GÖTEBORGS UNIVERSITET
HANDELSHÖGSKOLAN

WORKING PAPERS IN ECONOMICS

No 327

Market Structure and the Stability and Volatility of Electricity Prices

Mikael Bask and Anna Widerberg

November 2008

ISSN 1403-2473 (print)
ISSN 1403-2465 (online)

Institutionen för nationalekonomi med statistik
Handelshögskolan vid Göteborgs universitet
Vasagatan 1, Box 640, SE 405 30 Göteborg
031 786 0000, 031 786 1326 (fax)
www.handels.gu.se   info@handels.gu.se

Market Structure and the Stability and Volatility of Electricity Prices*

Mikael Bask†
Department of Economics, Stanford University‡
Landau Economics Building, 579 Serra Mall
Stanford, CA 94305-6072, U.S.A.
E-mail: mbask@stanford.edu

Anna Widerberg
Department of Economics, Göteborg University
P.O. Box 640, SE-405 30 Göteborg, Sweden
E-mail: anna.widerberg@economics.gu.se

November 4, 2008

Abstract

By using a novel approach in this paper, (λ, σ^2)-analysis, we have found that electricity prices most of the time have increased in stability and decreased in volatility when the Nordic power market has expanded and the degree of competition has increased. That electricity prices at Nord Pool have been generated by a stochastic dynamic system that most often has become more stable during the step-wise integration of the Nordic power market means that this market is less sensitive to shocks after the integration process than it was before this process. This is good news.

JEL codes: C14; C22; D43.

Keywords: Electricity Prices; GARCH Models; Lyapunov Exponents; Market Structure; Reconstructed Dynamics; Stability; Volatility.

---

* This paper has benefitted from presentations at various conferences and seminars as well as from comments and suggestions by the editor and two unknown referees. The usual disclaimer applies.

† Corresponding author.

‡ Address during the academic year 2008/09.

1 Introduction

During the 1990s, the Nordic countries have gone through an extensive deregulation of their electricity sectors and, at the same time, there has been an evolution from national markets to a multi-national electricity market. Thus, Denmark, Finland, Norway and Sweden have all reformed their electricity sectors and have today access to a common electricity market, which consists of two parts: (i) bilateral trade of contracts between operators; and (ii) the non-mandatory power exchange, Nord Pool.

The step-wise integration of the Nordic power market gives the opportunity to investigate several interesting questions that relate to the relationship between market structure and the behavior of electricity prices. First, because the aim of the deregulation and integration of the electricity sectors was to develop a competitive power market, which would benefit the consumers in the Nordic countries, one might ask whether the degree of competition has increased over time during this process? Bask et al [4] examine this question and also find that this, in fact, is the case.

What about the behavior of electricity prices? Has the increased competition at the Nordic power market affected the volatility of electricity prices in a systematic way? This issue is under scrutiny in this paper and using a data set that is similar to the one in Bask et al [4], we find that the volatility of electricity prices most often has decreased when the Nordic power market has expanded and the degree of competition has increased (see also Lundgren et al [26]). This finding is also consistent with the prediction of McLaren [27] who presents an oligopoly model for a storable commodity such as water, which to a large extent is used in electricity generation in the Nord Pool area.

The main contribution in this paper is that we take a step further trying to figure out the reason for this change in volatility of electricity prices. Of course, the natural candidate is that the dynamics that govern the evolution of electricity prices have changed, but it could also be the case that the nature of the shocks hitting the Nordic power market has changed. A problem, however, is that we do not know the dynamics that govern the evolution of electricity prices, but as will be clear below (especially in Section 3), this problem is only apparent since we can reconstruct the dynamics using only electricity prices and, thereafter, measure the stability of the reconstructed dynamics.

Concretely, what we do in this paper is to present (λ, σ^2)-analysis, which is a method to contrast the stability of a stochastic dynamic system (λ) with the volatility of a variable generated by this system (σ^2). Thus, if we focus on the development of Nord Pool, the method contrasts the stability of the dynamics that govern the evolution of electricity prices with the volatility of these prices. What we find is that the dynamic system generating electricity prices most often has increased in stability during the step-wise integration of the Nordic power market.

There are two dimensions of the findings in this paper. The first is theoretical. To the extent that approximating models of the energy sector are developed in an attempt to understand the behavior of energy prices and how they depend on the market structure, the stability of the dynamics that govern the evolution of energy prices should be added as an important dimension in which model and data should be matched. That is, our finding that electricity prices most of the time have increased in stability and decreased in volatility when the Nordic power market has expanded and the degree of competition has increased should be properties of a well-formulated model of the Nordic power market.

The second dimension is welfare. That electricity prices at Nord Pool most often have become more stable during the integration process means that the Nordic power market is less sensitive to shocks after this process than it was before it. This, however, does not mean that electricity prices never can change a lot from one day to the other, but it does mean that such large movements in prices are expected to happen less frequently. This is good news.

This is how the rest of this paper is organized: In Section 2, we first review the literature on the relationship between market structure and the price volatility for a storable commodity. Thereafter, we argue at an intuitive level why (λ, σ^2)-analysis is useful as a tool to contrast the stability of a stochastic dynamic system with the volatility of a variable generated by this system. In Section 3, we formally outline (λ, σ^2)-analysis and, in Section 4, we use this tool to investigate how the step-wise integration of the Nordic power market has affected the stability and volatility of electricity prices. Section 5 concludes the paper with a discussion.

2 Market structure and stability of prices

The aim of this section is to give a careful introduction to the subject of this paper, namely, the relationship between market structure and the stability and volatility of electricity prices. Therefore, in the first subsection, we review the literature on the relationship between market structure and the price volatility for a storable commodity. We do this to have a better idea of how the price volatility might be affected at a market with increased competition.

The second subsection is devoted to the relationship between the stability and volatility of prices or, more correctly, the relationship between the stability of a dynamic system and the volatility of prices generated by this system. We spend some time on this relationship since the techniques that we use to measure stability are not so well-known. Thus, what we do is to the give some intuition behind (λ, σ^2)-analysis that we formally outline in Section 3.

Market structure and volatility of prices Since roughly one half of the electricity generated in the Nord Pool area is produced by hydro plants, it is natural to examine what the literature on storable commodities can offer regarding the relationship between market structure and the volatility of commodity prices. However, since most studies within this literature assume perfect competition (see Deaton and Laroque [8]-[9] and Scheinkman and Schechtman [33]), this relationship is far from fully understood.

One exception is McLaren [27] who develops a model in which the price of the commodity decreases in volatility when the degree of competition increases. Specifically, McLaren [27] presents a model of oligopolistic commodity speculation with an entry barrier in which the speculators perform non-cooperative storage in an infinite-horizon game. The author takes explicit care of the non-negative commodity stock constraint in the analysis and derives a Markov perfect equilibrium of the model in closed form.

McLaren [27] emphasizes three properties of his model. First of all, there is less storage and more volatile prices than when the commodity market is competitive. Second, the oligopolistic equilibrium converges to the competitive equilibrium when the number of speculators becomes large. Third, an increase in the demand for the commodity lowers storage and increases the volatility of the commodity price.

Another paper is by Thille [37] who, among other things, investigates the effect of storage on the price volatility under alternative market structures. Specifically, Thille [37] presents a model of a duopoly in which the firms perform non-cooperative storage in an infinite-horizon game. A difference in this model compared to McLaren's [27] model is that production no longer is exogenous. Thus, the firms in Thille's [37] model have market power over both production and storage, and not only over storage as in McLaren's [27] model. In the latter model, production is modelled as random harvests.

Using numerical analysis, Thille [37] finds that the relative importance of demand and cost shocks determines whether the price volatility is higher or lower under imperfect competition than under perfect competition. The main findings are that when demand shocks dominate, the price volatility decreases when the degree of competition increases whereas when cost shocks dominate, the price volatility increases when the degree of competition increases. Thille [37] looks at three market structures in the analysis: monopoly, duopoly and perfect competition.

An early paper that examines the relationship between the degree of competition and the price volatility for a storable commodity, which also should be mentioned, is Newbery [29]. In his random harvests model, the monopolist stores more than firms do in a competitive market, which results in a less volatile commodity price when the commodity demand is linear in the price. However, when the price elasticity of demand is constant, the commodity price is less volatile under competition. Thus, the findings are ambiguous as in Thille [37]. Since Newbery's [29] model is a random harvests model, the firms have market power over storage only as in McLaren's [27] model.

