

Deep distributional time series models and the probabilistic forecasting of intraday electricity prices

Nadja Klein¹ | Michael Stanley Smith² | David J. Nott³

¹Research Center Trustworthy Data Science and Security, Department of Statistics at TU Dortmund, University Alliance Ruhr, Dortmund, Germany

²Melbourne Business School, University of Melbourne, Melbourne, Victoria, Australia

³Department of Statistics and Data Science, National University of Singapore, Singapore

Correspondence

Michael Stanley Smith, Melbourne Business School, University of Melbourne, 200 Leicester Street, Carlton, Melbourne, Victoria 3053, Australia.
Email: mikes70au@gmail.com

Summary

Recurrent neural networks (RNNs) with rich feature vectors of past values can provide accurate point forecasts for series that exhibit complex serial dependence. We propose two approaches to constructing deep time series probabilistic models based on a variant of RNN called an echo state network (ESN). The first is where the output layer of the ESN has stochastic disturbances and a Bayesian prior for regularization. The second employs the implicit copula of an ESN with Gaussian disturbances, which is a Gaussian copula process on the feature space. Combining this copula process with a nonparametrically estimated marginal distribution produces a distributional time series model. The resulting probabilistic forecasts are deep functions of the feature vector and marginally calibrated. In both approaches, Markov chain Monte Carlo methods are used to estimate the models and compute forecasts. The proposed models are suitable for the complex task of forecasting intraday electricity prices. Using data from the Australian market, we show that our deep time series models provide accurate short-term probabilistic price forecasts, with the copula model dominating. Moreover, the models provide a flexible framework for incorporating probabilistic forecasts of electricity demand, which increases upper tail forecast accuracy from the copula model significantly.

KEYWORDS

density forecasting, echo state network, electricity price forecasting, Gaussian copula process, marginal calibration, recurrent neural network

1 | INTRODUCTION

Deep models with rich feature vectors are often very effective in problems that require accurate forecasts (Goodfellow et al., 2016). These include financial applications, such as predicting equity risk premiums and returns (Feng et al., 2018; Gu et al., 2021) and bond returns (Bianchi et al., 2021). Another financial application where deep models have high potential is the forecasting of intraday electricity prices. Electricity prices exhibit a strong and complex nonlinear serial dependence, quite unlike security prices, and accounting for this is key to obtaining accurate forecasts (Manner et al., 2019; Nowotarski & Weron, 2018). Shallow neural networks (NNs) (Amjady, 2006), and more recently deep neural networks (DNNs) (Lago et al., 2018; Ugurlu et al., 2018), have been shown to capture these nonlinearities well and produce accurate point forecasts. However, it is the accurate forecasting of the entire distribution of prices—variously called prob-

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2023 The Authors. Journal of Applied Econometrics published by John Wiley & Sons Ltd.

abilistic, density or distributional forecasting—that is essential for both market operators and participants. Yet, to date, probabilistic forecasts of electricity prices using DNNs are rare. In this paper, we propose a number of time series probabilistic forecasting models that exploit and extend state-of-the-art deep models and apply them to data from the Australian market.

Day-ahead wholesale electricity markets operate throughout the world, including the United States and Europe. In these markets, generators and distributors bid for sale and purchase of electricity at an intraday resolution in an auction one day prior to transmission. The auction clearing price is called the electricity spot price; see Kirschen and Strbac (2018) for an introduction to such markets. Accurate price forecasts at an intraday resolution, one or more days ahead, are central to both the efficient operation of the market and profitability of participants. Particularly important are probabilistic forecasts of price, not just the mean, variance or other moments. This is because overall profitability of market participants is strongly affected by prices in the tails, which are very heavy in most wholesale markets.

Recurrent neural networks (RNNs) are DNNs tailored to capture temporal behaviour and are suitable for forecasting nonlinear time series (Goodfellow et al., 2016, ch. 10). However, RNNs typically have a large number of hidden weights and are difficult to train (i.e., estimate). Therefore, we use a variant of RNNs called echo state networks (ESNs) (Jaeger, 2007; Lukoševičius & Jaeger, 2009) that employ a form of regularization that makes them scalable to long series and computationally stable. We build statistical time series models based on ESNs using two approaches. The first extends that of Chatzis and Demiris (2011) and McDermott and Wikle (2017, 2019), who use ESNs within statistical models, with the output layer coefficients of the hidden state vector estimated using Bayesian methods. In our work, we include a shrinkage prior for the output layer coefficients to provide additional regularization, along with three different additive error distributions for the output layer: Gaussian, skew-normal and skew- t .

Our second approach is the main methodological contribution of the paper. It uses the implicit copula of the time series vector from a Gaussian probabilistic ESN of the type described above. By an ‘implicit copula’ we mean the copula that is implicit in a multivariate distribution and obtained by inverting Sklar’s theorem; see Smith (2022) for an introduction to implicit copulas. This copula is both a deep function of the feature vector and also a ‘Gaussian copula process’. We combine our proposed copula with a nonparametric marginal distribution for electricity prices, producing a time series model that captures the complex serial dependence in the series. The accurate estimation of the marginal distribution ensures ‘marginal calibration’, which is where the long run average of the predictive distributions of the time series variable matches its observed margin (Gneiting & Katzfuss, 2014). Importantly, the entire predictive distribution from the copula model is a deep function of the feature vector. This copula model extends the deep distributional regression methodology of Klein et al. (2021) to deep distributional time series and ESNs.

In both our deep time series models, the feature vector includes a rich array of past series values and possibly other variables. To regularize these, ESNs use sparse and randomly assigned fixed weights for the hidden layers of the DNN. For each of K random configurations of weights, Markov chain Monte Carlo (MCMC) is used to estimate the statistical model and compute Bayesian predictive distributions. The probabilistic forecasts are then ensembles of these predictive distributions over the K configurations of weights. We show in our empirical work that this ensemble provides for accurate uncertainty quantification.

We use our deep time series models to forecast intraday electricity prices in the Australian National Electricity Market (NEM). The NEM was established in 1998, comprises around 1% of Australian economic output, provides public data, and has a design that is typical of other day-ahead markets. It has five regional price series for which forecasting has been much studied; see Ignatieva and Trück (2016), Manner et al. (2019) and Han et al. (2020) for overviews. Using a feature vector with lagged prices from all five regions, we compute predictions at the half-hourly resolution for a 24 h horizon over an 8 month validation period during 2019. Serinaldi (2011), Gianfreda and Bunn (2018) and Narajewski and Ziel (2020) show the ‘generalized additive models for location, scale and shape’ (GAMLSS) methodology of Rigby and Stasinopoulos (2005) applied to time series provides accurate distributional forecasts of electricity prices, and we employ this as a benchmark. Using contemporary metrics, the probabilistic ESN with skew- t errors produces more accurate point and probabilistic forecasts than ESNs with either Gaussian or skew-normal errors, and also GAMLSS. However, the deep copula model produces more accurate probabilistic forecasts, particularly in the tails, and has superior forecast coverage. Thus, marginal calibration also improves calibration of the (conditional) predictive distributions.

Participants in the NEM are provided with probabilistic forecasts of electricity demand by the system operator. Recent studies (Shah & Lisi, 2020; Ziel & Steinert, 2016) suggest that using accurate demand forecasts may further improve time series forecasts of price. An advantage of deep models is that additional predictors are easily included in a flexible fashion as extra elements in the feature vector. We do so here using three quantiles of the 24 h ahead demand forecasts, and find

this increases forecast accuracy of the upper tail of the price distribution when using the deep copula model. Accurate forecasting of the upper tail of prices impacts the profitability of participants substantially (Christensen et al., 2012).

The paper is organized as follows. Section 2 provides an overview of electricity markets and price forecasting, with a focus on the NEM. Section 3 outlines ESNs, their probabilistic extension using additive disturbances, and Bayesian methods for their estimation and prediction. Section 4 outlines the implicit copula process and the proposed deep distributional time series model. Section 5 compares the deep time series and benchmark model forecasts, Section 6 considers the inclusion of demand forecast data, and Section 7 concludes. The Web Appendix provides key algorithms, computational details and additional empirical results.

2 | ELECTRICITY MARKETS, PRICE FORECASTING AND DATA

2.1 | Wholesale markets

Wholesale electricity markets include the European Power Exchange, US regional markets (such as the PJM interconnection and the Southwest Power Pool), and national markets in many countries including Australia, Chile and Turkey. Although the designs of these markets differ, they are largely 'day-ahead' markets where generators and distributors place bids for the sale and purchase of electricity at an intraday resolution up to one day prior to transmission (or 'dispatch'). The market is cleared at a wholesale spot price that reflects the marginal cost of supply at each intraday period. Prices also vary at different geographic reference nodes, creating multiple related price series. Markets are overseen by system operators, which match generation with short-term demand forecasts, impose constraints to ensure system stability (i.e., avoid load-shedding or blackouts), and enforce any price caps; see Kirschen and Strbac (2018) for an overview of wholesale markets.

Central to the operation of day-ahead markets is the intraday electricity spot price. For market participants, accurate short-term forecasts of the price are key to profitability. Because electricity is a flow commodity with a high cost of storage, arbitrage opportunities are limited. This fact, along with the complexities of transmission and that short-run demand is inelastic with respect to price, means that prices exhibit unique stylized characteristics; see Knittel and Roberts (2005), Karakatsani and Bunn (2008), Panagiotelis and Smith (2008) and Weron (2014) for summaries of these. From a time series perspective, this includes strong and complex nonlinear serial dependence, whereas from a distributional perspective, prices have very heavy tails, skew and often multiple modes that correspond to different regimes (Janczura & Weron, 2010) and economic equilibria (Smith & Shively, 2018).

2.2 | Electricity price forecasting

Many methods have been used for short-term forecasting of electricity spot prices; see Weron (2014) and Nowotarski and Weron (2018) for overviews of point and probabilistic forecasting methods, respectively. Although shallow neural networks (NNs) have long been popular for forecasting electricity prices, DNNs have the potential to be more accurate. For example, Lago et al. (2018) and Ugurlu et al. (2018) both found that RNNs provide more accurate point forecasts than a range of benchmark models. However, most previous usages of shallow and deep NNs have focused on point forecasts of prices, and only in a few cases also quantify predictive uncertainty using bootstrap or other Monte Carlo methods (Rafiei et al., 2016).

In contrast, a number of other methods have been used to construct probabilistic forecasts (Bunn et al., 2016; Huurman et al., 2012; Misiolek et al., 2006; Panagiotelis & Smith, 2008). One particularly promising avenue is to extend distributional regression methods to time series forecasting. For example, Gianfreda and Bunn (2018) and Narajewski and Ziel (2020) do so for German electricity prices using exogenous covariates, and Serinaldi (2011) does so for Californian and Italian electricity prices using historical prices and other variables. These papers report more accurate probabilistic forecasts. In Section 4, we develop a copula-based model that exploits the accuracy exhibited by deep models within a distributional time series forecasting setting, thereby combining the advantages of both approaches.

