Queen's Economics Department Working Paper No. 1211

A vector autoregressive model for electricity prices subject to long memory and regime switching

Niels Haldrup
Aarhus University and CREATES

Frank S. Nielsen
Aarhus University and CREATES

Morten Ørregaard Nielsen
Queen's University and CREATES

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

8-2009

A vector autoregressive model for electricity prices subject to long memory and regime switching*

Niels Haldrup†
Aarhus University and CREATES

Frank S. Nielsen
Aarhus University and CREATES

Morten Ørregaard Nielsen
Queen's University and CREATES

August 11, 2009

Abstract

A regime dependent VAR model is suggested that allows long memory (fractional integration) in each of the observed regime states as well as the possibility of fractional cointegration. The model is motivated by the dynamics of electricity prices where the transmission of power is subject to occasional congestion periods. For a system of bilateral prices non-congestion means that electricity prices are identical whereas congestion makes prices depart. Hence, the joint price dynamics implies switching between a univariate price process under non-congestion and a bivariate price process under congestion. At the same time, it is an empirical regularity that electricity prices tend to show a high degree of long memory, and thus that prices may be fractionally cointegrated.

Analysis of Nord Pool data shows that even though the prices are identical under non-congestion, the prices are not, in general, fractionally cointegrated in the congestion state. Hence, in most cases price convergence is a property following from regime switching rather than a conventional error correction mechanism. Finally, the suggested model is shown to deliver forecasts that are more precise compared to competing models.

Keywords: Cointegration, electricity prices, fractional integration, long memory, regime switching.

JEL Classification: C32.

*We are grateful to Javier Hualde and an anonymous referee for valuable suggestions and comments. Previous versions of this paper have been presented at the CREATES long memory symposium in Aarhus, July 2007, the ESRC econometrics study group meeting in Bristol, July 2007, the NBER-NSF time series conference in Iowa, September 2007, and seminars at Aarhus University. We are grateful to the participants for comments and suggestions. This work was partly done while F. S. Nielsen was visiting Cornell University, whose hospitality is gratefully acknowledged. We are also grateful to the Danish Social Sciences Research Council (FSE grant no. 275-05-220), the Social Sciences and Humanities Research Council of Canada (SSHRC grant no. 410-2009-0183), and the Center for Research in Econometric Analysis of Time Series (CREATES) funded by the Danish National Research Foundation for financial support.

†Please address correspondence to: Niels Haldrup, CREATES, School of Economics and Management, Aarhus University, Building 1322, DK-8000 Aarhus C, Denmark. email: nhaldrup@creates.au.dk.

1 Introduction

Over the past decade or so electricity markets have been strongly liberalized throughout the world. In particular, the Nordic power market consisting of Norway, Sweden, Finland, and Denmark has developed remarkably towards liberalization and the establishment of competitive market conditions, and today this market serves as a model for the restructuring of other power markets. The Nordic power market is characterized by a grid of physical exchanges of power across geographical regions where the actual exchange is constrained by the flow capacity. Naturally, this has implications for the way prices are formed. When there are no bilateral capacity restrictions, there is a free flow of power and prices will be identical. On the other hand, when there is congestion prices tend to depart to meet the supply and demand conditions subject to restricted access to power from other regions. In order to model electricity prices it is thus natural to consider regime dependent price processes reflecting the presence or absence of flow congestion. This particular feature of the market has been addressed in recent work by Haldrup & Nielsen (2006*a*, *b*). Another important property of electricity prices modeled in these works is the presence of long memory. Statistical tests strongly reject the price series being I(0) and I(1), whereas I(d) processes with d being fractional (see Granger & Joyeux (1980) and Hosking (1981)) better characterize the data.

The combination of fractional integration and regime switching gives rise to some challenges. Granger & Ding (1996), Diebold & Inoue (2001), and Granger & Hyung (2004), among others, argue that under certain conditions time series variables can spuriously have long memory when measured in terms of their fractional order of integration, when in fact the series exhibit non-linear features such as regime switching. In the model framework of Haldrup & Nielsen (2006*a*, *b*) separate long memory price dynamics is allowed in adjacent power regions depending upon whether the power exchange is subject to congestion or non-congestion. The model has some similarities to the Markov switching model defined by Hamilton (1989). However, because the defining property of e.g. a non-congestion state is that prices are identical, the state variable is observable as opposed to being a latent variable. Thus our model is not of the traditional Hamilton (1989) Markov switching type, but we still refer to it as a regime switching model since it does include switching between two separate regimes.

An important feature of the model is that the price processes in the different regimes can have different degrees of long memory, which gives rise to a number of interesting possibilities. For instance, consider the state with non-congestion and assume that the associated bivariate prices are fractionally integrated of a given order. It follows that prices are fractionally cointegrated in this case, i.e. extending the notion of Granger (1981, 1986) and Engle & Granger (1987), in the sense that individual prices are fractionally integrated but price differences are identically zero. Thus, an extreme form of cointegration occurs in this situation because the prices are identical and hence are governed by exactly the same price shocks. The price behavior in the congestion state can (and typically will) be very different. That is, the bivariate prices can be fractionally cointegrated in a more conventional way or the prices can appear not to cointegrate. Hence, the model can potentially exhibit state dependent fractional cointegration. By not appropriately conditioning on the congestion state, i.e. when having a model with no regime switching, the full sample estimates are likely to be a convex combination of the behavior in the individual states and hence misleading inference is likely to result.

The modeling approach used in Haldrup & Nielsen (2006*b*) is limited in the sense that the individual price series and the relative price series are analyzed separately as univariate models. When the focus of analysis is the potential (fractional) cointegration amongst multiple series a system approach is more natural, but clearly also more complex in the present context given the particular features the model should allow. In principle, the full set of price series should be modeled jointly, and, depending upon the market conditions, should shrink to a limited number of price series reflecting periods with non-congestion at some grid points.

We distinguish between price areas and geographical regions. Each geographical region corresponds to a physical exchange (e.g., West Denmark, South Norway, etc.) and is therefore constant over time. On the other hand, a price area is defined simply as an area with the same price and may therefore change over time. Thus, West Denmark and South Norway always constitute two geographical regions, but in the case of non-congestion the same price prevails in both geographical regions and they hence constitute just one price area in that case.

In this paper we model multiple price series jointly in a vector autoregression (VAR), which allows for fractionally integrated time series that potentially cointegrate in the congestion state. In the non-congestion state, prices are identical by definition and hence a univariate model for the price process is applied in this particular regime. Thus, our VAR model for fractionally cointegrated processes allows for the possibility of regime switching, and in particular differs from other specifications offered in the literature in the sense that our VAR model collapses to a pseudo-univariate model when a specific state arises. Our model is therefore directly motivated by the structure and functioning of the Nordic power market.

There are different reasons why the identification of separate price dynamics is important. The operation of electricity markets is similar to the operation of financial markets with electricity power derivatives being priced and traded in highly competitive markets and hence appropriate modeling of both means and variances is crucial. Furthermore, the price dynamics is of interest with respect to competition analysis of electricity markets where market delineation is a central issue, see e.g. Sherman (1989) and Motta (2004). Even though most power markets are highly liberalized there is still scope for regulating authorities to closely follow the market behavior, see also Fabra & Toro (2005). Under non-congestion there is obviously a single price existing in the market and the relevant market is defined as the geographical regions with identical prices. However, when there is congestion it is of interest to follow the price dynamics closely because suppliers can have a dominating position. The market delineation thus becomes less straightforward in this case. If the price dynamics appears to be very different there is scope for further examination of the market conditions by regulatory authorities.

In our empirical analysis we find that generally the behavior of electricity prices in geographical price regions are different across states. The analysis shows that it is important to condition on congestion/non-congestion as non-switching models can generate misleading conclusions with regard to the price dynamics. Three leading types of misclassification of the model dynamics may arise. First, non-switching models may indicate that the price series are fractionally cointegrated, whereas when conditioning on states this is only the case in the non-congestion state (in which prices are identical by definition). Secondly, the non-switching model could indicate that there is no fractional cointegration when in fact there is cointegration in the non-congestion state. Finally, there is the possibility of fractional cointegration in both regimes, but not in the non-switching model. Conditioning on the states is also important when looking at the adjustment coefficients as the non-switching models can lead to wrong conclusions about the convergence of regional prices towards equilibrium. One important finding of this paper is that fractional cointegration does not in general occur in the congestion state, and when it does the mechanism is relatively weak. Hence price convergence of geograhical prices is a result of regime switching rather than error correction in a more conventional sense.

The remainder of the paper is structured as follows: We next offer a brief description of the structure of the Nordic electricity market. Section 3 introduces the data and argues for the importance of allowing for long memory, regime switching, and seasonality when building a model to describe the regional price processes. In section 4 the VAR modeling framework with long memory and regime switching is presented. In section 5 the empirical results are presented, including some forecasting results which generally favor the suggested model. Section 6 concludes.

2 The operation of the Nordic power market

Within the Nordic countries (Denmark, Finland, Norway, and Sweden), major electricity reforms were implemented during the 1990s. The deregulation process started in Norway in 1991, continued in Sweden 1996, in Finland 1998, and was finally completed in Denmark in 2000. As part of the liberalization, the national electricity markets were opened up for cross-border trade by establishment of a common power exchange, Nord Pool. Today all member countries of the Nordic power market have adapted to the new competitive environment and the Nordic exchange serves as a model for the restructuring of other power markets throughout the world.[1]

The per capita consumption of electricity is very high in Norway and Sweden, slightly lower in Finland and at EU average in Denmark. The relatively high consumption level in the Nordic countries is caused by a relatively electricity intensive industrial production, a cold climate, and extensive use of electric heating in homes and offices, especially in Norway and Sweden. The sources of electricity power production are rather mixed in the Nordic area as a whole. The major energy source is hydropower supplying approximately 65% of total electricity in years with normal precipitation. On the national level the power generation systems differ significantly and are generally dominated by one or two technologies. In Norway the share of hydropower is close to 100%, in Sweden it is close to 50%, in Finland around 15% and in Denmark 0%. With respect to nuclear power the share is 50% in Sweden, 30% in Finland, and 0% in Denmark and Norway. Power generation from fossil fuels is of major significance in Denmark and Finland, minor in Sweden, and close to non-existent in Norway. In Denmark 15-20% of the power supply originates from wind power turbines.[2]

Because hydropower production is mainly found in the northern parts of the Nordic power web and thermal power plants are located in the south, the relatively cheap hydropower generation is transmitted to the heavily populated southern regions, which of course requires a well established power grid transmission capacity to facilitate the flow. When the reservoir levels are adequate, the less costly hydropower production causes low spot prices. In these cases national and cross-border transmission systems will be used to their capacity in order to level out price discrepancies across

[1] For a detailed description of the Nordic power market, see Nord Pool (2003*a*) or Amundsen & Bergman (2007).

[2] Increasing the relative production of electricity by renewable energy sources has considerable political focus in Denmark. According to official energy plans 50% of the Danish electricity production will come from wind power in 2030.

regions. On the other hand, when reservoir levels are low there will be a net flow from south to north, and the market will see relatively high prices for thermally generated electricity.

