IES INSIDER

ISSUE 31 March 2018

Reducing Artificial Price Volatility in the NEM

Hugh Bannister: CEO, IES

Reducing Price Volatility as a NEM Goal

In the early 1800s the steam engine was in its heyday when a French military engineer named Sadi Carnot turned his mind to the fundamental nature of a heat engine. He was influenced by his father who was much immersed in the design of water mills. Carnot senior had the (correct) intuition that, to maximise a water mill's power output, the flow through it should be arranged to minimise any shocks in the flow of water through the system.

Sadi Carnot applied this line of thinking to an abstract heat engine. He showed that, if the process was reversible (shock free), the efficiency of an engine did not depend on the nature of the working fluid; only on the temperatures of the high and low heat reservoirs between which the engine operated. Further, this was the highest efficiency achievable for a heat engine. Although not fully recognised until later, Carnot's startling analysis is now a fundamental building block of engineering, physics and the nature of matter, through the derived (and somewhat mysterious!) concept of entropy.

As with water mills and heat engines, we can also assert that a complex economic system such as an electricity market ought to be designed to be as shock-free as possible for maximum efficiency. In this article I will ty to put some flesh on the bones of this idea.

Real and Artificial Price Volatility

The NEM is one of the few energy-only markets in the world. Effective operation and investment depend on spot prices being able to reach high levels when supply is genuinely constrained. This results in the price volatility for which the NEM is well known and which is a fundamental part of its operation. Contracting can limit short-term financial exposure to price volatility and also, potentially at least, drive investment.

Price volatility can also be the unintended consequence of specific NEM rules. A well-known example addressed at length in a number of recent AEMC reviews and rule changes, is the averaging of 5 minute prices for settlement purposes. Incentives to game this logic by driving price spikes late in the half hour led to a "good faith bidding" rule change several years ago. More recently, in 2017 the AEMC made a rule change to implement 5-minute settlement, to be phased in over a period of years, which may reduce such artificial volatility but not eliminate it, as I will argue in this article.

In this article I will review two other examples of artificial shocks promoted by the current NEM Rules and propose fixes for them. They are:

- The offer structures in the NEM as they affect AEMO's ability to forecast demand and set prices robustly.
- The arrangement under 5-minute settlement in the NEM whereby settlement is based on a constant price over a settlement interval, with a step change in price between settlement intervals.

Forecasting 5-minute Demand and Price

We need to distinguish variability in forecast demand outcomes from variability in price outcomes. Clearly, the two are related but, when demand is tight, the nature of the NEM is to skew price variation very much to the high side. One such example is the approximately 1.5 hour period on 18 November, 2016 illustrated in Figure 1 following.

Figure 1: Unstable Prices in the NEM

Over the period there were four price spikes reaching near to the price cap, although at present they are smoothed by the averaging over half an hour for settlement, as shown by the green line. One cannot discern much response by non-scheduled load in the figure but, with a clearer price signal under 5-minute settlement, much more response and price volatility is likely.

While occasionally high spot prices are a necessary feature of the NEM, the volatility evident in this case could be the result of difficulties with AEMO's demand forecasting.

So could better demand forecasting stabilise prices and the resulting pattern of demand response? One option is to take a probabilistic view by recognising likely variation and to seek an unbiased estimate[1]. But should we seek to eliminate bias in demand or in price? The two are quite different because of the skewness in price, especially when supply is tight. Eliminating price bias seems to make more sense, but would that be AEMO playing in the market? That's supposed to be a strict no-no, although AEMO's 5-minute forecasting role is already influencing the market.

Let's take a close look at the situation from first principles.

Figure 2 following is a simplified version of system offers stacked up to meet the demand forecast, two cases of which are illustrated by the vertical lines. It is easy to see that, when the forecast is near a price change boundary, a small change in demand can change the price in a stepwise fashion. Forecast error becomes critical, even when that error is as small as it can be. When supply is tight and the system is using the higher price bands, this step change can be very large – of the order of the price cap. This is what we can observe in Figure 1.

