ANZIAM J. 57(2016), 369–383
doi:10.1017/S1446181115000322

A LOGNORMAL MODEL FOR DEMAND FORECASTING IN THE NATIONAL ELECTRICITY MARKET

J. MAISANO✉[1,3], A. RADCHIK[2,3] and T. LING[3]

(Received 20 October, 2014; accepted 4 June, 2015; first published online 15 February 2016)

Abstract

Many electricity market participants have a requirement to calculate the probabilistic risk measures, such as earnings at risk (EaR) and value at risk (VaR), for compliance reporting purposes. This requirement is currently hindered by the lack of analytical representations for forecasts of demand (load) and price curves; this motivates numerical simulation and models that need extensive calibration. In this paper, we derive an analytical representation of a state demand forecast which is the aggregated usage of all electricity consumers in a particular region (such as New South Wales or Victoria). We have used two probabilistic benchmarks from the Australian energy market operator as input, which are expressed as forecasted probability of exceedance.

Due to a number of considerations, including asymmetry of these quantiles with respect to the median, we have selected a series of truncated lognormal distributions with two parameters. The procedure of finding these parameters has been reduced to solving (for every half-hour) a single nonlinear equation. As a result, the two-year half-hourly forecast (expected curve) and demand volatility are found by explicit integration with the set of derived distributions. We have also tested an alternative method based on simplifying assumptions; using a nontruncated lognormal distribution, we found that under the test conditions this method produces an identical forward load and volatility curve.

2010 *Mathematics subject classification*: 00A69.

Keywords and phrases: electricity, demand forecasting, national electricity market.

1. Introduction

Until recently, most electricity grids were run by the states as an essential service utility. In the late 1980s, it was realized that markets could be more efficient in driving the operation of the power system and thus electricity markets and market operators were created.

[1]Trading Technology Australia Pty. Ltd., Australia; e-mail: joe@tta.com.au.
[2]Green Trading Systems Pty. Ltd., Australia; e-mail: aradchik@aksong.com.
[3]University of Technology Sydney, NSW, Australia; e-mail: Timothy.Ling-1@uts.edu.au.

The role of the market operator was to balance supply and demand in each grid node and provide energy dispatch instructions to all grid-connected generators as well as ensuring the security and quality of the power supply. Market operators soon developed into sophisticated command and control centres. They developed optimization of power flow (OPF) computer algorithms for real-time monitoring of voltages and currents in grid nodes, as well as running market systems that take offers from generators to reduce the cost of production by automatically issuing dispatch instructions. At this point the physical energy pool emerged.

Note that in different geographical regions, electricity markets have evolved differently. Some are mainly bilateral, with the pool covering the "overs and unders" (net pool), while others are "gross-pool" markets, where all power must be sold via the market. We are considering only the latter in this paper, since this approach is consistent with a optimized central dispatch mechanism. In southeast Australia, the national electricity market management company limited (NEMMCO) was created in 1996, and the national electricity market (NEM) commenced in 1998. In 2009, the Australian energy market operator (AEMO) was established to operate both the NEM and gas markets. The independent market operator (IMO) was created in 2005, which manages the energy market in western Australia. In this paper, we will use examples of the NEM, since it is a gross-pool energy-only market while the IMO is a net-pool market with capacity payments.

Economic drivers are as important for the development of energy markets as the laws of physics. Under the influence of the economic drivers, the pool model has been redefined as follows:

- the pool is the sole buyer and seller of electricity;
- the pool collects offers from the suppliers and bids from the consumers to determine the set of successful bidders whose offers and bids are accepted;
- the market operator finds the "optimum" pool price by solving a centralized economic dispatch model that takes into account the physical network constraints (physical state of the grid) to determine lowest possible spot nodal prices; and
- the solution must be unique to satisfy a supply–demand equilibrium given the physical state of the grid.

Given that demand is the primary driver of electricity price in the NEM, market participants have a requirement to forecast state demand, which is the aggregated usage of all electricity consumers in a particular region (such as New South Wales or Victoria). As demand (otherwise referred to as load, as it is the load on the electrical system) changes continuously over time, we require a forward load (demand) curve to describe the state energy consumption forecast.

Although changes in demand are continuous, bidding in the market follows a reverse auction process; therefore, bids (comprising generation capacity to be dispatched) must be set in discrete time (every five minutes in the case of the NEM). Ancillary service providers (fast-ramping generators) supply or reduce generation to counter the accidental deviations in supply–demand balance.

Prices are set every half-hour by averaging the six (five-minute) price calculations and customer demand and financial contracts are settled at these half-hourly price-demand levels. Therefore, the most useful structure of a forward load (demand) curve, especially one intended to model settlements of financial instruments, is a discrete-time (or piecewise) curve with breakpoints every half-hour.

2. Demand forecasts

The AEMO publishes two forecasts of demand in the NEM called STPASA and MTPASA (short-term and medium-term projected assessment of system adequacy, respectively). Strictly speaking, the role of these forecasts is the projected system adequacy of the NEM. System adequacy is measured in terms of how much system capacity is remaining at any point in time (to absorb any sudden spikes in demand). If system adequacy is insufficient, blackouts will occur. The AEMO has an operating charter that compels continuous monitoring of the adequacy of the system and, if certain metrics are not met, additional generation capacity must be introduced (see, for example, [2] for a description of the short-term forecasting process).