To summarize, the small theoretical literature that exists on the relationship between the degree of competition and the price volatility for a storable commodity does not give a unified answer. This should not come as a surprise since the models differ from each other in several respects (see Rui and Miranda [31], Vedenov and Miranda [40], and Williams and Wright [41] for more literature on this topic). But what about the empirical literature? What does it have to say on the relationship between market structure and the price volatility for a storable commodity?

First of all, there are few empirical studies that measure the degree of market power in the Nord Pool area (see Hjalmarsson [22] and Vassilopoulos [39]) and almost none of them studies the effect of market expansion on the degree of market power. One exception, however, is Bask et al [4] who investigate how the degree of market power has evolved during the integration process at the

Nordic power market using a conjectural variation method and they find that the degree of competition has increased when the common market has expanded. Unfortunately, we are not aware of any empirical study on the effects of market expansion on the degree of market power for some other power market than Nord Pool.

Then, what about the relationship between the degree of competition and the volatility of electricity prices at Nord Pool? Lundgren et al [26] take a close look at how the price dynamics have evolved during the integration process at the Nordic power market by estimating not only the price volatility, but also the intensity and size of the price jumps. To be able to accomplish their analysis, they divide their data set into three parts: (i) when only Norway and Sweden participate in Nord Pool; (ii) when also Finland has joined the common market; and (iii) when Denmark too participates in Nord Pool.

If we turn to their results, Lundgren et al [26] find that the volatility of electricity prices increased when Finland joined the Nordic power market, but that this increase mainly was driven by an increased size of the price jumps. However, the volatility of electricity prices decreased when Denmark joined the Nordic power market. In other words, Lundgren et al [26] do not find a clear-cut relationship between the degree of competition and the volatility of electricity prices. Anyhow, when Finland and Denmark joined the common market, the intensity of the price jumps decreased.

What we do in this paper (see Section 4 below) is that we fit a battery of volatility models to the time series in a data set that is similar to the one in Bask et al [4]. This means that we are able to draw conclusions regarding the degree of competition at Nord Pool and how this degree is related to the volatility of electricity prices. What we find is that the volatility of electricity prices most often has decreased when the Nordic power market has expanded and the degree of competition has increased.

Stability and volatility of prices The main contribution in this paper is that we not only distinguish between the stability and volatility of electricity prices, but that we also present a method that allows us to measure the stability of electricity prices. As will be clear below, more stable electricity prices is not synonymous for less volatile electricity prices. Instead, more stable electricity prices is a sloppy expression for a more stable stochastic dynamic system generating electricity prices.

In fact, what we argue in this paper is that one should contrast the stability of a stochastic dynamic system with the volatility of a variable generated by this system using what we call (λ, σ^2)-analysis. For example, think of macroeconomic models with the purpose of explaining aggregate features of an economy. At the heart of these models (that nowadays are very popular in the literature), there is an impulse-propagation mechanism in which impulses are shocks to the economy, while the propagation mechanism is the means by which these shocks lead to persistence over time. An economy that is less stable is, therefore, associated with a higher persistence of the shocks, meaning that even a small shock to the economy would have a large effect on it.

To further clarify the fundamental idea behind (λ, σ^2)-analysis, let σ^2 denote the conditional variance of a variable generated by a stochastic dynamic system and let λ denote the stability of this system. Then,

$$\sigma^2 = \sigma^2(\lambda, \varepsilon), \tag{1}$$

where ε is exogenous shocks to the dynamic system, meaning that the conditional variance (σ^2) is not only affected by the system's stability (λ), it is also affected by the shocks to the system (ε). Specifically, the conditional variance of a variable increases when the dynamic system decreases in stability, but also when the shocks' amplitude increases in size.

However, because of ε in (1), there is no one-to-one correspondence between

σ^2 and λ, which also motivates $\left(\lambda, \sigma^2\right)$-analysis. To better see this claim, imagine a very simple model of how the electricity price, s, is determined at the power market:

$$s_t = \alpha s_{t-1} + \varepsilon_t, \tag{2}$$

where ε is independently and identically distributed shocks with zero mean and finite variance. That is, the model in (2) is a linear autoregression of order one. If we focus on the deterministic part of the model, it is explosive when $|\alpha| > 1$, conservative when $|\alpha| = 1$ and stable when $|\alpha| < 1$. Further, if we assume that the model in (2) is stable, we can decompose the volatility of electricity prices as follows (see Bask and de Luna [1]):

$$\sigma^2 \equiv var\left(s_t\right) = \frac{var\left(\varepsilon_t\right)}{1-\alpha^2} \equiv \frac{var\left(\varepsilon_t\right)}{1-\left(\exp\left(\lambda\right)\right)^2}. \tag{3}$$

Thus, electricity prices are more volatile (σ^2) when the shocks are more volatile (ε), but also when the model is less stable (λ). Be aware that for a stable model, $\lambda < 0$ since $\lambda \equiv \log|\alpha|$, meaning that λ approaches zero from below when the model in (2) decreases in stability.

The stability concept outlined above can in a straightforward way be generalized to linear autoregressions of higher orders than one (see, among others, Bask and de Luna [1] for a careful demonstration and discussion of this claim). If the model is a linear autoregression of order n, one can characterize the stability of the model by looking at the modulus of the n eigenvalues that solve the eigenvalue problem for the model and this is because the eigenvectors of the model are orthogonal by definition.

First, if the autoregression is stable, the contraction of it in n-dimensional space is dominated by the largest eigenvalue. This also means that the largest eigenvalue can be used when comparing the stability of two linear autoregressions of order n. This, however, is not the only way to compare the stability of two autoregressions. An alternative is to look at the product of the modulus of the n eigenvalues since it describes the rate of contraction of an n-dimensional volume in n-dimensional space. That is, the autoregression with the faster contraction is more stable.

But we do not want to restrict the stability analysis to linear systems such as linear autoregressions. Instead, we would like to have a tool that is able to determine the stability of non-linear dynamic systems. Unfortunately, this also means that we cannot use the eigenvalues in the stability analysis since they are defined over linear systems. Fortunately, Bask and de Luna [1] argue that the Lyapunov exponents can be used in the stability analysis of non-linear dynamic systems and this is because the Lyapunov exponents for a non-linear system are the non-linear counterpart of the eigenvalues for a linear system.

We save the definition, motivation, estimation and inference of the Lyapunov exponents to Section 3, but a few paragraphs about the use of the Lyapunov exponents in the empirical literature should be spent already here. In the early 1980s, a stream of papers started to emerge on the search for chaotic dynamics in time series data. In the beginning, data were collected from different kinds of systems in the sciences, but later on, time series data in economics and finance were used when searching for chaotic dynamics (see Urbach [38] for an extensive although not complete list of papers within this literature).

Two reasons made this stream of papers possible. First, the dynamics that govern the evolution of a variable must be reconstructed since the dynamics are not known. In 1981, Takens [36] proved that it is possible to reconstruct the unknown dynamics using only a scalar time series. Remarkably, the reconstructed dynamics and the unknown dynamics hidden in the "black box" are equivalent in the sense that the reconstruction preserves topological information about the dynamic system such as the Lyapunov exponents. (Unfortunately, even though the reconstruction theorem by Takens [36] have made an impact on the sciences, it is still rather unknown in economics and finance.)

However, more was needed than the reconstruction theorem by Takens [36]. Specifically, a method to calculate the Lyapunov exponents from a scalar time series was necessary and the first such method was presented in 1985 by Wolf et al [42]. Thereafter, several methods have been proposed in the literature. Some of the methods only calculate the largest Lyapunov exponent whereas other methods calculate the whole spectrum of Lyapunov exponents. We will use one of the latter methods in this paper.

To give a general picture of the findings in the search for chaotic dynamics in time series data, these findings have been mixed. To be more precise and if we restrict attention to the findings in economics and finance, most if not all of the early studies demonstrated that several time series could be characterized by chaotic dynamics and that this meant that the predictive ability of these systems were strongly limited. It could be mentioned in this context that Bask et al [3] test whether electricity prices from Nord Pool has been characterized by chaotic dynamics, which also is detected in one of the time series in the data set (which is the same data set as in this paper).

However, various limitations in these studies were later on revealed (that, in fact, did not apply in Bask et al [3]) and one such limitation was the lack of statistical inference. In other words, that a distributional theory for the Lyapunov exponents was lacking. Further on, it turned out that some of the methods to calculate the Lyapunov exponents suffered from upward bias. Consequently, estimates of the largest Lyapunov exponent that had been reported to be positive in the literature, which is the operational definition of chaotic dynamics, could in fact be negative. Both limitations, especially the lack of a distributional theory for the Lyapunov exponents, had the consequence that the number of new studies within this field diminished over time.