Figure 2: Offer Stack with Inelastic Demand

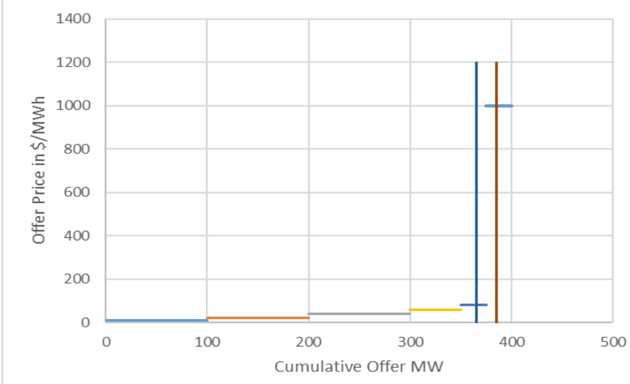

Note that this volatile price outcome is independent of whether or not there is any non-scheduled demand response to price. If we include some price sensitivity, which is likely to increase going forward (to the benefit of the market if well managed), we can analyse the possible outcomes with the aid of Figure 3 below.

Figure 3: Offer Stack with Elastic Demand

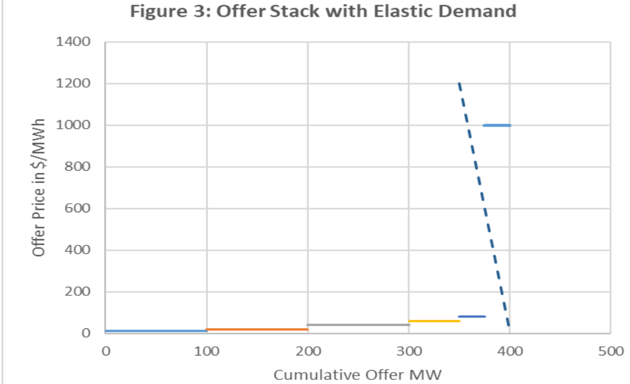

The backward sloping dashed line is the actual demand curve, showing some price sensitivity. We are trying to make a demand forecast that reflects the actual ex post demand.

The figure makes it clear that NO stable forecast is possible! If the forecast is fixed, we get one or other of the two

[1] Unbiased means that the average of many samples approaches the actual average

pricing cases shown in Figure 2, at prices of $80/MWh and $1,000/MWh in this simplified example. If we happen to hit the lower price, the demand outcome will be higher than forecast. If we hit the higher price, the demand outcome will be lower, as can be read for Figure 3. This is a recipe for price and demand instability, likely to be exacerbated when 5-minute prices rather than prices smeared over half an hour are used for settlement.

Under the current NEM offer arrangements, the problem is that only a discrete set of price outcomes is possible. These are unlikely to coincide with an equilibrium load. Generally, the mismatch passes unnoticed in the general system noise, but it can become visible at times of system stress and will become more so under 5-minute settlement.

Improving 5-minute Forecasting

A rule change proposal to require currently unscheduled loads to make offers into central dispatch was rightly rejected by AEMC in September 2017, but for AEMC to then suggest that AEMO could improve its demand forecasting by improving its neural network forecasting model[2] was disingenuous, given that an equilibrium forecast does not in general exist under the current NEM rules!

So what can be done to improve 5-minute forecasting? The first task is to ensure that an equilibrium forecast is actually possible, which means ensuring that an equilibrium price can be produced by the system. How can this be done?

One way is to attempt to model the demand curve shown in Figure 3. To do this using current modelling technology, one would model demand sensitivity with a series of blocks at different prices, analogously to the current offer blocks of generation and scheduled load. In other words, we could attempt to model the backward sloping demand curve of Figure 3 with a stepwise function. To cover the prices one would wish to have available, we would need many blocks, potentially of the order of 100 if the price interval is to be restricted to, say, $100/MWh.