However, rather than providing actual demand forecasts, PASA comprises the forecast error in terms of probability of exceedance (POE) for 10%, 50% and 90% confidence (MTPASA only quotes 10% and 50% POE). For compliance purposes, we require a forecast methodology based on an independent and transparent benchmark; therefore, we have based a demand forecast on the PASA shape and levels.

- STPASA file represents six trading days from the end of the trading day covered by the most recent pre-dispatch schedule with a half-hourly resolution. It contains 10%, 50% and 90% short-term probability of exceedance (STPOE) as well as short-term maximum available capacity (STMAC) forecast. It is updated every two hours.
- MTPASA file covers 24 months from the Sunday after the day of publication with a daily resolution. It contains 10% and 50% medium-term probability of exceedance (MTPOE) and maximum available capacity (MTMAC) forecast. It is updated weekly.

3. Constructing "ModPASA"

Our aim is to obtain the 2-year half-hourly demand forecast which would comply with both PASAs, that is, to be consistent with the STPASA profile and match levels of the MTPASA. Therefore, we develop the following procedure for construction of the forecast.

(1) We have taken the average of all most recent STPASA datasets (1 year history) to obtain a single time series which we use as a proxy for the forecast demand

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

shape over one year (as any particular forecast will only go out for six days). By this method, we obtain the full year of the forecast with the average forecasting error of the STPOE and STMAC included at each point.

(2) For each of the 10% and 50% STPOE and STMAC, we have divided by their respective daily maxima.

(3) The resultant (normalized) 10% and 50% time series are multiplied by their corresponding MTPOEs.

As a result of this "stretching" procedure we have obtained two-year projections for POE and maximum available capacities (MACs) while preserving the half-hourly profile of STPASA. *From now on we name this dataset* ModPASA. For further development, it is more conventional to work in terms of quantile time series rather than POEs. Therefore, for every half-hourly time period i, we have used

$$Q_{\alpha\%}(i) = \mathrm{POE}_{1-\alpha\%}(i).$$

We have found that for all ticks, $Q_{50\%} - Q_{10\%} \neq Q_{90\%} - Q_{50\%}$, which clearly indicates the asymmetry of the underlying distribution for demand values.

Since we only want to use STPASA for the "daily shape", we are not concerned with the overall levels of STPASA. Instead, we wish to rescale the STPASA to MTPASA levels while retaining the shape of STPASA (since MTPASA does not feature a daily shape). We do this by dividing each value of STPASA by its maximum daily value and then multiplying the resulting value by the corresponding MTPASA number. The same procedure applies to MAC and, using the following notation, it can be represented as

$$L_{\mathrm{MAX}} = \frac{\mathrm{STMAC}_i}{\mathrm{MAX}(\mathrm{STMAC}_{1-48})} \times \mathrm{MTMAC},$$

where

STMAC_i = maximum available capacity of STPASA for this half-hourly point,

MTMAC = maximum available capacity of MTPASA for this day,

$\mathrm{MAX}(\mathrm{STMAC}_{1-48})$ = maximum STMAC across the day.

For further calculations we introduce the following "rescaled" quantiles for convenience:

$$\begin{cases} y_1 = \dfrac{Q_{50\%}}{L_{\mathrm{MTMAC}}}, \\ y_2 = \dfrac{Q_{90\%}}{L_{\mathrm{MAX}}}. \end{cases} \tag{3.1}$$

We have deliberately omitted here the index i, whilst keeping in mind that the numerical procedure must be repeated for every single half-hourly interval of the ModPASA dataset (35 040 intervals in total, and 35 088 if a leap year is included).

4. Deriving the distribution and forward curves

Our choice for the distribution for demand values was dictated by the following considerations.

(1) The probability density function (PDF) should have two parameters in order to match two known ModPASA quantiles.
(2) The probability of zero demand occurrence should be zero.
(3) Support asymmetry, that is, dependent on ModPASA values to be skewed.
(4) Since a set of parameters would have to be found for every tick two years forward ($2 \times 2 \times 17\,520 = 70\,800$ values in total), it must be "computation friendly". This point is extremely important for a choice of distribution, for example, while the inverse Gamma distribution [3] is quite popular in financial literature, and could under certain values of parameters satisfy the conditions 1 to 4; solving this number of simultaneous transcendental equations is nontrivial and could be the focus of a separate paper.
(5) As we make the assumption that supply will always meet demand (that is, ancillary services are not invoked), all allowable values for demand should be contained within the interval $(0, L_{\text{MAX}})$. That imposes a truncation-caused normalization condition on the probability density function represented by the standard lognormal distribution; here, truncation is necessary as load changes from 0 to L_{MAX}, whilst for standard lognormal, the domain is $(0, \infty)$.

