



A probabilistic forecast methodology for volatile electricity prices in the Australian National Electricity Market

Cameron Cornell^a, Nam Trong Dinh^b, S. Ali Pourmousavi^{b,*}

^a The University of Adelaide, School of Computer and Mathematical Sciences, Adelaide, Australia

^b The University of Adelaide, School of Electrical and Mechanical Engineering, Adelaide, Australia

ARTICLE INFO

Keywords:

Electricity price forecasting
Probability forecasting
Australian National Electricity Market
Ensemble forecast
Quantile regression
Quantile regression forest
Autoregression

ABSTRACT

The South Australia region of the Australian National Electricity Market (NEM) displays some of the highest levels of price volatility observed in modern electricity markets. This paper outlines an approach to probabilistic forecasting under these extreme conditions, including spike filtration and several post-processing steps. We propose using quantile regression as an ensemble tool for probabilistic forecasting, with our combined forecasts achieving superior results compared to all constituent models. Within our ensemble framework, we demonstrate that averaging models with varying training-length periods leads to a more adaptive model and increased prediction accuracy. The applicability of the final model is evaluated by comparing our median forecasts with the point forecasts available from the Australian NEM operator, with our model outperforming these NEM forecasts by a significant margin.

© 2023 The Author(s). Published by Elsevier B.V. on behalf of International Institute of Forecasters. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Following the deregulation of most of the electricity industries in the 1990s, we saw a growing focus of research on electricity price forecasting (EPF). Decision makers who are properly equipped with reliable electricity price forecasts can adjust their bidding strategies and production/consumption schedule to maximise objectives in day-ahead and real-time trading. Additionally, the electricity market is unique in terms of commodity trade, as it is still economically and technically difficult to store electricity at large scales, thus requiring constant efforts to maintain a balance between production and consumption. The issue of large-scale storage combined with stochastic production and demand behaviour, which depends on the time of day, weather conditions, and

business activity, leads to extreme spot price volatility, generally unobserved in other commodity markets. As a result, this complex market structure has led to extensive EPF research in recent years (Lago, Marczasz, De Schutter, & Weron, 2021).

Recently, there has been a shift in focus from point forecasts (expected value of the spot price) to probabilistic forecasts (estimation of density/interval) (Nowotarski & Weron, 2018). Probabilistic prediction gained significant momentum after the 2014 Global Energy Forecasting Competition (GEFCom2014), which focused solely on probabilistic energy prediction (Hong et al., 2016). In many applications, such as risk management and bidding in the wholesale electricity market, a successful strategy depends not only on knowledge of expected price levels but also on predicted variability of prices within a given day. The precise method used to capture this variability varies across the literature. Generally, rather than estimating the entire density (Afrasiabi, Mohammadi, Rastegar, & Kargarian, 2019; Chai, Xu, & Jia, 2019),

* Corresponding author.

E-mail addresses: cameron.cornell@adelaide.edu.au (C. Cornell), trongnam.dinh@adelaide.edu.au (N.T. Dinh), a.pourm@adelaide.edu.au (S.A. Pourmousavi).

we limit ourselves to key characteristics of this distribution (Nowotarski & Weron, 2015). A common form of probabilistic forecasting is prediction intervals, where a proportion of the data is expected to lie inside this interval. Alternatively, it can be represented using q quantiles, where a proportion q of the data is below this quantile. Both methods can be considered a discretisation of the distribution function to simplify the estimation process. In this study, we focus on quantile estimates in our target range of 0.025–0.975. In particular, we focus on nine key quantiles: 0.025, 0.05, 0.1, 0.25, 0.5, 0.75, 0.9, 0.95, and 0.975.

For a comprehensive review of both point forecasting EPF methods and a more recent review of probabilistic EPF, we refer the reader to the influential papers by Weron (2014) and Nowotarski and Weron (2018). In this paper, we propose a model for probabilistic EPF based on forecast averaging. Many articles have previously concluded that ensemble point forecasting leads to a more accurate and robust prediction of electricity prices (Bordignon, Bunn, Lisi, & Nan, 2013; Nowotarski, Raviv, Trück, & Weron, 2014; Raviv, Bouwman, & Van Dijk, 2015). However, there is less comprehensive evidence to suggest the benefits of ensemble forecasting in the field of probabilistic forecasting methods. Nowotarski and Weron (2015) discussed the possible accuracy gains from averaging forecasts for probabilistic EPF, concluding that their ensemble method (dubbed quantile regression averaging (QRA)) leads to more accurate prediction intervals. The results of several forecasting competitions testify to the benefits of forecast averaging, with two of the four winners of the GEFCom2014 price track using some form of averaging (Hong et al., 2016), and the number one finding of the 2018 M4 forecast competition is ‘The improved numerical accuracy of combining’ (Makridakis, Spiliotis, & Assimakopoulos, 2020). In light of these significant advantages, in this paper, we seek a systematic way of combining sets of increasingly complex price predictions into one final forecast. Our general ensemble method is an extension to the method used by Wang et al. (2019) and QRA, where the contributions are as follows:

1. We first show that the probabilistic prediction of the price in the Australian National Electricity Market (NEM) can be improved for many quantiles by filtering the extreme spikes out of the training dataset. Specifically, quantiles 0.10 to 0.95 showed substantial improvements for most models.
2. Next, we demonstrate two post-processing steps to increase the accuracy of our forecasting methods. Many predictors exhibit ‘oscillatory’ prediction behaviour when trained on highly volatile data. By running a smoothing routine on our predictions, we see a minor increase in accuracy and improved interpretability. Next, we show how autoregression can be used at the probabilistic level to shift the entire estimated density up or down based on the market prices at the start of the 24-h prediction blocks. The advantage of this additive approach is that it allows us to capture additional temporal dependency when using models that generate probabilistic predictions, which are rarely designed in the literature with time series considerations.

3. During the training stage, we demonstrate that there are systematic accuracy gains from taking our forecast ensembles, not only across model types but also across the length of the training data (memory).

The paper is structured as follows. In Section 2, we describe our data and the extreme price volatility observed in the NEM, and we motivate the need for algorithms capable of capturing the complex risk profile observed in South Australia (SA) price series. Section 3 outlines our general methodology, with comprehensive coverage of our constituent models and an outline of the forecast combination framework. Different combination methods are presented in Section 4. In Section 5, we discuss the combination of models across different training lengths. Finally, the numerical results and the performance comparison from both statistical and economic points of view are presented in Sections 6 and 7. The paper is concluded in Section 8.

2. Data

2.1. The variables

The target variable in this study is the five-minute interval spot prices in the SA region of the Australian NEM.¹ The Australian NEM differs from many other countries, such as those depending on the NordPool, as there is no separation into a day-ahead and intraday market. Instead, the Australian NEM operates as a spot market in which prices are determined for each five-minute trading interval in advance. Generators can re-bid anytime until the bids are captured for dispatch for the trading interval (AEMO, 2021b). Prices are then determined by solving a security-constrained linear optimal power flow model (AEMO, 2023). Although bidding occurs continuously and can be submitted more than 24 h in advance, in this study, we restricted our forecast horizon to a maximum of 24 h for simplicity.

The exogenous variables that we use in our models are primarily historical prices, time, and weather data.

- **Historical electricity prices** consisting of five-minute spot prices in the SA region (matching our target) are provided by the Australian Energy Market Operator (AEMO). For all given predictor models, we have two lagged prices as model inputs, one from exactly 24 h ago and another from seven days ago (y_{t-288} and y_{t-2016} , respectively).
- **Temporal (time) information** is encoded with an indicator variable for each day of the week, as well as an integer variable between 1 (Time 00:00) and 288 (Time 23:55) corresponding to each of the five-minute observation times per day.

¹ Australian NEM publicly available data: <https://aemo.com.au/energy-systems/electricity/national-electricity-market-nem/data-nem/market-data-nemweb>.

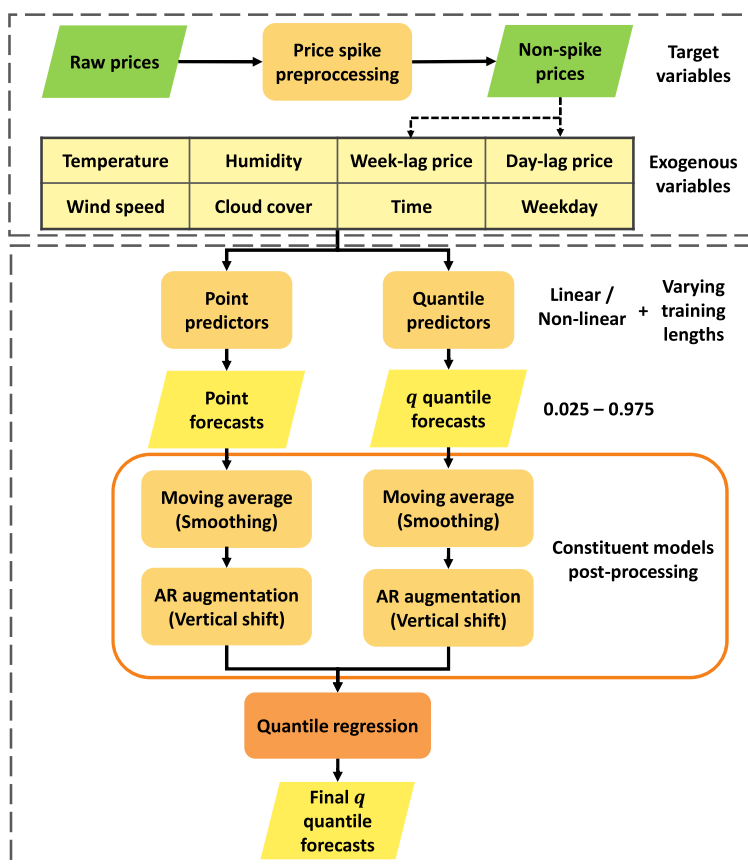


Fig. 1. Flowchart of the proposed Q-QRA prediction model.

- **Historical weather forecast data** were collected from World Weather Online (Online, 2022) for three different regions of SA: Hallett, Hornsdale, and Adelaide. The data comprise wind speed (km/h), temperature (°C), humidity (%), and cloud cover (%) for the testing period.
- **Nonlinear terms** were augmented from the inputs mentioned above for the considered linear models. These terms are an order six polynomial to the time integer variable to capture the cyclic daily structure, as well as quadratic terms for both temperature and wind speed.

The summary of exogenous variables is shown in Fig. 1.

2.2. Price volatility in the NEM

It has been noted that Australian electricity spot prices exhibit some of the highest volatility levels observed in modern power markets. This is particularly exacerbated within the SA region, where there is a high proportion of renewable energy and reliance on Victorian power production through two interconnectors. This volatility is observed on multiple levels. Not only is the time series generally volatile according to changing weather patterns and price history but also it exhibits extreme price spikes on an infrequent basis. The intense magnitude of these

Table 1

Summary statistics for the SA electricity spot price from 1/1/2018 to 31/12/2021.

