HEURISTIC TO DETERMINE THE CONTRACT POSITION

A REPORT FOR THE INDEPENDENT COMPETITION AND REGULATORY COMMISSION

24 JANUARY 2020

CONTENTS

Tables

Figures

1 INTRODUCTION

Frontier Economics has been engaged to advise the Independent Competition and Regulatory Commission (ICRC) to create a heuristic which determines an appropriate contracting position for an energy retailer in the Australian Capital Territory (ACT) which is robust enough to be used in future years.

1.1 Background

The ICRC is required to regulate prices for the small customers purchasing electricity from ActewAGL. To inform this the ICRC needs estimates of retailers' energy purchase costs. A common approach to estimating a retailer's energy purchase costs is to estimate the cost to the retailer of buying electricity from the wholesale electricity market, having regard to the hedging instruments that retailers typically purchase to manage the risks associated with volatile electricity wholesale prices. Under this approach, the estimate of a retailer's energy purchase costs depends on the number and type of hedging instruments purchased, which is commonly referred to as the retailer's contract position.

1.2 Frontier Economics' engagement

Frontier Economics has been engaged by the ICRC to provide advice on an appropriate contracting position for an energy retailer supplying small customers in the ACT. This will involve providing a heuristic to determine a prudent and efficient contracting position, which can be used to help set the energy purchase cost allowance for a retailer.

1.3 This report

This report sets out our advice to the ICRC on the contract position heuristic for retailers in the ACT distribution network area. This report is structured as follows:

- Section 2 provides an overview of the approach used to create the contracting heuristic.
- Section 3 discusses the half-hourly prices and half-hourly load used in our analysis.
- Section 4 discusses the contract prices used in our analysis.
- Section 5 discusses the assumed contract position.
- Section 6 provides our estimate of the contracting heuristic.

2 APPROACH TO CREATING THE HEURISTIC

Under the settlement rules in the National Electricity Market (NEM) retailers are responsible for purchasing electricity to meet the load of their customers in the wholesale electricity market. A retailer will pay, for each half hour, its customer's electricity load in that half-hour multiplied by the relevant regional reference price from the wholesale electricity spot market for that half hour. For customers in the ACT, the relevant regional reference price is the New South Wales regional reference price.

These settlement payments that retailers face can be extremely volatile. Electricity load for small customers can vary significantly from one half hour to the next, and electricity spot prices can be anywhere between the Market Price Cap (which is currently $14,700/MWh) and the market floor price (which is -$1,000/MWh). Since retailers will typically commit to supply their customers at a specified retail price for a period of time, this volatility in settlement payments can result in retailers paying more for electricity than they receive for that electricity through the retail price they have agreed with their customers. At worst, this exposes the retailers to the risk of financial failure.

To manage the risks associated with volatile load and spot prices, retailers will typically seek to hedge their exposure to spot prices by entering into hedging arrangements. There are a number of ways that retailers can hedge their exposure to spot prices. The most common are the following:

- Vertical integration through ownership of an electricity generator. A retailer that owns a generator has what is known as a natural hedge: when the spot price is high, the retailer will have to pay the high spot price for its customer's load but, as the owner of a generator, will also receive the high spot price for its electricity generation.
- Power purchase agreements with a generator. Power purchase agreements provide a similar hedging benefit to vertical integration, but they do so through contractual arrangements between a retailer and a generator, rather than through ownership.
- Financial derivatives. There are a range of financial derivatives that are available to retailers (and generators) to hedge their exposure to volatile spot prices. The most common are swap contracts (which effectively lock-in a spot price for the counterparties) and cap contracts (which effectively cap the spot price for a retailer). These are traded both on the stock exchange and over-the-counter between participants.

Retailer's energy purchase costs are typically taken to be the average cost to a retailer of purchasing electricity from the wholesale market for its customers, taking into account both the retailer's settlement payments to the Australian Energy Market Operator (AEMO) and the financial outcomes from the retailer's hedging arrangements.

Regulatory practice in Australia has typically focused on estimating the energy purchase cost for a benchmark retailer. In doing so, regulators have typically assumed that the benchmark retailer will make use of exchange-traded financial derivatives to hedge its exposure to spot prices. The assumption that a benchmark retailer will use exchange-traded financial derivatives is typically based on the following reasoning:

- Any retailer of a reasonable size should be able to hedge its exposure to wholesale spot prices using exchange-traded financial derivatives, while vertical integration and entering power purchase agreements can be impractical for retailers with a smaller retail position in a market or with a less certain retail position.

- Prices for exchange-traded financial derivatives are transparent, since they are traded on the ASX. In contrast, the costs of building generation plant or entering into power purchase agreements are less transparent.

In practice, it is clear that retailers in the NEM do adopt a mix of hedging strategies, including vertical integration and power purchase agreements. Retailers will presumably vertically integrate or enter into power purchase agreements because they think these strategies offer advantages that financial derivatives cannot; by excluding vertical integration and power purchase agreements from consideration, therefore, regulators will, if anything, tend to overstate the costs that retailers will face, or understate the risk management that retailers can achieve.

We follow this typical approach of assessing the contract position that retailers enter into based on an estimate of the least risk contracting position that a prudent retailer would enter into to supply electricity to their customers. The hedging contracts that we base this analysis on are quarterly base swaps, peak swaps and base caps, traded on ASXEnergy.

To estimate the contract position in this way, we need to answer four questions:

- What is the expected half-hourly load of the retailer's customers?
- What are the expected half-hourly spot prices that retailers will face?
- What is the cost of financial hedging contracts?
- What kind of hedging position is a prudent retailer likely to adopt?