Copulas have become increasingly popular in time series models of electricity prices because they can capture complex dependence, while also allowing for flexible margins. Smith et al. (2012), Ignatieva and Trück (2016), Manner et al. (2016), Pircalabu and Benth (2017) and Manner et al. (2019) use low-dimensional copulas to capture cross-sectional dependence between regional prices, price spikes and other energy series in power systems, whereas Smith (2015) and Smith and Shively (2018) use high-dimensional copulas to capture both serial and cross-sectional dependence jointly for multiple regional prices. However, these studies all employ copulas that are different to the copula process proposed here.




FIGURE 1 Boxplots of the logarithm of electricity prices in NSW for each hour of the day. *Note:* Five extreme outliers greater than 7.3 that correspond to price spikes have been removed to improve visualization. The price data are from the period 7-Jan-19 to 31-Dec-19.

2.3 | Australian electricity prices

Electricity generation in the NEM is an important component of economic activity, with 19.4bn Australian dollars of turnover during the 2018–2019 financial year (Australian Energy Regulator, 2019). Operations in the NEM are managed by the Australian Electricity Market Operator (AEMO), and since April 2006, it has had five regions that correspond to the power systems in the states of New South Wales (NSW), Queensland (QLD), Victoria (VIC), South Australia (SA) and Tasmania (TAS). Separate prices are set at a central location (or ‘node’) in each region, although they are dependent because the state-based power systems are interconnected by high voltage direct current lines.

Participating utilities place bids for the purchase (by distributors) and sale (by generators) of electricity at 5-min intervals up to one day prior to dispatch. The trading price is the average clearing price for this auction over six consecutive 5-min periods, so that it is observed at a half-hourly resolution. Re-bidding of prices is allowed before dispatch, although not the amount of energy. Price forecasting in the NEM has been studied extensively, with contributions by Higgs (2009), Panagiotelis and Smith (2008), Nowotarski et al. (2013), Ignatieva and Trück (2016), Rafiei et al. (2016), Smith and Shively (2018), Apergis et al. (2019) and Manner et al. (2019) among others. Intraday forecasts over a horizon of 24 h are used by market participants to develop effective strategies for bidding, re-bidding and managing risk. In our study, we employ the half-hourly trading prices (measured in Australian dollars per MW/h) observed during 2019 in the five regions of the NEM. Prices can be negative for short periods (when it is more cost-effective to sell into the market at a loss, rather than ramp down generation temporarily) although the floor price is $-\$1000$. There is also a maximum price that is adjusted annually on 1 July, which was $\$14,500$ prior to 1 July 2019 and $\$14,700$ afterwards.

Figure A in Part E of the Web Appendix plots prices in NSW during 4 weeks corresponding to the four seasons, illustrating the strong heterogeneity based on the time of day, day of the week and season. Table A in the Web Appendix provides a summary of the prices in each region. These have extreme positive skew, so that we follow most previous studies and work with the logarithm of price $Y = \log(\text{Price} + 1001)$. The addition of 1001 accounts for the minimum price, at which $Y = 0$. It is straightforward to construct density forecasts of the nominal price from those of Y using the Jacobian of this transformation. Figure B in the Web Appendix plots histograms of the five price series, along with (bounded) kernel density estimates (KDEs), showing that even the distribution of Y is both asymmetric and heavy-tailed for each series. Figure 1 gives boxplots of Y for NSW (which is the region with the most energy demand), broken down by hour of the day. It shows that the entire distribution varies substantially over the day. Equivalent plots for the day of the week, and month of the year (see Figures C and D of the Web Appendix) also reveal similar heterogeneity in the distribution of prices.

3 | DEEP TIME SERIES

RNNs represent complex temporal relationships between variables by allowing cycles and sequences in their hidden layers. However, RNNs are typically computationally expensive to estimate and can be numerically unstable for long series (Pascanu et al., 2013), both of which are issues here. One way to address this is to employ a variant called an Echo State Network (ESN) that decreases the number of weights that need to be trained (Jaeger, 2007). The key idea of ESNs is to only adapt the output layer in training, while keeping the weights of recurrent and input connections fixed yet randomly assigned. This approach is called a ‘reservoir computing’ method (Lukoševičius & Jaeger, 2009) and provides flexibility at a lower computational cost. However, classical ESNs rarely consider uncertainty quantification, and in this section we

outline probabilistic ESNs that do so. The approach is extended in Section 4 to allow for marginal calibration of the data distribution using a deep copula model.

3.1 | Gaussian probabilistic ESN

3.1.1 | Specification

We first specify the ensemble ESN proposed for spatio-temporal data by McDermott and Wikle (2017). For series with highly nonlinear dependence, these authors demonstrate that allowing the response equation to depend on the hidden states quadratically in an ESN can increase predictive accuracy, which we also find for the electricity price series examined here. Let $\{Y_t\}$ be a stochastic process, then the ESN with Gaussian disturbances $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$ takes the following form for $t \geq 1$:

$$\begin{aligned} \text{Response Equation:} \quad & Y_t = \beta_0 + \mathbf{h}_t' \boldsymbol{\beta}_1 + (\mathbf{h}_t')^{\circ 2} \boldsymbol{\beta}_2 + \varepsilon_t, \\ \text{Hidden State Equation:} \quad & \mathbf{h}_t = (1 - \kappa) \mathbf{h}_{t-1} + \kappa \tilde{\mathbf{h}}_t \\ & \tilde{\mathbf{h}}_t = g_h \left(\frac{\delta}{\lambda_V} V \mathbf{h}_{t-1} + U \mathbf{x}_t \right). \end{aligned} \quad (1)$$

Here, \mathbf{x}_t is an $(n_x \times 1)$ feature vector that includes unity for an intercept term, \mathbf{h}_t is a $(n_h \times 1)$ hidden state vector, ' $\circ 2$ ' denotes the element-wise square of a matrix, V, U are $(n_h \times n_h)$ and $(n_h \times n_x)$ matrices of hidden layer weights, and g_h is an activation function set to $g_h(x) = \tanh(x)$. The parameter $0 < \kappa \leq 1$ is called the 'leaking rate'. Following McDermott and Wikle (2017), we set $\kappa = 1$ after having checked that predictive performance is not improved when $\kappa < 1$, so that $\tilde{\mathbf{h}}_t = \mathbf{h}_t$. The feature vector \mathbf{x}_t contains past values of both the response and other series, which we specify in Section 5. The response Y_t is a nonlinear function not only of \mathbf{x}_t but also of all previous values $\mathbf{x}_1, \dots, \mathbf{x}_{t-1}$, so the stochastic process is not Markov.

The constant λ_V is the largest eigenvalue of V , and $0 < \delta < 1$ a scaling parameter, so that $(\delta/\lambda_V)V$ has spectral radius less than one. A spectral radius greater than one can result in unstable behaviour in the latent states (Lukoševičius & Jaeger, 2009). We follow McDermott and Wikle (2017) and set $\delta = 0.35$ and $n_h = 120$, although we found the forecasting results to be insensitive to variations in these settings.

The elements of the matrices $V = \{v_{il}\}, U = \{u_{ij}\}$ are assumed to be random, and distributed independently from mixtures of a uniform distribution and a point mass at zero. If $\mathcal{U}(a, b)$ denotes a uniform distribution over domain (a, b) , $\mathcal{B}(\pi)$ denotes a Bernoulli distribution with mean π , and δ_0 is the Dirac function at zero, then the elements of V and U are

$$\begin{aligned} v_{il} &= \gamma_{il}^v \mathcal{U}(-a_v, a_v) + (1 - \gamma_{il}^v) \delta_0, \quad \gamma_{il}^v \sim \mathcal{B}(\pi_v), \\ u_{ij} &= \gamma_{ij}^u \mathcal{U}(-a_u, a_u) + (1 - \gamma_{ij}^u) \delta_0, \quad \gamma_{ij}^u \sim \mathcal{B}(\pi_u). \end{aligned} \quad (2)$$

We follow McDermott and Wikle (2017) and set $a_v = a_u = \pi_v = \pi_u = 0.1$ after having checked the predictive performance under several alternative settings. This provides a sparse structure that differs from popular dropout regularization (Gal & Ghahramani, 2016) and provides computational and other advantages; see Gallicchio (2020).

Consider T time series observations $\mathbf{Y} = (Y_1, \dots, Y_T)'$ of the stochastic process with corresponding $(T \times n_x)$ matrix of feature values $X = [\mathbf{x}_1 | \dots | \mathbf{x}_T]'$. Denote $\boldsymbol{\xi} = \{V, U, \kappa, \delta\}$, $H_{\boldsymbol{\xi}}(X) = [\mathbf{h}_1 | \dots | \mathbf{h}_T]'$ as the $(T \times n_h)$ matrix of hidden state values, $B_{\boldsymbol{\xi}}(X) = [\mathbf{1}, H_{\boldsymbol{\xi}}(X) | H_{\boldsymbol{\xi}}(X)^{\circ 2}]$, $\mathbf{1}$ as a vector of ones, and $\boldsymbol{\beta} = (\beta_0, \boldsymbol{\beta}_1', \boldsymbol{\beta}_2')'$. Then (1) can be written as the linear model

$$\mathbf{Y} = B_{\boldsymbol{\xi}}(X) \boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_T)' \sim \mathcal{N}(\mathbf{0}, \sigma^2 I). \quad (3)$$

Given $\boldsymbol{\xi}, X$ and $\mathbf{h}_0 = \mathbf{0}$, the hidden state matrix $H_{\boldsymbol{\xi}}(X)$ is known without error as the hidden state vectors can be computed recursively. Only $\boldsymbol{\beta}$ and σ^2 (which we refer to as model parameters) require estimation, for which we use their Bayesian posterior distribution. Differing from McDermott and Wikle (2017), we regularize $\boldsymbol{\beta}$ by adopting the shrinkage prior

$$\boldsymbol{\beta} | \tau^2 \sim \mathcal{N}(\mathbf{0}, P(\tau^2)^{-1}), \quad \sigma^2 \sim \text{IG}(a, b), \quad (4)$$

where IG denotes an Inverse Gamma distribution. We found a ridge prior with $P(\tau^2) = \tau^2 I$, and hyper-prior $\tau^2 \sim \text{IG}(\tilde{a}, \tilde{b})$ with $a = b = \tilde{a} = \tilde{b} = 0.001$ to work well. The posterior of the model parameters $\boldsymbol{\vartheta} = (\boldsymbol{\beta}, \sigma^2, \tau^2)$ of this regularized

linear model can be computed using the standard MCMC sampler at Algorithm 1 in Part A of the Web Appendix. Other shrinkage priors of the form at (4) can also be used here as discussed in Klein and Smith (2019).

3.1.2 | Probabilistic forecasts

Most implementations of ESNs only draw a single set of weights from (2). Here, we follow McDermott and Wikle (2017, 2019) and simulate $K = 100$ matrices $\{V^k, U^k; k = 1, \dots, K\}$ from (2). An ensemble is then used to integrate over U, V when constructing the probabilistic forecasts. If $\xi^k = \{V^k, U^k, \kappa, \delta\}$, then the density forecast of Y_{T+h} at time T for each point $h = 1, \dots, h_1$ in the forecast horizon is the ensemble

$$f_{T+h|T}(y_{T+h}) \equiv \frac{1}{K} \sum_{k=1}^K p^k(y_{T+h}|X, \mathbf{y}), \quad (5)$$

where the subscript notation indicates $f_{T+h|T}$ is conditional on the filtration at time T .