Critics of this approach worry that it would requires AEMO to play in the market. Of course, AEMO plays in the market already by the very act of forecasting demand, but modelling price responsiveness (elasticity) would be a step beyond that.

Another occasionally suggested change that avoids this concern is to support more offer blocks for participants than the current ten per physical unit. Because the number of possible prices would remain finite, this is only a part solution. Even then, many bands would be required to narrow the gap between prices to a reasonable level, an inconvenience in a trading environment.

Yet another option is illustrated in Figure 4 below. We recognise that piecewise constant marginal costs are somewhat artificial, so we offer participants an option for price bands that they nominate to be interpreted slightly differently.

Figure 4: Modified Offer Stack with Elastic Demand

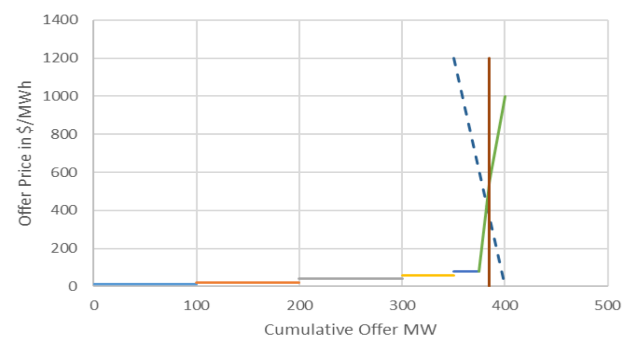

Figure 4 shows the last price band interpreted as a continuous range of prices moving linearly from the previous band's nominated price of $80/MWh to the currently nominated (last) price band of $1,000/MWh. This is shown as the green line sloping up to the right. If we now forecast demand as the vertical red line, we can solve for a specific price. If we are good enough, we can forecast the actual demand, shown where the offer price intersects with the dashed demand curve. By allowing offer prices to be continuous rather than discrete, we have solved the existence problem and can aspire to a more accurate forecast.

Technically, interpreting some or all offers as prices linearly ramping prices changes the bid functions in the objective of AEMO's dispatch and pricing engine to a quadratic form rather than a simple linear form. This is standard

[2] https://www.aemc.gov.au/sites/default/files/content/0bcaf68c-8449-4ce0-aaa6-da223ca6e01c/Final-Determination-ERC0203-Non-scheduled-generation-and-load.pdf Section 7.6

optimisation problem and relatively trivial to implement. AEMO's NEMDE developer could do it within a day, but in the real world it might take a little longer to justify, document and test.

From the participant perspective, an option for an offer band to have continuously varying prices gives greater flexibility. Participant system changes could be relatively modest. The MW and price offer structure would remain the same – only the interpretation by AEMO's dispatch and pricing engine would differ according to settings or defaults that a participant could nominate. e.g. all bid bands continuous; only the last one continuous. There would be more marginal offers in the final schedule than at present; not a major change.

With this approach, achieving some consistency and stability in AEMO's demand forecasting and pricing at the "top end" becomes a feasible goal, even under 5-minute settlement when large price changes are likely to drive increased unscheduled demand response.

Eliminating Stepwise Spot Price Changes

While we can reduce artificial volatility arising from current limitations on the structure of offers, as previously described, shocks to the system from stepwise spot price changes would remain from the settlement process. This is illustrated in Figure 5 below.