Year	SD	MAD	Mean	Median	Skew	Kurt
2018	346.01	20.94	92.13	82.89	33.69	1260.03
2019	452.23	22.76	83.04	83.50	25.86	777.44
2020	236.04	12.67	42.93	41.75	50.10	2914.50
2021	334.30	29.55	50.70	39.18	35.65	1409.57

^aHere, MAD is the median absolute deviation, median centred.

spikes ensures that while infrequent, they have a significant impact on the market operation and profitability of the market participants. To somewhat limit the effects of these spikes, there is a regulatory ceiling and floor prices, at (AUD) −\$1000 and \$15,100 per MWh in 2021 (AEMO, 2022). Table 1 shows price volatility statistics in the SA region over the out-of-sample testing period from 1/1/2018 to 31/12/2021.

The numbers presented in Table 1 suggest strong evidence of SA price variability. First, we see that the standard deviation (SD) is quite large, being several times higher than the mean and median prices. This contrasts with the median absolute deviations (MADs), which are generally lower than price levels. This behaviour can be explained by the extremely high kurtosis (fat tails), which ranges around a thousand. We also observe a

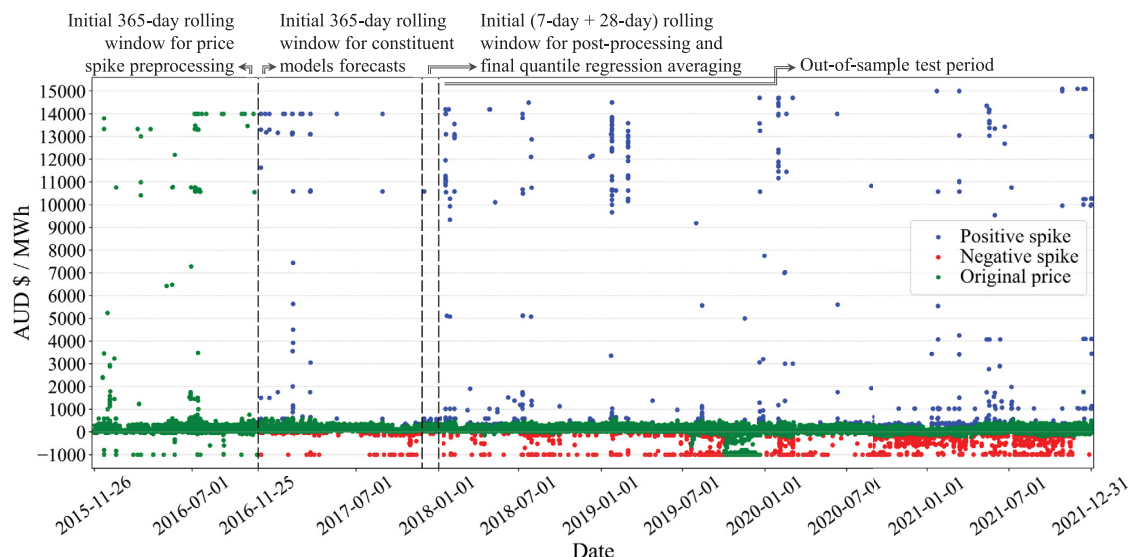


Fig. 2. Five-minute interval electricity spot prices in South Australia from 26/11/2015 to 31/12/2021. The vertical dashed lines mark the beginning of the rolling windows, which are identified and labelled within the figure. The colouration of points indicates the spike classifications generated from our pre-processing steps.

very heavy skew, indicating that our price distributions are top-heavy. The difference between the SD and MAD reflects the empirically observed behaviour: the majority of prices cluster around the distribution centre, with rare extreme outliers. To visualise this behaviour, Fig. 2 shows the price history over the four-year testing period, along with the initial rolling training windows and price classifications, as we will discuss in Section 3.

The difference between the implied volatility under the SD and MAD gives us a glimpse of the primary problem of the data. On the one hand, we see that the SD is a non-robust assessment of deviation from a non-robust centre and reports extremely high variation of the data. On the other hand, the MAD seems to be a robust measure of deviation from a robust centre and indicates quite mild variation in the data. As predictive models are trained by reducing variation statistics, it is critical that we consider the influence of these spikes and the robustness of any variation measure during the training process.

3. Proposed methodology

The proposed method is an extension of quantile regression averaging to incorporate not only point forecasts of electricity prices but also quantile estimates from multiple probabilistic forecasting models. Using this method, we can combine the benefits of forecast averaging with a capacity to represent nonlinear relationships at the upper- and lower-quantile levels of the price density. This is particularly important with our NEM dataset, as there is substantial skew and excess kurtosis (see Fig. 2). The goal of our ensemble model for probabilistic forecasting is to capture the exogenous conditions that lead to this spiking behaviour. In addition to the combination framework, we demonstrate several post-processing methods to improve prediction accuracy. We also show that ensembles can

be taken not only across models but also across different subsets of the training data.

The 24-h prediction process for a given quantile, q , is as follows. First, we begin by pre-processing the raw prices by filtering out the extreme price spikes. Using the processed prices as target variables, we then train a series of models to predict the 24-h-ahead electricity prices/desired quantiles based on a series of explanatory variables, such as the time of day and weather conditions. Next, we introduce two post-processing steps for these individual model predictions. The first step is to smooth out the prediction series to reduce variability. Then, we augment these predictions by utilising the autoregressive (AR) structure of the residual series to shift the predictions/quantiles based on whether prices are above or below expectation at the most recent observation. As mentioned, all models contain the 24-h lagged price for each prediction. Therefore, the baseline models are equivalent to a continuous 24-h-ahead rolling prediction. In this context, autoregression can be considered a post-processing step that incorporates information that is more recent than the previous 24-h prices used in the base models. In the end, we utilise quantile regression to combine this set of forecasts into our final quantile prediction. The proposed model is, thus, called Q-QRA. Fig. 1 shows a flowchart of the proposed prediction model. In the following subsections, different aspects of the proposed methodology are explained in detail.

3.1. Rolling forecast windows

Similar to the works by Hubicka, Marcjasz, and Weron (2019), Marcjasz, Serafin, and Weron (2018) and Serafin, Uniejewski, and Weron (2019), we consider a rolling window scheme in which we re-train our models daily and generate forecasts for the next 24 h. Although our models

are structured to be implemented in some arbitrary 24-h-ahead manner, for testing purposes, we only perform price prediction once a day at midnight. For the first forecast in the out-of-sample test period, i.e., the forecast for 1/1/2018, we use data from 26/11/2015 to process and obtain forecasts from the constituent models; see Fig. 2. Particularly, the initial 365 days starting from 26/11/2015 are used for pre-processing price spikes. Hence, price classifications are not recorded until 25/11/2016, as shown in Fig. 2. The subsequent 365 days are used for generating point/quantile forecasts from the constituent models. Note that since we train the constituent models on different lengths of training data, we visually indicate the size of the rolling window for the constituent forecast models to match the longest training length. We discuss the varying training lengths in Section 5. Finally, the last 35 days prior to the test date are allocated for post-processing and training the final Q-QRA model. To forecast for the next day, that is, 2/1/2018, we shift all the rolling windows forward by one day and repeat the process.

3.2. Price spike pre-processing

As discussed in Section 2, our price series exhibits extremely high levels of volatility and infrequent extreme price spikes. The presence of these spikes ensures the need for robust statistics and consideration of their impact on all levels of model development. As defined in the introduction, the target quantiles of our prediction routines within this study range from (0.025–0.975). As such, we focus on predicting the intra-day variation and mild spiking that occurs with relative frequency. In this context, the extreme spikes (visible in Fig. 2) are outliers, and below, we show that they can reduce the accuracy of predictions for our quantiles of interest. Therefore, we placed them outside the scope of our predictive goals and looked for spike filtration methods, in which several papers have shown improved accuracy in EPF (Conejo, Contreras, Espínola, & Plazas, 2005; Janczura, Trück, Weron, & Wolff, 2013; Weron & Misiorek, 2008). However, these studies focused on point prediction. Comparatively, we observed that the top two winners of the GEFCom2014 price track used some form of price filtration for their forecasts (Hong et al., 2016), indicating that these results extend to probabilistic forecasting.

We now outline our treatment of price spikes, in which the main step is to replace extreme prices with 'processed' values. The threshold to consider a given price level as a spike can be determined somewhat arbitrarily. In our treatment of spikes, we follow the guidelines found in review articles, e.g., Weron (2014), while extending these methods to form our classifier. Firstly, defining spikes in terms of a flat threshold leads to inconsistency across long time periods, as general price behaviour is observed to change substantially over time. Secondly, to set varying price thresholds, the measures of price levels should be derived from robust measures such as the median or quantiles. Otherwise, using the rolling average price to define the spikes will be influenced by the spikes.

Our spike classification method is based on different conditions, each designed to offset what we observed as

weaknesses in the other condition. The method combines different quantile levels at different rolling windows, capturing the degree to which the price magnitude is a statistical outlier. Our classification rule for price y_t at time t is described as follows:

$$\text{Positive spike} = (y_t > 0) \cap (y_t > Q_{t,a}^+ + Q_{t,m}^+), \quad (1a)$$

$$\text{Negative spike} = (y_t < 0) \cap (y_t < \frac{Q_{t,a}^- + Q_{t,m}^-}{2}). \quad (1b)$$

For our dataset, we have $Q_{t,a}^+$ and $Q_{t,a}^-$ as the 0.975 and 0.025 quantile price of one year of trailing data, whereas $Q_{t,m}^+$ and $Q_{t,m}^-$ represent the 0.975 and 0.025 quantile price of the month prior to the processing interval. To avoid the impact on the rolling window values of long-period negative prices near the end of 2019, we take the average of the trailing quantiles for the negative spike filtering. Fig. 2 displays the classification from this method in the corresponding colour. Once we have our classification, we replace spikes with the most recent non-spike price, an implementation of the 'neighbour' imputation used by Geman and Roncoroni (2006).

3.3. Point forecast

Our ensemble method is capable of incorporating quantile forecasts even though it is not limited to them. One of the original benefits of base QRA was to facilitate the use of extensive point prediction research and methods developed in the literature; see, e.g., Weron (2014). Given that some point prediction methods lack a probabilistic counterpart, we seek to incorporate point predictions into our ensemble model. This allows us to leverage their predictive capacity without having to directly convert the methods into a probabilistic format, which is unlikely to be feasible within the scope of aggregating many forecasts.

In this study, we selected the radial basis function (RBF) support vector machine (SVM) as the point predictor for our Q-QRA process. Although we have only one point predictor in this paper, our method can accommodate any number of point predictors. Our selection and inclusion of this specific method is largely illustrative and was not targeted as the most accurate point forecasting method.

SVM originates in the classification setting, where they can achieve complex nonlinear boundaries. To create these boundaries, the input data are nonlinearly mapped into higher-dimensional kernel space, creating a linear decision boundary in the kernel space. This method was extended by Drucker, Burges, Kaufman, Smola, and Vapnik (1996) to allow regression and has since seen relatively widespread use as a nonlinear regression tool. SVMs have been used successfully as primary point predictors in multiple hybrid EPF models (Che & Wang, 2010; Zhao, Dong, Xu, & Wong, 2008).

