MDPI

Article

Forecasting Volatility and Tail Risk in Electricity Markets

Antonio Naimoli [ID] and Giuseppe Storti *[ID]

Dipartimento di Scienze Economiche e Statistiche (DISES), Università di Salerno, Via Giovannni Paolo II, 132, 84084 Fisciano, Italy; anaimoli@unisa.it
* Correspondence: storti@unisa.it

Abstract: This paper investigates the benefits of jointly using several realized measures in predicting daily price volatility, Value-at-Risk, and Expected Shortfall in the Australian electricity markets of New South Wales, Queensland, and Victoria. We propose using Realized GARCH-type models with multiple measurement equations based on robust estimators to account for market microstructure noise and jumps in electricity price series. The model specifications that combine information from multiple realized measures improve the in-sample fit of the data. The out-of-sample analysis shows that use of the jump-robust *medRV* estimator significantly increases the accuracy of volatility forecasts, while in forecasting Value-at-Risk and Expected Shortfall at different risk levels, the standard GARCH(1,1) also performs remarkably well.

Keywords: volatility forecasting; value-at-risk; expected shortfall; realized GARCH; electricity prices

---

1. Introduction

The analysis of volatility of electricity spot prices is crucial for traders, portfolio managers, policy makers, and other market participants. The growing interest in modeling the dynamics of electricity prices has revealed several distinctive features typically not observed in financial assets due to the nonstorability of electricity. As electricity is not storable, and because of the inelasticity of supply and demand, electricity prices are known to be much more volatile than other commodity prices. In particular, the main stylized facts show that daily and intradaily electricity spot prices are usually characterized by seasonality, mean reversion, high volatility persistence, frequent price jumps and spikes of short duration, inverse leverage effects (electricity price volatility reacts more to positive shocks than to negative shocks), stationarity in both the price level and squared prices, and negative prices that are mainly related to the inability to dispose of electricity freely together with nontrivial start-up costs for generators (Bierbrauer et al. 2007; Byström 2005; Chan et al. 2008; Escribano et al. 2011; Fanone et al. 2013; Frömmel et al. 2014; Higgs and Worthington 2008; Knittel and Roberts 2005).

The popular GARCH-type framework is widely used to model and forecast the volatility of electricity prices (e.g., Bowden and Payne (2008); Escribano et al. (2011); Garcia et al. (2005); Hickey et al. (2012); Knittel and Roberts (2005); Liu and Shi (2013), among others). Standard GARCH models usually rely on daily data to estimate the latent conditional variance, using all current and past daily squared returns to provide expectations on future volatility. However, the daily return offers only a "weak" signal on the current level of volatility; therefore, this class of model is not able to capture rapid changes in the volatility level.

Since electricity cannot be physically stored directly, production and consumption need to be continuously balanced to smooth supply and demand shocks (Bierbrauer et al. 2007). In this direction, the liquidity of the electricity market has grown rapidly. The increased availability of high-frequency information has led to the development of new econometric methods for modeling and forecasting volatility, revealing that high-frequency data are much more informative about the price process not only at the intraday level, but also at

the daily level. The daily Realized Volatility (RV), defined as the summation of the squared intradaily price changes, provides an unbiased and highly efficient estimator of return volatility (Andersen et al. 2001, 2003; Barndorff-Nielsen and Shephard 2002).

Volatility modeling using intradaily price frequencies has received considerable attention, not only in the financial market but also in the electricity market. In this framework, most of the existing literature relies on HAR-type models to directly estimate time series of realized measures. Chan et al. (2008) used the HAR model of Corsi (2009) and the HAR-CJ model of Andersen et al. (2007) to estimate volatility and identify jumps in the electricity price process on five regions of the Australian market. However, their results show a modest improvement in volatility forecasts when total variation is separated into continuous and jump components, with no strong evidence that the HAR-type specifications outperform the EGARCH model. Haugom et al. (2011) added exogenous effects on the HAR model and the HAR-CJ model to assess the day-ahead predictions in the Nord Pool forward market, finding forecast improvements from the inclusion of exogenous effects. Haugom and Ullrich (2012) extended the HAR approach, including forward realized volatility as a predictor to improve spot price volatility forecasts for some electricity markets of the United States. Differently, focusing on the conditional variance of returns, Frömmel et al. (2014) referred to the Realized GARCH models (Hansen and Huang 2016; Hansen et al. 2012) with a single measurement equation to forecast the volatility on the Electricity Power Exchange market using both the RV and the intraday range as realized measures. The empirical results suggest that the RGARCH specifications outperform the EGARCH model in terms of forecasting accuracy, especially when intraday range is used, stating (among other potential determinants) that range-based measurements are more robust to microstructure noise bias.

The aim of this paper is to assess the benefits of jointly using different realized measures in fitting and forecasting electricity price volatility. Differently to Frömmel et al. (2014), we resort to the use of Realized GARCH models to combine information through multiple measurement equations and multiple realized estimators. In particular, to deal with market microstructure frictions and extreme jumps in electricity price series, in addition to the RV, we also refer to the robust estimators *Realized Kernel* (RK) (Barndorff-Nielsen et al. 2008) and *medRV* (MRV) (Andersen et al. 2012). Moreover, since accurate volatility modeling is crucial for risk management (Byström 2005; Chan and Gray 2006), we also focus on the ability of models to accurately predict Value-at-Risk (VaR) and Expected Shortfall (ES) at different risk levels.

Our empirical analysis on time series of spot prices sampled at 30-min intervals in the regional Australian power markets of New South Wales, Queensland, and Victoria reveals some interesting findings. First, the Realized Exponential GARCH (REGARCH) specifications, combining multiple realized measures, improve the in-sample fit over the standard GARCH(1,1) and the simple Realized GARCH (RGARCH) models. Second, the use of jump-robust MRV as a realized measure in the RGARCH model leads to a significant improvement in volatility forecasting by minimizing the QLIKE (Patton 2011), Mean Squared Error, and Mean Absolute Error loss functions. Finally, in evaluating the performances in forecasting VaR and ES at the risk levels of 1%, 2.5%, and 5% via the Quantile Loss function of (González-Rivera et al. 2004) and the class of strictly consistent loss functions (Fissler and Ziegel 2016), respectively, the GARCH(1,1) and some REGARCH specifications clearly outperform the simple RGARCH based on a single measurement equation. The Model Confidence Set (MCS) of Hansen et al. (2011) is used to assess the significance of differences in the predictive performances of the models under analysis.

The remainder of the paper is structured as follows: Section 2 reviews the Realized GARCH (Hansen et al. 2012) and the Realized Exponential GARCH (Hansen and Huang 2016). Section 3 presents the estimation procedure. Section 4 describes the data. Sections 5 and 6 show the in-sample and the out-of-sample results, respectively. Finally, Section 7 concludes with a summary of the results and some directions for future research.

2. Model Specifications

The GARCH model is by far the most widely used specification for fitting and forecasting financial volatility. Let r_t be the daily log-return at time t and $\mathcal{I}_{t-1}$ the information available up to time $t-1$; then, the GARCH(1,1) model takes the following form:

$$r_t = \sqrt{h_t} z_t \tag{1}$$

$$h_t = \omega + \beta \, h_{t-1} + \gamma \, r_{t-1}^2, \tag{2}$$

where $h_t = var(r_t | \mathcal{I}_{t-1})$ and $z_t \overset{iid}{\sim} (0,1)$. The positivity condition for the conditional variance requires γ and β to be non-negative constants and ω to be a (strictly) positive constant, while a necessary condition for the weak stationarity of GARCH(1,1) is $\gamma + \beta < 1$.