The density p^k in (5) is the Bayesian posterior predictive density computed for configuration ξ^k as follows. Let $X_{(t)} \equiv [\mathbf{x}_1 | \dots | \mathbf{x}_t]'$, then this posterior predictive density is

$$p^k(y_{T+h}|X, \mathbf{y}) = \iint p(y_{T+h}|X_{(T+h)}, \boldsymbol{\vartheta}) p(\mathbf{x}_{T+2}, \dots, \mathbf{x}_{T+h} | \boldsymbol{\vartheta}, X, \mathbf{y}) p(\boldsymbol{\vartheta} | X, \mathbf{y}) d\mathbf{x}_{T+2} \dots d\mathbf{x}_{T+h} d\boldsymbol{\vartheta}, \quad (6)$$

which is an integral over any unobserved feature values and the model parameter posterior.¹ From (1), the first term in the integrand is the density of a $\mathcal{N}(\mathbf{h}'_{T+h}\boldsymbol{\beta}_1 + (\mathbf{h}'_{T+h})^{\circ 2}\boldsymbol{\beta}_2, \sigma^2)$ distribution, where \mathbf{h}_{T+h} is computed through the recursion of the hidden state equation (therefore, \mathbf{h}_{T+h} is a deep function of $\mathbf{x}_1, \dots, \mathbf{x}_{T+h}$). The outer integral in $\boldsymbol{\vartheta}$ can be evaluated by averaging over draws from the posterior $p(\boldsymbol{\vartheta} | X, \mathbf{y})$ obtained from running the MCMC sampler. However, plugging in the posterior mean $E(\boldsymbol{\vartheta} | X, \mathbf{y})$ computed from the Monte Carlo sample is much faster, yet can be almost as accurate, so we follow this approach.²

In our empirical work, \mathbf{x}_t contains past values of both the price in the focal region and the other four regional prices, so that $\mathbf{x}_1, \dots, \mathbf{x}_{T+1}$ are observed at or before time T . However, some elements of $\mathbf{x}_{T+2}, \dots, \mathbf{x}_{T+h}$ are unobserved, and the integrals over these feature vectors in (6) are with respect to their unobserved elements only. The integrals are computed in a Monte Carlo manner by simulating all five series values sequentially from their predictive distributions. Algorithm 2 in Part B of the Web Appendix simulates from the ensemble density (5) for all five regional price series jointly and at each time point in the forecast horizon.

3.2 | Skew t probabilistic ESN

The assumption of Gaussian disturbances at (1) is inconsistent with the asymmetry in the empirical distribution of (logarithmic) prices. Thus, we also consider the skew- t distribution of Azzalini and Capitanio (2003) with location zero for the disturbances, with density $p(\varepsilon_t; \omega^2, \alpha, \nu) = \frac{2}{\omega} t_\nu \left(\frac{\varepsilon_t}{\omega} \right) T_{\nu+1} \left(\frac{\alpha \varepsilon_t}{\omega} \sqrt{\frac{\nu+1}{\nu + \varepsilon_t^2}} \right)$, where t_ν and T_ν are the student t density and distribution functions. The parameters ω^2, α, ν are scale, skew and degrees of freedom parameters, respectively, and we write $\varepsilon_t \sim \mathcal{ST}(0, \omega^2, \alpha, \nu)$. When $\nu > 30$ the density is effectively that of a skew-normal distribution, whereas when $\alpha = 0$ it is that of a t distribution. We fix values of ν , and use the prior at (4) for $\boldsymbol{\beta}$, so that the model parameters are $\{\boldsymbol{\beta}, \tau^2, \omega^2, \alpha\}$.

Bayesian estimation of the model parameters uses a conditionally Gaussian representation of the skew- t . This introduces latent variables ζ_t, w_t and expresses ε_t conditional on these values as $\varepsilon_t = \psi \zeta_t + \varepsilon_t$, where

$$\varepsilon_t | w_t \sim \mathcal{N}(0, \sigma^2/w_t), w_t \sim \mathcal{G}\left(\frac{\nu}{2}, \frac{\nu}{2}\right), \zeta_t | w_t \sim \mathcal{TN}_{[0, \infty)}(0, 1/w_t),$$

¹It is implicit that all densities in the integrand of (6) are conditional on weight configuration ξ^k .

²Implementation requires running the MCMC sampler and computing the posterior mean a total of K times, once for each hidden weight configuration ξ^1, \dots, ξ^K .

\mathcal{G} denotes a Gamma distribution, $\mathcal{TN}_{[0,\infty)}$ is a Gaussian truncated to $[0, \infty)$, and $\{\alpha, \omega\}$ are re-parameterized as $\{\psi, \sigma\}$ with $\psi = \alpha\sigma$ and $\sigma^2 = \omega^2 - \psi^2$. Integrating out ζ_t, w_t recovers the skew- t distribution $\varepsilon_t \sim \mathcal{ST}(0, \omega^2, \alpha, \nu)$. Following Frühwirth-Schnatter and Pyne (2010), the hyperpriors are $\psi \sim \mathcal{N}(0, D_0)$, $\sigma^2 \sim \mathcal{IG}(c_0, C_0)$ and $\tau^2 \sim \mathcal{IG}(b_0, B_0)$, with constants $D_0 = 1, c_0 = 2.5, b_0 = 1, B_0 = 0.005, C_0 = 0.5s_y^2$ and s_y^2 denoting the sample variance of Y_t .

Given ξ , the re-parameterized model parameters $\boldsymbol{\theta} = \{\boldsymbol{\beta}, \tau^2, \psi, \sigma^2\}$ are estimated using their Bayesian posterior computed by an MCMC sampler that generates $\boldsymbol{\zeta} = (\zeta_1, \dots, \zeta_T)'$ and $\boldsymbol{w} = (w_1, \dots, w_T)'$, similar to those proposed by Panagiotelis and Smith (2008) and Frühwirth-Schnatter and Pyne (2010). As with the Gaussian probabilistic ESN, prediction is based on Monte Carlo draws from the ensemble density at (5) produced using a minor adjustment of Algorithm 2. Details are given in Part B of the Web Appendix.

4 | DEEP DISTRIBUTIONAL TIME SERIES

Although the deep time series models in Section 3 provide probabilistic forecasts they have two drawbacks: (i) the feature vector only affects the mean of the response equation at (1), and (ii) the density forecasts are not calibrated in any manner. In this section, a copula model is outlined that is a deep distributional time series model in which the feature vector affects the entire predictive distribution and the probabilistic forecasts are marginally calibrated.

4.1 | Marginal calibration

Gneiting et al. (2007) discuss different forms of calibration of density forecasts, including predictive marginal calibration, which is defined as follows. For $t = T + 1, \dots, \infty$, assume a future observation of Y_t has true distribution $H_{t|T}(y)$ and forecast distribution $F_{t|T}(y)$, where the subscript indicates that both distributions are conditional on the filtration at time T . Then, if $\bar{H}(y) \equiv \lim_{h_1 \rightarrow \infty} \left\{ \frac{1}{h_1} \sum_{h=1}^{h_1} H_{T+h|T}(y) \right\}$ and $\bar{F}(y) \equiv \lim_{h_1 \rightarrow \infty} \left\{ \frac{1}{h_1} \sum_{h=1}^{h_1} F_{T+h|T}(y) \right\}$, the forecast distributions are called marginally calibrated if and only if $\bar{F}(y) = \bar{H}(y)$.

Gneiting et al. (2007) highlight that because $H_{t|T}$ is unknown for $t > T$, marginal calibration can be assessed in practice by comparing the empirical distribution function over the forecast horizon to the average of the corresponding distributional forecasts. In a study with a moving window (such as ours) and forecast origins $T = T_{\text{start}}, \dots, T_{\text{finish}}$, there are a total of $N = T_{\text{finish}} - T_{\text{start}} + 1$ sets of forecasts. Averaging over these and the forecast horizon of length h_1 , gives

$$\hat{H}(y) = \frac{1}{Nh_1} \sum_{T=T_{\text{start}}}^{T_{\text{finish}}} \sum_{t=T+1}^{T+h_1} \mathbf{I}(y_t \leq y), \text{ and } \hat{F}(y) = \frac{1}{Nh_1} \sum_{T=T_{\text{start}}}^{T_{\text{finish}}} \sum_{t=T+1}^{T+h_1} F_{t|T}(y), \tag{7}$$

where $\mathbf{I}(X)$ is an indicator function equal to one if X is true, and zero otherwise. The closer \hat{F} is to \hat{H} , the greater evidence for predictive marginal calibration.

4.2 | Copula model

We adopt a copula model for the joint distribution of $\mathbf{Y}_{(t)} = (Y_1, \dots, Y_t)'$, conditional on $X_{(t)} = [\mathbf{x}_1 | \dots | \mathbf{x}_t]'$ and weight configuration ξ , with the density decomposition

$$p(\mathbf{y}_{(t)} | X_{(t)}, \xi) = c_{\text{ESN}}(F_Y(y_1), \dots, F_Y(y_t) | X_{(t)}, \xi) \prod_{s=1}^t p_Y(y_s), \text{ for } t \geq 2, \tag{8}$$

where $\mathbf{y}_{(t)} = (y_1, \dots, y_t)'$ and $\mathbf{u}_{(t)} = (u_1, \dots, u_t)'$. The deep copula process has t -dimensional density $c_{\text{ESN}}(\mathbf{u}_{(t)} | X_{(t)}, \xi)$ specified below in Section 4.3. The marginal density p_Y and corresponding distribution function F_Y in (8) are time invariant, and are estimated nonparametrically from the training data.

We stress that even though in (8) the distribution $Y_t | \mathbf{x}_t$ is assumed marginally invariant with respect to \mathbf{x}_t , the distribution $\mathbf{Y}_{(t)} | X_{(t)}$ is still related to the matrix of feature vector values $X_{(t)}$ though the joint distribution. A consequence is that for $t > T$ the predictive density of Y_t can be heavily dependent on the feature vector \mathbf{x}_t , as outlined in Section 4.5.

4.3 | Specification of the deep Gaussian copula process

At (8), we employ an implicit copula process with density c_{ESN} constructed from the joint distribution of a second stochastic process $\{\tilde{Z}_s\}$ that follows the Gaussian probabilistic ESN at (1) with β integrated out under the prior (4). We call \tilde{Z}_s a ‘pseudo-response’ because it is not observed directly but is introduced only for specification of its implicit copula. The t observations $\tilde{\mathbf{Z}}_{(t)} = (\tilde{Z}_1, \dots, \tilde{Z}_t)'$ are conditionally distributed

$$\tilde{\mathbf{Z}}_{(t)} | X_{(t)}, \sigma^2, \tau^2, \xi \sim \mathcal{N} \left(\tilde{\mathbf{z}}_{(t)}; \mathbf{0}, \sigma^2 \left(I + \frac{1}{\tau^2} B_\xi(X_{(t)}) B_\xi(X_{(t)})' \right) \right), \quad (9)$$

with $\tilde{\mathbf{z}}_{(t)} = (\tilde{z}_1, \dots, \tilde{z}_t)'$ and the matrix $B_\xi(X_{(t)})$ is specified at (3) but without an intercept term (i.e., the first column of ones) because level is unidentified in a copula.