From the answers to these questions we can calculate the contract position that a retailer would enter into.

We address these questions in the sections that follow.

3 HALF-HOURLY SPOT PRICES AND HALF-HOURLY LOAD

This section addresses the first two questions we need to answer to estimate a prudent and efficient contract position:

- What is the expected half-hourly load of the retailer's small customers?
- What are the expected half-hourly spot prices that retailers will face?

We deal with these questions together because we believe it is important to forecast half-hourly spot prices and half-hourly load in a way that accounts for the correlation between prices and load. After all, this correlation is a key driver of the risks that retailers face; if load and prices are strongly correlated this means that retailers are purchasing more of their electricity at times of higher electricity prices.

Our approach to determining the expected half-hourly load and spot prices is to use historical data, combined with any expected future trends, to simulate a range of potential future outcomes. We do not seek to forecast a single specific outcome.

3.1 Historical data on half-hourly price and load

Our modelling of the contract position requires projections of half-hourly spot prices in the ACT and projections of customer load to be supplied by retailers in the ACT.

In our view, the best source of data about half-hourly patterns of spot prices, half-hourly patterns of customer load, and the correlation between the two, is historical data. The historical data on prices and load will reflect all of the complex factors that drive both spot prices and customer load, and the interactions between them. The alternative to relying on historical data is to develop forecasting models for both spot prices and customer load. However, it is extremely challenging to develop forecasting models that can capture all the complex factors that drive spot prices and customer load, and their relationship, at the half-hourly level. It is for this reason that we favour relying on historical data for this purpose.

Relying on historical data does not mean that we cannot adjust the historical data to account for any expected trends over time in load or prices. Indeed, the approach that we adopt in using historical data to develop estimates of future spot prices and customer load explicitly accounts for expected trends in both spot prices and customer load. We do this in the following way:

- We adjust historical load data to account for a trend towards a 'peakier' load shape (which implies that the difference between peak load and average load is increasing).
- We adjust historical spot prices to account for market expectations of quarterly average spot prices in future, as implied by the market price of forward contracts for electricity.

Our approach to making these adjustments is described in Section 3.3.2.

In using historical data to develop estimates of future customer load we also considered adjusting historical load to account for weather patterns. For instance, if the historical data on customer load includes a number of years with higher customer load because of higher than average summer temperatures that are not expected to persist in future (or lower customer load because of lower than average summer temperatures that are not expected to persist in future) then this can be accounted for by developing a model to 'weather normalise' historical data. However, models of this type are complex and often involve subjective judgements about which potential models best explain the historical data. For these reasons, and because any approach to adjust customer load should also be applied to adjust spot prices (to ensure the relationship between these is preserved), we have not sought to weather normalise the customer load or spot price data. As a result, we are implicitly assuming that the weather conditions that have occurred over the historical period over which we use data on customer load and spot prices are representative of the weather conditions that will occur in future.

The historical data that we use is provided to us by the ICRC, and consists of:

- For prices, the half-hourly spot prices for the NSW regional reference node, as published by AEMO.
- For customer load, the half-hourly net system load profile (NSLP) data that AEMO publish.

3.2 Selecting appropriate historical data

When using historical data on prices and load in this way, a useful starting point is to choose data on prices and load from an historical period that we think is likely to be most consistent with the forecast period. For example, the closure of coal-fired power stations may have substantial impacts on price levels and volatility. Likewise, the increasing adoption of rooftop solar PV is likely to materially affect load factors and prices over time.

The data that is provided by the ICRC is for 2014/15 to 2018/19. In our view, a time series of this length is appropriate for these purposes because it is likely to capture a broad range of potential market outcomes without becoming too out of date. This time period coincides with the ICRC's decision in its recent Electricity Model and Methodology Review[1] to use the most recent five years of data to determine customer load and spot prices.

Analysis of trends in historical data

Figure 1 shows the annual load factor for the NSLP for the ACT for the last five financial years. The load factor is a common measure of the 'peakiness' of customer load. It is calculated by dividing average annual load over the year by peak load in the year. The load factor can take any value between zero and one; a load factor of one represents a load that is perfectly flat (this is, the load is the same in each half-hour of the year). The lower the load factor, the peakier the customer load. We can see from **Figure 1** that customer load is quite peaky, having a value below 0.4. We can also see that there is a slight trend to a lower load factor from 2014/15 to 2018/19, shown by the dotted trend line. This represents a slight increase in in the peakiness of customer load.

Figure 2 shows the average daily load profile for the NSLP for the ACT for the last five financial years. These daily profiles represent the consumption pattern of customer load during an average day. These daily profiles have been normalised to the same annual consumption (that is, we have scaled total customer load in each year to the same annual value of 1 GWh) so that we can highlight differences in the *timing* of daily consumption, rather than differences in *total* annual consumption. These profiles are very similar between years, although there is an apparent trend towards relative reductions in load during the day, presumably as a result of increased rooftop solar PV.

