

Are multifractal processes suited to forecasting electricity price volatility? Evidence from Australian intraday data

MAWULI SEGNON ^a, CHI KEUNG LAU ^b, BERND WILFLING ^{a*},
RANGAN GUPTA ^c

^a *Westfälische Wilhelms-Universität Münster, Department of Economics (CQE),
Germany*

^b *University of Huddersfield, Department of Accountancy, Finance, and Economics,
UK*

^c *University of Pretoria, Department of Economics, South Africa*

(Date of this version: December 12, 2018)

Abstract

We analyze Australian electricity price returns and find that they exhibit multifractal structures. Consequently, we let the return mean equation follow a long memory smooth transition autoregressive (STAR) process and specify volatility dynamics as a Markov-switching multifractal (MSM) process. We compare the out-of-sample volatility forecasting performance of the STAR-MSM model with that of other STAR mean processes, combined with various conventional GARCH-type volatility equations (for example, STAR-GARCH(1,1)). We find that the STAR-MSM model competes with conventional STAR-GARCH specifications with respect to volatility forecasting, but does not (systematically) outperform them.

Keywords: Electricity price volatility; multifractal modeling; GARCH processes; volatility forecasting

JEL classification: C22, C52, C53

*Corresponding author. Tel.: +49 251 83 25040; fax: +49 251 83 25042. E-mail address: bernd.wilfling@wiwi.uni-muenster.de.

Are multifractal processes suited to forecasting electricity price volatility? Evidence from Australian intraday data

Abstract

We analyze Australian electricity price returns and find that they exhibit multifractal structures. Consequently, we let the return mean equation follow a long memory smooth transition autoregressive (STAR) process and specify volatility dynamics as a Markov-switching multifractal (MSM) process. We compare the out-of-sample volatility forecasting performance of the STAR-MSM model with that of other STAR mean processes, combined with various conventional GARCH-type volatility equations (for example, STAR-GARCH(1,1)). We find that the STAR-MSM model competes with conventional STAR-GARCH specifications with respect to volatility forecasting, but does not (systematically) outperform them.

Keywords: Electricity price volatility; multifractal modeling; GARCH processes; volatility forecasting

JEL classification: C22, C52, C53

1 Introduction

The deregulation of electricity markets and the rapid development of related financial products have unleashed an enormous interest in establishing econometric models that appropriately reflect the unique characteristics and dynamics of electricity price behavior. In particular, forecasting electricity price volatility has become a major task for analysts, energy companies and investors, due to its dominant role in pricing and risk management. As a result, a plethora of alternative volatility forecasting models have been proposed in the literature. (We give an overview in Section 2.)

In this paper, we apply an additional class of processes for modeling and forecasting electricity price volatility, which originally stem from the analysis of turbulent flows (Mandelbrot, 1974). Their adaptation to finance started with the work of Mandelbrot et al. (1997), which led to the first generation of multifractal models. Calvet and Fisher (2001, 2004) then introduced the second model generation, by specifying their Poisson multifractal model and its discretized version, the Markov-switching multifractal process.¹ To date, multifractal processes have been applied successfully in the volatility modeling and forecasting of different asset prices, including exchange rates (Calvet and Fisher, 2004; Lux, 2008), stock prices (Lux, 2008) and commodity prices (Lux et al., 2016; Wang et al., 2016; Segnon et al., 2017). Overall, these processes appear to be robust tools in capturing well-documented stylized facts of electricity price fluctuations and may therefore be considered as natural candidates for providing accurate volatility forecasts.

The main objective of this study is to investigate whether the multifractal framework — which accommodates both, long-term persistence in the volatility process, and structural breaks through regime-switching — (i) constitutes an appropriate tool for analyzing price volatility in deregulated electricity markets, and (ii) provides a comprehensive approach that outperforms alternative volatility models in terms of forecast-accuracy (Aggarwal et al., 2009).² To this end, we use the long memory smooth

¹For details on the genesis of multifractal models and applications in finance, see Lux and Segnon (2018).

²In a related context, Nowotarski et al. (2014) and Nowotarski and Weron (2015) consider forecast combinations of several individual volatility models. In this paper, we do not pursue this approach any further.

transition autoregressive (STAR) specification introduced by Hillebrand and Medeiros (2016) to model the mean equation of electricity price returns, while we represent the volatility equation by a Markov-switching multifractal (MSM) process. To the best of our knowledge, this study is the first to combine the STAR and MSM processes in a consolidated econometric framework, which we apply to Australian data. In contrast to previous studies, we do not confine ourselves to forecasting electricity price volatility only for the Australian New South, but provide an extended empirical analysis covering all five Australian regions. Overall, our analysis proceeds in two steps. (i) We investigate fractal structures of Australian electricity price returns (by computing the multifractal spectrum) via the so-called Multifractal Detrended Fluctuation Analysis (MFDFA) as proposed by Kantelhardt et al. (2002). (ii) We conduct a forecasting investigation in order to compare the volatility forecasting performance of our combined STAR-MSM model with the performance of several other combinations, in each of which the mean equation still follows the STAR process, but with the volatility equation governed by alternative GARCH-type processes.

Our analysis has two major findings. (i) Daily Australian electricity price returns exhibit multifractal structures. This empirical finding provides clear-cut motivation for using multifractal processes in the modeling of electricity-price dynamics. (ii) Incorporating multifractal structures into the volatility equation of the data-generating electricity-price process yields volatility forecasts that compete well with those of alternative conventional specifications, but does not (systematically) outperform them.

The remainder of the paper is organized as follows. In Section 2, we motivate our interest in Australian electricity markets and provide an overview of alternative forecasting models of electricity price dynamics. In Section 3, we analyze electricity price changes and present the results of the Multifractal Detrended Fluctuation Analysis. Section 4 presents our combined STAR-MSM model and briefly reviews the GARCH-type specifications used in our subsequent volatility forecasting analysis. In Section 5, we present our forecasting procedures and the results. Section 6 concludes.

2 Preliminaries and literature review

2.1 Characteristics of Australian electricity markets

The literature has extensively explored the price dynamics in the Australian electricity markets, due to its unique characteristics in the scale of the power system and the source of electricity generation. The National Electricity Market (NEM) is the world's largest interconnected power system running for more than 5,000 kilometers from North Queensland to Tasmania and central South Australia. In view of such a large interconnected power system, price volatility is high and uncertainty occurs for several reasons.

First, some generators do not follow the dispatch instructions issued by the Australian Energy Market Operator (AEMO) as a way to increase their revenue at the expense of the security of the power system and price uncertainty for end-consumers.³ Price also varies due to supply issues such as plant outages or constraints in the transmission network that limit transport capabilities. Another unique characteristic of NEM is its heavy dependence on coal fired generators. As the highest dependency compared to other developed countries, 84 percent of electricity was derived from coal in 1998 and 61 percent in 2014. The adoption of a carbon tax policy, implemented between July 2012 and July 2014, increased the cost of power production from coal.

Another policy uncertainty was induced by the introduction of a greenhouse gas emissions trading scheme by 2010. The legislation was initially passed by the House of Representatives, but rejected by the Senate. Apergis and Lau (2015) provide evidence that the Australian electricity market is a volatile one with a high degree of market power exercised by various generators, and one source of volatility stems from the failure to introduce a greenhouse gas emissions trading scheme by 2010. Therefore, it is important to forecast electricity price volatility in NEM, as its occurrence could be caused by demand shocks (e.g. extreme weather conditions), supply disruption (e.g. outages of transmission lines), or political uncertainty (e.g. rejection of greenhouse

³In July 2014, an electricity generator Snowy Hydro failed to comply with dispatch instructions issued by the Australian Energy Market Operator (AEMO), and the Australian Energy Regulator (AER) instigated proceedings against the company. Snowy Hydro paid total penalties of \$400,000, because of the potential hazard to public safety and material risk of damaging equipment.

gas emissions trading scheme).

2.2 Forecasting electricity price dynamics: A literature review

Up to date, a multitude of models for forecasting electricity price volatility have been proposed in the literature, which we can roughly divide into two groups. The first group of studies is based on non-traditional methodologies, like artificial intelligence and hybrid approaches, including (fuzzy) neural networks, fuzzy regression, cascaded architecture of neural networks and committee machines (e.g. Amjadi, 2006; Catalão et al., 2007; Vahidinasab and Kazemi, 2008; Amjadi and Hemmati, 2009; Amjadi, 2012).

The second group of studies, which are more relevant to our paper, include (i) traditional autoregressive time series models (Contreras et al., 2003; Garcia-Martos et al., 2007; Kristiansen, 2012), (ii) generalized autoregressive conditional heteroscedasticity (GARCH-type) models (Garcia et al., 2005; Chan and Gray, 2006; Koopman et al., 2007; Gianfreda, 2010; Cifter, 2013), (iii) jump-diffusion models (Huisman and Mahieu, 2003; Chan et al., 2008), (iv) autoregressive conditional hazard models (Christensen et al., 2012), and most recently (v), multivariate models (Raviv et al., 2015). Other authors combine traditional models with fundamentals (Karakatsani and Bunn, 2008; Huurman et al., 2012), use high-frequency data to forecast price volatility (Higgs and Worthington, 2005; Haugom et al., 2011; Gianfreda and Grossi, 2012; Haugom and Ullrich, 2012), and study the link between fundamentals, other energy markets and electricity price volatility (Goss, 2006; Gianfreda, 2010; Jonsson and Dahl, 2016).⁴

Although the literature has expressed a strong interest in studying the forecasting performance of electricity spot prices (Hong, 2015; Weron, 2014; Nowotarski et al., 2014; Nowotarski and Weron, 2015; Maciejowska et al., 2016), only a few empirical studies have focused on forecasting realized volatility in world-wide electricity markets. Nowotarski et al. (2014) perform a backtesting analysis using seven averaging and one selection scheme on day-ahead electricity prices in three major European and US markets. While the authors find that averaging forecasting techniques outper-

⁴For a detailed review of the different methodologies, see Möst and Keles (2010), Zareipour (2012) and Weron (2014).

forms counterpart models under normal market conditions, it fails to outperform an individual model in a more volatile environment or in the presence of price spikes. According to Clements et al. (2015), there is ample early literature dealing with price spikes, when forecasting spot electricity prices with various modeling approaches, including thresholds, Bernoulli and Poisson jump processes, heavy tailed error processes, Markov-switching, and diffusion models with time-varying intensity parameters. Examples include Misiolek et al. (2006), Knittel and Roberts (2005), Swider and Weber (2007), Higgs and Worthington (2008), among others.

