On models for estimating the effect on prices of CO_2 charges[1]

Andy Philpott
Electric Power Optimization Centre
The University of Auckland

Introduction

The Climate Change Office is seeking methodologies for estimating increases in electricity prices for consumers that would result from CO_2 charges on thermal fuels. Price increases that will be faced by consumers will be mainly due to increases in contract prices. However under the assumption that hedge-contract prices are expected spot prices, increments in contract price can be interpreted as increases in averaged spot price. This makes it important to use state-of-the-art models for predicting increments in spot prices.

The CO_2 charges will be different for different generating technologies and will have an effect on the electricity spot price through the offers of the thermal generators that incur the charges. However the spot price is determined by the combination of offers from thermal plants and renewable energy plants (who are not subject to the charge). This makes the determination of the spot price increment a challenging problem.

There are two general approaches to attacking this problem. The *bottom-up* approach endeavours to model the electricity price as arising from the actions of agents whose offers are made to optimize profit within a physical electricity system that is represented as accurately as possible within the model. In contrast, the *top-down* approach forms econometric relationships between price and input costs, and estimates the parameters of these models from historical data, typically using regression analysis. The resulting models can then be used to predict spot price increments based on estimates of fuel cost increases. These models have certain disadvantages in that it is difficult to estimate the effects of structural changes in the electricity market from historical data. Nevertheless they can serve as a useful benchmark.

This report focuses on the bottom–up approaches that have been suggested for attacking this problem, and comments on their strengths and weaknesses. We have included two appendices to this report. Appendix 1 describes some particular bottom-up models in more detail, and Appendix 2 presents some analysis based on a Nash-Cournot model of generation.

---

[1] This paper was commissioned by Norske Skog in October 2004, and is posted on the EPOC site with their permission.

Bottom-up models

The bottom-up models are essentially of two types. The first type of model seeks a *system-optimal* solution, as if the generation of electricity were centrally planned and we were seeking a plan of least cost. In this model the generators are assumed to make offers at their short-run marginal cost, and the dispatch delivers the perfectly competitive outcome. The major difficulty in this methodology is that price predictions based on offers at short-run marginal cost are likely to underestimate the true prices, because market participants often offer generation above short-run marginal cost. For example the generators who operate renewable energy plants under a carbon tax might seek to benefit from the increased costs of competitors and increase their offer prices.

To overcome this drawback the second type of model seeks a set of *user-optimal* policies that form an equilibrium to a non-cooperative game. In the context of electricity markets where the possibility exists to exercise market power, these models could be viewed as being more realistic than the system-optimal models, but constructing realistic game-theory models is much more complicated, both conceptually and computationally.

The most popular game-theory models assume generators offer quantities in a Nash-Cournot game. (An analysis of a simple one-shot model of this type is carried out in Appendix 2.) This approach is problematic for a number of reasons.

(1) A demand function must be estimated. This is often done by calibrating the model on historical price data to estimate an imputed demand elasticity. Since in the short term prices appear to respond to outages, climate forecasts, and agent perceptions, it is not clear that this approach will give reliable results.
(2) The models suppose that agents make single-quantity offers when in reality generators offer supply functions to meet an inelastic demand. The forms of the equilibria obtained are different for these two paradigms. Unfortunately supply-function equilibria are very difficult to compute (or even prove existence thereof) for realistic situations.
(3) We have assumed that electricity contract prices are equal to expected spot prices. It is not clear that this assumption is valid in a situation in which market power can be exercised. For example generators have an incentive to manipulate the spot price to improve their chances of obtaining a high contract price, then offer at low prices once contracted. This would tend to make contract prices exceed the expected spot price. (In a liquid market one might expect this discrepancy to be exploited by arbitrageurs, but even they might be wary of competitors who can influence prices.)
(4) Agents in electricity markets play a game that is repeated in every trading period. Repeated games produce outcomes that can be quite different from one-shot games. For example, by learning from past mistakes, and signaling to competitors over time, agents might converge on tacit collusion strategies that are not one-shot equilibria but give higher payoffs to both players. Modeling these effects accurately is very difficult.