Fortunately, Shintani and Linton [35] provide a distributional theory for the

Lyapunov exponents. Their framework together with the algorithms in Gencay and Dechert [17] and Kuan and Liu [25] to calculate the Lyapunov exponents mean that we have an asymptotically normal estimator of the Lyapunov exponents. This estimator is also used in Bask et al [3] as well as in this paper (see Section 4 below). However, we do not search for chaotic dynamics in electricity prices as in Bask et al [3]. Instead, we investigate whether they have changed in stability during the integration process at the Nordic power market.

As we already have pointed out, Bask and de Luna [1] argue that the Lyapunov exponents can be used in a stability analysis of a non-linear dynamic system. To be more precise, in the same way as the product of the modulus of the n eigenvalues for an n-dimensional linear system describes the rate of contraction of an n-dimensional volume, the average of the n Lyapunov exponents for an n-dimensional non-linear system describes the rate of contraction of an n-dimensional volume. The reason for this is that the Lyapunov exponents for a linear system are the logarithms of this system's eigenvalues.

What we find in this paper (see Section 4 below) is that the stability of electricity prices most often has increased when the Nordic power market has expanded and the degree of competition has increased. Thus, having the Nord Pool experience in focus, increased competition at the power market has been associated with an increasingly stable dynamic system generating decreasingly volatile electricity prices but without having a one-to-one correspondence.

3 Method: (λ, σ^2)-analysis

(λ, σ^2)-analysis is outlined in this section. Specifically, in the first four subsections, we carefully explain how our stability measure (λ) is defined, how it can be estimated and tested from data, but we also motivate why this measure is useful in a stability analysis. The final subsection is about the volatility measure

(σ^2).

Definition of λ As we have pointed out several times, Bask and de Luna [1] argue that the (spectrum of smooth) Lyapunov exponents can be used in the determination of the stability of a stochastic dynamic system. Specifically, assume that the dynamic system, $f : \mathbb{R}^n \to \mathbb{R}^n$, generates

$$S_{t+1} = f\left(S_t\right) + \varepsilon^s_{t+1}, \tag{4}$$

where S_t and ε^s_t are the state of the system and a shock to the system, respectively. For an n-dimensional system as in (4), there are n Lyapunov exponents that are ranked from the largest to the smallest exponent:

$$\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n, \tag{5}$$

and it is these quantities that provide information on the stability properties of the dynamic system f.

Assume temporarily that there are no shocks and consider how the dynamic system f amplifies a small difference between the initial states S_0 and S_0':

$$S_j - S_j' = f^j\left(S_0\right) - f^j\left(S_0'\right) \simeq Df^j\left(S_0\right)\left(S_0 - S_0'\right), \tag{6}$$

where $f^j\left(S_0\right) = f\left(\cdots f\left(f\left(S_0\right)\right)\cdots\right)$ denotes j successive iterations of the system starting at state S_0 and Df is the Jacobian of the system:

$$Df^j\left(S_0\right) = Df\left(S_{j-1}\right) Df\left(S_{j-2}\right) \cdots Df\left(S_0\right). \tag{7}$$

Then, associated with each Lyapunov exponent, λ_i, $i \in [1, 2, \ldots, n]$, there are nested subspaces $U^i \subset \mathbb{R}^n$ of dimension $n + 1 - i$ with the property that

$$\lambda_i \equiv \lim_{j\to\infty} \frac{\log_e \left\| Df^j\left(S_0\right)\right\|}{j} = \lim_{j\to\infty} \frac{1}{j} \sum_{k=0}^{j-1} \log_e \left\| Df\left(S_k\right)\right\|, \tag{8}$$

for all $S_0 \in U^i - U^{i+1}$. Due to Oseledec's multiplicative ergodic theorem, the limits in (8) exist and are independent of S_0 almost surely with respect to the measure induced by the process $\{S_t\}_{t=1}^{\infty}$ (see Guckenheimer and Holmes [20] for a careful definition of the Lyapunov exponents). Then, allow for shocks to the system, meaning that the aforementioned measure is induced by a stochastic process. The Lyapunov exponents have in this case been renamed to smooth Lyapunov exponents in the literature.

Motivation of λ The reason why the spectrum of smooth Lyapunov exponents provides information on the stability properties of a stochastic dynamic system may be seen by considering two starting values of a system, where the only difference is an exogenous shock at time $t = 0$ and that the shocks are identical at times $t > 0$ (see Bask and de Luna [1] for the origin of this discussion).

The largest smooth Lyapunov exponent, λ_1, measures the slowest exponential rate of convergence of two trajectories of the dynamic system starting at the two starting values at time $t = 0$. In fact, λ_1 measures the convergence of a shock in the direction defined by the eigenvector corresponding to this exponent. If the difference between the two starting values lies in another direction of $\mathbb{R}^n$, then the convergence is faster. Thus, λ_1 measures a "worst case scenario." (When $\lambda_1 > 0$, the trajectories diverge from each other and for a bounded stochastic dynamic system, this is an operational definition of chaotic dynamics.)

The average of the smooth Lyapunov exponents,

$$\lambda \equiv \frac{1}{n} \sum_{i=1}^{n} \lambda_i, \tag{9}$$

measures the exponential rate of convergence in a geometrical average direction. That is, the convergence of two trajectories of the dynamic system in the geometrical average of the directions defined by the eigenvectors corresponding to the different exponents. Thus, λ measures an "average scenario." We can, therefore, compare the stability of two stochastic dynamic systems via the smooth Lyapunov exponents since a shock has a smaller effect on the dynamic system with a smaller λ than for the system with a larger λ. Since we are dealing with dissipative systems, meaning that $\lambda < 0$ by definition, a dynamic system is more stable than another system if λ is more negative.

Potter [30] too claims that λ_1 can be used not only to categorize a time series as stable or unstable (in the sense it is chaotic), but also that it provides a measure of the convergence speed to (some sort of) equilibrium (see Shintani [34] who use λ_1 when looking at convergence speeds of exchange rates toward purchasing power parity).

Estimation of λ Since the actual form of the stochastic dynamic system f is not known, it may seem like an impossible task to determine the stability of the system. Fortunately, it is possible to reconstruct the dynamics using only a scalar time series and, thereafter, measure the stability of the reconstructed system. Therefore, associate the dynamic system f with an observer function, $g : \mathbb{R}^n \to \mathbb{R}$, that generates the following scalar time series:

$$s_t = g\left(S_t\right) + \varepsilon_t^m, \tag{10}$$

where $s_t \in S_t$ and ε_t^m are an observation in the time series and a measurement error, respectively. That is, the time series $\{s_t\}_{t=1}^N$ is observed, where N is the number of observations.

Specifically, the observations in a scalar time series contain information regarding unobserved state variables that can be used to define a state in present time. Therefore, let

$$T = \left(T_1, T_2, \ldots, T_M\right)' \tag{11}$$

be the reconstructed trajectory, where T_t is the reconstructed state and M is the number of states on the trajectory. Each T_t is given by

$$T_t = \{s_t, s_{t+1}, \ldots, s_{t+m-1}\}, \tag{12}$$

where m is the embedding dimension. Thus, T is an $M \times m$ matrix and the constants M, m and N are related as $M = N - m + 1$.

Takens [36] proved that the map

$$\Phi(S_t) = \left\{ g\left(f^0(S_t)\right), g\left(f^1(S_t)\right), \ldots, g\left(f^{m-1}(S_t)\right) \right\}, \tag{13}$$

which maps the n-dimensional state S_t onto the m-dimensional state T_t, is an embedding if $m > 2n$ (see below for the intuition behind this condition). This means that the map is a smooth map that performs a one-to-one coordinate transformation and has a smooth inverse. Moreover, a map that is an embedding preserves topological information about the unknown dynamic system such as the smooth Lyapunov exponents and, in particular, the map induces a function, $h : \mathbb{R}^m \rightarrow \mathbb{R}^m$, on the reconstructed trajectory,

$$T_{t+1} = h(T_t), \tag{14}$$

which is topologically conjugate to the unknown dynamic system f. That is,

$$h^j(T_t) = \Phi \circ f^j \circ \Phi^{-1}(T_t). \tag{15}$$

Thus, h is a reconstructed dynamic system that has the same smooth Lyapunov exponents as the unknown dynamic system f. (However, since the m-dimensional system h has a larger dimension than the n-dimensional system f, the number of smooth Lyapunov exponents that are spurious is $m - n$. This issue is discussed in detail in Dechert and Gencay [10]-[11] and Gencay and Dechert [18].)

