

Decoding the Australian Electricity Market:
New Evidence from Three-Regime Hidden Semi-Markov Model

Nicholas Apergis, Ph.D.
University of Piraeus, Greece
E-mail: napergis@unipi.gr

Giray Gozgor, Ph.D.
Istanbul Medeniyet University, Turkey
E-mail: giray.gozgor@medeniyet.edu.tr

Chi Keung Marco Lau, Ph.D.
University of Huddersfield, the United Kingdom
E-mail: c.lau@hud.ac.uk

Shixuan Wang, Ph.D.
(Corresponding Author)
University of Reading, the United Kingdom
E-mail: shixuan.wang@reading.ac.uk

Decoding the Australian Electricity Market: New Evidence from the Three-Regime Hidden Semi-Markov Model

Abstract

The hidden semi-Markov model (HSMM) is more flexible than the hidden Markov model (HMM). As an extension of the HMM, the sojourn time distribution in the HSMM can be explicitly specified by any distribution, either nonparametric or parametric, facilitating the modelling for the stylised features of electricity prices, such as the short-lived spike and the time-varying mean. By using a three-regime HSMM, this paper investigates the hidden regimes in five Australian States (Queensland, New South Wales, Victoria, South Australia, and Tasmania), spanning the period from June 8, 2008 to July 3, 2016. Based on the estimation results, we find evidence that the three hidden regimes correspond to a low-price regime, a high-price regime, and a spike regime. Running the decoding algorithm, the analysis systemically finds the timing of the three regimes, and thus, we link the empirical results to the policy changes in the Australian National Electricity Market. We further discuss the contributing factors for the different characteristics of the Australian electricity markets at the state-level.

Keywords: Australian electricity markets; hidden semi-Markov model; time-varying mean; price spikes; regime-switching.

JEL Classification Codes: Q41; Q48; L97; C13

1. Introduction

Unlike financial assets or other commodities, electricity has its own special features since it cannot be stored economically and requires immediate delivery. In a deregulated electricity market, the price can shoot up to an extremely high level and come back to normal level within a short period of time. These abrupt, short-lived and extreme fluctuations in electricity prices are referred to as spikes in the literature (Higgs and Worthington, 2008; Janczura and Weron, 2010; Mount et al., 2006). As pointed out by Mount et al. (2006), the price spike is the standard feature of the Australian electricity market since it uses a uniform price auction and imposes fewer restrictions on the supplier. It is worthwhile to note that the spikes in the electricity price are fundamentally different from the jumps in stock prices as “spikes are temporal level shifts that die out rather quickly and do not lead to sustainable higher price levels” (Huisman and Mahieu, 2003: 428). Therefore, the stochastic jump process used in the literature to model the sudden level shifts in stock prices cannot explicitly model the spikes in the electricity price.

Numerous attempts have been made to model the behaviour of the spikes in the electricity price. Among them, the hidden Markov model (HMM), a.k.a. the Markov regime-switching model, is one of natural candidates because of its elegance in accommodating the different regimes involved in the data generation process. Various versions of the HMM have been developed to compete with the goodness-of-fitting for the empirical data. Huisman and Mahieu (2003) model the electricity price by a three-regime HMM with constant transition probabilities. Their three regimes are defined as a normal regime, an initial jump regime, and a regime reverting back to a normal regime after the initial jump regime. Mount et al. (2006) employ a two-regime HMM of time-varying transition probability depending on the current market conditions. They show that the price spikes can be predicted accurately if the reserve margin and the system load are provided. Janczura and Weron (2010) conduct an empirical

comparison of alternate regime-switching models for electricity spot prices. They suggest that the best structure of the regime-switching models should set to be an independent three-regime model with time-varying transition probabilities, heteroscedastic diffusion-type base regime dynamics and shifted spike regime distributions.

Nevertheless, one limitation of the HMM is that the sojourn time distribution can only follow an exponential distribution implicitly, which sometimes does not fit the empirical data well, particularly for the spikes which typically last for a short period of time. The hidden semi-Markov model (HSMM) is more flexible than the traditional HMM because the sojourn time distribution in the HSMM can be explicitly specified by any distribution, either parametric or nonparametric, allowing us to model the very short sojourn time of the spikes. This is the novelty of the paper in relevance to the price spikes in the electricity markets.


In addition to the spikes, we also observe that the mean of the electricity prices is time-varying. As an example, the log electricity price in the New South Wales (NSW) market is plotted in Figure 1, showing that the mean level from 2012 to 2014 is significantly higher than that in other periods. To this end, we take consideration of the stylised features of electricity prices in Australia, namely spikes and the time-varying mean, and employ a three-regime HSMM to investigate the hidden regimes in the Australian National Electricity Market (NEM), spanning the period from June 8, 2008, to July 3, 2016. We find evidence that the first regime is the low-price regime, the second regime is the high-price regime, and the third regime is the spike regime.

Running the Viterbi (1967)'s algorithm, the analysis systemically finds the timing of the three regimes, and thus, we link the empirical results to the policy changes in the Australian electricity market. We discuss the policy implications and the contributing factors for the different characteristics of the Australian electricity markets at the state-level. Our findings are important from the policy perspective, since environmental policy reforms can

possibly provide policy uncertainty in relation to the climate change policy. For example, the persistence of the high-price regime in the period from 2012 to 2014 coincides with the environmental policy reforms adopted by the Australian government. The new energy policy can consider the development of a renewed reform agenda for the Energy Council, in a sense that addresses all these issues with a focus on outcomes, rather than being mired in process, as it has been the case so far.

The rest of the paper is organised as follows. Section 2 briefly reviews the previous literature. Section 3 explains the features of the data and illustrates the summary of descriptive statistics. Section 4 introduces the econometric methodology, while Section 5 provides the empirical results. Section 6 discusses the potential implications of the empirical findings and finally, Section 7 concludes the paper.

Figure 1. New South Wales – Electricity Price (at the Log-Level)



2. Brief Literature Review

There are a limited number of studies that investigate the market behaviour and the price dynamics (between and within) of Australia's NEM. For example, Hu et al. (2005) review the strategic bidding and the rebidding behaviour of Australia's NEM, spanning the period from

May 1, 2002 to May 31, 2003. The authors observe that large generators are more likely to use the capacity quotations rather than the price quotations to control market prices. The authors also show that larger generators do not suppress market prices during the peak periods by exploiting market rules; this especially leads to an inefficient market problem during the peak periods. Higgs (2009) also empirically analyses the degree of electricity market integration across the states of New South Wales, Queensland, South Australia, and Victoria for the period from January 1, 1999, to December 31, 2007. The empirical findings from the constant conditional correlation (CCC) model indicate that two pairs of the electricity market (New South Wales–Queensland and New South Wales–Victoria) are significantly interlinked at the higher degree when those are compared to the remaining pairs. The author also argues that the introduction of the Queensland and New South Wales Interconnector (QNI) on February 18, 2001 constitutes the main reason for the interconnections between the related pairs. However, Higgs (2009) concludes that despite the related findings for the interconnections between the related pairs, Australia’s NEM can be characterised by the limited inter-linkages across the states, meaning that Australia’s NEM at the regional level is relatively isolated for the period from January 1, 1999 to December 31, 2007.

In addition, Cutler et al. (2011) consider the 30-mins spot electricity price data in South Australia over the period September 2008–August 2010. They search for the potential implications of an expanded wind generation in South Australia and observe that the demand for electricity is the crucial factor that determines the electricity price (rather than supply). The wind power also significantly affects electricity prices since the authors find a negative correlation between the wind generation and the spot electricity prices. In addition, their findings illustrate that the wind power outputs can be related to the extreme prices, but it is difficult to provide a causal relationship between the spot electricity prices and the wind

generation in South Australia. In other words, finding a correlation does not indicate that there is also a significant causality between the related variables. These findings document the potential relationship between the demand and the spot prices at the regional level.

It is also important to note that the recent literature has extensively analysed the volatility spillovers across the energy markets. A strand of this literature investigates the volatility spillovers between energy (mainly oil) and stock markets (Arouri et al., 2012), while another strand focuses on the spillovers between energy and the commodity markets (mainly agricultural commodity and precious metals markets) (Du et al., 2011; Ewing and Malik, 2013; Gozgor and Memis, 2015). There is also a technical part of the literature that explicitly investigates whether the econometric methodologies that have been employed to analyse volatility spillovers are robust to the frequency of the data (Yarovaya et al., 2016), to the different time horizons (Gozgor et al., 2016) (e.g. splitting the data as the pre- and the post-crisis period), and to the presence of asymmetric effects across different (energy) markets (Barunik et al., 2015; Chang et al., 2010).