From an institutional point of view there is a common Nordic market for electricity; however, even though key market institutions are common this does not mean that the Nordic electricity market is an integrated market in the sense that "the law of one price" applies. The reason is that the transmission of power is subject to possible capacity constraints. The Nordic electricity market constitutes a number of distinct geographical regions different from the countries themselves and several price areas may coexist. Whenever the relevant interconnector capacity is insufficient, the Nord Pool area is divided into two or more price areas. The separate power regions consist of Sweden (SWE), Finland (FIN), West Denmark (WDK), East Denmark (EDK), North Norway (NNO), Mid Norway (MNO), and South Norway (SNO). Thus Denmark and Norway are each divided into multiple geographical regions in Nord Pool.[3] This division reflects the grid of physical exchanges of power and the bidding areas with respect to the pricing of electricity as we shall explain shortly. Not all physical exchanges are connected to each other and only bilateral connections exist. Figure 1 displays the actual electricity exchange points and interconnections.

Figure 1 about here

The power spot market[4] operated by Nord Pool Spot A/S is an exchange where market participants trade power contracts for physical delivery the next day. This is referred to as a day-ahead market. The spot market is based on an auction with bids for purchase and sale of power contracts of one hour duration covering the 24 hours of the following day. At the deadline for the collection of all buy and sell orders the information is gathered into aggregate supply and demand curves for each power delivery hour. From these supply and demand curves the equilibrium spot price - referred to as the system price - is calculated.[5] Therefore, the system price is determined under the assumption that no transmission constraint is binding, and thus in a situation where no grid congestions exist across neighboring interconnectors there will be a single identical price across the areas with no congestions.

The actual trade is not necessarily carried out at the system price. When there is insufficient transmission capacity in a sector of the grid, a grid congestion will arise and the market system will establish different price areas across the geographical division of the Nord Pool area. The Nordic market is then partitioned into separate bidding areas which therefore become separate price areas when the contractual flow between bidding areas exceeds the capacity allocated by the transmission system operators for spot contracts. Within each price area the buyers pay, and the generators are paid, the corresponding area price. The difference between the area prices in two adjacent and connected price areas determines the congestion charge. Because separate prices may coexist depending upon regional supply and demand conditions, the relevant market definition will vary with time. In practice, several price area combinations will occur. Some hours there will only be a single price area (given by the system price), other hours there will be two or more price areas.

[3]For the purpose of analysis of the Norwegian regions, only the SNO region is considered in the present paper.

[4]Since only the spot market will be relevant for the present study, only this market will be described here, see also Nord Pool (2003*b*). Nord Pool (2003*c*) describes the futures and forward markets of the Nordic power exchange which are used for price hedging and risk management.

[5]The system price is the reference price in the financial power contracts like futures, forwards, and options traded at Nord Pool.

3 Data

The data used in this paper are (log transformed) hourly electricity spot prices for the Nord Pool area: West Denmark (WDK), East Denmark (EDK), South Norway (SNO), Sweden (SWE), and Finland (FIN).[6] The data set is the same as that analyzed in Haldrup & Nielsen (2006*a*, *b*) and covers the period 3 January 2000 to 25 October 2003, including weekends and holidays. For EDK the sample period starts 1 October 2000. The data series are displayed in Figure 2. Some stylized facts about the data are reported in Haldrup & Nielsen (2006*b*).

Figure 2 about here

A pronounced characteristic of electricity markets is the abrupt and generally unanticipated extreme changes in spot electricity prices, suggesting fat-tailed distributions, see Escribano, Peña & Villaplana (2002), Haldrup & Nielsen (2006*a*, *b*), and Koopman, Ooms & Carnero (2007). In Haldrup & Nielsen (2006*b*) a range of tests document that prices are neither I(0) nor I(1). Estimating the memory parameter for fractionally integrated, FI(d), processes shows that the series generally exhibit long memory with d in the range 0.31-0.52 with the SNO area being most persistent and in fact being nonstationary. The remaining areas have point estimates of d in the stationary region. It should be noted, however, that these estimates do not allow for regime dependence.

Another important aspect of electricity prices is the very strong seasonal behavior characterizing the series. Seasonality is mainly driven from the demand side and appears as seasonal variation within the day, within the week, and over the year. However, the supply side also contributes to seasonal variation as electricity production is highly dependent upon weather conditions. In particular, the seasonal variation in precipitation affects water reservoir levels in the generation of hydropower, and seasonal variation in wind conditions also plays an increasing role due to the growing number of wind turbines, especially in West Denmark.

Figure 3 about here

In Figure 3 scatter plots of log prices for connected Nord Pool areas are shown. When there are no capacity contraints across neigboring regions the prices will be identical, whereas congestion makes prices differ. Observations on the 45° line therefore represent non-congestion hours, whereas observations off the 45° line represent congestion hours. It is especially this marked difference in observations that motivates the present analysis.

4 Modeling of regime dependent long memory in spot electricity prices

In this section, we present our econometric model which is specifically motivated by the main properties and features of the Nordic spot electricity market. In particular, based on the structure of Nord Pool described in section 2, we include the switching between congestion and non-congestion regimes with state dependent dynamics. The model should also reflect the rich dynamic features of the data in the form of seasonality and long memory.

[6] Mid and North Norway are also member areas of Nord Pool, but are left out from the present analysis because these areas coincide with South Norway for most of the year.

4.1 A univariate model

We here briefly discuss the univariate model setup used in Haldrup & Nielsen (2006*b*). The main features that the estimation model should allow include seasonality, long memory, and regime switching of the type described above. Assume that individual electricity prices across connected regions are fractionally integrated in the non-congestion state. This means that an extreme form of fractional cointegration will exist in this state because the prices are identical in the two areas and thus price differences will be identically zero. On the other hand, the behavior of the two individual price series in the congestion state can be very different. If prices are compared without considering the different regime possibilities it is unclear what to expect from the data. However, the mixing of the two processes is likely to produce price series with a behavior that is a convex combination of the two state processes.

Consider the following model specification, which we denote a regime switching multiplicative RS-SARFIMA[7] model

$$A_{s_t}(L)\left(1 - a_{s_t}L^{24}\right)\Delta^{d_{s_t}} y_t = \varepsilon_{s_t,t}, \qquad \varepsilon_{s_t,t} \sim NID\left(0, \sigma^2_{s_t}\right). \tag{1}$$

Here, $\Delta^{d_{s_t}}$ is the fractional difference operator defined by its binomial expansion in the lag operator (see e.g. Hosking (1981)), $A_{s_t}(L)$ is a lag polynomial, and $s_t \in \{c, nc\}$ denotes the regime (c: congestion, nc: non-congestion), determined by a Markov chain with transition probabilities

$$P = \begin{bmatrix} p_{11} & 1 - p_{11} \\ 1 - p_{22} & p_{22} \end{bmatrix}. \tag{2}$$

Thus, for example, p_{11} denotes the probability that a congestion state will follow a congestion state, i.e. $\Pr(s_t = c | s_{t-1} = c)$. Note that because identical prices mean that we are in a non-congestion state, all regimes are observable, which contrasts the standard Markov switching model of Hamilton (1989) where the regimes follow a *latent* Markov process.

The (univariate) series y_t may denote one of the two individual log price series or the associated log relative price. The series y_t has been corrected for deterministic seasonality prior to the estimation whilst allowing interaction with the two observable regimes, that is, the coefficients on the dummy variables are allowed to differ across states. When y_t denotes a log relative price, all parameters are put to zero when $s_t = nc$, including σ^2_{nc}. Estimation of the above model is by conditional maximum likelihood and is discussed in detail in Haldrup & Nielsen (2006*b*).

4.2 A bivariate model

A disadvantage of the model described above is that parameters are estimated separately for the three price series (two individual prices and one relative price), when in fact the three price series to a large extent are governed by the same price shocks. We therefore consider the following fractional error

[7] RS-SARFIMA: Regime Switching Seasonal Autoregressive Fractionally Integrated Moving Average.

correction model specification for a bivariate regime switching vector stochastic process

$$\Delta^{d_{s_t}} \begin{pmatrix} p_{1t} \\ p_{2t} \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} \Delta^{\gamma_{s_t}} (p_{1,t-1} - p_{2,t-1}) + \sum_{i=1}^{k} \Gamma_{s_t,i} \Delta^{d_{s_{t-i}}} \begin{pmatrix} p_{1,t-i} \\ p_{2,t-i} \end{pmatrix} + \varepsilon_{s_t,t}, \tag{3}$$

where $s_t \in \{c, nc\}$, $\varepsilon_{c,t} \sim NID_2(0, \Omega)$, $\varepsilon_{nc,t} \sim (1,1)' NID_1(0, \sigma^2)$, $\gamma_{nc} \equiv 0$, γ_c is a free parameter, and

$$\Gamma_{c,i} = \begin{bmatrix} \Gamma^c_{11,i} & \Gamma^c_{12,i} \\ \Gamma^c_{21,i} & \Gamma^c_{22,i} \end{bmatrix}, \ \Gamma_{nc,i} = \begin{bmatrix} \Gamma^{nc}_{11,i} & \Gamma^{nc}_{12,i} \\ \Gamma^{nc}_{11,i} & \Gamma^{nc}_{12,i} \end{bmatrix},$$

such that the lagged fractional differences reflect whether a particular observation is associated with a congestion or non-congestion state. Thus, d_{nc} is the common fractional integration order in the non-congestion state, whereas d_c is the common integration order of the two price areas in the congestion state. Note that correlation between the two shocks in the congestion state is accommodated through the off-diagonal elements of Ω. However, $\varepsilon_{c,t}$ and $\varepsilon_{nc,t}$ are not correlated since both are never present at the same time.

Notice that the non-congestion state bilateral prices are identical, $p_{1t} = p_{2t} = p_t$, and hence the bivariate setup collapses to a pseudo-univariate model, i.e.

$$\Delta^{d_{nc}} p_t = \sum_{i=1}^{k} \left(\Gamma^{nc}_{11,i}, \Gamma^{nc}_{12,i}\right) \Delta^{d_{s_{t-i}}} \begin{pmatrix} p_{1,t-i} \\ p_{2,t-i} \end{pmatrix} + \varepsilon_{nc,t}. \tag{4}$$

Essentially, the price process switches between being generated from the univariate and bivariate models, where switching takes place in accordance with the transition probabilities in (2).

We limit our study to the bivariate setup and disregard potential spill-overs from the other areas. From a theoretical point of view, it appears conceptually possible to extend the present bivariate model to the multivariate case, and thereby model spill-overs using more advanced dynamics. However, from a computational point of view this appears infeasible in practice as the number of regimes, and thereby the number of parameters, grows very fast. Indeed, in a multivariate setup with M geographical regions, there are 2^{M-1} different regimes.

A number of remarks are in order. Consider first the non-congestion state. In this regime the two price series are forced to be governed by the same process (4) and hence any conditional forecast for this regime will remain identical for both price series. This feature is not captured in the univariate model of Haldrup & Nielsen (2006*b*) and indeed requires our multivariate setup. Thus, in particular, forecasts of each price series in the non-congestion state may appear different when based on (1), whereas forecasts based on (3) or (4) will be identical for the two price series in the non-congestion state. Note that in the non-congestion state the prices are fractionally integrated of order d_{nc} and fractionally cointegrated in the sense that the two series are identical. This notion of (fractional) cointegration is somewhat different than originally suggested by Granger (1981) and Engle & Granger (1987).

Next, consider the congestion regime. We will discriminate between two situations, i.e. when p_{1t} and p_{2t} cointegrate or do not cointegrate. (i) Assume first the situation with fractional cointegration. In this case the individual price series are $FI(d_c)$, but the log relative price is $FI(\gamma_c)$, where $\gamma_c < d_c$.

(ii) When prices do not cointegrate in the congestion regime, i.e. $\gamma_c \geq d_c$, $(\alpha_1, \alpha_2)'$ has no obvious interpretation in terms of adjustment parameters.