The implicit copula of a Gaussian distribution is called a Gaussian copula, and is constructed for (9) by standardizing the distribution; see, for example, Smith (2022). Let $\mathbf{Z}_{(t)} = (Z_1, \dots, Z_t)' = \sigma^{-1} S \tilde{\mathbf{Z}}_{(t)}$, where $S = \text{diag}(\psi_1, \dots, \psi_t)$ is a diagonal scaling matrix with elements $\psi_s = (1 + \mathbf{b}'_s \mathbf{b}_s / \tau^2)^{-1/2}$, and \mathbf{b}_s is the s -th row of $B_\xi(X_{(t)})$. Then $\mathbf{Z}_{(t)} | X_{(t)}, \sigma^2, \tau^2, \xi \sim \mathcal{N}(\mathbf{0}, R)$, where

$$R = S \left(I + \frac{1}{\tau^2} B_\xi(X_{(t)}) B_\xi(X_{(t)})' \right) S,$$

has ones on the leading diagonal, and is a function of $(X_{(t)}, \xi, \tau^2)$. Then the copula has density $c_{\text{ESN}}(\mathbf{u}_{(t)} | X_{(t)}, \xi, \tau^2) = \frac{\phi(\mathbf{z}_{(t)}; \mathbf{0}, R)}{\prod_{s=1}^t \phi_1(z_s)}$, where $\mathbf{z}_{(t)} = \Phi_1^{-1}(\mathbf{u}_{(t)})$, $\mathbf{z}_{(t)} = (z_1, \dots, z_t)'$, ϕ_1 is the standard normal density, and $\phi(\cdot; \mathbf{0}, R)$ is the density of a $\mathcal{N}(\mathbf{0}, R)$ distribution. Notice that σ^2 does not feature in c_{ESN} , which is because copulas are scale free.

Because c_{ESN} is conditional on $X_{(t)}$, this is a Gaussian copula process on the feature space (Wilson & Ghahramani, 2010). Because the copula captures the dependence structure in $\{Z_t\}$, \mathbf{x}_t contains past values of this process (rather than $\{Y_t\}$), and the equivalent processes for the other four regional electricity prices. These can be computed easily as $Z_t = \Phi^{-1}(F_Y(Y_t))$. We write the copula density c_{ESN} also as a function of the unknown parameter τ^2 and employ it in the copula model at (8).

4.4 | Estimation

Given a configuration ξ , the only unknown copula parameter is τ^2 , for which we adopt the Weibull prior of Klein and Kneib (2016) with scale parameter $b_{\tau^2} = 2.5$. Direct estimation using the likelihood at (8) is difficult because evaluation of $c_{\text{ESN}}(\mathbf{u} | X, \xi, \tau^2)$ requires inversion of the matrix R , which is computationally infeasible for all but small sample sizes. This problem is solved by expressing the likelihood conditional on β as follows. For a sample of size T , we denote the observations as $\mathbf{y} \equiv \mathbf{y}_{(T)}$ and feature matrix as $X \equiv X_{(T)}$. Then the conditional likelihood is obtained by a change of variables from \mathbf{y} to $\mathbf{z} = (z_1, \dots, z_T)'$, with elements $z_t = \Phi_1^{-1}(F_Y(y_t))$, so that

$$p(\mathbf{y} | X, \beta, \xi, \tau^2) = p(\mathbf{z} | X, \beta, \xi, \tau^2) \prod_{t=1}^T \frac{p_Y(y_t)}{\phi_1(z_t)} = \phi(\mathbf{z}; \mathbf{0}, S B_\xi(X) \beta, S^2) \prod_{t=1}^T \frac{p_Y(y_t)}{\phi_1(z_t)},$$

which can be evaluated in $O(T)$ operations because S is diagonal. Part C of the Web Appendix provides an MCMC sampler to generate draws from the augmented posterior of (β, τ^2) , so that β is integrated out in a Monte Carlo manner and direct computation of R is avoided.

4.5 | Probabilistic forecasts

As with the probabilistic ESNs in Section 3, an ensemble is used to integrate over the distribution of U, V , so that the density forecast of Y_{T+h} is again given by (5). The ensemble components p^k (i.e., the Bayesian posterior predictive densities at (6)) are derived from the copula model where $\theta = \{\beta, \tau^2\}$. To do so, the first term in the integrand of (6) is obtained by

a change of variables from y_{T+h} to $z_{T+h} = \Phi_1^{-1}(F_Y(y_{T+h}))$, so that

$$\begin{aligned}
 p(y_{T+h}|X_{(T+h)}, \boldsymbol{\vartheta}) &= p(z_{T+h}|X_{(T+h)}, \boldsymbol{\beta}, \tau^2) \frac{p_Y(y_{T+h})}{\phi_1(z_{T+h})} \\
 &= \frac{1}{\psi_{T+h}} \phi_1 \left(\frac{\Phi_1^{-1}(F_Y(y_{T+h})) - \mu_{T+h}}{\psi_{T+h}} \right) \frac{p_Y(y_{T+h})}{\phi_1(\Phi_1^{-1}(F_Y(y_{T+h})))},
 \end{aligned}
 \tag{10}$$

where $\mu_{T+h} = \mathbf{b}_{T+h}\boldsymbol{\beta}$, $\psi_{T+h} = \left(1 + \frac{1}{\tau^2} \mathbf{b}_{T+h}\mathbf{b}'_{T+h}\right)^{-1/2}$ and $\mathbf{b}_{T+h} = (\mathbf{h}'_{T+h}, \mathbf{h}^2_{T+h})'$ is a row vector. Notice that the entire density at (10) is a nonlinear function of the feature vector \mathbf{x}_{T+h} via \mathbf{b}_{T+h} , so that Y_{T+h} is not marginally invariant of \mathbf{x}_{T+h} in the predictive distribution. As in Section 3.1.2, the posterior mean $E(\boldsymbol{\vartheta}|\mathbf{y})$ is computed from the Monte Carlo sample and plugged in for $\boldsymbol{\beta}, \tau^2$ in (10). Because \mathbf{x}_t comprises past values of $\{Z_t\}$ for both the focal price, and the other four regional prices, the integrals over $\mathbf{x}_{T+2}, \dots, \mathbf{x}_{T+h}$ in (6) are computed by simulating values for Z_{T+t} sequentially for all five regions. Algorithm 4 in Part C of the Web Appendix generates iterates from the ensemble density for the copula model.

5 | FORECASTING COMPARISON

We illustrate our deep model forecasts using half-hourly price data observed during 2019 in all five regions. The objective is to identify how the different distributional assumptions affect probabilistic forecasting accuracy, and compare with the benchmark GAMLSS method, which Serinaldi (2011) and Gianfreda and Bunn (2018) establish as a state-of-the-art distributional forecasting method for electricity prices. Competing models were fit using a moving window of 3 months training data, starting with the period 1 February to 30 April. (We start at February, rather than January, because the input vector \mathbf{x}_t defined below includes lagged values of the series.) This window was advanced monthly until the period 1 September to 30 November, resulting in eight fits per model. Forecasts were constructed for an 8-month evaluation period from 1 May to 31 December for a total of 245 days. The origin T was advanced every half hour, and out-of-sample probabilistic forecasts constructed for a horizon of $h_1 = 48$ half-hours.³ When constructing forecasts during each month, the model parameter estimates obtained from the three preceding months training data were used. The reason each model was only refit eight times (rather than for every one of 11,760 half-hourly forecast origins) was to reduce the computational burden of the study. A forecast horizon up to 24 h is necessary for the bidding and re-bidding process in the NEM.

5.1 | Probabilistic forecasting models

The following four deep time series models were applied to each of the five price series:

- RNN:** The Gaussian probabilistic ESN outlined in Section 3.1.
- RNNST:** The skew- t probabilistic ESN outlined in Section 3.2 with $\nu = 7$.
- RNNSN:** A skew-normal probabilistic ESN, approximated using RNNST with $\nu = 30$.
- RNNC:** The copula model outlined in Section 4 that uses the deep copula constructed from a Gaussian probabilistic ESN, along with bounded KDE margins.

For the RNN, RNNST and RNNSN, the feature vector contains the $n_x = 271$ elements

$$\mathbf{x}_t = (1, \underbrace{\mathbf{Y}_{\text{ALL},t-1}, \mathbf{Y}_{\text{ALL},t-2}, \dots, \mathbf{Y}_{\text{ALL},t-48}}_{\text{prices in previous 24 h}}, \underbrace{\mathbf{Y}_{\text{ALL},t-96}, \mathbf{Y}_{\text{ALL},t-144}, \dots, \mathbf{Y}_{\text{ALL},t-336}}_{\text{prices at same hour 2 to 7 days previously}})'$$

where $\mathbf{Y}_{\text{ALL},t} = (Y_{\text{NSW},t}, Y_{\text{QLD},t}, Y_{\text{SA},t}, Y_{\text{TAS},t}, Y_{\text{VIC},t})$ is the row vector of all five regional prices at hour t and the first entry is the intercept. For RNNC, the \mathbf{x}_t contains values of the transformed series $Z_t = \Phi_1^{-1}(F_Y(Y_t))$ for each of the five regions at the same lagged time points as above. The choice of inputs is motivated by previous studies (Higgs, 2009; Panagiotelis & Smith, 2008) that identify strong dependence in Australian electricity prices with those in the previous 24 h, and at

³Because there are 245 days in the evaluation period, this means there was a total of 245×48 forecast origins. At each forecast origin, a total of $48 \times 5 = 240$ probabilistic forecast densities were constructed for each model. Thus, there were a total of 2,822,400 probabilistic forecast densities for each model.

the same time during the previous 7 days. Price dependence is induced by serial dependence in both supply-side factors and in demand (Cottet & Smith, 2003). Complex cross-sectional serial dependence in price is also well-documented (Han et al., 2020; Higgs, 2009; Ignatieva & Trück, 2016; Panagiotelis & Smith, 2008) and is caused by inter-regional trade in electricity between the five regions.

Although the deep time series models above are trained separately for each price series, the inclusion of lagged prices from all regions as elements of \mathbf{x}_t forms five-dimensional multivariate deep times series models. From each model, multi-step ahead forecast densities for the five regional prices are constructed jointly via simulation using Algorithms 2 and 5 outlined in the Web Appendix. At each forecast origin, 2000 Monte Carlo draws from (5) are obtained for each half hour $h = 1, \dots, 48$ and region $r = 1, \dots, 5$.

The density p_Y was estimated for each series using the bounded KDE implemented in MATLAB function ‘ksdensity’, with bounds set equal to the admissible prices during the forecast period. These were a floor price of $-\$1000$ ($Y_{T+h} \geq 0$), and a maximum price of $\$14,500$ prior to 1 July ($Y_{t+h} \leq 9.582$) and $\$14,700$ on or after that date ($Y_{t+h} \leq 9.596$). This ensures that the predictive densities from the copula model at (6) are constrained to the same range, which is important for the accurate forecasting of prices near the bounds. In contrast, density forecasts from the other models are unconstrained, and can have substantial mass outside the bounds. Thus, we simply truncate their predictive densities to the admissible price region and normalize.