1 ICRC, *Electricity Model and Methodology Review 2018-19*, Report 5 of 2019. May 2019.

Figure 1: Load factor for customers in the ACT

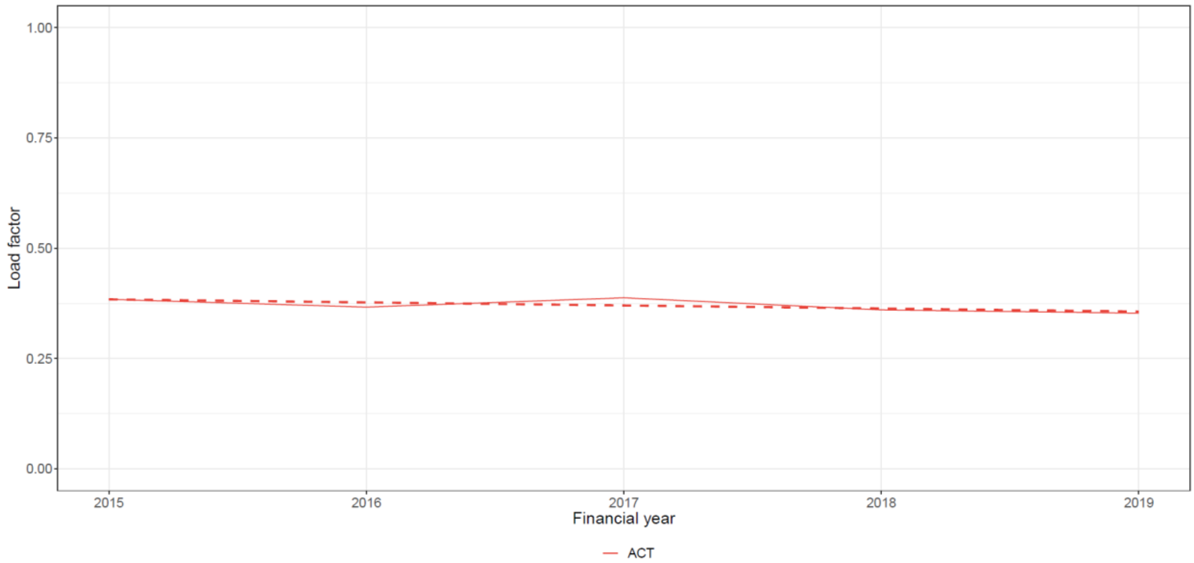

Source: Frontier Economics analysis of AEMO data

Figure 2: Average daily profile for customers in the ACT

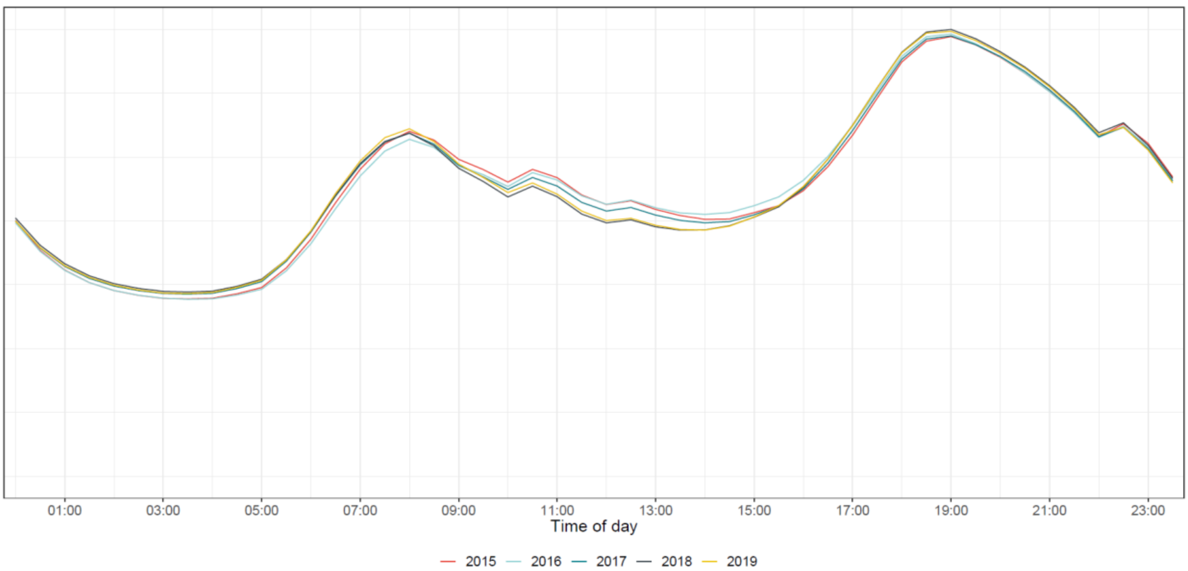

Source: Frontier Economics analysis of AEMO data

Figure 3 shows the average daily profile for NSW spot prices for the last five financial years. It is no surprise to see that there is greater volatility in daily patterns of spot prices than there is in daily patterns of customer load. As with **Figure 2**, these prices have been normalised to show differences in the *timing* of prices, rather than differences in the average *level* of prices. However, in each case we do see similar patterns of low overnight prices, a price spike tending to occur in the morning, and further high prices tending to occur in the mid-afternoon to evening. This general trend is a response to patterns of total demand and total supply for electricity; because demand tends to be higher in the mornings and the evenings, and supply increasingly tends to be higher during the middle of the day when solar is generating, prices tend to be highest in the mornings and evenings and lower during the day. Of course there is variability in these patterns of total demand and supply, with changes in supply (such as plant retirement or major network or generation outages reducing supply, or new investment adding to supply) and changes in demand (such as closure of industrial facilities or changes in weather conditions) all affecting observed historical prices.

In our view, the patterns in prices that we observe in **Figure 3** are sufficiently consistent that it is appropriate to use the patterns of prices observed in each of these historical years as a basis for estimating future patterns of prices.