Different from these studies, investigating whether there is a systematic price volatility transmission across the states (regions) in Australia's NEM is a relatively new research task. For instance, in their recent empirical analysis, Apergis et al. (2017b) show that there is a statistically significant convergence in the wholesale electricity prices across three Australian states: New South Wales, Queensland, and Victoria. This finding indicates that there are common market characteristics across these three states, which is associated with the carbon tax regimes in the related states. A more significant factor would be a greater degree of interconnectedness via greater transfer capacity on interstate interconnectors, linking QLD, NSW and VIC relative to the interconnectors linking TAS and SA to VIC. In particular, the lack of congestion on transmission branches between different markets would lead to a wholesale price convergence across these markets. Furthermore, Apergis et al.

(2017a) analyse the price volatility spillovers across four Australian states (i.e. New South Wales, Queensland, South Australia, and Victoria) by using the intraday 5-minutes electricity price data, spanning the period December 8, 1998, to May 5, 2016. The authors also consider the presence of the asymmetric characteristics in terms of the price volatility, known as the ‘good volatility’ and ‘bad volatility’. Their findings highlight that Australia’s NEM is asymmetrically interconnected with the presence of the different degrees.

As for modelling the spike behaviour in the electricity price, HMM is one of the widely used models due to its elegance in accommodating different regimes mixing in the data generation process. Numerous versions of the HMM have been developed to compete with the goodness-of-fitting. One of the early studies is conducted by Deng (2000) who considers a two-regime HMM with the setting of one base regime driven by an AR(1) process and another spike regime driven by the same AR(1) process shifted by an exponential distributed random variable. Huisman and De Jong (2003) propose to use a simple independent spike two-regime model for deseasonalised log-prices. Their base regime is also driven by an AR(1) process, but the spike regime in their setting is a normal distributed random variable whose mean and variance were higher than those of the base regime. In a follow-up study, De Jong (2006) modified the basic independent spike two-regime model with the Poisson-driven spike regime. The results show that the previous model with Gaussian spikes is nearly as good as the new model with Poisson spikes. In an alternative setting, Husiman and Mahieu (2003) employ an HMM with three regimes, which contain a normal regime, an initial jump regime, and a regime reverting back to a normal regime after the initial jump regime. This setting is implicitly set the restriction of the so-called ‘one day spike’ since the reverting regime has to be immediately followed by the jump regime.

Diebold et al. (1994) argue that it is restrictive to treat transition probabilities as fixed parameters. A loose assumption could be that the transition probabilities are time-varying and

depend on exogenous variables. Mount et al. (2006) associate the transition probability in a two-regime HSMM with current market conditions and show that the price spikes can be predicted accurately if the reserve margin and the system load are provided. Janczura and Weron (2010) empirically compare a range of regime-switching models for electricity spot prices in terms of the performance in goodness-of-fitting and in reproducing the special features of electricity prices. The suggested best model should set to be an independent three-regime model with time-varying transition probabilities, heteroscedastic diffusion-type base regime dynamics and shifted spike regime distributions.

As a matter of fact, the sojourn time in the HMM can only follow an exponential distribution, which is too restrictive for some empirical data. The HSMM is an extension of HMM by explicitly specifying the sojourn time distribution, which facilitates a better fitting for the short-lived spikes. It is noteworthy to indicate that the HSMM was initially developed in 1980 for the speech recognition. Since then, it has been applied in many scientific areas, including handwriting recognition, human activity recognition, functional magnetic resonance imaging (MRI) brain mapping (Yu, 2010). In terms of the applications in the economics and the finance literature, Bulla and Bulla (2006) employ the two-regime HSMM¹ to model 18 series of the United States daily sector returns and they found that the HSMM outperform the hidden Markov models (HMM) in terms of the better description about the stylised facts of daily returns. By using two-regime HSMM and three-regime HSMM, Lau et al. (2017) provide some insights into the characteristics of white precious metal markets. Liu and Wang (2017a) use the three-regime HSMM to decode the Chinese stock market returns. The authors observe that the time-varying distribution of the Chinese stock market returns depends on the hidden regimes, which turns out to be three market conditions, namely the bear, the bull, and the sidewalk markets. Liu and Wang (2017b) further use the same method

¹ The two-regime HSMM is also referred to as the two-state HSMM. Similarly, the three-regime HSMM is also known as the three-state HSMM. In order to avoid the confusion with states (regions) in Australia, we choose to use the term “regimes” rather “states”.

and compare the market conditions across eight countries. However, this robust method, HSMM, has still not been used by the papers in the energy field. To the best of our knowledge, our paper provides the first analysis to use the HSMM in the energy economics and the energy finance literature.

To conclude the literature review, we observe that there is still a gap in the literature for using HSMM to understand the spikes and the time-varying mean in Australia's NEM at the state (regional) level. Considering that the weekly data span the period from June 8, 2008, to July 3, 2016, and using the HSMM modelling framework, our paper examines the timing and transitions of hidden regimes in the Australian NEM, and further discuss the contributing factors and policy implications.

3. Data and Descriptive Statistics

The paper employs weekly wholesale electricity prices, spanning the period June 8th, 2008 to July 3th, 2016 as the earliest data available for the weekly volume weighted average spot prices is on June 8th, 2008². The sample consists of five Australian regions, i.e. Queensland (QLD), New South Wales (NSW), Victoria (VIC), South Australia (SA), and Tasmania (TAS). In line with the literature, the variable of interest is the logarithm of the electricity prices that can be written as:

$$p_t = \ln(P_t) \tag{1}$$

Where $\ln(P_t)$ is the natural logarithm of the electricity price at time t . There are 422 active observations available in total. At this point, Table 1 presents the summary statistics of the electricity prices³ across the five electricity markets, which contains the sample means,

² The related data can be downloaded from <http://www.aer.gov.au/wholesale-markets/wholesale-statistics/weekly-volume-weighted-average-spot-prices>.

³ We present the descriptive statistics of the original prices, rather than the log price. This is because they are more straightforward for the reader.

medians, maximums, minimums, standard deviations, skewness, kurtosis, and the Jarque-Bera (JB) statistic.

Table 1. Summary Statistics (NEM Regions)

	QLD	NSW	VIC	SA	TAS
Mean	48.36	45.51	42.65	57.41	48.49
Median	36.00	36.00	34.00	40.00	38.50
Maximum	420.00	627.00	619.00	693.00	405.00
Minimum	14.00	20.00	15.00	6.00	0.00
Standard Deviation	43.84	40.99	40.04	75.94	42.90
Kurtosis	35.11	107.93	110.68	40.87	24.52
Skewness	5.02	8.78	8.79	5.82	4.23
Jarque–Bera Test	19,903	199,011	209,299	27,604	9,397
P-value	0.00	0.00	0.00	0.00	0.00

Notes: Weekly volume weighted average prices (\$/MWh). All data range from 08/06/2008 to 03/07/2016.

The weekly volume weighted average spot prices for the five markets ranges from \$42.65/MWh (TAS) to \$57.41/MWh (SA). The standard deviations of the spot electricity prices range from \$40.04 (VIC) to \$75.94 (SA). Therefore, South Australia is the most volatile market, while Victoria is the most stable market. Table 1 shows clear evidence that the unconditional distributions depart from normality (as implied by the significant JB test statistics and p-values).

Another observation from Table 1 is that the maximum value of the price can be considerably higher than the mean and the median. This is caused by the spikes in the price. Due to the presence of spikes, the distribution of prices has a large positive skewness and substantially high kurtosis. These issues further indicate that the distribution of the price is positively skewed with right tail outliers, which motivates us to use the HSMM method. In the next section, we explain the details of the HSMM.