In the GAMLSS benchmark method, the response Y_t has a specified parametric distribution, where the parameters are modelled as flexible functions of covariates. We used the R package ‘gamlss’ but found the software unstable when the ESN output layer terms were used as covariates and also for many of the candidate distributions supported by the package. However, the four parameter Johnson’s SU (JSU) distribution with \mathbf{x}_t as linear covariates worked well, and we include forecasts from this model. Serinaldi (2011) also found the JSU distribution to work well for electricity prices, which are a challenging series to fit because of the frequent price spikes. Details on our GAMLSS implementation and experiments are given in Part D of the Web Appendix.

5.2 | Measuring forecast accuracy

For point forecasts, the mean absolute error (MAE) and root mean square error (RMSE) are used to measure accuracy. For a univariate probabilistic forecast distribution F , there are a growing number of measures of accuracy, many of which are listed by Nowotarski and Weron (2018), and we compute the following four. The first is the quantile score function $QS_\alpha(F^{-1}(\alpha), y) = 2(I(y < F^{-1}(\alpha)) - \alpha)(\alpha - y)$, where $I(A) = 1$ if A is true, and zero otherwise. This is also called the ‘pinball loss’ (Gneiting, 2011). Although the entire quantile score function is of interest, we also consider its value at $\alpha = 0.05$ and $\alpha = 0.95$ as measures of lower and upper tail accuracy, respectively. The second measure is the continuous ranked probability score (CRPS) (Gneiting et al., 2007) defined as $CRPS(F, y) = \int_0^1 QS_\alpha(F^{-1}(\alpha), y) d\alpha$ for observed value y , where the integral is computed numerically. It is a proper scoring rule that measures overall probabilistic forecast accuracy.

The third measure is a loss function proposed by Fissler and Ziegel (2016) that is based on the ‘value-at-risk’ and ‘expected shortfall’, which are popular measures of financial tail risk for low values (i.e., losses). To measure tail risk associated with high values (the primary source of price risk in electricity markets), we employ the restatement of the loss function for the upper tail given by Nolde and Ziegel (2017). For $Y \sim F$ and α close to 1, let $EL(\alpha) \equiv E(Y|Y \geq F^{-1}(\alpha))$ be the ‘expected longrise’ (analogous to the expected shortfall measure of lower tail risk). Then from (Nolde & Ziegel, 2017, Prop. 3), the loss function is

$$S_\alpha(F, EL, y) = I(y \geq F^{-1}(\alpha)) \left(-G_1(F^{-1}(\alpha)) + G_1(y) - G_2(EL(\alpha))(F^{-1}(\alpha) - y) \right) + (1 - \alpha) \left(G_1(F^{-1}(\alpha)) - G_2(EL(\alpha))(EL(\alpha) - F^{-1}(\alpha)) + G_3(EL(\alpha)) \right), \quad (11)$$

where $G_1(x) = x$ and $G_2(x) = G_3(x) = \exp(x)$. We compute this score for $\alpha = 0.975$, and lower values indicate more accurate upper tail forecasts of the price distribution. The fourth and last measure is the empirical coverage of the predictive distributions at the 95% level.

For each of the first three metrics, we report the weighted average of their values over the evaluation period and over regions, with weights equal to the proportion of total electricity consumption in each region.⁴ We call these

⁴These weights are 0.3687 (NSW), 0.2355 (VIC), 0.2818 (QLD), 0.0624 (SA) and 0.0516 (TAS).

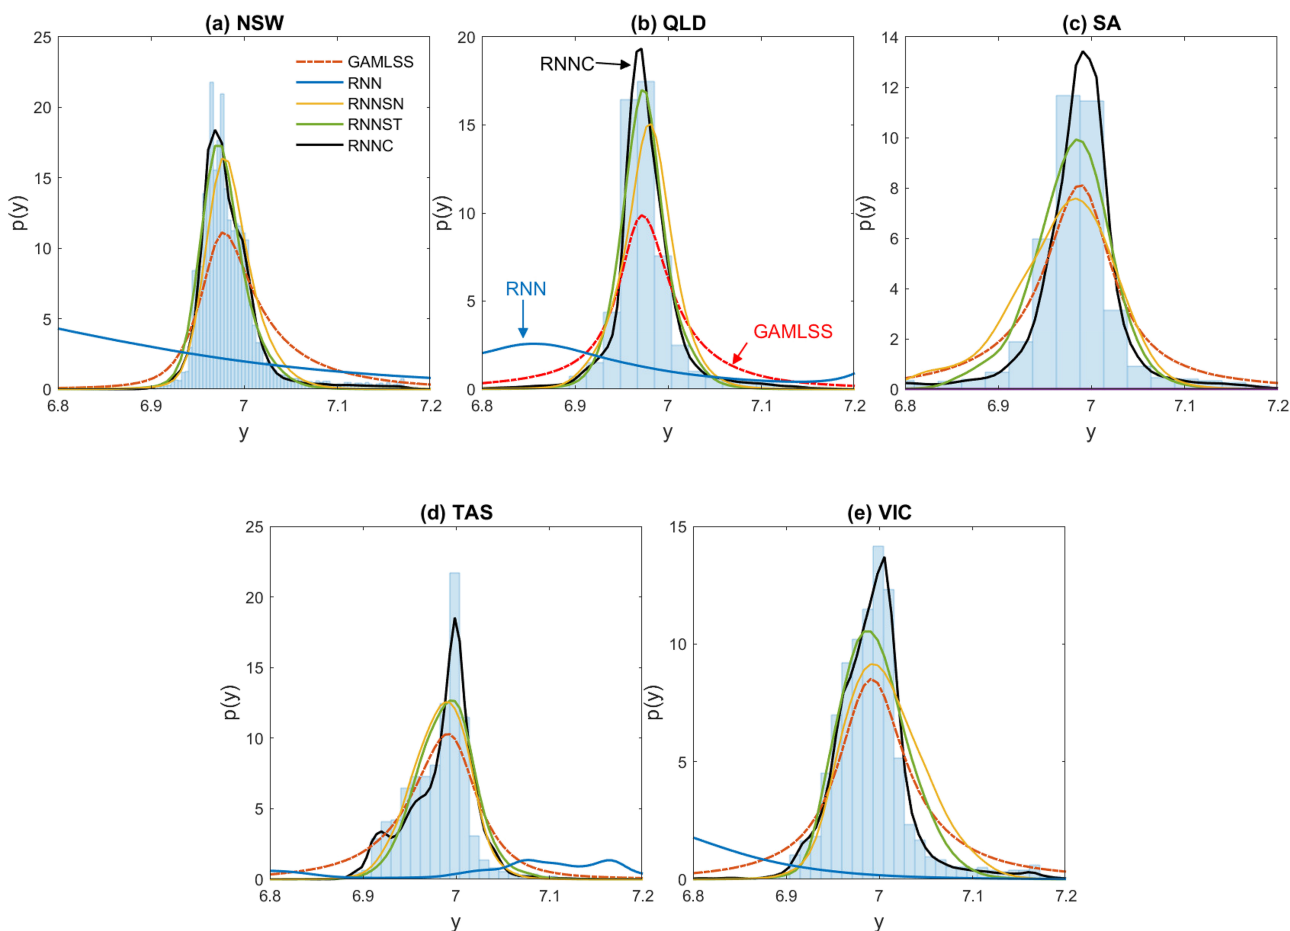


FIGURE 2 Out-of-sample predictive marginal calibration. *Note:* Histograms are of the out-of-sample observations of Y during the evaluation period (1 May–31 December 2019). The average predictive density for each method is given for the four deep time series models, plus the GAMLSS benchmark. The deep time series copula model (RNNC) produces average densities close to the bounded KDEs fit to the in-sample data found in Fig. B of the Web Appendix by construction. The panels correspond to forecast prices in (a) NSW, (b) QLD, (c) SA, (d) TAS and (e) VIC.

‘system-weighted’ average metric values, and they reflect the differing importance of price forecasting accuracy across regions in the NEM. For the fourth metric, the coverage is with respect to the price distributions in all five states.

5.3 | Empirical results

5.3.1 | Predictive distributions

Figure 2 presents histograms of the out-of-sample observations of Y for each region during the evaluation period. Also plotted are the average predictive densities \hat{F} for the five forecasting methods, allowing a visual comparison of out-of-sample marginal calibration. The copula model (RNNC) produces forecasts that are most accurately calibrated, followed by RNNNSN and RNNST. However, GAMLSS and RNN both exhibit poor predictive marginal calibration. The exceptionally wide intervals for RNN are caused by extremely inaccurate density forecasts at horizons close to 24 h ahead as characterized further below.

To visualize the predictive densities, Figure 3 plots the 95% prediction intervals for forecasts made during the month of July 2019. Results from the five methods for forecasts 1, 6 and 24 h ahead are given in separate panels. RNN forecasts are negatively biased, and for some days lack ‘sharpness’ (Gneiting et al., 2007). In contrast, skew-normal disturbances (RNNNSN) can produce a heavier upper tail through right skew but at the expense of overly sharp densities. De-coupling of the level of asymmetry from the kurtosis through skew- t disturbances (RNNST) allows for asymmetric predictive densities