Figure 3: Average daily profile for NSW spot prices

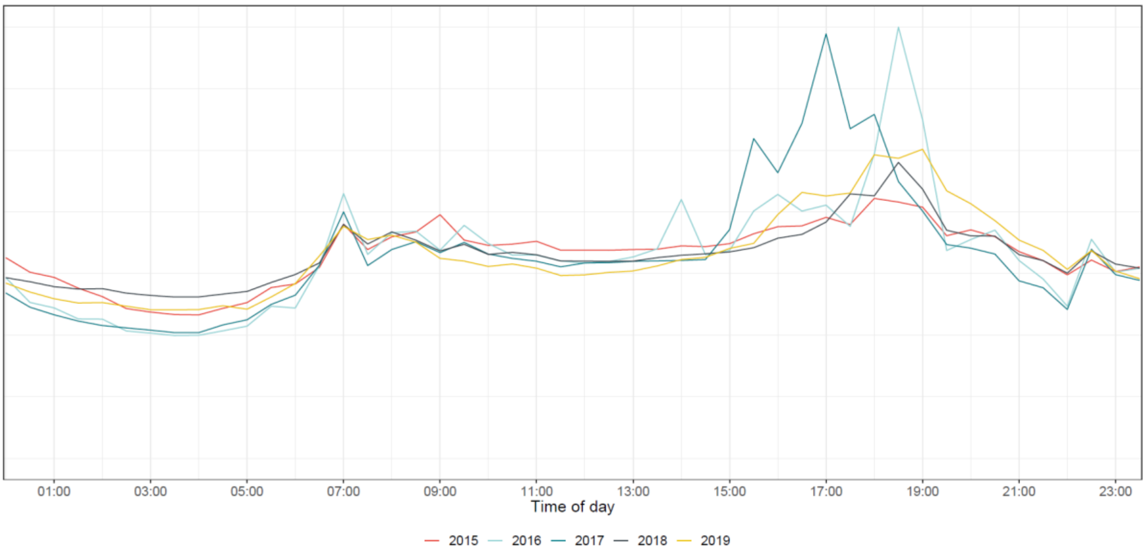

Source: Frontier Economics analysis of AEMO data

Figure 4 combines the historical customer load data and spot price data to report the load premium (calculated as the load-weighted price divided by the time-weighted price) for each of the last five financial years. We can see from **Figure 4** that the load premium over 2014/15 to 2018/19 was reasonably constant. There was a material increase in the load premium for 2015/16, because the spike in evening prices in 2015/16 was particularly strongly correlated with peak demand for small customers (as seen in **Figure 2** and **Figure 3** in 2015/16 both load and prices peak at around 7 in the evening). However, we consider that this is part of expected variability in load and price outcomes.

Figure 4: Load premium for customers in the ACT, based on NSW spot prices

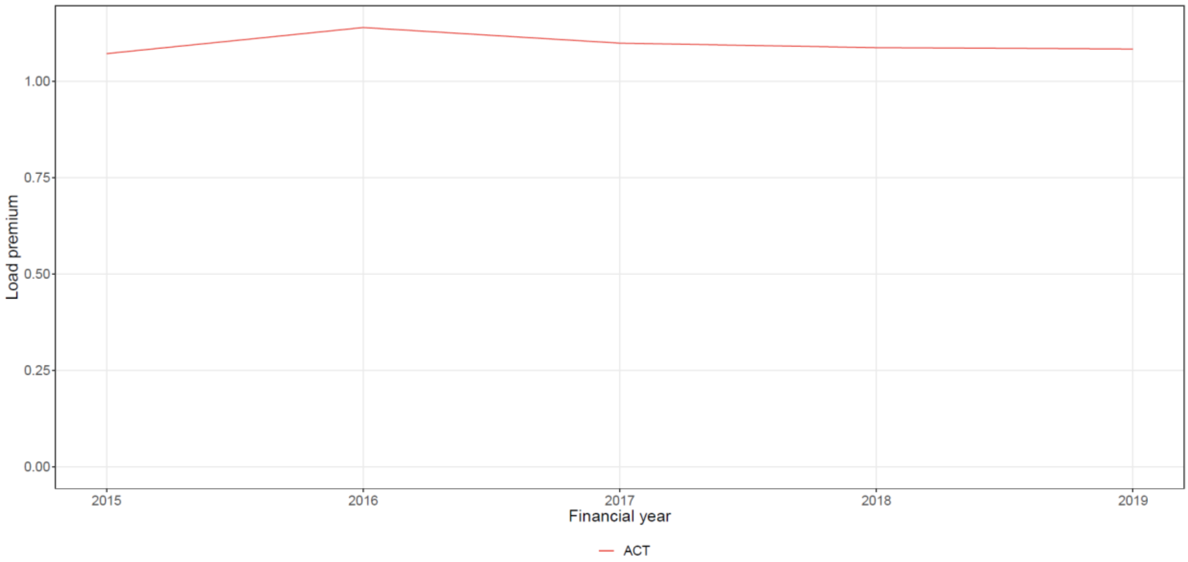

Source: Frontier Economics analysis of AEMO data

3.3 Projecting half-hourly load and spot prices

3.3.1 Monte Carlo simulation

Rather than take a single one of the historical years from 2014/15 to 2018/19 as representative of outcomes in future years (we focus on 2019/20 to 2023/24), we perform a Monte Carlo simulation on the five years of half-hourly load and price data. In our view there are two benefits of using a Monte Carlo analysis:

- Any single year will be subject to unique market conditions that are unlikely to be repeated. This creates the risk that any single year may not be representative of conditions that might be expected in the future. However, using a Monte Carlo approach based, in this case, on five years of data increases the likelihood of basing our analysis on a representative set of conditions.
- Using a Monte Carlo analysis allows us to create a distribution of market conditions, providing some insight into the expected contracting position.