4. Econometric Methodology: The Hidden Semi-Markov Model (HSMM)

The sojourn time in the HMM implicitly follows a geometric distribution, which is restrictive for many empirical applications. For example, it is unlikely that such a type of implicit

sojourn time distribution has been always appropriate for empirical data, such as speech segment duration, the length of the segments of the DNA sequences, and the length of branching zones in plants (Gu édon, 2003). Such a type of implicit sojourn time distribution is also unlikely to be suitable for the time lengths of the regimes in electricity prices, particularly for the short-lived spikes which typically last for a short period of time. The HSMM generalises the HMM by explicitly specifying the sojourn time distribution. After being initially developed for the speech recognition in 1980, the HSMM has been successfully applied in various scientific areas, including handwriting recognition, human activity recognition, functional magnetic resonance imaging (MRI) brain mapping, and network anomaly detection (Yu, 2010). For the general description of semi-Markov models, please refer to the book of Kulkarni (2016).

The HSMM is constituted by a pair of two stochastic processes, namely an output process (a.k.a. observation process) $\{X_t\}$ and a regime process (a.k.a. state process) $\{S_t\}$. The regime process $\{S_t\}$ is a semi-Markov chain with finite regime space $\{1, \dots, J\}$, and it is unobservable but can be decoded from the observable regime process $\{X_t\}$. A semi-Markov chain⁴ is an embedded first-order Markov chain, which is constructed by a first-order Markov chain and sojourn time distributions (a.k.a. occupancy distributions, duration distributions). Firstly, the transition probabilities are defined in Equation (2):

$$\gamma_{i,j} = P(S_{t+1} = j | S_{t+1} \neq i, S_t = i) \text{ with } \sum_{j \neq i} \gamma_{i,j} = 1 \text{ and } \gamma_{i,i} = 0 \quad (2)$$

The transition probability matrix (TPM) can be arranged as a $J \times J$ matrix with $\gamma_{i,j}$ in row i and column j . The diagonal entries in the TPM are zeros in the HSMM because the self-transition (transition from i to i) probabilities are zero.

More importantly, a sojourn distribution is associated to each regime in the embedded first-order Markov chain

⁴ We only consider non-absorbing regimes and exclude absorbing regimes.

$$d_j(u) = P(S_{t+u+1} \neq j, S_{t+u-v} = j, v = 0, \dots, u - 2 | S_{t+1} = j, S_t \neq j) \quad (3)$$

Due to this ‘bonding’, the sojourn time in an HSMM can explicitly follow any arbitrary distributions, either nonparametric or parametric. It should be noted that the sojourn time in an HMM can only follow the geometric distribution, which sometimes cannot fit the data well, especially for the very short sojourn time in the spike regime of the AEM. The release of assumption on the sojourn time gives the flexibility of an HSMM to capture the complexity in the sojourn time distribution of real data.

There is a complication with the last visited regime. Since the assumption that the last observation always coincides with the exit from a regime is unrealistic, it is necessary to consider the right-censoring setting and use the survival function of the sojourn time in regime j

$$D_j(u) = \sum_{v \geq u} d_j(v) \quad (4)$$

In the semi-Markov chain, the Markovian property is converted at the level of the first-order embedded Markov chain, rather than at each time point in a Markov chain.

The output process $\{X_t\}$ generates observations by conditional distributions (a.k.a. component distributions, emission distributions) depending on the current regime at time t

$$b_j(y) = P(X_t = y | S_t = j) \text{ with } \sum_y b_j(y) = 1 \quad (5)$$

The way of defining output process indicates the assumption that the observations at time t only depend on the underlying semi-Markov chain at time t .

For the sake of convenience, we need to introduce some notations. The upper-case notations X_t and S_t stand for the variables at time t , and the lower-case notations x_t and s_t represent the realised values at time t . The observation sequence from time 0 to time T , $\{X_0 = x_0, \dots, X_T = x_T\}$, is abbreviated as \mathbf{X}_0^T . Similarly, the regime sequence from time 0 to time T , $\{S_0 = s_0, \dots, S_T = s_T\}$, is abbreviated as \mathbf{S}_0^T .

4.2. Estimation Method: EM Algorithm

Unlike the HMM, the likelihood of the HSMM cannot be represented conveniently in a tractable form. Gu élon (2003) provided the *complete-data likelihood*, i.e. the observations \mathbf{x}_1^T as well as the unobserved sequence \mathbf{s}_1^{T+u} , for the right-censored HSMM, which is shown in Equation (6):

$$L(\theta) = \sum_{s_1, \dots, s_T} \sum_{u_{T+}} L_c(\mathbf{s}_1^{T+u}, \mathbf{x}_1^T | \theta) \quad (6)$$

$$L_c(\mathbf{s}_1^{T+u}, \mathbf{x}_1^T | \theta) = P(\mathbf{X}_1^T = \mathbf{x}_1^T, \mathbf{S}_1^T = \mathbf{s}_1^T, \mathbf{S}_{T+1}^{T+u-1} = s_T, S_{T+u} \neq s_T, | \theta)$$

where θ stands for the complete set of parameters, \sum_{s_1, \dots, s_T} denotes the summation over every possible regime sequence, and $\sum_{u_{T+}}$ represents the sum over every supplementary sojourn time after T . As shown in Equation (6), the completed regime sequence stops at $T + u$, rather than T without right-censoring (Gu élon, 2003).

However, the likelihood function is difficult to be calculated in Equation (6) as it needs to consider all possibilities of the regime sequence, which is unrealistic. Following Bulla and Bulla (2006), we estimate the parameters in the HSMM by using the Expectation-Maximisation (EM) algorithm, which is an iterative two-step procedure to increase the likelihood monotonically until the stopping criterion is satisfied. Here are the main steps of the algorithm:

- 1) Give an initial guess $\theta^{(0)}$ of the parameter set.
- 2) **E-step**: compute the Q -function, which is the conditional expectation for the complete-data likelihood.

$$Q(\theta, \theta^{(k)}) = E\{L_c(\mathbf{s}_1^{T+u}, \mathbf{x}_1^T | \theta) | \mathbf{X}_1^T = \mathbf{x}_1^T, \theta^{(k)}\}$$

where $\theta^{(k)}$ is the k^{th} iteration of the estimate of the parameter set θ .

- 3) **M-step**: update the $\theta^{(k+1)}$, which is to maximise the Q -function with respect to θ .

$$\theta^{(k+1)} = \arg \max_{\theta} Q(\theta, \theta^{(k)})$$

4) Check whether the stopping criterion is met. If not, go back to the E-step.

In the E-step, the aim is to calculate the *Q-function* given that the parameter set $\theta^{(k)}$ is calculated in the k^{th} iteration of the M-step. In the M-step, the aim is to update the estimated parameter set $\theta^{(k+1)}$ given that the *Q-function* is calculated in the E-step. By the iterative procedure of two steps, the likelihood increases monotonically until the stopping criterion is met (Baum et al., 1970; Dempster et al., 1977). The stopping criterion is that the successive change of likelihood is less than an arbitrarily small number. Note that the EM algorithm is not guaranteed to reach the global maxima, and the limit of this sequence of estimates, i.e. $\{\theta^{(0)}, \theta^{(1)}, \dots\}$, reaches a local maximum of the likelihood function. Interested readers can refer to Bulla and Bulla (2006) for the mathematical details of the EM algorithm for the HSMM. They provided the convenient expression of the *Q-function* in the E-step by separating it into four components, namely the initial probability, the transition probability, the sojourn time distribution, and the observation probability. By this separation, the four components can be maximised individually in the M-step.