Another strand of the recent literature focuses on forecasting the probability of such spike events, instead of simply forecasting the level of spot prices. These authors apply a multivariate point process to model the occurrence and size of extreme price events and conclude that physical infrastructure is the main influential factor in determining the transmission of price spikes in interconnected regions of the Australian electricity market (Clements et al., 2015). Their econometric approaches also lead to improved forecast indicators, such as forecasts and estimates of risk measures (for example value-at-risk and expected shortfall).

Qu et al. (2016) note that most of the early research uses the generalized autoregressive conditional heteroskedastic (GARCH) model introduced by Bollerslev (1986) and related models for estimating and forecasting electricity price volatility. The authors use a group of logistic smooth transition heterogeneous autoregressive (LSTHAR) models to forecast realized volatility in the Australian New South Wales (NSW) electricity market. Their modeling approaches allow simultaneously capturing of long memory behavior, as well as sign and size asymmetries, and provide improved volatility forecasts.

3 Preliminary data analysis

Currently, there are five Australian states — Queensland (QLD), New South Wales (NSW), Victoria (VIC), South Australia (SA) and Tasmania (TAS) — operating via a nationally interconnected grid. Since December 1998, the Australian Energy Market Operator (AEMO) has been responsible for operating Australia’s electricity markets and power systems, and the main domestic network is known as the National Electricity

Market (NEM). In 1998, NEM started operating as a wholesale market for the supply of electricity to retailers and end-users in Queensland, New South Wales, Victoria and South Australia, whereas Tasmania joined the NEM only in 2005.⁵

Exchange between electricity producers and retailers is facilitated through a spot and future market operated by the Australian Energy Market Operator, where the output from all generators is aggregated and scheduled to meet the demand of end-use customers. In our analysis, we use 5-minute intraday data, and the electricity dispatched price was obtained from AEMO with prices quoted as Australian dollars per megawatt hour (MWh).⁶ In each 24-hour period, there are 288 trading intervals, and we transform our 1087776 intraday observations, covering the time period from 1 January 2006 00:00 until 4 May 2016 23:55 into 3777 daily prices by averaging intraday prices. The spot price is the unit price received by generators, by selling electricity to the pool, where the output from all generators is aggregated and scheduled to meet demand.

In a first step, we compute the (daily) realized variance (RV) using 5-minute (intraday) data. We denote the daily logarithmic dispatch prices by p_t , for $t = 0, 1, 2, \dots$. To formally establish the realized variance at date t , we partition the daily log dispatch electricity price process $\{p_t\}$ further, by observing $n + 1$ equidistantly spaced (log) intraday prices $p_{t:0}, p_{t:1}, \dots, p_{t:n}$ (where $p_{t:0} = p_{t-1:n}$), and then define the realized variance at date t as

$$RV_t = \sum_{i=1}^n (p_{t:i} - p_{t:i-1})^2. \quad (1)$$

The stochastic properties of the realized variance RV_t from Eq. (1), in particular its use as a consistent estimator of the so-called integrated variance from continuous-time (jump) diffusion models of the log electricity price, have been discussed extensively in the literature (see McAleer and Medeiros, 2008, for an in-depth overview). However, in our subsequent empirical analysis, we consider RV_t as a proxy of the true (but practically unobservable) daily volatility, and use RV_t for assessing the accuracy of our

⁵The NEM operates on the world's longest interconnected power system with a distance of around 5,000 kilometers. The annual turnover of electricity traded is more than \$10 billion, so as to meet the demand of more than eight million end-user consumers.

⁶The data can be downloaded from <http://www.aemo.com.au/Electricity/National-Electricity-Market-NEM/Data-dashboard>.

volatility forecasts in Section 5.⁷

Figure 1 about here

Figure 2 about here

Table 1 about here

Figures 1 and 2 display the daily electricity price returns (defined as the difference between two successive daily log prices, $x_t = p_t - p_{t-1}$) and the daily realized variances RV_t along with their autocorrelation functions for the five Australian states. Descriptive statistics of the time series for all five states are reported in Table 1. We observe negatively skewed electricity price returns in New South Wales, Queensland, Tasmania and Victoria, and positively skewed price returns in South Australia. All return series exhibit excess kurtosis, thus conflicting with the Normal distribution. Table 1 also displays the tail index of the price returns, which we computed via the Hill estimator. All tail indices range between 1.3 and 2.1, indicating that the returns series for all five states exhibit heavy tails with (i) finite means, but infinite variances for New South Wales, Tasmania and Victoria (tail indices between 1 and 2), and (ii) finite means and finite variances for Queensland and South Australia (tail indices larger than 2). In line with these latter findings, the Jarque-Bera test (JB in Table 1) rejects the null hypothesis of normally distributed price returns for all five Australian states.

In order to analyze long memory properties in the data, we employ the Hurst index, computed via the detrended fluctuation analysis (DFA) described in Weron (2002). For all five states, the Hurst index values obtained for the price returns (Hurst index¹ in Table 1) are rather close to 0, indicating strongly anti-persistent (or mean-reverting) dynamics. We note that this anti-persistence in Australian electricity return data contrasts with the price return dynamics frequently observed for other commodities (where Hurst exponents are typically close to 0.5). For the realized variances, we obtain

⁷The concept of approximating unobservable daily volatility by the realized variance obtained from intraday data was proposed by Andersen and Bollerslev (1998). See also Marcucci (2005) and Reher and Wilfling (2016), who recently use this approach in the context of volatility forecasting.

Hurst exponents (Hurst index² in Table 1) substantially larger than 0.5, indicating the presence of long memory in return volatility.

Table 2 about here

The Ljung-Box Q-tests out to lag 5 (Q(5) in Table 1) reject the null hypothesis of no autocorrelation in the electricity price returns and the Engle tests for heteroscedasticity (ARCH-tests) indicate significant ARCH effects in the returns for all five states. The Phillips-Perron tests (PP and PP*) in Table 2 reject the null hypothesis of a unit root for all states at the 1% level. These (stationarity) results for the return series are confirmed by the Kwiatkowsky-Phillips-Schmidt-Shin tests (KPSS and KPSS*) in Table 2, that are unable to reject the null hypothesis of stationarity.

Table 3 about here

In order to analyze the (unconditional) variance processes of the price returns, we apply a structural break test according to Sansó et al. (2004), which is based on the modified iterated cumulative squares (ICSS) algorithm. This procedure accounts for potential heteroscedasticity and excess kurtosis in the data. As reported in Table 3, the algorithm detects two break points in the (unconditional) return variances for New South Wales, 5 for Queensland, one for South Australia and Victoria, and two break points for Tasmania. The dates, at which the break points occur, are highlighted in Table 3. These breakpoints will become relevant in our forecasting analysis in Section 5.

Figure 3 about here

Figure 4 about here

Finally, we analyze the fractal properties of the electricity price returns via the multifractal detrended fluctuation analysis (MFDFA), as developed by Kantelhardt et al. (2002). The MFDFA is based on the computation of local root mean square (RMS) for non-overlapping multiple segment sizes and allows us (i) to estimate the multifractal

spectrum of power law exponents from the electricity price returns, and (ii) to compare its characteristics with those of monofractal time series. We first compute the q -order Hurst exponent as slopes of regression lines for each q -order RMS. The results are depicted in Figure 3. We observe that the slopes of the regression lines vary with q -order, i.e. are q -order dependent. The q -order Hurst exponents (H_q) for all five states decrease with increasing segment sample size, indicating that the small segments are able to distinguish between the local periods with high and low volatility. Figure 4 displays the multifractal spectra for the five Australian states, each of which resembles a large arc. This observation is in contrast with the "small arc" spectra typically observed for monofractal time series.