3.4. Linear quantile regression model

Linear quantile regression (QR) has seen widespread application following its introduction by Koenker and Hallock (2001). It functions as an extension to the least

absolute deviation (LAD) regression, where the minimisation of the residual loss function, \mathcal{L}_1 , results in an estimate of the conditional median. This extension relies on a further alteration of the loss function from a simple residual to the pinball loss function shown below:

$$\mathcal{L}_q(y_t, \hat{y}_{t,q}) = \begin{cases} (y_t - \hat{y}_{t,q})q & \hat{y}_{t,q} \leq y_t \\ (\hat{y}_{t,q} - y_t)(1 - q) & y_t < \hat{y}_{t,q}. \end{cases} \quad (2)$$

A regression model trained to minimise the expected pinball loss \mathcal{L}_q for a given quantile $q \in (0, 1)$ can be shown to result in predictions of the conditional quantile (Koenker & Hallock, 2001). Predictions are generated by a linear combination of the m predictor variables, as shown below:

$$\hat{y}_{t,q} = \mathbf{X}_t^\top \hat{\boldsymbol{\beta}}_q, \quad (3)$$

where $\mathbf{X}_t = (1, x_{t,1}, x_{t,2}, \dots, x_{t,m})^\top$ denotes the input vector, and $\hat{\boldsymbol{\beta}}_q = (\hat{\beta}_{q,0}, \hat{\beta}_{q,1}, \hat{\beta}_{q,2}, \dots, \hat{\beta}_{q,m})^\top$ denotes the estimated coefficient vector of the model, which is determined by the following minimisation process:

$$\hat{\boldsymbol{\beta}}_q = \arg \min_{\boldsymbol{\beta}_q} \sum_{i=1}^l \mathcal{L}_q(y_i, \mathbf{X}_i^\top \boldsymbol{\beta}_q), \quad (4)$$

on the l observations in the training data. The main limitation of this method is the inability to inherently represent nonlinearities in the data. While we can manually include identifiable nonlinearities, e.g., *Temperature*², and a complex hour structure, it is impossible to capture all such effects exhaustively in this way. This effect is particularly detrimental to our dataset, as we observe complex behaviour regarding the outer distributional quantiles. Therefore, it requires the use of nonlinear models for generating predictive quantiles to improve prediction accuracy.

3.5. Nonlinear quantile regression model

To more accurately capture the risk profile observed in our NEM dataset, we consider nonlinear models for probabilistic forecasting. The associated models typically fall under many different labels, such as data mining, machine learning, non-parametric models, etc.

The nonlinear quantile regression method that we use in this study is the quantile regression forest (QRF) algorithm introduced by Meinshausen and Ridgeway (2006). This works in an almost identical manner to the commonly utilised machine learning algorithm of random forests (RF). However, rather than predicting conditional expectations, we estimate the conditional quantile. QRF has recently been applied to forecasting problems, with three of the nine final-round contestants from the GEF-Com2017 competition using some forms of quantile forests (Hong, Xie, & Black, 2019). This competition, however, was for load forecasting. To our knowledge, QRF has not yet been used for EPF purposes; hence, we provide an outline of this method below.

RF and QRF are bootstrap-aggregated (bagged) versions of a decision tree. A decision tree is an intuitive model that captures nonlinear relationships between response and predictor variables. To do so, it recursively

splits the data into smaller subsets based on binary splits along an explanatory variable. The final product is the original model space separated into multiple distinct rectangular subspaces, each corresponding to a terminal node in the decision tree. To predict the value of a new datapoint, the model allocates the point to the appropriate subspace based on the binary rulings applied to the explanatory variables. Our prediction is the mean of the values in this subspace. We can think of this prediction as a weighted sum of the training data y_i :

$$\hat{y}_t = \sum_{i=1}^l \omega_i(\mathbf{X}_t) y_i, \quad (5)$$

where $\omega_i(\mathbf{X}_t)$ is $1/s$ for the s points in the relevant terminal node, 0 elsewhere.

The 'forest' aspect of RF is related to the use of bootstrap aggregation to reduce model variance. By training many decision trees based on slightly perturbed data samples, the final prediction can be the average output of each tree. Using the linear combination framework, we have the prediction from our RF as the following:

$$\hat{y}_t = \sum_{i=1}^l \bar{\omega}_i(\mathbf{X}_t) y_i, \quad (6)$$

where $\bar{\omega}_i(\mathbf{X}_t)$ is the average of $\omega_i(\mathbf{X}_t)$ across all trees in the forest.

The extension to QRF is straightforward. The formal definition of a conditional quantile is $Q_q(Y_t | \mathbf{X}_t) = \inf(y_t \in \mathbb{R} : q \leq F(y_t | \mathbf{X}_t))$. From this we can see that a sufficient value to estimate this quantile is $F(y_t | \mathbf{X}_t)$ evaluated at finite points. However, $F(y_t | \mathbf{X}_t) = P(Y_t \leq y_t | \mathbf{X}_t) = E(\mathbf{1}_{(Y_t \leq y_t)} | \mathbf{X}_t)$. Just as the conditional expectation can be estimated using our weighted combination of y_i , so can this indicator expectation. Using the RF prediction algorithm where y_i is replaced by the indicator function for $y_i \leq y_t$, we arrive at an estimate:

$$\hat{F}(y_t | \mathbf{X}_t) = \sum_{i=1}^l \bar{\omega}_i(\mathbf{X}_t) \mathbf{1}_{(y_i \leq y_t)}. \quad (7)$$

Using this function and the definition of $Q_q(Y_t | \mathbf{X}_t)$, we can arrive at the prediction $\hat{y}_{t,q}$ for QRF as follows:

$$\hat{y}_{t,q} = \inf(y_t \in \mathbb{R} : q \leq \hat{F}(y_t | \mathbf{X}_t)). \quad (8)$$

The key benefit of using a QRF over a linear QR model is the ability to capture complex behaviour across each quantile. As the QRF model does not estimate any regression coefficient, simply returning values based on empirical quantiles of data points with similar explanatory variable conditions, it has the potential to capture the complex skew risk that we observed in the NEM price series.

3.6. Constituent model post-processing

3.6.1. Prediction smoothing

The SA price series exhibits high volatility. Consequently, the final predictions are often extremely unstable. This is particularly exacerbated with the outer

quantiles of the QRF models, where the model predictions often exhibit oscillatory behaviour. To remedy this for application purposes and to improve prediction accuracy, we run a smoothing routine over the predicted values. This is particularly helpful before the AR step. Consider a series of consecutive prices closely related such that this period exhibits high autocorrelation. If the predictions during this time are frequently oscillating, we will observe weakly correlated residuals, obscuring the true AR behaviour behind unstable predictions. The technique we use for smoothing is a simple 12th-order-centred moving average. Although this may seem extreme given the highly transient nature of prices, we observe that beyond oscillatory behaviour, the prediction volatility is minor within a given hour.

3.6.2. Density autoregression

Within our base forecasting routines, we continuously predict prices 24 h ahead from the current time τ , i.e., all values $t \in [\tau + 1, \tau + 288]$ for price series with a five-minute sampling rate. For each of these forecasts, we have both 24-h and seven-day lagged prices as input; that is, to forecast noon tomorrow, we are using noon prices today. While this ensures there is always historical price data to perform forecasting in the base models, it fails to utilise the strong AR price behaviour for the earlier predictions within our 24-h blocks. The predicted values, e.g., $\hat{y}_{\tau+1}$, are not using the obviously helpful price y_τ at time τ . To capture this key information, we run an AR model on the residual series obtained by subtracting the predicted prices from the true price series. In a point prediction scenario, this is equivalent to assuming a decomposition of prices at general time t into

$$y_t = f(\mathbf{X}_t) + \epsilon_t, \quad (9)$$

where $f(\mathbf{X}_t)$ is the output of a function that captures price dependency, and ϵ_t is a temporally correlated residual series detailing short-term price trends. This residual series could be modelled by any univariate time series technique, such as ARMA or ARIMA. However, we found equivalent results from simple AR models. The final forecast of the price k steps ahead of the current time τ would be

$$\hat{y}_{\tau+k} = \hat{f}(\mathbf{X}_{\tau+k}) + \hat{\epsilon}_{\tau+k}, \quad (10)$$

where $\hat{f}(\mathbf{X}_{\tau+k})$ is the prediction from a regression model, and $\hat{\epsilon}_{\tau+k}$ represents the predicted residual value from autoregression. Informally, when prices are above/below our baseline forecasts at the beginning of the prediction cycle (24-h block), we expect prices to remain above/below these levels for some time. All AR models considered are mean (zero) reverting such that $\hat{\epsilon}_{\tau+k} \rightarrow 0$ as $k \rightarrow \infty$. This means that our forecast for $\hat{y}_{\tau+k} \rightarrow \hat{f}(\mathbf{X}_{\tau+k})$ as $k \rightarrow \infty$. Che and Wang (2010) utilised the AR residual values to increase the robustness of parameter variation in point forecasts. However, to the best of our knowledge, this has not yet been applied to probabilistic forecasting, which has multiple forecast ranges. Thus, for the augmentation of our quantile predictions, we choose to define this residual series, ϵ_t , as the difference between observed prices and the post-smoothing median prediction ($q =$

0.5) for that given model. When making the AR-adjusted predictions for k steps ahead of the current time, we vertically shift the entire set of predicted quantiles from a constituent model by the value $\hat{\epsilon}_{\tau+k}$.

$$\hat{y}_{\tau+k,q}^p = \hat{y}_{\tau+k,q}^s + \hat{\epsilon}_{\tau+k}, \quad (11)$$

where $\hat{y}_{\tau,q}^p$ is the post-processing and $\hat{y}_{\tau,q}^s$ the post-smoothing of $\hat{y}_{\tau,q}$ from the constituent model. For example, the vertical shift of all post-smoothing quantile predictions in a QRF model is determined by the AR residual prediction of the post-smoothing 0.5 quantile prediction of that QRF. Intuitively, our quantile predictions $\hat{y}_{\tau+k,q}^s$ describe the density of the prices, with the short-term autoregression simply re-centring the quantile predictions with more recent information. As discussed in Section 3.1, we elect to make 24-h-ahead forecasts daily at midnight; hence, the re-centring occurs daily at midnight.

Overall, this combination of post-processing approaches allows us to combine a non-inherently time series model (such as our QRF) for long-term, non-recursive forecasting with a simple, inherently temporal model for short-term trends.

4. Forecast averaging with quantile regression

Numerous ensemble methods have been tested within the EPF literature. Among the point forecast methods, simple averaging is a common and effective method. Ordinary least squares (OLS) averaging was introduced by Crane and Crotty (1967), with a follow-up paper that generated significant research traction in this area (Granger & Ramanathan, 1984). Many papers sought to address the limitations of this method, such as the volatility of beta estimates due to the serial correlation of errors and non-robust (\mathcal{L}_2) fitting routine. The simplest remedy to the second problem was the LAD averaging developed by Nowotarski et al. (2014), who eventually developed quantile regression averaging. LAD averaging can be seen as the first implementation of QR for averaging, as this is equivalent to a 0.5 quantile regression. Our research follows the direction of these developments, and the main ensemble tool in this study is quantile regression.

4.1. QRA model

Quantile regression averaging (QRA) is a probabilistic forecast combination method introduced by Nowotarski and Weron (2015). It estimates conditional quantiles by applying QR to a set of point forecasts. The improved performance of QRA against its constituent models was verified with several follow-up papers (Kath & Ziel, 2021; Kostrzewski & Kostrzewska, 2019; Maciejowska, Nowotarski, & Weron, 2016; Nowotarski & Weron, 2018; Uniejewski & Weron, 2021). However, its most significant success came from the GEFCom2014 forecasting competition, where the top two winning teams for the price track used some variant of QRA (Gaillard, Goude, & Nedellec, 2016; Maciejowska & Nowotarski, 2016).