4.3. Decoding Method: Viterbi Algorithm

After estimating the parameters in the HSMM, the next step is to decode the most likely sequence of regimes in the semi-Markov chain. There are two popular decoding algorithms (Zucchini et al., 2016), namely the global decoding and the local decoding algorithms, which are based on different rules. The global decoding determines the most likely sequence of regimes for the whole sample period simultaneously, whilst the local decoding determines the most likely regime for each time point individually. In this paper, we focus on the global decoding. The problem is formulated by Equation (7):

$$\begin{aligned}
& \max_{s_1, \dots, s_T} P(\mathbf{S}_1^T = \mathbf{s}_1^T, \mathbf{X}_1^T = \mathbf{X}_1^T) \\
&= \max_{s_t} \left\{ \max_{s_{t+1}, \dots, s_T} P(\mathbf{S}_{t+1}^T = \mathbf{s}_{t+1}^T, \mathbf{X}_{t+1}^T = \mathbf{X}_{t+1}^T \mid S_{t+1} \neq s_t, S_t = s_t) \right. \\
&\quad \left. \times \max_{s_1, \dots, s_{t-1}} P(S_{t+1} \neq s_t, \mathbf{S}_0^t = \mathbf{s}_0^t, \mathbf{X}_0^t = \mathbf{x}_0^t) \right\} \quad (7)
\end{aligned}$$

The exhaustive attack method is only feasible for short sequence data. The Viterbi algorithm (Guédon, 2003), at the computational complexity level $O(JT(J + T))$ in time for the worst case and $O(JT)$ in space, is an efficient dynamic programming algorithm to maximisation problems by taking advantage of the Markov property of the HSMM. Let us define

$$\alpha_j(t) = \max_{s_1, \dots, s_{t-1}} P(S_{t+1} \neq j, S_t = j, \mathbf{S}_0^{t-1} = \mathbf{s}_0^{t-1}, \mathbf{X}_0^t = \mathbf{x}_0^t) \quad (8)$$

Then, Equation (8) can be written as

$$\begin{aligned}
& \max_{s_1, \dots, s_T} P(\mathbf{S}_1^T = \mathbf{s}_1^T, \mathbf{X}_1^T = \mathbf{X}_1^T) \\
&= \max_j \left\{ \max_{s_{t+1}, \dots, s_T} P(\mathbf{S}_{t+1}^T = \mathbf{s}_{t+1}^T, \mathbf{X}_{t+1}^T = \mathbf{X}_{t+1}^T \mid S_{t+1} \neq s_t, S_t = s_t) \alpha_j(t) \right\} \quad (9)
\end{aligned}$$

We can construct the following recursion for $t = 1, \dots, T - 1; j = 1, \dots, J$:

$$\begin{aligned}
\alpha_j(t) &= \max_{s_1, \dots, s_{t-1}} P(S_{t+1} \neq j, S_t = j, \mathbf{S}_0^{t-1} = \mathbf{s}_0^{t-1}, \mathbf{X}_0^t = \mathbf{x}_0^t) \\
&= b_j(x_t) \max \left[\left[\max_{1 \leq u \leq t} \left\{ \prod_{v=1}^{u-1} b_j(x_{t-v}) \right\} d_j(u) \max_{i \neq j} \{p_{i,j} \alpha_i(t-u)\} \right], \left\{ \prod_{v=1}^t b_j(x_{t-v}) d_j(t \right. \right. \\
&\quad \left. \left. + 1) \pi_j \right\} \right] \quad (10)
\end{aligned}$$

where π_j is the initial probability for regime j . The right-censoring of the sojourn time in the last visited regime is the special case that $t = T; j = 1, \dots, J$:

$$\begin{aligned}
\alpha_j(T) &= \max_{s_1, \dots, s_{T-1}} P(S_T = j, \mathbf{S}_0^{T-1} = \mathbf{s}_0^{T-1}, \mathbf{X}_0^T = \mathbf{x}_0^T) \\
&= b_j(x_T) \max \left[\left[\max_{1 \leq u \leq T} \left\{ \prod_{v=1}^{u-1} b_j(x_{T-v}) \right\} D_j(u) \max_{i \neq j} \{p_{i,j} \alpha_i(T-u)\} \right], \left\{ \prod_{v=1}^T b_j(x_{T-v}) D_j(T+1) \pi_j \right\} \right] \quad (11)
\end{aligned}$$

The recursion above largely facilitates the computation of the most likely trajectory for each regime up to time t . At time T , the optimal regime sequence inferred from the observation sequence is $\max_j \{\alpha_j(T)\}$, which is selected as the corresponding trajectory for the solution in the Viterbi algorithm.

Here, we briefly demonstrate the principle why the Viterbi algorithm can largely facilitate the computational implementation. Essentially, two back-pointers are recorded for each time t and each regime j . The first backpointer is the optimal preceding regime, and the second backpointer is the optimal preceding time of transition from this preceding regime. In the ‘backtracking’ stage, the optimal regime sequence is tracing backward via those backpointers from time point at T to the initial time point at 1. Hence, an efficient implementation can be realised through the forward and backward recursion⁵.

4.4. Practical Setting

- We choose two settings for the sojourn time distributions. The first setting is to use nonparametric distribution with an upper limit. The advantage of nonparametric distribution is that there is no need to arbitrarily specify the parametric form of the sojourn time distribution. Additionally, we can set an upper limit for the sojourn time distribution of the spike regime. In our implementation, the upper limit for the spike regimes is set to be 2 weeks⁶ for QLD, NSW, VIC, and SA. However, the EM algorithm cannot converge for TAS with the upper limit of 2 weeks in the spike regime. We investigate the reason of the fail in convergence and find that there was a major 6-month outage of the Basslink interconnector⁷, resulting in a continuous period of abnormally high price in TAS over 2015-2016. Hence, the upper limit is modified to be 30 weeks for TAS. For the sake of comparison, we also employ the

⁵ The pseudo-code of the forward-backward algorithm for the HSMM is provided in Guédon (2003).

⁶ We have also tried higher upper limits, such as 3, 4, and 5 weeks, and there is no obvious difference in the empirical results.

⁷ For details, visit (https://en.wikipedia.org/wiki/2016_Tasmanian_energy_crisis).

second setting of a parametric distribution, namely the negative binomial distribution, which is used by Bulla and Bulla (2006).

- The number of regimes is set to be three. This is our choice based on the fact that the inspection of the empirical histogram of the data is multimodal, which is featured with two modes and the right-tailed outliers. Additionally, we aim to capture the low-price regime, the high-price regime, and the spike regimes.
- There are two motivations for us to set the conditional distribution as the Gaussian distribution. Firstly, a mixture of Gaussian distributions⁸ can generally give good fitting for multimodal distributions (McLachlan and Peel, 2004). Secondly, using the Gaussian distribution as conditional distribution can facilitate straightforward interpretation. This choice has been used by in the literature of HMM (Huisman and De Jong, 2003) and also in the literature of HSMM (Lau et al., 2017; Liu and Wang, 2017a; Liu and Wang, 2017b).
- Our implementation is based on the R package “HSMM” (Bulla and Bulla, 2013).

5. Empirical Results

Our empirical analysis contains three main steps. In the first step, we use the EM algorithm⁹ to estimate the parameter set θ in the HSMM. The parameter set includes the parameters in the conditional distributions, the transition probabilities, and the sojourn time distributions. Table 2 shows the estimation results of the three-regime HSMM. To illustrate the feature of bimodal distributions with right outliers, Figure 2 displays the densities for the three regimes along with the empirical densities, from the nonparametric setting. It is more convenient to show the nonparametric sojourn time distributions in Figure 3, and we did not report the

⁸ Note that Hidden (semi-) Markov Models are different from Gaussian Mixture Models. Hidden (semi-)Markov Models have temporal dependence captured by (semi-)Markov chain, while Gaussian Mixture Models do not have temporal dependence.

⁹ We have tried different initial values of parameters in order to check that the converged result is the global maximum rather than the local maximum.

sojourn time distribution in the nonparametric setting for Regimes 1 and 2 in Table 2 because they are large vectors.

Given the estimated parameters, the second step is to employ the Viterbi algorithm to globally decode the most likely sequence of regimes, revealing the information of the timing and transitions of regimes. The sequence of the regimes in the parametric setting is similar to that in the nonparametric setting. For the sake of conciseness, we present the most likely sequence of the regimes, from the nonparametric setting, for the five States over the sample period in Figure 4.

In order to fully reveal the information of sojourn time, we further collect the *ex post* sojourn time information in the third step. Based on the most likely sequence of regimes decoded by the Viterbi algorithm, each time-point of transition can be identified. Then we cut the regime sequence into pieces between the time-points of transition and count the sojourn time for each piece. Then, the number of weeks and the number of times are collected to calculate the average sojourn time and frequency of the three regimes. Table 3 reports the sojourn time information, *ex post*, from the global decoding. We notice that there is a slight difference between the nonparametric setting and the parametric setting. The reason is that the parametric setting has the assumption that the sojourn time follows a negative binomial distribution, whilst the probability of the sojourn time in the nonparametric setting is clustered in some integers, which can be seen in Figure 3.