A Discussion of System-Optimal Models

Since in most normal circumstances generators who are optimizing will offer their energy at no less than marginal cost, system-optimal models can be used to obtain bounds on the increment in electricity prices that would result from carbon charges. Let *p*(SRMC) be the price obtained by running a model based on short-run marginal cost offers without a CO_2 tax, and let *p*(SRMC+ CO_2) be the price obtained by running a model based on short-run marginal cost offers with a CO_2tax. Similarly let *p*(T) be the true price without a CO_2 tax, and let *p*(T+ CO_2) be the true price with such a tax.

We know that in the absence of unusual circumstances that might depress short term prices (e.g. many overcontracted generators) we expect short-run marginal cost offers to give lower prices than those observed, and so

$$p(\mathrm{SRMC}) \leq p(\mathrm{T})$$

and

$$p(\mathrm{SRMC}+ CO_2) \leq p(\mathrm{T}+ CO_2).$$

Observe that this does not necessarily mean that

$$p(\mathrm{SRMC}+ CO_2) - p(\mathrm{SRMC}) \leq p(\mathrm{T}+ CO_2) - p(\mathrm{T}).$$

Furthermore, if the cost of generation of some agents increases by some factor then we cannot assert that the clearing price *p* cannot increase beyond this factor. In other words it is not necessarily true that

$$p(\mathrm{T}+ CO_2) - p(\mathrm{T}) \leq p(\mathrm{SRMC}+ CO_2) - p(\mathrm{SRMC}).$$

This inequality depends on the offer behaviour of the generators. In Appendix 2 we show using a one-shot Cournot model that prices in equilibrium might increase by more than the factor or by less than the factor depending on the demand elasticity.

A (weak) bound that we do have from this analysis is given by the inequality

$$p(\mathrm{T}+ CO_2) - p(\mathrm{T}) \geq p(\mathrm{SRMC}+ CO_2) - p(\mathrm{T}).$$

The application of this inequality for prediction is not straightforward as it requires an estimate of *p*(T). We can only really apply it to historical information to determine a lower bound *p*(SRMC+ CO_2) - *p*(T) on the change in price due to a carbon tax if it had been applied during the period being investigated. Observe that this lower bound might be close to zero or even negative.

One proposed estimate for the increment in price is *p*(SRMC+ CO_2) - *p*(SRMC). This might serve as an *estimate* of the price increment but it is not an upper bound.

Choice of System-Optimal Models

The effect of a carbon charge will depend on the levels of demand for electricity and the supply of power from renewable sources. In a wet year it is likely that the effect of the carbon charge (when less coal is burnt) will be less than during a dry year. To accurately assess these effects, multi-period models are needed that trade off use of stored water against the use of thermal fuel.

A discussion of different multi-period models is given in Appendix 1. These vary in their degree of simplicity. Since the output of simple models is easiest to understand, it makes sense to begin any study with the two-reservoir SPECTRA models and proceed to a more detailed model (e.g. with more reservoirs) if it is warranted.

It is noteworthy that all of the models described in Appendix 1 can be configured[2] to represent two reservoirs linked by the HVDC cable, and run in a mode in which all generators offer at short-run marginal cost. This allows a certain degree of validation, by cross checking the results obtained from the same data provided to different models. (The undertakings by the Electricity Commission to make data publicly available is a welcome move in the validation of these models.) It would be prudent of the Climate Change Office to perform such a validation before adopting recommendations drawn from a single implementation. Indeed it would also be prudent to triangulate the predictions of these models with the results of a simple top-down model calibrated to historical data.

[2] For DUBLIN this reconfiguration would require the equilibrium computation to be replaced by a system optimization.

