Review of modelling work: NEM Rule Changes and the Snowy Region

Professor Richard Green

Executive Summary

The Australian Energy Market Commission has appointed me to advise on:

- The generic modelling approaches that could have been adopted in seeking to understand and assess the economic impacts of the different Rule change proposals [with respect to network congestion in the Snowy Region];

- The appropriateness and robustness of the modelling approach adopted by the Commission and its appointed economic modelling consultants in terms of:
  - Theoretical underpinnings; and
  - Practical application.

Two generators within the Snowy Region currently face different operating conditions, due to transmission constraints, but would receive the same price, were it not for a derogation from the market rules. Three proposed rule changes would make this derogation permanent, create a second zone, or place each generator within a (different) larger zone – the base case is to end the derogation. Modelling work has been commissioned to assess the economic impacts of these different changes (or non-change), both by the Commission and by a group of generators proposing one of the rule changes.

I have read background documents on the Snowy Region rule changes, the Commission's Draft Rule Determination of 19 January 2007, including the modelling appendix by Frontier, and the draft of the report to be released with the Commission's Final Rule Determination. Frontier have answered a number of my questions via conference calls with Commission staff or via email. My review of ROAM Consulting's work has been based on reading their report of 3 April 2007, and their responses (informed by a teleconference) to questions posed by the Commission. I presented a draft report to the Commission in August 2007, and received helpful feedback. The report presented here remains my own work and responsibility, however.

Questions asked by the Commission

The Commission asked me to answer several detailed questions about the work performed by its consultants, Frontier Economics (Frontier), and about the work performed by Roam on behalf of the group of Southern Generators:

1. What are the generic economic modelling techniques that could have been adopted in seeking to understand what impact implementing one of the relevant Rule change proposals might have on pursing the NEM objective?

The system of electricity generators and transmission lines is a complex one. Except in special cases, it will not be possible to predict the impact of a change in the market rules by logical reasoning. A numerical simulation must be used instead.

There are five key components to such a simulation.

a) The timescale for the simulation, including the periods to be studied, and whether the model runs for every hour in chronological order, or a sample of representative hours
b) The characteristics (capacity, operating costs, ownership) of each generator on the system, and the levels of demand at different times over the period of interest
c) A description of the transmission system, including the locations of generators and loads, equations to determine the flow across each line (or boundary) likely to be congested, and to determine the level of transmission losses
d) A description of the market rules, that takes bids from generators (and loads, if applicable), and calculates the operating patterns, costs, prices and hence profits that will result
e) An assumption about each market participant's behaviour – which variables it can choose, whether it attempts to influence the market price with these choices, and how it believes other participants will respond to its actions

For each of these components, the modeller has a number of choices to make. The key thing is that the model should represent the most important features of the market, from the point of view of the question to be analysed, should do so in sufficient detail to be realistic, but should also be simple enough to create a model that can be solved without excessive computing resources. Wherever possible, the modeller should be able to provide an intuition for the results observed, if only after the fact. If it is not possible to provide an intuition for a particular effect, it may be an artefact of the modelling approach, rather than something that we would expect to see in reality.

2. What is the evidence for and against the view that the modelling approach adopted by Frontier on behalf of the AEMC is "fit for purpose" in helping inform the Commission's assessment of the competing Rule change proposals?

Each of the modelling choices made by Frontier is within the mainstream of economic modelling of electricity markets, as described above. I therefore believe that the model is "fit for purpose". During the process of preparing this report, I became aware of a number of issues where I thought the modelling, or its presentation, might usefully be altered, and Frontier has responded well to these suggestions.

3. What limitations or simplifications have been used by Frontier in applying its general modelling approach to the specific issues being addressed in this context? Are any such simplifications or limitations based on sound reasoning, e.g. data limitations?

A number of simplifications have been made by Frontier – this is an inevitable part of the process of producing a tractable model. The constraint equations used to represent the network effectively give a full nodal treatment, and this is what is needed in this area. Frontier restrict most of the strategic generators to a small strategy space, with few options (two or three levels of available capacity). Each additional option carries a heavy burden in computing requirements, and it would be infeasible to add more, but a larger set of options might have revealed other equilibria.

The model appears to ignore outages, which could inflate the modelled cost saving from changed patterns of Snowy output – less thermal output is needed at peaks, and more at off-peak periods – outages would increase the cost of this replacement generation, but are unlikely to have a significant impact on the peak costs. Ignoring start-up costs at peak times could reduce the modelled cost saving from additional Snowy output at these times. Note that these two effects (which might in fact be insignificant) are in opposite directions – the impacts of modelling choices can go either way.

Other simplifications, such as the use of a limited number of representative periods, a "notional generator" that caps market prices at $2,500 per MWh, and a fixed water budget for the Snowy stations, are sensible.

4. Are the assumptions used in the modelling (generally and specifically) clearly defined and documented, and appropriately derived from external data? Are the modelling assumptions applied accurately and consistently across all modelling scenarios?

5. Is there clear documentation on what external data sources are used in the modelling? Are there appropriate safeguards to ensure that such external data are used accurately and consistently across all modelling scenarios? Is there any primary evidence of data being used inaccurately or inconsistently?

Frontier's modelling reports clearly spell out the assumptions made in the modelling, and the sources used for the data. Each of these appears to be an appropriate source for the type of data obtained from it. I do not have the local knowledge to suggest any other sources that might have been used instead. I have no reason to doubt that the assumptions have been applied consistently across the scenarios analysed, except where the report specifically mentions differences in approach. A fuller answer to this part of the questions would require me to audit the computer code, assuming it was in a language I knew, which goes beyond the scope of the work envisaged. Frontier might be asked to run their model for known sets of past bids, to ensure that it replicates the market results that these produced. Their validation process has shown that the model generally produces market prices within the range actually observed.

