

Electricity Derivatives Benchmark Pricing via Simulation

Version: 1.0.2

Trading Technology Australia

Thursday 24th December, 2009

1 Overview

One of the problems of the electricity industry is that of data. In one sense, we have too much data, with half-hourly (or five-minute) prices, intra-day profiles and strong seasonality. On the other hand, we have too little data and it is not always possible to find liquid derivative prices in the market for some instruments.

A related problem can be the availability of a reliable pricing model for pricing electricity derivatives. Such a model is generally only as good as its calibration, and many industry participants do not have the budget to install, understand and continue to calibrate complex and expensive software to price derivatives and produce risk reports.

Trading Technology Australia (TTA) has committed considerable resources to the research of numerical simulation engines for the forecasting and simulation of spot electricity price paths. Our most recent results in this area were released in the document 'Beacon Price / Load Forecasting and Simulation Engine - Product Overview' (February 2009)¹.

The next step was to mark the expected spot price curve to futures contract levels, and since July 2007 we have offered the TTA / d-cyphaTrade Mark to Futures Forward Curve² on a subscription basis.

Our next innovation has been to calibrate the simulation engine to the forward market, taking us much further than curve levels and giving us the distributions around the expected spot prices and the expected prices for various electricity derivatives. For cap pricing, for example, we require the correct distributions of possible spot prices around the expected mean, while for a quarterly swaption price we require the distributions of the average quarterly settlement price.

The approach we have taken is to calibrate the engine using only historically observed data and market traded futures settlement prices. From here we imply all other prices and price distributions. For this approach to work, we make the following key assumption:

That the market prices forward derivatives in line with the observed behaviour of historical spot prices in the NEM. Such behaviour includes:

- *The intra-day shapes and peak / off peak parities;*
- *The frequency of observed price spikes;*
- *The distributions of spot prices for each time period and day type; and*

¹Available at www.tta.com.au/wp-content/uploads/2009/03/beaconoverview_20090205.pdf

²See www.d-cyphatrade.com.au/products/data_centre/packages

- *The variance of average prices over a period of time (e.g. quarterly peak).*

In this paper we will attempt to prove that the market, in general, prices derivatives based on historically-observable price behaviour (albeit at forward-looking levels). If this can be shown, then we would argue that to use such an approach for ‘benchmark pricing’ and risk reporting is the most reliable and accurate technique currently available to Australian energy market participants.

1.1 Benchmark Pricing

By benchmark price, we mean the calculated price of a derivative that is:

- Transparent in its methodology, independently generated and widely available to the market;
- Based on actual traded market prices; and
- Is consistent with prices for similar derivatives;

Where a derivative is traded on an exchange, the benchmark price would be the settlement price of the security, however not all instruments are exchange traded, leaving us with only broker-sourced data (the quality of which depends on the volumes traded by that broker) or other, usually averaged and non-traded data sources (such as AFMA data).

The benchmark price may not match any particular quoted price, as there are other factors involved such as who is quoting and what their motivations are. It also may not match a price calculated by any participant’s internal system as they would most likely be using a different methodology.

However in the situation where the holder of the instrument had to mark that instrument to market, the benchmark would be the closest thing available to a market price, as it would be independently generated (or at least generated using a methodology from an independent party).

2 \$300 Cap Futures

We have calibrated our simulation engine to the market using only the peak and flat quarterly futures contracts. This calibration ensures that the levels of the price paths, on average, will match the levels of these contracts.

Cap pricing, on the other hand, does not depend on price levels, but rather the frequency and severity of price ‘spikes’. For a \$300 cap, a spike can be thought of as any price greater than \$300. By pricing a cap with Beacon, we can check whether it has the correct frequency and extent of price spikes. Beacon will simulate spikes along various paths based on the historically observed probability of such spikes. Unless there is additional information to suggest otherwise (for example, impending drought conditions), it is reasonable to use the historical frequencies.

The methodology we have used here involves calibrating the engine on a monthly basis, from December 2008 to October 2009 and then re-pricing the four 2010 quarterly flat cap futures contracts. As we are in a backtest situation, we have the traded cap futures settlement prices available to us via d-cyphaTrade³.

We are using two sets of inputs:

1. The peak and flat futures contract levels; and
2. The half-hourly simulated spot price paths, which include price spikes.

We are using these two inputs (one observable, one model-based) to calculate the cap price. Because this price is itself observable, we can make a judgement on whether the three sets of information are consistent.

Figure 1 shows the Q1 2010 flat cap prices and compares the calculated prices against the actual settlement prices for each month through 2009. Figure 2 shows the same for the Q2 contract.

³Note that the quoted prices for 18/9/09 have been used in place of prices for 30/9/09 due to availability




Figure 1: Q1 2010 \$300 Caps




Figure 2: Q2 2010 \$300 Caps

As we can see from figures 1 and 2, the calculated price tracks the quoted price (albeit with a certain error margin). The tracking appears better for the Q1 contract. We can thus conclude that the three sets of data (namely futures settlements, historically-observed spike probability and cap futures settlements) are reasonably consistent (for example, there do not appear to be arbitrage opportunities between the futures and caps).