Under congestion the adjustment coefficients, $(\alpha_1, \alpha_2)'$, may give an indication of whether the specific price areas adjust towards equilibrium. Specifically, if $\alpha_1 > 0$ then p_{1t} is moving away from equilibrium (non-congestion), whereas if $\alpha_2 > 0$ then p_{2t} is moving towards equilibrium. Note that the full stability of the model requires that the entire system dynamics is included in the calculation, but in any case the values of α_1 and α_2 give a rough idea of the system dynamics under a ceteris paribus assumption. An alternative interpretation of the adjustment coefficients follows from the market setup and varying costs of electricity production in different geographical regions. For example, if there is no congestion between SNO and WDK, prices are identical and electricity flows from the cheaper area (usually SNO because of the hydropower) to the more expensive area (WDK). However, if there is congestion, prices in WDK will be higher reflecting the higher costs of electricity production. This increase in price in WDK corresponds to $\alpha_1 > 0$ in the WDK-SNO bivariate model, i.e. a move away from equilibrium. Importantly, this is not due to system instability but rather reflects that electricity is more expensive to produce in WDK compared to SNO. Hence the estimated error correction model need not be given a standard interpretation.

The model analyzed in this paper is unique in the literature on regime switching and/or (fractionally) cointegrated models since it collapses to a pseudo-univariate model in one of the regimes. The error correction model specification (3) reflects the particular structure and features of the market design. For discussions of representation theory in the context of (non-switching) fractional cointegration, see Granger (1986), Davidson (2002), Robinson & Yajima (2002), Davidson, Peel & Byers (2006), and Johansen (2008).

4.3 Estimation

In our case, congestion and non-congestion are observed states such that regimes are known, and the maximum likelihood estimates of the transition probabilities in (2) are

$$\hat{p}_{11} = \frac{n_{c,c}}{n_{c,c} + n_{c,nc}}, \; \hat{p}_{22} = \frac{n_{nc,nc}}{n_{nc,c} + n_{nc,nc}}, \tag{5}$$

where n_{ij} is the number of times we observe regime i followed by regime j for $i, j \in \{c, nc\}$.

Estimation of the remaining parameters of the two states is done by quasi conditional maximum likelihood. The regime-specific log-likelihood functions, omitting the constant, are

$$\begin{aligned} l_c(d_c, \theta_c) &= -\frac{\sum_t \mathbf{1}\{s_t = c\}}{2} \log|\Omega| - \frac{1}{2}\sum_t \operatorname{trace}\left(\Omega^{-1}\varepsilon_{s_t,t}\mathbf{1}\{s_t = c\}\varepsilon'_{s_t,t}\right), \\ l_{nc}(d_{nc}, \theta_{nc}) &= -\frac{\sum_t \mathbf{1}\{s_t = nc\}}{2} \log\sigma^2 - \frac{1}{2}\sum_t \left(\sigma^{-2}\varepsilon_{s_t,t}\mathbf{1}\{s_t = nc\}\varepsilon'_{s_t,t}\right), \end{aligned}$$

where $\mathbf{1}\{A\}$ is the indicator function of the event A. The full-sample log-likelihood function is given by

$$l(d_c, d_{nc}, \theta) = -\frac{T}{2}\log(2\pi) + l_c(d_c, \theta_c) + l_{nc}(d_{nc}, \theta_{nc}), \tag{6}$$

Table 1: Estimated transition probabilities (mean durations of states in hours)

Link	$\hat{p}_{11}$ $(c \to c)$	$\hat{p}_{12}$ $(c \to nc)$	$\hat{p}_{21}$ $(nc \to c)$	$\hat{p}_{22}$ $(nc \to nc)$
EDK-SWE	0.7848 (4.65)	0.2152	0.0131	0.9869 (76.57)
WDK-SWE	0.8216 (5.60)	0.1784	0.1259	0.8740 (7.94)
WDK-SNO	0.9247 (13.28)	0.0753	0.1221	0.8779 (8.19)
SNO-SWE	0.9478 (19.16)	0.0523	0.0462	0.9538 (21.64)
SWE-FIN	0.8505 (6.51)	0.1495	0.0210	0.9790 (48.78)

Notes: The table presents estimated transition probabilities for each bivariate model based on (5). Numbers in parentheses are estimated mean durations of states (in hours).

which is maximized numerically.[8]

Finally, we remark that our model framework assumes that states are observable and that the cointegrating vector in the congestion state, $\beta = (1, -1)$, is given. Therefore, asymptotic distribution theory for the remaining parameters will be standard under suitable regularity conditions on the errors $\varepsilon_{s_t,t}$, such as serial independence and moment conditions, see e.g. Tanaka (1999). In particular, Gaussianity of the errors is not a necessary condition for the asymptotic distribution theory, but is used only to derive the likelihood function. This property of the estimation methodology is especially important in dealing with the fat tails present in the data.

5 Empirical results

Prior to estimation, each log price series had deterministic seasonality removed by regression on a constant, a time trend, dummy variables for hour-of-day, day-of-week, month-of-year, and a holiday dummy. For the switching models the parameter estimates of the deterministics are allowed to differ across states. We include lags $1, ..., 8$, and 12 to capture within-the-day effects, and we also include a 24th lag to capture the daily stochastic seasonality.[9]

5.1 Estimation of transition dynamics

Since the states are observable, estimates of the transition probabilities for each state are easily calculated according to (5) and are reported in Table 1. It is clear that some grid points are more subject to congestion than others. This fact may be explained by demand and supply fluctuations, but there is also the possibility that congestion may be caused by exploitation of market power and hence calling for further economic analysis of the sources of congestion.

The estimated transition probabilities indicate a high degree of persistence of the states. The probability of staying in the congestion regime, $\hat{p}_{11}$, is highest for the grid point SNO-SWE, i.e. 0.9478, whereas it is lowest for EDK-SWE link, 0.7848. This corresponds to a mean duration of 19.16 and 4.65 hours, respectively. In general, the probability of staying in the non-congestion regime, $\hat{p}_{22}$,

[8] We have used the fractional integration estimates from (1) as our starting values. For the remaining parameters, i.e. autoregressive and variance-covariance terms etc., we find the starting values by letting the fractional integration parameters be fixed and maximizing the log-likelihood with respect to the remaining parameters. We did not notice significant dependence on the choice of starting values in any of our models.

[9] Note that for the univariate model (1) we have here chosen a richer dynamics compared to Haldrup & Nielsen (2006*b*), and hence the estimation results are not exactly identical.

Table 2: Estimates of the fractional integration and cointegration parameters

				Switching					
	No switching			Non-congestion			Congestion		
Model	$\hat{d}_1$	$\hat{d}_2$	$\hat{\gamma}$	$\hat{d}_1^{nc}$	$\hat{d}_2^{nc}$	$\hat{\gamma}^{nc}$	$\hat{d}_1^{c}$	$\hat{d}_2^{c}$	$\hat{\gamma}^{c}$
Panel A: EDK-SWE									
Univariate	0.42 (0.014)	0.47 (0.015)	0.21 (0.017)	0.47 (0.015)	0.48 (0.014)	0	0.01 (0.016)	−0.02 (0.015)	0.57 (0.046)
VAR	0.57 (0.009)		0.23 (0.047)	0.46 (0.003)		0	0.11 (0.004)		0.00 (0.006)
Panel B: WDK-SWE									
Univariate	0.28 (0.018)	0.43 (0.014)	0.30 (0.017)	0.29 (0.021)	0.34 (0.017)	0	0.28 (0.021)	0.51 (0.022)	0.33 (0.019)
VAR	0.61 (0.019)		0.55 (0.009)	0.25 (0.018)		0	0.32 (0.023)		0.10 (0.026)
Panel C: WDK-SNO									
Univariate	0.28 (0.018)	0.54 (0.014)	0.30 (0.016)	0.25 (0.035)	0.33 (0.014)	0	0.29 (0.020)	0.56 (0.020)	0.33 (0.017)
VAR	0.65 (0.011)		0.30 (0.005)	0.30 (0.004)		0	0.46 (0.004)		0.30 (0.006)
Panel D: SNO-SWE									
Univariate	0.54 (0.014)	0.43 (0.014)	0.31 (0.015)	0.52 (0.017)	0.49 (0.016)	0	0.36 (0.025)	0.22 (0.022)	0.33 (0.018)
VAR	0.67 (0.007)		0.59 (0.010)	0.65 (0.003)		0	0.26 (0.004)		0.24 (0.015)
Panel E: SWE-FIN									
Univariate	0.43 (0.014)	0.39 (0.013)	0.30 (0.017)	0.41 (0.014)	0.43 (0.013)	0	0.40 (0.012)	0.39 (0.013)	0.36 (0.022)
VAR	0.61 (0.008)		0.28 (0.014)	0.35 (0.005)		0	0.03 (0.003)		0.02 (0.007)

Notes: The table presents quasi maximum likelihood estimates for the models (1) and (3). Subscripts denote the geographical region and superscripts denote the state. Note that $d_1 = d_2 = d$ is assumed in the VAR model (3). Robust standard errors based on the sandwich formula are given in parentheses.

is higher, estimated at $0.8740 - 0.9869$, corresponding to mean durations of $7.94 - 76.57$ hours.

5.2 Estimation of fractional integration and cointegration parameters

Table 2 presents estimates of the fractional integration order d for a number of different cases. The models estimated under the heading "No switching" use pooled data, i.e. the data is not separated by congestion and non-congestion periods. The results presented under the heading "Switching" refer to the corresponding estimates when data is partitioned into congestion and non-congestion periods, where we use superscripts c or nc to denote estimates under the congestion and non-congestion regimes, respectively. "Univariate" and "VAR" refer to the models (1) and (3), respectively. The estimates of d_1 and d_2 are the fractional orders for the first and second regions, respectively, whereas the estimate γ is the fractional integration order of the log relative price. Note that, for the swiching model, $\gamma^{nc} \equiv 0$ in the non-congestion state because the individual price series are identical and hence the series are fractionally cointegrated in an extreme form. Furthermore, observe that in the VAR model it is imposed that $d_1 = d_2 = d$.

Figures 4 and 5 about here

Kernel spectral density estimates of the residuals from the non-switching and switching models together with the observed (deseasonalized) price series are presented in Figures 4 and 5, respectively. Generally, when comparing the residuals for both the non-switching and switching models with the observed counterparts, the long memory feature of the specific area prices appears to be appropriately captured by the models. For the WDK-SNO connection, however, there seems to be some long memory left in the residuals of the SNO area price for the switching model. Possibly this stems from forcing the fractional integration orders to be identical in the congestion regime, and thereby estimating an integration order that is lower than the individual fractional integration order of SNO. Another justification for the regime switching model is that there are fewer outliers in the residuals when compared to the residuals from the non-switching models, but tails are still more heavy than normally distributed residuals. However, as discussed in section 4.3 this is not critical to the estimation procedure.

Consider now the East Denmark-Sweden connection exhibited in Panel A of Table 2, and consider initially the pooled data set without regime switching. For the univariate model the point estimates of d for the two regions are very similar, 0.42 and 0.47, and the estimate of γ associated with the relative price is somewhat lower, i.e. 0.21. When we use the VAR model the point estimate of d for the individual price series is 0.57 and the estimate of γ is 0.23. The results indicate that when data is not classified according to regimes, there is evidence of fractional cointegration between the series. Now, the question is whether this result is caused by the non-congestion state dominating the sample or whether both regimes contribute to the cointegration finding. In the regime switching case, the non-congestion estimates clearly indicate cointegration (as expected) with point estimates of d^{nc} at 0.46 for the VAR model and similar point estimates for the univariate model. In the congestion case, the regime switching results for the univariate model do not make sense because $\hat{\gamma}^c > \max\{\hat{d}_1^c, \hat{d}_2^c\}$. This finding may be caused by adopting a univariate modeling approach when joint modeling is more appropriate. In fact, for the VAR case the point estimate of d^c is 0.11 and there is indication of (weak) fractional cointegration since the relative price is FI(0).