Many studies have documented that realized volatility measures based on intradaily returns can greatly improve the accuracy of volatility forecasts. Differently from standard GARCH approach, the Realized GARCH (RGARCH) of Hansen et al. (2012) employs both low (daily returns) and high (realized volatility measures) frequency information to model the dynamics of daily volatility. Following the log-linear specification, the RGARCH is given by

$$r_t = \sqrt{h_t} \, z_t \tag{3}$$

$$log(h_t) = \omega + \beta \, log(h_{t-1}) + \gamma \, log(x_{t-1}) \tag{4}$$

$$log(x_t) = \xi + \varphi \, log(h_t) + \tau(z_t) + u_t, \tag{5}$$

where x_t is a realized volatility measure and $\tau(z_t) = \tau_1 z_t + \tau_2(z_t^2 - 1)$ is the leverage function, with $z_t \overset{iid}{\sim} (0,1)$ and $u_t \overset{iid}{\sim} (0, \sigma_u^2)$ being mutually independent.

Therefore, the RGARCH provides a joint modeling framework of the return and realized volatility, replacing squared returns with a more informative volatility estimator to capture the conditional variance dynamics. The model is completed by the measurement equation in (5), which allows us to define the link between the (ex-post) realized measure and the (ex-ante) latent conditional variance. In addition, replacing the measurement equation into the GARCH equation, the model implies an AR(1) representation for the log-conditional variance, namely,

$$log(h_t) = (\omega + \xi\gamma) + (\beta + \varphi\gamma) log(h_{t-1}) + \gamma \, w_{t-1}, \tag{6}$$

where $w_t = \tau(z_t) + u_t$ and $E(w_t) = 0$, with the restriction $(\beta + \varphi\gamma) < 1$ to ensure the stationarity of the process (Hansen et al. 2012; Li et al. 2019).

Hansen and Huang (2016) proposed the Realized Exponential GARCH (REGARCH) model while allowing the inclusion of multiple realized measures of volatility and also including an explicit leverage term in the GARCH equation, providing further flexibility in modeling the dependence between returns and volatility.

In its general formulation, considering a vector of K realized measures $\boldsymbol{x}_t = (x_{1,t}, \cdots, x_{K,t})'$, the REGARCH is defined as

$$r_t = \sqrt{h_t} \, z_t \tag{7}$$

$$log(h_t) = \omega + \beta \, log(h_{t-1}) + \delta(z_{t-1}) + \boldsymbol{\gamma}' \, \boldsymbol{u}_{t-1} \tag{8}$$

$$log(x_{j,t}) = \xi_j + \varphi_j \, log(h_t) + \tau_j(z_t) + u_{j,t} \qquad j = 1, ..., K, \tag{9}$$

where $\boldsymbol{u}_t = (u_{1,t}, ..., u_{K,t})'$ and $\boldsymbol{\gamma} = (\gamma_1, ..., \gamma_K)'$. In the empirical analysis, as usual, a quadratic form for the leverage function $\delta(z_t)$—that is, $\delta(z_t) = \delta_1 z_t + \delta_2(z_t^2 - 1)$, is used.

3. Estimation and Inference

The GARCH and R(E)GARCH models are estimated using a maximum likelihood (ML) approach. In particular, we assume a standardized Student-t distribution for innovations $z_t \overset{iid}{\sim} t(0,1,\nu)$ to model both fat-tail and excess kurtosis observed in return series.

For REGARCH specifications, due to the presence of K measurement equations, following [Hansen and Huang (2016)](#), we assume that $\boldsymbol{u}_t \overset{iid}{\sim} \mathcal{N}_K(\boldsymbol{0}, \Sigma)$, where $\mathcal{N}_K(\boldsymbol{0}, \Sigma)$ denotes a K-variate Normal distribution with mean $\boldsymbol{0}$ and variance-covariance matrix Σ. Therefore, the log-likelihood function for this class of models is given by

$$\mathcal{L}(r, u; \boldsymbol{\theta}) = -\frac{1}{2} \sum_{t=1}^{T} -\mathcal{A}(\nu) + log(h_t) + (1+\nu) \, log\left(1 + \frac{r_t^2}{h_t(\nu-2)}\right) \tag{10}$$

$$-\frac{1}{2} \sum_{t=1}^{T} Klog(2\pi) + log(|\Sigma|) + \boldsymbol{u}_t' \Sigma^{-1} \boldsymbol{u}_t \, , \tag{11}$$

where $\mathcal{A}(\nu) = log\left[\Gamma\left(\frac{\nu+1}{2}\right)\right] - log\left[\Gamma\left(\frac{\nu}{2}\right)\right] - \frac{1}{2} log[\pi(\nu-2)]$. The log-likelihood accounts for the contribution of realized measures by (11) and the contribution of returns by (10).

On the other hand, as the RGARCH is characterized by a single realized measure, Equation (11) reduces to

$$-\frac{1}{2} \sum_{t=1}^{T} log(2\pi) + log(\sigma_u^2) + \frac{u_t^2}{\sigma_u^2}, \tag{12}$$

with $u_t \overset{iid}{\sim} \mathcal{N}(0, \sigma_u^2)$.

Finally, as standard GARCH models do not allow for measurement equations, estimation of model parameters is performed by focusing only on the partial log-likelihood of returns in Equation (10).

It is worth noting that the GARCH and R(E)GARCH specifications are not directly comparable in terms of the maximized global log-likelihood. However, since the contribution to the log-likelihood value of returns is the same for both classes of models, the partial log-likelihood of the returns in (10) enables us to compare the empirical fit of the conventional GARCH with that of R(E)GARCH-type models.

4. The Data

The dataset used in this study consists of half-hourly spot prices from the Australian electricity market. In Australia, the Australian Energy Market Operator (AEMO) manages the National Energy Market (NEM), which interconnects Queensland (QLD), New South Wales (NSW), the Australian Capital Territory (ACT), South Australia (SA), Victoria (VIC), and Tasmania (TAS), as well as the Wholesale Electricity Market (WEM) in Western Australia. In particular, the empirical analysis focuses on 30-min spot prices of the NSW, QLD, and VIC for the period between 1 January 2012 and 31 December 2019, where the prices are in Australian dollars per megawatt hour. Continuously recorded half-hourly spot prices are publicly available at https://www.aemo.com.au (accessed on 12 May 2021).

The summary statistics reported in Table 1 highlight some stylized facts of price and return dynamics. Differently to financial assets, electricity prices can be negative when the supply of electricity temporarily exceeds demand, resulting in a minimum half-hourly intradaily price lower than zero for the three observed markets. A further characteristic of the electricity prices is that the maximum price is more than 200 times larger than the average price, implying high variability and a high degree of skewness and kurtosis. These features can easily be seen in Figure 1, which shows the time series of 30-min spot prices in the considered Australian electricity markets. Each market exhibits negative prices and several spikes with a magnitude much higher than their average price. This emphasizes a further characteristic of electricity prices: the existence of extreme jumps.

As negative prices are a rare occurrence, to compute log-returns, we excluded days with nonpositive prices (Haugom and Ullrich 2012; Qu et al. 2018). Further, to account for intraday seasonal patterns in the raw data, following Haugom and Ullrich (2012), returns are demeaned using half-hourly median returns $\mu^*_{t,j} = \bar{r}_{mn,dy,hr}$, where $\bar{r}_{mn,dy,hr}$ denotes the median return for day t in month mn, on day of the week dy, and in half-hour periods $j = hr$. Therefore, in our empirical study, we used demeaned intraday returns $r^*_{t,j} = r_{t,j} - \mu^*_{t,j}$ for computing realized volatility measures.

