Submission on Reserve Generation Policy Proposal

Electric Power Optimization Centre

University of Auckland

Professor Andy Philpott

Dr Geoffrey Pritchard

Dr Golbon Zakeri

www.esc.auckland.ac.nz/Organisations/EPOC

June 30, 2003

Executive Summary

1. The Electricity Commission's task is complicated, and its mission should be precisely defined before its *modi operandi* are determined.

2. The establishment of the Commission should take care not to rule out innovative solutions to the supply-security problem – in particular, demand-side participation.

3. In meeting objectives the Commission should be mindful of the (possibly unforeseen) effects of their decisions on electricity market participant behaviour.

4. Some sophisticated mathematical modelling will be required to ensure that the decisions of the Commission are well-founded.

Contents

1. Introduction.

2. The Commission's objectives.

3. Contracting decisions.

4. Possible components of the reserve portfolio.

5. Operating decisions.

6. Funding the Commission.

7. Alternative approaches.

Appendix: Modelling.

1. Introduction

The Discussion Paper published on May 20, 2003 by the Ministry of Economic Development lists a range of new policy measures to improve security of electricity supply. An Electricity Commission will contract for reserve generation and set terms and conditions for its use. This submission prepared by the Electric Power Optimization Centre at the University of Auckland examines this part of the proposal and the design issues that it raises.

The proposal to form an Electricity Commission has emerged out a sequence of winters that have precipitated national calls for savings of electricity. The current form of electricity market has been unable to guarantee the level of supply security that the government desires. The Commission is intended to provide this level of security.

The Discussion Paper has specified that the Commission should provide sufficient resources to withstand a 1 in 60 year event. To achieve this the Commission will contract by tender for a quantity of security reserve generation to be set aside. The proposal is that the Commission will pay the capital costs and maintenance costs plus the fuel costs for any generation that the plant is called upon to deliver. The Commission will have the right to offer the power into the market at a price to be chosen by them, and will recover any costs from a levy to be paid on every MWHr of energy generated.

The remainder of this document is laid out as follows. In the next section we examine the Commission's objectives. Our approach is to attempt to place these in an optimization framework to illuminate the tradeoffs between cost and security. In this framework the Commission's task is separated into two problems: the first is to decide what reserve capacity to contract for, and the second is how to offer this capacity to the wholesale market. In section 3 we formulate the contracting decision, and in section 4 we look at possible components that might be included in a portfolio of contracts held by the Commission. Section 6 looks at the costs of the Commission and the proposal to fund these with an industry levy, and section 5 discusses how the Commission might offer its contracted generation to the wholesale electricity market. Finally in section 7 we canvas some alternative approaches to providing reserve generation in the New Zealand wholesale market.. These approaches vary in the degree of intervention in the wholesale market, from very little intervention to central operation of electric power transmission and supply. The report is completed with an appendix on electricity market models.

2. The Commission's objectives

The proposal in the Discussion Paper describes in broad terms a policy for the Commission to provide the necessary reserve generation in the market as well as a set of instructions for offering this generation into the market. We see these as being distinct (albeit connected) problems that form part of a two-stage decision process: first decide the reserve capacity to contract and then determine an offering policy given this capacity.

In this document we discuss this two-stage decision problem from the perspective of optimization modeling. The motivation for this rests in the observation that the Commission must have at its disposal a mechanism for comparing different capacity contracts it is offered. For example, one tender might be for a coal plant at an annual cost of $200K/MW and fuel cost of $40/MWHr, and a second tender might be for an oil plant at $100K/MW per year but with a fuel cost of $120/MWHr. It is not clear which of these provides better value for alleviating the risk of a supply crisis.

Optimization models provide a useful framework for trading off the costs and benefits of competing proposals, and choosing the best with respect to some objective. It is desirable that the Commission chooses contracts that together are as effective as possible, if only to minimize the levy required. Its task is not easy. Essentially the Commission must solve a portfolio optimization problem at each contract point that determines how to provide security for electricity supply at least cost.