The West Denmark-Sweden link in Panel B is an interesting case where there seems to be no fractional cointegration in the non-switching models. However, looking at the models where we condition on congestion/non-congestion, we see that there is fractional cointegration in the non-congestion state. In the VAR model, there is in fact cointegration in both states. That is, the results from the non-switching models (which are clearly misspecified) are thus some combination of their regime switching counterparts. It is clear that by not taking regime switching into account we falsely conclude that there is no sign of fractional cointegration, when in fact regime dependent fractional cointegration exists.

The West Denmark-South Norway link with estimates in Panel C are similar to the West Denmark-Sweden link in Panel B, so (weak) fractional cointegration occurs in the congestion state in the VAR model. Note that there is also evidence of fractional cointegration for the VAR model when not conditioning on regime switching.

As seen from Panel D, presenting estimates for the link between South Norway and Sweden, no (or extremely weak) evidence of fractional cointegration is found for the models without regime switching. However, when conditioning on states, it is seen that it is only in the non-congestion state that cointegration occurs. Interestingly, based on the VAR model, prices in this state seem

Table 3: Estimated adjustment coefficients

Series	No switching $\hat{\alpha}_1$	No switching $\hat{\alpha}_2$	Switching $\hat{\alpha}_1$	Switching $\hat{\alpha}_2$
EDK-SWE	**4.5295***	**0.0289**	**1.4269***	**−0.1990***
	(0.2789)	(0.1965)	(0.2267)	(0.1007)
WDK-SWE	−0.3163*	0.0049	**0.0135**	**−0.0277***
	(0.0184)	(0.0043)	(0.0378)	(0.0088)
WDK-SNO	**−0.4309***	**0.0042**	**0.1338***	**−0.1053***
	(0.0213)	(0.0035)	(0.0149)	(0.0032)
SNO-SWE	−0.2373	3.2256*	0.6729*	0.0754
	(0.1375)	(0.1952)	(0.0820)	(0.1128)
SWE-FIN	**0.3699**	**−5.7785***	1.6996*	−0.0365
	(0.4661)	(0.5639)	(0.0971)	(0.1545)

Notes: Subscripts denote the geographical region. Numbers in bold face refer to situations with indication of fractional cointegration based on the VAR estimates of d and γ reported in Table 2. Robust standard errors are given in parentheses. An asterisk denotes significance at the 5% level.

non-stationary whilst relative prices are stationary.

Finally, for the Sweden-Finland link in Panel E there is some evidence of fractional cointegration in the non-switching models. For both the univariate and VAR models there is cointegration in the non-congestion state, whereas there is no cointegration (univariate model) or all series seem to be I(0) (VAR model) in the congestion state. Hence, the non-congestion state seems to dominate the pooled data set.

5.3 Estimation of adjustment coefficients

An advantage of the regime switching VAR model (3)-(4) compared to univariate models is that estimates of the adjustment coefficients in the congestion state, i.e. the parameters $(\alpha_1, \alpha_2)'$, can be obtained. The adjustment coefficients indicate (ceteris paribus) the price move directions in response to a particular gap between the area prices under congestion. An alternative interpretation of the adjustment coefficients can be given in our model compared to standard error correction models where price changes respond to disequilibrium. This follows from the market setup and varying costs of electricity production in different geographical regions. For example, if an inexpensive electricity supply from another geographical region is suddenly stopped due to congestion, prices are expected to be higher until non-congestion is restored, which may result in adjustment parameters indicating a move away from equilibrium defined as the case where area prices are identical. Parameter interpretation is of course an issue here, because we force the cointegrating vector to be $(1, -1)$ and the parameters α_1, α_2, and γ do not have the usual interpretation in the congestion state if in fact there is no cointegration present in that state.

In Table 3 the estimated adjustment coefficients $(\hat{\alpha}_1, \hat{\alpha}_2)'$ associated with the VAR models are reported for the switching and non-switching cases. Numbers in boldface font indicate situations where, based upon the d and γ estimates, some degree of fractional cointegration is suggested by the results in Table 2. In the regime switching models, boldface indicates situations where there appears to be cointegration in the congestion state.

Consider first the East Denmark-Sweden connection. When we do not condition on regime switching, prices in East Denmark move away from the steady state solution with identical area prices, whereas Swedish price adjustment appears to be insignificant. When we condition on regime switching, prices in both East Denmark and Sweden appear to depart from steady state. This contradicts the standard interpretation of error correction adjustment. Recall that if there is no congestion between EDK and SWE, prices are identical and electricity flows from the cheaper area (usually SWE because of the hydropower and nuclear electricity production) to the more expensive area (usually EDK because of the majority of electricity production stemming from thermal plants). Therefore, when congestion occurs, prices in East Denmark will usually be higher and thus reflect the higher marginal cost of electricity production in East Denmark compared to Sweden. This increase in price in East Denmark corresponds to $\alpha_1 > 0$ in the EDK-SWE bivariate model, i.e. a move away from equilibrium. Importantly, this is not due to system instability but rather due to electricity being more expensive to produce in East Denmark compared to Sweden.

Next, we look at the West Denmark-Sweden link. In this case, no cointegration was found for the non-switching model and for the switching case there was cointegration in both states for the VAR model. Conditioning on regime switching both prices tend to depart and thus further extending the price gap. However, the adjustment seems weak in this case and only the Swedish adjustment parameter is significant.

For the West Denmark-South Norway connection we found signs of cointegration for both the non-switching and switching models. When not conditioning on regime switching the adjustments parameters have the conventional signs, albeit the adjustment coefficient for the South Norway area is small and insignificant. However, this finding may be spurious because, when we condition on regimes, both area prices depart from equilibrium and hence the price gap is widened following the argument previously given: When congestion occurs, prices in West Denmark will be higher reflecting the higher costs of electricity production. If demand continues to increase in West Denmark during the congestion more expensive generators will be taken into use and thus increasing marginal cost of production even further. This increase in price in West Denmark corresponds to $\alpha_1 > 0$ in the WDK-SNO bivariate model, i.e. a further increase in the price gap. Again, this is not due to system instability but rather due to electricity being more expensive to produce in West Denmark compared to South Norway.

The South Norway-Sweden and Sweden-Finland cases are similar in the sense that no cointegration was found in the congestion state, and therefore the interpretation of the adjustment coefficients is less interesting.

To sum up, appropriate modeling of the regime switching feature is seen to have a major impact on the electricity price dynamics. In addition to giving estimates of the adjustment parameters specific to particular states, conditioning on congestion or non-congestion allows interpretation of the adjustment coefficients, which is different from standard error correction models. In particular, we have found evidence that, when fractional cointegration takes place, the price gap is widened under congestion. All in all, what ensures price convergence is in fact the switching mechanism towards non-congestion where prices are identical rather than error correction in a more conventional sense during congestion periods.

5.4 Forecasting spot electricity prices

In this section, we consider forecasting of spot electricity prices for up to 24 hours. Because the calculation of analytical forecast bands for the k-step ahead forecast requires 2^{k+1} steps, it is not computationally possible to do the forecasting exercise analytically (although the formulae are available, e.g. Davidson (2004)). Therefore, we consider Monte Carlo stochastic simulation forecasting, see Davidson (2004). The method implemented here differs from the one used by Davidson (2004) because our states are observable. Our forecast simulation differs from the setup used in Haldrup & Nielsen (2006*b*), where the individual price series are estimated separately, and therefore leads to different forecasts when in fact the individual prices under non-congestion are governed by exactly the same price shocks.

The forecasting exercise is implemented by simulating the model 24 periods ahead assuming independent draws from the estimated residuals. The states are also simulated 24 periods ahead using the estimated transition probabilities in Table 1. The median and the 95% forecast error bands for each period are extracted from $10,000$ simulated forecasts.

Figures 6-10 about here

Figures 6-10 display the forecasting results for both the univariate model and the VAR model. Each figure contains 2 panels displaying the results for the non-switching model and 2 panels for the switching model. In each panel, the diamonds depict the (deseasonalized) observed values covering the last 24 in-sample observations as well as 24 out-of-sample observations. Each panel also has three solid and three dotted lines. The three solid lines are median forecasts and 95% error bands for the VAR model, whereas the three dotted lines are the equivalent forecasts and bands for the univariate model. Notice that we have displayed the non-switching and switching models in separate panels because the (deseasonalized) observed values are different for these two cases. In Table 4 the mean absolute forecast errors (MAFE) for the different models are reported where the forecasted values are the simulated median price forecasts. Mean squared forecast errors were also calculated and yielded qualitatively very similar results which are not presented.

Figure 6 displays the forecasts for the East Denmark-Sweden physical link. First, considering the non-switching models, we observe that the forecasts from the univariate model slightly outperform the VAR model, which is also confirmed in terms of MAFE when looking at Panel A of Table 4. However, the confidence bands for the VAR model are tighter than for the univariate model. Focusing on the regime switching models, the median forecasts for both models are very close to the actually observed (deseasonalized) values. The 95% error bands for the EDK log price series are tighter for the VAR model, and for the SWE log price series they are better initially and similar for later hours. The MAFE for the switching VAR model is smaller than that of the misspecified non-switching model.

In Figure 7 the forecasts for the West Denmark-Sweden physical link are displayed. Overall, we notice again that the 95% error bands are tighter for the bivariate model than for the univariate model. In the non-switching case neither model performs particularly well for the WDK price series, but for the SWE price series they both perform considerably better. This is also confirmed when looking at Panel B of Table 4. In the regime switching case, the VAR model outperforms the univariate model for the first 5 hours for the WDK price series and is pretty close to the actual observed series.

Table 4: Mean absolute forecast error (MAFE)

	No switching		Switching	
	Univariate	VAR	Univariate	VAR
Panel A: EDK-SWE				
EDK	0.0639	0.0882	0.0555	0.0520
SWE	0.0502	0.0645	0.0559	0.0555
Panel B: WDK-SWE				
WDK	0.2518	0.2256	0.1945	0.1929
SWE	0.0538	0.0502	0.2862	0.1165
Panel C: WDK-SNO				
WDK	0.2518	0.4847	0.2546	0.2119
SNO	0.0359	0.0342	0.3405	0.1105
Panel D: SNO-SWE				
SNO	0.0368	0.0352	0.0810	0.0809
SWE	0.0538	0.1523	0.1225	0.0909
Panel E: SWE-FIN				
SWE	0.0549	0.0521	0.0408	0.0406
FIN	0.0833	0.0990	0.0845	0.0747

Subsequently, it degenerates to its unconditional mean. Regarding the Swedish price series none of the two models perform adequately which is also confirmed in Panel B of Table 4.

Figure 8 considers the West Denmark-South Norway physical link. In the non-switching case the VAR model produces very good forecasts of the South Norway price series and slightly outperforms the univariate model. When considering the switching case the VAR model clearly outperforms the univariate model, see also Panel C of Table 4.

Figure 9 displays the forecast results for the South Norway-Sweden physical link. Without regime switching the VAR model underestimates the observed price series for Sweden, whereas for the univariate model the median forecast is close to the actually observed price series. With regime switching both models are very close to the actually observed price series. The VAR model outperforms the univariate model in terms of tightness of the confidence band. Panel D of Table 4 shows that in the regime switching case, the VAR model is again superior to the univariate model.