6. Are the materials provided to the Commission summarising the modelling results accurate and comprehensive representations of the underlying modelling results? Are there appropriate internal checks and processes in place to ensure that this is the case?

Frontier's reports include information on all of the main variables of interest – system costs, power flows, average prices, and the pattern of output from different power stations. The results appear to be internally consistent. Once again, a detailed check of Frontier's working methods is beyond the scope of this report.

I held in-depth discussions with Frontier about the way in which results are reported. Frontier's model is able to find a number of Nash Equilibria for some (or all) demand levels – some involve high output levels, while others are based on several generators withdrawing capacity from the market. Frontier reports the average of all the equilibria it finds for each demand point. This is an objective measure, but does not itself usually represent a state that would be observed in equilibrium. I was worried that it might lead to misleading comparisons when the number of equilibria for a demand point changes between scenarios. A company's output might appear to change when an additional equilibrium is discovered, even if it was very similar to other equilibria already revealed, so that there was no particular reason for the company to move away from those. Frontier provided data on the number of equilibria found by the model for each year and scenario, which showed that the number of equilibria varied by less than 10% between scenarios, on average, over the course of a year. While the varying number of equilibria for a particular demand point might affect the results presented for it, I doubt that the effect will be significant over the whole year. I am therefore happy with Frontier's decision to present averaged results.

Frontier give a lot of detail about their results for demand period 29, but I believe that its high contribution to the overall cost saving could come from a high weighting when periods are combined over the year, rather than from a particularly high saving per hour. I recommended presenting the cost savings by period on a per-hour basis, and Frontier now do so.

I recommended that Frontier be asked to run the model under the assumption of competitive bidding by all participants. This would allow any differences between their results and those of other consultants (if also based on competitive bidding) to be ascribed to differences in the underlying model, and not to the way in which market power is analysed. Frontier did this and the modelling results showed the same ranking across scenarios as the strategic modelling.

In respect of modelling results provided to the AEMC by or on behalf of stakeholders, the Commission is seeking advice on the following questions:

1. What is the evidence, based on the available documentation, for and against these modelling results also being viewed by the Commission as "fit for purpose" in helping inform the Commission's assessment of the competing Rule change proposals?

I have reviewed the report by ROAM Consulting, prepared for the group of Southern Generators. This model differs in a number of respects from Frontier's model, most notably in only considering strategic behaviour on the part of Snowy Hydro, and requiring all thermal generators to bid their available capacity (after outages) at cost. I am not sure that this is an appropriate assumption for the National Electricity Market, in which two generators (apart from Snowy) each own about 10% of the industry's capacity. I expect that these firms would have the ability to bid in a strategic manner at peak times. (Several other firms own between 7% and 5% of capacity – I would need to run a model to know whether their strategic bids would be significantly different from marginal cost.)

A slightly odd feature of the model is the algorithm used to maximise Snowy Hydro's profits. The company will only change its bid from the typical bid to a strategic one if "the Snowy Hydro production revenue (in $/MWh) exceeded the 'typical' bid revenue (in $/MWh) by an adjustable margin" (Main report, page 5). I wonder whether this requirement would lead Snowy to miss some opportunities to exploit its market power.

Peaking generators are bid at their Long-Run Marginal Cost. This is a sensible assumption – these stations need to recover their full costs over a short period of operation, and the resulting bids may not be significantly different from the five times Short-Run Marginal Cost assumed by Frontier. However, I believe that it is inappropriate to also use LRMC, instead of SRMC, when calculating the cost of dispatch – changes in the output of these plants would then have a disproportionate impact on the calculated system costs.

With no multiple equilibria, the results are relatively straightforward to report, and I believe that they are appropriately set out and described.

2. What are the main gaps in the available documentation in terms of explaining how the modelling results were derived?

I found that the ROAM report was generally well-documented, with enough information to understand how their results were derived.

<u>Conclusion</u>
I have been asked to review two modelling reports related to rule changes proposed for the National Electricity Market. I believe that the Frontier model is "fit for purpose". I discussed a number of possible changes to the way in which the results were presented, and am satisfied with Frontier's response to these points. I also recommended running the model once for each scenario, based on price-taking behaviour, as a base case, which would also allow an easier comparison with ROAM's model, were they to undertake the same exercise. In the case of ROAM's model, I am concerned that the way in which they only adjust Snowy Hydro's behaviour from a typical pattern if it leads to a sufficient increase in the company's revenue per MWh, rather than attempting to equalise marginal revenue across periods, gives an unusual representation of profit-maximising behaviour by this company. I also believe that at least two, and possibly more, other generators might usefully have been modelled as strategic players.

Richard Green
23 August 2007

Review of modelling work: NEM Rule Changes and the Snowy Region

Professor Richard Green

The Australian Energy Market Commission has appointed me to advise on:

- The generic modelling approaches that could have been adopted in seeking to understand and assess the economic impacts of the different Rule change proposals [with respect to network congestion in the Snowy Region];

- The appropriateness and robustness of the modelling approach adopted by the Commission and its appointed economic modelling consultants in terms of:
  - Theoretical underpinnings; and
  - Practical application.

Two generators within the Snowy Region currently face different operating conditions, due to transmission constraints, but would receive the same price, were it not for a derogation from the market rules. Three proposed rule changes would make this derogation permanent, create a second zone, or place each generator within a (different) larger zone – the base case is to end the derogation. Modelling work has been commissioned to assess the economic impacts of these different changes (or non-change), both by the Commission and by a group of generators proposing one of the rule changes.