On the other hand, we refer to daily close-to-close log-returns as daily price changes. Although the mean and median are close to zero, electricity returns are highly volatile, as evident from the standard deviation, which varies between 0.216 and 0.314. In addition, the negative skewness and excess kurtosis clearly reveal the non-Gaussian nature of the r_t distribution.
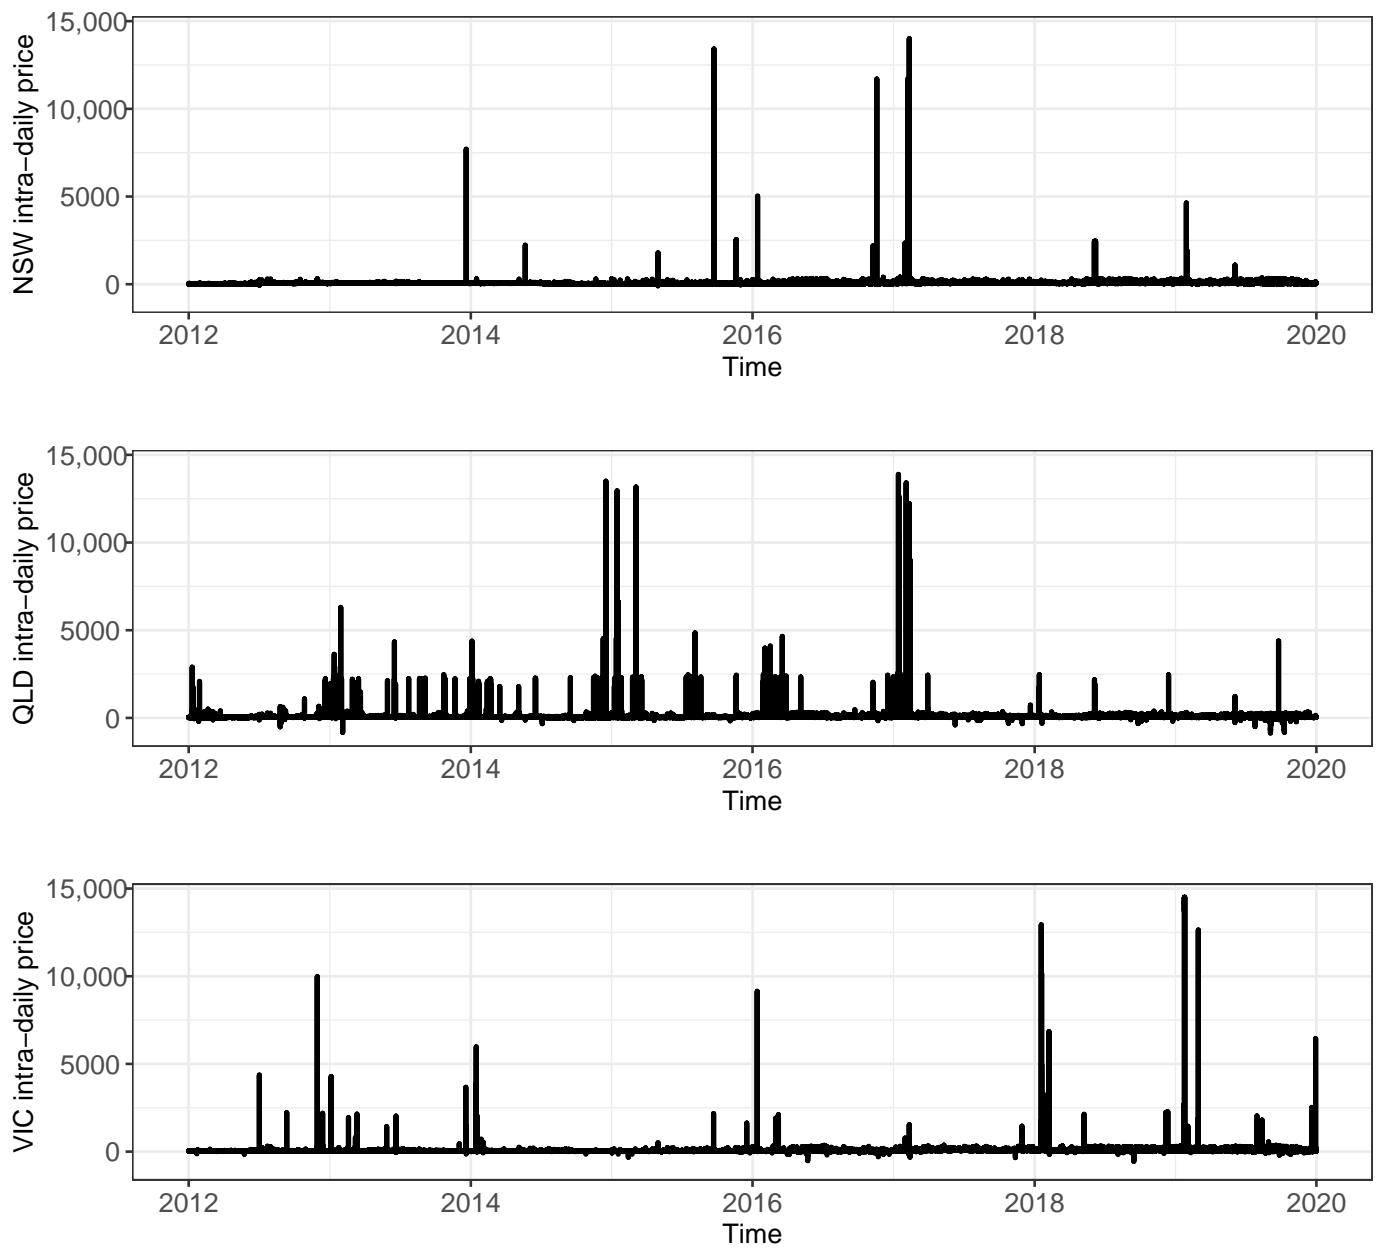
Finally, descriptive statistics point out that the 30-min RV is, as expected, characterized by positive skewness and a strong excess kurtosis because of the many peaks and troughs in the series. This can easily be seen in Figure 2, which shows the RV behavior for the three electricity markets, where each plot has the same y-axis scale to facilitate comparison across markets. The QLD market reveals several periods of high volatility and the most extreme levels of RV compared to NSW and VIC, which instead show more moderate and less recurrent spikes.

As price volatility is characterized by different dynamics under different market conditions, in addition to the *Realized Volatility* (RV) (Andersen et al. 2003), to account for market microstructure noise and jumps in our empirical analysis, we also referred to the robust estimators *Realized Kernel* (RK) (Barndorff-Nielsen et al. 2008) and *medRV* (MRV) (Andersen et al. 2012), computed at a frequency of 30 min.

Table 1. Summary statistics.

	NSW									
	n	mean	sd	min	$Q_{0.25}$	median	$Q_{0.75}$	max	skew	kurt
Raw 30-min intradaily prices	140,256	62.557	107.055	−89.210	39.368	52.870	72.850	14,000	88.971	10,006.338
Positive 30-min intradaily prices	139,824	62.559	107.099	1.530	39.390	52.870	72.850	14,000	89.107	10,020.409
Daily close-to-close log-returns	2913	0.000	0.216	−1.695	−0.074	0.000	0.073	1.251	−0.181	8.363
30-min Realized Volatility	2913	1.187	2.688	0.017	0.156	0.440	1.276	53.502	9.544	133.944
	QLD									
	n	mean	sd	min	$Q_{0.25}$	median	$Q_{0.75}$	max	skew	kurt
Raw 30-min intradaily prices	140,256	66.314	190.754	−859.850	37.950	53.010	69.140	13,882.77	39.978	2179.302
Positive 30-min intradaily prices	135,024	66.698	192.134	0.170	38.310	53.030	69.020	13,882.77	40.298	2192.214
Daily close-to-close log-returns	2813	0.000	0.258	−2.000	−0.081	−0.002	0.089	2.002	-0.227	12.238
30-min Realized Volatility	2813	3.558	10.313	0.021	0.245	0.718	2.102	151.031	6.389	55.106
	VIC									
	n	mean	sd	min	$Q_{0.25}$	median	$Q_{0.75}$	max	skew	kurt
Raw 30-min intradaily prices	140,256	64.012	175.75	−554.62	34.860	49.830	79.730	14,500	63.019	4603.573
Positive 30-min intradaily prices	133,824	65.056	179.333	0.150	35.760	50.310	80.610	14,500	62.123	4448.796
Daily close-to-close log-returns	2788	0.000	0.314	−2.118	−0.130	0.001	0.127	1.837	−0.140	4.774
30-min Realized Volatility	2788	1.858	3.921	0.022	0.248	0.746	1.865	62.105	6.632	60.529

Summary statistics of raw 30-min intradaily prices, positive 30-min intradaily prices (excluding days with negative prices), daily close-to-close log-returns and 30-min Realized Volatility for NSW (top), QLD (middle), and VIC (bottom). *n*—number of observations; sd—standard deviation; min—minimum; $Q_{0.25}$—first quartile; $Q_{0.75}$—third quartile; max—maximum; skew—skewness; kurt—excess kurtosis. The sample period is from 1 January 2012 to 31 December 2019.