Finally, Figure 10 displays the forecasts for the Sweden-Finland connection. Here the VAR model seems to do much better than the univariate model in terms of forecasting the observed price series. This is also the conclusion drawn from Panel E of Table 4, where it is seen that the switching model outperforms the non-switching model. Furthermore, we again observe that the forecast confidence bands from the univariate models are wider than those from the VAR model in the non-switching case and indeed a lot wider in the switching case.

To conclude, the regime switching VAR model proposed in this paper seems to provide overall better forecasts compared to its univariate and non-switching counterparts. In general, forecast confidence bands are more narrow for the switching VAR. In 7 out of 10 cases the switching VAR model delivers smaller MAFE compared to the non-switching VAR model, and in 10 out of 10 cases the switching VAR model outperforms the switching univariate model in terms of MAFE.

6 Conclusion

In this paper we have proposed a multivariate modeling framework for spot electricity prices within the Nord Pool power grid which enables us to describe the complex price dynamics characterizing this market. When the actual transmission of electricity is constrained by the flow capacity, congestion will occur and hence the presence or absence of transmission bottlenecks may have implications for the price dynamics. Moreover, it is an empirical regularity that electricity prices exhibit long memory in the form of fractional integration which may be regime dependent. Our multivariate fractional cointegration model is new and is motivated by these particular features, and thus allows us to explicitly take into account the fact that, in non-congestion periods, prices are the same across geographical regions and therefore also governed by exactly the same price shocks.

From our empirical analysis it is clear that conditioning on congestion or non-congestion states has a major impact on the dynamics of the electricity prices, and this feature is well described by the VAR model for both estimation and forecasting. In fact, when not conditioning on the specific states very misleading conclusions may be drawn with respect to the potential fractional cointegration properties of the data and the adjustment mechanism describing the price behavior. We find that what ensures price convergence is in fact the switching mechanism towards non-congestion where prices are identical, rather than error correction occurring in a more conventional sense. We believe this is the first empirical example demonstrating that the standard interpretation of error correction models may break down when in fact a dynamic non-linear feature characterizes the data.

There are three possible types of misclassification of the model dynamics in the empirical analysis. First, non-switching models may indicate that the price series are fractionally cointegrated, whereas when conditioning on states this is only the case in the non-congestion state (which is cointegrated by definition). Second, the non-switching model could indicate that there is no fractional cointegration when in fact there is cointegration in the non-congestion state, and finally there is the possibility of fractional cointegration in both regimes, but not in the non-switching model.

We also emphasize the appropriateness of our VAR model in terms of forecasting, where more narrow forecast confidence bands are delivered. In 7 out of 10 cases a smaller MAFE is obtained from the switching VAR model compared to a non-switching VAR model, and in 10 out of 10 cases the regime switching VAR model outperforms its univariate counterpart in terms of MAFE.

For future research we would like to point to the fact that some geographical regions are indirectly connected, e.g. West Denmark and East Denmark are indirectly connected through Sweden, so there are regimes where West Denmark and East Denmark constitute the same price area. The effects of these indirect links between geographical regions and how they potentially affect the price dynamics is therefore of major interest. A detailed analysis similar to the analysis presented in this paper including indirect links is conceptually straightforward using a higher-dimensional model, but computationally the analysis is difficult and left for future research.

References

Amundsen, E. S. & Bergman, L. (2007), 'Integration of multiple national markets for electricity: the case of Norway and Sweden', *Energy Policy* **35**, 3383–3394.

Davidson, J. (2002), 'A model of fractional cointegration, and tests for cointegration using the bootstrap', *Journal of Econometrics* **110**, 187–212.

Davidson, J. (2004), 'Forecasting Markov-switching dynamic, conditionally heteroscedastic processes', *Statistics & Probability Letters* **68**, 137–147.

Davidson, J., Peel, D. & Byers, J. (2006), 'Support for governments and leaders: fractional cointegration analysis of poll evidence from the UK, 1960-2004', *Studies in Nonlinear Dynamics & Econometrics* **10**, issue 1, article 3.

Diebold, F. X. & Inoue, A. (2001), 'Long memory and regime switching', *Journal of Econometrics* **105**, 131–159.

Engle, R. F. & Granger, C. W. J. (1987), 'Co-integration and error correction: representation, estimation, and testing', *Econometrica* **55**, 251–276.

Escribano, Á., Peña, J. I. & Villaplana, P. (2002), 'Modeling electricity prices: international evidence', *Universidad Carlos III, Departamento de Economía, Working Paper* **02-27**.

Fabra, N. & Toro, J. (2005), 'Price wars and collusion in the Spanish electricity market', *International Journal of Industrial Organization* **23**, 155–181.

Granger, C. W. J. (1981), 'Some properties of time series data and their use in econometric model specification', *Journal of Econometrics* **16**, 121–130.

Granger, C. W. J. (1986), 'Developments in the study of cointegrated economic variables', *Oxford Bulletin of Economics and Statistics* **48**, 213–228.

Granger, C. W. J. & Ding, Z. (1996), 'Varieties of long memory models', *Journal of Econometrics* **73**, 61–77.

Granger, C. W. J. & Hyung, N. (2004), 'Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns', *Journal of Empirical Finance* **11**, 399–421.

Granger, C. W. J. & Joyeux, R. (1980), 'An introduction to long-memory time series models and fractional differencing', *Journal of Time Series Analysis* **1**, 15–29.

Haldrup, N. & Nielsen, M. Ø. (2006*a*), 'Directional congestion and regime switching in a long memory model for electricity prices', *Studies in Nonlinear Dynamics & Econometrics* **10**, issue 3, article 1.

Haldrup, N. & Nielsen, M. Ø. (2006*b*), 'A regime switching long memory model for electricity prices', *Journal of Econometrics* **135**, 349–376.

Hamilton, J. D. (1989), 'A new approach to the economic analysis of nonstationary time series and the business cycle', *Econometrica* **57**, 357–384.

Hosking, J. R. M. (1981), 'Fractional differencing', *Biometrika* **68**, 165–176.

Johansen, S. (2008), 'A representation theory for a class of vector autoregressive models for fractional processes', *Econometric Theory* **24**, 651–676.

Koopman, S. J., Ooms, M. & Carnero, M. A. (2007), 'Periodic seasonal Reg-ARFIMA-GARCH models for daily electricity spot prices', *Journal of the American Statistical Association* **102**, 16–27.

Motta, M. (2004), *Competition Policy, Theory and Practice*, Cambridge University Press, Cambridge, UK.

Nord Pool (2003*a*), 'Derivatives trade at Nord Pool's financial market', *Working Paper, www.nordpool.no* .

Nord Pool (2003*b*), 'The Nordic power market: electricity power exchange across national borders', *Working Paper, www.nordpool.no* .

Nord Pool (2003*c*), 'The Nordic spot market: the world's first international spot power exchange', *Working Paper, www.nordpool.no* .

Robinson, P. M. & Yajima, Y. (2002), 'Determination of cointegrating rank in fractional systems', *Journal of Econometrics* **106**, 217–241.

Sherman, P. (1989), *The Regulation of Monopoly*, Cambridge University Press, Cambridge, UK.

Tanaka, K. (1999), 'The nonstationary fractional unit root', *Econometric Theory* **15**, 549–582.

Figure 1: Map of the Nord Pool area

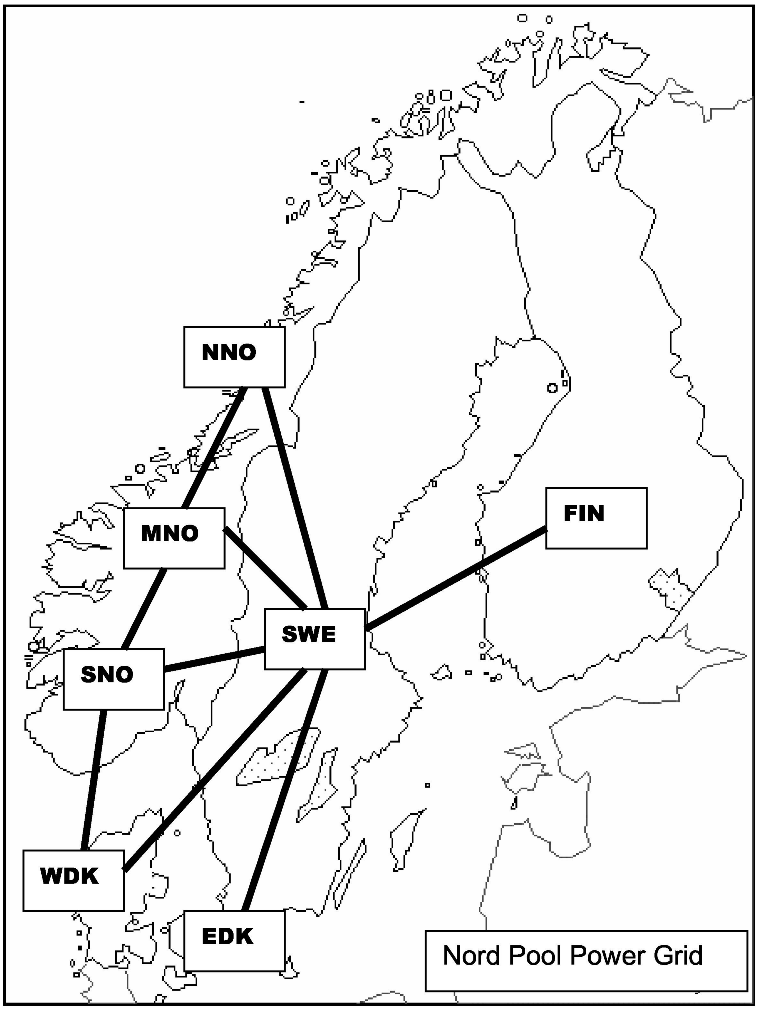

Figure 2: Hourly log spot electricity prices for the Nord Pool area covering the period 3 January 2000 to 25 October 2003