Then, to be able to estimate the smooth Lyapunov exponents, one has to estimate h. However, since

$$h : \begin{pmatrix} s_t \\ s_{t+1} \\ \vdots \\ s_{t+m-1} \end{pmatrix} \longrightarrow \begin{pmatrix} s_{t+1} \\ s_{t+2} \\ \vdots \\ v\left(s_t, s_{t+1}, \ldots, s_{t+m-1}\right) \end{pmatrix}, \tag{16}$$

the estimation of h reduces to the estimation of v:

$$s_{t+m} = v\left(s_t, s_{t+1}, \ldots, s_{t+m-1}\right). \tag{17}$$

Moreover, since the Jacobian of h at the reconstructed state T_t is

$$Dh\left(T_t\right) = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ \frac{\partial v}{\partial s_t} & \frac{\partial v}{\partial s_{t+1}} & \frac{\partial v}{\partial s_{t+2}} & \cdots & \frac{\partial v}{\partial s_{t+m-1}} \end{pmatrix}, \tag{18}$$

a feedforward neural network is recommended to estimate the above derivatives to derive the smooth Lyapunov exponents and this is because Hornik et al [23] have shown that a map and its derivatives of any unknown functional form can be approximated arbitrarily accurately by such a network.

Having shown how to estimate λ, an intuitive explanation of Takens' [36] embedding theorem (or reconstruction theorem as we called it in Section 2) is in place. For the sake of the argument, assume that $\mathbf{M}_1 \subset \mathbf{M}$ and $\mathbf{M}_2 \subset \mathbf{M}$ are two subspaces of dimension n_1 and n_2, respectively, where $\mathbf{M} \in \mathbb{R}^m$ is an m-dimensional manifold representing phase space for the reconstructed dynamic system. In general, two subspaces intersect in a subspace of dimension $n_1 + n_2 - m$, meaning that when this expression is negative, there is no intersection of the two subspaces. Therefore, and of greater interest, the self-intersection of an n-dimensional manifold with itself fails to occur when $m > 2n$ (see Sauer et al [32] for generalizations of Takens' [36] embedding theorem).

A problem in this context is that the dimension of the "true" dynamic system is not known, meaning that the required embedding dimension according to Takens' [36] embedding theorem is not either known. However, this problem can be solved indirectly by making use of a generic property of a proper reconstruction, namely, that the dynamics in original phase space must be completely unfolded in reconstructed phase space. In other words, if the embedding dimension is too low, the dynamics are not completely unfolded, meaning that distant states in original phase space are close states in reconstructed phase space and are, therefore, false neighbors in phase space.

There are at least two methods to calculate the required embedding dimension from a scalar time series: (i) false nearest neighbors; and (ii) the saturation of invariants on the reconstructed dynamics such as the saturation of the Lyapunov exponents. The first method is based on the aforementioned property of a proper reconstruction, meaning that by increasing the embedding dimension enough, the dynamics are completely unfolded when there are no false neighbors in reconstructed phase space (see Kennel et al [24]).

The second method, the saturation of invariants on the reconstructed dynamics, is based on the fact that when the dynamics are completely unfolded, the Lyapunov exponents and other invariants (such as entropy and fractal dimension) are independent of the embedding dimension. If, however, the dynamics are not completely unfolded in reconstructed phase space, these invariants depend on the embedding dimension. Therefore, by increasing the embedding dimension enough, the dynamics are completely unfolded when an invariant stops changing (see Fernández-Rodríguez et al [15] for an example regarding the largest Lyapunov exponent and a statistical test for chaotic dynamics).

Inference of λ First of all, we are not aware of any distributional theory for λ. However, Shintani and Linton [35] show that a neural network estimator of the smooth Lyapunov exponents like the one above is asymptotically normal. Our conjecture is, therefore, that asymptotic normality also holds for a neural network estimator of λ since the eigenvectors corresponding to the smooth Lyapunov exponents are pairwise orthogonal.

An alternative to Shintani and Linton [35] is to derive an empirical distribution for λ via bootstrapping (see Bask and Gencay [2] and Gencay [16] for the case of λ_1). However, as is demonstrated in Ziehmann et al [43], this approach is not straightforward for multiplicative ergodic statistics since the limits in (8) may not exist under bootstrapping.

Measuring σ^2 We will not spend time here on how to measure the volatility of a variable generated by a stochastic dynamic system. The reason is that an extensive literature already exists on this topic. Instead, we will return to this issue when measuring the volatility of electricity prices (see Section 4 below).

4 Nord Pool experience

The Nord Pool experience is in focus in this section. In the first two subsections, we give a short review of the integration process at the Nordic power market as well as provide some descriptive statistics of our data set. In the two final subsections, we use (λ, σ^2)-analysis when examining how the step-wise integration of the Nordic power market has affected the stability and volatility of electricity prices.

Integration process Nord Pool is a multi-national power exchange, joining the Nordic countries. Norway was the first of the Nordic countries to deregulate their power market and Nord Pool was established in 1993, but then under the name Statnett Marked AS. Sweden started their deregulation process too in the first half of the 1990s and went step-wise to a deregulated power market. At the turn of 1996, the joint Norwegian-Swedish power exchange market, Nord Pool ASA, was launched.

Finland started a power exchange of its own, EL-EX, in 1996, but joined Nord Pool in 1997. Denmark Nord Pool Consulting was established in 1998 and western Denmark joined the common market in 1999. When eastern Denmark joined in 2000, the Nordic power market became fully integrated. For specific dates in the integration process, see Table 1.

Region	Date of entry
Norway	January 1, 1993
Sweden	January 1, 1996
Finland	December 29, 1997
Western Denmark	July 1, 1999
Eastern Denmark	October 1, 2000

Table 1: Dates in the integration process at the Nordic power market.

Data set used and descriptive statistics The data set used is spot electricity prices from Nord Pool. Specifically, it is the daily average of the hourly system price during the period January 1, 1993, to December 31, 2005. (The price is calculated from a bidding process before any bottlenecks are discovered. Around half of the time, the bidding process price is the price used in the entire region.) The data set is analyzed both as one time series but also split in parts with the natural breakpoints when a new region is joining the common market. The daily average of the hourly system price is presented in Figure 1a and basic descriptive statistics of this time series are presented in Table 2.

[Figure 1a about here.]

Statistic	Quantity
Number of observations	4745
Minimum	14.8
Maximum	831
Mean	174
Median	158
Variance	7310
Skewness	1.62
Kurtosis	7.53
Dickey-Fuller unit root test	-4.55^{***}

Table 2: Descriptive statistics of the daily average of the hourly system price during the period January 1, 1993, to December 31, 2005. *** Significant at the 1 per cent level (critical value: -3.44).

As is typical with financial time series data, the time series in Figure 1a does not look stationary and the autocorrelation function (ACF) in Figure 1b also indicates non-stationary, while the partial autocorrelation function (PCF) in Figure 1c is harder to analyze.

[Figure 1b about here.]

[Figure 1c about here.]

The Dickey-Fuller unit root test in Table 2 shows, however, that the time series qualifies as a stationary time series, but the test statistic is close to the critical value. A positive skewness indicates that the distribution of prices has an asymmetric distribution with a longer tail on the right side and a positive kurtosis indicates a distribution with fat tails.

Instead of using the daily average of the hourly system price in the analysis, we use the logarithmic return of this price. The new time series is presented in Figure 2a and basic descriptive statistics of it are presented in Table 3.

[Figure 2a about here.]

Statistic	Quantity
Number of observations	4744
Minimum	-0.774
Maximum	1.19
Mean	0.000189
Median	-0.00450
Variance	0.00936
Skewness	0.895
Kurtosis	19.0
Dickey-Fuller unit root test	-23.4^{***}

Table 3: Descriptive statistics of the logarithmic returns of the daily average of the hourly system price during the period January 1, 1993, to December 31, 2005. *** Significant at the 1 per cent level.

The shape of the time series together with the Dickey-Fuller unit root test in Table 3 and the ACF in Figure 2b indicate that the time series is stationary (see also the PCF in Figure 2c). Moreover, a positive skewness and a positive kurtosis indicate that the distribution of logarithmic returns of prices has an asymmetric distribution with a longer tail on the right side as well as fat tails.

[Figure 2b about here.]

[Figure 2c about here.]

We proceed the analysis by dividing the time series into five parts with the natural breakpoints when a new region is joining the common market. See the Appendix for descriptive statistics of each of the five time series.