Figure 1. 30-min intradaily prices. Time series of 30-min electricity spot prices for NSW (**top**), QLD (**middle**), and VIC (**bottom**). The sample period is from 1 January 2012 to 31 December 2019.

Figure 2. 30-min Realized Volatility. Time series of 30-min Realized Volatility for NSW (**top**), QLD (**middle**), and VIC (**bottom**). The sample period is from 1 January 2012 to 31 December 2019.

5. In-Sample Analysis

In this section, we present the in-sample results for NSW in Table 2, QLD in Table 3, and VIC in Table 4. For each electricity market, we report the in-sample fit of the GARCH(1,1); three RGARCH and REGARCH models based on the single realized estimators RV, RK, and MRV; three REGARCH models with two measurement equations for the pairs (RV,RK), (RV,MRV), and (RK,MRV); and the REGARCH(RV,RK,MRV) combining all three realized estimators considered. Note that to simplify the presentation of the results, the RGARCH models are estimated using the autoregressive representation in order to make the estimated coefficients comparable with those of the REGARCH models.

Table 2. NSW in-sample parameter estimates.

	GARCH(1,1)	RG(RV)	RG(RK)	RG(MRV)	REG(RV)	REG(RK)	REG(MRV)	REG(RV,RK)	REG(RV,MRV)	REG(RK,MRV)	REG(RV,RK,MRV)
ω	0.000	−0.069	−0.069	−0.058	−0.061	−0.061	−0.054	−0.060	−0.063	−0.063	−0.062
β	0.863	0.981	0.981	0.984	0.987	0.986	0.988	0.987	0.986	0.986	0.986
γ_1	0.137	0.320	0.321	0.283	0.234	0.235	0.205	0.688	0.220	0.198	0.669
γ_2	–	–	–	–	–	–	–	−0.457	**0.018**	**0.039**	−0.449
γ_3	–	–	–	–	–	–	–	–	–	–	**0.014**
τ_1	–	0.142	0.144	0.148	0.200	0.202	0.205	0.202	0.197	0.199	0.199
τ_2	–	0.057	0.058	0.058	0.049	0.049	0.046	0.047	0.050	0.050	0.048
ξ_1	–	2.362	2.355	2.277	2.476	2.471	2.403	2.477	2.486	2.480	2.486
φ_1	–	0.864	0.864	0.862	0.873	0.874	0.873	0.872	0.876	0.876	0.875
$\delta_{1,1}$	–	–	–	–	0.210	0.213	0.218	0.211	0.209	0.212	0.210
$\delta_{2,1}$	–	–	–	–	0.026	0.025	0.024	0.025	0.026	0.025	0.025
ξ_2	–	–	–	–	–	–	–	2.472	2.426	2.424	2.482
φ_2	–	–	–	–	–	–	–	0.872	0.877	0.878	0.874
$\delta_{1,2}$	–	–	–	–	–	–	–	0.213	0.217	0.218	0.212
$\delta_{2,2}$	–	–	–	–	–	–	–	0.024	0.023	0.023	0.024
ξ_3	–	–	–	–	–	–	–	–	–	–	2.427
φ_3	–	–	–	–	–	–	–	–	–	–	0.876
$\delta_{1,3}$	–	–	–	–	–	–	–	–	–	–	0.218
$\delta_{2,3}$	–	–	–	–	–	–	–	–	–	–	0.022
$\sigma_{u,1}^2$	–	0.740	0.740	0.812	0.724	0.723	0.794	0.724	0.724	0.723	0.723
$\sigma_{u,2}^2$	–	–	–	–	–	–	–	0.722	0.792	0.793	0.722
$\sigma_{u,3}^2$	–	–	–	–	–	–	–	–	–	–	0.792
$\rho_{1,2}$	–	–	–	–	–	–	–	0.998	0.963	0.962	0.998
$\rho_{1,3}$	–	–	–	–	–	–	–	–	–	–	0.963
$\rho_{2,3}$	–	–	–	–	–	–	–	–	–	–	0.962
ν	5.370	3.766	3.731	3.686	4.289	4.258	4.225	4.330	4.302	4.269	4.341
$\ell(r)$	1643.532	1591.066	1588.091	1584.072	1657.235	1655.619	1655.049	1659.573	1656.897	1655.252	1659.227

Parameter estimates for the sample period 1 January 2012 to 31 December 2019. Coefficients that are not significant at the 5% level are shown in bold. Key to table: RG—Realized GARCH; REG—Realized Exponential GARCH.

Table 3. QLD in-sample parameter estimates.

	GARCH(1,1)	RG(RV)	RG(RK)	RG(MRV)	REG(RV)	REG(RK)	REG(MRV)	REG(RV,RK)	REG(RV,MRV)	REG(RK,MRV)	REG(RV,RK,MRV)
ω	0.002	−0.276	−0.273	−0.277	−0.319	−0.313	−0.336	−0.318	−0.301	−0.299	−0.297
β	0.669	0.910	0.911	0.906	0.904	0.906	0.897	0.904	0.910	0.910	0.912
γ_1	0.331	0.364	0.362	0.338	0.314	0.310	0.297	0.955	0.628	0.581	1.255
γ_2	−	−	−	−	−	−	−	−0.642	−0.316	−0.271	**−0.628**
γ_3	−	−	−	−	−	−	−	−	−	−	−0.316
τ_1	−	0.184	0.185	0.191	0.196	0.196	0.194	0.196	0.199	0.200	0.199
τ_2	−	0.070	0.070	0.066	0.080	0.080	0.088	0.078	0.073	0.074	0.071
ξ_1	−	2.886	2.878	2.817	3.103	3.090	3.041	3.096	3.104	3.100	3.094
φ_1	−	1.011	1.012	1.035	1.024	1.021	1.028	1.021	1.019	1.020	1.015
$\delta_{1,1}$	−	−	−	−	0.245	0.245	0.256	0.245	0.245	0.245	0.246
$\delta_{2,1}$	−	−	−	−	0.031	0.032	**0.019**	0.031	0.029	0.030	0.029
ξ_2	−	−	−	−	−	−	−	3.086	3.054	3.059	3.085
φ_2	−	−	−	−	−	−	−	1.018	1.019	1.022	1.013
$\delta_{1,2}$	−	−	−	−	−	−	−	0.246	0.256	0.255	0.246
$\delta_{2,2}$	−	−	−	−	−	−	−	0.031	**0.014**	**0.015**	0.029
ξ_3	−	−	−	−	−	−	−	−	−	−	3.044
φ_3	−	−	−	−	−	−	−	−	−	−	1.015
$\delta_{1,3}$	−	−	−	−	−	−	−	−	−	−	0.256
$\delta_{2,3}$	−	−	−	−	−	−	−	−	−	−	**0.014**
$\sigma_{u,1}^2$	−	1.499	1.497	1.689	1.488	1.488	1.675	1.489	1.486	1.485	1.486
$\sigma_{u,2}^2$	−	−	−	−	−	−	−	1.488	1.656	1.657	1.484
$\sigma_{u,3}^2$	−	−	−	−	−	−	−	−	−	−	1.657
$\rho_{1,2}$	−	−	−	−	−	−	−	0.999	0.982	0.981	0.999
$\rho_{1,3}$	−	−	−	−	−	−	−	−	−	−	0.982
$\rho_{2,3}$	−	−	−	−	−	−	−	−	−	−	0.981
ν	4.041	3.015	3.001	2.850	3.497	3.487	3.405	3.509	3.555	3.533	3.567
$\ell(r)$	995.431	922.376	920.371	899.299	996.507	995.206	987.042	998.823	1004.208	1000.576	1006.403

Parameter estimates for the sample period 1 January 2012 to 31 December 2019. Coefficients that are not significant at the 5% level are shown in bold. Key to table: RG—Realized GARCH; REG—Realized Exponential GARCH.