The Commission's first task should be to decide on precisely what its objective is. From the Discussion Paper, several objectives can be inferred:

- Minimize the probability of a power shortage, or ensure that this probability does not exceed 1 in 60.
- Minimize the industry levy, or ensure that this levy does not exceed 0.5c/kWh.
- Minimize intervention in the electricity market; allow the market to work unimpeded in normal years.
- Minimize excessive spot price volatility.

These desiderata allow the Commission to formulate its portfolio optimization problem in many different ways. For example, three different formulations are as follows:

P1: minimize Expected National Economic Loss from a Power Shortage
+ Total Industry Levy

P2: minimize Total Industry Levy
subject to Prob(Power Shortage) ≤ 1/60

P3: minimize Prob(Power Shortage)
subject to Total Industry Levy ≤ 0.5c/KWHr

We observe that all of these formulations are consistent with the purpose of the Commission, yet may yield different decisions when solved. The formulation P1 is a cost-benefit model that seeks to quantify the expected loss from a shortage of power and trade this off against a levy that can be used to alleviate it. If it could be adequately modelled and solved[1], then this model should guide the Commission. The current specification of a 1 in 60 year event as being the appropriate security threshold requires some rigorous analysis in this context, since in their current form the proposals

[1] Modelling the effect on the national economy is not straightforward, requiring a general equilibrium model that accounts for intersectoral flows of goods and services.

advanced by the Discussion Paper make the tradeoff between security and cost at 0.5c/KWHr versus a 1 in 60 year event. It may be the case that the objective of P1 is better served by some other tradeoff.

From now on in this report we shall focus on the latter two formulations P2 and P3, which encapsulate the essence of the Discussion Paper's proposals. Observe that with the parameters chosen, these models may give different emphases. The formulation P2 will tend to produce inexpensive decisions that may be close to the maximum risk level, while the objective of formulation P3 is to minimize the risk, while staying within a budget constraint. Of course both of these formulations are different representations of the bi-objective problem that seeks to find efficient solutions with respect to both objectives (of cost and security).

Treating the Commission's task as an optimization model has some added benefits. In principle, additional constraints can be added to the portfolio model to ensure certain conditions are met. For example to minimize the distortion of any incipient cap contract market, the Commission might seek to avoid paying too much for any particular contract, by constraining it to select no contract that is priced too high above its theoretical option value.

The decision variables in these optimization problems are the components of the Commission reserve portfolio, along with the offering decisions on how to operate the plant corresponding to these contracts. We shall now look at each of these aspects in turn.

3. Contracting decisions

The Commission will contract by tender for a quantity of security reserve generation to be set aside. The proposal is that the Commission will pay the capital costs and maintenance costs plus the fuel costs for any generation that the plant is called upon to deliver. The Commission will have the right to offer the power into the market at a price to be chosen by them. (Since the Commission offers into the market in an uncertain future, the fuel cost per MWHr can be settled at the time of the contract but not the total fuel cost.)

The contract terms will require sufficient fuel to be available to deliver security of supply in a 1 in 60 year. The level of this fuel availability might be open to some dispute, since it depends on the operating policy of the Commission, which in turn might depend on the other components of its reserve portfolio. This will require some careful modelling.

The Commission must choose a portfolio of contracts from a larger collection of contracts tendered by the market. The tender process can have an influence on the outcome. The simplest model is to call for sealed tenders and then to select the optimal portfolio from submitted tenders using an optimization model. However there is an opportunity here to design a more elaborate repeated tender process in which earlier tenders are published and bidders are invited to resubmit sealed tenders in a final round.

This has advantages that are well known in the auction design literature (see Klemperer).

How often should the Commission seek contracts? The term structure of the contracts might vary. Since the portfolio of contracts is intended to provide security for dry winters, it seems inadvisable to offer these for less than one year. Although one could conceive of a contract for say, June, July, August of any given year, it is likely that the Commission will want control of its plant for hydro firming purposes well before these months. Although there may be different terms of the contracts, it might be useful to allow re-bidding of at least some of the contracts once a year. It seems as if a natural time to do this is in the second half of each year for the following year.