Mathematically, the model is identical to traditional quantile regression. However, rather than a set of explanatory variables such as temperature and seasonal features, we use the point forecasts \hat{y}_t from multiple predictive models as the explanatory variables:

$$\hat{y}_{t,q} = \hat{y}_t^T \hat{\beta}_q^{\text{QRA}}, \quad (12)$$

where $\hat{y}_t = (\hat{y}_{t,1}^p, \dots, \hat{y}_{t,n}^p)^T$ from our n point prediction models, and $\hat{\beta}_q^{\text{QRA}}$ represents the parameter vector for the QRA model. In QRA, the point predictions are the outputs from the SVM and the 0.5 quantile forecasts (median) of each quantile regression model after post-processing, i.e., including smoothing and autoregression.

4.2. Q-QRA model

To arrive at a probabilistic combination tool capable of handling nonlinear volatility features, we extend the model input to include quantile forecasts from the constituent models. Rather than combining point forecasts for probabilistic prediction, this model is a direct combiner of probabilistic forecasts. To do so, the prediction process is mathematically identical to QRA. However, instead of supplying point estimates, we provide estimates of the desired quantile from other probabilistic forecasting models:

$$\hat{y}_{t,q} = \hat{y}_t^T \hat{\beta}_q^{\text{Q-Q}}, \quad (13)$$

where $\hat{\beta}_q^{\text{Q-Q}}$ represents the parameter vector for the Q-QRA model. This specification deviates from most ensemble algorithms, since now each input contains a different feature set for each q -quantile ($\hat{y}_t \rightarrow \hat{y}_{t,q}$). A review paper on probabilistic forecast averaging tested many quantile forecast combination routines and found that the above specification (labelled QRA-T) achieved the best results (Wang et al., 2019). Our proposed method demonstrates the generalisation capacity of this framework and shows that we can also gain the benefits found in traditional QRA by including point forecasting models. Hence, $\hat{y}_{t,q} = (\hat{y}_{t,q,1}^p, \dots, \hat{y}_{t,q,p}^p, \hat{y}_{t,1}^p, \dots, \hat{y}_{t,n}^p)^T$ consists of p quantile prediction models and n point prediction models. Note that the point forecast inputs of our Q-QRA model only include the output values from the SVM.

Using QR for amalgamating probabilistic forecasts seeks to address the gap between forecast combination techniques and the increasingly complex individual methods. We observed that many papers generate sets of predictive quantiles from nonlinear models, only to use simple arithmetic averaging to make final predictions. This fails to use the relative performance and correlation of each model in a combining methodology. This issue is addressed in our model. During our testing phase, the QRA and Q-QRA are re-trained every day at midnight, with one month of trailing training data.

5. Selection of training period length

The highly dynamic and non-stationary nature of NEM electricity prices leads to a myriad of problems for forecasters, one of which is the selection of the appropriate

training period lengths. Longer training periods contain a larger volume of information upon which predictions can be obtained. However, in periods where market behaviour departs from historical performance, e.g., due to changes in the market rules, the longer training period will contain mainly inaccurate information for current predictions.

We see this as a realisation of the common variance–bias trade-off. Short-term models display greater reactivity to market change and shifting variable relationships. The cost of this reactivity is a greater tendency to overfit due to the reduced sample size. Longer-term models have a larger sample size and display less overfitting. However, if there are significant changes to the observed market behaviour (such as the consequences of Covid-19), the models will react slowly, thus showing poor performance during this adjustment period. In this sense, short-term models can be considered to be low-bias, while longer-term models are low-variance.

With the rapid shift of the energy generation mix to a greater emphasis on renewable energy and the introduction of increasingly complex technologies into the power grid, such as distributed energy resources, we expect many such structural changes to occur in the coming years. Considering these conditions, we suggest that EPF practitioners keep in mind this inherent trade-off and consider systematic solutions such as those detailed in this paper.

Our proposed method for this problem is the combination of models across different training sets. By including models of both short and long training periods in our ensemble, we can leverage the benefits of both stability and reactivity in our forecasts. To do so, in our rolling prediction routines, the QRF models are trained on historical data of one month (Short-QRF), three months of data (Medium-QRF), or one year of data (Long-QRF). The rolling re-training nature of our combining methods ensures that during periods of change, where shorter-term models excel, a larger proportion of the forecast will be attributed to these models (larger co-efficient in the Q-QRA). Conversely, during periods of stability, we expect a higher allocation to the longer-term models. We also expect that different quantiles will have different relationships between accuracy and the length of training data. Under this method, each quantile can use the training data that leads to the most accurate prediction for that quantile.

Recent research results indicate that there are accuracy gains to be found under these kinds of combination methods, with results showing improved accuracy for simple arithmetic averaging, as well as weighted average schemes based on model performance (Hubicka et al., 2019; Marcjasz et al., 2018; Serafin et al., 2019). Our implementation relies on inserting these different window models into our already existing model combination scheme, Q-QRA, thus evading additional steps/hyperparameters while retaining the performance of different training-length combinations.

6. Forecasting results

6.1. Evaluation methods

In this paper, we follow the probabilistic forecast evaluation criteria discussed by Gneiting, Balabdaoui, and

Raftery (2007) and Pinson, Nielsen, Møller, Madsen, and Kariniotakis (2007). Namely, we seek to maximise sharpness subject to calibration. Sharpness refers to the concentration of the forecast distributions, with sharper distributions providing narrower prediction intervals. Calibration, or reliability, ensures that nominal coverage is consistent with empirical coverage. As explained by Gensler, Sick, and Vogt (2018), an important attribute of a successful scoring method is whether it is a ‘proper’ scoring method. Mathematically, this means that the expected loss under this scoring process is minimised by predictions that are exactly equal to the probabilities/quantiles found in the true data-generating distribution. The pinball loss function in (2) is a proper scoring method (Gensler et al., 2018), and we use it to compare the performance of our predictive models in an out-of-sample manner. Please note that although there are better scoring methods for point predictors, we use the pinball loss function with a 0.5 quantile for the SVM for consistency with the quantile predictors. The 0.5 quantile pinball loss can also be regarded as half of the mean absolute error.

Another metric that can be used to evaluate the sharpness and reliability of the probabilistic forecast is the continuous ranked probability score (CRPS) (Gneiting & Raftery, 2007; Hersbach, 2000; Lauret, David, & Pinson, 2019; Nowotarski & Weron, 2018). This score calculates the difference between the cumulative distribution functions (CDFs) of the observation and the prediction:

$$\text{CRPS} = \int_{-\infty}^{\infty} (\hat{F}_t(u) - \mathbf{1}_{(u \geq y_t)})^2 du, \quad (14)$$

where $\hat{F}_t(u)$ denotes the predictive CDF constructed using the nine key quantiles analysed in this study. It is derived based on the classic CDF approach described by Lauret et al. (2019).

While both the pinball loss function and CRPS can be used to jointly measure the sharpness and reliability of the quantile forecast, it is necessary to consider a more intuitive metric when assessing the quality of the proposed model. To this extent, we consider the prediction interval coverage probability (PICP) (Lauret et al., 2019), which allows us to measure the actual coverage of the prediction interval (PI). Each PI is determined as the interval between the upper and lower quantiles centred around the median. For example, the 95% PI is defined by the upper 0.975 quantile and the lower 0.025 quantile. Consequently, we want the outcome of the PICP to be as close as possible to the nominal coverage rate. For a $\text{PI}^{(\alpha)}$ with a nominal coverage rate $(1 - \alpha)100\%$, the PICP score is given by

$$\text{PICP} = \frac{1}{H} \sum_{t=1}^H c_t, \quad (15)$$

where

$$c_t = \begin{cases} 1, & y_t \in \text{PI}^{(\alpha)} \\ 0, & y_t \notin \text{PI}^{(\alpha)} \end{cases}, \quad (16)$$

and H represents the number of test samples.

In the following subsections, we first investigate the effects of the proposed spike treatment on the constituent

predictive models by comparing the pinball loss evaluations of the raw and processed series. Then, we show that the prediction can be improved by applying smoothing and AR augmentation on the outputs of each model. Also, we evaluate the effect of the dynamic training length at different periods over the considered horizon. Lastly, we show the performance of different ensemble routines and draw a comparison with the available forecast model from AEMO.

6.2. Spike processing

To demonstrate the harmful effects of extreme price spikes and investigate the benefits of our spike processing routines, we train all models twice, once on the raw data and then on the processed price series. Fig. 3 shows the relative improvements when training on the processed prices compared with the raw counterparts. The results demonstrate that even when using inherently robust methods such as linear quantile regression, the presence of extreme spikes in the training data reduces the accuracy of predictions within most quantiles. An observed exception to this is the extremely lower quantiles, where the raw models achieved a higher degree of accuracy (negative improvement values). This is due to the frequent occurrence of negative spikes in the later years, thus entering our target quantiles. Overall, these results indicate a moderate but consistent benefit to extreme spike filtration. For those desiring a wider target quantile range, it is perhaps best to use the filtered series for the inner quantiles and reserve the original, spike-inclusive series for the extreme ranges. However, for consistency, the remaining training routines are performed using the filtered series.

6.3. Smoothing and autoregression

Upon obtaining the outputs of the constituent models, the prediction values are smoothed out from a one-hour-window moving average and vertically shifted using AR augmentation. The effects of this post-processing step can be seen in Fig. 4, which shows minor improvements for all quantile levels, generally around 1%–2.5%.

6.4. Constituent models

Table 2 shows the average pinball losses of each constituent model at different quantile levels from the four-year testing period. At first glance, the Short-QRF model, which is trained on one month of rolling data, appears to have the worst prediction accuracy. Especially when compared to the Long-QRF model, with the one-year rolling model outperforming in all tested quantiles. However, Fig. 5 shows that the superiority of the Long-QRF model is simply a result of a lower average value across the entire assessment period. It can be seen that the short-memory model systematically beat the longer-term model at around the beginning of 2020, which was the start of Covid-19. Then later that year, once sufficient data were gathered to ‘learn’ about the impact of the pandemic, the long-memory model performed slightly

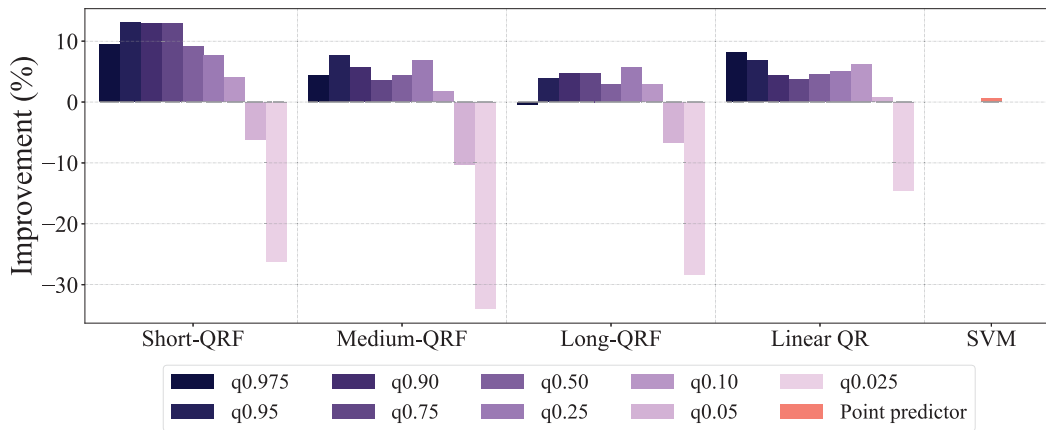


Fig. 3. Improvements in pinball losses (% reduction) of constituent models after removing spikes using the pre-processing steps. The colour gradient indicates the different quantiles, read left to right in descending quantile order.