Table 4. VIC in-sample parameter estimates.

	GARCH(1,1)	RG(RV)	RG(RK)	RG(MRV)	REG(RV)	REG(RK)	REG(MRV)	REG(RV,RK)	REG(RV,MRV)	REG(RK,MRV)	REG(RV,RK,MRV)
ω	0.000	−0.035	−0.034	−0.026	−0.038	−0.039	−0.028	−0.039	−0.040	−0.040	−0.041
β	0.874	0.987	0.987	0.990	0.987	0.987	0.990	0.987	0.986	0.986	0.986
γ_1	0.126	0.208	0.207	0.173	0.201	0.202	0.165	0.404	0.311	0.289	0.538
γ_2	–	–	–	–	–	–	–	**−0.201**	−0.106	−0.085	**−0.222**
γ_3	–	–	–	–	–	–	–	–	–	–	−0.110
τ_1	–	0.095	0.097	0.097	0.106	0.105	0.107	0.106	0.107	0.106	0.107
τ_2	–	0.082	0.082	0.079	0.029	0.031	0.024	0.030	0.030	0.031	0.030
ξ_1	–	2.325	2.319	2.203	2.359	2.356	2.239	2.355	2.358	2.357	2.353
φ_1	–	1.011	1.009	1.008	1.011	1.010	1.008	1.009	1.009	1.009	1.006
$\delta_{1,1}$	–	–	–	–	0.126	0.128	0.129	0.126	0.127	0.129	0.127
$\delta_{2,1}$	–	–	–	–	0.076	0.076	0.073	0.075	0.075	0.075	0.075
ξ_2	–	–	–	–	–	–	–	2.349	2.224	2.229	2.347
φ_2	–	–	–	–	–	–	–	1.006	0.994	0.996	1.004
$\delta_{1,2}$	–	–	–	–	–	–	–	0.128	0.131	0.131	0.129
$\delta_{2,2}$	–	–	–	–	–	–	–	0.075	0.072	0.072	0.074
ξ_3	–	–	–	–	–	–	–	–	–	–	2.220
φ_3	–	–	–	–	–	–	–	–	–	–	0.991
$\delta_{1,3}$	–	–	–	–	–	–	–	–	–	–	0.131
$\delta_{2,3}$	–	–	–	–	–	–	–	–	–	–	0.071
$\sigma_{u,1}^2$	–	0.885	0.887	0.976	0.885	0.887	0.976	0.885	0.886	0.888	0.886
$\sigma_{u,2}^2$	–	–	–	–	–	–	–	0.887	0.979	0.980	0.887
$\sigma_{u,3}^2$	–	–	–	–	–	–	–	–	–	–	0.979
$\rho_{1,2}$	–	–	–	–	–	–	–	0.998	0.965	0.964	0.998
$\rho_{1,3}$	–	–	–	–	–	–	–	–	–	–	0.965
$\rho_{2,3}$	–	–	–	–	–	–	–	–	–	–	0.964
ν	5.098	4.176	4.157	4.058	4.448	4.441	4.312	4.473	4.464	4.448	4.488
$\ell(r)$	23.397	36.249	34.943	21.483	54.123	53.602	40.085	55.446	57.727	56.104	59.150

Parameter estimates for the sample period 1 January 2012 to 31 December 2019. Coefficients that are not significant at the 5% level are shown in bold. Key to table: RG—Realized GARCH; REG—Realized Exponential GARCH.

The empirical results show that the β estimates of the GARCH(1,1) for NSW, QLD, and VIC are 0.863, 0.669, and 0.874, respectively, confirming the fact that QLD is the most-nervous electricity market. This aspect is also highlighted by the estimate of γ_1 taking the highest value for QLD. Focusing on the volatility persistence, for the GARCH(1,1), we observe $(\beta + \gamma_1) \approx 1$; whereas, for the R(E)GARCH-type models, β (the parameter which summarizes the persistence of volatility for this class of models) is, on average, about 0.96. Thus, volatility decay is faster for models including high-frequency information. Moreover, a larger γ_1 for individual realized measures in R(E)GARCHs with respect to the standard GARCH(1,1) is a sign that intraday information provides a stronger signal on future volatility than squared returns.

Analyzing the estimated parameters of the measurement equations, it turns out that although the realized measures x_t are an upward biased measure of the conditional variance ($\xi > 0$), they are approximately proportional to h_t, as suggested by estimates of φ close to 1.

Furthermore, differently to what is usually observed for financial stock returns showing $\tau_1 < 0$ and $\tau_2 > 0$ as well as $\delta_1 < 0$ and $\delta_2 > 0$, here, we find that all parameters of the leverage functions are positive with $\tau_1 > \tau_2$ and $\delta_1 > \delta_2$. This implies an inverse leverage effect, as positive shocks in electricity prices lead to a larger increase in volatility than negative shocks, confirming the finding in Knittel and Roberts (2005); Frömmel et al. (2014); Qu et al. (2018), among others.

Overall, the estimated measurement error variance confirms that the QLD electricity market is the most volatile, showing the highest values of σ_u^2. Additionally, the parameter estimates for σ_u^2 are approximately the same for RV and RK, but are larger for MRV, highlighting that a frequency of 30 min is not sufficient to cancel out the effects of price jumps. In addition, focusing on the individual realized measures, an inverse relationship emerges between the coefficient γ and the variance σ_u^2 of the residual measurement error. This is because γ reflects the "amount" of information about volatility variation; so, the larger the coefficient, the more accurate the realized measure.

Regarding the ρ parameters, as expected, their estimates are very close to 1, indicating that RV, RK, and MVR are highly positively correlated, implying that some γ are negative or not significant for REGARCH models using two or three realized measures.

The estimated ν parameter ranges between 2.85 and 5.37, with the upper bound always provided by the GARCH(1,1), confirming the existence of leptokurtosis in the conditional distribution of returns.

Finally, as it makes no sense to compare the full log-likelihood for different realized measures, we only report the partial log-likelihood of the return component $\ell(r)$. Not surprisingly, there is a clear improvement from RGARCH to REGARCH specifications, which provide the highest maximized values. The lowest values for $\ell(r)$ always occur in the RGARCH(MRV), while, overall, the standard GARCH(1,1) performs quite well.

6. Out-of-Sample Analysis

In this section, one-day-ahead forecasts of volatility, Value-at-Risk (VaR), and Expected Shortfall (ES) are generated with the rolling window method with daily re-estimation. In particular, the first estimation window is of 2000 observations, leading to different out-of-sample periods for the electricity markets: 26/06/2017–31/12/2019 (913 days) for NSW; 01/08/2017–31/12/2019 (813 days) for QLD; 06/09/2017–31/12/2019 (788 days) for VIC.