Generators and market participants might tender a complete generating plant to the Commission, or some fraction of a plant. In the latter case they might retain low cost units to offer into the market and tender higher cost plant. More than this, they might restrict their tranches to half a unit capacity and allow the Commission to offer the other half, or even allow the Commission to choose the proportion of capacity to add to the portfolio. In general a generator (or demand-side participant) might offer a collection of different tranches of reserve capacity all at different prices, each of which can be partially dispatched. (Observe that this is different from dispatch of energy in SPD as the reserve need not be dispatched in order of increasing price.) From a modelling perspective, there appear to be no good reasons for precluding any of these options (including all of them makes the portfolio optimization problem a mixed integer (nonlinear) programming problem.

4. Possible components of a reserve portfolio

The Discussion Paper appears to take a very limited view of what the Commission's reserve portfolio might comprise. It is envisaged that the Commission will contract to secure control over a certain type of generation plant, and thereby improve the dry-year preparedness of the industry as a whole. However, inventive Commissioners will soon discover a range of other types of contracts that might serve the Commission's purposes well; we list some of these below. There is no good reason to dismiss any of these possibilities out of hand, and while some may ultimately not be used, the Commission should at least be able to consider them.

1. **Contracts over peaking plant**

This is the type of contract envisaged in the Discussion Paper. In normal years, this plant does not generate at all.

2. **Contracts over hydro-firming plant**

There is some concern that by the time a winter has been identified as "very dry", and peaking plant has begun to run, it might be too late to easily rectify the problem. The Commission might better be able to minimize the risk of shortage if it also had control over some "hydro-firming" thermal plant, which could be made to generate in the late summer and autumn if that were deemed desirable. This would include running in many years that, with hindsight, would turn out to be "normal". (It could be argued that the usual market pressures will be sufficient to ensure such operation, without the

Commission’s involvement. However, this seems not to occur in practice, and indeed the market does not have supply security as an explicit objective.)

3. **Contracts over hydroelectric plant**
The rationale for this would not be to withhold the contracted hydro capacity from the market completely, but to gain some measure of control over its operation, and thereby serve the Commission’s objectives. In dry winters, the Commission would probably wish to see stored hydro-electricity consumed more cautiously than a private owner would, since they have different objectives (shortage-avoiding vs. revenue-maximizing). Note that this reasoning would also apply to seasonal pumped-storage reservoirs, should any be developed – with or without the Commission's involvement – in New Zealand.

4. **Options over new plant**
The steady, low-risk cash flow from a Commission contract would be of great value to a firm seeking financing for the development of a new power station. Such contracts might give the Commission great influence over the future development of the industry. The Commission could, for example, secure for itself the right to nominate the size or location of the new plant, or when construction should start – and thereby ensure that the plant is built larger and earlier than might otherwise be the case.

5. **Contracts over transmission**
While transmission is not the Commission's core business, it will be required to determine whether the grid is capable of meeting the unusual demands placed on it in a very dry winter. If the Commission's modelling indicates that this is not the case, then it might (in the absence of any other investors) find itself obliged to become involved in funding the construction of new transmission lines.

6. **Contracts offering partial control**
There is no reason why a Commission contract should have to be, in effect, an operating lease for an entire power station. The contract could, for example, place only some fraction of the plant's capacity into the reserve portfolio, with the rest being offered into the wholesale market by its owner as usual. Alternatively, a joint operating agreement might be negotiated, with the owner retaining control but agreeing to certain constraints imposed by the Commission.

7. **Demand-side participation**
Finally, there is much concern that the demand side of the market has so far not been included in this package. Indeed in the Q&A document published on May 20, reduction in demand was explicitly precluded from providing security of supply. Demand-side bids for curtailable load would seem to be a sensible thing to include in a reserve portfolio. For one thing, it seems that they could be provided at much shorter notice, and sooner, than new generating plant. Moreover, in many cases it might be less expensive to fund these than to fund a new power station. Furthermore, it is likely that there will be a number of industries seeking to contract with the Commission for curtailable load, who will be able to transfer the risk of curtailment to other parties using the contract payment. For example, a lumber mill might purchase call options on amounts of lumber to be delivered, which can be exercised in the case of electricity

curtailment. Another example might see consumers (of all sizes, from the smallest to the largest) being offered the opportunity to “opt out” of the levy, by compensating them their share of the levy in return for a contractual obligation to save 10% every dry winter.