Figure 5: Linear Dispatch and Stepped Prices in the NEM

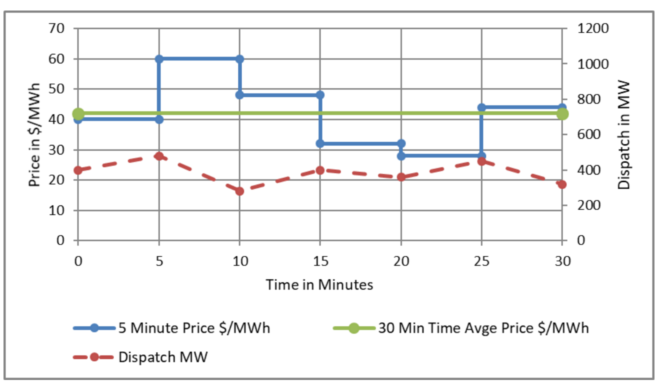

5-minute dispatch prices are show in blue and the scheduled trajectory of a scheduled generator is shown in as the dashed red line. A present, it is assumed that price is constant over the dispatch interval and set 5 minutes ahead (ex ante), but averaged over the half hour for settlement as shown by the flat green line. Under 5-minute settlement, the blue prices will be used to value the energy metered within the same 5 minutes; the price signal will be much clearer than at present. Further, the stepwise price transitions can be very large; of the order of the price cap and much more than in the example shown.

With improved communications and fast acting options such as batteries becoming more prevalent, we can expect to see many more unscheduled loads switching at price boundaries. Such generally destabilising behaviour would require correction with additional regulation FCAS or, worse, with restrictive operational rules.

How should this issue be managed? One way is to impose a rule that loads cannot ramp more than 20% of their capability in any given minute, as suggested by AEMO in a submission to the 5-minute settlement consultation. Another approach is to take the AEMC stance; yes, it's a risk but we will deal with it under separate system security/ frequency control reviews.

In my view, such approaches would be heading in the wrong direction. Much better is to fix the problem at source. A moment's reflection reveals the source of the problem is the artificial step change in price at the dispatch interval boundary. Costs and prices do not in reality change in such steps, they should move smoothly (say linearly) from the current level to the level set 5 minutes ahead. This is illustrated in Figure 6 below.

Figure 6: Linear Dispatch and Linear Prices in the NEM

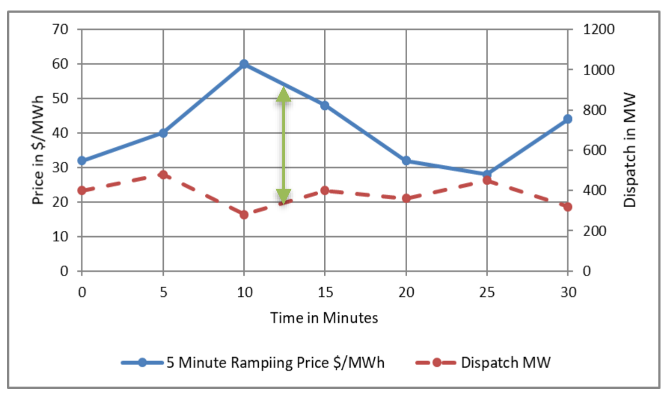

If a unit follows its linear schedule as shown by the red dashed lines, the settlement amount in each 5 minutes is the product of the price and MW (as illustrated by the vertical green arrow) summed/integrated over each 5 minutes.

A unit or load will not precisely follow a linear trajectory within each 5 minutes, so revenue metering must provide measurements that will support accurate settlement. I will show how this can be done, but it's useful first to review the settlement amounts in the idealised case where a linear schedule is followed.

Let x1, p1 be the MW and price at the start of the interval.

Let x2, p2 be the MW and price at the end of the interval.

Let T be the time interval in hours (1/12).

Under proposed arrangements for 5-minute settlement, the idealised settlement amount will be:

Settlement amount = p2 * T * (x1 + x2) / 2 (1)

That is, the 5-minute ex ante dispatch price p2 applies to the energy in the linear MW profile, which is T*(x1+x2)/2.