The out-of-sample forecasting performance of the models is evaluated by considering different loss functions. First, the ability to accurately forecast volatility is assessed by the QLIKE (Patton 2011), Mean Squared Error (MSE), and Mean Absolute Error (MAE):

$$QLIKE = \frac{1}{T} \sum_{t=1}^{T} \left(log(\hat{h}_t) + \frac{x_t}{\hat{h}_t} \right),$$

$$MSE = \frac{1}{T} \sum_{t=1}^{T} \left(x_t - \hat{h}_t \right)^2,$$

$$MAE = \frac{1}{T} \sum_{t=1}^{T} \left| x_t - \hat{h}_t \right|,$$

where $\hat{h}_t$ is the 1-step-ahead conditional variance forecast and x_t is the volatility proxy at time t. Since the magnitude of bias due to microstructure noise and jumps tend to vanish at low frequencies, to consider different market scenarios in our forecasting study, we refer to the RV computed at the frequencies of 30 min, 2 h, and 6 h as volatility proxies.

Next, we evaluate the out-of-sample forecasting ability of the models considering one-step-ahead VaR and ES forecasts generated for three different risk levels: 1%, 2.5%, and 5%. The adequacy of VaR forecasts is assessed through the Quantile Loss (*QL*) function (González-Rivera et al. 2004; Koenker 2005)

$$QL_t(\alpha) = (\alpha - l_t)(r_t - VaR_t(\alpha)), \tag{13}$$

where $l_t = I_{(r_t < VaR_t(\alpha))}$. The *QL* is a strictly consistent scoring rule for VaR prediction. Further, being an asymmetric loss function, it is particularly suited to assess quantile risk measures as it imposes a higher penalty, with weight $(1 - \alpha)$, for observations below the α-quantile level, namely, when we observe returns exceeding the VaR.

Regarding the ES, as there is no loss function for which the ES is the solution for minimizing the expected loss, it lacks the mathematical property called *elicitability* (see e.g., Fissler and Ziegel (2016); Fissler et al. (2016); Gneiting (2011); Weber (2006) among others).

However, as ES turns out to be jointly elicitable with VaR, we rely on the class of (strictly) consistent scoring function (Fissler and Ziegel 2016) to evaluate the ability of the proposed models to jointly forecast VaR and ES

$$FZ_t(r_t, v_t, e_t, \alpha, G_1, G_2) = (l_t - \alpha) \left(G_1(v_t) - G_1(r_t) + \frac{1}{\alpha} G_2(e_t) v_t \right) \tag{14}$$
$$- G_2(e_t) \left(\frac{1}{\alpha} l_t r_t - e_t \right) - \mathcal{G}_2(e_t),$$

with v_t and e_t, the VaR and ES, respectively, while $G_1(\cdot)$ is weakly increasing, $G_2(\cdot)$ is strictly increasing and strictly positive, and $\mathcal{G}_2'(\cdot) = G_2(\cdot)$.

Although several strictly consistent scoring rules for the pair (VaR, ES) can be obtained as special cases of the family of functions in (14), following Patton et al. (2019), we assume VaR and ES to be strictly negative and $ES_t(\alpha) \leq VaR_t(\alpha) < 0$, with $G_1(x) = 0$ and $G_2(x) = -1/x$, resulting in the zero-degree homogeneous loss function

$$FZ_t^0 = \frac{1}{\alpha ES_t(\alpha)} l_t (r_t - VaR_t(\alpha)) + \frac{VaR_t(\alpha)}{ES_t(\alpha)} + log(-ES_t(\alpha)) - 1. \tag{15}$$

Finally, the Model Confidence Set (MCS) of Hansen et al. (2011) is used to assess the significance of differences in the predictive performances of the models under analysis considering the confidence levels of 75% and 90%. In particular, in the MCS implementation, we have considered a *Range* statistic and 5000 bootstrap resamples generated by means of a block-bootstrap procedure, where the optimal block length has been estimated using the method described in Patton et al. (2009).

Table 5 shows the out-of-sample model comparison based on average losses for QLIKE, MSE, and MAE using different volatility proxies. Values in boldface indicate the models that return the minimum average loss, while those shaded in gray and light-gray are associated with models that are included in the 75% and 90% MCS, respectively.

Table 5. Out-of-sample loss functions comparison: QLIKE, MSE, and MAE.

	NSW								
	RV 30 min			**RV 2 h**			**RV 6 h**		
	QLIKE	**MSE**	**MAE**	**QLIKE**	**MSE**	**MAE**	**QLIKE**	**MSE**	**MAE**
GARCH(1,1)	37.436	9.780	1.349	28.219	2.153	0.936	18.741	1.902	0.673
RG(RV)	21.209	9.706	1.338	15.163	2.113	0.926	9.718	1.880	0.665
RG(RK)	21.089	9.704	1.337	15.039	2.112	0.925	9.627	1.879	0.665
RG(MRV)	**19.637**	**9.675**	**1.333**	**14.087**	**2.103**	**0.921**	**8.936**	**1.872**	**0.661**
REG(RV)	24.369	9.747	1.345	17.995	2.134	0.933	11.668	1.893	0.670
REG(RK)	24.319	9.746	1.345	17.936	2.134	0.933	11.623	1.892	0.670
REG(MRV)	23.237	9.731	1.343	17.219	2.130	0.931	11.098	1.889	0.668
REG(RV,RK)	24.445	9.749	1.345	18.112	2.136	0.933	11.761	1.894	0.671
REG(RV,MRV)	24.573	9.749	1.346	18.129	2.135	0.933	11.768	1.894	0.671
REG(RK,MRV)	24.417	9.747	1.345	17.996	2.134	0.933	11.665	1.893	0.670
REG(RV,RK,MRV)	24.661	9.751	1.346	18.243	2.136	0.934	11.854	1.894	0.671
	QLD								
	RV 30 min			**RV 2 h**			**RV 6 h**		
	QLIKE	**MSE**	**MAE**	**QLIKE**	**MSE**	**MAE**	**QLIKE**	**MSE**	**MAE**
GARCH(1,1)	49.133	11.313	1.556	32.334	4.062	0.998	24.867	4.524	0.834
RG(RV)	22.080	11.168	1.532	13.859	3.978	0.974	9.934	4.442	0.813
RG(RK)	21.985	11.167	1.532	13.767	3.977	0.973	9.881	4.441	0.813
RG(MRV)	**19.534**	**11.135**	**1.523**	**11.887**	**3.958**	**0.965**	**8.586**	**4.427**	**0.807**
REG(RV)	28.145	11.238	1.546	18.101	4.020	0.987	13.192	4.481	0.825
REG(RK)	28.155	11.238	1.545	18.093	4.020	0.987	13.179	4.481	0.825
REG(MRV)	27.310	11.233	1.543	17.183	4.015	0.985	12.670	4.479	0.823
REG(RV,RK)	28.059	11.238	1.545	18.095	4.020	0.987	13.153	4.481	0.824
REG(RV,MRV)	29.172	11.243	1.547	19.262	4.024	0.989	13.653	4.483	0.825
REG(RK,MRV)	28.966	11.241	1.547	19.033	4.023	0.988	13.584	4.482	0.825
REG(RV,RK,MRV)	29.160	11.243	1.547	19.304	4.024	0.989	13.628	4.483	0.825
	VIC								
	RV 30 min			**RV 2 h**			**RV 6 h**		
	QLIKE	**MSE**	**MAE**	**QLIKE**	**MSE**	**MAE**	**QLIKE**	**MSE**	**MAE**
GARCH(1,1)	28.051	28.570	2.532	20.956	13.423	1.789	11.038	6.806	0.951
RG(RV)	15.714	28.238	2.481	11.087	13.176	1.739	5.457	6.686	0.911
RG(RK)	15.818	28.236	2.481	11.129	13.173	1.739	5.506	6.684	0.911
RG(MRV)	**14.946**	**28.206**	**2.471**	**10.404**	**13.164**	**1.730**	**5.066**	**6.675**	**0.904**
REG(RV)	16.962	28.300	2.491	12.103	13.220	1.749	6.086	6.711	0.918
REG(RK)	17.074	28.301	2.491	12.162	13.219	1.749	6.139	6.710	0.918
REG(MRV)	16.099	28.276	2.483	11.359	13.214	1.741	5.620	6.703	0.913
REG(RV,RK)	16.930	28.301	2.491	12.097	13.222	1.749	6.076	6.713	0.918
REG(RV,MRV)	17.454	28.308	2.494	12.520	13.218	1.752	6.385	6.714	0.921
REG(RK,MRV)	17.550	28.307	2.493	12.550	13.215	1.752	6.418	6.712	0.921
REG(RV,RK,MRV)	17.424	28.311	2.494	12.528	13.222	1.752	6.372	6.716	0.921

The table shows the average loss of QLIKE, mean squared error (MSE), and mean absolute error (MAE) using different volatility proxies, for NSW (top), QLD (middle), and VIC (bottom). The best model is reported in bold while models ∈ 75% and ∈ 90% MCS are shaded in gray and light-gray, respectively. Key to table: RG—Realized GARCH; REG—Realized Exponential GARCH.