A possible objection to inclusion of curtailable load contracts in the Commission’s reserve portfolio, is the inability to generate revenue from this to offset the costs of the contract. We see this as being a temporary impediment that will attenuate with the emergence of demand-side exchanges. Here the Commission will be able to recover some of its costs during periods with high prices, by selling the curtailable load back to the market.

In summary the Commission should be broad minded in its approach to ensuring security of supply. Building or contracting existing generating plant seems to be the focus of the proposal. However one can imagine that new technologies might emerge which can offer security of supply in unforeseen ways, and it makes sense not to preclude these from the Commission’s range of options.

5. Operating decisions

In order to determine an optimal reserve portfolio the Commission must determine an offering policy for the plant over which they have control. The specification of an appropriate operational model has two aspects. The first is that the outcomes of a reserve portfolio depend on how this generation is offered into the market as well as the actions of other participants in conjunction with this. A key ingredient in the Commission’s contracting arrangements will be models that will enable them to investigate the market effects of holding a given portfolio and offering this to the market in an optimal way. This will be discussed in section 6.

The second aspect is that the security-of-supply constraint must be evaluated. It is not enough to monitor security in hindsight and adjust the contracting level. The Commission will need a model to evaluate whether a candidate contract portfolio delivers 1-in-60 year security. There are several ways of constructing such a model. The simplest is to simulate a given contract portfolio and operating policy through 60 historical inflow sequences. Observe that the policies of other agents in the market need to be simulated as well. They will also depend on the contract portfolio of the Commission (firstly because they may have contracted reserve capacity to the Commission and so have less to play with, and secondly because the Commission’s contract portfolio and offering strategy will affect the behaviour of market participants to maximize their profits.)

The current proposal for an operating policy is to set a (high) price P and offer the generation at this price. The issue of determining the appropriate level of P is a delicate one. The Discussion Document takes the position that P should be chosen high enough so that it does not distort investment signals in the spot market. A high premium is placed on transparency and certainty for the choice of this price.

If P is public knowledge, then there appears to be a incentive for other generators to increase their offer price to \$(P-0.01)/MWHr. The incentives on hydro-generators are also different from those in the absence of a price cap. In the presence of a price cap, a hydro generator might offer consistently at P-0.01 throughout a dry period until it runs out of water. The value of water can never exceed P, so the Commission plant will not run until the water level is zero. This seems to be contrary to the purpose of the Commission.

The realism of the example in the preceding paragraph is open to debate. A mathematical equilibrium analysis shows that in a game consisting of a single trading period, identical hydro-electric generators should offer price-quantity offers that increase linearly from marginal cost to the price cap if they are seeking to maximize their profit subject to an uncertain demand. This Nash equilibrium gives prices for lower levels of demand that are significantly less than the cap P, as generators seek to gain more market share by offering at lower prices. It is not clear whether this behaviour makes sense in the repeated game (over many trading periods) where generators can learn from competitors that undercutting each other is contrary to all of their interests in the long term. These effects have been studied in detail in the game-theory literature (see e.g. Fudenberg and Levine).

To make it more difficult for existing generators to extract profit from strategic bidding, a high premium on certainty (in the Commission's offering policies) might be exactly what is **not** required here. Even a secret but fixed offer price will soon become transparent to market participants with good models of the dispatch system.