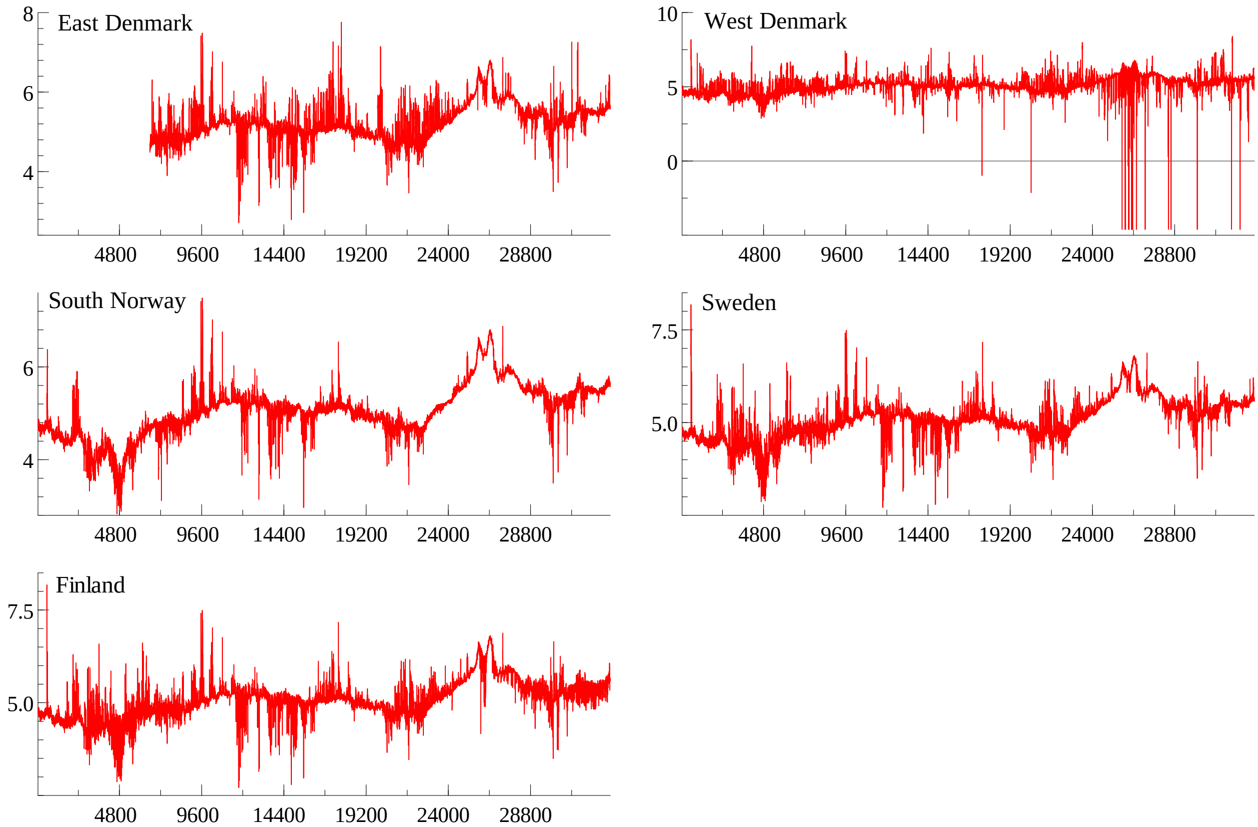

Figure 3: Scatter plots of hourly log prices across Nord Pool regions

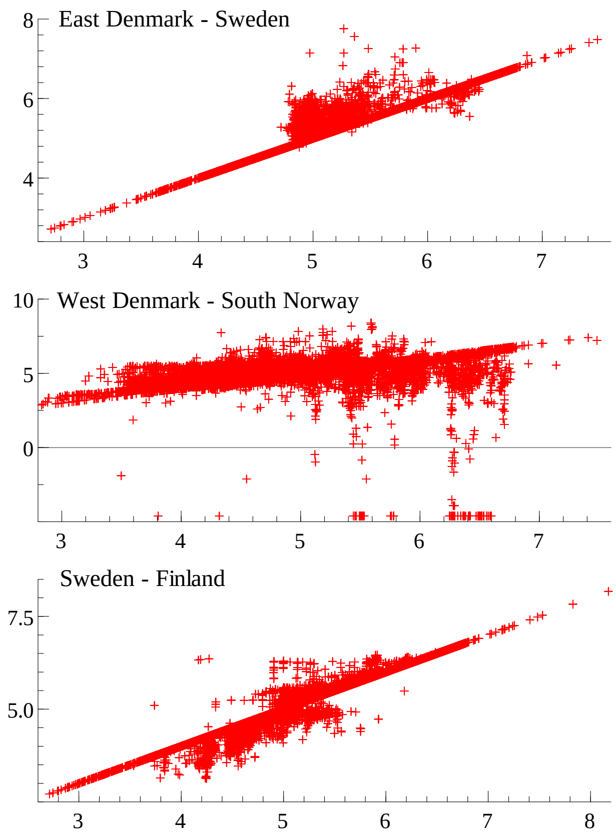

Figure 4: Kernel spectral density estimates of electricity price series and residuals without regime switching

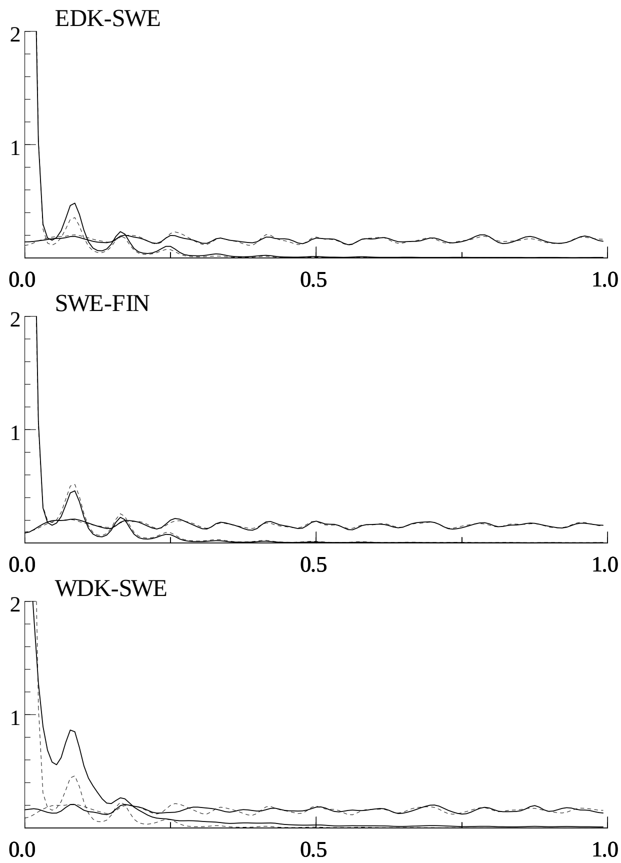

Note: In each panel the solid and dotted lines constitute the first and second area in the physical link, respectively. The spectral density curves with the most mass at the zero frequency are for the observed deseasonalized price series.

Figure 5: Kernel spectral density estimates of electricity price series and residuals with regime switching

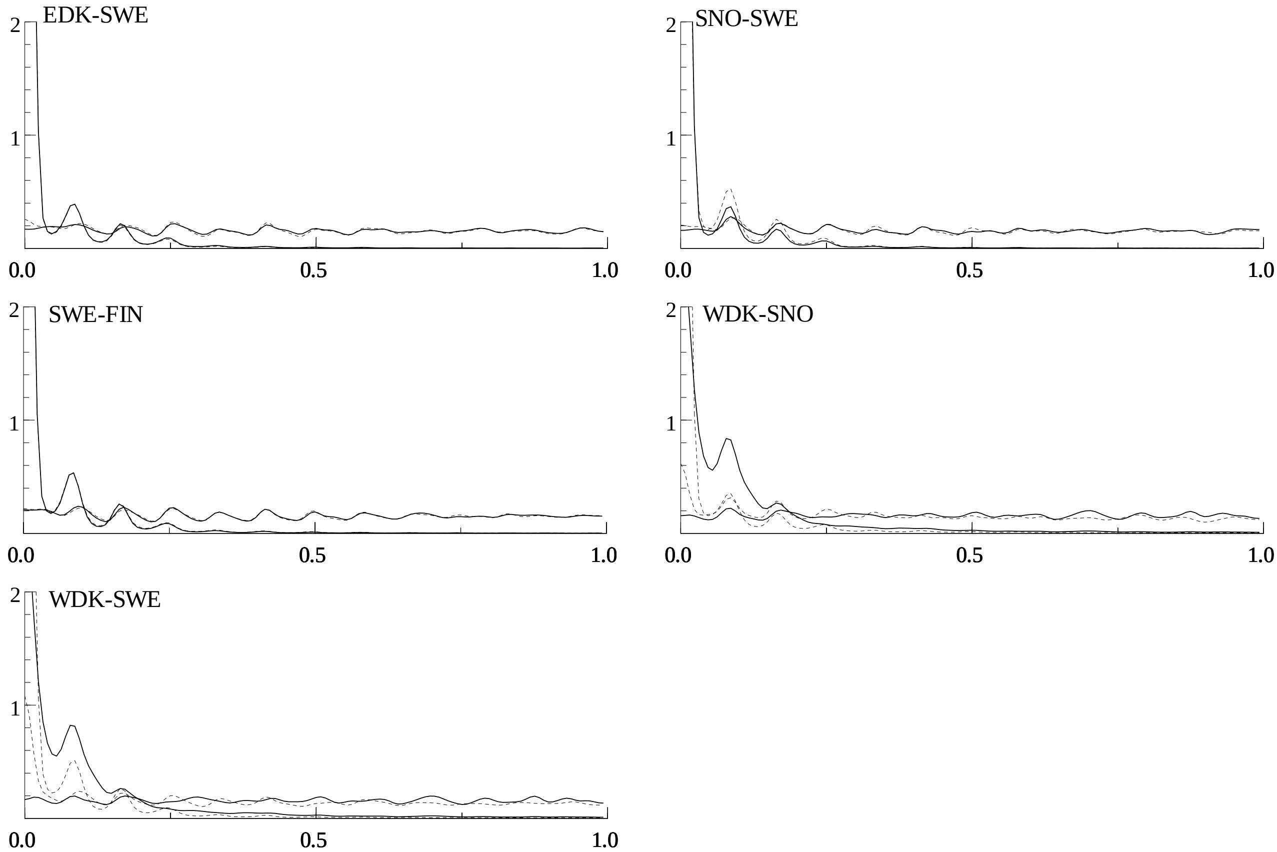

Note: In each panel the solid and dotted lines constitute the first and second area in the physical link, respectively. The spectral density curves with the most mass at the zero frequency are for the observed deseasonalized price series.

Figure 6: Forecasts for the EDK-SWE physical link for the non-switching (NoRS) and switching (RS) models