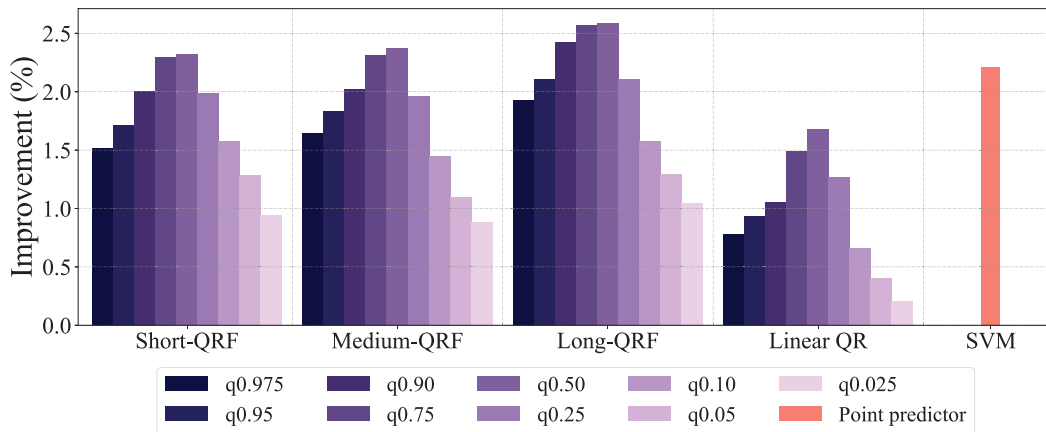


Fig. 4. Improvements in pinball losses (% reduction) of constituent models after the post-processing steps of AR adjustment and forecast smoothing. The colour gradient indicates the different quantiles, read left to right in descending quantile order.

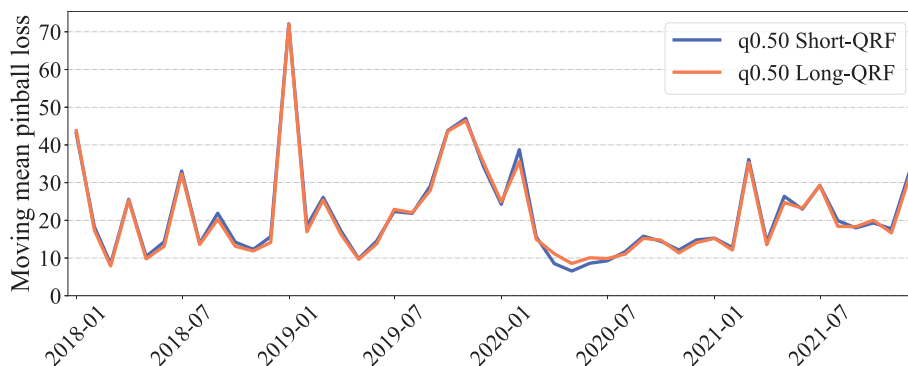


Fig. 5. Pinball loss 28-day moving average for 0.5 quantile (median) forecasts by the Short-QRF and Long-QRF models, spanning from 1/1/2018 to 31/12/2021.

more accurately than the shorter model. Therefore, despite common sentiment in non-time-series-based regression, more data does not necessarily imply more accuracy.

These findings suggest that the benefit of combining models with different training lengths may not simply be the variance reduction seen in traditional ensemble methods but an increased capacity for adaptation when

Table 2

Mean pinball loss associated with the probabilistic forecasts of each constituent model. Boldface values indicate the top-performing model for a specific quantile.

Quantile	Linear QR	Short-QRF	Medium-QRF	Long-QRF
0.025	9.87	10.19	10.06	10.17
0.05	11.36	11.71	11.56	11.63
0.10	13.64	13.92	13.71	13.80
0.25	17.96	18.25	17.90	18.01
0.50	20.99	21.55	21.06	21.18
0.75	20.31	20.78	20.34	20.41
0.90	17.11	17.15	16.91	16.88
0.95	14.85	14.74	14.54	14.49
0.975	12.98	12.94	12.73	12.67
SVM	22.94			

Table 3

Mean pinball loss associated with the probabilistic forecasts of each averaging routine, together with the score for the ‘best constituent’, representing the highest-performing individual model for each quantile (the boldface values from Table 2). Boldface values indicate the top-performing model for a specific quantile.

Quantile	Best constituent	QRA	Q-QRA
0.025	9.87	10.0	9.59
0.05	11.36	11.52	11.14
0.10	13.64	13.72	13.29
0.25	17.90	17.69	17.45
0.50	20.99	20.42	20.42
0.75	20.31	19.76	19.68
0.90	16.88	16.82	16.49
0.95	14.49	14.63	14.25
0.975	12.67	12.86	12.48

significant changes occur in the market. An ensemble routine would ideally leverage longer-memory models during periods of stability and shift model weights towards shorter training periods under periods of market changes.

In addition, among the constituent models we see in Table 2, the nonlinear QR models (QRF) achieve the best performance in the upper outer quantiles (0.9, 0.95, and 0.975 quantiles) where linear QR struggles—whereas, by a small margin, linear QR achieves the highest scores at the 0.50 and 0.75 quantiles. Although the low quantiles display results favouring linear QR, these are somewhat skewed due to the frequent occurrence of negative spikes, which are not accurately captured due to the process of spike filtration. Overall, the relative performance of these models supports our hypothesis that much of the nonlinearities occur specifically in the outer distributional quantiles.

6.5. Different ensemble routines

To investigate the possible gains from averaging over the models, we compare the probabilistic forecasts from two ensemble models (QRA & Q-QRA) against the best-performing constituent model for each quantile. Table 3 displays the average pinball loss for these ensembles and constituents. Since the QRA routine only leverages the median forecasts from linear and QRF models, as well as point predictions from constituent models, the improvements are generally concentrated at the inner quantiles

Table 4

Continuous ranked probability score (CRPS) of the probabilistic forecasts for the individual models together with the averaging routines. The boldface value indicates the top-performing model.

Short-QRF	Medium-QRF	Long-QRF	Linear QR	QRA	Q-QRA
39.26	38.91	38.51	38.14	37.17	37.02

Table 5

Number of daily dispatch periods (out of 288) where the Diebold–Mariano test revealed statistically significant discrepancies in the forecasting accuracy between the QRA and Q-QRA models at a 5% significance level.

Quantile	0.025	0.05	0.1	0.25	0.5	0.75	0.9	0.95	0.975
# Significant	218	228	269	204	–	57	91	160	194

and fail to capture the nonlinearities at the outer quantiles, thus resulting in worse predictions. For the Q-QRA model, by contrast, the results are commensurate with much of the literature surrounding the benefits of forecast averaging. The proposed Q-QRA model attains the minimum pinball losses at all investigated quantiles. A similar pattern is observed in Table 4 for the CRPS metric, with the Q-QRA model again achieving the lowest losses.

To validate these observations, we used the Diebold–Mariano test (Diebold & Mariano, 1995) to evaluate the statistical significance of the disparity in pinball loss observed in Table 3 between the Q-QRA and QRA methods. We found that for all non-median quantiles assessed in Table 3, the difference in forecast performance was statistically significant at a 5% level of significance in a one-sided test (the median forecasts are identical across these models by construction). Beyond this aggregated test that assessed all dispatch periods together, we also ran the test on each dispatch period individually. Table 5 records the number of daily dispatch periods (out of 288) where the Q-QRA model had a statistically significantly lower pinball loss. There is notable variation in the number of passed periods across the relevant quantiles. However, even with this variation, the minimum number of periods passed greatly exceeds what would be expected if the QRA and Q-QRA models had identical forecast performance. Under the null hypothesis of equal forecast accuracy, the Q-QRA is expected to pass between seven and 22 periods with 95% confidence. This range is derived from $np \pm 1.96\sqrt{np(1-p)}$, which is determined using a normal approximation for the binomial result (either passing or not passing within each of the 288 dispatch periods). The lower half of the quantiles had a greater number of periods passing the test, with the upper half of the quantiles showing a larger performance difference for the outer quantiles. These results are consistent with Table 3, with the number of passed periods generally relating to the size of the difference in the pinball score.

In addition to pinball losses, we also compare the coverage performance of each ensemble routine. The relative performance, indicated by the PICP shown in Table 6, is consistent with the pinball loss results. Overall, the results show that empirical coverage was below target for all models and prediction intervals. Nonetheless, our proposed Q-QRA model consistently achieved the smallest

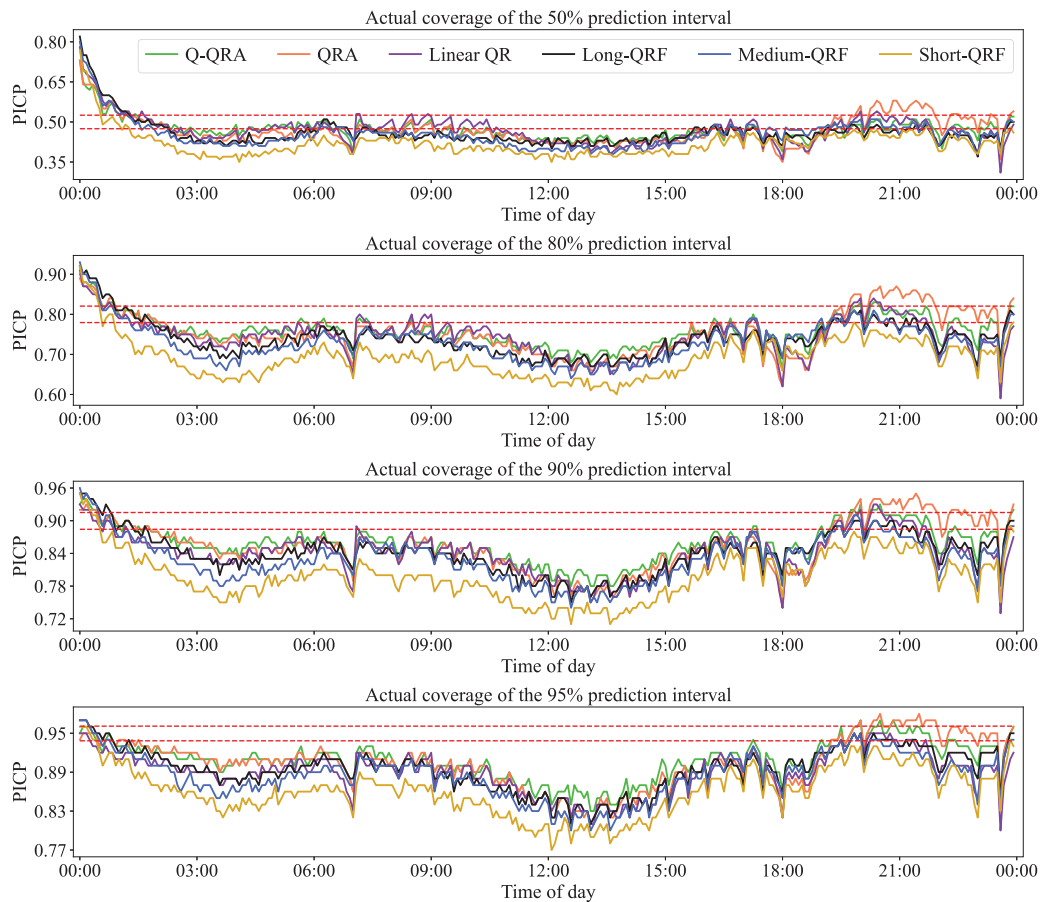


Fig. 6. Prediction interval coverage probability (PICP) for all models across the 288 dispatch periods in the 24-h cycle. The horizontal red dashed lines serve as visual aids, indicating consistency bars associated with a 95% confidence level. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 6

Prediction interval coverage probability (PICP) of the probabilistic forecasts for the individual models together with the averaging routines. Boldface values indicate the top-performing model for a specific interval.