A second issue here is that the choice of operating policy will have an influence on the revenue that a plant can generate, as well as on security of supply. One can illustrate this using a simple model for pricing cap contracts. Consider a contract for 1 MW of capacity at \$K/MW per year with a fuel cost of \$c/MWHr. Suppose the Commission sets their offer price at \$P/MWHr. Thus the contract gives them the right to buy 1MWHr at any time at \$c and offer it at \$P. The contract will be in the money whenever the spot price p(t) exceeds \$P. Then the value of the contract is

$$v = \sum_{p(t)\geq P}(p(t)-c),$$

where the sum is made over all the hours in the contract period (here a year). This is like a cap contract at strike price c, except that it is only exercised at price P>c. In a perfectly competitive market, a market participant would pay the larger amount

$$V = \sum_{p(t)\geq c}(p(t)-c),$$

for an annual right to buy 1 MWHr of energy at \$c. This would be the market price of such a cap contract - if V is enough to cover the capital and maintenance cost (K) of the plant then the owner can sell such a cap and fund the plant construction.

Under the policy of offering at \$P, the Commission can only raise v from the market for their contract. The balance V-v must be sought from a levy of the industry. This levy would be less if the Commission were to offer at \$c/MWHr. Thus there will be significant pressure for them to do this, both to lower the levy it must charge, as well as lowering wholesale prices by being dispatched at a lower price.

In summary there are several competing interests. On one hand the Commission seeks to avoid distortions of the investment market by offering at a high price. The Commission is not by design a profit maximizer, but its participation in the spot market means that it must confront the choice whether to attempt to game the spot price. Irrespective of its policy here, other participants will game the price, and so a fixed strategy by the Commission is likely to perform poorly (e.g. because of undercutting by competitors) and result in a higher than expected levy to cover the contract costs. Indeed if it is solving formulation P2, then its objective is to keep the levy as small as possible (by maximizing profit) without violating the security constraint.

On the other hand if the Commission adopts formulation P3 then it will seek to minimize the probability of a shortage subject to keeping the levy within a budget. For example, if hydro lake levels look to be under threat, even though prices are low, the Commission might want to offer some of its generation at less than the plant's fuel cost, so as to ensure dispatch. This is consistent with the purposes of the Commission, though it might produce a higher level of security than is required at the consumer's expense.

Finally, the Commission might choose to ignore market distortions and behave as a competitive fringe, by offering at marginal cost. This will provide some elasticity in the residual demand faced by other generators, which will attenuate any market power they might have. In this case the Commission will be dispatched more often than it expects, but will generate lower spot prices, require a lower industry levy, and be less susceptible by lobbying from the demand side to lower its prices. Observe that this approach carries with it the inherent risks of market distortion in base load generation, with the main concern being that under this approach subsidized Commission generation may displace commercial base-load generation.

To sum up this section, the offering strategy should be thought of as a second stage decision in a two-stage stochastic optimization problem, the first stage of which determines the portfolio components. The second stage offering decisions and contracting decisions should be chosen, in view of the first, to meet the objectives set down by the Commission. It is likely that intervention in the market to meet objectives other than maximizing profit will affect the behaviour of other participants in possibly unforeseen ways.

6. Funding the Commission

The proposed mechanism for funding the Commission is to impose a levy (of up to 0.5c/kWh) on all electricity. This seems a fair and reasonable mechanism. The levy can be collected from generators who pass the costs on to consumers by adding a premium to their offer prices. This means that every offer price will increase by up to $5/MWhr. For some large users of electricity this levy will amount to a significant increase in cost.

A possible alternative might be to impose the levy only on electricity generated at wholesale prices above some threshold. This might help to minimize the distortionary effects during summers and normal winters, when the threshold would rarely be reached.

Under models P2 and P3, the Commission has an incentive to keep the levy they charge as low as possible. This represents an opportunity for the Commission to use their plant to make profits that will offset the levy. Observe however that this might have perverse consequences. For example a Commission generating plant that is dispatched K MWHr on the margin might increase its offer price by, say, $1/MWHr. The increase in profit (assuming the same dispatch) for the Commission is $K, which can be used to offset the levy. The decrease in cost per MWHr consumed will be $K/M <1, where M is the total market load. However, the increase in Commission offer price has increased the spot price by $1, resulting in a net loss for each consumer.