Looking at quarters 3 and 4, we see a slightly different picture, which appears in figures 3 and 4.

In this situation there is not a close tracking between the quoted and calculated cap futures prices. Both the futures and cap futures prices are actual settlements, although the futures contracts are the more liquid of the two. Assuming that the futures prices are reliable, this means that either the prices of Q3 and Q4 cap futures are inconsistent with futures contracts, or that the market is pricing in a different spike frequency than that observed historically.

We can offer three reasons to suggest that the calculated prices are more 'reliable' on a fair-pricing basis:

1. The Q3 and Q4 quotes remain virtually unchanged throughout the year, despite variations in futures prices. This suggests that they are not heavily traded and are therefore not indicative of 'true' prices.
2. The possibility that the market is pricing in a different spike frequency did not appear in the more liquid Q1 contract, nor in the Q2. Unless there was a specific reasoning for expecting more spikes, we would be safer using observed probabilities.
3. The pricing difference decreases as we get closer to the contract start, which is consistent with both the quoted and calculated contract prices converging to settlement (as cap futures contracts ultimately settle off spot).




Figure 3: Q3 2010 \$300 Caps




Figure 4: Q4 2010 \$300 Caps

3 Future Options

Turning to the future options, we use a similar methodology to that of the cap futures. Once again we calibrate the simulation engine on a monthly basis (from December 2008 to October 2009) and then re-price the calendar 2010 future option each time. We again compare against the traded future option settlement prices available to us via d-cyphaTrade.

As with the cap futures, we are using the same two sets of inputs:

1. The peak and flat futures contract levels; and
2. The half-hourly simulated spot price paths, which include price spikes.

Figures 5, 6 and 7 show the performance of the \$35 strike, \$40 strike and \$45 strike call options respectively. Tracking is excellent for the lowest strike but does deteriorate as the strike levels increase.




Figure 5: \$35 Strike Calls




Figure 6: \$40 Strike Calls




Figure 7: \$45 Strike Calls

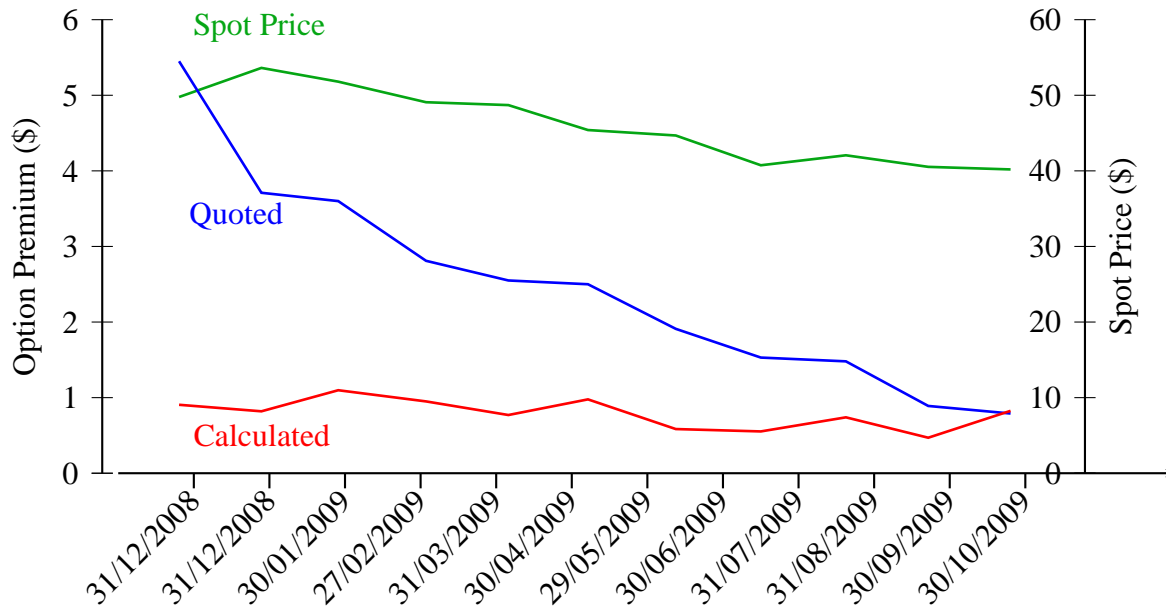


Figure 8: Calendar 2010 option - At the money prices

Figure 8 is a variation on the above. Rather than using a single strike level, in this case we select the closest strike to the current futures price on each observation month. This gives us a series of At-the-Money (ATM) strikes for which we can plot the quoted price against the calculated price.

As the figure shows, the calculated price is initially well off the quoted price, but as time goes by the two converge to the settlement price. As a guide to the price of the underlying, we have also plotted the futures price in green (and on the right-hand-side axis). It is interesting to note that the calculated figure does not vary greatly across the year of observations while the quoted option price comes down dramatically.