PI	Linear QR	Short-QRF	Medium-QRF	Long-QRF	QRA	Q-QRA
50%	0.469	0.420	0.454	0.460	0.470	0.470
80%	0.746	0.693	0.732	0.741	0.756	0.759
90%	0.841	0.800	0.831	0.841	0.854	0.858
95%	0.891	0.862	0.886	0.895	0.905	0.908

gap between the PICP score and the nominal coverage. Fig. 6 displays the PICP for these intervals throughout the day to see whether there is seasonality in the reliability metrics. To provide intuition around the scale of these trends, the plots include consistency bars with a 95% confidence level (Bröcker & Smith, 2007). These bars indicate the expected variation of the PICP from the nominal coverage due to the inherent randomness caused by finite test data. Given that the observations within each specific five-minute interval are spaced 24 h apart, we consider low serial correlation for these intervals. Consequently, we calculate the consistency bars using the binomial distribution, which is identified as the null distribution for

non-serially correlated coverage indicators (Pinson, McSharry, & Madsen, 2010). Note that these lines only serve as visual aids; a formal evaluation of the coverage is provided through the Kupiec results detailed below. We see that reliability is excessively high at the start of the 24-h cycle as the AR boosting ‘re-centres’ the density, which likely improves coverage. The quantiles are not trained conditional on the AR improvements. Hence, they are unaware of the excess coverage during these periods and could have shrunk the quantiles inward to reflect the increased accuracy of this short-term prediction augmentation. (Post-processing solely affects reliability, without altering sharpness.) During the rest of the dispatch period, the PICP is consistently below target, roughly matching our previous PICP results that were aggregated across time periods. To formally analyse reliability across the dispatch periods, we conduct a two-sided hypothesis test to assess whether empirical coverage is equal to nominal coverage. Specifically, we employ the Kupiec (1995) coverage test at the 5% level of significance for each quantile (0.025, 0.05, 0.1, 0.25, 0.5, 0.75, 0.9, 0.95, and 0.975) and for every dispatch period across all models. For the Q-QRA model, the test was passed in 15%–47% (27.8% on average) of the dispatch periods across the nine quantiles.

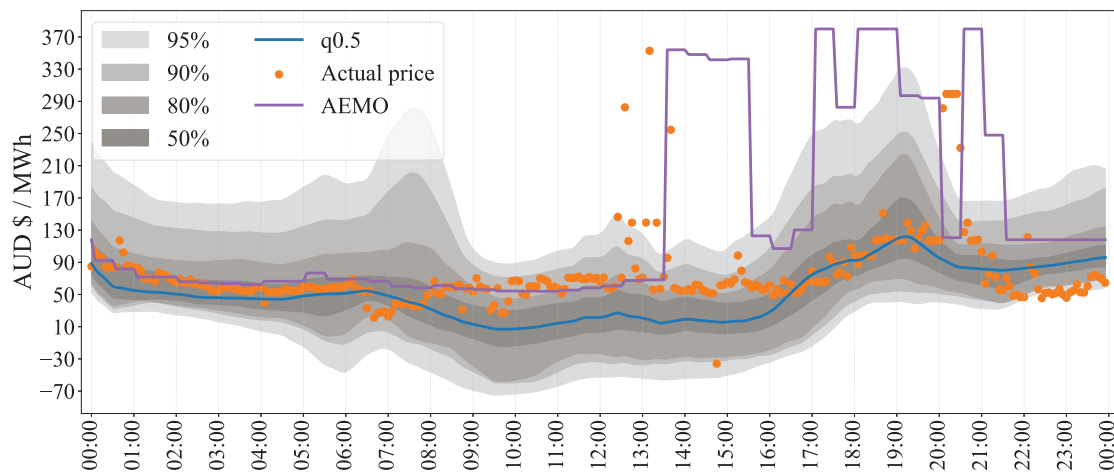


Fig. 7. Electricity price forecasts in the South Australia region for December 1, 2021. The AEMO's predicted values are presented as a solid purple line, while the Q-QRA model's median forecast is marked with a blue line. The predictive density is illustrated through various degrees of shading. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

For the QRA model, the test was passed in 15%–44% (26.7% on average) of the cases, which approximately coincides with the 15%–20% cases (across 24 h and 90% or 50% PI) reported in Table 1 of [Uniejewski and Weron \(2021\)](#) for the Polish and Scandinavian markets. The forecasts for the constituent models were generally lower, particularly for Short-QRF, which achieved the lowest score. In general, the test was passed more times for higher quantiles than for lower ones. Overall, all forms of reliability analysis suggest systematic under-coverage for all model types. However, the level of this under-coverage approximately matches other instances of EPF coverage analysis. This suggests that coverage remains a unique challenge and stands as a strong candidate for future research. In terms of the relative performance of reliability metrics, the Q-QRA can be seen generally to achieve the highest reliability for most of the prediction cycles and intervals.

6.6. AEMO pre-dispatch price comparison

In addition to providing the spot prices from the historical dispatch intervals, AEMO must prepare and publish a pre-dispatch schedule containing 30-minute pre-dispatch (forecast) wholesale prices until the end of the next trading day. According to [AEMO \(2021a\)](#), these pre-dispatch prices represent the expected bids at the last trading interval in each 30-minute period; i.e., pre-dispatch prices are only available at 00:00, 00:30, etc. Although these 24-h-ahead forecasts are updated every half hour, we select and evaluate the day-ahead forecasts made at 00:00 every day to align with the prediction horizon of our Q-QRA implementation. Note that our proposed model can perform a rolling forecast every 5 min. However, as mentioned above, our results are based on daily predictions, so that all 24 h of predictions are included in the evaluation metrics. Due to the availability of prediction data from AEMO, only 12 months of data from the beginning of 2021 until the end of the year are used.

Table 7

Mean and median absolute prediction error for the AEMO pre-dispatch forecast and the 0.5 quantile (median) forecast from the proposed Q-QRA model in 2021.

Predictor	Mean AE	Median AE
AEMO pre-dispatch	223.82	18.17
Q-QRA median	41.12	17.26

[Table 7](#) shows the comparative performance between the point forecast from AEMO and the median forecast of the proposed model in terms of the mean absolute error (AE) and median AE (in AUD\$/MWh). The results show that our median forecast outperforms AEMO in the half-hourly prediction period. The high mean AE of the prediction of AEMO can be explained in [Fig. 7](#), which displays the probabilistic electricity price forecasts for the next 24 h on the first day of December 2021. It can be seen that although the AEMO model (purple line) thrives in the first few hours, its predictions become significantly unstable in the second half of the forecast horizon. As a result, the AEMO model has a mean AE 12 times higher than its median AE and five times higher than the mean AE of our point forecast (blue line).

7. Economic evaluation

Although achieving a lower prediction error is an important goal in price prediction, the true impact of the accuracy of the forecast must be evaluated in decision-making processes ([Hong et al., 2020](#)). For example, [Chitsaz, Zamani-Dehkordi, Zareipour, and Parikh \(2018\)](#) developed a mixed-integer linear programming (MILP) problem for battery operation using price forecasts. Specifically, the model was developed to maximise the savings for a microgrid through energy arbitrage of a central battery system. In our model, however, instead of considering a central battery, we assume that each household has a home battery system, e.g., a Tesla Powerwall

Table 8
Optimisation parameters.

Γ	90	Battery round-trip efficiency (%)
E^{\min}	0	Battery minimum energy (kWh)
E^{\max}	13.5	Battery maximum energy (kWh)
E^{init}	0	Battery initial energy (kWh)
P^{\max}	5	Battery power rating (kW)
Δt	0.5	Time interval resolution (h)
λ_t	0.097	Network usage charge (AUD\$/kWh)

(Tesla, 2022). This is to ensure that we consider variability in different energy profiles.

In Australia, there are several newly emerged electricity retailers, e.g., Amber Electric (Electric, 2022), allowing residential prosumers to buy and sell electricity at spot prices. Therefore, in our model, we consider residential households to be the decision makers that operate their battery system to minimise their electricity bill. To provide a simple case study, we generate price forecasts once a day at midnight for one day ahead, making it easier for the prosumers to manage and monitor their home battery system. Since price forecasts are updated at midnight, the optimisation model only needs to be solved once a day to find the optimal scheduling of the home battery for the following day.

7.1. Prosumers' electricity cost minimisation problem

We assume that each household has a rooftop solar photovoltaic (PV) system, and that the generated electricity can be used to satisfy household consumption, stored in the home battery, or sold back to the grid. Using the home battery, electricity prosumers can exercise energy arbitrage not only to reduce their electricity bill but also to make a profit by selling electricity back to the grid during high price intervals. The optimisation model for each prosumer for each day is as follows:

$$\min_{\theta} \text{Cost} = \sum_{t \in T} \delta_t (n_t^+ - n_t^-) + \lambda_t n_t^+ \quad (17a)$$

s.t.