Finally, the design should not rule out the possibility that the Commission might on-sell cap contracts to the market to transfer some of the price risk to purchasers.

7. Alternative approaches to providing security of supply

The proposals in the Discussion Paper are aimed at reducing the occurrence of "dry-year" electricity crises. Although they are open for discussion, these are relatively concrete proposals providing specific new regulatory instruments for providing this dry-year security. In this section we briefly canvas alternative approaches, varying from the least interventionist to the most.

a) No change in policy

Under this regime, the government would continue to rely on the wholesale market to provide security through the wholesale spot price mechanism.

b) Reserve constrained dispatch (1)

Under this regime, the Commission would monitor the level of the hydro lakes and restrict total release from these lakes when their levels were considered to be too low in view of observed weather conditions. The restrictions would be in the form of a security constraint, added to the market dispatch software SPD, placing an upper bound in each half hour on generation from hydro electricity. Different bounds could be chosen for each reservoir and trading period, subject to the constraint that they add to the total maximum hydro release. Alternatively a weekly bound could be provided to each reservoir with financial penalties imposed for exceeding this. However the lightest regulation would have a bound on total release for each half hour, to be divided between the stations based on price.

A regime based on this proposal would not involve the Commission acting as a market participant. Reserve constrained dispatch mirrors the current wholesale market structure for spinning reserve and grid security constraints, whereby Transpower imposes constraints on dispatch and line flows to ensure the security of electricity supply in the short term. The operational costs of such a scheme (involving the addition of simple constraints to SPD) would be minimal. On the other hand there would be some modeling overhead in determining when and how to impose these constraints.

Whenever one adds constraints to a linear program like SPD to ensure some outcome there is typically an increase in dispatch cost, accompanied by an increase in wholesale spot prices. This proposal would tend to increase spot prices. This is because hydro generation that would have been dispatched will now be constrained, and higher-cost reserve thermal plant will be necessary to meet demand.

c) Reserve constrained dispatch (2)

Under this regime, the Commission would monitor the level of the hydro lakes and restrict total release from these lakes as in (b). However in this case the restrictions would be in the form of a security constraint added to SPD, which placed a **lower** bound on generation from thermal power stations.

Similar to b), a regime based on this proposal would not involve the Commission acting as a market participant. A lower bound that constrains thermal generation with no constraints on hydro would have some interesting consequences on the dispatch in cases where the constraint was binding. In the absence of other constraints, the marginal generator in this case is likely to be a hydro-electricity generator. This means that the spot price of electricity (that at the margin) will be less than the offer price of the constrained-on thermal plant. Some compensation for the constrained-on plant will need to be found. In some cases this will exceed the rental earned by the system operator, and so revenue injection to the market is necessary to fund these payments. In summary, a model such as this will lead to lower wholesale prices than (a) but require extra funding (possibly in the form of a levy) to meet constrained-on payments.

d) Commission makes a market in cap contracts

Under this regime, the Commission would participate as an intermediate agent in the contract market. Their function would be to assist the negotiation of contracts between generators and purchasers to encourage the building or retaining of peaking plant, when the contract market fails to deliver these without Commission intervention.

An example might be an offer of an appropriate cap to a group of large consumers, at a contract price that is less than that demanded by generators. The cash flow from this cap, with an extra contribution from a levy, could then be used to purchase the cap from the owners of the peaking station (at a level that is sufficient to pay for its capital, and operating and maintenance costs.) This model is similar to the current proposal, except that the role of the Commission is to act merely as a financial intermediary to provide some liquidity in the contracts market rather than to offer as a generator-class market participant in the wholesale market.

e) Commission contracts reserve plant with constrained offering

This is similar to the regime proposed by the Discussion Document. Under that regime, the Commission contracts with generators to provide reserve capacity that the Commission offers to the market according to some fixed rules. Here the generator retains control of the plant but is constrained in their offering strategy. For example a generator might be (lightly) constrained by the Commission specifying the maximum offer price for contracted plant. When prices are below this level, this leaves the generator free to offer the generation to the market however it sees fit.