<u>Qualifications of the consultant</u>

I have been working on the economics of electricity transmission pricing for many years. My PhD thesis (Cambridge, 1995) contained some modelling of different pricing rules; I edited a journal issue that compared the transmission pricing systems of several OECD countries (and wrote the summarising paper) in 1997; and contributed to a consultancy report to the UK government on the economic impacts of charging for transmission losses in 2004. Earlier this year, the *Journal of Regulatory Economics* has published my work on the economic impact of three different transmission pricing rules within England and Wales, considering the welfare of generators and consumers, and incorporating strategic behaviour on the part of the two largest generating companies. Another paper predicting future transmission prices in the UK has been accepted by the *IEEE Transactions on Power Systems*. While I would not automatically expect the modelling commissioned by the AEMC and other stakeholders to follow the same approaches that I have used, I believe that my past work in this field gives me the expertise to assess their techniques and results.

I have worked for the Office of Electricity Regulation (UK), and advised the Commission de Regulation d'Electricite (France) and the Trade and Industry Committee of the House of Commons (UK) on transmission issues. I am on the Academic Advisory Panel of the Competition Commission (UK).

Questions asked by the Commission

The Commission asked me to answer several detailed questions about the work performed by its consultants, Frontier Economics (Frontier), and about the work performed by Roam on behalf of the group of Southern Generators:

1. What are the generic economic modelling techniques that could have been adopted in seeking to understand what impact implementing one of the relevant Rule change proposals might have on pursing the NEM objective?

2. What is the evidence for and against the view that the modelling approach adopted by Frontier on behalf of the AEMC is "fit for purpose" in helping inform the Commission's assessment of the competing Rule change proposals?

3. What limitations or simplifications have been used by Frontier in applying its general modelling approach to the specific issues being addressed in this context? Are any such simplifications or limitations based on sound reasoning, e.g. data limitations?

4. Are the assumptions used in the modelling (generally and specifically) clearly defined and documented, and appropriately derived from external data? Are the modelling assumptions applied accurately and consistently across all modelling scenarios?

5. Is there clear documentation on what external data sources are used in the modelling? Are there appropriate safeguards to ensure that such external data are used accurately and consistently across all modelling scenarios? Is there any primary evidence of data being used inaccurately or inconsistently?

6. Are the materials provided to the Commission summarising the modelling results accurate and comprehensive representations of the underlying modelling results? Are there appropriate internal checks and processes in place to ensure that this is the case?

In respect of modelling results provided to the AEMC by or on behalf of stakeholders, the Commission is seeking advice on the following questions:

1. What is the evidence, based on the available documentation, for and against these modelling results also being viewed by the Commission as "fit for purpose" in helping inform the Commission's assessment of the competing Rule change proposals?

2. What are the main gaps in the available documentation in terms of explaining how the modelling results were derived?

Method of working

I have read background documents on the Snowy Region rule changes, the Commission's Draft Rule Determination of 19 January 2007, including the modelling appendix by Frontier, and the draft of the report to be released with the Commission's

Final Rule Determination. Frontier have answered a number of my questions via conference calls with Commission staff or via email. My review of ROAM Consulting's work has been based on reading their report of 3 April 2007, and their responses (informed by a teleconference) to questions posed by the Commission. I presented a draft report to the Commission in August 2007, and received helpful feedback. The report presented here remains my own work and responsibility, however.

In the rest of the report, I answer these questions in the order they were asked.

1. Generic Modelling Techniques

The system of electricity generators and transmission lines is a complex one. Except in special cases, it will not be possible to predict the impact of a change in the market rules by logical reasoning. A numerical simulation must be used instead.

There are five key components to such a simulation.

a) The timescale for the simulation, including the periods to be studied, and whether the model runs for every hour in chronological order, or a sample of representative hours
b) The characteristics (capacity, operating costs, ownership) of each generator on the system, and the levels of demand at different times over the period of interest
c) A description of the transmission system, including the locations of generators and loads, equations to determine the flow across each line (or boundary) likely to be congested, and to determine the level of transmission losses
d) A description of the market rules, that takes bids from generators (and loads, if applicable), and calculates the operating patterns, costs, prices and hence profits that will result
e) An assumption about each market participant's behaviour – which variables it can choose, whether it attempts to influence the market price with these choices, and how it believes other participants will respond to its actions

Most of the early studies of electricity markets combined elements a), b), d) and e), but did not include transmission effects. Where the markets themselves minimised the impact of transmission (either losses or congestion) on their price determination process, this was an acceptable simplification. Some more recent studies have included transmission effects, and in some markets, where these explicitly form a large part of the price determination process, this is essential.

For component a), the choice of the year to study will depend on the purpose of the study – for example, modelling the results of investment requires a study looking some years into the future. The further forward the model looks, the more likely that some important factors differ from those assumed in the modelling – one approach to this is to run the model for several different scenarios that can cover a range of outcomes.

It is possible to run a model either for individual hours, or for a number of hours in sequence. The first alternative leads to a simpler model, in which plants can be "stacked" in merit order by price (for a market simulation) or cost (for an optimising run). The second alternative is more complicated, but allows for a realistic treatment of start-up costs, which tend to raise the cost (and hence price) at peak times. It also allows analysis of decisions taken to avoid having to shut down, including (in Australia) sometimes bidding negative prices overnight.

The modeller must also choose how many hours to run the model for. The model could be run for each hour of a year (individually or chronologically), or it could be run for representative hours. If the chronological approach is taken, the model would usually be run for at least a day at a time, in order to capture a cycle of start-up and shut-down decisions; however, it would be quite acceptable to run it for a sample of representative days over the year. If the model is not run for every hour of the year (which, for a complex model, could need prohibitive amounts of computer time), then the results should be weighted appropriately when they are presented for the whole year. It may be sensible to run more hours (or days) at times when the system is likely to be under stress – these will almost certainly include the demand peaks.[1] There may be significant differences between the day of highest demand, and that of, say, fourth-highest, which would not be the case between the days around the lowest demand levels.