Stability of electricity prices We estimate the smooth Lyapunov exponents for each time series using the algorithm proposed in Gencay and Dechert [17] and Kuan and Liu [25] using 4, 8, 12, 16 and 20 inputs to the neural network, respectively, where the number of hidden units runs from 1 unit to 20 units and, thereafter, we calculate our stability measure, λ. It is the λ that minimizes the

Schwarz Information Criterion for each time series that we report in Table 4. (We have used NETLE, a program developed by R. Gencay, C.-M. Kuan and T. Liu, when estimating the smooth Lyapunov exponents.)

Region	Stability
Norway (1/1/1993-12/31/1995)	−0.268 12 inputs 5 hidden units
	Increase
Norway and Sweden (1/1/1996-12/28/1997)	−0.359 8 inputs 2 hidden units
	Decrease
Norway, Sweden and Finland (12/29/1997-6/30/1999)	−0.168 12 inputs 3 hidden units
	Increase
Norway, Sweden, Finland and western Denmark (7/1/1999-9/30/2000)	−0.273 8 inputs 1 hidden unit
	Increase
Norway, Sweden, Finland and Denmark (10/1-2000-12/31/2005)	−0.275 8 inputs 5 hidden units

Table 4: Stability and change in stability of the logarithmic returns of the daily average of the hourly system price during the integration process at the Nordic power market.

The general picture is that the integration process at the Nordic power market most of the time has been associated with more stable electricity prices. However, we do not test for a change in stability since a distributional theory for λ is lacking.

Volatility of electricity prices Since there is no obvious choice of model to fit to our data set to be able to determine whether electricity prices have become more or less volatile during the integration process at the Nordic power market, we fit four time series models to the data set, where all of them belong to the generalized autoregressive conditional heteroskedastic (GARCH) model family (see Bollerslev [5] and Engle [13] for the seminal contributions in this area).

The first model that we fit to our data set is the GARCH (1,1) specification (see Enders [12] and Hamilton [21] for details on this and the other three specifications we use when measuring the volatility of a time series):

$$\sigma_t^2 = \beta_0 + \beta_1 \varepsilon_{t-1}^2 + \beta_2 \sigma_{t-1}^2, \tag{19}$$

where σ_t^2 is the conditional and time-varying variance, and ε_t^2 is the squared innovation with $\varepsilon_t = z_t \sigma_t$ and $z_t \sim IID\,(0,1)$. Our volatility measure is β_2.

However, we are not primarily interested in the volatility level, but, instead, to test whether there has been a change in volatility of electricity prices when the Nord Pool area has expanded in size with a new region. A straightforward way to perform such a test is to fit the following model to two adjacent time series:

$$\sigma_t^2 = \beta_0 + \beta_1 \varepsilon_{t-1}^2 + \beta_2 \sigma_{t-1}^2 + \beta_3 \sigma_{t-1}^2 I_{t \geq t_0}, \tag{20}$$

where $I_{t \geq t_0} = 1$ when $t \geq t_0$ and zero elsewhere, and t_0 is the time point when a new region is joining the common market.

To be more concrete by giving an example, what we do is to fit the specification in (19) to the time series when only Norway participates in what would become Nord Pool and find that $\beta_2 = 0.258$, which is the volatility of electricity prices during this period. Thereafter, we fit the specification in (19) to the time series when both Norway and Sweden participate in Nord Pool and find that $\beta_2 = 0.217$. Finally, to test if there has been a change in volatility of electricity prices when Sweden joined the common market, we fit the specification in (20) to the joint time series and find that the decrease in volatility is not significant at any conventional significance level. Our findings are reported in Table 5 (that also includes the findings for the other three GARCH specifications).

Region	Volatility GARCH	Volatility E-GARCH	Volatility I-GARCH	Volatility GJR-GARCH
Norway (1/1/1993-12/31/1995)	0.258	0.974	0.230	0.668
	Decrease	Decrease	Decrease	Decrease
Norway and Sweden (1/1/1996-12/28/1997)	0.217	0.935	0.220	0.146
	Increase	Decrease	Increase	Increase
Norway, Sweden and Finland (12/29/1997-6/30/1999)	0.410	0.686	0.612	0.160
	Increase	Decrease	Increase	Increase*
Norway, Sweden, Finland and western Denmark (7/1/1999-9/30/2000)	0.765	0.148	0.905	1.02
	Decrease	Increase**	Decrease*	Decrease
Norway, Sweden, Finland and Denmark (10/1-2000-12/31/2005)	0.198	0.943	0.174	0.0716

Table 5: Volatility and change in volatility of the logarithmic returns of the daily average of the hourly system price during the integration process at the Nordic power market. ** and * Significant change at the 5 and 10 per cent level, respectively.

According to the table, there has not been any significant change in volatility of electricity prices during the integration process at the Nordic power market.

The second model that we fit to our data set is the E-GARCH (1,1) specification and the reason is that this model allows positive and negative shocks to have different impact on the conditional volatility (see Nelson [28]):

$$\log \sigma_t^2 = \beta_0 + \beta_1 \cdot \frac{\varepsilon_{t-1}}{\sigma_{t-1}} + \beta_2 \cdot \left| \frac{\varepsilon_{t-1}}{\sigma_{t-1}} \right| + \beta_3 \log \sigma_{t-1}^2, \tag{21}$$

where β_3 is our volatility measure. Then, if we adopt the same idea as above to test whether there has been a change in volatility of electricity prices when the Nord Pool area has expanded in size with a new region, we fit the following model to two adjacent time series:

$$\begin{aligned} \log \sigma_t^2 &= \beta_0 + \beta_1 \cdot \frac{\varepsilon_{t-1}}{\sigma_{t-1}} + \beta_2 \cdot \left| \frac{\varepsilon_{t-1}}{\sigma_{t-1}} \right| + \\ & \quad \beta_3 \log \sigma_{t-1}^2 + \beta_4 \log \sigma_{t-1}^2 I_{t \geq t_0}. \end{aligned} \tag{22}$$

According to Table 5, the step-wise integration of the Nordic power market is most of the time associated with less volatile electricity prices. However, only one of the changes in volatility is significant (at the 5 per cent level).

However, as discussed, among others, in Dacorogna et al [7], there is often a long memory in the volatility of financial time series and since we have no reason to exclude this possibility for electricity prices, the third model that we fit to our data set is the I-GARCH (1,1) specification (see Engle and Bollerslev [14]):

$$\sigma_t^2 = \beta_0 + (1 - \beta_1)\,\varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \tag{23}$$

where β_1 is our volatility measure. Then, again, to test whether there has been a change in volatility of electricity prices during the step-wise integration of the Nordic power market, we fit the following model to two adjacent time series:

$$\sigma_t^2 = \beta_0 + (1 - \beta_1)\,\varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 + \beta_2 \sigma_{t-1}^2 I_{t \geq t_0}. \tag{24}$$

According to Table 5, only one of the changes in volatility of electricity prices is significant (at the 10 per cent level).

The fourth model that we fit to our data set is the Glosten-Jagannathan-Runkle GARCH (1,1) specification (see Glosten et al [19]):

$$\sigma_t^2 = \beta_0 + \beta_1 \varepsilon_{t-1}^2 + \beta_2 \sigma_{t-1}^2 + \beta_3 \varepsilon_{t-1}^2 I_{\varepsilon_{t-1} < 0}, \tag{25}$$

where $I_{\varepsilon_{t-1}<0} = 1$ when $\varepsilon_{t-1} < 0$ and zero elsewhere, and our volatility measure is β_2. The advantage of the GJR-GARCH (1,1) specification is that it provides a simple way to allow for asymmetries in the conditional volatility. Then, to test whether there has been a change in volatility of electricity prices when the Nord Pool area has expanded in size with a new region, we fit the following model to two adjacent time series:

$$\sigma_t^2 = \beta_0 + \beta_1 \varepsilon_{t-1}^2 + \beta_2 \sigma_{t-1}^2 + \beta_3 \varepsilon_{t-1}^2 I_{\varepsilon_{t-1}<0} + \beta_4 \sigma_{t-1}^2 I_{t \geq t_0}. \tag{26}$$

According to Table 5, only one of the changes in volatility of electricity prices is significant (at the 10 per cent level).

Then, which model is the best model? According to the information criteria that are proposed in Brooks and Burke [6], which are modified versions of the Akaike and Schwarz Information Criteria designed for GARCH models, the best model is the E-GARCH (1,1) specification in (21). (To come to this conclusion, we have fitted all models to the whole and unbroken time series with electricity prices.) Thus, the general picture is that the integration process at the Nordic power market most often has been associated with more stable and less volatile electricity prices, even though a one-to-one correspondence between the stability and volatility measures is lacking. Be aware that the lack of a one-to-one correspondence between these measures motivates (λ, σ^2)-analysis.