The Monte Carlo simulation is used to generate a year of half-hourly data by randomly drawing one day of data, from the pool of available historical days, for each day of the forecast year. This random drawing of data is done from a pool of like days (where days are classified according to day type – weekday/weekend – and quarter). The Monte Carlo simulation is then performed 100 times to get a distribution of simulated years, which allows us to choose a number of simulated years from within this distribution to use in the modelling.

For example, the first simulated year will be generated as follows:

- The first day of 2019/20 is 1 July 2019, which is a Monday. Since this is a Monday in Q3, the half-hourly load and spot data for the first day of 2019/20 will determined by randomly drawing a day's half-hourly data from all the Q3 weekdays that occurred in 2014/15 to 2018/19.
- The second day of 2020 is 2 July 2019, which is a Tuesday. Since this is a Tuesday in Q3, the half-hourly load and spot data for the second day of 2019/20 will also be determined by randomly drawing a day's half-hourly data from all the Q3 weekdays that occurred in 2014/15 to 2018/19.

- And so on for the 365 days that make up 2019/20, having regard, for each day, to its type and its quarter.

This process will then be repeated 99 further times, resulting in 100 simulated years, each made up of a random daily sequence of actual price and customer load outcomes that occurred over the period 2014/15 to 2018/19. For each of these 100 simulated years, load and prices are drawn at the same time (i.e. from the same historical day) so that the correlation between load and prices is maintained.

3.3.2 Accounting for trends

The 100 simulated years that result from the standard Monte Carlo process described above do not account for the ongoing effect of any expected trends in customer load or prices over the period 2019/20 to 2023/24. Given that we are developing a heuristic which is robust enough to be used in future years, it is important to consider these trends. To account for these trends we make some adjustments to the standard Monte Carlo process, as described below.

Accounting for trends in customer load

As discussed, the relative reductions in load during the day that we saw in **Figure 2** are consistent with increasing amounts of rooftop solar PV from year to year. We expect that the installation of rooftop solar PV will continue in coming years, and that this trend towards relative reductions in load during the day will continue. To account for this expected change in load, we adjust the half-hourly load profile.

The adjustment we make to the half-hourly load profile is an adjustment to the five years of historical data that is an input into the Monte Carlo process. We make an adjustment to the inputs to the Monte Carlo process rather than the outputs of the Monte Carlo process because we do not want to reduce the variability in the customer load profiles for the 100 simulated years that are an output of the Monte Carlo process.

The adjustment that we make to the five years of historical data is an adjustment to account for the impact of additional rooftop solar PV panels that are expected to be installed between the historical year and the forecast year. For instance, when we use historical data from 2014/15 as a basis for estimating load in 2019/20 we would like to account for the impact of 5 extra years of rooftop solar PV panels; when we use historical data from 2015/16 as a basis for estimating load in 2019/20 we would like to account for the impact of 4 extra years of rooftop solar PV panels; and so on.

The adjustment that we make to each year of historical load data is based on the trend in the historical load factor that we observe in **Figure 1**. One of the things that drives this reduction in load factor seems to be the installation of rooftop solar PV panels. However, there are also likely to be other changes that affect load factor; for instance, more extreme weather, increased installation of air conditioning and a greater mix of working households would also be likely to drive a reduction in load factor. We are not able to disentangle the various factors that drive a reduction in load factor, but assume that *half* of the historical reduction in load factor is due to the increase in PV generation, with the remainder due to other factors. While we cannot be sure that half of the historical reduction in load factor is due to the increase in PV generation, we find that the results we get by using this assumption are plausible.

The assumption that half of the historical reduction in load factor is due to an increase in PV generation provides us with a target load factor for each forecast year. For instance, when we use historical data from 2014/15 as a basis for estimating load in 2019/20 we adjust the half-hourly load from 2014/15 so that the load factor reduces by an amount that reflects five years of half the historical trend reduction. We make the adjustment to the half-hourly load (and consequently to the load factor) by deducting a half-hourly solar generation shape from the half-hourly consumption load until the target load factor is achieved. We do this for each of the historical years, to account for the difference in solar PV installations between each historical year and 2019/20.

Of course, one of the implications of applying this trend approach is that we expect the load factor to decrease from 2019/20 to 2020/21, and so on. This means that when adopting the Monte Carlo process for each of the forecast years (2019/20 to 2023/24) we need to make the appropriate adjustments to half-hourly load input into the Monte Carlo process. So, for instance, while we account for five years of trend reduction in load factor due to solar installations when inputting the historical 2014/15 data into the Monte Carlo analysis for 2019/20, we account for six years of trend when inputting the historical 2014/15 data into the Monte Carlo analysis for 2020/21, and so on.

This means that we are running the Monte Carlo process described above for each forecast year. The result is 100 simulated years for each of the 5 forecast years, or 500 simulated years in total. The way that we use these simulated years to determine the heuristic is described in Section 5.

Accounting for future spot prices

We also make a further adjustment to the half-hourly spot prices. We consider that historical half-hourly spot prices provide the best source of information about patterns of half-hourly spot prices and how these are correlated with half-hourly load, but historical average spot prices are not necessarily a good predictor of the future average level of spot prices at the NSW regional reference node. There is no reason, for instance, that NSW spot prices during 2014/15 will, on average, be the same as average NSW spot prices for 2020/21.

