

DECISION AIDS FOR SCHEDULING AND HEDGING (DASH) IN DEREGULATED ELECTRICITY MARKETS: A STOCHASTIC PROGRAMMING APPROACH TO POWER PORTFOLIO OPTIMIZATION

Suvrajeet Sen
Lihua Yu
Talat Genc

Raptor Laboratory, SIE Department,
University of Arizona,
Tucson, AZ 85721, U.S.A.

ABSTRACT

The DASH model for Power Portfolio Optimization provides a tool which helps decision-makers coordinate production decisions with opportunities in the wholesale power market. The methodology is based on a stochastic programming model which selects portfolio positions that perform well on a variety of scenarios generated through statistical modeling and optimization. When compared with a commonly used fixed-mix policy, our experiments demonstrate that the DASH model provides significant advantages over several fixed-mix policies.

1 INTRODUCTION

Deregulation is an evolving process. In many states (including Arizona), the major electricity producers have the responsibility of meeting a certain “native load” which constitutes the regulated portion of the business. Beyond this regulated native load, a power producer may buy or sell power in the wholesale electricity market in a manner that the producer finds profitable. Prior to the emergence of electricity markets, profitability was determined simply by the ability of a power producer to convert fuel into electricity in a least-cost manner. Hence minimization of generation costs provided the appropriate strategy. With the emergence of wholesale electricity markets, a utility can manage its power production and revenue potential by trading within this market. A forward (contract) for power is a financial instrument that allows a power producer to buy or sell power for delivery on a future (maturity) date at a price that is agreed upon several months earlier. As weather patterns, economic activity, and market prices evolve, these power portfolios can be rebalanced so as to maximize expected profitability, while appropriately balancing risk exposure. In this environment, judicious decision-making can mean the difference between survival and demise of a power company.

The DASH model for Power Portfolio Optimization provides a tool which helps decision-makers coordinate production decisions with opportunities in the wholesale power market. Before providing the technical details of our approach, we provide a brief outline of some of the major determinants of profitability in electricity markets. Following this description, we describe statistical models that are used for scenarios used within the stochastic programming model. The latter model consists of a financial sub-model and a generation sub-model which are used to determine the profitability of any portfolio position. We also describe an alternative investment strategy based on a certain type of “fixed-mix” policy which is commonly used by electricity traders. This strategy provides a “base-case” against which we compare the results of a stochastic programming model. Our results are based on data obtained from Pinnacle West Capital, which is a holding company for Arizona Public Service, the largest investor owned electric utility in Arizona. In order to maintain confidentiality of their data, our results will be presented in terms of percentage gain. The backcasting experiment, which covers a five month operating period from January 2001 through May 2001, shows a monthly advantage of approximately 7% in favor of the stochastic programming approach. The DASH model has also been tested against a variety of synthetic scenarios. These experiments reveal the robustness of the forward decisions recommended by DASH.

2 SCOPE OF THE DASH MODEL

To begin with, we outline the manner in which we expect the decision process to unfold. At the start of each month, financial analysts/traders for the producer wish to reevaluate/rebalance their power portfolio. At this point, they may invoke some decision model (e.g. DASH) which recommends the mix of power products that the producer ought to hold. While the decision model itself may be dynamic (as in DASH), the trader only commits to a recommenda-

tion for the current month. After the appropriate rebalancing trades are executed, the traders wait and observe the market until the end of the month, at which point, they update the decision model by “rolling the horizon” forward, and providing up-to-date information to the decision model which then provides an updated recommendation for the next month. While it is possible to use the DASH model at decision-epochs that are less than a month long, the portfolios within DASH are represented at monthly intervals.

Market modeling is another feature incorporated within DASH. In some cases, power producers trade electricity in multiple markets. For example, a California utility may trade in Palo Verde (AZ) and the California-Oregon Border (COB). For the sake of this model however, we will consider only one market for electricity. In addition to electricity, the model also allows interactions with one natural gas market. On the generation-side, the unit commitment decisions are made on a weekly basis, and allow us to incorporate heat-rates, start-up costs, minimum downtimes, etc. The current model does not accommodate hydro generation, although this extension is currently under consideration. Some of the modeling issues related with the above features are discussed below.

2.1 Electricity Demand

In a completely deregulated market, the traditional notion of load takes a back-seat to demand-curves relating prices and quantities. However the extent of deregulation is in a state of flux in most states in the U.S. For instance in Arizona, retail tariffs are regulated by the state Corporation Commission and are held constant over long periods of time. Electric utilities are required to serve the “native load” that arises from their customers at regulated retail rates. There are several different demand models that have been studied in conjunction with current the DASH model, including time-series that use temperature as one of the main factors. In more humid climates, we expect that humidity will play an important role as well (Feinberg 2002).

2.2 The Wholesale Electricity Market

The current market model allows electricity forward contracts, and spot market activity. While the current model does not accommodate options, these can be included without adding to the computational burden of the current model. While prices in the electricity market (especially the spot market) vary on an hourly basis, we have discretized time according to a sixteen hour “on-peak” period, and an eight hour “off-peak” period for each day.