$$E_t = E^{\text{init}} + \sum_{h=1}^t (P_h^{\text{ch}} - \frac{1}{\Gamma} P_h^{\text{dis}}) \Delta t \quad \forall t \in T, \quad (17b)$$

$$P_t = P_t^{\text{ch}} - P_t^{\text{dis}} \quad \forall t \in T, \quad (17c)$$

$$P_t^{\text{ch}} \leq M P_t^{\text{b}} \quad \forall t \in T, \quad (17d)$$

$$P_t^{\text{dis}} \leq M(1 - P_t^{\text{b}}) \quad \forall t \in T, \quad (17e)$$

$$E^{\min} \leq E_t \leq E^{\max} \quad \forall t \in T, \quad (17f)$$

$$-P^{\max} \leq P_t \leq P^{\max} \quad \forall t \in T, \quad (17g)$$

$$\sum_{t \in T} P_t^{\text{ch}} \leq E^{\max}, \quad (17h)$$

$$d_t - g_t + P_t \Delta t = n_t^+ - n_t^- \quad \forall t \in T, \quad (17i)$$

$$n_t^+ \leq M n_t^{\text{b}} \quad \forall t \in T, \quad (17j)$$

$$n_t^- \leq M(1 - n_t^{\text{b}}) \quad \forall t \in T, \quad (17k)$$

$$E_t, P_t, P_t^{\text{ch}}, P_t^{\text{dis}}, n_t^+, n_t^- \geq 0 \quad \forall t \in T, \quad (17l)$$

$$P_t^{\text{b}}, n_t^{\text{b}} \in \{0, 1\} \quad \forall t \in T, \quad (17m)$$

where $\theta = \{E_t, P_t, P_t^{\text{ch}}, P_t^{\text{dis}}, n_t^+, n_t^-, P_t^{\text{b}}, n_t^{\text{b}}\}$, and T represents the set of time intervals for one day ahead. The first term in (17a) represents the energy payment at wholesale prices, δ_t , for the net energy exchanged with the grid. If the net energy, $n_t^+ - n_t^-$, is positive, the prosumer buys from the grid. Conversely, if the net energy is negative, the prosumer gets paid at wholesale prices for selling electricity. In addition to paying for electricity at wholesale prices, prosumers are responsible for paying network usage charges, λ_t , as shown in the second term in (17a), which is only applied on the imported energy (Ausgrid, 2022; Electric, 2022). Constraint (17b) represents the energy stored, E_t , in the battery over the course of one day. Similar to Dinh et al. (2022), we consider a round-trip efficiency, Γ , on the battery (dis)charging cycles. We define the charging, P_t^{ch} , and discharging, P_t^{dis} , power in (17c). To avoid simultaneous charging and discharging, we consider complementarity constraints utilising a binary variable, P_t^{b} , and a sufficiently large constant, M , in (17d) and (17e). Constraint (17f) limits the stored energy within the battery's nominal capacity. Constraint (17g) restricts the maximum (dis)charging power of the battery to be lower than its rated power. To mitigate battery degradation caused by excessive (dis)charging, constraint (17h) limits the battery to a maximum of one cycle per day. Constraint (17i) separates the household consumption, d_t , solar PV generation, g_t , and the (dis)charging energy of the home battery, $P_t \Delta t$, into imported n_t^+ and exported n_t^- electricity. Similar to the battery (dis)charging power, we restrict n_t^+ and n_t^- from simultaneously taking non-zero values using complementarity constraints in (17j) and (17k). Lastly, we define the sign and type of the variables in (17l) and (17m), respectively.

In this optimisation model, wholesale prices are the main source of uncertainty and can be obtained using the proposed probabilistic forecasts. While there are other sources of uncertainty, such as household consumption, for an easy evaluation of the impact of price uncertainty, we assume that other uncertain parameters are perfectly known. To this end, we convert the objective function in (17a) to an expected cost minimisation problem using expected prices (Lemos-Vinasco, Bacher, & Møller, 2021; Zareipour, Canizares, & Bhattacharya, 2010), as follows:

$$\min_{\theta} E[\text{Cost}] = \sum_{t \in T} E[\delta_t] (n_t^+ - n_t^-) + \lambda_t n_t^+. \quad (18)$$

If the optimisation is solved using point forecasts, as in AEMO pre-dispatch prices, we can directly substitute the expected prices with these point forecasts. However, to obtain the expected prices from the probabilistic forecasts, we leverage the predictive CDF, outlined in Section 6.1, to generate a set of equiprobable samples. We then calculate the expected price as follows:

$$E[\delta_t] = \sum_{s \in S} (\pi_s \hat{\delta}_{t,s}), \quad (19)$$

where S represents the set of price forecast samples, and π_s denotes the probability of sample $s \in S$. The solutions obtained from the proposed MILP problem are the optimal scheduled power of the battery systems and the net energy exchanged with the grid for each interval.

Table 9

Comparison of weekly average electricity costs across ten prosumers. Negative values indicate profit, with 'Perfect prediction' indicating the theoretical upper bound achievable with perfect prediction of wholesale prices. Boldface values indicate the most profitable prediction scheme for each prosumer.

Prosumer #	Perfect prediction	AEMO pre-dispatch	Q-QRA median	QRA	Q-QRA
1	−6.31	−0.16	−1.45	−1.55	−1.45
2	8.61	14.90	13.92	13.90	13.82
3	−1.81	4.63	3.19	3.11	3.06
4	−6.51	0.04	−0.95	−1.02	−1.12
5	−4.43	2.06	0.95	0.89	0.83
6	0.27	6.69	5.30	5.19	5.18
7	−5.14	1.38	0.12	0.04	−0.03
8	−3.14	3.21	2.21	2.13	2.05
9	−1.62	4.61	3.74	3.69	3.64
10	−2.73	3.70	2.83	2.70	2.66

We calculate the ground-truth electricity cost of each household as follows:

$$\text{Cost}^* = \sum_{t \in T} \delta_t(n_t^{*+} - n_t^{*-}) + \lambda_t n_t^{*+}, \quad (20)$$

where n_t^{*+} and n_t^{*-} represent the realised imported energy and realised exported energy, respectively. We use the ground-truth results for comparison.

7.2. Optimisation setup

- **Energy profiles** of ten prosumers were obtained from the Solar Home dataset containing half-hourly solar PV generation and electricity consumption data (Ausgrid, 2012). Since the data resolution is 30 min, we consider the price forecasts at the last trading interval of each 30-minute period, similar to in Section 6.6.
- The **optimisation period** is the first week of each month in 2021 (12 weeks in total), in which we obtain energy profiles of the prosumers, wholesale electricity prices, and network usage charges (Ausgrid, 2022).
- **Battery data** were sourced from the Tesla Powerwall user manual.² We summarise the battery and other optimisation parameters in Table 8.

7.3. Optimisation results

To evaluate the economic benefit of the proposed Q-QRA model, we compare the ground-truth electricity cost of all prosumers using five different sets of wholesale price forecasts, namely, wholesale prices from perfect prediction, AEMO pre-dispatch prices, median forecasts of the Q-QRA model, quantile forecasts from the QRA model, and quantile forecasts from the Q-QRA model. For each set of probabilistic price forecasts, we use the predictive CDF to generate 100 equally probable expected price samples. Hence, the probability of each sample in (19) is set to $\pi_s = 0.01$.

Table 9 compares the weekly average electricity cost of ten prosumers in different price forecast models. Since

the model using perfect prediction of wholesale prices would always produce the lowest electricity cost, for an easy comparison, we highlighted the second-lowest electricity cost. The negative cost values indicate the profit of prosumers under the proposed retailer scheme. This profit is achieved mainly by using the home battery to store electricity when prices are low and to export electricity back to the grid when prices are high. Note that prosumers receive payments from the retailer for exporting excess electricity to the grid (Electric, 2022). The results in Table 9 show that the last three forecast models outperform the AEMO model, with the proposed Q-QRA model providing the least-cost solutions for most prosumers. Overall, the results demonstrate that using the price forecasts from the proposed Q-QRA model leads to a minimum cost for most prosumers, due to the combined benefits of higher prediction accuracy and consideration of a large range of potential scenarios.

8. Conclusion

In this paper, we demonstrated the use of ensemble forecasting methods to predict the 24-h-ahead electricity price distribution in the SA region of the Australian NEM. We observed that even when predictions are made using traditionally 'robust' models, such as quantile regression, higher accuracy in the inner quantiles can be gained by filtering extreme spikes occurring in the dataset. However, in cases where the extreme outer quantiles are concerned, training on the unfiltered series would be more suitable. Next, we demonstrated two post-processing steps, smoothing and autoregression, that can be applied in a probabilistic setting to attain higher predictive accuracy. We showed that probabilistic forecasting models could be combined across model classes and training data subsets to create a more accurate prediction process. In terms of the ensemble method, the results from the proposed Q-QRA model indicated a systematic improvement at all quantile levels compared to the constituent models. We further evaluated the superiority of the Q-QRA model in a battery scheduling optimisation problem using real-world data from ten prosumers. The results showed that using Q-QRA forecasts led to the lowest electricity costs for most prosumers.

² Tesla Powerwall 2 Datasheet: https://www.tesla.com/sites/default/files/pdfs/powerwall/Powerwall%20AC_Datasheet_en_AU.pdf.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT to improve readability and language. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Acknowledgment

This research was supported by the Adelaide Summer Research Scholarship (ASRS), Faculty of Sciences, Engineering and Technology, University of Adelaide.