Appendix 1: Models for medium-term hydro-thermal planning

Models for hydro-thermal planning typically focus on handling uncertainties in inflows to reservoirs. The objective is to meet a known demand over some time horizon at least expected thermal cost using a mix of thermal and hydro generation. This corresponds to a central-planning paradigm where profit maximization for generators is not considered. The output of these models gives an indication of the likely outcomes should all agents offer at short-run marginal cost. One must be careful here as the short-run marginal cost of a hydro generator is approximately zero, and under this assumption a dispatch system like the NZEM will dispatch all hydro successively until water storage is zero. On the other hand a centrally-planned system will impute an expected value to stored water that comes from both a thermal cost and (if there is insufficient thermal capacity to meet demand) an estimated shortage cost. Indeed in some jurisdictions (e.g. Chile) the regulator requires the offers from hydro plant to be the values computed from a centrally operated model.

The models alluded to above are optimization models. They compute a policy for each generator over a time horizon. They do not by themselves construct a distribution of outcomes that would occur if the policies were implemented. To do this a model is required that simulates the policies over the time horizon with sampled data. Preferably the inputs should be "out-of-sample" data, that were not used in constructing the optimal policies in the first place. Simulations allow modelers to estimate the probability of shortage or spill, and the probability distribution of price outcomes.

An alternative set of models endeavour to represent strategic behaviour of agents within a market situation. The most common approach is to model generators as playing a Nash-Cournot game in quantities, whereby a postulated demand function serves to set the clearing price in each trading period. A variation on this seeks supply-function equilibria in markets either with or without elastic demand. Though supply functions are a better approximation to how markets actually operate, they are computationally much more difficult than Cournot models, and so are less popular.

In what follows we shall outline the features of the models of which we have some knowledge. This is not an exhaustive list, but will cover all of the hydro-thermal models being considered by the Climate Change Office. There is some difficulty in assessing some of these models as there are few publications in the open literature that give details of they work.

SPECTRA

SPECTRA is a two-reservoir model designed specifically for the New Zealand system. The hydro reservoirs are aggregated into a single reservoir for each island joined by a single line representing the HVDC link. Water value surfaces are computed using a technique called "constructive stochastic dynamic programming" (CSDP) or "dual dynamic programming" (DDP) due to Dr Grant Read and his collaborators (see e.g. [9]). Once water values are obtained, SPECTRA simulates the system using the generation policy determined by the water values along with historical inflow data to obtain storage trajectories and probabilities of spill and shortage. If more detail is required during the simulation, heuristics are available in SPECTRA that disaggregate each reservoir release into separate reservoirs.

In the two-reservoir model SPECTRA computes water value surfaces for every pair of reservoir levels. With random inflows these polyhedral surfaces become quite complicated, and one can imagine that the computations over many stages might require some approximations. Furthermore it is clear that such a constructive approach is difficult to extend to higher dimensions, both for conceptual as well as computational reasons.

The SPECTRA system has been widely tested and used over the last 20 years within the NZED and ECNZ for water valuation, and so is likely to have acquired some degree of reliability, at least within its capabilities. Of course there exists the possibility that the code (which is not in the public domain) conceals some undetected bugs that have not been revealed by testing. The danger of undetected bugs is higher with complicated code with a relatively small user base. SPECTRA constructs polyhedral surfaces, the computational representation of which can be very complicated, so the probability of undetected errors is high. Validation with other codes (like SDDP-PSR) on the same data sets would increase confidence that the code(s) perform correctly.

SDDP

SDDP stands for Stochastic Dual Dynamic Programming, a sampling-based technique developed by Mario Pereira of Power Systems Research Inc. in Brazil. By utilizing a linear programming approximation of the dispatch problem in each stage, SDDP computes optimal policies using a multi-stage Benders decomposition approach that is specifically tailored for hydro-thermal scheduling. The linear programming formulation allows the model to use a detailed representation of the transmission grid, and unlike SPECTRA, a large number of reservoirs.

It is important to discriminate between what I am calling (generic) SDDP, the technique published by Pereira and Pinto in their 1991 Mathematical Programming paper [5], and the implementation of this algorithm developed by Pereira and colleagues at Power Systems Research (PSR). For this reason I have chosen to describe the PSR implementation in a separate section below.

The generic form of SDDP assumes that inflows are serially independent from week to week, and that the transmission system is modeled using DC load flow without losses. This means that the cost-to-go function is a convex polyhedral function that can be represented by "cuts" derived from the dual solutions to linear programs. Moreover cuts can be shared between subproblems at the same stage.