In comparison with the Discussion Document proposal, regimes like (d) and (e) involve less intervention by the Commission in the wholesale market, where prices (when they are below the threshold) will be determined solely by the actions of the other participants. On the other hand they give the Commission less control over conservation of stored water in anticipation of a potentially dry winter. This seems to require a level of early intervention that is beyond the capability of (d) and (e) to deliver.

f) Centrally operated electricity market

The greatest control over security of supply is afforded by a centrally-operated electricity market, in which a single operator is assigned the task of dispatching all plant to meet current demand, while constraining the dispatch so as to ensure security of supply over a dry winter.

Appendix A Modelling

The decisions and tradeoffs facing the Commission areas are not easy. They will require the analysis of large amounts of information, and the solution of complicated optimization and market equilibrium models. Although there will be some trade offs to be made between model simplicity and veracity, there will be significant pressure on the Commission to work with state-of-the-art modelling techniques. In this appendix we review some of the modeling problems that the Commission will need to address.

Verifying security of supply

The Commission will need appropriate tools to determine whether its contract portfolio and operating strategy is likely to meet the 1 in 60 supply constraint.

The simplest approach simulates candidate policies using historical inflow data. If the policy is run over each year of 60 years of inflow data without shortages then one is tempted to deduce that it meets a 1 in 60 year security constraint.

There are two weaknesses with this approach. If the candidate policy has been determined by using the 60 years of data (or even parameters estimated from it) then this method will tend to give a lower (optimistic) estimate for the probability of shortage, as the policy has been adapted to the 60 years of data.

A second weakness in this approach is that in assuming that each historical year is an independent sample, the method ignores possible long-range climatological effects such as the ENSO. The observed weather conditions at the time that the reserve portfolio is decided should have an influence the choice of portfolio, if the following year's inflows are correlated with them.

An alternative approach is to construct synthetic series of inflows using an appropriate hydrological model. This model must account for long-range climatological effects as well as the random effects of weather on inflows. Such a model can be calibrated to historical data (although there are probably too few data points to estimate long range effects confidently) and used to generate appropriate scenarios to use in a simulation of candidate policies. The advantage of this approach is that millions of scenarios can be constructed, giving more accurate estimates of the probability of shortage arising from a give policy.

Contract portfolio optimization

We focus on the formulation P2. This seeks to minimize the levy subject to meeting the 1 in 60 year security constraint. This is a two-stage stochastic optimization problem. In the first stage the Commission selects a portfolio, and in the second stage it operates this optimally.

The first stage decisions in this model are to determine which offered contracts (or parts of a contract) to include in the reserve portfolio. The first-stage cost is the total offered cost of this contract. Observe that these costs need not be additive. For example a generator might offer a bundle of two plants together for a different amount than the sum of each when offered separately.

The portfolio can be expressed a vector *z* of 0's and 1's for indivisible contracts, and a vector *x* of numbers between 0 and 1 for the remaining contracts. Some of these contracts will be mutually exclusive (as they might be different combinations of plant). In the absence of security constraints the least-cost portfolio can be computed using standard mixed integer programming software.

Observe however that any choice of portfolio *(x,z)* places capacity constraints on what the Commission might offer to the market in the second stage. Therefore the true cost of the portfolio (assuming model P2) to be recovered from the levy is the payment for *(x,z)* minus the profit that will be earned from offering generation to the market subject to these capacity constraints as well as the constraint that a shortage is avoided with higher probability than (1/60).

In summary the solution to the first-stage contracting problem affects what is possible in the second stage. Also, the second stage solution affects the amount of levy to recover, and so it must influence the Commission's contracting decisions. It is pertinent to look at how one might solve the second-stage problem with a given portfolio *(x,z)*.

Operating model

The second stage decisions is the offering policy chosen using a given portfolio *(x,z)* to maximize the profit that the Commission can make from offering the generation while satisfying the 1 in 60 year security constraint. This is a very challenging problem to solve, because the offering decisions affect the behaviour of other market participants.