4 Econometric modeling

4.1 The STAR-MSM model

We start by splitting up the data-generating process of our date- t electricity price return x_t into a mean and a volatility equation by writing

$$x_t = \mu_t + \epsilon_t. \quad (2)$$

In Eq. (2), μ_t represents the mean process and ϵ_t the volatility process with zero mean and variance σ_t^2 . Since we focus on modeling and forecasting the volatility of electricity price returns rather than return levels, we use a unique mean process throughout this paper. In particular, we apply a special variant of the long memory smooth transition autoregressive (STAR) model, as suggested by Hillebrand and Medeiros (2016), which is able to reflect long memory properties in the data. The STAR specification of the mean process can be formalized as

$$(1 - L)^d x_t = \phi_0(x_{t-1}, \eta_0) + \sum_{i=1}^p \phi_i(x_{t-1}, \eta_i) (1 - L)^d x_{t-i} + \epsilon_t. \quad (3)$$

In Eq. (3), $(1-L)^d$ is the fractional differencing operator with parameter $d \in (-0.5, 0.5)$, defined as $(1 - L)^d = \sum_{k=0}^{\infty} \frac{\Gamma(k-d)L^k}{\Gamma(-d)\Gamma(k+1)}$, where $\Gamma(\cdot)$ denotes the Gamma function. Moreover, for $i = 0, 1, \dots, p$, we consider the nonlinear functions $\phi_i(x_{t-1}, \eta_i) = \phi_{i0} +$

$\phi_{i1}h[\gamma(x_{t-1} - c)]$ with $h(z) = [1 + \exp(z)]^{-1}$ and parameter vector $\eta_i = (\phi_{i0}, \phi_{i1}, \gamma, c)'$.

As a benchmark for representing the volatility process ϵ_t from Eq. (2), we use the Markov-switching multifractal (MSM) model that adopts a multiplicative structure (Calvet and Fisher, 2001, 2004, 2008; Lux, 2008). In particular, the centered returns are modeled as

$$x_t - \mu_t = \epsilon_t = \sigma_t^2 \cdot u_t \quad (4)$$

with $\{u_t\}$ being an i.i.d. standard normal process and

$$\sigma_t^2 = \sigma^2 \prod_{i=1}^k M_t^{(i)}. \quad (5)$$

In Eq. (5), the instantaneous volatility σ_t^2 equals the product of a positive scale parameter σ^2 and k random volatility components $M_t^{(1)}, \dots, M_t^{(k)}$. The dynamics of the MSM model arise from renewing each multiplier $M_t^{(i)}$ at date t with probability γ_i (with probability $1 - \gamma_i$ the multiplier remains unchanged), where all new multipliers at all levels of the hierarchy are independently sampled from a common base distribution M . In our empirical analysis below, we follow Calvet and Fisher (2004) and model the base distribution M as a binomial random drawing of the two distinct values m_0 and $2 - m_0$ ($1 \leq m_0 \leq 2$) each with probability 0.5 (implying $\mathbb{E}(M_t^{(i)}) = 1$ for all $i = 1, \dots, k$ at every date t).

Calvet and Fisher (2001) suggest a specific structure of the transition probabilities γ_i , under which our discrete-time MSM converges towards a Poisson multifractal process in continuous time. However, since the continuous-time limit does not play any role in our analysis, we use the transition probabilities

$$\gamma_i = 2^{i-k}, \quad (6)$$

as proposed by Lux (2008). The probabilities from Eq. (6) entail a parsimonious overall model specification and thus greatly facilitate parameter estimation. We note that the hierarchical structure with a spectrum of low and high frequencies of the renewing components induces temporal dependence in the volatility process that — with increasing k — might generate similarity to long memory over a large number of lags. We obtain optimal forecasts in the MSM model via Bayesian updating of the

conditional probabilities for the unobserved volatility states.

4.2 Alternative volatility equations

4.2.1 GARCH-type models

Since the seminal papers of Engle (1982) and Bollerslev (1986), GARCH (generalized autoregressive conditional heteroscedasticity) models have become the most prominent tool for modeling and forecasting the time-varying conditional variance σ_t^2 in Eq. (4). In order to account for specific volatility characteristics (asymmetries, long memory), a multitude of alternative GARCH specifications has been suggested in the literature. While Hentschel (1995) establishes a nesting framework for many of these distinct GARCH models, we focus on six GARCH specifications in our forecasting analysis below and briefly review their probabilistic structure and the corresponding forecasting formulae.

The best-known specification is the standard GARCH(1,1) model, for which the conditional variance is parameterized as

$$\sigma_t^2 = \omega + \alpha \epsilon_{t-1}^2 + \beta \sigma_{t-1}^2, \quad (7)$$

with $\omega > 0$, $\alpha > 0$, $\beta > 0$ and $\alpha + \beta < 1$ (Bollerslev, 1986). For the GARCH(1,1) model, h -step ahead forecasts of the conditional variance σ_t^2 are given by the formula

$$\hat{\sigma}_{t+h}^2 = \bar{\sigma}^2 + (\alpha + \beta)^h (\sigma_t^2 - \bar{\sigma}^2), \quad (8)$$

where $\bar{\sigma}^2 = \omega / (1 - \alpha - \beta)$ denotes the unconditional variance of the process.

The GJR-GARCH(1,1) model, suggested by Glosten et al. (1993), extends the standard GARCH model by capturing asymmetric leverage effects, a phenomenon frequently observed in real-world data. Its conditional variance is specified as

$$\sigma_t^2 = \omega + \alpha \epsilon_{t-1}^2 + \gamma \epsilon_{t-1}^2 I(\epsilon_{t-1} < 0) + \beta \sigma_{t-1}^2, \quad (9)$$

where $I(\cdot)$ denotes the indicator function, which takes on the value 1 if the market is shocked by bad news ($\epsilon_{t-1} < 0$), and is 0 otherwise (in the case of good news). In Eq. (9), the parameter γ quantifies the magnitude of the asymmetric leverage effect.

For the GJR-GARCH(1,1) model the h -step ahead forecasts are given by

$$\hat{\sigma}_{t+h}^2 = \bar{\sigma}^2 + \left(\alpha + \beta + \frac{\gamma}{2} \right)^h (\sigma_t^2 - \bar{\sigma}^2), \quad (10)$$

where $\bar{\sigma}^2 = \omega / (1 - \alpha - \beta - \gamma/2)$ is the unconditional variance of the process.

Another GARCH variant to account for asymmetric effects of bad and good news on conditional volatility is Nelson's (1991) exponential GARCH (EGARCH) model. Setting $z_t = \epsilon_t / \sigma_t$ and assuming that z_t follows a standard normal distribution, we define the logarithmic conditional variance in an EGARCH(1,1) model as

$$\ln(\sigma_t^2) = \omega + \alpha (|z_{t-1}| - \mathbb{E}[|z_{t-1}|]) + \gamma z_{t-1} + \beta \ln(\sigma_{t-1}^2). \quad (11)$$

We note that the EGARCH(1,1) specification does not impose any restrictions on the parameters from Eq. (11), since it models $\ln(\sigma_t^2)$ instead of σ_t^2 . The parameter γ is able to capture asymmetric leverage effects, since in the case of $\gamma < 0$ positive (negative) values of z_{t-1} have a decreasing (an increasing) effect on conditional volatility. For the EGARCH(1,1) model, the h -step ahead forecast is given

$$\ln(\hat{\sigma}_{t+h}^2) = \bar{\sigma}^2 + \beta^h [\ln(\sigma_t^2) - \bar{\sigma}^2], \quad (12)$$

where $\bar{\sigma}^2 = (\omega - \gamma/\sqrt{2/\pi}) / (1 - \beta)$.

The asymmetric power ARCH (APARCH) model, introduced by Ding et al. (1993), allows us to model the leverage effect and a few other characteristics frequently encountered in the dynamics of real-world financial returns. Here, the conditional standard deviation σ_t raised to the power $\delta > 0$ is modeled as

$$\sigma_t^\delta = \omega + \alpha (|\epsilon_{t-1}| - \gamma \epsilon_{t-1})^\delta + \beta \sigma_{t-1}^\delta, \quad (13)$$

where the parameter γ again represents a leverage coefficient. Within this model class, the h -step ahead forecast formula is given by

$$\hat{\sigma}_{t+h}^\delta = \kappa + (\alpha c + \beta)^h (\sigma_t^\delta - \kappa), \quad (14)$$

where $\kappa = \omega(1 - \alpha c - \beta)^{-1}$ is the unconditional variance to the power δ and

$$c = \frac{1}{\sqrt{2\pi}} [(1 + \gamma)^\delta + (1 - \gamma)^\delta] 2^{\frac{\delta-1}{2}} \Gamma\left(\frac{\delta+1}{2}\right).$$

In our preliminary data analysis in Section 3, we employed the Hurst index to indicate that long memory is likely to be present in the volatility process of electricity price changes. In order to incorporate long memory into the family of GARCH processes, Baillie et al. (1996) proposed the fractionally integrated GARCH (FIGARCH) model, the formulation of which is based on a fractional differentiation parameter $d \in [0, 1]$ (FIGARCH(1, d ,1)). Without going into technical details, the conditional variance σ_t^2 of the FIGARCH model class can be expressed as

$$\sigma_t^2 = \frac{\omega}{1 - \beta} + \eta_1 \epsilon_{t-1}^2 + \eta_2 \epsilon_{t-2}^2 + \dots, \quad (15)$$

with $\omega > 0, \beta < 1$, and where the parameters η_1, η_2, \dots can be computed recursively from a system of equations, which depend on the fractional differentiation parameter d .⁸ From Eq. (17) an obvious 1-step ahead forecast of σ_t^2 is given by

$$\hat{\sigma}_{t+1}^2 = \omega(1 - \beta)^{-1} + \eta_1 \epsilon_t^2 + \eta_2 \epsilon_{t-1}^2 + \dots \quad (16)$$

Using $\hat{\sigma}_{t+1}^2$ from Eq. (16), we write the h -step ahead forecasts of the FIGARCH(1, d ,1) model as

$$\hat{\sigma}_{t+h}^2 = \omega(1 - \beta)^{-1} + \sum_{i=1}^{h-1} \eta_i \hat{\sigma}_{t+h-i}^2 + \sum_{j=0}^{\infty} \eta_{h+j} \epsilon_{t-j}^2. \quad (17)$$

The infinite sum on the right side of Eq. (17) clearly reflects the long memory properties of volatility forecasts in the FIGARCH model.