5 Discussion

Denmark, Finland, Norway and Sweden have gone through an extensive deregulation of their electricity sectors during the 1990s and, at the same time, there has been an evolution from national markets to a multi-national electricity market. Since the aim of the deregulation and integration of the electricity sectors in these countries was to develop a competitive power market, Bask et al [4] examine whether the degree of competition has increased over time during this process and they also find that this, in fact, is the case.

What we have found in this paper is that the step-wise integration of the Nordic power market most of the time also has been associated with less volatile electricity prices. The natural question to pose and answer is, therefore, why these prices have decreased in volatility over time? Of course, the natural candidate is that the dynamics that govern the evolution of electricity prices have changed, but it could also be the case that the nature of the shocks hitting the Nordic power market has changed.

By using a novel approach in this paper, we have found that the increased competition at the Nordic power market most of the time also has been associated with more stable electricity prices. That electricity prices at Nord Pool have increased in stability means that the dynamic system generating these prices has increased in stability during the step-wise integration of the Nordic power market. This, in turn, means that this market is less sensitive to shocks after the integration process than it was before this process. This is good news.

To summarize, what we have learned about Nord Pool regarding the relationship between market structure and the behavior of electricity prices are the following: (i) the degree of competition has increased over time during the integration process; (ii) the volatility of electricity prices has most often decreased over time; and (iii) the stability of the dynamic system generating electricity prices has most often increased over time. For this reason, one might ask: are there any theoretical justifications for these findings?

First of all, there is a small theoretical literature on the relationship between market structure and the price volatility for a storable commodity such as water, which to a large extent is used in electricity generation in the Nord Pool area. One study that our findings are consistent with is by McLaren [27] who develops an oligopoly model in which the price of the storable commodity decreases in volatility when the degree of competition increases. When it comes to the stability of the dynamics that govern the evolution of commodity prices, the literature is silent on its relationship with market structure and the price volatility.

Anyhow, we argue that when approximating models of the energy sector are developed in an attempt to understand the behavior of energy prices and how they depend on the market structure, the stability of the dynamics that govern the evolution of energy prices should be added as an important dimension in which model and data should be matched. That is, our finding that electricity prices most of the time have increased in stability and decreased in volatility when the Nordic power market has expanded and the degree of competition has increased should be properties of a well-formulated model of the Nordic power market.

To conclude, we believe that the empirical work in this paper has demonstrated that (λ, σ^2)-analysis can be a useful tool in empirical research. It is, therefore, our hope that future research will provide us with more applications of this tool than the Nord Pool experience.

References

[1] Bask, M. and de Luna, X. (2002). Characterizing the Degree of Stability of Non-Linear Dynamic Models. *Studies in Nonlinear Dynamics and Econometrics*, **6** (1) art. 3.

[2] Bask, M. and Gencay, R. (1998). Testing Chaotic Dynamics via Lyapunov Exponents. *Physica D*, **114**, 1-2.

[3] Bask, M., Liu, T. and Widerberg, A. (2007). The Stability of Electricity Prices: Estimation and Inference of the Lyapunov Exponents. *Physica A*, **376**, 565-572.

[4] Bask, M., Lundgren, J. and Rudholm, N. (2008). Market Power in the Expanding Nordic Power Market. Forthcoming in *Applied Economics*.

[5] Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, **31**, 307-327.

[6] Brooks, C. and Burke, S.P. (2003). Information Criteria for GARCH Model Selection. *European Journal of Finance*, **9**, 557-580.

[7] Dacorogna, M.M., Gencay, R., Müller, U., Olsen, R.B. and Pictet, O.V. (2001). *An Introduction to High-Frequency Finance*. Academic Press.

[8] Deaton, A. and Laroque, G. (1992). On the Behaviour of Commodity Prices. *Review of Economic Studies*, **59**, 1-23.

[9] Deaton, A. and Laroque, G. (1996). Competitive Storage and Commodity Price Dynamics. *Journal of Political Economy*, **104**, 896-923.

[10] Dechert, W.D. and Gencay, R. (1996). The Topological Invariance of Lyapunov Exponents in Embedded Dynamics. *Physica D*, **90**, 40-55.

[11] Dechert, W.D. and Gencay, R. (2000). Is the Largest Lyapunov Exponent Preserved in Embedded Dynamics? *Physics Letters A*, **276**, 59-64.

[12] Enders, W. (2003). *Applied Econometric Time Series*. 2nd Edition. Wiley.

[13] Engle, R.F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, **50**, 987-1007.

[14] Engle, R.F. and Bollerslev, T. (1986). Modelling the Persistence of Conditional Variances. *Econometric Reviews*, **5**, 1–50.

[15] Fernández-Rodríguez, F., Sosvilla-Rivero, S. and Andrada-Félix, J. (2005). Testing Chaotic Dynamics via Lyapunov Exponents. *Journal of Applied Econometrics*, **20**, 911-930.

[16] Gencay, R. (1996). A Statistical Framework for Testing Chaotic Dynamics via Lyapunov Exponents. *Physica D*, **89**, 261-266.

[17] Gencay, R. and Dechert, W.D. (1992). An Algorithm for the n Lyapunov Exponents of an n-Dimensional Unknown Dynamical System. *Physica D*, **59**, 142-157.

[18] Gencay, R. and Dechert, W.D. (1996). The Identification of Spurious Lyapunov Exponents in Jacobian Algorithms. *Studies in Nonlinear Dynamics and Econometrics*, **1** (3) art. 2.

[19] Glosten, L.R., Jagannathan, R. and Runkle, D.E. (1993). On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance*, **48**, 1779-1801.

[20] Guckenheimer, J. and Holmes, P. (1983). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. (Applied Mathematical Sciences, Vol. 42), Springer-Verlag: Berlin.

[21] Hamilton, J.D. (1994). *Time Series Analysis*. Princeton University Press: Princeton, New Jersey.

[22] Hjalmarsson, E. (2000). *Nord Pool: A Power Market without Market Power*. Göteborg University Working Paper in Economics No. 28.

[23] Hornik, K., Stinchcombe, M. and White, H. (1990). Universal Approximation of an Unknown Mapping and its Derivatives using Multilayer Feedforward Networks. *Neural Networks*, **3**, 551-560.

[24] Kennel, M.B., Brown, R. and Abarbanel, H.D.I. (1992). Determining Embedding Dimension for Phase-Space Reconstruction using a Geometrical Construction. *Physical Review A*, **45**, 3403-3411.

[25] Kuan, C.-M. and Liu, T. (1995). Forecasting Exchange Rates Using Feedforward and Recurrent Neural Networks. *Journal of Applied Econometrics*, **10**, 347-364.

[26] Lundgren, J., Hellström, J. and Rudholm, N. (2008). *Multinational Electricity Market Integration and Electricity Price Dynamics*. HUI Working Paper No. 16.

[27] McLaren, J. (1999). Speculation on Primary Commodities: The Effects of Restricted Entry. *Review of Economic Studies*, **66**, 853-871.

[28] Nelson, D.B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, **59**, 347-370.

[29] Newbery, D.M. (1984). Commodity Price Stabilization in Imperfect or Cartelized Markets. *Econometrica*, **52**, 563-578.

[30] Potter, S.M. (2000). Nonlinear Impulse Response Functions. *Journal of Economic Dynamics and Control*, **24**, 1425-1446.

[31] Rui, X. and Miranda, M.J. (1996). Solving Nonlinear Dynamic Games via Orthogonal Collocation: An Application to International Commodity Markets. *Annals of Operations Research*, **68**, 89-108.

[32] Sauer, T., Yorke, J.A. and Casdagli, M. (1991). Embedology. *Journal of Statistical Physics*, **65**, 579-616.

[33] Scheinkman, J.A. and Schechtman, J. (1983). A Simple Competitive Model with Production and Storage. *Review of Economic Studies*, **50**, 427-441.

[34] Shintani, M. (2006). A Nonparametric Measure of Convergence Towards Purchasing Power Parity. *Journal of Applied Econometrics*, **21**, 589-604.

[35] Shintani, M. and Linton, O. (2004). Nonparametric Neural Network Estimation of Lyapunov Exponents and a Direct Test for Chaos. *Journal of Econometrics*, **120**, 1-33.

[36] Takens, F. (1981). Detecting Strange Attractors in Turbulence. In *Dynamical Systems and Turbulence* (Lecture Notes in Mathematics, Vol. 898) by Rand, D.A. and Young, L.S., eds., Springer-Verlag: Berlin, 366-381.

[37] Thille, H. (2006). Inventories, Market Structure, and Price Volatility. *Journal of Economic Dynamics and Control*, **30**, 1081-1104.