To understand how this might work consider first a model in which the Commission has no security constraint to ensure. Then it would offer its generation to the wholesale market as a price-making generator, using the same models that are currently used by (some) generators. For example it might seek to compute a dynamic Cournot equilibrium (see e.g. Scott and Read) over the period of the contract and then follow the policy that this recommends for its plant.

In years where there is a surplus of energy, it may be the case that such a policy in concert with the policies of other generators meets demand in all but the most extreme scenarios. It might be desirable, if this constraint is over satisfied, to rebalance the reserve portfolio by reducing *x* or *z*, releasing reserve capacity back into the market.

In years in which there might be a shortage of energy, the Commission must ensure that the probability of shortage is low. It then acts as an agent with a dynamic constrained strategy: to keep the probability of shortage below a certain level. In the Scott and Read context this would involve all generators optimizing current and future profit by offering quantities at a given reservoir level, while the Commission offers subject to the constraint that the dispatches resulting lead to reservoir levels in the next period that have a low probability of shortage.

The other generators in the market will recognize this and offer accordingly, for example if energy is in short supply, the Commission will ensure dispatch by offering low. The optimal response of other generators is to offer as if they were meeting a shifted demand function. This will in general result in lower prices.

In some situations when evaluating a candidate reserve portfolio prior to contracting, there will be no offering policy that will meet the security constraint. This means that the portfolio must be expanded. How to do this in an optimal way is not clear. One must change x and z so that the portfolio is feasible for the second-stage problem.

In summary, the Commission requires a model that can be used to evaluate candidate contract portfolios when viewed as a two-stage optimization problem. This requires the solution to a dynamic stochastic game theory model over the term of the contract. Adjusting the first stage contracting decisions in a "what-if" type analysis will allow the Commission to investigate a small number of portfolio choices and choose the best.

Stochastic linear programming

An alternative approach to the above is to treat the second-stage problem as a multi-stage stochastic linear programming problem. There are a number of codes that have been developed to solve these models; SDDP (Pereira and Pinto) is a well-known application to hydro-electricity generation. In models such as SDDP thermal generators offer at their fuel cost, and the objective is to meet demand with minimum cost. Failure to meet demand is modelled with high penalty costs.

Stochastic linear programming models assume a centrally-run electricity system, and so they do not attempt to model participant behaviour. The coordination of the hydro and thermal plant in the system is directed towards security of supply first and cost second. In this sense the output gives a worst case. If security is unattainable from a stochastic linear program, then it will not be delivered by any other choice of reservoir releases, including those arising from generator's choices of offers in a wholesale market.

It follows that stochastic linear programming models can be used to identify if enough plant is available to avoid an electricity shortage if it were dispatched centrally. This might be used as guide to inform the Commission as to whether new plant is definitely needed in the generation mix. However unless the Commission acquires control of all the generating plant in the market, they cannot be sure by using only these models that they have enough capacity to meet the security constraint. Of course, as another test they can determine an optimal dispatch policy for the plant that they do control, assuming fixed policies for other agents, but this will only determine security of supply under this scenario of behaviour.

A compromise between the equilibrium models and stochastic linear programming is provided by simulation-based optimisation models (see e.g. BOOMER developed by Pritchard et al.) These models sample from appropriate distributions for generators offers and loads and then compute an optimal offers for a selected generator.

References

Fudenberg, D. and D. Levine. *The Theory of Learning in Games*, MIT Press, 1998.

Klemperer, P.D. What really matters in auction design. www.paulklemperer.org, 2001.

Pereira, M.V.F. and L.M.V.G. Pinto. Multi-stage stochastic optimization applied to energy planning, *Mathematical Programming*, 52, 1991.

Pritchard, G., Philpott A.B., Neame P.J., and Zakeri, G., The electricity sellers' Psi, *Proceedings of the 36th Annual Conference of ORSNZ*, 278-287, 2001.

Scott T., and E. G. Read. Modelling Hydro Reservoir Operation in a Deregulated Electricity Market, *International Transactions in Operations Research*, 3, 1996.