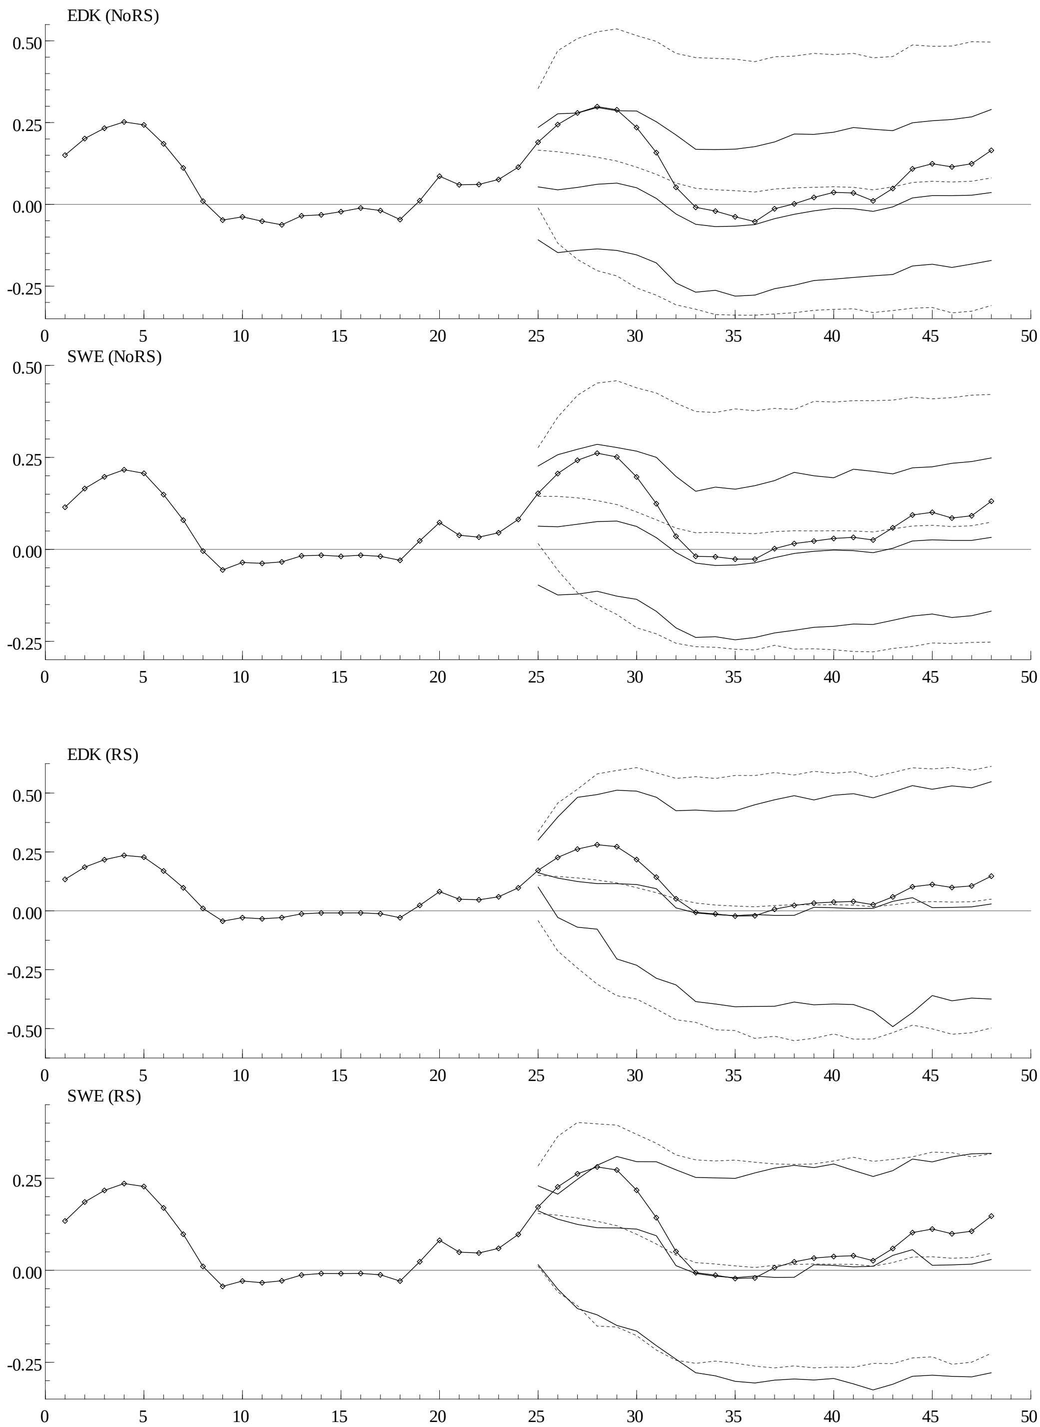

Note: In each panel, the solid line with the diamonds are the actually observed (deseasonalized) price series covering the last 24 in-sample observations as well as the 24 out-of-sample observations. Each panel also has a three solid lines and three dotted lines. The three solid and dotted lines are median forecasts and error bands for the VAR model and univariate model, respectively.

Figure 7: Forecasts for the WDK-SWE physical link for the non-switching (NoRS) and switching (RS) models

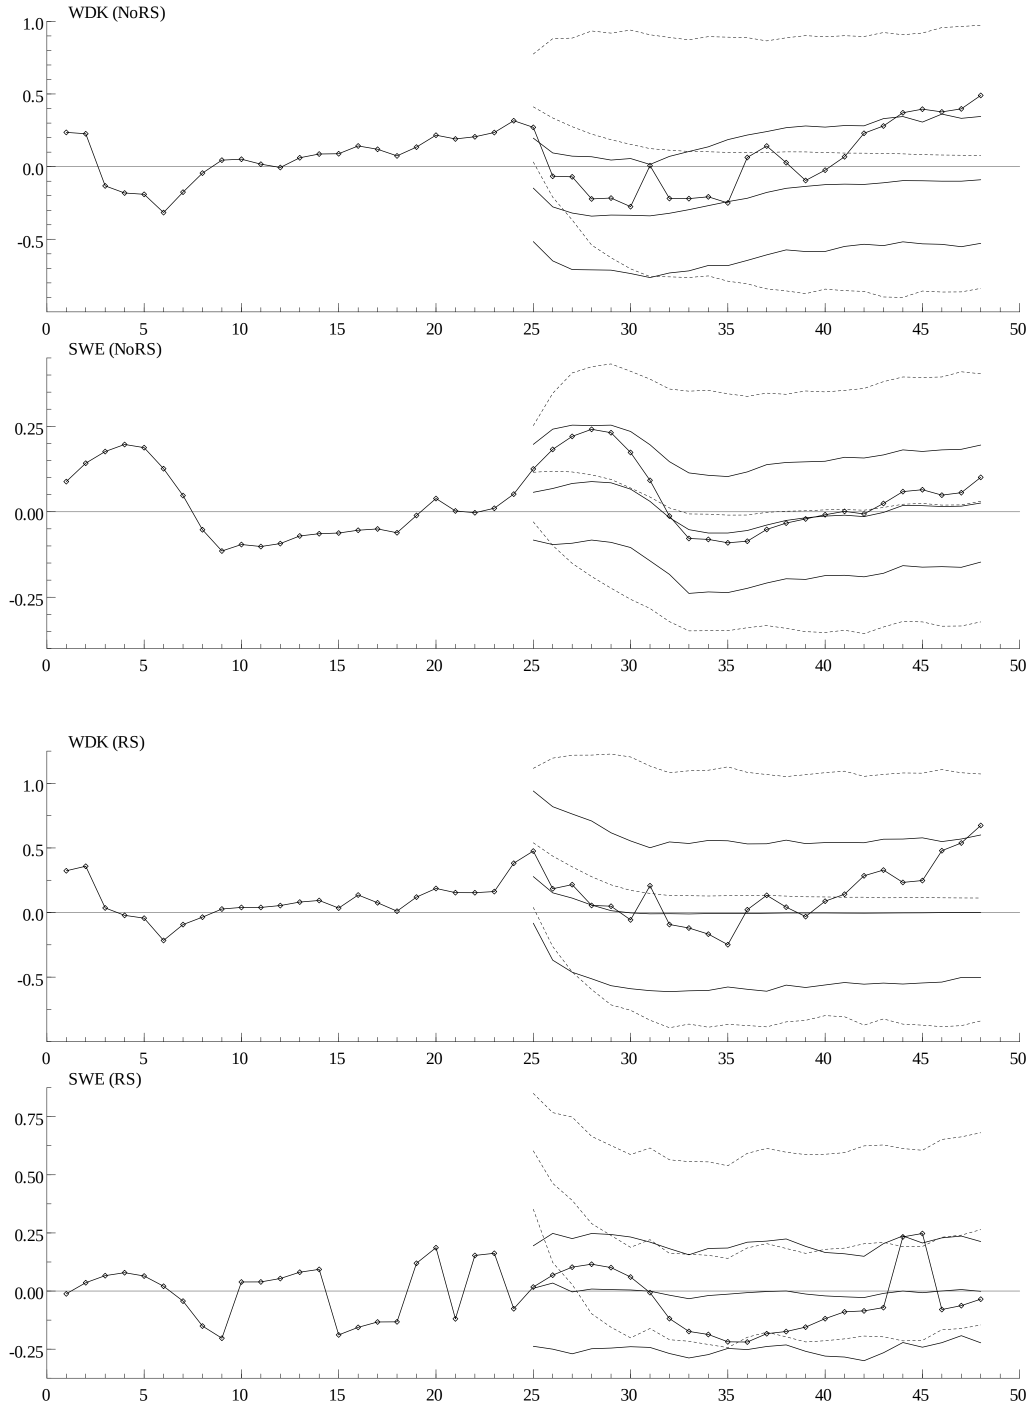

Note: In each panel, the solid line with the diamonds are the actually observed (deseasonalized) price series covering the last 24 in-sample observations as well as the 24 out-of-sample observations. Each panel also has a three solid lines and three dotted lines. The three solid and dotted lines are median forecasts and error bands for the VAR model and univariate model, respectively.

Figure 8: Forecasts for the WDK-SNO physical link for the non-switching (NoRS) and switching (RS) models

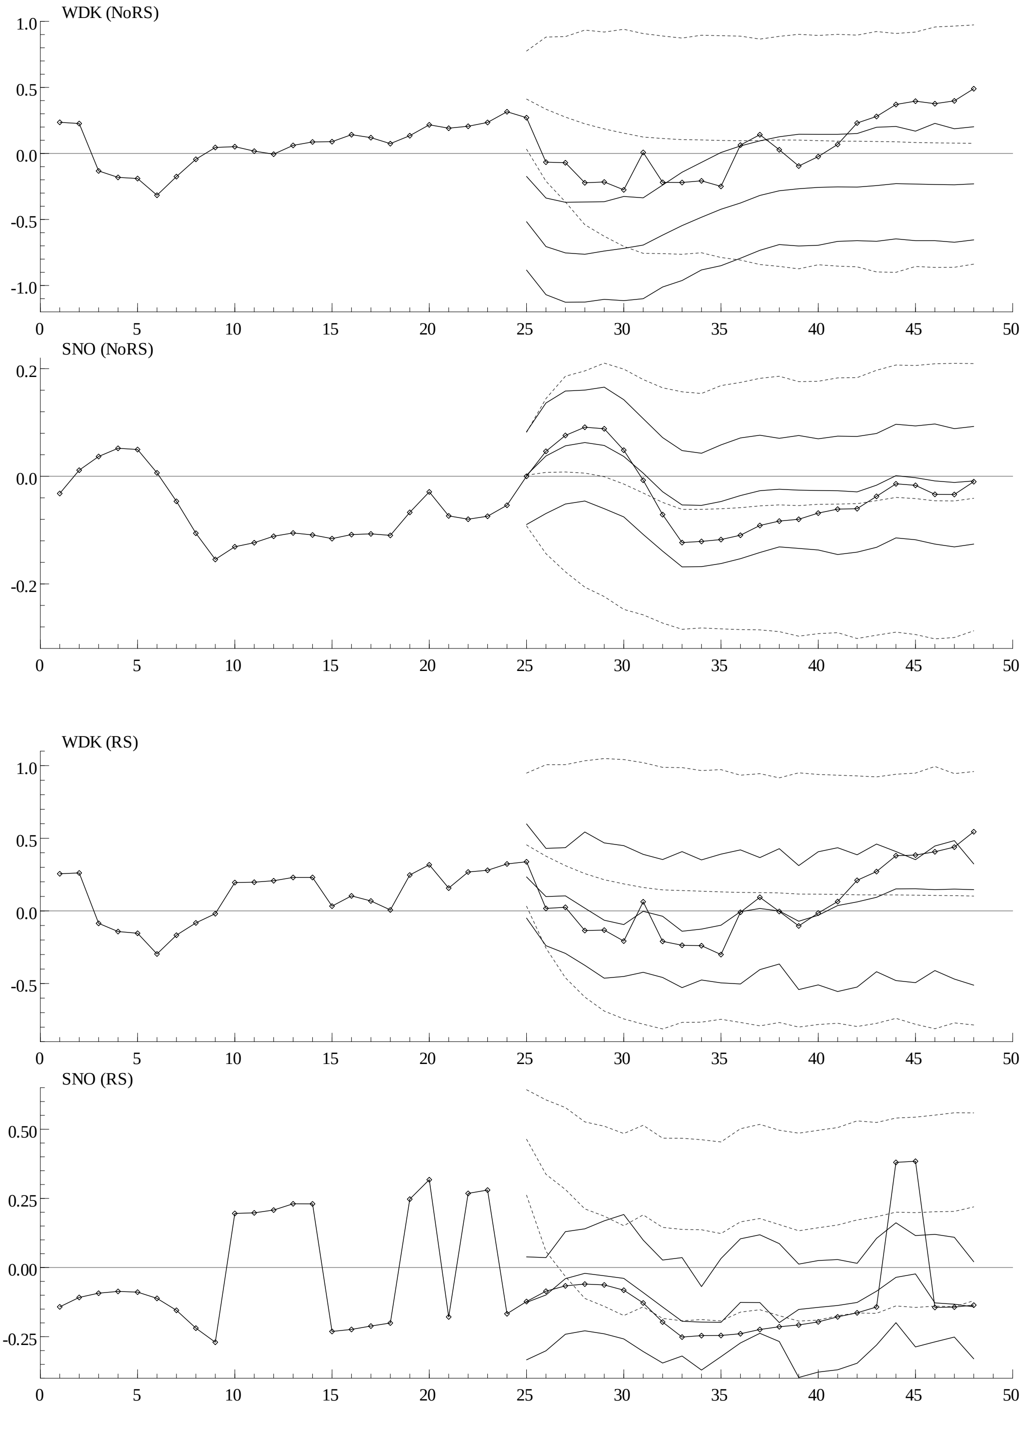

Note: In each panel, the solid line with the diamonds are the actually observed (deseasonalized) price series covering the last 24 in-sample observations as well as the 24 out-of-sample observations. Each panel also has a three solid lines and three dotted lines. The three solid and dotted lines are median forecasts and error bands for the VAR model and univariate model, respectively.