What we can say from the 2008 and early 2009 observations of the quoted price is that the option price was trading well above the figure that historical annual volatility (and current futures prices) implied. In this case, it is important to note that in December 2008 (just before the first observation), the federal government released its White Paper on the Carbon Pollution Reduction Scheme (CPRS) and a start date of July 2010 was proposed. Since futures contracts do not pass through the carbon price, this would have increased the 2010 (and later) calendar year futures prices. Given the uncertainty as to whether the scheme would be introduced, market participants preferred options over futures, which pushed up the call option prices. The scheme was deferred in May 2009.

At this level of calibration, we have not changed the simulations engine's volatility at an annual level (although this remains an option). For this reason, we have continued to price a fair value, rather than the artificially high market levels created (presumably) by the proposed CPRS.

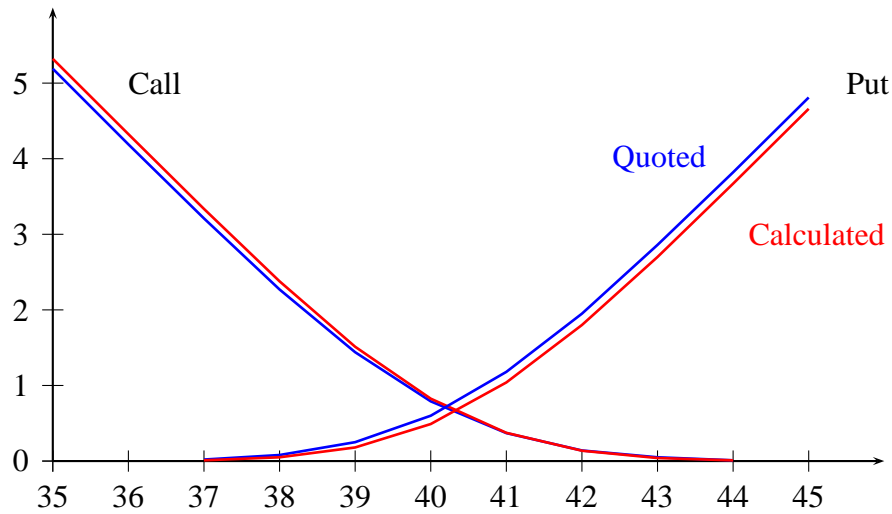


Figure 9: Call and Put option premiums - Quoted vs Calculated - 30th Oct 2009

In figure 9 we have plotted call and put option premiums at various strike levels, both quoted and calculated. The quote date was the 30th October 2009. The figure shows that the simulation engine can approximate all strike levels quoted and this in turn suggests that the options are trading in accordance with historical price volatility on an annual basis.

4 Conclusion

In conclusion we can say that the test results show that the market pricing of electricity derivatives is not out of line with the historically observed price behaviour. Forward-looking information is provided by the futures contracts in the calibration process, as is any expected shifts in peak / off peak relationships.

The cap futures pricing test suggests that our historically-observed spot price volatility (at the half-hourly level) and frequency and extent of price spikes are consistent with the contract prices being offered in the market. Similarly, matching the more recent future options quotes suggests that the volatility of average prices over a period of time in the simulated paths are also consistent with market prices.

While we have used futures, cap futures and future options contracts in this backtest, there is no reason that we could not use OTC swaps, OTC caps and OTC Swaptions in a similar calibration. While there are some limitations to the transferrability of price levels between swaps and futures (due to differences in the contracts such as credit risk, carbon price pass-through and margin accounts), we have seen that once levels are set, the price behaviour of either should be correct, as all contracts eventually settle against the half-hourly spot market.

Once futures to OTC price margining is taken into account, we propose this technique as the most reliable way of pricing the following derivatives:

- OTC swaps with sculpted load and / or price;
- OTC Caps (any strike levels);
- OTC Collars (any strike levels);
- Cap Futures;
- OTC Swaptions;
- Future Options;
- Captions;
- Asian Options; and
- Virtually any derivative contract that settles against the regional spot price and can be priced numerically on an expected return basis.

Since the simulated paths are inherently half-hourly, the derivatives can be priced for the following time periods:

- Peak;
- Base (Flat);
- Off Peak;
- Overnight; or
- Any other custom interval.

Also, instruments could be priced at any frequency including:

- Common frequencies (Monthly, Quarterly, Annual, Financial Year); and
- Custom and odd-dated frequencies (start / end date).

With the addition of proper volatility scaling (in order to convert market-implied option volatilities to 10, 30 or 90 day volatilities), we can perform an additional calibration on the engine that will allow risk reporting on portfolios including:

- Value at Risk;
- Earnings at Risk;
- Cashflow at Risk; and
- Potential Credit Exposure (incorporating credit ratings).

For further information on this topic, please contact:

Joe Maisano
Director, Energy Markets
Trading Technology Australia Pty Ltd (TTA)
Email: joe@tta.com.au
Direct: +61 2 8281 5304