In our view, the best available public information about the average level of NSW spot prices for 2020/21 is the contract prices published by ASXEnergy. These contract prices – particularly the prices of base swaps – provide the market's view on what will be the average spot price for 2020/21. Given this, for each simulated year, we assume that the average level of prices is consistent with the relevant ASXEnergy futures prices. Specifically, for each simulated year we scale the half-hourly prices so that the time-weighted average price in each quarter is equal to the relevant quarterly base swap prices for the forecast year from ASXEnergy[2] (less an assumed contract premium of 5 per cent on the underlying prices, consistent with that used by the ICRC in its decision in its recent Electricity Model and Methodology Review). The ICRC has asked us to use the 23-month average of ASXEnergy contract prices for quarterly base swap prices (up to 17 January 2020) to scale spot prices for each quarter of the five forecast years. This is consistent with the time period used by the ICRC in its decision in its recent Electricity Model and Methodology Review.

Because contracts for 2022/23 and 2023/24 are not currently trading we are not able to rely on ASXEnergy contract prices for these years; instead we assume that prices for 2022/23 and 2023/24 will be the same as prices in 2021/22, which is the last year for which we have reliable price data from ASXEnergy.

This approach to generating half-hourly price forecasts results in:

- The appropriate average *level* of spot prices (i.e. the time-weighted quarterly average price is consistent with the 23-month average ASXEnergy prices that the ICRC uses).
- The appropriate *half-hourly profile* of spot prices (i.e. the half-hourly profile of prices, and load, are consistent with historical outcomes).

[2] An alternative approach would be to attempt to scale half-hourly prices having regard to each of base swaps, peak swaps and cap prices. However, the scaling process would require subjective judgements about how to simultaneously scale to each of these prices. Given there would be little on which to base these subjective judgements our preference is to scale only to base swap prices, which is a mechanical process.

4 CONTRACT PRICES

This section addresses the third question we need to answer to estimate a contract position:

- What is the cost of financial hedging contracts?

As discussed, our approach to estimating the contract position that retailers use to hedge is based on a hedging position that a prudent retailer would face in supplying electricity to their customers, having regard to the hedging contracts that a prudent retailer is likely to enter into. The hedging contracts that we base this analysis on are ASXEnergy contracts. There are three main types of electricity contracts that are traded on ASXEnergy:

- Base swaps for each quarter.
- Peak swaps for each quarter.
- Base $300 caps for each quarter.

These contracts trade for a number of years in advance. Prices are published by ASXEnergy for each contract for each trading day.

ASXEnergy contract prices are shown in **Table 1**, for the 23-month time weighted average price, up to 17 January 2020.

Table 1: 23-month time weighted average[3] ASXEnergy base swaps for NSW

	FINANCIAL YEAR	QUARTER			
		Q3	Q4	Q1	Q2
TIME WEIGHTED	2020	$70.78	$68.22	$93.50	$74.05
	2021	$68.19	$67.61	$82.10	$66.58
	2022	$70.26	$70.56	$77.56	$69.80
	2023	$70.26	$70.56	$77.56	$69.80
	2024	$70.26	$70.56	$77.56	$69.80

Source: Frontier Economics analysis of ASXEnergy data

[3] Because contract prices from each day are given equal weight, the time weighted average in this context is the same as a simple average. This contrasts to a trade weighted average, where the contract prices for each day are weighted by the number of trades on that day.

5 CONTRACT POSITION

This section addresses the final question we need to answer to determine a contract position heuristic:

- What kind of hedging position is a prudent retailer likely to adopt?

We use our portfolio optimisation model – *STRIKE* – to determine the efficient mix of hedging products that a prudent retailer would likely adopt. *STRIKE* calculates an efficient frontier, which represents the contracting positions that provide the lowest energy purchase cost for a given level of risk (where risk is measured by the standard deviation of the energy purchase cost).

STRIKE applies a Minimum Variance Portfolio (MVP) approach to the task of hedging a retailer's exposure to wholesale spot prices. An MVP approach seeks to identify a group of assets that provide the lowest possible risk for a given rate of expected return. In this context, an MVP approach seeks to identify the contract position for a retailer supplying the customer load that provides the lowest possible risk for a given energy purchase cost.

STRIKE applies the MVP approach as follows. The model incorporates an estimate of a retailer's exposure to the wholesale spot market, which is determined by the retailer's load and by wholesale spot prices. There is an expected return and a variance associated with this. *STRIKE* also incorporates the types of hedging products that are typical in the electricity industry. These contracts – swaps and caps – generate cashflows that also have an expected return and a variance. Instead of assessing the expected return and associated risk for each asset in isolation, *STRIKE* applies the concepts of portfolio theory to evaluate the contribution of each asset to the risk of the portfolio as a whole. Based on this approach, *STRIKE* calculates efficient hedging strategies.

In order to determine a hedging position for a retailer in the ACT, we make use of the following inputs in *STRIKE*:

- Forecast spot prices and load, as discussed in Section 3. As we discussed, we have developed 100 simulated years of half-hourly spot prices and load for each forecast year from 2019/20 to 2023/24. Our view is that an efficient retailer's hedging position should have regard to the uncertainty associated with what kind of year 2019/20 will be. For example, will 2019/20 be a year with high prices and high load corresponding, so that the load-weighted price is high? Or will 2019/20 be a year with low prices and high load corresponding, so that the load-weighted price is low? The lowest risk hedging position will differ for these two years. To account for this uncertainty, we input 7 simulated years into *STRIKE* for each forecast year, for a total of 35 years. So, for instance, for 2019/20 we input into *STRIKE* the 7 simulated years from our Monte Carlo simulation that represent the 99th, 95th, 75th, 50th, 25th, 5th and 1st percentile, when the 100 simulated years are ranked according to load-weighted price. We also input into *STRIKE* the 7 simulated years representing these same percentiles for each of the other forecast years. Because this full set of 35 years is input into *STRIKE* at the same time, this means that the resulting hedging position provides the lowest risk across the combination of simulated years with low and high load-weighted prices. The resulting hedging position may not be the one that provides the lowest risk for any single simulated year, but is the one that provides the lowest risk across all the 35 simulated years included. In other words, it is the hedging position that provides the lowest risk accounting for the fact that retailers do not know in advance what kind of year they will be facing.
- Contract prices, as discussed in Section 4.