The results in Table 5 clearly indicate that the RGARCH relying on the jump-robust estimator *medRV* provides more accurate volatility forecasts than other models, always minimizing the loss functions considered for all proxies and each electricity market. Furthermore, it is the only model entering MCS at any confidence level. No other models enter MCS, with the exception of the RGARCH(RV) and RGARCH(RK) using the 30-min RV and the RGARCH(RK) using the 2-hour RV, which are included in the 90% MCS for the MSE loss function and the VIC electricity market. On the other hand, the GARCH(1,1) appears to be the worst competitor as it produces the highest loss in every possible scenario. The largest discrepancies between the GARCH(1,1) and R(E)GARCH models occur for the QLIKE loss function, confirming that the QLIKE is more powerful in rejecting poorly performing predictors (Liu et al. 2015; Patton 2011). Overall, the simple RGARCH structure leads to substantial improvements in the accuracy of volatility forecasts over the REGARCH models based on one or more realized measures.

The scenario is completely reversed when forecasting VaR and ES. The results in Table 6 point out that REGARCH specifications provide lower average loss values for QL and FZ^0 than RGARCH models, which are always excluded from the MCS at any risk level. At the same time, the GARCH(1,1) model also turns out to be a good competitor, minimizing the loss functions in most cases. In particular, it is the only model that always enters the MCS at the most extreme 1% risk level both for VaR and ES, and for NSW and QLD, no other model belongs to the MCS. Although the REGARCH(RV,RK,MRV) model minimizes QL and FZ^0 for different risk scenarios, models based on a single or combination of two variables almost always enter the MCS, especially when RV and RK are considered to explain volatility dynamics. Furthermore, the MCS shows that the differences between the models are more pronounced in forecasting ES. Finally, moving towards less extreme risk levels, such as 5%, there is less discrimination between models.

Table 6. One-step ahead Value-at-Risk and Expected Shortfall loss functions comparison: QL and FZ^0.

	NSW					
	$\alpha = 0.05$		$\alpha = 0.025$		$\alpha = 0.01$	
	QL	FZ^0	QL	FZ^0	QL	FZ^0
GARCH(1,1)	20.366	−0.912	11.260	**−0.797**	**5.286**	**−0.630**
RG(RV)	22.130	−0.786	12.954	−0.605	6.380	−0.362
RG(RK)	22.185	−0.781	13.037	−0.597	6.454	−0.350
RG(MRV)	22.338	−0.777	13.206	−0.588	6.727	−0.329
REG(RV)	18.725	−0.965	10.871	−0.783	5.596	−0.522
REG(RK)	18.735	−0.964	10.911	−0.780	5.613	−0.519
REG(MRV)	18.862	−0.958	11.062	−0.771	5.663	−0.511
REG(RV,RK)	18.627	−0.971	**10.774**	−0.790	5.575	−0.525
REG(RV,MRV)	18.706	−0.966	10.853	−0.785	5.586	−0.524
REG(RK,MRV)	18.733	−0.964	10.910	−0.781	5.607	−0.521
REG(RV,RK,MRV)	**18.620**	**−0.971**	10.777	−0.791	5.569	−0.527
	QLD					
	$\alpha = 0.05$		$\alpha = 0.025$		$\alpha = 0.01$	
	QL	FZ^0	QL	FZ^0	QL	FZ^0
GARCH	15.427	**−1.015**	**8.793**	**−0.860**	**4.552**	**−0.597**
RG(RV)	17.524	−0.853	10.866	−0.628	5.584	−0.350
RG(RK)	17.552	−0.851	10.909	−0.624	5.605	−0.347
RG(MRV)	17.670	−0.836	11.146	−0.595	5.918	−0.298
REG(RV)	15.400	−0.987	9.410	−0.763	5.011	−0.461
REG(RK)	15.408	−0.987	9.424	−0.762	5.009	−0.462
REG(MRV)	**15.358**	−0.986	9.523	−0.754	5.053	−0.454
REG(RV,RK)	15.388	−0.987	9.410	−0.762	5.027	−0.457
REG(RV,MRV)	15.415	−0.986	9.301	−0.769	4.993	−0.461
REG(RK,MRV)	15.424	−0.985	9.301	−0.767	4.991	−0.463
REG(RV,RK,MRV)	15.405	−0.985	9.301	−0.767	5.009	−0.456

Table 6. *Cont.*

| | VIC | | | | | |
| | $\alpha = 0.05$ | | $\alpha = 0.025$ | | $\alpha = 0.01$ | |
	QL	FZ^0	QL	FZ^0	QL	FZ^0
GARCH	28.343	−0.299	16.870	−0.118	8.699	0.163
RG(RV)	28.262	−0.342	17.208	−0.150	8.856	0.089
RG(RK)	28.315	−0.340	17.227	−0.149	8.867	0.090
RG(MRV)	28.477	−0.324	17.523	−0.123	9.287	0.144
REG(RV)	27.462	−0.380	16.643	−0.192	8.651	0.057
REG(RK)	27.506	−0.378	16.650	−0.192	8.661	0.058
REG(MRV)	27.653	−0.363	16.890	−0.168	8.964	0.102
REG(RV,RK)	27.436	−0.380	16.637	−0.193	8.641	0.056
REG(RV,MRV)	27.322	−0.389	16.409	−0.208	8.375	0.030
REG(RK,MRV)	27.438	−0.385	16.460	−0.205	8.437	0.035
REG(RV,RK,MRV)	**27.280**	**−0.391**	**16.391**	**−0.209**	**8.347**	**0.028**

The table shows the loss values of *QL* and *FZ*^{0} at the risk levels of $\alpha = \{0.05; 0.025; 0.01\}$, for NSW (top), QLD (middle), and VIC (bottom). The best model is reported in bold while models $\in 75\%$ and $\in 90\%$ MCS are shaded in gray and light-gray, respectively. Key to table: RG—Realized GARCH; REG—Realized Exponential GARCH.