The key modelling choices for component b) are which plants to include, and what data to use on their performance. For the modelling work performed by the Commission, which looks at the near future, this boils down to the question of whether new capacity should be expected in the market.[2] There are basically two alternatives here – the first is to include those projects currently under construction (which are very likely to join the market at around their planned commissioning dates), and to extrapolate current trends for any plant that could be completed in time for the later study periods, even if not currently under construction. This extrapolation does require judgment – if one large plant is about to be completed, this could well be the signal that it is not necessary to start a second one; whereas if there has been a large amount of investment in renewable generators, supported by special incentives that are going to continue, then further investment may be expected. The second alternative is to use the least-cost investment solution from a system planning model. This has the advantage of not requiring any discretion from the modeller. The disadvantage is that if current projects under construction are incompatible with that least-cost expansion plan, then this approach may produce a set of power stations that is different from those that will shortly appear on the system.

One variant would be to consider more than one investment scenario, particularly if the modelling is for a period several years ahead – however, the effort is less likely to be worthwhile for short-term predictions, such as those commissioned by the AEMC.

Operating cost data are important, as they determine the merit order in which plants would be used in an optimised system with no transmission or other operating constraints. A key question here is whether the merit order would change in response to conceivable variations in fuel prices. In the UK, the relative prices of gas and coal

---

[1] If peaking and base load generators are at different points on the network, with the peaking stations closer to the load, flows on the transmission system might increase at off-peak times, and it could be important to model these in detail if constraints are likely to bind.
[2] In some studies, looking at the United States or Europe, the treatment of imports from another market can involve difficult modelling choices.

have been such that (combined cycle) gas-fired stations were cheaper than coal stations for much of 2004, but were more expensive in 2005. If there is any chance that the merit order might change in response to fuel prices, and particularly if this would affect the pattern of flows on the transmission system, it would be advisable to perform the modelling work for two fuel price scenarios.

If hydro-electric generators with storage are important, then their reservoir capacities must be specified, together with any restrictions on their minimum and maximum generation levels over a time period that may be necessary to keep river systems healthy. Run-of-river generators have little control over their output, but the expected levels at different times of year must be part of the dataset.

The modeller must decide how much of the capacity in the system is actually available for use. Even at peak times, some capacity will suffer from forced outages, while at off-peak times, a significant proportion of capacity will be out of service for planned maintenance. There are two approaches to representing this. One is to de-rate every unit by a set percentage, representing the forced outage rate at peak times, and the level of maintenance scheduled for other times, so that a 100 MW unit might be given an available capacity of, say, 93 MW at the peak times, and 80 MW at others. The second approach is to model individual outages as random events, so that each unit would have a fixed chance of being out of service at particular times. The percentage might vary with the type of unit, and would vary with the time of year. This approach is more realistic than the first, for most problems with power plants do lead to outages rather than to de-ratings, but the results of a simulation can be sensitive to the exact mix of unavailable plants. If the other elements of the simulation are sufficiently simple, the simulation can be run many times for different outage patterns, and the results used as a Monte Carlo simulation. If the simulation is already a complex one, it may not be possible to repeat it – in this case, if the aim of the simulation is to compare different rules, or policies, it is probably best to ensure that the same out-turn pattern of outages is used for each run.

If the model uses a chronological approach, solving for several periods sequentially, then inter-temporal operating constraints can also be considered. At the very least, the cost of starting up a unit should be included, but a more complex model could also include limitations on ramp rates, minimum generation levels, and minimum running or shut-down times. Including start-up costs will typically increase peak marginal costs, and hence prices, and reduce off-peak marginal costs and prices.

Demand levels for the periods to be studied should be taken from the system operator's forecasts, or a similar source. A key modelling choice is whether demand should respond to the market price, falling when prices are higher than normal, and possibly rising when prices are low. In a model with a short-term focus, it might be appropriate only to include loads that explicitly submit demand-side bids to the market, treating them as a kind of negative generation. In other models, a demand function with a negative slope may be included. At times of high demand, this will shift out, showing that higher combinations of price and load are expected, but if the price is particularly high, the load will be reduced from a normal peak level. If demand does respond to price, the modeller must choose the extent of this response, or its elasticity. In the short term, electricity demand is notoriously inelastic, such that a 10 per cent increase in price might only lead to a one per cent reduction in quantity

– an elasticity of (minus) 0.1. In the medium term, demand is more responsive, and so a higher elasticity would be appropriate if the intention of the modelling is to show how demand would respond to a sustained increase in prices. An elasticity of (minus) 0.3 has been used by the Office of Gas and Electricity Markets in the UK, for example.

Small-scale embedded generators are often netted off the gross load and do not participate fully in the wholesale market, selling their output via a contract to a local retailer. A market model can either include this generation, and the corresponding load, or can exclude both, modelling only the net load considered by the market operator. The advantage of including embedded generators explicitly is that if their operating patterns respond to market prices, this can be treated in the same way as any other (price-taking) generator. If price-responsive embedded generators are not explicitly included in the model, then their response should be included within the overall response of load to market prices.

For component c), the transmission system, the key question is the level of detail at which the system is modelled. It is possible to include every line and transformer, but the model might become complex to solve. Most modellers choose to simplify the network to a number of zones, together with the interconnections between them. The zones should be chosen so that they did not contain congested lines. The connections between zones should be modelled in such a way as to show whether there is congestion. In the simplest case, it may be adequate to set the maximum transfer between two zones (which may differ according to the direction of flow) without having to determine the source of the flow within a zone. In other cases, when the lines within a zone are meshed, and there are several lines between the zones, output from one generator within the zone may have a much greater impact on the constraint than output from another. If the system operator, or some other source, publishes constraint equations that relate the permissible output from different locations, then these can be used. Otherwise, the zones chosen may be too aggregated to give reliable results, and it might be better to choose a greater level of detail. If the connections between zones are complex, it may be necessary to use load flow equations to determine the flows on particular lines – typically using the DC approximation.

