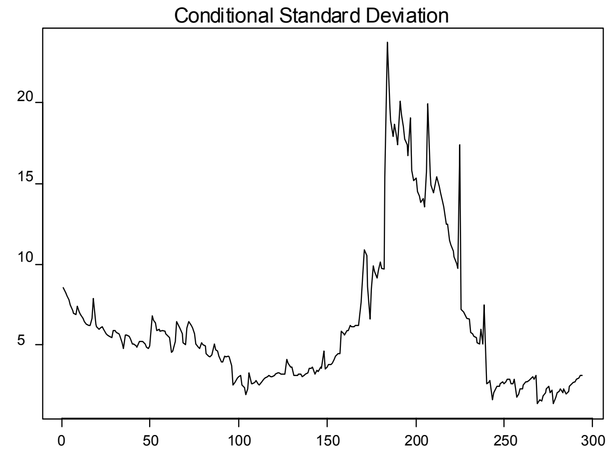

Figure 6. Conditional Standard Deviation with GARCH modelling for Spot Prices in Periods 25 and 35.

Period 25 **Period 35**

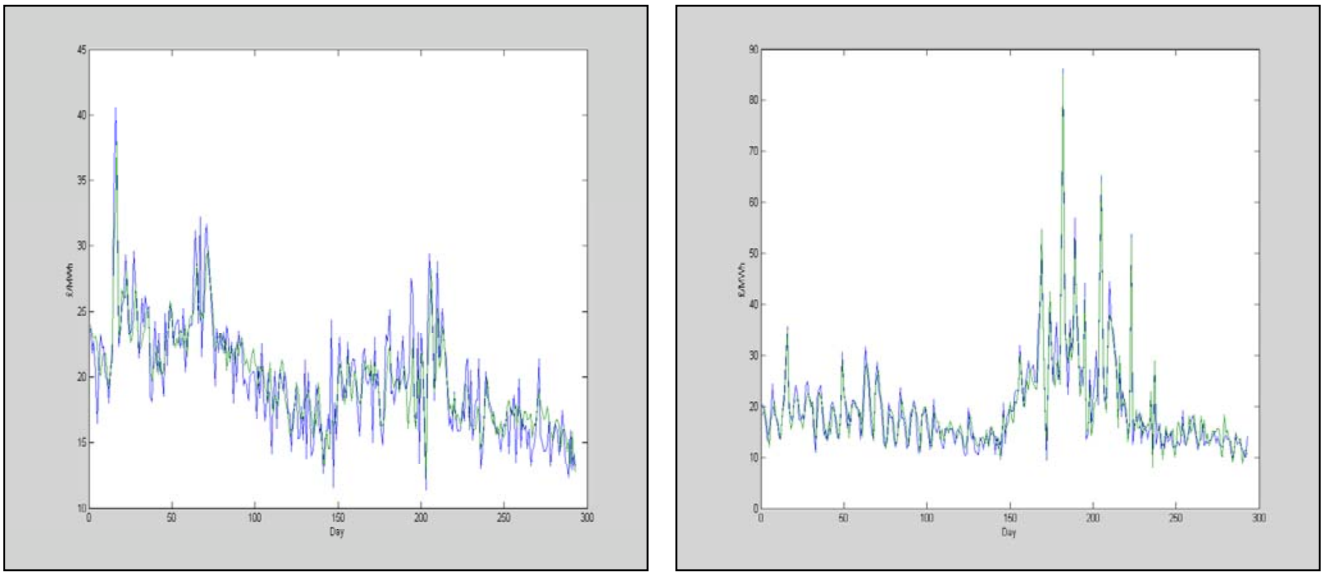

Figure 7. Actual (Dark Line) vs. Fitted Spot Prices in Periods 25 and 35 with Markov Regime-switching Regression Modelling.

Period 25 **Period 35**

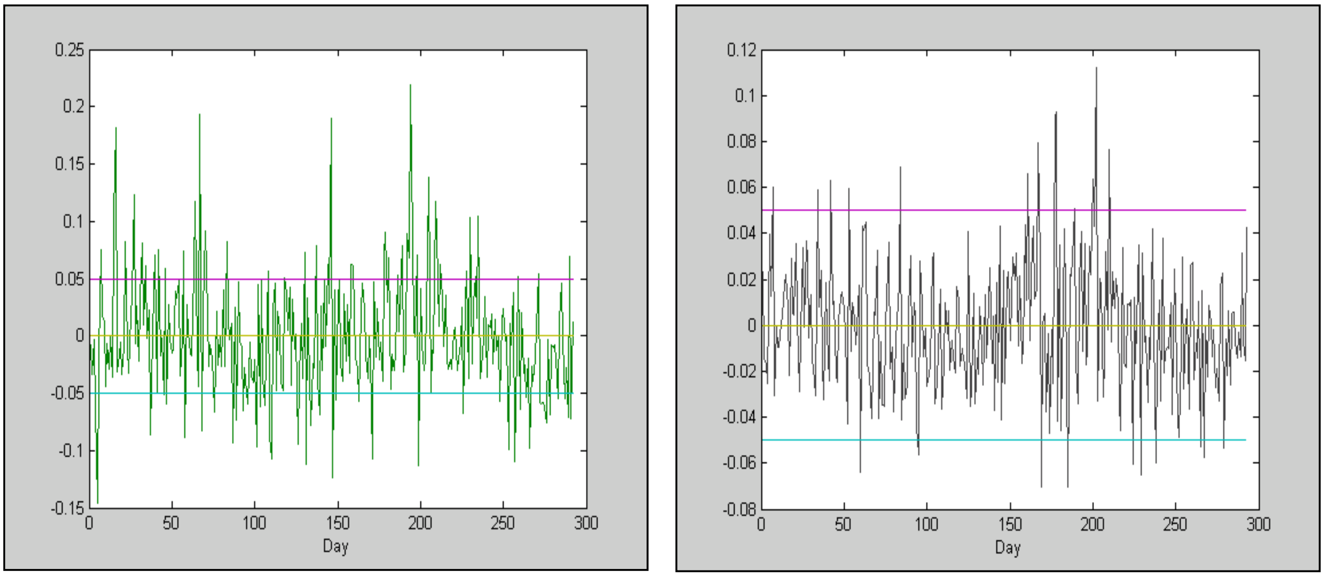

Figure 8. Relative Residuals in Periods 25 and 35 from Markov Regime-switching Regression Modelling.