4.2.2 The Markov-Switching GARCH Model

Besides the (single-regime) GARCH-type models discussed so far, a large strand of theoretical and empirical literature on two-regime Markov-switching GARCH (MS-GARCH) models has evolved, taking into account sudden (stochastic) changes in the

⁸Besides the stated inequality conditions, the parameters ω and β need to satisfy further technical restrictions (see Baillie et al., 1996).

level of the conditional variance. In most applications, it is assumed that the data-generating process of the conditional variance is affected by a non-observable two-state first-order Markov-process S_t , that is, at any point in time, we either have $S_t = 1$ or $S_t = 2$. The stochastic switches between Regimes 1 and 2 are usually modeled by constant transition probabilities and for each of the two distinct regimes, the conditional variance equation is governed by a regime-specific set of parameters. Recently, Reher and Wilfling (2016) present an overview of alternative Markov-switching GARCH approaches and implement of a rich class of Markov-switching GARCH models.

As described in Reher and Wilfling (2016), it is generally difficult to obtain analytic expressions for multiple state-ahead volatility forecasts within a Markov-switching GARCH framework. An exception is the two-regime Markov-switching GARCH(1,1) model proposed by Klaassen (2002), which models the conditional variance as

$$\sigma_{t-1}^2\{\epsilon_t|\tilde{S}_t\} = \omega_{S_t} + \alpha_{S_t}\epsilon_{t-1}^2 + \beta_{S_t}\mathbb{E}_{t-1}\left[\sigma_{t-1}^2\{\epsilon_{t-1}|\tilde{S}_{t-1}\}|S_t\right], \quad (18)$$

where \tilde{S}_t denotes the regime path $\{S_{t-1}, S_{t-2}, \dots\}$, and $\omega_{S_t} > 0$, $\alpha_{S_t}, \beta_{S_t} \geq 0$. The expectation on the right side of Eq. (18) makes use of the regime path \tilde{S}_t and is conditioned on the information set $\mathfrak{S}_{t-1} = \{\epsilon_{t-1}, \epsilon_{t-2}, \dots\}$ and S_t . Under this setup, Klaassen (2002) obtains the h -step-ahead volatility forecast of his MSGARCH(1,1) model as

$$\hat{\sigma}_{t,t+h}^2 = \sum_{i=1}^2 \Pr(S_{t+h} = i|\mathfrak{S}_t) \hat{\sigma}_{t,t+h}^{2(i)}, \quad (19)$$

where (i) $\Pr(S_{t+h} = i|\mathfrak{S}_t)$ represents the conditional probability of being in Regime i at date $t+h$, and (ii) $\hat{\sigma}_{t,t+h}^{2(i)}$ is the h -step-ahead volatility forecast in Regime i made at date t , which can be computed recursively as

$$\hat{\sigma}_{t,t+h}^{2(i)} = \omega^{(i)} + (\alpha^{(i)} + \beta^{(i)}) \mathbb{E}\left(\sigma_{t,t+h-1}^{2(i)}|\delta_{t+h}\right). \quad (20)$$

5 Volatility-forecasting performance

In this section, we evaluate the quality of volatility forecasts for the alternative models presented in Section 4. To this end, we divided each of our five regional data sets into appropriate in-sample and out-of-sample periods and then applied a daily rolling

window, in order to fix the number of observations used for the estimation of the respective models over time. We estimated all econometric specifications with the maximum likelihood techniques as suggested in the original articles. To initialize the rolling window for each region, we separated the in-sample from the out-of-sample period at the first break point detected in the five sub-data sets via the modified ICSS algorithm as proposed by Sansó et al. (2004). Table 3 reports all break points detected via this algorithm in the date format "dd/mm/yyyy". Thus, the initializing in-sample (out-of-sample) periods are given by (i) 01/01/2006 – 01/11/2009 (02/11/2009 – 04/05/2016) for New South Wales and Tasmania, (ii) 01/01/2006 – 15/02/2010 (16/02/2010 – 04/05/2016) for Queensland, (iii) 01/01/2006 – 10/02/2010 (11/02/2010 – 04/05/2016) for South Australia, and (iv) 01/01/2006 – 22/04/2010 (23/04/2010 – 04/05/2016) for Victoria.

5.1 Forecast evaluation criteria

For each state, we computed the model-specific volatility forecasts for the four alternative forecast horizons $h = 1, 5, 10, 20$ trading days. In a first step, we assess forecast accuracy on the basis of the two most frequently used measures, the root mean squared error (RMSE) and the mean absolute error (MAE), which are given respectively by

$$\text{RMSE} = \sqrt{T^{-1} \sum_{t=1}^T (\hat{\sigma}_{t,M}^2 - \text{RV}_t)^2}, \quad (21)$$

$$\text{MAE} = T^{-1} \sum_{t=1}^T |\hat{\sigma}_{t,M}^2 - \text{RV}_t|, \quad (22)$$

where $\hat{\sigma}_{t,M}^2$ denotes the volatility forecast for date t (given the forecast horizon h), M the specific model from which the forecast is obtained (e.g. $\hat{\sigma}_{t,\text{GARCH}(1,1)}^2$), RV_t is the actually observed realized variance at date t , and T denotes the number of out-of-sample observations.

In a second step, we make statistical inferences about the relative forecasting performance of our alternative volatility models by reporting the results of the Equal Predictive Ability (EPA) test suggested by Diebold and Mariano (1995), and the Superior Predictive Ability (SPA) test from Hansen (2005). The EPA test enables us

to compare the forecasting accuracy of two competing models (say M_1 and M_2) by considering the loss differential

$$d_t = g(e_{t,M_1}) - g(e_{t,M_2}), \quad (23)$$

where $e_{t,M_1} = \hat{\sigma}_{t,M_1}^2 - \text{RV}_t$ and $e_{t,M_2} = \hat{\sigma}_{t,M_2}^2 - \text{RV}_t$ are the model-specific forecast errors at date t and the loss function $g(\cdot)$ either denotes the squared error loss $g(e_{t,M_i}) = e_{t,M_i}^2$ or the absolute error loss $g(e_{t,M_i}) = |e_{t,M_i}|$. Then, the null hypothesis of the EPA test states that there is no difference in the forecast accuracy between two competing models:

$$H_0 : \mathbb{E}(d_t) = 0 \quad \text{for all } t. \quad (24)$$

For large sample sizes, an appropriate test statistic of the EPA test is given by

$$\text{EPA} = \frac{\bar{d}}{\sqrt{1/T \sum_{k=-N}^N \hat{\gamma}(k)}}, \quad (25)$$

where $\bar{d} = 1/T \sum_{t=1}^T d_t$, N is the nearest integer larger than $T^{1/3}$ and

$$\hat{\gamma}(k) = \frac{1}{T} \sum_{t=|k|+1}^T (d_t - \bar{d})(d_{t-|k|} - \bar{d}).$$

Following Diebold and Mariano (1995), the test statistic EPA is approximately standard normally distributed under the null hypothesis in large samples.⁹

In contrast to the EPA test, the SPA test suggested by Hansen (2005) allows us to compare a benchmark forecast model M_0 with K competitive forecast models M_1, \dots, M_K under a given loss function. In line with Eq. (23), we define the loss differential between the benchmark model M_0 and the alternative model $M_k \in \{M_1, \dots, M_K\}$ as

$$d_{t,M_k} = g(e_{t,M_0}) - g(e_{t,M_k}). \quad (26)$$

Based on these K loss differentials, we can state the null hypothesis that the benchmark

⁹We refer the reader to Diebold (2015), who provides a detailed overview and an in-depth discussion of the EPA framework.

model M_0 is not inferior to any of the other K competing models as

$$H_0 : \max\{\mathbb{E}(d_{t,M_1}), \dots, \mathbb{E}(d_{t,M_K})\} \leq 0 \quad \text{for all } t. \quad (27)$$

In order to express the test statistic of the SPA test, we define the sample mean of the k th loss differential as $\bar{d}_{M_k} = 1/T \sum_{t=1}^T d_{t,M_k}$ and consider the estimated variance $\widehat{\text{Var}}(\sqrt{T} \cdot \bar{d}_{M_k})$ for $k = 1, \dots, K$. We note that this latter variance is estimated by using a bootstrap method (Hansen, 2005). One way to test the null hypothesis from Eq. (27) is now to consider the test statistic

$$\text{SPA} = \max \left\{ \frac{\sqrt{T} \bar{d}_{M_1}}{\widehat{\text{Var}}(\sqrt{T} \cdot \bar{d}_{M_1})}, \dots, \frac{\sqrt{T} \bar{d}_{M_K}}{\widehat{\text{Var}}(\sqrt{T} \cdot \bar{d}_{M_K})} \right\}, \quad (28)$$

the p -values of which can be obtained via a stationary bootstrap procedure.

Table 4 about here

5.2 Out-of-sample forecasting results

Table 4 reports the root mean squared (RMSE) and mean absolute forecast errors (MAE) for all seven volatility specifications across the five Australian states, where the forecast horizons were chosen as 1, 5, 10 and 20 trading days. *Prima facie*, the forecast errors appear to be rather similar for all volatility models, except for the MSGARCH model, for which the forecast errors are substantially higher for the (long) forecast horizon $h = 20$ in 4 of 5 Australian states. This rough forecast-error analysis appears to suggest that six out of seven volatility models exhibit quasi indistinguishable forecasting performance across all Australian states and over all forecast horizons. In order to gain deeper insight, our next step consists of conducting a disaggregate state-specific error analysis of Table 4, in which we (i) highlight noteworthy forecasting performance results, and (ii) address the particular performance of the MSM model.