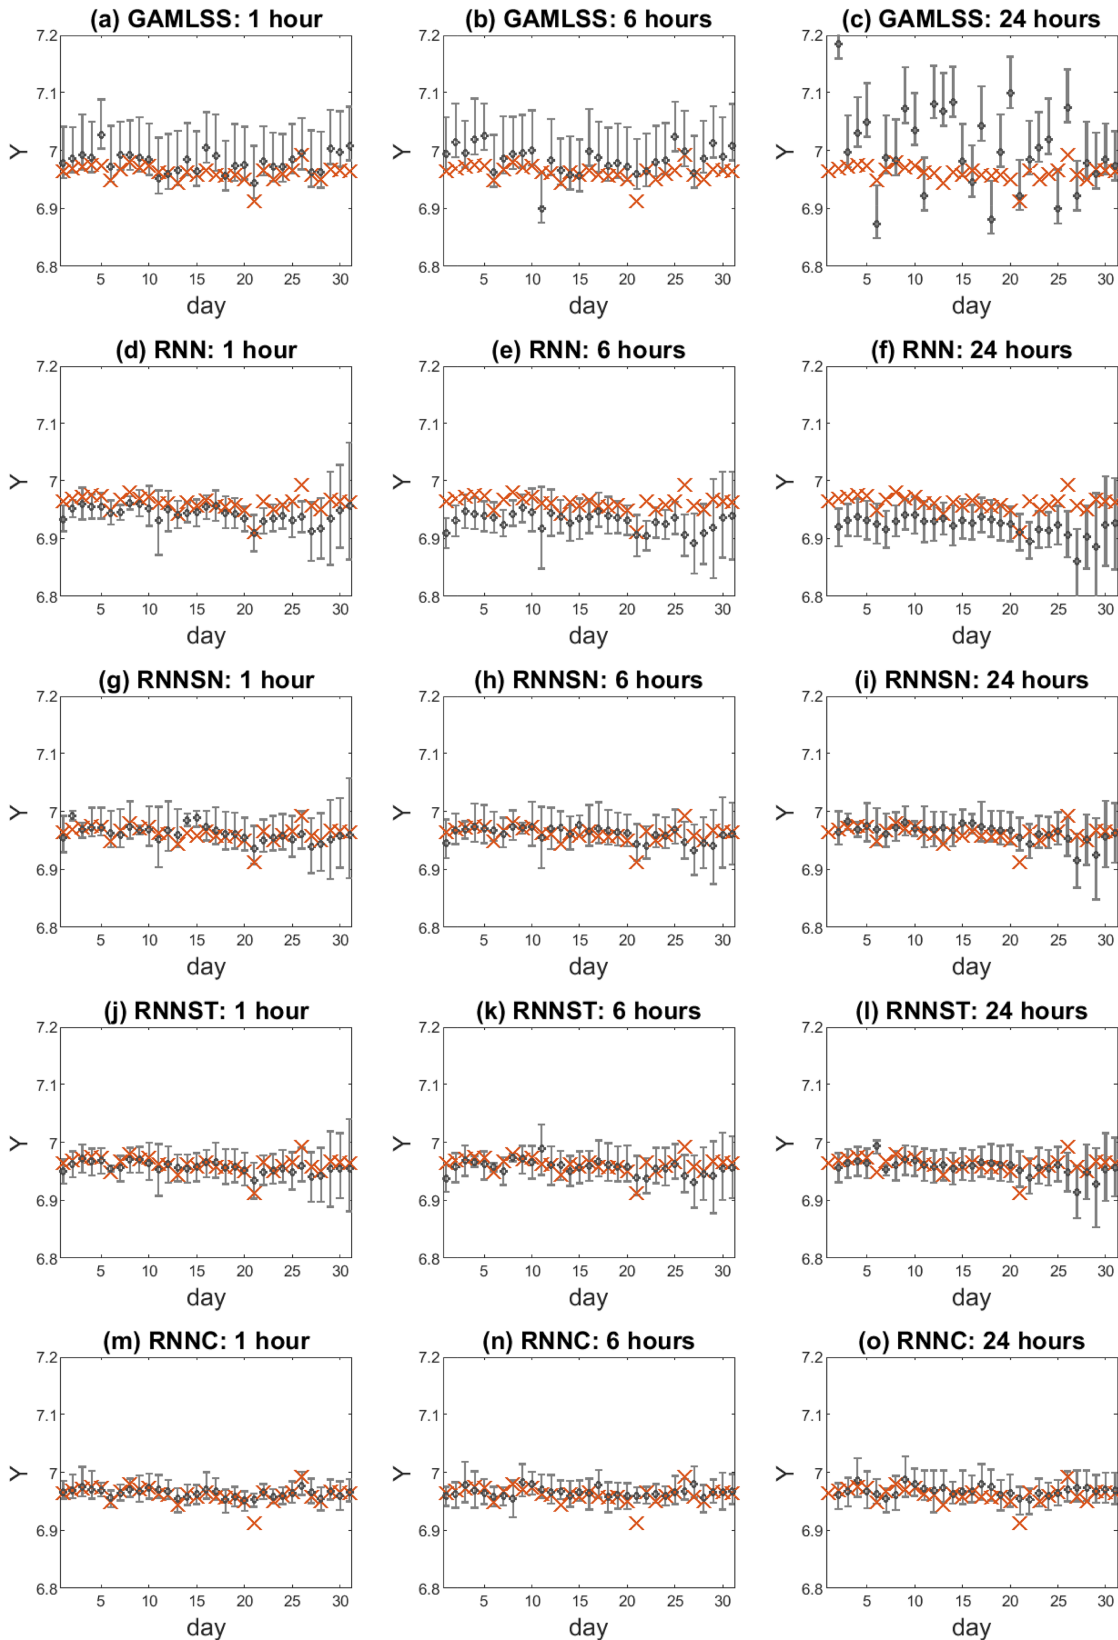


FIGURE 3 Prediction intervals for NSW prices at 12:00 during July 2019. *Note:* Panels correspond to forecasts from the five methods made at 1, 6 and 24 hours ahead. In each panel, 95% prediction intervals are denoted by a vertical bar, with the predictive mean given by a grey diamond. The true values are denoted with a red cross. Results are presented on the logarithmic scale for clarity.

with heavy tails. RNN, RNNSN and RNNST all have probabilistic forecasts that are nearly homoscedastic, because they are ensembles of homoscedastic Bayesian predictive densities. In contrast, RNNC allows for a much heavier upper tail for electricity prices and heteroscedasticity; the latter is because it is a distributional model. Last, the performance of the benchmark GAMLSS model (which is also a distributional model) declines substantially over the forecast horizon.

TABLE 1 Point forecast accuracy of models during the evaluation period.

Model	Hours ahead in the forecast horizon						
	0.5 h	1 h	2 h	3 h	6 h	12 h	24 h
MAE							
GAMLSS	0.0226***	0.0314***	0.0408***	0.0468***	0.0591***	0.0692***	0.0766***
RNN	0.0249***	0.0319***	0.0411***	0.0494***	0.0600***	0.0722***	0.0800***
RNNSN	0.0221***	0.0219*	0.0239	0.0252	0.0262	0.0268**	0.0278***
RNNST	0.0189	0.0210	0.0231	0.0242	0.0250	0.0255***	0.0257***
RNNC	0.0169	0.0194	0.0224	0.0238	0.0252	0.0244	0.0239
RMSE							
GAMLSS	0.0340 ⁺	0.0461 ⁺⁺	0.0610 ⁺⁺⁺	0.0623 ⁺⁺⁺	0.0644 ⁺⁺⁺	0.0653	0.0661
RNN	0.0537	0.0616**	0.0712	0.0797	0.0935	3.7173***	16.5496***
RNNSN	0.0538	0.0572	0.0603	0.0615	0.0623	0.0629	0.0638
RNNST	0.0528 ⁺	0.0571	0.0603	0.0612	0.0623	0.0630	0.0635
RNNC	0.0538	0.0573	0.0613	0.0638	0.0657	0.0632	0.0629

Note: Both the system-weighted mean absolute error (MAE) and root mean squared error (RMSE) are reported. Results are given for the four deep time series models, plus GAMLSS benchmark model. Results of a two-sided Diebold–Mariano test of the equality of the mean error of the RNNC with each of the other models are reported. Rejection of the null hypothesis of equal means at the 10%, 5% and 1% levels are denoted by ‘*’, ‘**’ and ‘***’ if favourable to the RNNC model, or by ‘+’, ‘++’ and ‘+++’ if unfavourable.

TABLE 2 Probabilistic forecast accuracy during the evaluation period.

Model	Hours ahead in the forecast horizon						
	0.5 h	1 h	2 h	3 h	6 h	12 h	24 h
CRPS							
GAMLSS	0.0180***	0.0252***	0.0335***	0.0392***	0.0508***	0.0601***	0.0652***
RNN	0.0203***	0.0235***	0.0299***	0.0364***	0.0449***	0.0574***	0.4061***
RNNSN	0.0174***	0.0177**	0.0191	0.0201	0.0207	0.0212**	0.0219***
RNNST	0.0157***	0.0172***	0.0187***	0.0195	0.0201	0.0206***	0.0208***
RNNC	0.0135	0.0154	0.0177	0.0189	0.0199	0.0193	0.0187
Joint upper tail loss at $\alpha = 0.975$							
GAMLSS	31.6394*	32.4717***	33.2367***	33.9396***	35.9352***	40.0541***	39.7960***
RNN	30.8568	31.5025	33.3171	35.5605	39.2732**	41.9374***	33.4037***
RNNSN	32.1052**	31.2797	31.6342	31.6492	31.7099	31.5852***	31.3780***
RNNST	31.3493*	31.4414	31.6911	32.0056*	31.8607	32.2372***	31.8694***
RNNC	30.1002	30.6244	31.0090	31.0934	31.1121	30.6238	29.7661
Quantile score at $\alpha = 0.95$							
GAMLSS	0.0056***	0.0069***	0.0082***	0.0093***	0.0124***	0.0190***	0.0198***
RNN	0.0050	0.0057	0.0078***	0.0107***	0.0155***	0.0201***	1.0808***
RNNSN	0.0056	0.0050	0.0054	0.0055	0.0057	0.0058*	0.0060***
RNNST	0.0048	0.0050	0.0054	0.0057**	0.0057*	0.0062***	0.0063***
RNNC	0.0036	0.0043	0.0048	0.0050	0.0052	0.0050	0.0046
Quantile score at $\alpha = 0.05$							
GAMLSS	0.0049***	0.0080***	0.0129***	0.0167***	0.0242***	0.0260***	0.0298***
RNN	0.0047	0.0053	0.0057***	0.0061	0.0069***	0.0096***	0.5602***
RNNSN	0.0046**	0.0055	0.0056	0.0062	0.0061	0.0064	0.0065
RNNST	0.0048***	0.0055*	0.0059**	0.0061	0.0064	0.0063	0.0066
RNNC	0.0039	0.0046	0.0053	0.0058	0.0061	0.0059	0.0057

Note: System-weighted metrics are reported for the four deep time series models, plus the GAMLSS benchmark, at different points in time during a 24 h forecast horizon. Lower values suggest greater accuracy. Results of a two-sided Diebold–Mariano test of the equality of the mean metric value for the RNNC with that of each of the other models are reported. Rejection of the null hypothesis of equal means at the 10%, 5% and 1% levels are denoted by ‘*’, ‘**’ and ‘***’ if favourable to the RNNC model, or by ‘+’, ‘++’ and ‘+++’ if unfavourable.

5.3.2 | Out-of-sample forecast accuracy

We compute the metrics outlined in Section 5.2 for each of the five methods. Table 1 reports point forecast accuracy, and the asymmetric deep time series models RNNNSN, RNNST and RNNC are more accurate than RNN and GAMLSS in terms of MAE at all horizons but not by the RMSE metric at short horizons. At horizons of 12 or more hours, RNNC produces the most accurate point forecasts. However, the main objective is accurate probabilistic forecasting, and Table 2 reports metrics for this. RNN is the poorest method by all measures, which is unsurprising given the asymmetry and heavy tails in even logarithmic prices. GAMLSS is also very poor overall. RNNST is more accurate than RNNNSN, demonstrating that allowing for heavy tails increases accuracy over-and-above allowing for asymmetry. However, it is RNNC that provides superior forecasts, both in terms of the entire density (as measured using CRPS) and the tails (as measured by the quantile scores and upper tail loss function). The improvement is greatest at a horizon of between 12 and 24 h, which is also the most important for informing bidding strategies in the NEM.