[38] Urbach, R.M.A. (1999). *Footprints of Chaos in the Markets: Analyzing Non-Linear Time Series in Financial Markets and Other Real Systems.* Financial Times/Prentice Hall.

[39] Vassilopoulos, P. (2003). *Models for the Identification of Market Power in Wholesale Electricity Markets.* Discussion Paper, University of Paris-Dauphine.

[40] Vedenov, D.V. and Miranda, M.J. (2001). Numerical Solution of Dynamic Oligopoly Games with Capital Investment. *Economic Theory*, **18**, 237-261.

[41] Williams, J.C. and Wright, B.D. (1991). *Storage and Commodity Markets.* Cambridge University Press: Cambridge, Massachusetts.

[42] Wolf, A., Swift, J.B., Swinney, H.L. and Vastano, J.A. (1985). Determining Lyapunov Exponents From a Time Series. *Physica D*, **16**, 285-317.

[43] Ziehmann, C., Smith, L.A. and Kurths, J. (1999). The Bootstrap and Lyapunov Exponents in Deterministic Chaos. *Physica D*, **126**, 49-59.

Appendix

Descriptive statistics Descriptive statistics of each of the five time series are presented in Table A.1a-e.

Statistic	Quantity
Number of observations	1094
Minimum	-0.672
Maximum	0.701
Mean	0.000344
Median	-0.000762
Variance	0.0126
Skewness	0.260
Kurtosis	11.0
Dickey-Fuller unit root test	-10.7^{***}

Table A.1a: Descriptive statistics of the logarithmic returns of the daily average of the hourly system price during the period January 1, 1993, to December 31, 1995. *** Significant at the 1 per cent level.

Statistic	Quantity
Number of observations	727
Minimum	-0.277
Maximum	0.524
Mean	0.0000261
Median	-0.00192
Variance	0.00521
Skewness	1.04
Kurtosis	8.01
Dickey-Fuller unit root test	-9.59^{***}

Table A.1b: Descriptive statistics of the logarithmic returns of the daily average of the hourly system price during the period January 1, 1996, to December 28, 1997. *** Significant at the 1 per cent level.

Statistic	Quantity
Number of observations	550
Minimum	-0.718
Maximum	0.684
Mean	-0.00121
Median	-0.00921
Variance	0.0127
Skewness	0.431
Kurtosis	11.0
Dickey-Fuller unit root test	-8.67^{***}

Table A.1c: Descriptive statistics of the logarithmic returns of the daily average of the hourly system price during the period December 29, 1997, to June 30, 1999. *** Significant at the 1 per cent level.

Statistic	Quantity
Number of observations	457
Minimum	-0.697
Maximum	0.992
Mean	0.000587
Median	-0.00514
Variance	0.0109
Skewness	1.47
Kurtosis	23.6
Dickey-Fuller unit root test	-6.13^{***}

Table A.1d: Descriptive statistics of the logarithmic returns of the daily average of the hourly system price during the period July 1, 1999, to September 30, 2000. *** Significant at the 1 per cent level.

Statistic	Quantity
Number of observations	1916
Minimum	-0.774
Maximum	1.19
Mean	0.000470
Median	-0.00523
Variance	0.00778
Skewness	1.57
Kurtosis	30.2
Dickey-Fuller unit root test	-15.3^{***}

Table A.1e: Descriptive statistics of the logarithmic returns of the daily average of the hourly system price during the period October 1, 2000, to December 31, 2005. *** Significant at the 1 per cent level.

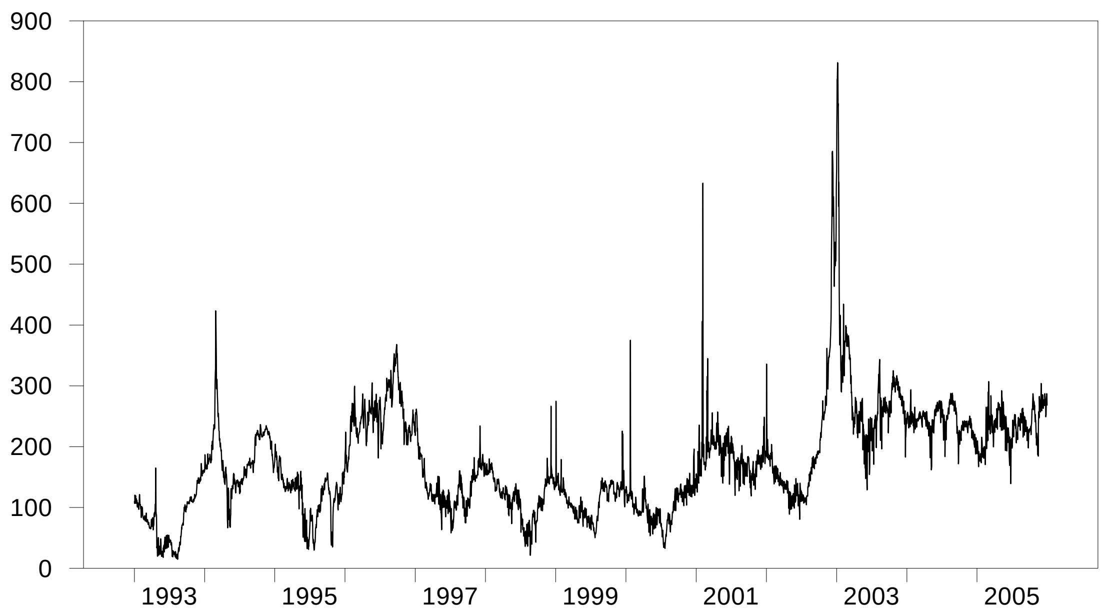

Figure 1a: Daily average of the hourly system price during the period January 1, 1993, to December 31, 2005.

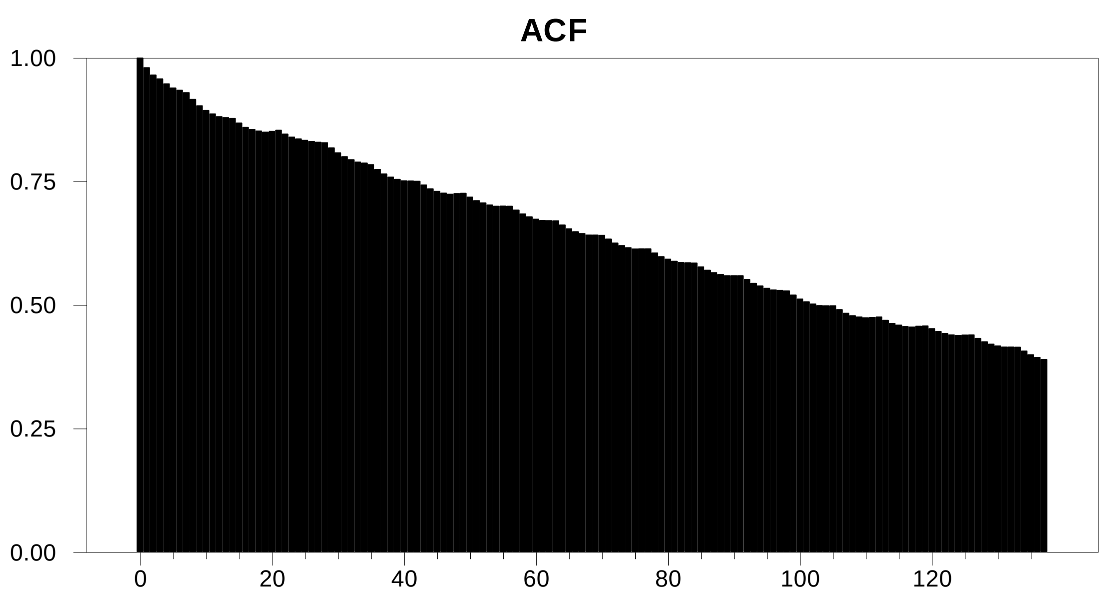

Figure 1b: ACF for the daily average of the hourly system price during the period January 1, 1993, to December 31, 2005.