We have chosen a truncated lognormal distribution (TLD), as it meets the above criteria and can be integrated analytically.

$$f(\sigma, \mu; x) = \frac{2}{Zx\sigma\sqrt{2\pi}} e^{-(\ln x - \mu)^2/2\sigma^2}, \quad x \in [0, 1). \tag{4.1}$$

Here, we have *truncated* the distribution by restricting the domain to $[0, L_{\text{MAX}})$ (by rescaling (3.1) to $[0, 1)$). We have then normalized the distribution to make x have a probability density confined within the domain $[0, 1)$. The normalization constant Z and parameters σ and μ are yet to be defined. Also, as a consequence of (3.1), we know that the 50% quantile (y_1) is actually the median; therefore, we can assume that

$$e^{\mu} \approx y_1. \tag{4.2}$$

Note that while generally e^{μ} is the median of a lognormal distribution, since we have a truncated distribution, we say that it is *approximately* the median.

Using the Wolfram Online Integrator (integrals.wolfram.com/index.jsp), we obtained the corresponding cumulative distribution function (CDF) for the yth quantile by integration of (4.1) using

$$F(y) = \frac{1}{Z}\left\{\operatorname{erf}\left(\frac{1}{\sigma\sqrt{2}} \ln \frac{y}{y_1}\right) + 1\right\}. \tag{4.3}$$

Here, erf is an error function [1] defined as

$$\operatorname{erf}(z) \equiv \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2}\, dt.$$

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

The forward load curve (FLC) is the first moment of the PDF (4.1) and, keeping in mind (3.1) and the need to solve for all time periods t_i,

$$\mathrm{FLC}(t_i) = \frac{L_{\mathrm{MAX}}}{Z(t_i)} y_1(t_i) e^{\sigma^2(t_i)/2}\bigg[1 - \mathrm{erf}\bigg\{\frac{\sigma(t_i)}{\sqrt{2}}\bigg(\frac{\ln y_1(t_i)}{\sigma^2(t_i)} + 1\bigg)\bigg\}\bigg]$$

and the second moment

$$M_2 = \frac{L_{\mathrm{MAX}}^2}{Z(t_i)} y_1^2(t_i) e^{2\sigma^2(t_i)}\bigg\{1 - \mathrm{erf}\bigg(\frac{\ln y_1(t_i)}{\sigma(t_i)\sqrt{2}} + \sigma(t_i)\sqrt{2}\bigg)\bigg\}.$$

Since we have a separate volatility for each time period, we have a volatility curve; therefore, the forward load volatility curve (FLVC) is

$$\mathrm{FLVC} = \sqrt{\mathrm{FLC}^2 - M_2}.$$

Now, all that is left is to find a normalization constant Z and a parameter σ. We will proceed with the change of variables

$$v = \frac{1}{\sqrt{2\sigma^2}}. \tag{4.4}$$

For two given ModPASA quantiles and normalization, condition (4.3) yields the system of simultaneous equations

$$\begin{cases} \mathrm{erf}(-v \ln y_1) = Z - 1, \\ \mathrm{erf}\Big(v \ln \dfrac{y_2}{y_1}\Big) = 0.9Z - 1. \end{cases} \tag{4.5}$$

From the first equation in (4.5), $Z = 1 - \mathrm{erf}(v \ln y_1)$ and the resulting equation is

$$G(v) = \mathrm{erf}\Big(v \ln \frac{y_2}{y_1}\Big) + 0.9\,\mathrm{erf}(v \ln y_1) + 0.1 = 0. \tag{4.6}$$

Equation (4.6) is solved for every value of $y_1(t_i)$ and $y_2(t_i)$, where $i \in (1, 35\,040)$ (that is, half-hourly intervals for two years). This equation has two roots, v_1 and v_2 (as shown in Figure 1), both of which give valid solutions. We choose the root by backtesting the forward curve against historical data – one of the roots results in volatilities that are clearly too high (see Figure 2(a) and (b) – in this case we have re-averaged the PASA forecast to historical maxima rather than MTPASA in order to find the best fit against actual load).

We need to prove that equation (4.6) can have a maximum of two roots. For any quantile $Q(y_2)$, $y_2 \neq y_1$. (Note that AEMO publishes two different POEs for PASA, which are used in construction of the system of simultaneous equations (4.6). Therefore, the case $y_2 = y_1$ is meaningless as equation (4.6) turns into two independent equations with the same unknown.)

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

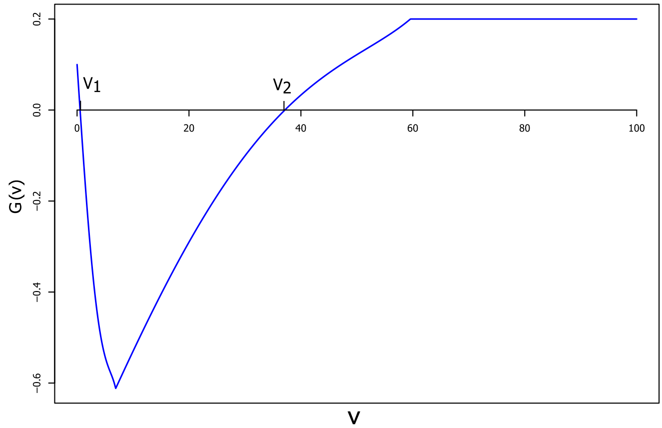

FIGURE 1. Solutions for equation (4.6) (NSW).