Figure 9: Forecasts for the SNO-SWE physical link for the non-switching (NoRS) and switching (RS) models

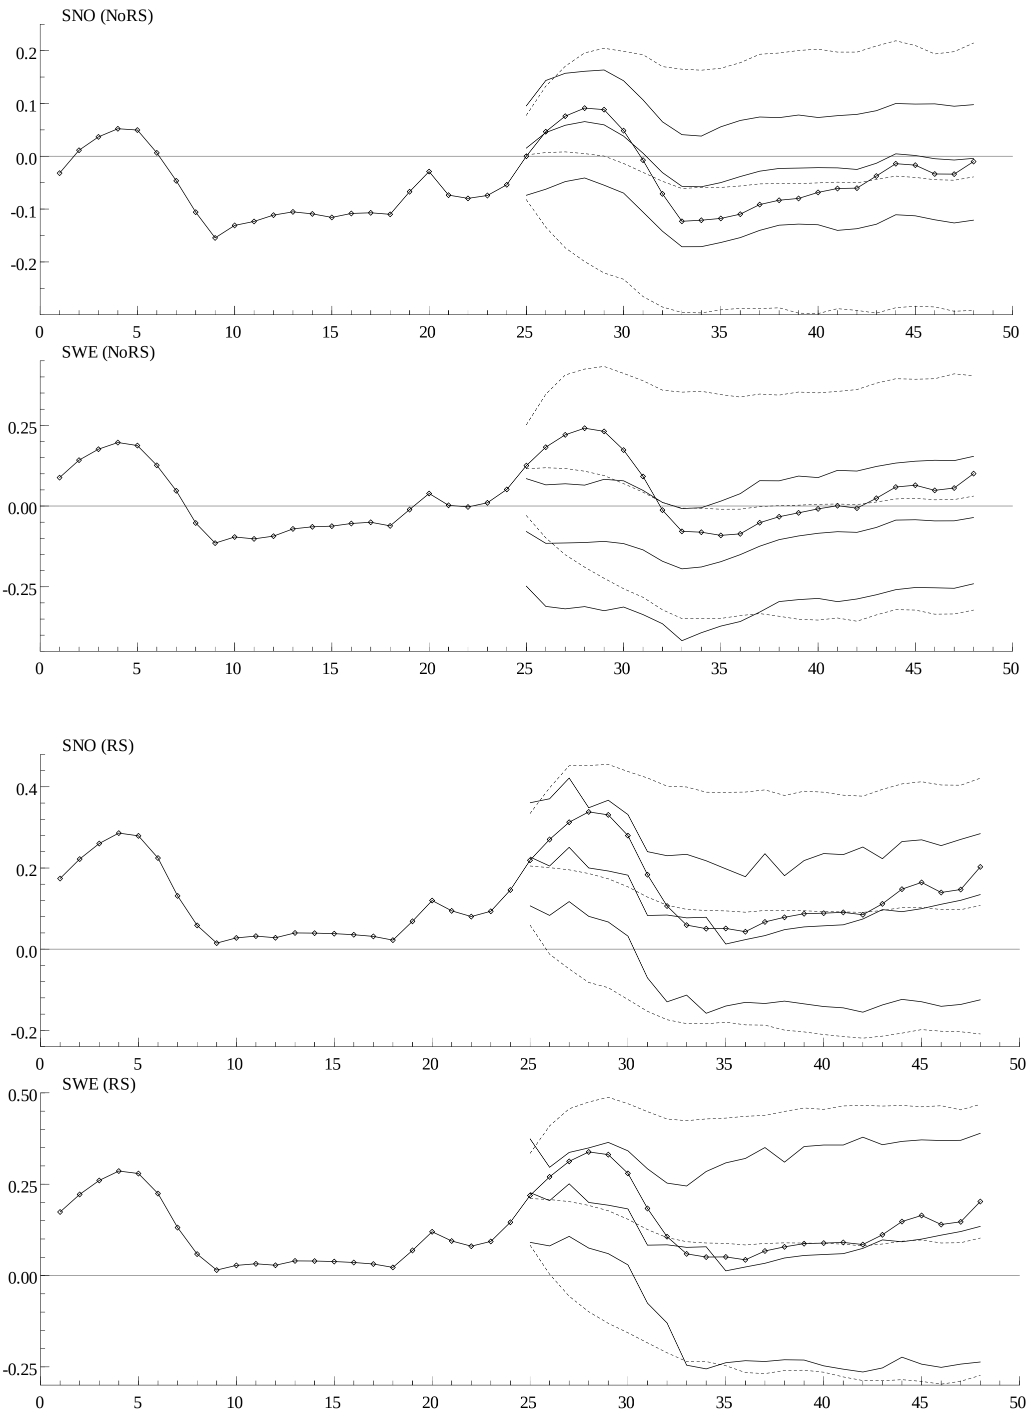

Note: In each panel, the solid line with the diamonds are the actually observed (deseasonalized) price series covering the last 24 in-sample observations as well as the 24 out-of-sample observations. Each panel also has a three solid lines and three dotted lines. The three solid and dotted lines are median forecasts and error bands for the VAR model and univariate model, respectively.

Figure 10: Forecasts for the SWE-FIN physical link for the non-switching (NoRS) and switching (RS) models

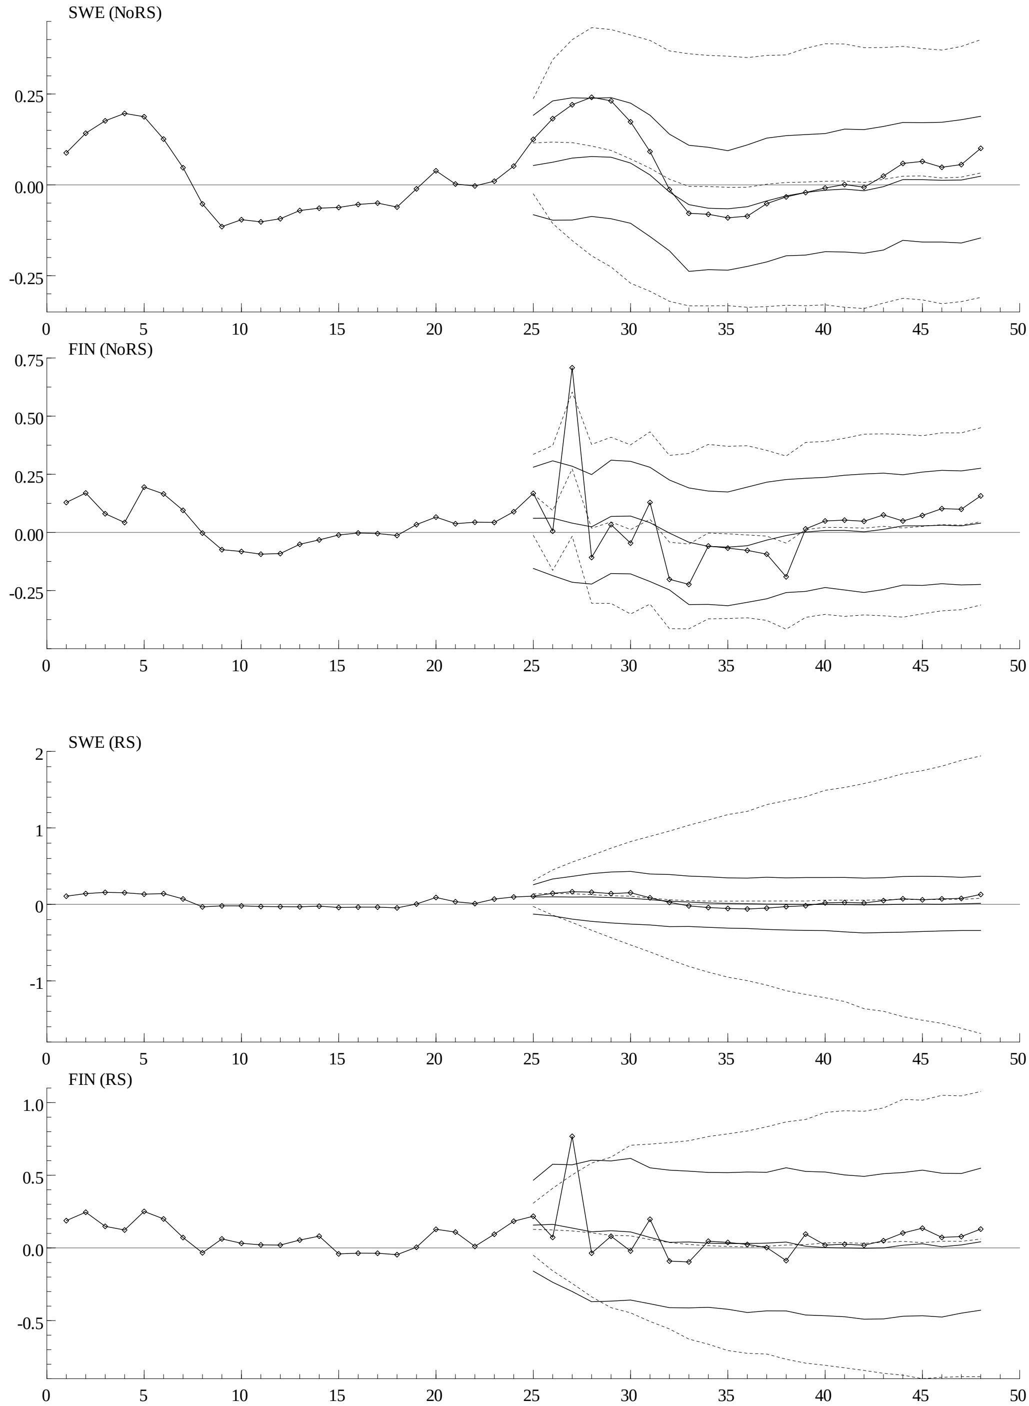

Note: In each panel, the solid line with the diamonds are the actually observed (deseasonalized) price series covering the last 24 in-sample observations as well as the 24 out-of-sample observations. Each panel also has a three solid lines and three dotted lines. The three solid and dotted lines are median forecasts and error bands for the VAR model and univariate model, respectively.