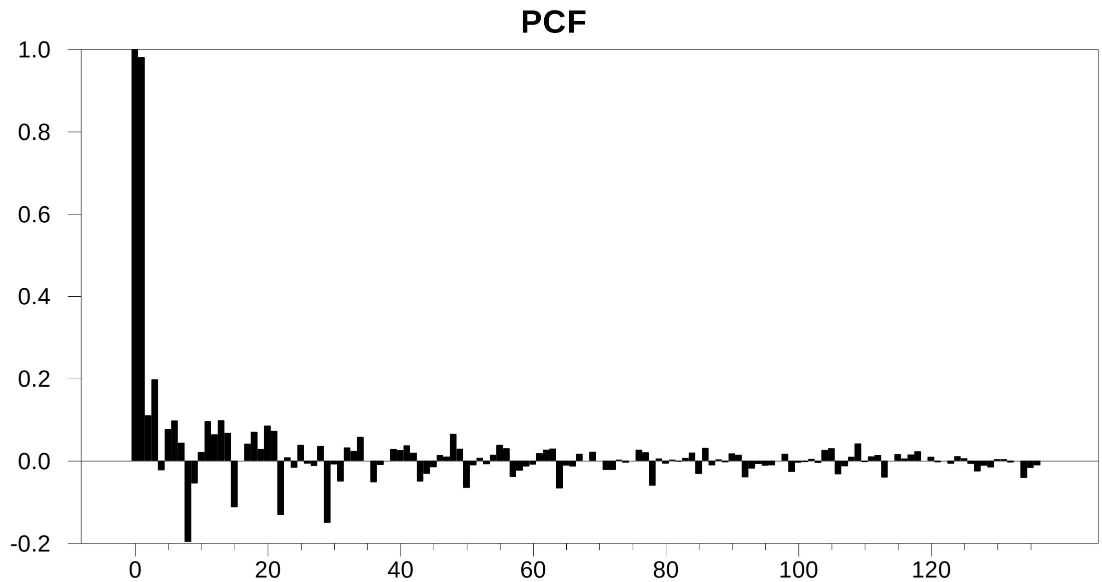

Figure 1c: PCF for the daily average of the hourly system price during the period January 1, 1993, to December 31, 2005.

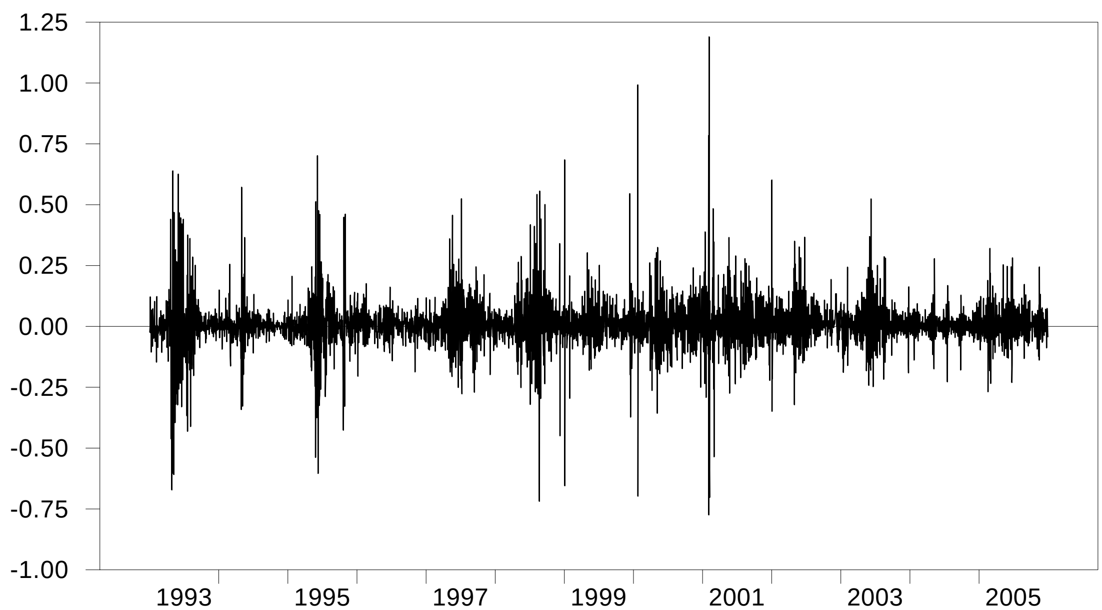

Figure 2a: Logarithmic returns of the daily average of the hourly system price during the period January 1, 1993, to December 31, 2005.

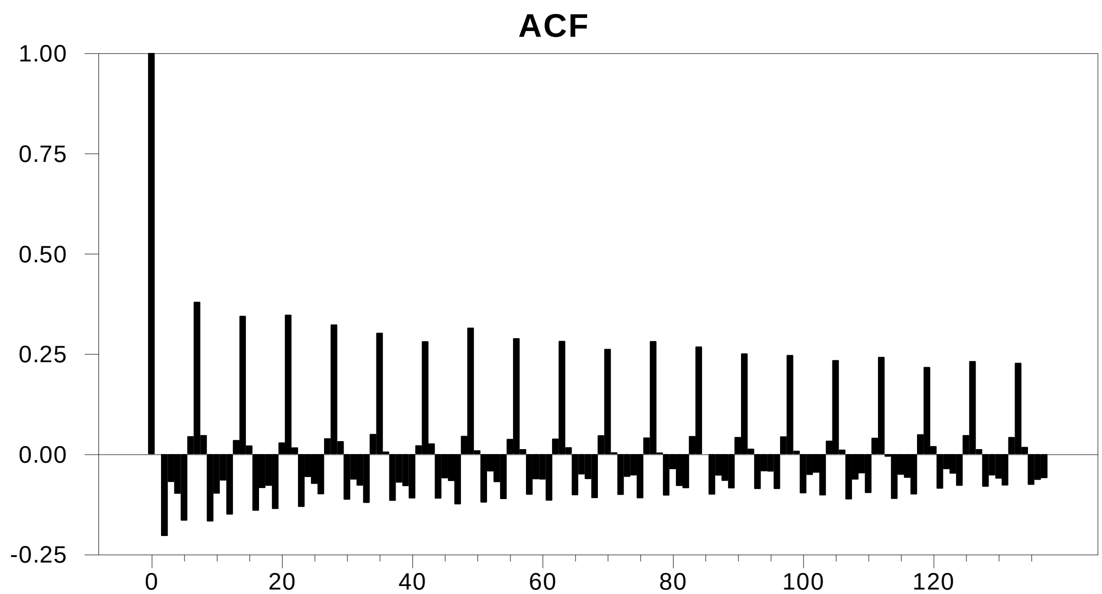

Figure 2b: ACF for the logarithmic returns of the daily average of the hourly system price during the period January 1, 1993, to December 31, 2005.

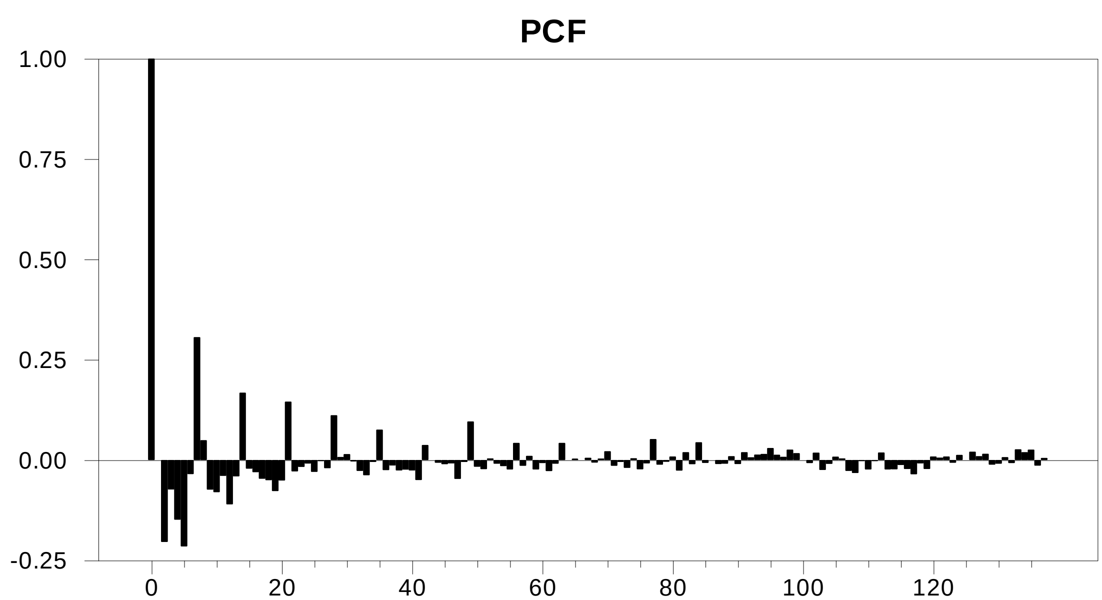

Figure 2c: PCF for the logarithmic returns of the daily average of the hourly system price during the period January 1, 1993, to December 31, 2005.