For component d), the market rules, the key modelling choice is the level of detail to adopt. In the National Electricity market, for example, each generator can choose up to ten prices at which offer different slices of capacity to the market. Should the modeller use the same number of prices, or will offering all of a unit's capacity at the same price be close enough as an approximation to the market rules? It will be necessary to specify the way in which prices are set in the market for each time period, and what rules apply if prices vary between zones. If it is impossible to resolve all the transmission congestion by inter-zonal price differences, how does the market deal with it?

Component e) can be the most important in explaining the different results obtained by modellers. What variables do generators choose, and what do they take into account when they do so? For small generators, the assumption of competitive behaviour is the most appropriate – they will offer their full capacity at their marginal cost, and hence generate whenever the market price exceeds this cost. For a hydro generator, which receives its “fuel” free of charge, a marginal cost can still be expressed as the shadow value of its water – if the generator runs whenever the market price is above this level, the unit should maximise its profits while just using its available water.

Larger generators, and companies owning several stations, can affect the market price with their bidding strategies. The usual modelling assumption is to designate a number of “strategic” bidders, each of which will aim to maximise its profits, taking its impact on the market price(s) into account, and taking the other bidders’ behaviour as given. In other words, there is no “leadership” behaviour, in which one company expects others to follow it when they set their prices. A Nash Equilibrium occurs when each company is maximising its profits, taking into account its effect on the market price(s), but taking the other firms’ strategies as given.

One key question for the modeller is which generators should be considered as strategic. There is no hard and fast dividing line, but it will be important to be consistent – if two companies have similar capacities, they should either both be treated as strategic, or neither should. The key test for whether a company should be considered as strategic is its ability to raise prices profitably – if the company tries to raise the market price, will its output fall so much that the change is unprofitable?

The second key question is the generators’ decision variable – what do they choose. There are three standard approaches to this question. The most common method in electricity market modelling is to assume that firms choose the quantity that they offer to the market, and that the market price is then determined by the point at which the strategic generators’ quantity is just equal to the load less the output of non-strategic generators, each of which may depend on the price.[3] This is the Cournot modelling assumption. The advantage of this method is that small changes in a generator’s quantity choice rarely lead to discontinuous changes in its profits, which makes it easier to find the equilibrium. In the absence of transmission effects, the equilibrium should also be unique. The disadvantage comes from the assumption, implicitly made by each strategic firm, that the other firms’ output levels will not change in response to a change in its bids. Short of assuming collusion between the firms, this is the assumption that will give the highest equilibrium prices. Each firm is free to attempt to sacrifice output and raise the market price without fearing that others will take the opportunity to raise their own output, which would make a bigger sacrifice necessary and make the price increase less likely to be profitable. Furthermore, the assumption is unrealistic to the extent that firms offer prices to an electricity market, and not just quantities.

The second modelling assumption is thus that each firm chooses the price at which it is willing for its output to be sold. This fits the procedures of a typical electricity market more closely, but has the disadvantage that a small change in the firm’s price can lead to a discontinuous change in its profits, if this allows the firm to undercut a rival and replace that firm’s output in the market. Pure price-setting models (the Bertrand assumption) are thus rarely used for empirical modelling in electricity markets.

[3] If neither the non-strategic generators’ output, nor the load, depends on price, then a pure quantity-setting model may well not have a solution.

The third modelling assumption, and the most realistic, is that firms offer some combination of prices and quantities to the market. One variant would be for each firm to offer all of the capacity of a given unit at a price of its choice – the strategy then consists of these price offers. Another is to set the price of each unit at its marginal cost, but not necessarily to offer all of its capacity to the market – the strategic decision is how much capacity to withhold. A third variant is for the firms to offer a schedule of both prices and quantities – a supply function. This is the most complex to compute, but gives the firms the greatest degree of flexibility. The other two choices – setting a number of prices, and setting a number of quantities – have some of the same advantages and disadvantages of the pure price and pure quantity strategies described earlier.

The modeller also needs to specify whether the firms' short-term incentives are affected by their contract holdings. A generator that has already signed contracts to sell electricity at a fixed price, whether for physical delivery or through a financial contract for differences, has less to gain from an increase in the spot market price. This will reduce its incentive to raise the price, and may even give the generator an incentive to reduce the spot market price, should it have sold more in advance than it expects to produce. In some cases, the volumes of contracts may be known, particularly if vesting contracts set up at the time of restructuring are still in force. In other cases, the modeller will have to make an assumption about the level of contract holdings, perhaps in the form of a percentage of the firm's capacity, or of the output it would produce in an equilibrium where no generator had market power (the competitive outcome, or system optimum).

Finally, any other government policies that may affect the generators' incentives should also be included in the model. These might include taxes on emissions (which would properly be included in the relevant stations' operating costs) or emissions trading schemes, or requirements on retailers to procure a certain percentage of their power from green sources.

This section of the report has listed many choices that the modeller must make. The key thing is that the model should represent the most important features of the market, from the point of view of the question to be analysed, should do so in sufficient detail to be realistic, but should also be simple enough to create a model that can be solved without excessive computing resources. Wherever possible, the modeller should be able to provide an intuition for the results observed, if only after the fact. If it is not possible to provide an intuition for a particular effect, it may be an artefact of the modelling approach, rather than something that we would expect to see in reality.