Starting our state-specific analysis with New South Wales, we find that the standard GARCH model performs best at the 5-day forecast horizon and beyond under both error criteria (RMSE, MAE). Here, the MSM model (i) outperforms the EGARCH model at all forecast horizons under the RMSE and at the 1-day horizon under the

MAE criterion, and (ii) slightly outperforms the APARCH model at the 5-day horizon under the RMSE criterion (while the reverse is true in all other cases). For Queensland, the standard GARCH model provides a good forecasting performance (relative to all other models) across all forecast horizons, and under both error criteria. The MSM model performs rather poorly and appears to be broadly dominated by all other volatility models. For South Australia, the GJR and the MSGARCH models alternately exhibit minimal RMSEs and MAEs, while the standard GARCH model outperforms all other specifications at the 10-day and 20-day forecast horizons under both error criteria. The MSM specification slightly outperforms the FIGARCH model in 5 out of 8 constellations (5-, 10-, 20-day forecast horizons under the RMSE; 1-, 5-day horizons under the MAE), while the reverse is true for the remaining 3 cases. For Tasmanian data, the MSGARCH model performs best at the 1-, 5-, and 10-day forecast horizons both under the RMSE and MAE criteria, but performs poorest for the 20-day forecast horizons. In this state, the MSM specification consistently outperforms the FIGARCH model at each forecast horizon under both error criteria. Finally, we note that the results for Victorian data appear to be inconclusive in that at each forecast horizon, we find a distinct best-performing model. Here, the MSM performance appears close to that of any other specification in most of the 8 cases.

Overall, it turns out to be difficult to identify a particular volatility model that unambiguously outperforms all alternative models over all forecast horizons and under both error criteria. In this spirit, the MSM specification competes with the GARCH alternatives, but does not systematically outperform them. Nevertheless, we emphasize two major advantages of the MSM model over the GARCH alternatives. (i) The MSM model is more parsimonious than all GARCH-type specifications (with only two parameters having to be estimated). (ii) Owing to computational burdens, the estimation of MSM specifications with a large number of volatility components ($k > 10$) becomes unfeasible. However, the future increase in computational processing power is likely to render the estimation of MSM models with many volatility components practicable. In Section 4.1, we noted that MSM specifications with large k are able to reproduce long-range dependence and therefore constitute promising candidate models for producing accurate long-horizon forecasts.

Table 5 about here

Table 6 about here

In order to compare the volatility forecasting performance among the distinct specifications on a statistical basis, we apply EPA and SPA tests as presented in Section 5.1. We report the test results in Tables 5 and 6, in the computation of which we used 5000 bootstrap samples to find p -values of the SPA tests. For the EPA tests in Table 5, we choose the MSM model as the benchmark specification (Model 2), against which we compare all other specifications. Only in very few cases does the MSM model significantly outperform any of the other specifications at the 5% level. In particular, this is the case for (i) the EGARCH model in New South Wales for the horizons $h = 5$ (squared error loss) and $h = 1$ (absolute error loss), (ii) the FIGARCH model in South Australia for the horizon $h = 1$ (absolute error loss) and in Tasmania for all horizons under both error losses, and (iii) the MSGARCH model in New South Wales, South Australia, Victoria for the horizon $h = 20$ under both error losses. In total, this amounts to a significant EPA volatility-forecasting outperformance of the MSM model over any other specification in only 7 % (17 out of 240) of the cases analyzed.

We emphasize that the EPA analysis from Table 5 only provides pairwise performance comparisons of the MSM model with each of the other specifications. It does not provide an overall evaluation of the MSM volatility forecasting performance, when compared with all other volatility specifications simultaneously. Statistical evidence on this latter issue is reported by the results of the SPA tests in Table 6, where the null hypothesis states that the benchmark model is not inferior to any of the other six competing models. When considering the MSM specification as the benchmark model, the SPA tests in Table 6 always reject the null hypothesis across all 5 states and for all forecast horizons under both error losses (that is, in 40 out of 40 tests) at the 5 % level. This yields the robust result that for each of the 40 volatility forecast settings (that is, across 5 states, 4 forecast horizons, 2 error loss functions) there is at least one competing specification that significantly outperforms the MSM model. Viewed from this angle, it is interesting to note that the standard GARCH model performs best with 13 out of 40 (13/40) rejections of the null hypothesis at the 5 % level, followed by

MSGARCH (18/40), APARCH (25/40), GJR (32/40), EGARCH (35/40), FIGARCH (37/40), and MSM (40/40).

Our major result is that we are not able to identify a unique model that systematically outperforms all alternative specifications across the 5 Australian states at every forecast horizon. This finding is in line with the results of previous forecasting studies using data from deregulated electricity markets. In a comparative study, comprising 47 published works on forecasting electricity prices, Aggarwal et al. (2009) concludes that

... there is no systematic evidence of out-performance of one model over other models on a consistent basis.

In similar vein, Weron (2014) reports mixed forecasting performance of one-and-the-same model in different studies.

Finally, two general comments on our modeling and forecasting strategy are in order. (i) The specification of our transition probabilities (following Lux, 2008) is extremely parsimonious. (In fact, it does not contain any parameter to be estimated.) It is well-documented, that this convenient form of the transition probabilities greatly facilitates estimation and performs well in the modeling of stock and commodity markets. It is conceivable, however, that this specification confines the flexibility of the model, and thus does adequately account for some empirical characteristics in deregulated electricity markets. It might be an interesting line of future research to apply the formally more complex transition probabilities specified in Calvet and Fisher (2004), which — through two additional parameters — allow for a wide range of persistence.

(ii) Our major motivation to adopt the MSM model is based on the empirical evidence of multifractality in the data. We emphasize that there is ongoing scientific debate about the reliability of the statistical approaches to detecting multifractality in real-world data. Barndorff-Nielsen and Prause (2001) and Lux (2004), *inter alia*, argue that the scaling-function and Hölder-spectrum analysis (Figures 3 and 4) may misdiagnose the presence of multifractal structures. According to Lux (2004), it is possible to obtain a non-linear shape of the empirical scaling function after randomization of the temporal data structure, while Barndorff-Nielsen and Prause (2001) illustrate how

— in the absence of true scaling — empirical scaling may emerge under heavy-tailed data distributions.

6 Conclusion

In this paper we analyze electricity price returns using a unique intraday data set covering 5-minute electricity prices for five Australian states. In a multifractal detrended fluctuation analysis, we find that electricity price returns exhibit multifractal behavior and therefore model the returns as a smooth transition autoregressive process with a Markov-switching multifractal volatility component (STAR-MSM model).

We implement a forecasting analysis, in which we compare the out-of-sample volatility-forecasting performance of our STAR-MSM model with that of several alternative conventional GARCH specifications. Although multifractal structures are statistically verifiable in the Australian data, incorporating them into the volatility processes of electricity price returns does not appear to substantially improve volatility forecasts, compared to those obtained from conventional GARCH models. A useful line of future research could therefore entail the specification of alternative multifractal volatility processes with improved predictive content for electricity prices.

References

- Aggarwal, S.K., Saini, L.M., Kumar, A., 2009. Electricity price forecasting in deregulated markets: A review and evaluation. *International Journal of Electrical Power & Energy Systems* 31, 13-22.
- Amjady, N., 2006. Day-ahead price forecasting of electricity markets by new fuzzy neural network. *IEEE Transactions on Power Systems* 21, 887-996.
- Amjady, N., 2012. Short-term electricity price forecasting. In J. P. S. Catalão (Ed.), *Electric Power systems: advanced forecasting techniques and optimal generation scheduling*. CRC, Press.
- Amjady, N., Hemmati, M., 2009. Day-ahead price forecasting of electricity markets by a hybrid intelligent system. *European transactions on Electrical Power* 19, 89-102.

- Andersen, T.G., Bollerslev, T., 1998. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885-905.
- Apergis, N., Lau, M.C.K., 2015. Structural breaks and electricity prices: Further evidence on the role of climate policy uncertainties in the Australian electricity market. *Energy Economics* 52, 176-182.
- Baillie, R.T., Bollerslev, T., Mikkelsen, H.O., 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 74, 3-30.
- Barndorff-Nielsen, O.E., Prause, K., 2001. Apparent scaling. *Finance and Stochastics* 5, 103-113.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307-327.
- Calvet, L., Fisher, A., 2001. Forecasting multifractal volatility. *Journal of Econometrics* 105, 27-58.
- Calvet, L., Fisher, A., 2004. Regime-switching and the estimation of multifractal processes. *Journal of Financial Econometrics* 2, 44-83.
- Calvet, L., Fisher, A., 2008. *Multifractal Volatility: Theory, Forecasting and Pricing*. New York, Academic Press.
- Catalão, J.P.S., Mariano, S.J.P.S., Mendes, V.M.F., Ferreria, L.A.F.M., 2007. Short-term electricity prices forecasting in a competitive market: A neural network approach. *Electric Power Systems Research* 77, 1297-1304.
- Chan, K.F., Gray, P., 2006. Using extreme value theory to measure value-at-risk for daily electricity spot prices. *International Journal of Forecasting* 22, 283-300.
- Chan, K.F., Gray, P., Campen, B.V., 2008. A new approach to characterizing and forecasting electricity price volatility. *International Journal of Forecasting* 24, 728-743.
- Christensen, T.M., Hurn, A.S., Lindsay, K.A., 2012. Forecasting spikes in electricity prices. *International Journal of Forecasting* 28, 400-411.
- Cifter, A., 2013. Forecasting electricity price volatility with Markov-switching GARCH model: Evidence from the nordic electric power market. *Electric Power Systems Research* 102, 61-67.