Figure F in the Web Appendix plots the full quantile score functions (pinball loss). The poor calibration of the RNN and GAMLSS probabilistic forecasts is clear, as is the superiority of the asymmetric deep time series models (in particular RNNC). To compare the probabilistic forecasts for the two most accurate models (RNNC and RNNST), Figure 4 plots the expected shortfall ($E(Y|Y < F^{-1}(\alpha))$ for $\alpha = 0.025$) and expected longrise ($E(Y|Y > F^{-1}(\alpha))$ for $\alpha = 0.975$) at 12:00 during July 2019 for both models. The RNNC produces predictive forecasts that exhibit substantial heteroscedasticity and are sharper (i.e., lower spread between the longrise and shortfall) than the RNNST. These differences in the two deep time series models are because RNNC is a distributional time series model, whereas RNNST is not. Finally, Table 3 reports the coverage of the predictive distributions at the 95% level, and there is a substantial difference between the methods, with the predictive distributions from RNNC exhibiting superior coverage at all horizons.

6 | INCORPORATING PROBABILISTIC FORECASTS OF DEMAND

We now show how to extend the deep time series models to include demand forecasts as additional information in a flexible fashion, allowing for greater price forecast accuracy.




FIGURE 4 Tail risk for NSW prices at 12:00 during July 2019. *Note:* Panels (a)-(c) correspond to forecasts made 1, 6 and 24 hours ahead for the RNNC (in blue) and RNNST (in red). In each panel, both the the expected longrise $E(Y|Y > Q(\alpha))$ at $\alpha = 0.975$ and the expected shortfall $E(Y|Y < Q(\alpha))$ at $\alpha = 0.025$ are plotted. Results are presented on the logarithmic scale for clarity.

Model	Hours ahead in the forecast horizon						
	0.5 h	1 h	2 h	3 h	6 h	12 h	24 h
GAMLSS	0.9078	0.8148	0.7278	0.6928	0.6476	0.5838	0.5208
RNN	0.8715	0.8609	0.7906	0.7173	0.6336	0.5698	0.6124
RNNNSN	0.8496	0.8386	0.8354	0.8187	0.8123	0.7998	0.7759
RNNST	0.8550	0.8424	0.8353	0.8216	0.8171	0.7955	0.7667
RNNC	0.9279	0.9012	0.8791	0.8741	0.8620	0.8594	0.8706

Note: Results are reported for the four deep time series models, plus the GAMLSS benchmark, at different points in time in the 24 h forecast horizon. Values closer to 0.95 indicate improved calibration of the probabilistic forecasts.

TABLE 3 Coverage of 95% prediction intervals during the evaluation period.

6.1 | Role of demand forecasts

Demand for electricity is almost perfectly inelastic to price in the short term because individual users face fixed tariffs. As demand varies over time, the spot price traces out the average supply curve. This explains the occurrence of price ‘spikes’ because electricity supply curves are typically kinked, and for short periods of time demand can exceed the location of this kink; see Clements et al. (2015) & Smith and Shively (2018) and Ziel and Steinert (2016). An implication is that including accurate demand forecasts in a price forecasting model may increase its accuracy, particularly for the upper tail. However, incorporating demand forecasts into most existing time series models for price is difficult for two reasons. First, aspects of probabilistic demand forecasts beyond the point forecast may increase the accuracy of price forecasts. Second, demand forecasts are likely to be related to future prices in a highly nonlinear fashion, particularly when combined with past price values. However, both difficulties can be addressed by incorporating some quantiles of the demand forecast distribution as additional elements of the feature vector of a deep time series model.

6.2 | NSW demand forecasts

Methods for accurate short-term probabilistic forecasting of demand are well established, including for NSW (Cottet & Smith, 2003) and AEMO makes real time demand forecasts publicly available. The forecasts are of the 10th, 50th (i.e., median) and 90th percentiles of demand at a half-hourly resolution over a horizon of one week. We employ demand forecasts for NSW that are exactly 24 h ahead, to forecast NSW prices 24 h ahead. Thus, only demand forecasts truly available at the forecast origin T were used. We label the three percentile forecasts of NSW demand as D10, D50 and D90. Figure E in Part E of the Web Appendix plots actual demand against D50 during 2019, showing the high degree of point forecast accuracy. However, empirical coverage during 2019 suggests some minor miscalibration, with 12.1%, 62.5% and 95.6% of demand observations falling below D10, D50 and D90, respectively.

6.3 | NSW price forecasting study

The forecasting study was extended, where the NSW demand forecasts were included in the feature vectors of our two best performing deep time series models, RNNST and RNNC, for forecasting NSW price. For RNNST

$$\mathbf{x}_t = (1, Y_{NSW,t-1}, Y_{NSW,t-2}, \dots, Y_{NSW,t-48}, Y_{NSW,t-96}, Y_{NSW,t-144}, \dots, Y_{NSW,t-336}, D10_t, D50_t, D90_t)',$$

so that $n_x = 58$. The feature vector for RNNC contained the equivalent lagged values of the transformed series $Z_{NSW,t} = \Phi_1^{-1}(F_Y(Y_{NSW,t}))$, and the three demand forecasts. For comparison, the two models without the demand forecasts (so $n_x = 55$) are also used.

Table 4 reports the average value of the forecast metrics over the evaluation period. Including the demand forecasts as features decreases the accuracy of the point forecasts and the CRPS metric for both models. However, upper tail forecast accuracy increased significantly for the RNNC. This is important in this application because short periods of very high prices often dominate the profitability of participants in the NEM. To illustrate, Figure 5 plots the density forecast where the inclusion of the demand forecast information greatly increased accuracy (other examples are given in Figure G of the Web Appendix).

TABLE 4 NSW price forecast accuracy when including demand forecast information.

Model	MAE	RMSE	CRPS	JS	QS95	QS05	C95
	Point and Density Accuracy			Tail Forecast Accuracy			
RNNC+D	0.0192	0.0314	0.0146	30.6902	0.0045	0.0025	82.6%
RNNC	0.0188+++	0.0311+++	0.0144+++	31.0350**	0.0048***	0.0023	84.5%
RNNST	0.0250***	0.0401	0.0221***	39.4411***	0.0121***	0.0049***	46.2%
RNNST+D	0.0256***	0.0395	0.0222***	39.6431***	0.0123***	0.0046***	53.9%

Note: Accuracy of 24 h ahead NSW (logarithmic) price forecasts during the evaluation period. Models that include 24 h ahead demand forecast information as features are labelled as ‘+D’. The metrics are described in the text, and include the joint score (JS) at $\alpha = 0.975$, quantile score (QS) at $\alpha = 0.95, 0.05$, and the empirical coverage of the 95% prediction intervals (C95). Results of a two-sided Diebold–Mariano test of the equality of the mean metric values of the RNNC+D with that of the other methods are reported. Rejection of the null hypothesis of equal means at the 10%, 5% and 1% levels are denoted by ‘*’, ‘**’ and ‘***’ if favourable to the RNNC+D model, or by ‘+’, ‘++’ and ‘+++’ if unfavourable.

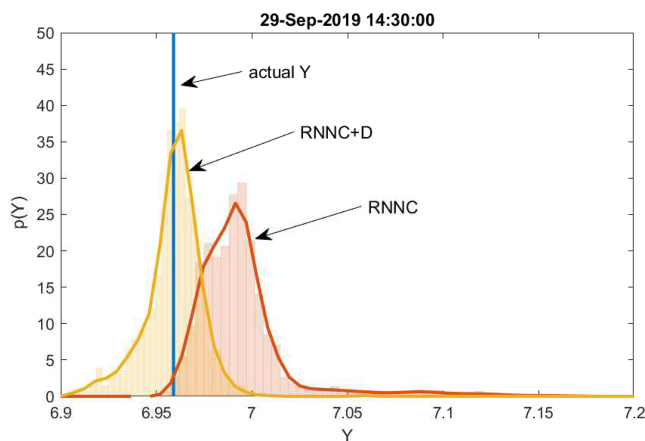


FIGURE 5 Density forecast of NSW price with and without demand forecast information. *Note:* Density forecast is of NSW electricity price on 29 September 2019 at 14:30, where the inclusion of demand forecast information greatly improved accuracy (as measured by CRPS). The density forecast without demand forecast information is labelled as ‘RNNC’, and with as ‘RNNC+D’. The actual price is marked with a blue vertical line.

Finally, we also experimented with employing the ESN at (1) without the quadratic term in the output layer but found that this decreased forecast accuracy significantly. Moreover, including different ridge priors for the linear and quadratic terms in the output layer had little effect on forecast accuracy; see Part F of the Web Appendix for further details.

7 | CONCLUSION

This paper makes contributions to both the time series and energy economics literatures. We propose new deep time series models that exploit the reservoir computing techniques found in ESNs for high frequency time series. Recasting a quadratic ESN as a statistical time series model, with a shrinkage prior for regularization of the output layer weights, allows for uncertainty quantification and probabilistic forecasting. Our first approach is to allow for different error distributions for the output layer, extending the Bayesian methodology of McDermott and Wikle (2017, 2019). However, our main methodological contribution is the proposal of a new deep *distributional* time series model. This is obtained by constructing the implicit copula of a Gaussian probabilistic ESN, and extends the deep distributional regression method of Klein et al. (2021) to time series. This copula is a deep process on the feature space, allowing for highly adaptive nonlinear serial dependence, and generalizing existing echo state Gaussian processes (Chatzis & Demiris, 2011). When combined with a nonparametric estimate of p_Y , it allows for marginal calibration. Our empirical work suggests that accurate in-sample marginal calibration also results in more accurate calibration of both the out-of-sample marginal and predictive distributions.

Our deep time series models are not Markov. Those defined in Section 3 have additive homoscedastic errors and are stationary. Moreover, it is straightforward to show (e.g., Smith, 2022) that the implicit copula of a stationary process is also stationary. When this copula is combined with an invariant margin p_Y , as in Section 4, so is the resulting copula time series model. Each density forecast at (5) is an ensemble of Bayesian predictive posterior densities, which are well-established in probabilistic forecasting (Harrison & Stevens, 1976). Because the posterior predictive density from the copula model at (10) is a function of the feature vector, the forecasts from this model are conditionally heteroscedastic.

Although the distribution of intraday electricity prices have a sizable predictable component, these time series are complex (Ignatieva & Trück, 2016; Panagiotelis & Smith, 2008). Key to their accurate modelling and forecasting is to capture jointly three features: nonlinear serial dependence, high levels of asymmetry and kurtosis, and strong time-variation in the entire distribution (the latter feature includes heteroscedasticity and time-varying tail risk). All our proposed deep time series models account for the first feature, while allowing for skew-normal or skew- t disturbances also accounts for the second. However, only the deep copula model allows for all three, producing the most accurate probabilistic forecasts. Recent models (Shah & Lisi, 2020; Ziel & Steinert, 2016) suggest that including demand forecasts may further increase probabilistic forecasting accuracy. We find this to be the case for the upper tail of the 24 h ahead NSW electricity prices when including demand forecasts as features in the deep copula model. Moreover, the inclusion of other available exogenous variables as features is straightforward and may also improve forecast accuracy as studies by Gianfreda and Bunn (2018), Gianfreda et al. (2020) and others suggest.

Last, we highlight two areas for future work. First, estimation of our deep time series models using variational methods, rather than MCMC, has the potential to speed up computations. Second, there are other applications of our deep time series models. One example is macroeconomic forecasting, where large regularized time-varying parameter models are

popular (Bitto & Frühwirth-Schnatter, 2019; Carriero et al., 2019; Huber et al., 2021). The deep time series models provide a flexible alternative, as they can incorporate feature vectors of many economic variables regularized through the reservoir structure of the ESN, plus Bayesian shrinkage of the output layer coefficients. Moreover, the copula model produces time-varying and asymmetric probabilistic forecasts, which is a feature of this problem.