For the function (obtained from (4.6))

$$G(v) = \mathrm{erf}\Big(v \ln \frac{y_2}{y_1}\Big) + Q\,\mathrm{erf}(v \ln y_1) + 1 - Q$$

to have two roots, its derivative must have only one extremum. By differentiating,

$$G'(v) = -2v \ln^2 \frac{y_2}{y_1} e^{-[v \ln(y_2/y_1)]^2} - 2Qv \ln^2 y_1 e^{-(v \ln y_1)^2} = 0.$$

Since volatility is always greater than 0, we have either

$$e^{(v \ln y_1)^2 - [v \ln(y_2/y_1)]^2} = Q\Big[\frac{\ln y_1}{\ln(y_1/y_2)}\Big]^2$$

or

$$\begin{aligned} v &= \pm\sqrt{\frac{1}{(\ln y_1)^2 - [\ln (y_2/y_1)]^2} \ln\Big[Q\Big\{\frac{\ln y_1}{\ln(y_1/y_2))}\Big\}^2\Big]} \\ &= \pm\sqrt{\frac{1}{\ln y_2(2\ln y_1 - \ln y_2)} \ln\Big[Q\Big\{\frac{\ln y_1}{\ln(y_1/y_2)}\Big\}^2\Big]}. \end{aligned} \tag{4.7}$$

Again, since volatility is always greater than 0, the derivative can have up to one root that corresponds to the "+" sign in the right-hand side of equation (4.7), which means that $G(v)$ can have only one extremum and, therefore, a maximum of two roots.

In practice, we have not come across a case (solving for 70 000 points with real-world data) where $G(v)$ does not solve for a root. However, we cannot guarantee that $G(v)$ will have at least one root, which is a weakness of this model.

5. Simplified model

Although the arbitrary market price cap superimposes a fixed cut-off to our distribution (caused by limited generation available for every half-hour), we will

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

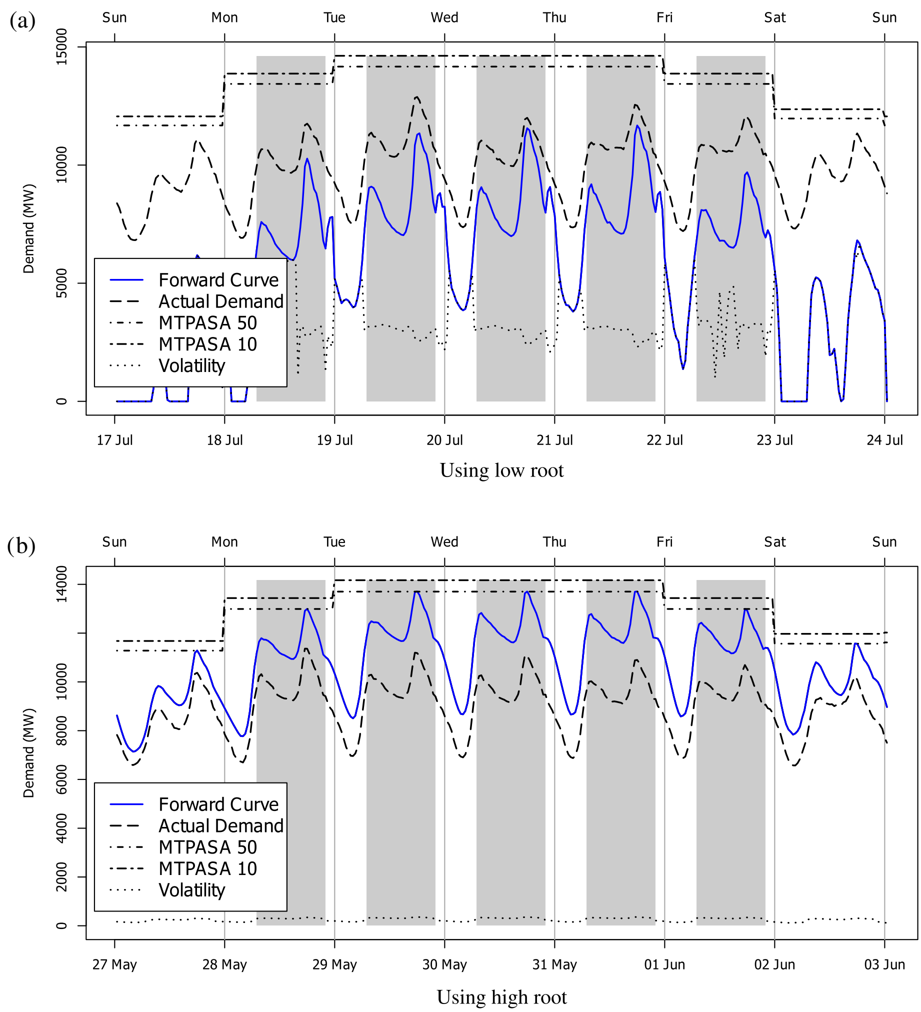

FIGURE 2. Comparison of low and high roots over time (colour available online).

propose the hypothesis that the probability of exceeding this market cap, given the market calibrated quantiles we are using, is negligible. By making this assumption, we can remove the truncation of the lognormal distribution and, therefore, the normalization factor Z, which in turn removes the need to solve for Z and, thus, greatly simplifies the model.