- Clements, A.E., Herrera, R., Hurn, A.S., 2015. Modelling interregional links in electricity price spikes. *Energy Economics* 51, 383-393.
- Contreras, J., Espinola, R., Nogales, F.J., Conejo, A.J., 2003. ARIMA models to predict next-day electricity prices. *IEEE Transactions on Power Systems* 18, 1014-1020.
- Diebold, F.X., 2015. Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of Diebold-Mariano tests. *Journal of Business and Economic Statistics* 33, 1-10.
- Diebold, F.X., Mariano, R.S., 1995. Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253-263.
- Ding, Z., Granger, C., Engle, R., 1993. A long memory property of stock market returns and a new model. *Journal of Empirical Finance* 1, 83-106.
- Engle, R., 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica* 50, 987-1007.
- Garcia, R.J., Contreras, J., Akkeren, M.V., Garcia, J.B.C., 2005. A GARCH forecasting model to predict day-ahead electricity prices. *IEEE Transactions on Power Systems* 20, 867-874.
- Garcia-Martos, R.C., Rodriguez, J., Sanchez, M.J., 2007. Mixed models for short-run forecasting of electricity prices: application for the spanish market. *IEEE Transactions on Power Systems* 22, 544-551.
- Gianfreda, A., 2010. Volatility and volume effects in European electricity Markets. *Economic Notes* 1, 47-63.
- Gianfreda, A., Grossi, L., 2012. Volatility models for electricity prices with intraday information. Available at SSRN: <https://ssrn.com/abstract=2188148> or <http://dx.doi.org/10.2139/ssrn.2188148>.
- Glosten, L., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and volatility of the nominal excess return on stocks. *Journal of Finance* 46, 1779-1801.
- Goss, B., 2006. Liquidity, volume and volatility in US electricity futures: The case of Palo Verde. *Applied Financial Economic Letters* 2, 43-46.
- Hansen, P.R., 2005. A test for superior predictive ability. *Journal of Business and Economic Statistics* 23, 365-380.

- Haugom, E., Ullrich, C.J., 2012. Forecasting spot price volatility using the shortterm forward curve. *Energy Economics* 34, 1826-1833.
- Haugom, E., Wesrgaard, S., Solibakke, P.B., Lien, G., 2011. Realized volatility and the influence of market measures on predictability: Analysis of Nord Pool forward electricity data. *Energy Economics* 33, 1206-1215.
- Hentschel, L., 1995. All in the family: Nesting symmetric and asymmetric GARCH models. *Journal of Financial Economics* 39, 71-104.
- Higgs, H., Worthington, A.C., 2005. Systematic features of high-frequency volatility in Australian electricity markets: Intraday patterns, information arrival and calendar effects. *The Energy Journal* 26, 23-40.
- Higgs, H., Worthington, A.C., 2008. Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The Australian wholesale spot electricity market. *Energy Economics* 30, 3172-3185.
- Hillebrand, E., Medeiros, M.C., 2016. Nonlinearity, breaks, and long-range dependence in time-series models. *Journal of Business and Economic Statistics* 34, 23-41.
- Hong, T., 2015. Energy forecasting: Past, present, and future. *foresight. Journal of Applied Forecasting* 32, 43-48.
- Huisman, R., Mahieu, R., 2003. Regime jumps in electricity prices. *Energy Economics* 25, 425-434.
- Huurman, C., Ravazzolo, F., Zhou, C., 2012. The power of weather. *Computational Statistics* 56, 3793-3807.
- Jonsson, E., Dahl, R.E., 2016. Regime shifts in electricity prices in USA and EU. <http://www.usaee.org/usaee2016>.
- Kantelhardt, J.W., Zschiegner, S.A., Koscielny-Bunde, E., Havlin, S., Bunde, A., Stanley, H.E., 2002. Multifractal detrended fluctuation analysis of nonstationary time series. *Physica A* 316, 87-114.
- Karakatsani, N.V., Bunn, D.W., 2008. Modelling the volatility of spot electricity prices. *International Journal of Forecasting* 24, 764-785.
- Klaassen, F., 2002. Improving GARCH volatility forecasts. *Empirical Economics* 27, 363-394.
- Knittel, C.R., Roberts, M.R., 2005. An empirical examination of restructured electricity prices. *Energy Economics* 27, 791-817.

- 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.
- Kristiansen, T., 2012. Forecasting Nord Pool day-ahead prices with an autoregressive model. *Energy Policy* 49, 328-332.
- Lux, T., 2004. Detecting multi-fractal properties in asset returns. *International Journal of Modern Physics* 15, 481-491.
- Lux, T., 2008. The Markov-switching multifractal model of asset returns: GMM estimation and linear forecasting of volatility. *Journal of Business and Economic Statistics* 26, 194-210.
- Lux, T., Segnon, M., 2018. Multifractal models in finance: Their origin, properties and applications. In S.-H. Chen, M. Kaboudan, and Y.-R. Du (Eds.), *The Oxford Handbook on Computational Economics and Finance*, pp. 204-248. Oxford University Press, New York.
- Lux, T., Segnon, M., Gupta, R., 2016. Forecasting crude oil price volatility and value-at-risk: Evidence from historical and recent data. *Energy Economics* 56, 117-133.
- Maciejowska, K., Nowotarski, J., Weron, R., 2016. Probabilistic forecasting of electricity spot prices using factor quantile regression averaging. *International Journal of Forecasting* 32, 957-965.
- Mandelbrot, B.B., 1974. Intermittent turbulence in self similar cascades; divergence of high moments and dimension of the carrier. *Journal of Fluid Mechanics* 62, 331-358.
- Mandelbrot, B.B., Fisher, A., Calvet, L., 1997. A multifractal model of asset returns. Cowles Foundation Discussion Papers 1164, Cowles Foundation for Research in Economics, Yale University.
- Marcucci, J., 2005. Forecasting stock market volatility with regime-switching GARCH models. *Studies in Nonlinear Dynamics & Econometrics* 9(4), 1-53.
- McAleer, M., Medeiros, M.C., 2008. Realized volatility: A Review. *Econometric Reviews* 27, 10-45.
- Misiorek, A., Trück, S., Weron, R., 2006. Point and interval forecasting of spot electricity prices: Linear vs. non-linear time series models. *Studies in Nonlinear Dynamics & Econometrics* 10.

- Möst, D., Keles, D., 2010. A survey of stochastic modelling approaches for liberalized electricity markets. *European Journal of Operational Research* 207, 543-556.
- Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* 59, 347-370.
- Nowotarski, J., Raviv, E., Trück, S., Weron, R., 2014. An empirical comparison of alternative schemes for combining electricity spot price forecasts. *Energy Economics* 30, 1116-1157.
- Nowotarski, J., Weron, R., 2015. Computing electricity spot price prediction intervals using quantile regression and forecast averaging. *Computational Statistics* 30, 791-803.
- Qu, H., Chen, W., Niu, M., Li, X., 2016. Forecasting realized volatility in electricity markets using logistic smooth transition heterogeneous autoregressive models. *Energy Economics* 54, 68-76.
- Raviv, E., Bouwman, K.E., van Dijk, D., 2015. Forecasting day-ahead electricity prices: utilizing hourly prices. *Energy Economics* 50, 227-239.
- Reher, G., Wilfling, B., 2016. A nesting framework for Markov-switching GARCH modelling with an application to the German stock market. *Quantitative Finance* 16, 411-426.
- Sansó, A., Arragó, V., Carrion, J.L., 2004. Testing for change in the unconditional variance of financial time series. *Revista de Economía Financiera* 4, 32-53.
- Segnon, M., Lux, T., Gupta, R., 2017. Modeling and forecasting the volatility of carbon dioxide emission allowance prices: A review and comparison of modern volatility models. *Renewable and Sustainable Energy Reviews* 69, 692-704.
- Swider, D.J., Weber, C., 2007. Extended ARMA models for estimating price developments on day-ahead electricity markets. *Electric Power Systems Research* 77, 583-593.
- Vahidinasab, V., Kazemi, S.J., 2008. Day-ahead forecasting in restructured power systems using artificial neural networks. *Electric Power Systems Research* 78, 1332-1342.
- Wang, Y., Wu, C., Yang, L., 2016. Forecasting crude oil market volatility: A Markov switching multifractal volatility approach. *International Journal of Forecasting* 32, 1-9.

- Weron, R., 2002. Estimating long-range dependence: finite sample properties and confidence intervals. *Physica A: Statistical Mechanics and its Applications* 312, 285-299.
- Weron, R., 2014. Electricity price forecasting: A review of the state-of-the-art with look into the future. *International Journal of Forecasting* 30, 1030-1081.
- Zareipour, H., 2012. Short-term electricity market prices: A review of characteristics and forecasting methods. In A. Sorokin, S. Rebennack, P.M. Pardalos, N.A. Iliadis, and M.V.F. Pereira (Eds.), *Handbook of Networks in Power Systems I*, pp. 89-121, Springer.