The algorithm as described in [5] constructs a sequence of forward and backward passes. In each forward pass it samples fifty inflow sequences conditionally to give a "horse-tail" scenario tree. The backward pass then constructs fifty cuts for each stage that give a polyhedral function that is a lower bounding approximation of the true cost-to-go function. The next forward passes optimize each stage accounting for the effect of the decision on the cost-to-go represented by these functions. The algorithm can be shown to converge (see [4]) with probability 1 as long as:

(1) cuts are computed in every stage;
(2) there is a strictly positive probability that the cuts computed in each stage include the sampled inflow outcome.

SDDP-PSR

The implementation of SDDP at Power Systems Research has added a number of features not mentioned in [5]. Here we list a few that we have identified from looking at PSR documents on the PSR homepage (http://www.psr-inc.com.br/).

1. Presampled scenarios: The "horse-tail" scenario tree that is sampled in each iteration of SDDP is pre-sampled in SDDP-PSR, and is not changed from iteration to iteration.
2. Line losses: In SDDP-PSR fixed line losses are allowed in the DC-load flow model in each stage. This is acceptable as long as power can be shed at any node. Otherwise a linear programming model is not a valid approximation because it can lead to non-physical circulating branch flows. So SDDP-PSR must either work with no losses, or admit a relaxed model that allows power to be shed at any node (to ensure nonnegative prices). (An alternative relaxation admits circulating branch flows). Observe that including integer variables in the dispatch model to prevent this behaviour is not an option, as this destroys the convexity of the stage problems.
3. Correlated inflows: In SDDP-PSR the inflow sequences can be serially dependent. SDDP-PSR represents these as periodic autoregressive models of order p (PAR(p) models). This represents an advance on the SPECTRA models which have been extended to cover AR(1) inflows (see [8]) but cannot handle more lags. However there is some controversy amongst hydrologists that these models represent the true correlation structure seen in observed sequences. (We note that this criticism can be leveled at all the models included in this Appendix.) Of course in a simulation using synthetic inflow sequences, one should use sequences that most accurately represent reality, even if these must be approximated to carry out the optimization. In the optimization process SDDP-

PSR handles inflow models with quite complicated inflow dependencies by augmenting the state space to include previous inflow observations. Care must be taken here to avoid nonlinear transformations of the data (e.g. an assumption that log(inflow) is PAR(p)) that might destroy the convexity of the cost-to-go function.

4. Convergence criteria: SDDP-PSR produces a lower bound on the cost of an optimal policy. It also constructs a statistical estimate of the cost of the policy it computes. This estimate has a probability distribution that is asymptotically normal (by the central limit theorem). We may estimate the variance of this distribution and apply a Student t test to the hypothesis that the true cost of an optimal policy is within a given tolerance of the lower bound. (In their papers PSR assume a known variance so their confidence interval for the upper bound is narrower than the correct one). SDDP-PSR terminates when the lower bound falls in a 95% confidence interval for the upper bound. Observe that if the variance is high then termination might occur when the candidate policy is far from optimal.
5. It is important to observe that any policy from SDDP generates an estimate of the probability distribution of its cost, along with a lower bound on its expected cost. This should be borne in mind if making policy decisions based on this.

Variations on SDDP

1. CUPPS

The cutting-plane and partial sampling algorithm (CUPPS) was developed by Chen and Powell at Princeton's Castle Lab [1]. Its main area of application has been for logisitics planning under uncertainty. CUPPS is a sampling-based method (similar to SDDP) that performs only forward sampling passes, computing cuts along the way.

2. AND

Abridged Nested Decomposition (AND) was developed by Birge and Donohue [2]. It was applied by Birge to the Colombian power system in the late 90's. AND is a sampling-based method (similar to SDDP) tailored for scenario trees that are bushy rather than long and narrow.

3. ReSa

ReSa is a reduced sampling method produced by Hindsberger and Philpott [3]. It is similar to SDDP, but does not compute cuts at every stage in every iteration. In some instances this can speed up the computation.