We rewrite the density function in (4.1) as a standard lognormal

$$f(\sigma, \mu; x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-(\ln x-\mu)^2/2\sigma^2}, \quad x > 0,$$

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

and, from (4.2),

$$\mu = \ln y_1. \tag{5.1}$$

Next, using equation (5.1), we get the cumulative distribution function

$$\begin{aligned} F(y) &= \frac{1}{2} + \frac{1}{2}\,\mathrm{erf}\Big(\frac{\ln y - \ln \mu}{\sigma\sqrt{2}}\Big) \\ &= \frac{1}{2} + \frac{1}{2}\,\mathrm{erf}\Big(\frac{\ln y - \ln y_1}{\sigma\sqrt{2}}\Big). \end{aligned} \tag{5.2}$$

Now equating (5.2) with y_2 yields

$$\begin{aligned} F(y) &= 0.9 \\ &\Rightarrow \frac{1}{2} + \frac{1}{2}\,\mathrm{erf}\Big[\frac{\ln y - \ln y_1}{\sigma\sqrt{2}}\Big] = 0.9 \\ &\Rightarrow \mathrm{erf}\Big[\frac{\ln(y_2/y_1)}{\sigma\sqrt{2}}\Big] = 0.8 \\ &\Rightarrow \sigma = \frac{\ln(y_2/y_1)}{\mathrm{erf}^{-1}(0.8)\sqrt{2}}. \end{aligned} \tag{5.3}$$

The simplified forward load curve (SFLC) is the first moment of the PDF; therefore,

$$\mathrm{SFLC} = e^{\mu+\sigma^2/2} = e^{\ln y_1} \times e^{\sigma^2/2} = y_1 e^{\sigma^2/2}.$$

Keeping in mind (3.1) and the need to solve for all time periods t_i,

$$\mathrm{SFLC}(t_i) = L_{\mathrm{MAX}}\, y_1(t_i) e^{(1/2)[\ln\{y_2(t_i)/y_1(t_i)\}/\,\mathrm{erf}^{-1}(0.8)\sqrt{2}]^2}. \tag{5.4}$$

6. Results – FLC

Figures 3 and 4 show the resultant expected curves against the original 50% and 10% MTPASA POEs. Note that actual data is shown as indicative only. First, a forecast is not necessarily expected to predict actual demand, rather to give an expected demand, and ideally also a variance measure. Secondly, the published MTPASA levels that we have calibrated to appear to be systematically overstating demand – this is caused by the input data, not the expected curve methodology.

Figures 5–12 show sections of the forward curve on a weekly basis for selected weeks, and for regions in New South Wales (NSW) and Victoria (VIC).

Based on the methodology, our criterion for a curve of the correct magnitude is that the maximum daily demand for any day in any week exactly matches the MTPASA 50% POE line for that day. This is because the MTPASA 50% represents the expected value (50% POE) of the maximum daily demand for each day forecast. The shape of the demand profile in any day within any week is shaped by the rolling STPASA short-term forecast, and will differ for each day.