References

- AEMO (2021a). *Pre-dispatch: Technical Report*, URL: https://www.aemo.com.au/-/media/files/electricity/nem/5ms/procedures-workstream/stakeholder-consultation/dispatch-procedures/so_op_3704---predispatch---marked-up.pdf.
- AEMO (2021b). *Spot market operations timetable*. URL: https://www.aemo.com.au/-/media/files/electricity/nem/security_and_reliability/dispatch/spot-market-operations-timetable.pdf.
- AEMO (2022). *Market price cap*. URL: <https://www.aemc.gov.au/news-centre/media-releases/2022-23-market-price-cap-now-available>.
- AEMO (2023). *Constraint formulation guidelines*. URL: https://aemo.com.au/-/media/files/stakeholder_consultation/consultations/nem-consultations/2022/cfg-and-scvpf/final/constraint-formulation-guidelines-v12---final_.pdf.
- Afrasiabi, M., Mohammadi, M., Rastegar, M., & Kargarian, A. (2019). Probabilistic deep neural network price forecasting based on residential load and wind speed predictions. *IET Renewable Power Generation*, 13(11), 1840–1848. <http://dx.doi.org/10.1049/iet-rpg.2018.6257>, arXiv:https://ietresearch.onlinelibrary.wiley.com/doi/pdf/10.1049/iet-rpg.2018.6257. URL: <https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-rpg.2018.6257>.
- Ausgrid (2012). *Solar home electricity data*. URL: <https://www.ausgrid.com.au/Industry/Our-Research/Data-to-share/Solar-home-electricity-data>.
- Ausgrid (2022). *Network prices*. URL: <https://www.ausgrid.com.au/Industry/Regulation/Network-prices>.
- Bordignon, S., Bunn, D., Lisi, F., & Nan, F. (2013). Combining day-ahead forecasts for british electricity prices. *Energy Economics*, 35, 88–103. <http://dx.doi.org/10.1016/j.eneco.2011.12.001>.
- Bröcker, J., & Smith, L. A. (2007). Increasing the reliability of reliability diagrams. *Weather and Forecasting*, 22, 651–661. <http://dx.doi.org/10.1175/WAF993.1>, https://journals.ametsoc.org/view/journals/wefo/22/3/waf993_1.xml.
- Chai, S., Xu, Z., & Jia, Y. (2019). Conditional density forecast of electricity price based on ensemble ELM and logistic EMOS. *IEEE Transactions on Smart Grid*, 10(3), 3031–3043. <http://dx.doi.org/10.1109/TSG.2018.2817284>.
- Che, J., & Wang, J. (2010). Short-term electricity prices forecasting based on support vector regression and Auto-regressive integrated moving average modeling. *Energy Conversion and Management*, 51(10), 1911–1917. <http://dx.doi.org/10.1016/j.enconman.2010.02.023>, URL: <https://www.sciencedirect.com/science/article/pii/S0196890410000890>.
- Chitsaz, H., Zamani-Dehkordi, P., Zareipour, H., & Parikh, P. P. (2018). Electricity price forecasting for operational scheduling of behind-the-meter storage systems. *IEEE Transactions on Smart Grid*, 9(6), 6612–6622. <http://dx.doi.org/10.1109/TSG.2017.2717282>.
- Conejo, A. J., Contreras, J., Espinola, R., & Plazas, M. A. (2005). Forecasting electricity prices for a day-ahead pool-based electric energy market. *International Journal of Forecasting*, 21(3), 435–462. <http://dx.doi.org/10.1016/j.ijforecast.2004.12.005>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207004001311>.
- Crane, D. B., & Crotty, J. R. (1967). A two-stage forecasting model: Exponential smoothing and multiple regression. *Management Science*, 13(8), B–501.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263. <http://dx.doi.org/10.1080/07350015.1995.10524599>.
- Dinh, N. T., Pourmousavi, S. A., Karimi-Arpanahi, S., Kumar, Y. P. S., Guo, M., Abbott, D., et al. (2022). Optimal sizing and scheduling of community battery storage within a local market. In *e-Energy '22, Proceedings of the thirteenth ACM international conference on future energy systems* (pp. 34–46). New York, NY, USA: Association for Computing Machinery. <http://dx.doi.org/10.1145/3538637.3538837>.
- Drucker, H., Burges, C. J. C., Kaufman, L., Smola, A., & Vapnik, V. (1996). Support vector regression machines. In M. Mozer, M. Jordan, & T. Petsche (Eds.), *Advances in neural information processing systems*. Vol. 9. MIT Press, URL: <https://proceedings.neurips.cc/paper/1996/file/d38901788c533e8286cb6400b40b386d-Paper.pdf>.
- Electric, A. (2022). *Amber Electric*. URL: <https://www.amber.com.au/>.
- Gaillard, P., Goude, Y., & Nedellec, R. (2016). Additive models and robust aggregation for GEFCom2014 probabilistic electric load and electricity price forecasting. *International Journal of Forecasting*, 32(3), 1038–1050. <http://dx.doi.org/10.1016/j.ijforecast.2015.12.001>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207015001545>.
- Geman, H., & Roncoroni, A. (2006). Understanding the fine structure of electricity prices. *Journal of Business*, 79(3), 1225–1261, URL: <http://www.jstor.org/stable/10.1086/500675>.
- Gensler, A., Sick, B., & Vogt, S. (2018). A review of uncertainty representations and metaverification of uncertainty assessment techniques for renewable energies. *Renewable and Sustainable Energy Reviews*, 96, 352–379.
- Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society. Series B. Statistical Methodology*, 69(2), 243–268. <http://dx.doi.org/10.1111/j.1467-9868.2007.00587.x>, arXiv:https://academic.oup.com/jrsssb/article-pdf/69/2/243/49794500/jrsssb_69_2_243.pdf.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359–378. <http://dx.doi.org/10.1198/016214506000001437>.
- Granger, C. W., & Ramanathan, R. (1984). Improved methods of combining forecasts. *Journal of Forecasting*, 3(2), 197–204.
- Hersbach, H. (2000). Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather and Forecasting*, 15(5), 559–570. [http://dx.doi.org/10.1175/1520-0434\(2000\)015<0559:DOTCRP>2.0.CO;2](http://dx.doi.org/10.1175/1520-0434(2000)015<0559:DOTCRP>2.0.CO;2), URL: https://journals.ametsoc.org/view/journals/wefo/15/5/1520-0434_2000_015_0559_dotcrp_2_0_co_2.xml.
- Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., & Hyndman, R. J. (2016). Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. *International Journal of Forecasting*, 32(3), 896–913. <http://dx.doi.org/10.1016/j.ijforecast.2016.02.001>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207016000133>.
- Hong, T., Pinson, P., Wang, Y., Weron, R., Yang, D., & Zareipour, H. (2020). Energy forecasting: A review and outlook. *IEEE Open Access Journal of Power and Energy*, 7, 376–388. <http://dx.doi.org/10.1109/OAJPE.2020.3029979>.
- Hong, T., Xie, J., & Black, J. (2019). Global energy forecasting competition 2017: Hierarchical probabilistic load forecasting. *International Journal of Forecasting*, 35(4), 1389–1399.
- Hubicka, K., Marcjasz, G., & Weron, R. (2019). A note on averaging day-ahead electricity price forecasts across calibration windows. *IEEE Transactions on Sustainable Energy*, 10(1), 321–323. <http://dx.doi.org/10.1109/TSTE.2018.2869557>.

- Janczura, J., Trück, S., Weron, R., & Wolff, R. C. (2013). Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling. *Energy Economics*, 38, 96–110. <http://dx.doi.org/10.1016/j.eneco.2013.03.013>, URL: <https://www.sciencedirect.com/science/article/pii/S0140988313000625>.
- Kath, C., & Ziel, F. (2021). Conformal prediction interval estimation and applications to day-ahead and intraday power markets. *International Journal of Forecasting*, 37(2), 777–799. <http://dx.doi.org/10.1016/j.ijforecast.2020.09.006>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207020301473>.
- Koenker, R., & Hallock, K. F. (2001). Quantile regression. *Journal of Economic Perspectives*, 15(4), 143–156.
- Kostrzewski, M., & Kostrzewska, J. (2019). Probabilistic electricity price forecasting with Bayesian stochastic volatility models. *Energy Economics*, 80, 610–620. <http://dx.doi.org/10.1016/j.eneco.2019.02.004>, URL: <https://www.sciencedirect.com/science/article/pii/S0140988319300544>.
- Kupiec, P. (1995). *Techniques for verifying the accuracy of risk measurement models: Finance and Economics Discussion Series*, (95–24), Board of Governors of the Federal Reserve System (U.S.), URL: <https://EconPapers.repec.org/RePEc:fip:fedgfe:95-24>.
- Lago, J., Marcjasz, G., De Schutter, B., & Weron, R. (2021). Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 293, Article 116983. <http://dx.doi.org/10.1016/j.apenergy.2021.116983>, URL: <https://www.sciencedirect.com/science/article/pii/S0306261921004529>.
- Lauret, P., David, M., & Pinson, P. (2019). Verification of solar irradiance probabilistic forecasts. *Solar Energy*, 194, 254–271. <http://dx.doi.org/10.1016/j.solener.2019.10.041>, URL: <https://www.sciencedirect.com/science/article/pii/S0038092X19310382>.
- Lemos-Vinasco, J., Bacher, P., & Möller, J. K. (2021). Probabilistic load forecasting considering temporal correlation: Online models for the prediction of households' electrical load. *Applied Energy*, 303, Article 117594. <http://dx.doi.org/10.1016/j.apenergy.2021.117594>, URL: <https://www.sciencedirect.com/science/article/pii/S0306261921009685>.
- Maciejowska, K., & Nowotarski, J. (2016). A hybrid model for GEFCom2014 probabilistic electricity price forecasting. *International Journal of Forecasting*, 32(3), 1051–1056. <http://dx.doi.org/10.1016/j.ijforecast.2015.11.008>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207015001430>.
- Maciejowska, K., Nowotarski, J., & Weron, R. (2016). Probabilistic forecasting of electricity spot prices using Factor Quantile Regression Averaging. *International Journal of Forecasting*, 32(3), 957–965. <http://dx.doi.org/10.1016/j.ijforecast.2014.12.004>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207014001848>.
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, 36(1), 54–74. <http://dx.doi.org/10.1016/j.ijforecast.2019.04.014>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207019301128>. M4 Competition.
- Marcjasz, G., Serafin, T., & Weron, R. (2018). Selection of calibration windows for day-ahead electricity price forecasting. *Energies*, 11, 2364. <http://dx.doi.org/10.3390/en11092364>.
- Meinshausen, N., & Ridgeway, G. (2006). Quantile regression forests. *Journal of Machine Learning Research*, 7(6).
- Nowotarski, J., Raviv, E., Trück, S., & Weron, R. (2014). An empirical comparison of alternative schemes for combining electricity spot price forecasts. *Energy Economics*, 46, 395–412. <http://dx.doi.org/10.1016/j.eneco.2014.07.014>, URL: <https://www.sciencedirect.com/science/article/pii/S0140988314001716>.
- Nowotarski, J., & Weron, R. (2015). Computing electricity spot price prediction intervals using quantile regression and forecast averaging. *Computational Statistics*, 30(3), 791–803. <http://dx.doi.org/10.1007/s00180-014-0523-0>.
- Nowotarski, J., & Weron, R. (2018). Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and Sustainable Energy Reviews*, 81, 1548–1568. <http://dx.doi.org/10.1016/j.rser.2017.05.234>, URL: <https://www.sciencedirect.com/science/article/pii/S1364032117308808>.
- Online, W. W. (2022). Historical weather forecast data. URL: <https://www.worldweatheronline.com/>.
- Pinson, P., McSharry, P., & Madsen, H. (2010). Reliability diagrams for non-parametric density forecasts of continuous variables: Accounting for serial correlation. *Quarterly Journal of the Royal Meteorological Society*, 136(646), 77–90. <http://dx.doi.org/10.1002/qj.559>.
- Pinson, P., Nielsen, H. A., Möller, J. K., Madsen, H., & Kariniotakis, G. N. (2007). Non-parametric probabilistic forecasts of wind power: required properties and evaluation. *Wind Energy*, 10(6), 497–516. <http://dx.doi.org/10.1002/we.230>, arXiv: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/we.230>, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/we.230>.
- Raviv, E., Bouwman, K. E., & Van Dijk, D. (2015). Forecasting day-ahead electricity prices: Utilizing hourly prices. *Energy Economics*, 50, 227–239.
- Serafin, T., Uniejewski, B., & Weron, R. (2019). Averaging predictive distributions across calibration windows for day-ahead electricity price forecasting. *Energies*, 12, 2561. <http://dx.doi.org/10.3390/en12132561>.
- Tesla (2022). Tesla PowerWall. URL: https://www.tesla.com/en_au/powerwall.
- Uniejewski, B., & Weron, R. (2021). Regularized quantile regression averaging for probabilistic electricity price forecasting. *Energy Economics*, 95, Article 105121. <http://dx.doi.org/10.1016/j.eneco.2021.105121>, URL: <https://www.sciencedirect.com/science/article/pii/S0140988321000268>.
- Wang, Y., Zhang, N., Tan, Y., Hong, T., Kirschen, D. S., & Kang, C. (2019). Combining probabilistic load forecasts. *IEEE Transactions on Smart Grid*, 10(4), 3664–3674. <http://dx.doi.org/10.1109/TSG.2018.2833869>.
- Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30(4), 1030–1081. <http://dx.doi.org/10.1016/j.ijforecast.2014.08.008>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207014001083>.
- Weron, R., & Misiorek, A. (2008). Forecasting spot electricity prices: A comparison of parametric and semiparametric time series models. *International Journal of Forecasting*, 24(4), 744–763. <http://dx.doi.org/10.1016/j.ijforecast.2008.08.004>, URL: <https://www.sciencedirect.com/science/article/pii/S0169207008000952>. Energy Forecasting.
- Zareipour, H., Canizares, C. A., & Bhattacharya, K. (2010). Economic impact of electricity market price forecasting errors: A demand-side analysis. *IEEE Transactions on Power Systems*, 25(1), 254–262. <http://dx.doi.org/10.1109/TPWRS.2009.2030380>.
- Zhao, J. H., Dong, Z. Y., Xu, Z., & Wong, K. P. (2008). A statistical approach for interval forecasting of the electricity price. *IEEE Transactions on Power Systems*, 23(2), 267–276. <http://dx.doi.org/10.1109/TPWRS.2008.919309>.