Tables and Figures

Table 1: Descriptive statistics for electricity price returns

	NSW	QLD	SA	TAS	VIC
Nb. of observations	3776	3776	3776	3776	3776
Minimum	-4.034	-4.145	-3.985	-3.298	-3.993
Maximum	4.064	4.045	4.396	2.674	3.618
Mean	4.605E-05	3.020E-4	2.120E-4	3.496E-4	2.079E-4
Std	0.335	0.406	0.468	0.289	0.335
Skewness	-0.489	-0.201	0.343	-0.366	-0.161
Kurtosis	42.943	26.599	25.065	31.476	36.308
Tail index	1.388	2.085	2.007	1.948	1.721
JB	2.512E+5 (0.001)	8.765E+4 (0.001)	7.668E+4 (0.001)	1.277E+5 (0.001)	1.746E+5 (0.001)
Hurst index ¹	0.157	0.147	0.118	0.178	0.138
Hurst index ²	0.812	0.859	0.750	0.850	0.760
Q(5)	329.221 (0.000)	343.307 (0.000)	382.147 (0.000)	324.342 (0.000)	341.068 (0.000)
ARCH-test	458.950 (0.000)	429.267 (0.000)	526.464 (0.000)	153.509 (0.000)	382.584 (0.000)

Note: Hurst index¹ denotes the Hurst values for electricity price returns, Hurst index² for the realized variances. p -values are in parantheses. The five Australian states are abbreviated as NSW (Nouth South Wales), QLD (Queensland), SA (South Australia), TAS (Tasmania), VIC (Victoria).

Table 2: Unit root and stationarity tests for electricity price returns

States	$H_0 : I(1)$		$H_0 : I(0)$	
	PP	PP*	KPSS	KPSS*
New South Wales	-233.362	-233.398	0.016	0.021
Queensland	-195.861	-195.910	0.014	0.014
South Australia	-258.811	-258.871	0.017	0.021
Tasmania	-184.912	-184.912	0.022	0.036
Victoria	-140.435	-140.462	0.009	0.010

Note: PP and PP* are the Phillips-Perron adjusted t-statistics of the lagged dependent variable in a regression with (i) intercept and time trend, and (ii) intercept only. The respective critical values at the 1% level are -3.960 and -3.432. KPSS and KPSS* are the Kwiatkowski, Phillips, Schmidt, and Shin test statistics using residuals from regressions with (i) intercept and time trend, and (ii) intercept only. The respective critical values at the 1% level are 0.216 and 0.739.

Table 3: Structural breaks in the variance processes of electricity price returns

State	No. of break points	Date (dd/mm/yyyy)	Standard deviation
New South Wales	2	01/01/2006 – 01/11/2009	0.374
		02/11/2009 – 22/02/2010	1.080
		23/02/2010 – 04/05/2016	0.208
Queensland	5	01/01/2006 – 15/02/2010	0.501
		16/02/2010 – 13/11/2014	0.251
		14/11/2014 – 20/03/2015	0.834
		21/03/2015 – 02/04/2015	0.257
		03/04/2015 – 28/03/2016	0.637
		29/03/2016 – 04/05/2016	0.637
South Australia	1	01/01/2006 – 10/02/2010	0.574
		11/02/2010 – 04/05/2016	0.382
Tasmania	2	01/01/2006 – 01/11/2009	0.354
		02/11/2009 – 22/02/2010	0.150
		23/02/2010 – 04/05/2016	0.247
Victoria	1	01/01/2006 – 22/04/2010	0.430
		23/04/2010 – 04/05/2016	0.246

Note: The bold dates represent the structural break points obtained from the modified iterated cumulated squares algorithm suggested by Sansó et al. (2004).

Table 4: Root mean squared errors (RMSE) and mean absolute errors (MAE)

Model	Forecast horizon (trading days)							
	1				5			
	1	5	10	20	1	5	10	20
	RMSE				MAE			
	New South Wales							
MSM	31.833	31.852	31.872	31.768	5.857	5.886	5.925	5.967
GARCH	31.803	31.839	31.850	31.742	5.836	5.824	5.835	5.841
GJR	31.744	31.848	31.865	31.760	5.827	5.849	5.888	5.925
EGARCH	33.452	31.857	31.874	31.769	6.240	5.886	5.921	5.962
APARCH	31.837	31.853	31.872	31.763	5.851	5.875	5.904	5.922
FIGARCH	31.782	31.842	31.855	31.754	5.839	5.848	5.874	5.903
MSGARCH	31.820	31.849	31.880	41.041	5.828	5.838	5.961	54.323
	Queensland							
MSM	118.556	118.570	118.564	118.585	27.251	27.263	27.273	27.386
GARCH	118.535	118.546	118.518	118.499	27.221	27.164	27.104	27.082
GJR	118.533	118.550	118.528	118.518	27.218	27.181	27.138	27.147
EGARCH	118.541	118.567	118.558	118.578	27.234	27.229	27.223	27.333
APARCH	118.550	118.556	118.541	118.548	27.229	27.196	27.157	27.212
FIGARCH	118.549	118.568	118.561	118.577	27.243	27.234	27.234	27.334
MSGARCH	118.551	118.551	118.539	118.543	27.202	27.199	27.190	27.256
	South Australia							
MSM	67.834	67.848	67.856	70.455	18.093	18.107	18.136	18.974
GARCH	67.816	67.834	67.839	70.420	18.069	18.065	18.068	18.864
GJR	67.789	67.823	67.847	70.436	18.056	18.076	18.105	18.926
EGARCH	67.819	67.848	67.857	70.458	18.073	18.105	18.137	18.979
APARCH	67.820	68.429	68.168	71.668	18.073	18.453	51.186	43.110
FIGARCH	67.828	67.850	67.856	70.456	18.105	18.113	18.134	18.963
MSGARCH	67.828	67.832	67.862	72.751	18.049	18.036	18.380	171.060
	Tasmania							
MSM	21.333	21.357	21.373	21.423	8.675	8.700	8.711	8.763
GARCH	21.320	21.346	21.358	21.403	8.653	8.666	8.666	8.708
GJR	21.319	21.344	21.355	21.401	8.651	8.660	8.658	8.701
EGARCH	21.325	21.350	21.364	21.410	8.658	8.680	8.684	8.726
APARCH	21.328	21.347	21.357	21.396	8.662	8.673	8.667	8.693
FIGARCH	21.341	21.364	21.380	21.431	8.685	8.711	8.723	8.775
MSGARCH	21.310	21.325	21.326	45.982	8.620	8.624	8.600	10.348
	Victoria							
MSM	58.437	58.444	58.454	61.463	10.305	10.325	10.369	11.218
GARCH	58.424	58.431	58.434	61.427	10.284	10.265	10.255	11.018
GJR	58.418	58.430	58.446	61.450	10.279	10.292	10.314	11.125
EGARCH	58.254	58.444	58.455	61.464	10.305	10.319	10.359	11.210
APARCH	58.433	58.443	68.761	67.490	10.297	10.306	13.916	13.880
FIGARCH	58.431	58.442	58.450	61.454	10.299	10.315	10.353	11.194
MSGARCH	58.430	58.432	58.429	71.413	10.260	10.248	10.435	75.927

Note: The volatility models are abbreviated as MSM (Markov-switching multifractal), GARCH (generalized autoregressive conditional heteroscedasticity), GJR (Glosten-Jagannathan-Runkle GARCH), EGARCH (exponential GARCH), APARCH (asymmetric power ARCH), FIGARCH (fractionally integrated GARCH), MSGARCH (Markov-switching GARCH).

Table 5: Equal predictive ability (EPA) tests (p -values)

Model 1	Model 2	Forecast horizon (trading days)															
		1				5				10				20			
		Squared error loss				Absolute error loss											
New South Wales																	
GARCH	MSM	0.995	0.999	1.000	1.000	1.000	1.000	1.000	1.000	1.000							
GJR		0.946	0.971	1.000	1.000	1.000	1.000	1.000	1.000	1.000							
EGARCH		0.081	0.036	0.176	0.258	0.039	0.499	0.833	0.845								
APARCH		0.198	0.421	0.554	0.995	0.999	1.000	1.000	1.000								
FIGARCH		0.959	1.000	0.997	0.999	1.000	1.000	1.000	1.000								
MSGARCH		0.993	0.923	0.151	0.010	1.000	1.000	0.221	0.009								
Queensland																	
GARCH	MSM	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000								
GJR		1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000								
EGARCH		0.965	0.870	0.907	0.920	1.000	1.000	1.000	1.000								
APARCH		0.985	1.000	1.000	1.000	1.000	1.000	1.000	1.000								
FIGARCH		0.939	0.819	0.771	0.921	0.998	1.000	1.000	1.000								
MSGARCH		0.963	1.000	1.000	1.000	1.000	1.000	1.000	1.000								
South Australia																	
GARCH	MSM	1.000	0.996	1.000	1.000	1.000	1.000	1.000	1.000								
GJR		0.990	0.908	0.999	0.997	1.000	1.000	0.998	1.000								
EGARCH		0.988	0.399	0.214	0.073	1.000	0.707	0.380	0.240								
APARCH		0.956	0.171	0.155	0.156	0.994	0.114	0.089	0.120								
FIGARCH		0.936	0.114	0.491	0.335	0.000	0.081	0.578	0.771								
MSGARCH		0.982	1.000	0.447	0.047	1.000	1.000	0.058	0.020								
Tasmania																	
GARCH	MSM	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000								
GJR		1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000								
EGARCH		1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000								
APARCH		1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000								
FIGARCH		0.000	0.000	0.000	0.012	0.000	0.000	0.001	0.024								
MSGARCH		1.000	1.000	1.000	0.061	1.000	1.000	1.000	0.056								
Victoria																	
GARCH	MSM	1.000	0.998	1.000	1.000	1.000	1.000	1.000	1.000								
GJR		0.998	0.914	1.000	1.000	1.000	1.000	1.000	1.000								
EGARCH		0.728	0.480	0.398	0.257	0.499	0.993	1.000	0.992								
APARCH		0.897	0.952	0.158	0.159	1.000	1.000	0.147	0.155								
FIGARCH		0.997	0.968	0.933	0.931	0.993	1.000	1.000	0.998								
MSGARCH		1.000	1.000	0.999	0.036	1.000	1.000	0.162	0.032								

Note: p -values obtained for the null hypothesis that, for a given forecast horizon, there is no difference in forecast accuracy between Model 1 and Model 2 versus the one-sided alternative that the forecasts from Model 1 are inferior to those from Model 2. The volatility models are abbreviated as MSM (Markov-switching multifractal), GARCH (generalized autoregressive conditional heteroscedasticity), GJR (Glosten-Jagannathan-Runkle GARCH), EGARCH (exponential GARCH), APARCH (asymmetric power ARCH), FIGARCH (fractionally integrated GARCH), MSGARCH (Markov-switching GARCH).