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

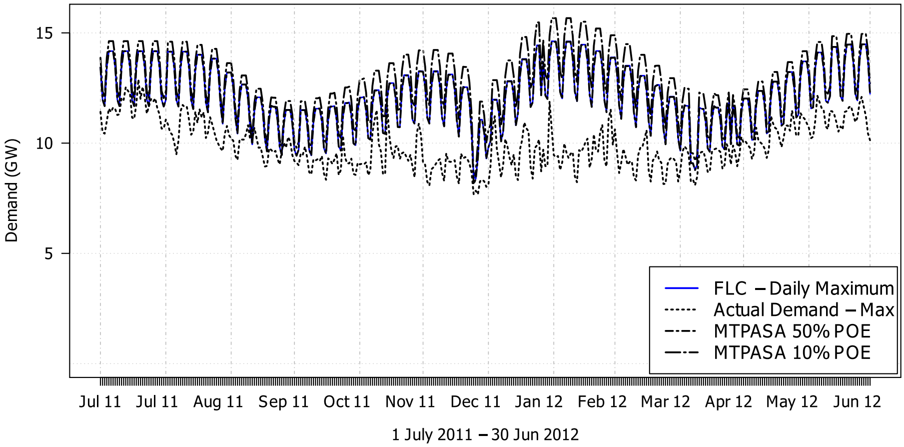

FIGURE 3. NSW backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

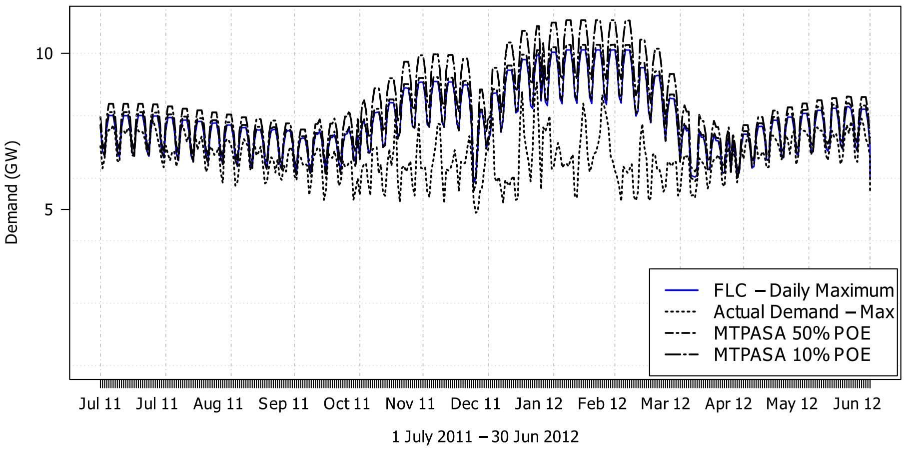

FIGURE 4. VIC backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

7. Results – SFLC

The simplified forward load curve (5.4) produces a virtually identical curve to the FLC ($R^2 = 0.999\,999\,999\,987$ for the test period of 1 year). Similarly, if we compare volatility curves given by equations (4.4) and (5.3), we get identical curves ($R^2 = 0.999\,999\,999\,998\,97$).

This result supports the hypothesis that the truncated portion of the density function is negligible. Given that the SFLC results in a far simpler representation, avoids the need for iterative root solving and always solves for a value, it appears to be a better candidate for the load curve modelling.

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

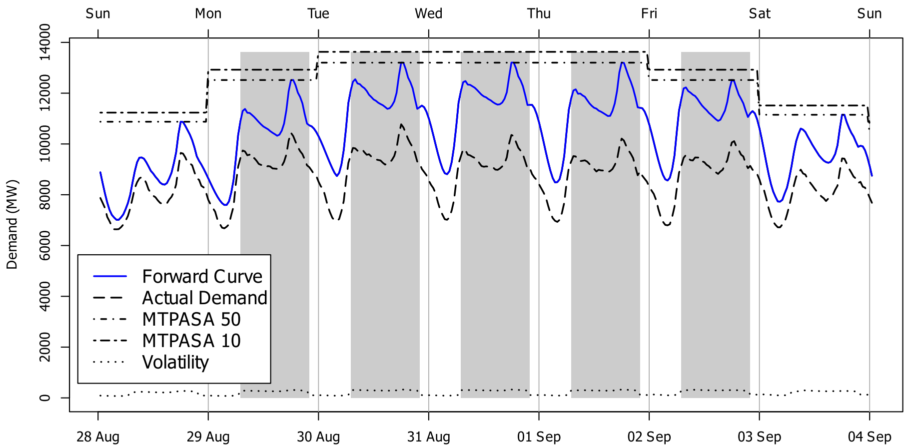

FIGURE 5. NSW backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

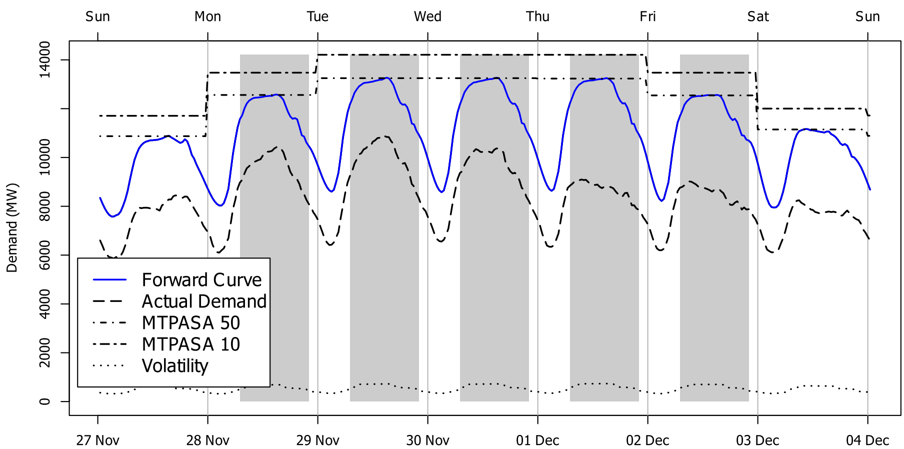

FIGURE 6. NSW backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

8. Conclusions

We have demonstrated a way of reconstituting multiple distributions based on two vectors of quantiles. For every time period, the solution satisfies the definition of the measures. We have made two observations while working with the PASA forecasts.