Nash-Cournot models

DUBLIN

DUBLIN is a two-reservoir model designed specifically for the New Zealand system. The hydro reservoirs are aggregated into a single reservoir for each island joined by a single line representing the HVDC link. In DUBLIN the reservoirs are assumed to be owned by different agents who offer electricity to the market as players in a Nash-Cournot game. A third set of agents who operate thermal plant can also be included. Observe that in DUBLIN the thermal agents offer strategically only in the single-period game (they do not consider the impact of their strategy on later periods).

The theory of DUBLIN was developed by Grant Read and Tristram Scott in his PhD thesis [7] as part of the Electricity Modelling Research Group (EMRG) at Canterbury and the approach has been promoted by EMRG in a number of areas. His implementation was extended by Mark Craddock while at Putnam-Hayes-Bartlett. Implementations of these models have been developed further and maintained by Orbit (DUBLIN) and PA-Consulting (called EMPACT). Details of these implementations are not available, and the account below is based on Scott's thesis and Scott and Read [8].

As in SPECTRA, the water-value surfaces in Cournot models are created using a "constructive dynamic programming" approach. The papers describing the method speak of adding "demand curves for release" to "demand curves for storage", but this terminology is misleading, as there is only one real demand curve, that describing demand for electricity as a function of price. In each stage of this game, agents offer quantities of electricity in such a way to maximize their individual profit, where the clearing price is determined *a la Cournot* from the demand curve and the total amount offered.

With serially independent inflows, the maximum expected future profit that can be earned by each player in this game at any point in time t can be represented as a function of each of the hydro-generators' reservoir levels. This payoff can be interpreted as a future value of water for each generator (as a function of its own level and that of its competitor). It is then straightforward (in principle at least) to compute a Nash-Cournot equilibrium in a one-shot game at period $t-1$ in which each player tries to maximize his current and expected future profit. This computation is carried out for each possible pair of reservoir levels to compute the maximum expected future profit that can be earned by each player in this game in the previous time period.

Although the DUBLIN model has obvious attractions, its use requires some care. A key input parameter that drives the results is the demand function, and estimating the form of this is very difficult for electricity markets. Moreover when demand is inelastic (like it is in the short term) the prices obtained from these models are very sensitive to the elasticity, so accurate estimation is important. Some of the other drawbacks of this model are mentioned in the body of this report.

References

[1] Chen, Z.L. and Powell, W.B. "Convergent cutting-plane and partial sampling algorithm for multi-stage stochastic linear programs with recourse", JOTA, 102, (1999) 497-524.

[2] Donohue, C.J. "Stochastic network programming and the dynamic vehicle allocation problem", PhD dissertation, 1996, University of Michigan.

[3] Hindsberger, M. and Philpott, A.B. ReSa: a method for solving multi-stage stochastic linear programs, submitted to EJOR, 2003.

[4] Linowsky K. and Philpott, A.B. "On the convergence of sampling-based methods decomposition algorithms for multi-stage stochastic programs", to appear in JOTA, available from http://www.esc.auckland.ac.nz/Organisations/EPOC/SDDP.html

[5] Pereira M.V.F. and Pinto L.M.V.G. "Multi-stage stochastic programming applied to energy planning", Math Programming 52 (1991) 359-375.

[6] Pereira M.V.F. , Campodonico, N. and Kelman R. "Application of stochastic dual DP and extensions to hydrothermal scheduling", PSRI TR 012/99 (1999).

[7] Scott, T.J. Hydro-reservoir management for an electricity market with long-term contracts, PhD thesis, 1998, University of Canterbury.

[8] Scott, T.J. and Read E.G. "Modelling hydro-reservoir operation in a deregulated electricity sector" International Transactions in OR, 3, (1996), 243-253.

[9] Yang, M. and Read, E.G. A constructive Dual DP for a reservoir model with correlation, Working paper, EMRG-WP-99-01, University of Canterbury,1999,.(see http://www.mang.canterbury.ac.nz/research/emrg/wp-99-01.pdf)