2. Fitness for purpose

Each of the modelling choices made by Frontier is within the mainstream of economic modelling of electricity markets, as described above. I therefore believe that the model is "fit for purpose". In reading Frontier's reports, I became aware of a number of issues where I thought the modelling, or its presentation, might usefully be altered, and Frontier has responded well to these suggestions.

3. Simplifications made by Frontier

A number of simplifications have been made by Frontier – this is an inevitable part of the process of producing a tractable model. In this part of the report, I discuss those that appear more important to me, together with their likely implications.

The model is described as a zonal representation of the electricity network, but incorporates constraint equations, supplied by Nemmco, which incorporate station-specific coefficients. In other words, each station's output can have a different impact on a given constraint. This means that the model is effectively nodal when it comes to these constraints. This is the level of detail required for this problem. These constraints do contain some approximations,[4] but these come from Nemmco, and not from the modellers.

Most generators are limited to a small strategy space – they can offer their units at marginal cost with output levels equal to 90%, 80% or in some cases 70% of capacity. This is a restriction, but may not be a binding one. There are two aspects – the range of strategies considered, and the number of possible strategies within this range. Has Frontier ever run the model with a wider strategy space (e.g. including the option of 60%) or a finer grid (e.g. including 85%) to check that the additional strategies would not be used in equilibrium? It is quite possible that strategies outside the range considered would not be used in equilibrium. However, Frontier show that the model has multiple equilibria in at least some periods, and it is possible that a wider range of strategies, or a finer grid, would reveal more equilibria. A finer grid would also lead to more accurate characterisation of the model's Nash equilibria. At present, if the model shows that a (two-choice) generator offers 80% of its capacity in the equilibrium, we know that this will be more profitable for it than offering 90%, but do not know whether the most profitable level to offer is actually 80%, 75% or 85% (it is unlikely to be closer to the other strategy than this). Increasing the number of strategies that the various generators can choose raises the complexity of the model, and the time it will take to solve, dramatically. Frontier has told me that the modelling of the scenarios, as currently specified, involves 12 million dispatch operations. Adding a single extra strategy for each strategic firm would raise this to 55 billion dispatch operations, which clearly implies a disproportionate effort. Given the limitations of computer resources, there is probably nothing that can be done about this issue.

The report states that this modelling does not take account of generator outages. At peak times, this should not distort the results for the strategic generators, compared to a model in which every generator is de-rated by a fixed percentage. If they are only choosing to offer 80% of their capacity in equilibrium, it is irrelevant whether they could offer 100% (with no outages) or 90% (with a 10% forced outage rate). There is a potential distortion for the non-strategic generators, who will be seen to offer more capacity than they would normally have available after taking forced outages into account – I doubt that this will be have a significant effect. I am more concerned at the effect at the "other" model periods. The rule changes that lead to lower operating costs do so because Snowy raises its output at peak times. While this requires the Snowy stations to generate less in the other periods, the system marginal cost in these

[4] When the network boundaries are changed between scenarios, the constraint equations change, and it is sometimes possible for the same set of station bids to give an unconstrained network in one scenario, but to violate a constraint in the second.

periods is lower than at the peaks. It therefore costs less to replace a MW of Snowy output with thermal generation at these times, than at the peak times. The assumptions that all non-strategic generators are fully available in the other periods, and that the strategic generators offer 90% of their capacity, reduces the system marginal cost, relative to a situation in which there are higher levels of outages. Modelling higher outages would raise the cost of replacement power off-peak, and hence reduce the savings from the rule changes. If the system marginal cost curve at these times is flat, this effect will not be significant, but it could usefully be the subject of sensitivity analysis, if this has not already been done.

The model solves each period individually, rather than solving for each hour in a representative day in sequence. This makes the model more tractable, but reduces the realism of the way in which thermal generators' costs are modelled. In particular, start-up costs at the time of daily demand peaks can add significantly to the marginal cost of generation. If the rule changes under analysis lead to higher levels of Snowy output at these times, then the true cost saving (and hence benefit from the rule change) would be greater than a model based on individual periods would show. Hydro generators do not suffer from inter-temporal constraints in the same way as thermal plants (apart from the need to observe their annual water limits, which is included in the model). This means that the operational possibilities facing the Snowy stations are not changed by solving the model for individual periods – however, if the market prices would change with a sequential solution, the economic possibilities would change.[5]

Note that the two previous paragraphs have described one feature of the model that might lead to an over-estimate of the cost savings resulting from implementing some of the proposed rule changes, and one feature that might lead to an under-estimate. I am not suggesting that such features should be counted in any formal way, but want to point out that modelling decisions can have impacts in opposite directions.

The model is only solved for a number of representative periods, rather than for every hour of the year. Demand conditions in which constraints in the Snowy region have typically appeared are over-represented, compared to than those in which the region is unconstrained. This is a sensible simplification which is unlikely to have any significant impact on the results.

The model contains a "notional generator" which bids at a price of $2,500 per MWh and caps the market price at this level. This is because the strategic generators would otherwise find it profitable to set the price at VOLL ($10,000 per MWh) far more frequently than occurs in reality. The notional generator is an effective way to avoid these exceptionally high prices. If generators purely aimed to maximise their short-run profits, they would want to (and be able to) drive the market price above this level, but doing so too often would invite political or regulatory action. I believe that capping the prices directly in this way is an acceptable way of dealing with this issue.

[5] The direction of such a change is unclear. However, if the marginal cost of thermal generation is higher, the market price might also be higher. The higher the margin between price and marginal cost, the less worthwhile any sacrifice of output that further increases this margin. This might imply that Snowy's output would increase further, relative to a model that did not incorporate start-up costs.