5.1 Results of Conditional Distribution: Evidence of Three Regimes

In each state, ‘Regime 1’ is featured with both a low mean and a small variance; therefore, ‘Regime 1’ corresponds to the low-price regime. ‘Regime 2’ has both a high mean and small variance, indicating that ‘Regime 2’ refers to the high-price regime. It can be inferred that ‘Regime 3’ is the spike regime because of its conditional distribution and sojourn time distribution. In ‘Regime 3’, the mean is abnormally high, and the variance is considerably

higher than those in the other two regimes. Our estimation result of the spike regime is consistent with Huisman and De Jong (2003) who set the spike regime in their model as a normal distributed random variable whose mean and variance were higher than those of the base regime. In order to better convince the reader, we provide further evidence of the spike regime in terms of the sojourn time in Section 5.2.

Figures 2(A) to 2(E) illustrate the probability densities for the three regimes from the nonparametric setting¹⁰. It is clear that the distribution is the bimodal with the number of the outliers (mainly the right-tail outliers), verifying our choice that the number of regimes is set to be three. The three regimes clearly stand out of the empirical histogram. The area of the left mode corresponds to the low-price regime (highlighted in green), the area of the right mode refers to the high-price regime (highlighted in purple), while the outliers are placed on the spike regime (highlighted in red). We notice that the density function of the spike regime is very flat, which marginally moves above the horizontal axis. This is due to the fact that the frequency of the spike regime (reported in Table 3) is very small, which is typically no more than 6% in the nonparametric setting and less than 9% in the parametric setting. Hence, it is reasonable to observe the flat density of the spike regime as it only covers the right-tail outliers, which is not frequently occurred. Interestingly, the shape of the histogram in the case of the TAS State is different from the other four States, which is featured with the higher density for the high-price regime and the lower density for the low-price regime, while the other four states display contrasting situations.

5.2 Results of Sojourn Time Distribution: Further Evidence of the Spike Regime

The *ex post* sojourn time information in Table 3 provides further evidence to support that ‘Regime 3’ is the spike regime. Except for TAS, the average sojourn time in ‘Regime 3’ is no more than 2 weeks in the nonparametric setting and less than 3 weeks in the parametric

¹⁰ The probability densities for the three regimes from the parametric setting can be provided upon request.

setting. Unlike the definition of jumps, the spike is an abrupt and substantial price change, but more importantly, it will revert to a normal level after a short period of time. Hence, the very short sojourn time we observed in our results exactly matches the characteristics of the spike, which are defined as “temporal level shifts that die out rather quickly” (Huisman and Mahieu, 2003: 428). Based on the two facts: i) the abnormally high mean and large variance of ‘Regime 3’ (in Table 2); and ii) the very short sojourn time of ‘Regime 3’ (in Table 3), it is believed that ‘Regime 3’ should be deemed as the spike regime.

It should be highlighted that the average sojourn time of the spike regime in TAS is a special case. Due to the major 6-month outage of the Basslink interconnector in late 2015, TAS was isolated from the transmission grid of the mainland states in the NEM, resulting in a continuously long period of abnormally high prices. This particularly long sojourn time of the spike regime only happened once, which can be observed in Figure 4(E). This is also the reason that the EM algorithm cannot converge for TAS if the upper limit of the sojourn time distribution is set to be 2 weeks in the spike regime.

5.3 Results of Transition Probabilities

In both settings, the transition probability from the high-price regime to the spike regime ($\gamma_{2,3}$) is generally higher than the transition probability from the low-price regime to the spike regime ($\gamma_{1,3}$) for all five Australian states. In the nonparametric setting, the VIC has the highest transition probability from the high-price regime to the spike regime, $\gamma_{2,3} = 61.72\%$, while SA and TAS have a relatively higher transition probability from the low-price regime to the spike regime, $\gamma_{1,3} = 45.56\%$ and 40.01% , respectively. In the parametric setting, the VIC has the highest transition probabilities to the spike regime, including both $\gamma_{1,3} = 52.85\%$ and $\gamma_{2,3} = 59.86\%$. It can be concluded that the spikes can occur in a comparatively easier manner in the VIC. We will revisit this finding with the discussion in Section 6.2. Comparing results in both settings, the nonparametric setting is more reasonable than the

parametric setting since the transition probability from the low-price regime to the spike regime ($\gamma_{1,3}$) is always non-zero.

5.4 Results of the Regime Sequence

For simplicity, we only plot the most likely sequence of regimes from the nonparametric setting¹¹ in Figure 4. Four interesting periods are highlighted:

- From January 2009 to May 2012, the QLD, the NSW, the VIC and the SA are mainly in the low-price regime, with the several short periods of the high-price regime and some spikes. In the TAS, there are some high-price regimes and two spike regimes before mid-2009. Afterwards, the TAS is similar to the other four States until 2012. A contributing factor in the low-price outcomes would be a general trend reduction in average and peak scheduled demand observed over this period by many states in the NEM, especially over the period 2011-2012 relative to the earlier years, i.e. 2007-2008 (and 2010 in the case of VIC and QLD). The low electricity regime was the result of energy efficiency programs, structural change of electricity intensive industries, and the responses of residential consumers to high electricity prices since 2010 (Saddler, 2013)¹².
- Between June 2012 and June 2014, the QLD, the NSW, the VIC and the SA are in the high-price regime. We will revisit this special period with a detailed discussion in Section 6.1. Note that the TAS enters this remarkable high-price regime earlier in January 2012 and exits later in October 2015.

¹¹ The most likely sequence of regimes from the parametric setting can be requested upon a request.

¹² For example, Saddler (2013) found evidence that Victorian retailer energy efficiency obligation schemes have contributed to the reduction in demand.

- Over the period July 2014 to mid-2015, the QLD, the NSW, the VIC and the SA all have the period of the low-price regime, while the TAS is in the high-price regime. It can be observed that the timing of regimes in the TAS is not synchronised with the other four states, which will be revisited with discussion in Section 6.3.
- From May 2016 and onwards, the QLD, the NSW, the VIC, and the SA enter the high-price regime, and there are some spikes at the end of the sample period. Due to the major outage of the Basslink interconnector, the TAS enters an abnormally long spike regime earlier in October 2015 and leaves in May 2016. It is an untypically long spike which lasts for 28 weeks.

Table 2. Estimation Results of the Three-Regime HSMM for the Australian States

<i>Nonparametric Setting</i>																
		QLD			NSW			VIC			SA			TAS		
		<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>	<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>	<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>	<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>	<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>
Conditional Distribution	<i>Mean</i>	3.291	3.992	5.148	3.339	3.895	5.041	3.277	3.830	4.983	3.388	4.082	4.951	3.292	3.787	4.854
	<i>Variance</i>	0.036	0.057	0.234	0.019	0.051	0.374	0.032	0.052	0.337	0.043	0.084	1.628	0.058	0.037	0.256
TPM	<i>from/to</i>															
	<i>Regime 1</i>	0.00%	85.38%	14.62%	0.00%	85.69%	14.31%	0.00%	59.99%	40.01%	0.00%	54.44%	45.56%	0.00%	87.40%	12.60%
	<i>Regime 2</i>	54.77%	0.00%	45.23%	64.13%	0.00%	35.87%	38.28%	0.00%	61.72%	45.30%	0.00%	54.70%	78.02%	0.00%	21.98%
	<i>Regime 3</i>	28.92%	71.08%	0.00%	52.39%	47.61%	0.00%	45.09%	54.91%	0.00%	28.99%	71.01%	0.00%	39.36%	60.64%	0.00%
Sojourn Time Distribution	<i>P(Spike Sojourn=1)</i>			66%			61%			58%			66%			25%
	<i>P(Spike Sojourn=2)</i>			34%			39%			42%			34%			19%
	<i>P(Spike Sojourn=3)</i>															26%
	<i>P(Spike Sojourn=28)</i>															30%
	<i>P(Spike Sojourn=others)</i>			0%			0%			0%			0%			0%
<i>Parametric Setting</i>																
		QLD			NSW			VIC			SA			TAS		
		<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>	<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>	<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>	<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>	<i>Regime1</i>	<i>Regime2</i>	<i>Regime3</i>
Conditional Distribution	<i>Mean</i>	3.322	4.021	4.601	3.341	3.890	4.635	3.322	3.878	4.678	3.387	4.043	4.843	3.255	3.726	4.543
	<i>Variance</i>	0.044	0.035	0.476	0.019	0.049	0.645	0.040	0.028	0.436	0.046	0.065	1.209	0.044	0.043	0.574
TPM	<i>from/to</i>															
	<i>Regime 1</i>	0.00%	75.65%	24.35%	0.00%	100.00%	0.00%	0.00%	47.15%	52.85%	0.00%	56.61%	43.39%	0.00%	100.00%	0.00%
	<i>Regime 2</i>	50.68%	0.00%	49.32%	59.96%	0.00%	40.04%	40.14%	0.00%	59.86%	59.46%	0.00%	40.54%	48.08%	0.00%	51.92%
	<i>Regime 3</i>	34.60%	65.40%	0.00%	61.17%	38.83%	0.00%	32.59%	67.41%	0.00%	43.12%	56.88%	0.00%	19.66%	80.34%	0.00%
Sojourn Time Distribution	\hat{r}	0.505	0.156	1.275	0.793	0.065	0.127	0.511	0.505	0.083	0.898	0.163	0.545	2.219	0.082	0.123
	\hat{p}	0.045	0.033	0.454	0.083	0.008	0.047	0.029	0.068	0.102	0.072	0.018	0.236	0.165	0.006	0.011