(1) The values of the 10% and 50% POEs produced by AEMO contradict AEMO's assertion that the load distributions per time period are normal, as the POEs demonstrate asymmetry.
(2) When backtested against the observed load, MTPASA appears to regularly overstate the load, to the point that the 50% POE is more like a 5% or 6% POE. The performance of MTPASA as a forecast is outside the scope of this

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

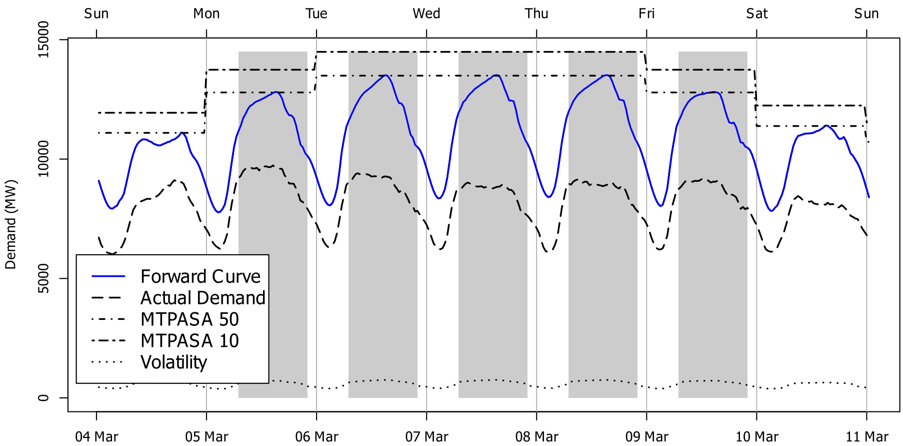

FIGURE 7. NSW backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

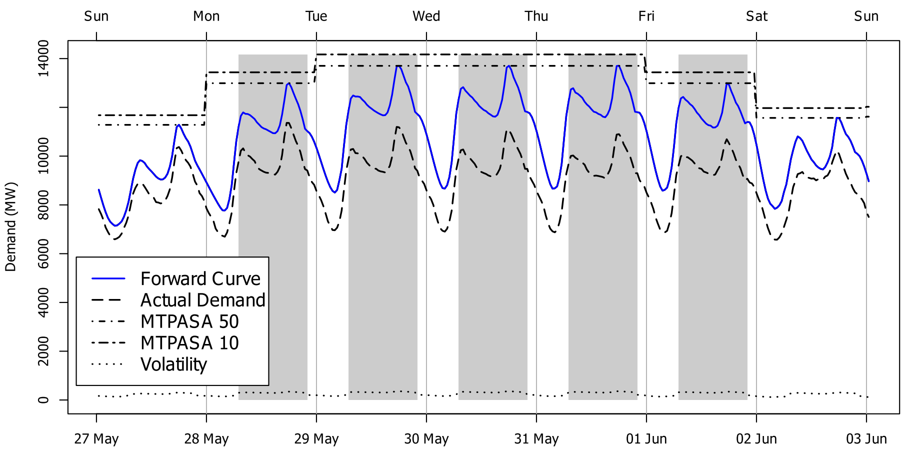

FIGURE 8. NSW backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

paper, and it does not affect the methodology described here, since the aim of the methodology is to effectively "shape" MTPASA for compliance reporting purposes. Whether MTPASA is a reliable forecast of actual demand is a matter for AEMO's forecasting process.

The MTPASA forecasts are based on economical and seasonal indicators, mainly considering estimates from the transmission network service providers (TNSPs), and it is only quoted on a daily basis. This methodology adds two types of information to the forecast.

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

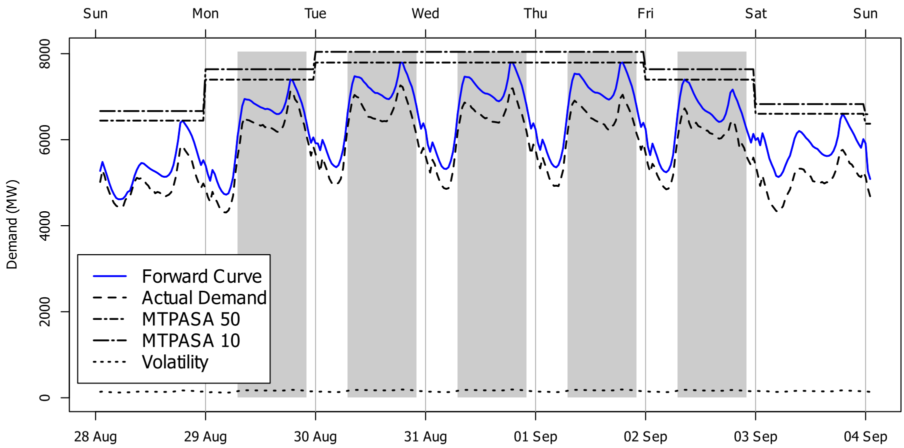

FIGURE 9. VIC backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

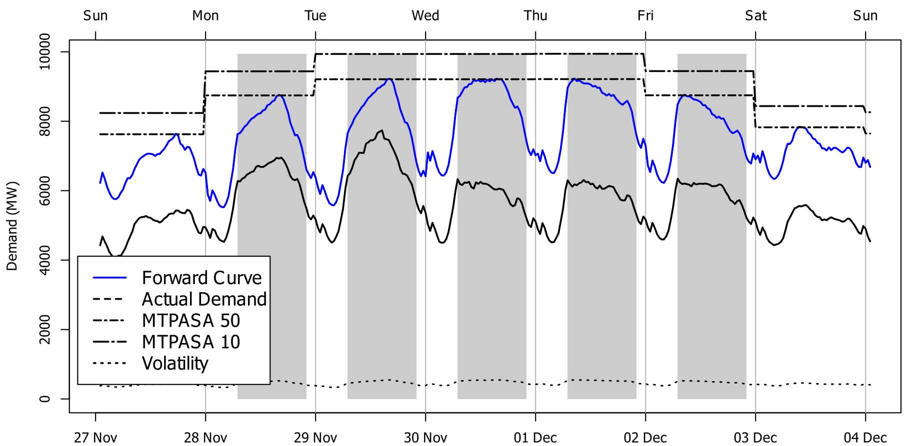

FIGURE 10. VIC backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

(1) The shaping process increases the frequency from daily to half-hourly.
(2) The resultant distributions support probabilistic risk measures, such as earnings at risk (EaR) or insurance-based hedging agreements based on load triggers. EaR is possible when making the assumption that the portfolio in question exhibits the same volatility profile as state load.