Snowy Hydro is given a fixed water budget for the year, amounting to 4.9 TWh of generation, and appears to be able to produce this much electricity from either of its two main generation schemes. I have been told that this is realistic, in that Snowy Hydro can divert water from one power station system to the other. There does not seem to be a role for pumping water to increase the generation budget – the questions are whether the pumped storage station can add much to its output by pumping, and whether is already included in the 4.9 TWh.

<u>4. Assumptions made by Frontier</u>
<u>5. Data sources used by Frontier</u>
It is easiest to answer these questions together. Frontier's modelling reports clearly spell out the assumptions made in the modelling, and the sources used for the data. Each of these appears to be an appropriate source for the type of data obtained from it. I do not have the local knowledge to suggest any other sources that might have been used instead. I have no reason to doubt that the assumptions have been applied consistently across the scenarios analysed, except where the report specifically mentions differences in approach. Since changing assumptions between scenarios requires the modeller to re-code part of their programme, I am confident that Frontier's preferred approach, as well as that requested by the Commission, would be to minimise changes between scenarios.

Frontier's model of imperfect competition is based on the partial withdrawal of capacity by the various strategic generators. This is a generally accepted way of modelling imperfect competition in electricity, but does tend to lead to higher prices, and lower output levels, than alternative methods. Given the assumption, I have no reason to doubt that the modelling is performed correctly. Since the same method is used throughout the studies, any bias induced by it should have little impact on the comparison between scenarios.

At one point in preparing this report, I was under the impression that Frontier was assuming significantly less investment in Queensland than ROAM Consulting. Frontier do not explicitly mention the Kogan Creek plant, which will be commissioned during the modelling period, and I had supposed that it was not included in the investment forecasts from their cost-minimising system planning model, WHIRLYGIG. Since ROAM Consulting explicitly mention the plant, and its presence in Queensland would likely reduce flows from Victoria to New South Wales, I had wondered whether this could be a factor accounting for the differences in results between the models. In practice, Frontier have included this plant in their modelling as a "committed" investment, and so this does not account for any differences in the model results – their report now makes this clear.

A full answer to these questions would require me to audit Frontier's programme codes. This would only be a sensible activity if the codes were in a computer language that I already understood, and would also increase the cost of the project. I am assuming that the Commission only wishes for "desk research", confirming the plausibility of the assumption and results in the documents supplied to me. However, Frontier could be asked to run their model for a given set of power station bids observed in the past, in order to check that it gives the market results at that time period (or to report the results of such a check, which they have probably done themselves). Frontier have told me that their model successfully replicates the distribution of out-turn prices seen in the National Electricity Market, which is an important condition for it to meet.

6. Reporting of results

Frontier's reports include information on all of the main variables of interest – system costs, power flows, average prices, and the pattern of output from different power stations. The results appear to be internally consistent. Once again, a detailed check of Frontier's working methods is beyond the scope of this report.

I had lengthy discussions with Frontier on one issue concerning the way in which their results are reported. Frontier's model is able to find a number of Nash Equilibria for some (or all) demand levels. Typically, there might be one equilibrium (or set of equilibria) in which most of the strategic generators offer a large proportion of their capacity, and it is not profitable for any of the others to withhold their capacity, either. There can be a second equilibrium (or set), however, in which a number of generators are withholding capacity, the transmission system is congested, and it is more profitable for each generator (among those with the choice) to keep the system congested than to expand output and relax the constraint. Other permutations of output and congestion are of course possible in other system conditions – sometimes a firm that expands output is the one that congests an interconnector.

Frontier reports the unweighted average of all the equilibria that it discovers, although I could not find this stated explicitly in their draft report. Giving this unweighted average is an objective way of reporting the results. In most cases, however, this average does not reflect a state that the system will actually reach in equilibrium at that point in time

I was concerned that this averaging could lead to misleading conclusions, in particular if the number of equilibria for a given demand point was to change across scenarios. For example, the model might find one "competitive" equilibrium, with high output levels, and one "uncompetitive" equilibrium, with low output levels, for a particular demand point and scenario. In another scenario, the model might find the competitive equilibrium, and two, quite similar, uncompetitive equilibria. Taking the weighted average of these three would suggest that the second scenario would lead to lower output levels (for the generator withdrawing capacity) than the first. However, it would also be possible to take the view that if the generators were going to play a competitive equilibrium, little would change between the scenarios, and that there would also be little change between scenarios if the generators were to play an uncompetitive equilibrium. The mere existence of a second uncompetitive equilibrium does not make it more likely that the generators will choose to act in an uncompetitive manner, although if they are involved in a learning process, it may be easier for them to converge on an uncompetitive outcome if there are more of them to find!

I therefore discussed alternative ways of presenting the results with Frontier, such as presenting the most and least "competitive" equilibria separately. Their view was that generators do appear to switch between more and less competitive bidding behaviour at different times, and presenting the average results is the best way of reflecting this. Furthermore, they provided me with information on the number of equilibria the model produced for each year and scenario.[6] This showed that the number of equilibria varied by less than 10% between scenarios, on average, over the course of a year. While the varying number of equilibria for a particular demand point might affect the results presented for it, I doubt that the effect will be significant over the whole year. I am therefore happy with Frontier's decision to present averaged results.

Frontier give quite a lot of detail about the results obtained for demand period 29, since they typically contribute the most to the cost savings observed under the various rule changes. Focusing on a single period to provide intuition about what is happening in the model, both in that and in similar periods, is a sensible way to report the results. I was not sure that period 29 actually gives unusually high savings on a per-hour basis, however. From figure B1, it is possible to tell that when the results are aggregated from individual demand points to a total for the year, demand point 29 is given a much higher weight than most of the other peak demand periods (though *not* demand point 30). I recommended that Frontier should also report the period-by-period results for cost savings on a per-hour basis, and they now do so.