As discussed, *STRIKE* calculates an efficient frontier, which represents the various different contracting positions that provide the lowest energy purchase cost for a given level of risk. Contract positions that have a higher risk tend to have a lower energy purchase cost, and contract positions that have a lower risk tend to have a higher energy purchase cost. The contract position that we think a prudent retailer would aim for is based on the most conservative contracting position on the efficient frontier, which is the point on the efficient frontier with the lowest risk (but highest cost).

Outlined in **Figure 5** are the resulting contract positions at the conservative point having regard to all forecast years 2019/20 to 2023/24. For each quarter (the vertical panels) and each peak/off-peak period (the horizontal panels), the charts show the following:

- The distribution of half-hourly load for the 48 half-hours of the day (shown by the box plots in the 'Load' panel).
- The distribution of half-hourly spot prices for the 48 half-hours of the day (shown by the box plots in the 'Spot price' panel). The half-hourly spot price can be as high as the Market Price Cap of $14,700. In order that daily patterns of prices can be distinguished, the price chart is truncated at a spot price of $750/MWh. We use the full set of spot prices (including prices above $750/MWh) in our analysis.
- The quantity of swaps and caps at the conservative point of the efficient frontier (shown by the coloured areas in the 'Load' panel).

It should also be noted that the conservative point on the efficient frontier reflects the contract position that achieves the lowest risk for the 35 simulated years that are input into *STRIKE*. In the event that different simulated years were input into *STRIKE*, the model would find a different contract position that achieves the lowest risk. In particular, if it were assumed, for instance, that next year will have an unusually large number of very high price events that coincided with high load, then the model would certainly find a different contract position that achieves the lowest risk. That load forecasts and price forecasts (and their correlation) are important to the costs that retailers face in supplying regulated customers is why we use the best available information to develop load forecasts and price forecasts that are consistent with the observed peakiness of historic load and historic prices (and the observed correlation between them).