ACKNOWLEDGEMENTS

Nadja Klein was supported by the Deutsche Forschungsgemeinschaft (DFG, German research foundation) through the Emmy Noether grant KL 3037/1-1. David Nott is affiliated with the Institute of Operations Research and Analytics at the National University of Singapore. The authors thank Sonnia Fuenteseca for invaluable research assistance in constructing the demand forecast data, and Patrick McDermott and Chris Wikle for sharing their code. They also thank two referees and the handling editor for their encouragement and comments that helped improve the paper.

OPEN RESEARCH BADGES



This article has been awarded Open Data Badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. Data is available at <http://qed.econ.queensu.ca/jae/datasets/klein003/>.

DATA AVAILABILITY STATEMENT

The data that supports the findings of this study are available in the supplementary material of this article, as well as at the Journal of Applied Econometrics Data Archive found at <http://dx.doi.org/10.15456/jae.2022327.1200725211>.

REFERENCES

- Amjady, N. (2006). Day-ahead price forecasting of electricity markets by a new fuzzy neural network. *IEEE Transactions on Power Systems*, 21(2), 887–896.
- Apergis, N., Gozgor, G., Lau, C. K. M., & Wang, S. (2019). Decoding the Australian electricity market: New evidence from three-regime hidden semi-Markov model. *Energy Economics*, 78, 129–142.
- Australian Energy Regulator (2019). State of the energy market report: data update.
- Azzalini, A., & Capitanio, A. (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society: Series B*, 65(2), 367–389.
- Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond risk premiums with machine learning. *The Review of Financial Studies*, 34(2), 1046–1089.
- Bitto, A., & Frühwirth-Schnatter, S. (2019). Achieving shrinkage in a time-varying parameter model framework. *Journal of Econometrics*, 210(1), 75–97.
- Bunn, D., Andresen, A., Chen, D., & Westgaard, S. (2016). Analysis and forecasting of electricity price risks with quantile factor models. *The Energy Journal*, 37(1), 101–122.
- Carriero, A., Clark, T. E., & Marcellino, M. (2019). Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors. *Journal of Econometrics*, 212(1), 137–154.
- Chatzis, S. P., & Demiris, Y. (2011). Echo state Gaussian process. *IEEE Transactions on Neural Networks*, 22(9), 1435–1445.
- Christensen, T. M., Hurn, A. S., & Lindsay, K. A. (2012). Forecasting spikes in electricity prices. *International Journal of Forecasting*, 28(2), 400–411.
- Clements, A., Herrera, R., & Hurn, A. (2015). Modelling interregional links in electricity price spikes. *Energy Economics*, 51, 383–393.
- Cottet, R., & Smith, M. (2003). Bayesian modeling and forecasting of intraday electricity load. *Journal of the American Statistical Association*, 98(464), 839–849.
- Feng, G., He, J., & Polson, N. G. (2018). Deep learning for predicting asset returns. arXiv preprint arxiv:1804.09314.
- Fissler, T., & Ziegel, J. F. (2016). Higher order elicibility and Osband's principle. *The Annals of Statistics*, 44(4), 1680–1707.
- Frühwirth-Schnatter, S., & Pyne, S. (2010). Bayesian inference for finite mixtures of univariate and multivariate skew-normal and skew-t distributions. *Biostatistics*, 11(2), 317–336.
- Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of The 33rd International Conference on Machine Learning*, 48, pp. 1050–1059. PMLR.
- Gallicchio, C. (2020). Sparsity in reservoir computing neural networks. In *2020 International Conference on INnovations in Intelligent Systems and Applications (INISTA)*, pp. 1–7.
- Gianfreda, A., & Bunn, D. (2018). A stochastic latent moment model for electricity price formation. *Operations Research*, 66(5), 1189–1203.
- Gianfreda, A., Ravazzolo, F., & Rossini, L. (2020). Comparing the forecasting performances of linear models for electricity prices with high res penetration. *International Journal of Forecasting*, 36(3), 974–986.
- Gneiting, T. (2011). Quantiles as optimal point forecasts. *International Journal of forecasting*, 27(2), 197–207.

- Gneiting, T., Balabdaoui, F., & Raftery, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society Series B*, 69(2), 243–268.
- Gneiting, T., & Katzfuss, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1, 125–151.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- Gu, S., Kelly, B., & Xiu, D. (2021). Autoencoder asset pricing models. *Journal of Econometrics*, 222(1), 429–450.
- Han, L., Kordzakhia, N., & Trück, S. (2020). Volatility spillovers in Australian electricity markets. *Energy Economics*, 90, 104782.
- Harrison, P. J., & Stevens, C. F. (1976). Bayesian forecasting. *Journal of the Royal Statistical Society: Series B*, 38(3), 205–228.
- Higgs, H. (2009). Modelling price and volatility inter-relationships in the Australian wholesale spot electricity markets. *Energy Economics*, 31(5), 748–756.
- Huber, F., Koop, G., & Onorante, L. (2021). Inducing sparsity and shrinkage in time-varying parameter models. *Journal of Business and Economic Statistics*, 39(3), 669–683.
- Huurman, C., Ravazzolo, F., & Zhou, C. (2012). The power of weather. *Computational Statistics & Data Analysis*, 56(11), 3793–3807.
- Ignatieva, K., & Trück, S. (2016). Modeling spot price dependence in Australian electricity markets with applications to risk management. *Computers & Operations Research*, 66, 415–433.
- Jaeger, H. (2007). Echo state network. *Scholarpedia*, 2(9), 2330.
- Janczura, J., & Weron, R. (2010). An empirical comparison of alternate regime-switching models for electricity spot prices. *Energy economics*, 32(5), 1059–1073.
- Karakatsani, N. V., & Bunn, D. W. (2008). Intra-day and regime-switching dynamics in electricity price formation. *Energy Economics*, 30(4), 1776–1797.
- Kirschen, D. S., & Strbac, G. (2018). *Fundamentals of Power System Economics*. John Wiley & Sons.
- Klein, N., & Kneib, T. (2016). Scale-dependent priors for variance parameters in structured additive distributional regression. *Bayesian Analysis*, 11(4), 1071–1106.
- Klein, N., Nott, D. J., & Smith, M. S. (2021). Marginally calibrated deep distributional regression. *Journal of Computational and Graphical Statistics*, 30(2), 467–483.
- Klein, N., & Smith, M. S. (2019). Implicit copulas from Bayesian regularized regression smoothers. *Bayesian Analysis*, 14(4), 1143–1171.
- Knittel, C. R., & Roberts, M. R. (2005). An empirical examination of restructured electricity prices. *Energy Economics*, 27(5), 791–817.
- Lago, J., De Ridder, F., & De Schutter, B. (2018). Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied Energy*, 221, 386–405.
- Lukoševičius, M., & Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3), 127–149.
- Manner, H., Fard, F. A., Pourkhanali, A., & Tafakorli, L. (2019). Forecasting the joint distribution of Australian electricity prices using dynamic vine copulae. *Energy Economics*, 78, 143–164.
- Manner, H., Türk, D., & Eichler, M. (2016). Modeling and forecasting multivariate electricity price spikes. *Energy Economics*, 60, 255–265.
- McDermott, P. L., & Wikle, C. K. (2017). An ensemble quadratic echo state network for non-linear spatio-temporal forecasting. *Stat*, 6(1), 315–330.
- McDermott, P. L., & Wikle, C. K. (2019). Deep echo state networks with uncertainty quantification for spatio-temporal forecasting. *Environmetrics*, 30(3), e2553.
- Misiorek, A., Trueck, S., & Weron, R. (2006). Point and interval forecasting of spot electricity prices: Linear vs. non-linear time series models. *Studies in Nonlinear Dynamics & Econometrics*, 10(3). <https://doi.org/10.2202/1558-3708.1362>
- Narajewski, M., & Ziel, F. (2020). Ensemble forecasting for intraday electricity prices: Simulating trajectories. *Applied Energy*, 279(1), 115801.
- Nolde, N., & Ziegel, J. F. (2017). Elicitability and backtesting: Perspectives for banking regulation. *The Annals of Applied Statistics*, 11(4), 1833–1874.
- Nowotarski, J., Tomczyk, J., & Weron, R. (2013). Robust estimation and forecasting of the long-term seasonal component of electricity spot prices. *Energy Economics*, 39, 13–27.
- Nowotarski, J., & Weron, R. (2018). Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and Sustainable Energy Reviews*, 81, 1548–1568.
- Panagiotelis, A., & Smith, M. (2008). Bayesian density forecasting of intraday electricity prices using multivariate skew t distributions. *International Journal of Forecasting*, 24(4), 710–727.
- Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13*, pp. III–1310–III–1318.
- Pircalabu, A., & Benth, F. E. (2017). A regime-switching copula approach to modeling day-ahead prices in coupled electricity markets. *Energy Economics*, 68, 283–302.
- Rafiei, M., Niknam, T., & Khooban, M.-H. (2016). Probabilistic forecasting of hourly electricity price by generalization of elm for usage in improved wavelet neural network. *IEEE Transactions on Industrial Informatics*, 13(1), 71–79.
- Rigby, R. A., & Stasinopoulos, D. M. (2005). Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society Series C*, 54(3), 507–554.
- Serinaldi, F. (2011). Distributional modeling and short-term forecasting of electricity prices by generalized additive models for location, scale and shape. *Energy Economics*, 33(6), 1216–1226.

- Shah, I., & Lisi, F. (2020). Forecasting of electricity price through a functional prediction of sale and purchase curves. *Journal of Forecasting*, 39(2), 242–259.
- Smith, M. S. (2015). Copula modelling of dependence in multivariate time series. *International Journal of Forecasting*, 31(3), 815–833.
- Smith, M. S. (2022). Implicit copulas: An overview. *Econometrics and Statistics*. in press. <https://doi.org/10.1016/j.ecosta.2021.12.002>
- Smith, M. S., Gan, Q., & Kohn, R. J. (2012). Modelling dependence using skew t copulas: Bayesian inference and applications. *Journal of Applied Econometrics*, 27(3), 500–522.
- Smith, M. S., & Shively, T. S. (2018). Econometric modeling of regional electricity spot prices in the Australian market. *Energy Economics*, 74, 886–903.
- Ugurlu, U., Oksuz, I., & Tas, O. (2018). Electricity price forecasting using recurrent neural networks. *Energies*, 11(5), 1255.
- 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.
- Wilson, A. G., & Ghahramani, Z. (2010). Copula processes. *Advances in Neural Information Processing Systems*, 23, 2460–2468.
- Ziel, F., & Steinert, R. (2016). Electricity price forecasting using sale and purchase curves: The X-Model. *Energy Economics*, 59, 435–454.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of the article.

How to cite this article: Klein, N., Smith, M. S., & Nott, D. J. (2023). Deep distributional time series models and the probabilistic forecasting of intraday electricity prices. *Journal of Applied Econometrics*, 38(4), 493–511. <https://doi.org/10.1002/jae.2959>