In this paper, we have tested two models, the FLC, which is based on the fact that load (demand) is confined within 0 and L_{MAX}, while the second model, the SFLC, assumes that L_{MAX} is high enough so that the probability of exceeding it is negligible. Since the two models produce near-identical results, the SFLC model appears to be a highly accurate approximation for the FLC. Given that the simpler model offers

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press

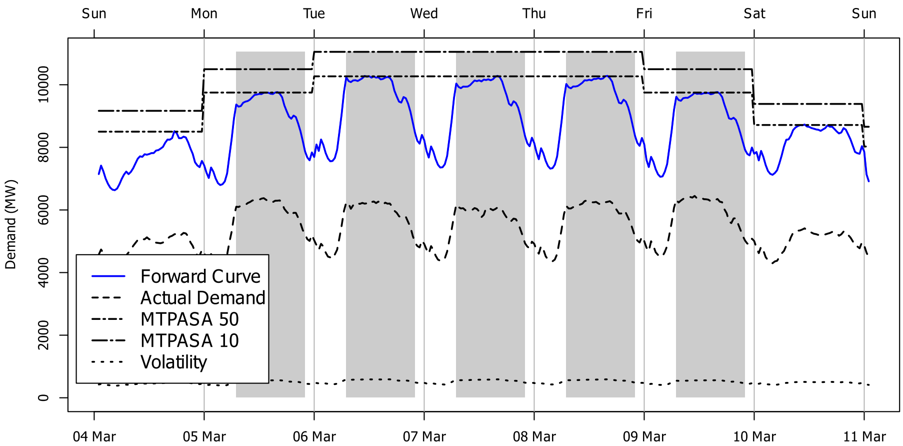

FIGURE 11. VIC backtest, FY 2012 – MTPASA generated 21 June 2011 (colour available online).

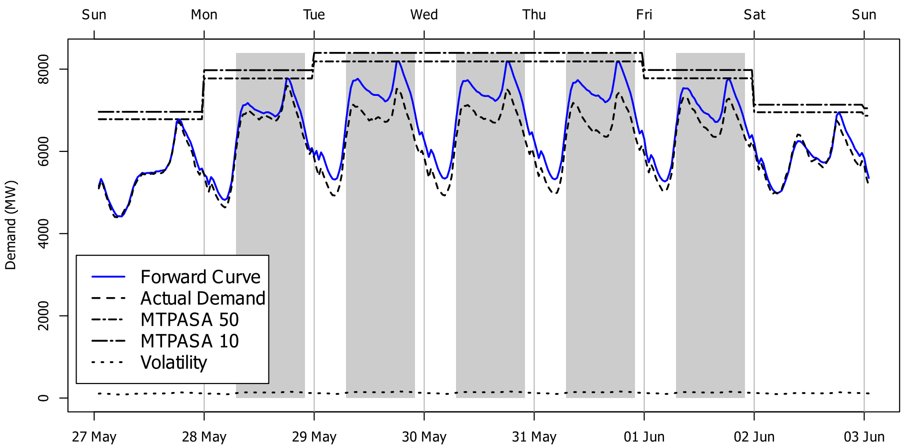

FIGURE 12. VIC backtest, FY 2012 – MTPASA generated on 21 June 2011 (colour available online).

significant performance advantages, it would be the preferable load curve model for practical applications.

Acknowledgements

We would like to thank Ben Vanderwaal from Roam Consulting for helpful feedback on the draft of this paper and a fruitful discussion on the MTPASA process. Timothy Ling acknowledges the support of the Australian Research Council's Discovery Projects funding scheme (DP130103315).

References

[1] M. Abramowitz and I. Stegun, *Handbook of mathematical functions with formulas, graphs, and mathematical tables* (Dover Publications, New York, 1972); http://people.math.sfu.ca/~cbm/aands/intro.htm.

[2] AEMO, "Short term PASA process description", Technical Report, AEMO Electricity Market Performance, ver. 006 (2012); http://www.aemo.com.au/Electricity/Market-Operations/Dispatch/STPASA-Process-Description.

[3] V. Witkovsky, "Computing the distribution of a linear combination of inverted gamma variables", *Kybernetika* **37** (2001) 79–90; http://www.researchgate.net/publication/236323890_Computing_the_distribution_of_a_linear_combination_of_inverted_gamma_variables.

https://doi.org/10.1017/S1446181115000322 Published online by Cambridge University Press