Figure 5: Contracting position compared to load and prices

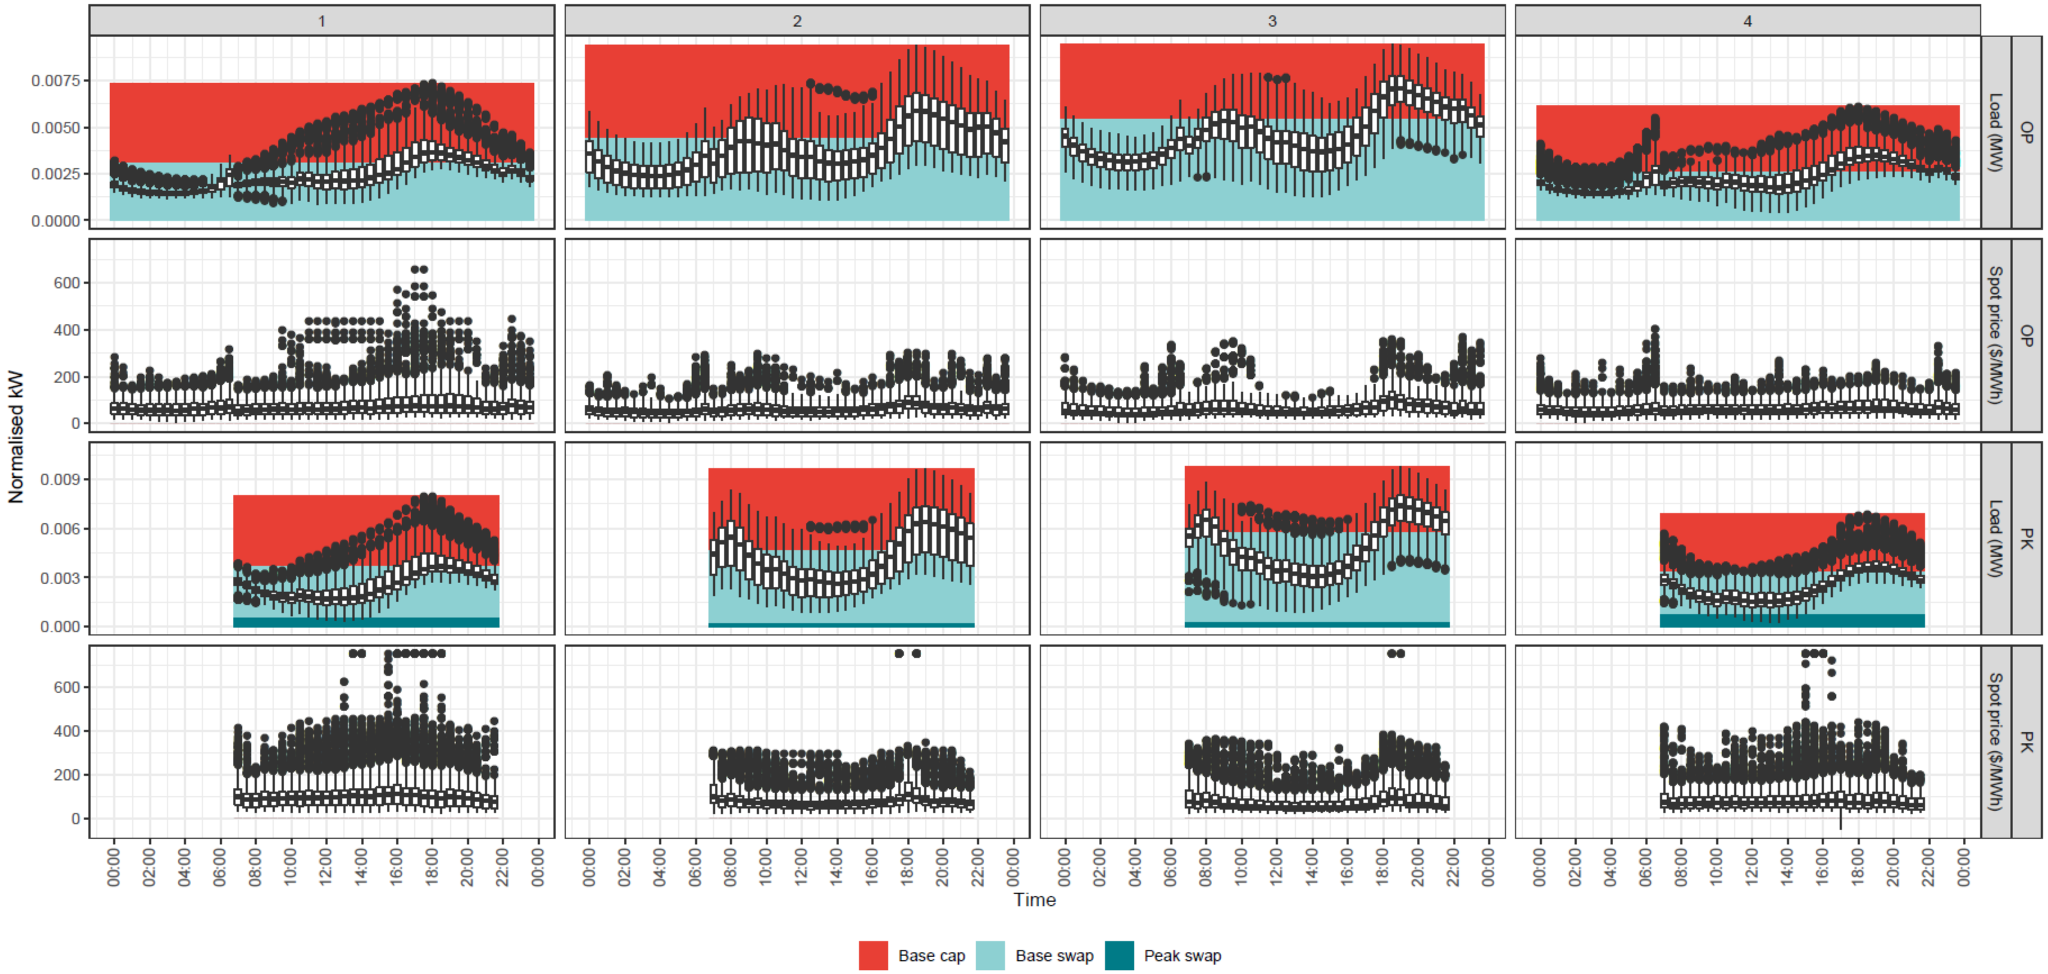

Source: Frontier Economics

6 CONTRACTING HEURISTIC

Based on the data discussed in Section 3 through Section 5, this section reports the contracting heuristic that we have estimated.

6.1 Contracting heuristic

We determine a contracting heuristic by calculating the volume of base swaps, peak swaps and caps by quarter, expressed in relation to load.

The contracting heuristic values are presented in **Table 2**. The volume of base swaps is expressed as a percentile of load for all the half-hourly intervals in the quarter. The volume of peak swaps is expressed as a percentile of load for all the *peak period* half-hourly intervals in the quarter, less the volume of base swaps. The volume of caps is expressed as a percentage of load in the highest demand half-hourly interval in the quarter, less the volume of base and peak swaps.

For example, in Quarter 3:

- The base swap contract volume is set to equal the 70^{th} percentile of half hourly load in Quarter 3.
- The peak swap contract volume is set to equal the 65^{th} percentile of half hourly load in peak periods in Quarter 3, less the base contract volume for Quarter 3.
- The cap contract volume is set to equal the peak demand for Quarter 3, less the base contract volume and peak contract volume for Quarter 3.

These contracting volumes change by quarter due to the differing shapes for the load profile, spot prices and their correlation.

Table 2: Contract level percentiles

QUARTER	BASE SWAP VOLUME, EXPRESSED AS A PERCENTILE OF HALF-HOURLY LOAD IN THE QUARTER	PEAK SWAP VOLUME, EXPRESED AS A PERCENTILE OF HALF-HOURLY LOAD IN PEAK PERIODS IN THE QUARTER LESS BASE SWAP VOLUME	CAP VOLUME, EXPRESSED AS A PERCENTAGE OF THE PEAK HALF-HOURLY DEMAND IN THE QUARTER LESS BASE AND PEAK SWAP VOLUMES
3	70th	65th	100%
4	65th	80th	100%
1	80th	85th	100%
2	70th	65th	100%

Source: Frontier Economics.

frontier economics

BRISBANE | MELBOURNE | SINGAPORE | SYDNEY

Frontier Economics Pty Ltd
395 Collins Street Melbourne Victoria 3000

Tel: +61 (0)3 9620 4488
www.frontier-economics.com.au

ACN: 087 553 124 ABN: 13 087 553