I thought of another set of results which could usefully be reported, and would not be hard to calculate, since WHIRLYGIG already performs some of the calculations as part of its investment planning process. These would be based on competitive bidding by every generator, and are used in the system optimisation that underlies the investment model. Even if we do not expect every generator to bid its marginal cost, a set of results that shows the impact of the different rules, given this assumption, can act as an unambiguous starting point for comparing different models. I recommended that Frontier should be asked to run the model using this bidding assumption for each of the scenarios – the results would indicate the "pure" efficiency gains from the rule changes. If two models produce different results for this comparison, then the difference should be due to the way in which they model the system running in an optimised manner. If the models produce different results for cases involving market power, and we have not compared their results with competitive bidding, it is not possible to tell whether the difference is due to the way in which the electricity system is modelled, or market power is studied. Frontier has done this and obtained the same ranking across scenarios as in the strategic analysis.

7. Fitness for purpose of stakeholder modelling

I have also reviewed the report by ROAM Consulting, prepared for the group of Southern Generators. This model differs in a number of respects from Frontier's model, most notably in only considering strategic behaviour on the part of Snowy Hydro, and requiring all thermal generators to bid their available capacity at cost.

The assumption that Snowy is the only strategic generator may be a weakness of the model. Two other generators in the National Electricity Market have capacity shares of around 10%, and a further six have shares of 5% to 7%. I would expect the larger two generators to have the ability to profitably bid in a strategic manner in high-demand periods. It is also possible that the medium-size generators might have some ability to bid strategically – I would need to run a model to see whether this was likely

[6] A scenario with two demand points, one for which the model produced two equilibria and one for which it found three, would have five equilibria in total.

to have a significant impact on the market. While a high forced outage rate might be equivalent in its effects to the strategic withdrawal of capacity, ROAM state that their outage rates are consistent with those used by Nemmco in its 2006 Minimum Reserve Level calculations. The data for these give an average equivalent forced outage rate (including de-ratings) for non-peaking plant of just under 5%. This is not going to be equivalent to the strategic behaviour considered by Frontier, with capacity withholding of between 10% and 30%.

The model is solved for every half-hour over the year, which ensures that no one period can have a disproportionate impact on the reported results. The documentation does not mention any use of inter-temporal constraints on generators' output patterns, apart from hydro generators' water limitations. The increase in computational effort involved in modelling every half-hour is offset by the fact that strategic behaviour is only considered for Snowy Hydro, reducing the number of runs required per half-hour.

There can be no question of multiple equilibria – only one company is acting strategically, and so its profit-maximising bids will provide the sole equilibrium for each half-hour period. This simplifies the reporting of results. ROAM also provide a case based on "typical Snowy bidding" – it is not clear how close this case would be to pure price-taking behaviour, as Snowy Hydro's typical behaviour may include some actions to take advantage of market power. In particular, if its dispatch pattern is designed to exploit the existing rules, then we should expect it to change in response to rule changes, and so the "typical bidding" results *may* be suspect in the other scenarios.

There is one slightly odd feature of the algorithm used to maximise Snowy Hydro's profits. The company will only change its bid from the typical bid to a strategic one if "the Snowy Hydro production revenue (in $/MWh) exceeded the 'typical' bid revenue (in $/MWh) by an adjustable margin" (Main report, page 5). I wonder whether this requirement would lead Snowy to miss some opportunities to exploit its market power. Conceptually, Snowy should be aiming to equalise its marginal revenue across periods – the change in revenue from releasing a little more water should be the same, whether it is in a peak or an off-peak period. A test based on a (varying) benchmark production revenue could well give different results. When explaining differences between their results and Frontier's, ROAM state that "there is decreased incentive for Snowy Hydro to operate strategically in the BAU case, as the pool price is sufficiently high to produce high returns on a 'normal' level of generation" whereas "pool prices are lower in the typical bidding case. This presents a greater benefit for Snowy Hydro by bidding strategically, and Snowy Hydro responds to this opportunity" (Main report, pp 22-3). However, if the pool prices are lower, this may imply that there are more generators able to increase output in response to reductions by Snowy Hydro, which would reduce Snowy Hydro's marginal revenue and thus make output reductions less profitable than in the higher price case. The exact outcome depends upon the shape of the supply curve of the other generators.

Peaking generators are bid at their Long-Run Marginal Cost. This is a sensible assumption – these stations need to recover their full costs over a short period of operation, and the resulting bids may not be significantly different from the five times Short-Run Marginal Cost assumed by Frontier. However, I believe that it is inappropriate to also use LRMC, instead of SRMC, when calculating the cost of dispatch – changes in the output of these plants would then have a disproportionate impact on the calculated system costs.

With no multiple equilibria, the results are relatively straightforward to report, and I believe that they are appropriately set out and described.

8. Documentation of stakeholder modelling

I found that the ROAM report was generally well-documented, with enough information to understand how their results were derived.

9. Conclusion

I have been asked to review two modelling reports related to rule changes proposed for the National Electricity Market. I believe that the Frontier model is "fit for purpose". I discussed a number of possible changes to the way in which the results were presented, and am satisfied with Frontier's response to these points. I also recommended running the model once for each scenario, based on price-taking behaviour, as a base case, which would also allow an easier comparison with ROAM's model, were they to undertake the same exercise. In the case of ROAM's model, I am concerned that the way in which they only adjust Snowy Hydro's behaviour from a typical pattern if it leads to a sufficient increase in the company's revenue per MWh, rather than attempting to equalise marginal revenue across periods, gives an unusual representation of profit-maximising behaviour by this company. I also believe that at least two, and possibly more, other generators might usefully have been modelled as strategic players.

Richard Green
23 August 2007