Notes: we use the expectation-maximization (EM) algorithm to estimate the model. There are two settings for the sojourn time distribution: 1) nonparametric setting with an upper limit; 2) parametric setting of negative binomial distribution used by Bulla and Bulla (2006). The negative binomial distribution has the probability mass function $P(X = x|r, p) = \binom{x+r-1}{x} p^x (1-p)^r$, where x is the number of success, r is the number of failures, p is the probability of success, and the quantity in parentheses is the

binomial coefficient. The number of States is set to be three as the empirical histogram of the data featured with the two modes and the right-tailed outliers. The conditional distributions are set to be Gaussian distributions. The diagonal entries of the TPM in HSMM are always zero. We did not report the sojourn time distribution in the nonparametric setting for Regime 1 and 2 because they are large vectors, which are displayed in Figure 3.


Table 3. Sojourn Time Information, *ex post*, from Global Decoding

<i>Nonparametric Setting</i>															
	QLD			NSW			VIC			SA			TAS		
	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>
<i>No. of Weeks</i>	214	188	20	192	217	13	205	200	17	215	184	23	161	218	43
<i>No. of Times</i>	20	29	15	20	23	10	10	13	12	16	23	18	17	19	6
<i>Average No. of Weeks</i>	10.70	6.48	1.33	9.60	9.43	1.30	20.50	15.38	1.42	13.44	8.00	1.28	9.47	11.47	7.17
<i>Frequency</i>	50.71%	44.55%	4.74%	45.50%	51.42%	3.08%	48.58%	47.39%	4.03%	50.95%	43.60%	5.45%	38.15%	51.66%	10.19%
<i>Parametric Setting</i>															
	QLD			NSW			VIC			SA			TAS		
	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>	<i>Regime 1</i>	<i>Regime 2</i>	<i>Regime 3</i>
<i>No. of Weeks</i>	235	151	36	190	213	19	246	154	22	216	176	30	119	247	56
<i>No. of Times</i>	20	24	13	17	21	8	12	18	16	16	15	11	8	14	7
<i>Average No. of Weeks</i>	11.75	6.29	2.77	11.18	10.14	2.38	20.50	8.56	1.38	13.50	11.73	2.73	14.88	17.64	8.00
<i>Frequency</i>	55.69%	35.78%	8.53%	45.02%	50.47%	4.50%	58.29%	36.49%	5.21%	51.18%	41.71%	7.11%	28.20%	58.53%	13.27%


Notes: the *ex post* sojourn time information is obtained in the following procedure. Firstly, the Viterbi algorithm is employed to global decode the most likely sequence of regimes. Secondly, the time points of transition are identified in the regime sequence. Then we cut the regime sequence into pieces between the time points of transition and count the sojourn time for each piece. Lastly, the number of weeks and number of times are collected to calculate the average sojourn time and frequency of the three regimes.

Figure 2. Probability Densities for the Three-Regime HSSM


A) Queensland (QLD)




B) New South Wales (NSW)




C) Victoria (VIC)



D) South Australia (SA)



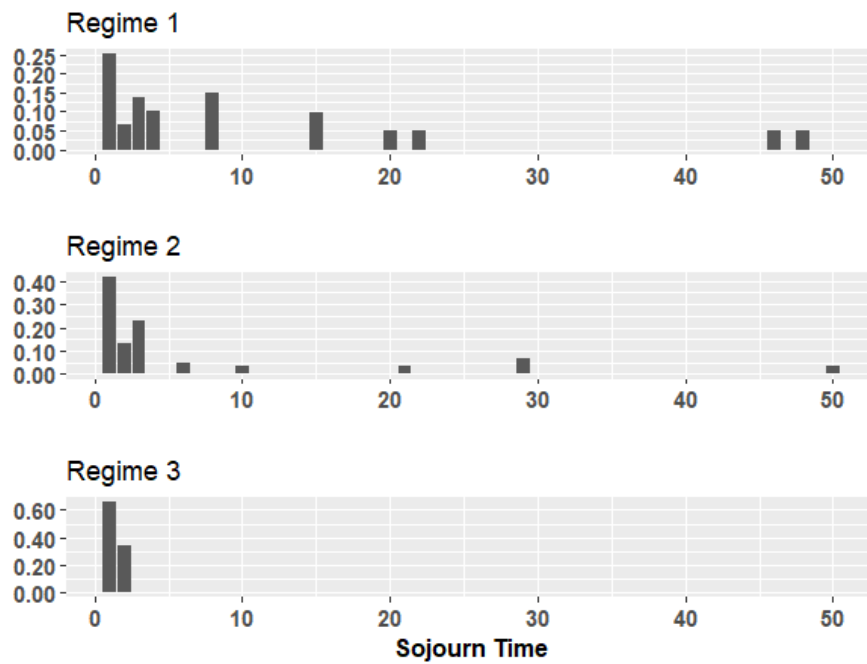
E) Tasmania (TAS)




F

Figure 3. Nonparametric Sojourn Time Distribution for the Three-Regime HSSM


A) Queensland (QLD)



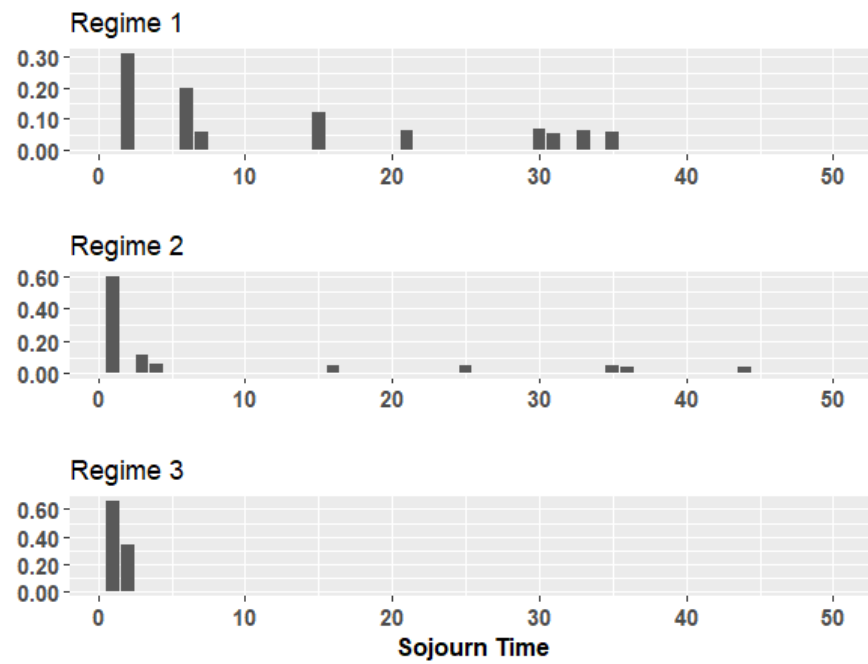
B) New South Wales (NSW)



C) Victoria (VIC)



D) South Australia (SA)



E) Tasmania (TAS)






Figure 4. Regimes Decoding Results for the Three-Regime HSSM


A) Queensland (QLD)




B) New South Wales (NSW)




C) Victoria (VIC)



D) South Australia (SA)



E) Tasmania (TAS)



6. Discussion and Implications

In this section, we link the empirical results to the policy changes in the Australian electricity market and further discuss the policy implications and the contributing factors for the different characteristics of the Australian electricity markets at the state-level.