Table 6: Superior predictive ability (SPA) tests (p -values)

Benchmark model	Forecast horizon (trading days)							
	1				5			
	1	5	10	20	1	5	10	20
	Squared error loss				Absolute error loss			
	New South Wales							
MSM	0.040	0.001	0.000	0.000	0.000	0.000	0.000	0.000
GARCH	0.155	0.830	1.000	1.000	0.057	1.000	1.000	1.000
GJR	0.912	0.004	0.000	0.000	0.630	0.000	0.000	0.000
EGARCH	0.085	0.000	0.000	0.000	0.044	0.000	0.000	0.000
APARCH	0.053	0.002	0.001	0.000	0.000	0.000	0.000	0.000
FIGARCH	0.088	0.170	0.054	0.001	0.019	0.000	0.000	0.000
MSGARCH	0.072	0.010	0.000	0.004	0.450	0.001	0.000	0.000
	Queensland							
MSM	0.019	0.000	0.000	0.000	0.000	0.000	0.000	0.000
GARCH	0.215	0.878	1.000	1.000	0.004	1.000	1.000	1.000
GJR	0.981	0.000	0.000	0.000	0.007	0.000	0.000	0.000
EGARCH	0.207	0.000	0.000	0.000	0.000	0.000	0.000	0.000
APARCH	0.034	0.000	0.001	0.000	0.000	0.000	0.000	0.000
FIGARCH	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MSGARCH	0.085	0.156	0.003	0.000	1.000	0.000	0.000	0.000
	South Australia							
MSM	0.023	0.000	0.000	0.000	0.000	0.000	0.000	0.000
GARCH	0.143	0.674	0.927	0.870	0.001	0.000	1.000	1.000
GJR	1.000	0.863	0.000	0.000	0.255	0.000	0.000	0.000
EGARCH	0.081	0.000	0.000	0.000	0.000	0.000	0.000	0.000
APARCH	0.117	0.180	0.113	0.130	0.001	0.060	0.055	0.084
FIGARCH	0.044	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MSGARCH	0.112	0.656	0.552	0.015	0.745	1.000	0.003	0.001
	Tasmania							
MSM	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
GARCH	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.000
GJR	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
EGARCH	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
APARCH	0.000	0.000	0.000	1.000	0.000	0.000	0.000	1.000
FIGARCH	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MSGARCH	1.000	1.000	1.000	0.065	1.000	1.000	1.000	0.041
	Victoria							
MSM	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
GARCH	0.165	0.777	0.555	0.742	0.000	0.000	0.906	0.872
GJR	0.253	0.733	0.000	0.000	0.001	0.000	0.000	0.000
EGARCH	0.748	0.000	0.000	0.000	0.362	0.000	0.000	0.000
APARCH	0.060	0.000	0.191	0.258	0.000	0.000	0.096	0.128
FIGARCH	0.007	0.000	0.000	0.000	0.000	0.000	0.000	0.000
MSGARCH	0.182	0.329	0.923	0.000	0.645	1.000	0.000	0.000

Note: p -values obtained for the null hypothesis that the benchmark model is not inferior to any of the other competing models. The volatility models are abbreviated as MSM (Markov-switching multifractal), GARCH (generalized autoregressive conditional heteroscedasticity), GJR (Glosten-Jagannathan-Runkle GARCH), EGARCH (exponential GARCH), APARCH (asymmetric power ARCH), FIGARCH (fractionally integrated GARCH), MSGARCH (Markov-switching GARCH).




Figure 1: Daily electricity price returns (left panels) and autocorrelation functions (right panels)




Figure 2: Daily realized variances (left panels) and autocorrelation functions (right panels)

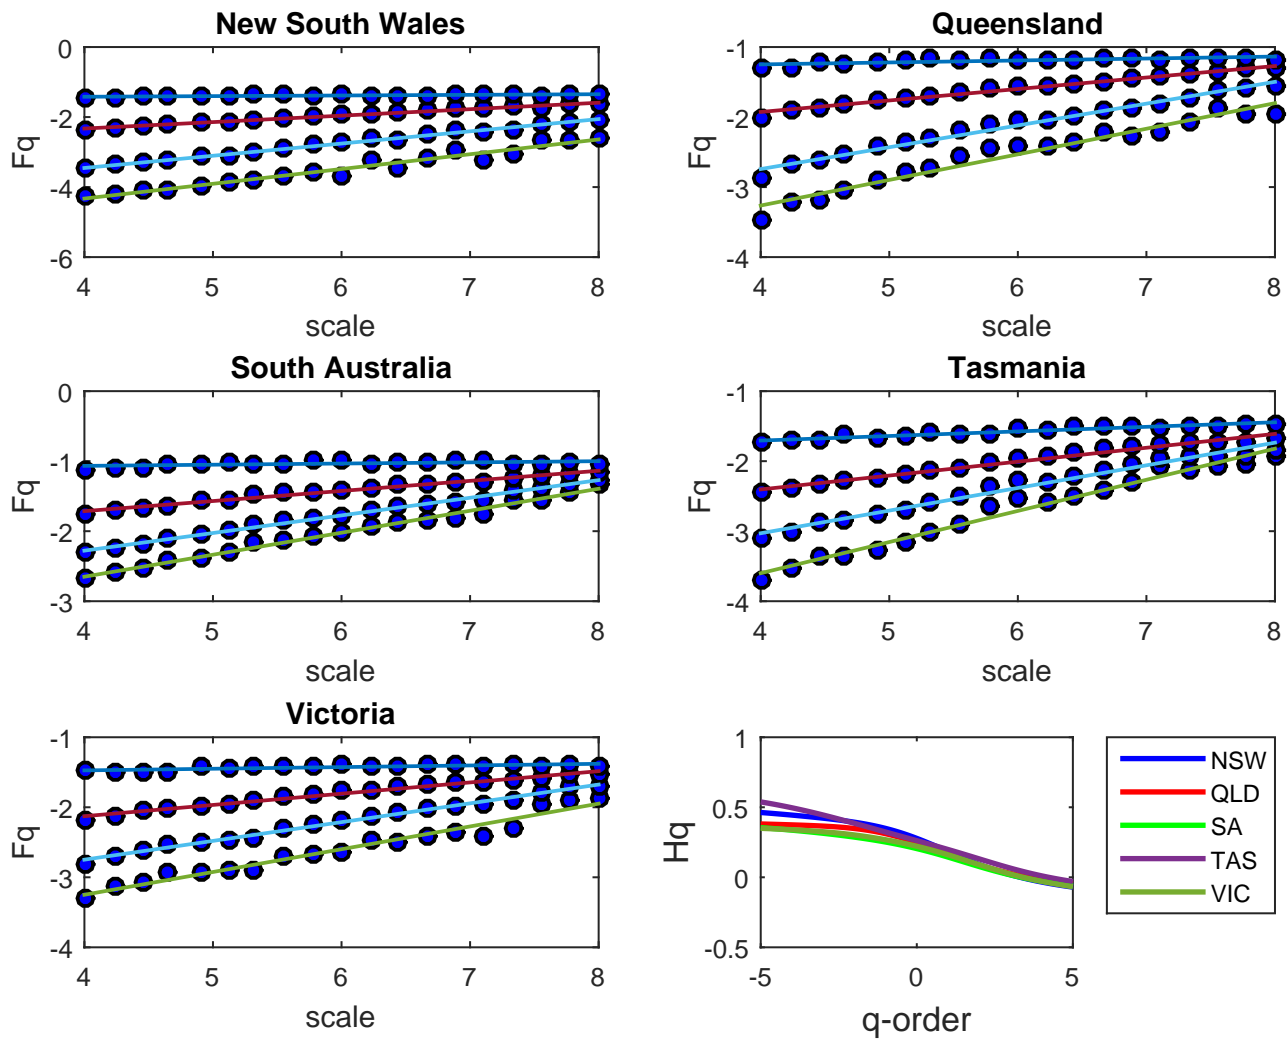


Figure 3: The scaling functions (F_q) and the q-order Hurst exponents (H_q) for the electricity price returns




Figure 4: Multifractal spectra of the electricity price returns