Under a linear ramped price settlement regime and in the idealised case of a linear MW trajectory, a little high school calculus yields the following idealised settlement amount:

Settlement amount =

T * (p1*x1/3 + p2*x2/3 + p1*x2/6 + p2*x1/6) (2)

This formula has a pleasing symmetry, even though it appears more complex than (1). The average interval price is now (p1 + p2)/2 instead of p2. However, the settlement amount is NOT in general the average price multiplied by the total energy due to correlation between price and load within the interval. The difference can be regarded as a ramping payment/charge.

Over a long period, the cumulative outcomes from (1) and (2) would not likely differ greatly. These assertions can be tested with spreadsheet examples. However, eliminating artificial price shocks removes the temptation for bad behaviour at dispatch boundaries.

In practice, settlement must be based on actual rather than the scheduled pattern of generation or load. How do we meter this? With programmable electronic metering we can meter for a ramped price relatively easily.

An electronic meter takes instantaneous readings of voltage and current from each phase of an instrument transformer at very small intervals –several hundred times for each 50 Hz cycle. The firmware multiplies these two values and accumulates/averages them over the desired interval, in our case 5 minutes. This gives an accurate measure of the average power/energy metered over the 5 minutes. This value is stored and uploaded later for settlement, where the appropriate settlement price is applied to each interval of metered energy. To put this into a simple formula:

Let i be the i^{th} of n readings (n-1 intervals between them).

Let x(i) be the instantaneous MW (voltage times current per phase);

The following quantity is accumulated in the meter[3].

MW_flat = [sum_over_i (x(i))] / (n -1) (3)

Then the settlement amount under proposed 5-minute settlement is:

Settlement amount = T * p2 * MW_flat.

To provide for ramped price settlement, we need to accumulate and average an additional value, which is the measured instantaneous power weighted by its fraction along the dispatch interval. As a simple formula, this is:

MW_ramp = [sum_over_i (alpha(i) * x(i))] / (n-1) (4)

We can show that the ramped price settlement amount is:

Settlement amount =

T * (p1*MW_flat + (p2–p1) * MW_ramp) (5)

Note that the change to the meter firmware program is relatively simple; accumulate, store and upload the additional quantity defined by (4). The firmware in most modern meters could easily be re-programmed to do this.

Note also from (5) that the meter does not need to know about the actual dispatch prices to be used for settlement. We just need to know MW_flat and MW_ramp for each interval; prices are applied only at settlement time.

The logic above can be tested with spreadsheet examples.

Of course such a change would affect not only AEMO's settlement logic but also participant systems. These are in the process of being re-vamped to accommodate 5-minute settlement. The improved settlement logic just described

[3] In practice, the needs to be a small adjustment at the boundary of the dispatch/settlement period. This adjustment is omitted here for simplicity.

could be incorporated along with the much larger 5-minute settlement changes.

Conclusions

A casual reader of this article might well ask - what's wrong with a bit of price volatility? Surely these are solutions looking for a problem!

If loads remain passive in the face of artificially volatile prices, we could no doubt get by, much as we have done until now. But if loads become more price sensitive, as all pointers suggest they will, there will be much spurious load management activity that will need additional regulation FCAS to correct. Regulation FCAS is no longer cheap and won't get any cheaper, so it makes sense to reduce the burden on regulation FCAS where practical to do so.

An even worse outcome is that AEMO/AEMC put distributed resources into operational straightjackets to minimise these stability problems or, worse again, arrange to have distributors and large retailers take over the function of managing them. These are stark choices.

Perhaps we still have much to learn from Sadi Carnot and his dad; get rid of the shocks and you can hope to approach a high level of efficiency.

Postscript

For PROPHET users, we will implement an option to model continuous offers so you can play with it and understand what this option may offer you.

From time to time we will also publish on our website or in NEO some analysis of how a modified settlement logic as described in this article would work.

For those interested to follow up these ideas, you are welcome to contact us using the details following.

CONTACT

Hugh Bannister	CEO, IES
+61 (0)2 8622 2210	+61 (0)411 408 086
hbannister@iesys.com	www.iesys.com