7. Conclusions

This paper uses half-hourly spot prices from the Australian electricity markets of New South Wales, Queensland, and Victoria to predict volatility and manage risk in energy markets. In this framework, we extend the literature on modeling the conditional variance of returns using the Realized GARCH approach by combining information from multiple realized robust and nonrobust measures to capture the key features of electricity prices such as extreme jumps and the inverse leverage effect. Our empirical analysis underlines the following points. First, specifications with multiple realized measures outperform those based on a single realized measure as well as the standard GARCH(1,1), resulting in a remarkably better fit of the data in-sample. Second, the *medRV* jump-robust measure significantly increases the accuracy of out-of-sample volatility forecasts. In particular, the simple Realized GARCH based on a single measurement equation for the jump-robust *medRV* estimator always minimizes the set of loss functions considered—i.e., QLIKE, MSE, and MAE—and is the only model that enters the MCS under all circumstances addressed. Finally, in contrast to volatility forecasting, when assessing the predictive ability of the models in terms of VaR and ES, it emerges that the standard GARCH is highly competitive especially for the extreme risk level of 1%. Similarly, REGARCH models based on one or more realized measures outperform the simple RGARCH, which shows—in this case—the worst results in minimizing the loss functions at any risk level. Further, the MCS highlights greater discrimination between models in predicting ES. Electricity market participants aim to constantly pursue optimal trading limits in order to adequately allocate capital and to cover potential losses if trading limits are violated. This is also because overcapitalization implies idle capital that could undermine the profitability of energy industries; at the same time, undercapitalization could cause financial difficulties when they are unable to honor trading contracts. Therefore, accurately predicting VaR and ES is essential for effective energy risk management, as they are the most commonly used measures for establishing optimal trading limits. One aspect that has not been considered here, and that is worth examining in future research, is how the inclusion of exogenous factors, such as weather conditions, would affect electricity price volatility. In addition, volatility inter-relationships between various regions and energy markets are interesting future research areas, as is extending the results to energy markets other than Australia. Finally, as this study has mainly focused on modeling one-day-ahead volatility, exploiting the properties of the Realized GARCH-type models, a natural extension would be to predict price volatility at a longer horizon.

References

Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold. 2007. Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *The Review of Economics and Statistics* 89: 701–20. [CrossRef]

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys. 2001. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* 96: 42–55. [CrossRef]

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys. 2003. Modeling and forecasting realized volatility. *Econometrica* 71: 579–625. [CrossRef]

Andersen, Torben G., Dobrislav Dobrev, and Ernst Schaumburg. 2012. Jump-robust volatility estimation using nearest neighbor truncation. *Journal of Econometrics* 169: 75–93. [CrossRef]

Barndorff-Nielsen, Ole E., Peter Reinhard Hansen, Asger Lunde, and Neil Shephard. 2008. Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica* 76: 1481–536.

Barndorff-Nielsen, Ole E., and Neil Shephard. 2002. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64: 253–80. [CrossRef]

Bierbrauer, Michael, Christian Menn, Svetlozar T. Rachev, and Stefan Trück. 2007. Spot and derivative pricing in the EEX power market. *Journal of Banking & Finance* 31: 3462–85.

Bowden, Nicholas, and James E. Payne. 2008. Short term forecasting of electricity prices for miso hubs: Evidence from arima-egarch models. *Energy Economics* 30: 3186–97. [CrossRef]

Byström, Hans N. 2005. Extreme value theory and extremely large electricity price changes. *International Review of Economics & Finance* 14: 41–55.

Chan, Kam Fong, and Philip Gray. 2006. Using extreme value theory to measure value-at-risk for daily electricity spot prices. *International Journal of forecasting* 22: 283–300. [CrossRef]

Chan, Kam Fong, Philip Gray, and Bart van Campen. 2008. A new approach to characterizing and forecasting electricity price volatility. *International Journal of Forecasting* 24: 728–43. [CrossRef]

Corsi, Fulvio. 2009. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7: 174–96. [CrossRef]

Escribano, Alvaro, J. Ignacio Peña, and Pablo Villaplana. 2011. Modelling electricity prices: International evidence. *Oxford Bulletin of Economics and Statistics* 73: 622–50. [CrossRef]

Fanone, Enzo, Andrea Gamba, and Marcel Prokopczuk. 2013. The case of negative day-ahead electricity prices. *Energy Economics* 35: 22–34. [CrossRef]

Fissler, Tobias, and Johanna F. Ziegel. 2016. Higher order elicitability and osband's principle. *Annals of Statistics* 44: 1680–707. [CrossRef]

Fissler, Tobias, Johanna F. Ziegel, and Tilmann Gneiting. 2016. Expected shortfall is jointly elicitable with value at risk-implications for backtesting. *arXiv* arXiv:1507.00244.

Frömmel, Michael, Xing Han, and Stepan Kratochvil. 2014. Modeling the daily electricity price volatility with realized measures. *Energy Economics* 44: 492–502. [CrossRef]

Garcia, Reinaldo C., Javier Contreras, Marco Van Akkeren, and João Batista C. Garcia. 2005. A garch forecasting model to predict day-ahead electricity prices. *IEEE Transactions on Power Systems* 20: 867–74. [CrossRef]

Gneiting, Tilmann. 2011. Making and evaluating point forecasts. *Journal of the American Statistical Association* 106: 746–62. [CrossRef]

González-Rivera, Gloria, Tae-Hwy Lee, and Santosh Mishra. 2004. Forecasting volatility: A reality check based on option pricing, utility function, value-at-risk, and predictive likelihood. *International Journal of Forecasting* 20: 629–45. [CrossRef]

Hansen, Peter Reinhard, and Zhuo Huang. 2016. Exponential garch modeling with realized measures of volatility. *Journal of Business & Economic Statistics* 34: 269–87.

Hansen, Peter Reinhard, Zhuo Huang, and Howard Howan Shek. 2012. Realized garch: A joint model for returns and realized measures of volatility. *Journal of Applied Econometrics* 27: 877–906. [CrossRef]

Hansen, Peter R., Asger Lunde, and James M. Nason. 2011. The model confidence set. *Econometrica* 79: 453–97. [CrossRef]

Haugom, Erik, and Carl J. Ullrich. 2012. Forecasting spot price volatility using the short-term forward curve. *Energy Economics* 34: 1826–33. [CrossRef]

Haugom, Erik, Sjur Westgaard, Per Bjarte Solibakke, and Gudbrand Lien. 2011. Realized volatility and the influence of market measures on predictability: Analysis of nord pool forward electricity data. *Energy Economics* 33: 1206–15. [CrossRef]

Hickey, Emily, David G. Loomis, and Hassan Mohammadi. 2012. Forecasting hourly electricity prices using armax–garch models: An application to miso hubs. *Energy Economics* 34: 307–15. [CrossRef]

Higgs, Helen, and Andrew Worthington. 2008. Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The australian wholesale spot electricity market. *Energy Economics* 30: 3172–85. [CrossRef]

Knittel, Christopher R., and Michael R. Roberts. 2005. An empirical examination of restructured electricity prices. *Energy Economics* 27: 791–817. [CrossRef]

Koenker, Roger. 2005. *Quantile Regression*. In *Econometric Society Monographs*; Cambridge: Cambridge University Press.

Li, Yu-Ning, Yi Zhang, and Caiya Zhang. 2019. Statistical inference for measurement equation selection in the log-realgarch model. *Econometric Theory* 35: 943–77. [CrossRef]

Liu, Heping, and Jing Shi. 2013. Applying arma–garch approaches to forecasting short-term electricity prices. *Energy Economics* 37: 152–66. [CrossRef]

Liu, Lily Y., Andrew J. Patton, and Kevin Sheppard. 2015. Does anything beat 5-min rv? a comparison of realized measures across multiple asset classes. *Journal of Econometrics* 187: 293–311. [CrossRef]

Patton, Andrew, Dimitris N. Politis, and Halbert White. 2009. Correction to "Automatic block-length selection for the dependent bootstrap" by D. Politis and H. White. *Econometric Reviews* 28: 372–75. [CrossRef]

Patton, Andrew J. 2011. Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics* 160: 246–56. [CrossRef]

Patton, Andrew J., Johanna F. Ziegel, and Rui Chen. 2019. Dynamic semiparametric models for expected shortfall (and value-at-risk). *Journal of Econometrics* 211: 388–413. [CrossRef]

Qu, Hui, Qingling Duan, and Mengyi Niu. 2018. Modeling the volatility of realized volatility to improve volatility forecasts in electricity markets. *Energy Economics* 74: 767–76. [CrossRef]

Weber, Stefan. 2006. Distribution-invariant risk measures, information, and dynamic consistency. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics* 16: 419–41. [CrossRef]