6.1. Discussion on the Policy Reforms

From our empirical results, we observe that the persistence of the high-price regime between 2012 and 2014 coincides with the environmental policy reforms adopted by the Australian States. For the period from 2009 to 2010, the Australian government could not implement the key 2007 election commitment relating to the introduction of the greenhouse gas emissions trading scheme until 2010. The legislation underpinning the scheme was passed by the House

of Representatives, but it has been rejected by the Senate. Eventually, the greenhouse gas emissions trading scheme was adopted in July 2012 to become ineffective in July 2014. For example, Nelson et al. (2012) argue that the lack of policy certainty in relation to the climate change policy effectively prevented firms from investing in projects mitigating carbon emissions and in investing in non-fossil energy sources. They provide the empirical evidence that delaying the provision of the policy certainty resulted in firms investing too heavily in open-cycle gas turbines, investments that minimise the risk associated with the investment of the capital. This led to the significant increase in wholesale electricity prices, thus, imposing the largely deadweight loss cost on society. Indeed, the greenhouse gas emissions trading scheme was only adopted in July 2012, while we can see that Regime 2 was already present even prior to July 2012 in several states, albeit for short periods.

Reports from the Australian government in 2013 indicate a 7.7 percent decline in greenhouse gas emissions in the NEM system per year as resulted in the operation of the carbon tax regime (Commonwealth of Australia, 2013). However, the effect of the carbon tax on emissions has been similar to a tax on exports or a tax on the import-competing industries. Although carbon emissions may decline after the carbon tax was imposed (which justifies the presence of less spikes during the era of the carbon tax regime), it could be counted as an environmental gain given that the ultimate effect was that emissions simply rose overseas (Meng et al., 2014). The net effect was a pure deadweight cost to the economy. In Australia, as part of the compensation package for the carbon tax, the Australian Government lowered some average income tax rates, but actually increased marginal tax rates for around 2 million taxpayers. Instead of mitigating the adverse effects of the carbon tax on the labour market, these changes to the personal tax system have likely exacerbated those effects. That had exactly the opposite policy of what policymakers would do if they were trying to capture a ‘double dividend’ from environmental taxation.

The Australian environmental policy experience has been very limited. Many opportunities to increase environmental policy capacity have been sidestepped. The view, according to which, investments in environmental protection represent a win-win for both business and society has enjoyed only limited penetration in Australia's environmental policy design (Young, 2000). In 2014, energy policymakers repealed the carbon price, wound back the Renewable Energy Target and established the Emissions Reduction Fund (ERF), which has contracted for more than 100 million tons of CO² emission reductions, but it largely sidestepped the reforms needed to address emerging energy trends, such as low demand growth, the rise of distributed wind power generation, the boom in domestic solar power and the dramatic growth of coal seam gas. In Figure 3, we observe that wholesale prices did not revert back to their pre-carbon tax 2012 mean levels after it was repealed in 2014. This evidence indicates that a carbon tax has changed the price change dynamics across the region (Apergis and Lau, 2015). The upshot was that the period 2013-16 left the energy industry with huge uncertainty about what is in store, at a time when it craves reassurance more than ever. The country needs an energy policy that should give priorities to build on current climate change policies in order to generate a stable, long-term approach that will lead the transition to a low-emissions economy, to revive energy market reforms through the Energy Council, and to maximise the value of Australia's gas resources and ensure continuity of supply. Such policies do require galvanising cooperation across states and territory governments and energy groups. Moreover, the new energy policy should develop a renewed reform agenda for the Energy Council, in a sense that addresses all these issues with a focus on outcomes, rather than being mired in process, as it has been so far.

6.2. Discussion on the Price Spikes

The spikes in energy-only power markets are normal events that take place whenever the systems run short of reserves and get close to load shedding (Higgs and Worthington, 2008).

According to both parametric and nonparametric results, the transition probability to the spike regime in VIC is generally highest among all five states¹³ (see nonparametric and parametric results in Table 2). One contributing factor for Victoria to have a higher transition probability to the spike regime is that Victoria has a relative paucity of baseload or intermediate gas-fired generation plant (apart from the NewPort power station) in comparison with other mainland NEM States.¹⁴ Higgs et al. (2015) found evidence that the generation mix in wholesale electricity influences wholesale price volatility, and there is an increasing price volatility for coal, open-cycle gas turbine, and hydro generation. In line with Higgs et al. (2015), the use of coal affects not only price volatility, but also the transition probability to the spike.

Another potential factor would be represented by the changes in the bidding behaviour of the generation companies as we saw with Stanwell after they were directed by the QLD Government to moderate the impact of their bidding on observed increases in QLD wholesale market prices, and this strategic changes in bidding behaviour of generation companies may have wider impact on the behaviour of generation companies in other states , including VIC. Moreover, the shift to renewables is expected to moderate wholesale price increases via the merit order effect as outlined in Bell et al. (2017), for example.

6.3. Discussion on Tasmania

According to our empirical results, Tasmania has different characteristics in terms of hidden regimes. For example, we observe that there is a higher density for the high-price regimes and a lower density for the low-price regime in Tasmania. Furthermore, there is a relatively longer sojourn in the spike regime in Tasmania, and the timing of hidden regimes in

¹³ A high transition probability to the spike regime in Victoria indicates that it is easier for Victoria to have spikes, but it does not necessarily imply a large number of times in the spike regime (there are only 12 times of spikes realised in Victoria by the nonparametric setting).

¹⁴ With more than 65 million tons of brown coal burnt every year, Victoria is the number one brown coal consumer (i.e. brown coal contributes 92 percent of electricity generation in Victoria). Higgs et al. (2015) find the evidence for the generation mix in wholesale electricity.

Tasmania is not synchronized with the other four states. We can suggest that there are several aspects in Tasmania, which give rise to the different characteristics in terms of the hidden regime.

The first issue is related to geographical reasons. Tasmania is geographically separated from the mainland and there are limited inter-linkages in the electricity market in Tasmania. This could be among the main reasons for the different features of its electricity market. Secondly, the market position in Tasmania presents a number of potential risks to the investors. For example, the electricity production in the region uses a lot of wind for electricity generation (Bell et al., 2017). Thirdly, from a wholesale market perspective, the price behaviour might be linked (from 2012-13 onwards) to the observed aggressive bidding by Hydro Tasmania that occurred over the period 2013-2014 that led to hydro storages being run down significantly and lower inflows occurring after that period. This is linked, in turn, to the special nature of the generation structure in Tasmania - mainly hydro and wind power - and their sensitivities to rainfall and wind conditions (Minshull, 2017). Lastly, the relatively longer sojourn in the spike regime is mainly due to the major 6-month outage of the Basslink interconnector that also occurred effectively in late 2015, islanding TAS from the mainland states in the NEM. Therefore, due to the four potential reasons above, the electricity market in Tasmania contains the different timing of regimes when it is compared with other states.

Through the empirical results from the HSMM, we reveal many new insights about the hidden regimes in the Australian NEM. Those insights are potentially helpful for the policy-makers and market participants. It is worthwhile to note that the limitation of this study is that the policy discussion is based on our observations and a more formal analysis (e.g. causality test) is needed to establish such an association/relationship. However, the formal causality analysis goes beyond the scope of this work and is left for future research.

7. Conclusion

In this paper, we took the consideration of the stylised features of the spikes and the time-varying mean in the Australian electricity prices and employed a three-regimes HSMM to analyse the hidden regimes in five Australian states (New South Wales, Queensland, South Australia, Tasmania, and Victoria) for the period from June 8, 2008 to July 3, 2016. As for the estimation results, we find evidence that the three hidden regimes correspond to a low-price regime, a high-price regime, and a spike regime. With the exception of Tasmania, the sojourn time in the spike regime is typically very short, no more than two weeks in the nonparametric setting and less than three weeks in the parametric setting. Victoria has higher transition probabilities to the spike regime. Running the decoding algorithm, the analysis systemically finds the timing of the three regimes, and thus, we link the empirical results to the policy changes in the Australian electricity market. The persistence of high-price regime during 2012 and 2014 coincides with the environmental policy reforms adopted by the Australian States. There are higher transition probabilities to the spike regime in Victoria, and Tasmania has different characteristics in terms of hidden regimes. The limitation of this study is that the policy discussion is based on our observations and a more formal analysis (e.g. causality test) is needed to establish such an association/inference. Future work can focus on employing formal econometric methods to study the causal relationship between the policy changes and the transmission of different hidden regimes.

Acknowledgements

The research of the corresponding author was supported by the Economic and Social Research Council (UK) [grant number ES/J50001X/1] and the Royal Economics Society Junior Fellowship.

References

- Apergis, N., & Lau, M.C.K. (2015). Structural breaks and electricity prices: Further evidence on the role of climate policy uncertainties in the Australian electricity market. *Energy Economics*, 52, 176–182.
- Apergis, N., Barunik, J., & Lau, M.C.K. (2017a). Good volatility, bad volatility: What drives the asymmetric connectedness of Australian electricity markets? *Energy Economics*, 66, 108–116.
- Apergis, N., Fontini, F., & Inchauspe, J. (2017b). Integration of regional electricity markets in Australia: A price convergence assessment. *Energy Economics*, 62, 411–418.
- Arouri, M.E.H., Jouni, J., & Nguyen, D.K. (2012). On the impacts of oil price fluctuations on European equity markets: Volatility spillover and hedging effectiveness. *Energy Economics*, 34 (2), 611–617.
- Barunik, J., Kocenda, E., & Vacha, L. (2015). Volatility spillovers across petroleum markets. *Energy Journal*, 36 (3), 309–329.
- Baum, L.E., Petrie, T., Soules, G., & Weiss, N. (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics*, 41 (1), 164–171.
- Bell, W.P., Wild, P., Foster, J., & Hewson, M. (2017). Revitalising the wind power induced merit order effect to reduce wholesale and retail electricity prices in Australia. *Energy Economics*, 67, 224–241.
- Bulla, J., & Bulla, I. (2006). Stylized facts of financial time series and hidden semi-Markov models. *Computational Statistics and Data Analysis*, 51 (4), 2192–2209.
- Bulla, J., & Bulla, I. (2013). *HSMM: Hidden Semi Markov Models*. <https://CRAN.R-project.org/package=hsmm> R package version 0.4.

- Chang, C.L., McAleer, M., & Tansuchat, R. (2010). Analyzing and forecasting volatility spillovers, asymmetries and hedging in major oil markets. *Energy Economics*, 32 (6), 1445–1455.
- Commonwealth of Australia (2013). *Australian National Greenhouse Accounts – Quarterly Update of Australia’s National Greenhouse Gas Inventory December Quarter 2012*, Australian Government.
- Cutler, N.J., Boerema, N.D., MacGill, I.F., & Outhred, H.R. (2011). High penetration wind generation impacts on spot prices in the Australian national electricity market. *Energy Policy*, 39 (10), 5939–5949.
- De Jong, C. (2006). The nature of power spikes: A regime-switch approach. *Studies in Nonlinear Dynamics & Econometrics*, 10 (3).
- Dempster, A.P., Laird, N.M., & Rubin, D.B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B (Methodological)*, 39 (1), 1–38.
- Deng, S. (2000). *Stochastic Models of Energy Commodity Prices and Their Applications: Mean-reversion with Jumps and Spikes*. Berkeley, CA: University of California Energy Institute.
- Diebold, F.X., Lee, J.H., & Weinbach, G.C. (1994). Regime switching with time-varying transition probabilities. *Business Cycles: Durations, Dynamics, and Forecasting*, 1, 144–165.
- Du, X., Yu, C.L., & Hayes, D.J. (2011). Speculation and volatility spillover in the crude oil and agricultural commodity markets: A Bayesian analysis. *Energy Economics*, 33 (3), 497–503.
- Ewing, B.T., & Malik, F. (2013). Volatility transmission between gold and oil futures under structural breaks. *International Review of Economics and Finance*, 25, 113–121.

- Gozgor, G., & Memis, C. (2015). Price volatility spillovers among agricultural commodity and crude oil markets: Evidence from the range-based estimator. *Agricultural Economics–Czech*, 61 (5), 214–221.
- Gozgor, G., Lau, C.K.M., & Bilgin, M.H. (2016). Commodity markets volatility transmission: Roles of risk perceptions and uncertainty in financial markets. *Journal of International Financial Markets, Institutions and Money*, 44, 35–45.
- Guádon, Y. (2003). Estimating hidden semi-Markov chains from discrete sequences. *Journal of Computational and Graphical Statistics*, 12 (3), 604–639.
- Higgs, H. (2009). Modelling price and volatility inter-relationships in the Australian wholesale spot electricity markets. *Energy Economics*, 31 (5), 748–756.
- Higgs, H., & Worthington, A. (2008). Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The Australian wholesale spot electricity market. *Energy Economics*, 30 (6), 3172–3185.
- Higgs, H., Lien, G., & Worthington, A.C. (2015). Australian evidence on the role of interregional flows, production capacity, and generation mix in wholesale electricity prices and price volatility. *Economic Analysis and Policy*, 48, 172–181.
- Hu, X., Grozev, G., & Batten, D. (2005). Empirical observations of bidding patterns in Australia's National Electricity Market. *Energy Policy*, 33 (16), 2075–2086.
- Huisman, R., & De Jong, C. (2003). Option pricing for power prices with spikes. *Energy Power Risk Management*, 7 (11), 12–16.
- Huisman, R., & Mahieu, R. (2003). Regime jumps in electricity prices. *Energy Economics*, 25 (5), 425–434.
- Janczura, J., & Weron, R. (2010). An empirical comparison of alternate regime-switching models for electricity spot prices. *Energy Economics*, 32 (5), 1059–1073.

- Kulkarni, V.G. (2016). *Modeling and Analysis of Stochastic Systems*. Third Edition, Boca Raton, FL: CRC Press.
- Lau, M.C.K., Vigne, S.A., Wang, S., & Yarovaya, L. (2017). Return spillovers between white precious metal ETFs: The role of oil, gold, and global equity. *International Review of Financial Analysis*, 52, 316–332.
- Liu, Z., & Wang, S. (2017a). Decoding Chinese stock market returns: Three-state hidden semi-Markov model. *Pacific-Basin Finance Journal*, 44, 127–149.
- Liu, Z., & Wang, S. (2017b). Understanding the Chinese stock market: International comparison and policy implications. *Economic and Political Studies*, 5 (4), 441–455.
- McLachlan, G., & Peel, D. (2004). *Finite Mixture Models*. New York: John Wiley & Sons.
- Meng, S., Sirwardana, M., & McNeil, J. (2014). The impact of the Australian carbon tax on industries and households. *Margin: The Journal of Applied Economic Research*, 8 (1), 15–37.
- Minshull, L. (2017). Securing Tasmania's Energy Future: How Tasmania can Become Energy Self-reliant and an Exporter of Renewable Energy. *The Australia Institute Tasmania Discussion Paper*, Manuka: The Australia Institute Tasmania.
- Mount, T.D., Ning, Y., & Cai, X. (2006). Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Economics*, 28 (1), 62–80.
- Nelson, T., Kelley, S., & Orton, F. (2012). A literature review of economic studies on carbon pricing and Australian wholesale electricity markets. *Energy Policy*, 49, 217–224.
- Saddler, H. (2013). Power Down: Why is Electricity Consumption Decreasing? *The Australia Institute Tasmania Discussion Paper No. 14*, Manuka: The Australia Institute Tasmania.

- Viterbi, A. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13 (2), 260–269.
- Yarovaya, L., Brzezczynski, J., & Lau, M.C.K. (2016). Intra-and inter-regional return and volatility spillovers across emerging and developed markets: Evidence from stock indices and stock index futures. *International Review of Financial Analysis*, 43, 96–114.
- Young, S. (2000). *The Emergence of Ecological Modernisation: Integrating the Economy and the Environment?* London: Routledge.
- Yu, S. Z. (2010). Hidden semi-Markov models. *Artificial intelligence*, 174 (2), 215–243.
- Zucchini, W., MacDonald, I.L., & Langrock, R. (2016). *Hidden Markov Models for Time Series: An Introduction Using R*. Boca Raton, FL: CRC Press.