2.2.1 Forward Contracts for Power

For the purposes of our model, forward contracts will be assumed to be “monthly”, so that planning for period t re-

fers to some month t in the future. Note that the megawatts committed (bought or sold) to the market in period j influences the total electricity generated during period t , $t > j$. To facilitate profit-making, trading decisions must consider future load projections and generation capacity, both of which are subject to uncertainty. If the decisions for the delivery month (t) could be treated independently of other months, then one could develop a model that could treat each delivery month independently. However, such an assumption might expose the firm to far greater risk level than might be acceptable. This is because (financial) risk exposure of a firm depends on the mix of instruments in its portfolio at any point in time. Hence, it is not sufficient to simply consider profitability for a delivery month; the collection of forwards held at any point in time is an important determinant of risk exposure.

The current price of any forward contract is usually assumed to be known. However, forward prices for each delivery month will evolve over time until the delivery month. As one might expect, this evolution is uncertain on the decision-making date. In the current version of the DASH model, we use a non-parametric approach in which historical data is used to create a vision for the future (e.g., the next six months). This vision is based on creating a number of scenarios of “returns” (percentage change in prices) which may be revealed in the future. The actual process of developing these scenarios is discussed in the next section.

2.2.2 Spot Market for Wholesale Power

As with forward contracts, “on-peak” and “off-peak” power have different price trajectories, and are modeled separately. However, there are two important observations in modeling the spot market. The time scale for spot prices can be hourly. In the interest of computational tractability, we treat spot market on a daily basis, and allow it to fluctuate according to the sixteen-hour “on-peak” and eight-hour “off-peak” periods. Also, the spot prices for each day (d) during the month (t) must be correlated to the forward prices associated with the scenario (s) that unfolds.

2.3 Unit Commitment

The technological constraints of this socio-technical model arise in the unit- commitment problem. Traditionally, unit commitment models are used to determine a short-term (weekly) power generation schedule. While they have also been used to estimate annual production costs, the deterministic nature of the original models (e.g. Bertsekas et al. 1983) do not lend themselves to mid- and long-term analysis. More recently, these models have been extended to accommodate uncertainty in load forecasts, fuel prices, etc. (Takriti, Birge, and Long 1996), (Takriti, Krasenbrink, and Wu 2000), and (Nowak and Romisch 2000). Recent ad-

vances in unit commitment models are summarized in the edited volume by Hobbs et al (2001).

The models mentioned above are typically focused on a short-term scheduling issue (a week or two at most). Due to the medium-term nature (i.e., one year) of many financial instruments, it is difficult to measure their impact using short-term models. Our approach integrates the unit-commitment model with financial decision-making by including the forwards and spot market activity within the scheduling decision model.

3 STATISTICAL INPUT MODELS

With the exception of the Unit Commitment model, all features discussed in the previous section are represented by statistical models. The main purpose of these statistical models is to help generate a finite number of scenarios which are represented in the form of a tree. A scenario models the evolution of information during the decision process (Birge and Louveaux 1997). It is important to emphasize that our procedures are a combination of statistical methods and heuristics that maintain tractability of the decision model.

3.1 Modeling Electricity Demand

Our load data represents an eleven-year period (1990 – 2000) of hourly loads in an APS service area. Since each day is modeled by “on-peak” and “off-peak” segments, we begin by transforming the hourly data into averages for “on” and “off” peak segments. The hours 6 a.m. to 10 p.m. are considered on-peak, and the remaining hours are considered “off peak.” In order to give the reader a sense of the load data, Figure 1 provides a three year sequence of “on peak” loads. The “off peak” loads also portray similar cyclical and seasonal trends, and these are confirmed by the Kendal-Tau and Turning Point tests.




Figure 1: On Peak Load Data for 3 Years

Based on seasonality of loads as depicted above, we partition the data for a year into four groups. The first has a decreasing trend, the next an increasing trend and soon. For each group/partition, we develop the following model:

$$L_d = \alpha_0 + \alpha_1 d + \varepsilon_d \quad (1)$$

$$\varepsilon_d = \sum_{i=1}^7 \beta_i \varepsilon_{d-i} + \eta_d + \beta_8 \eta_{d-1}, \quad \eta_d \sim N(0,1) \quad (2)$$

where d stands for days of the year. In order to create load scenarios from such a model, we generate standard normal random numbers as suggested in (2). One shortcoming of this process is that we treat load as independent of wholesale prices. We plan to incorporate a correlation between the two in our future work.

For the data set we investigated, the de-trended load (for both on peak and off peak segments) followed ARIMA(7,0,1) for each partition. This is consistent with the study of Dupacova, Growe-Kuska and Roemisch (2000) who examined hourly loads (which can be considered as high frequency data) and concluded that SARIMA(7,0,9) \times (0,1,0) was an appropriate model for hourly loads.

In Figure 2, we provide plots of the remaining residuals, the autocorrelation, partial autocorrelation functions, p-values of Ljung-Box statistics, and qq-normal plot of residuals. Both ACF and PACF of residuals are in the bounds and portmanteau test validates this with the qq-normal plots as well. These diagnostics validate the sufficiency of the approach.

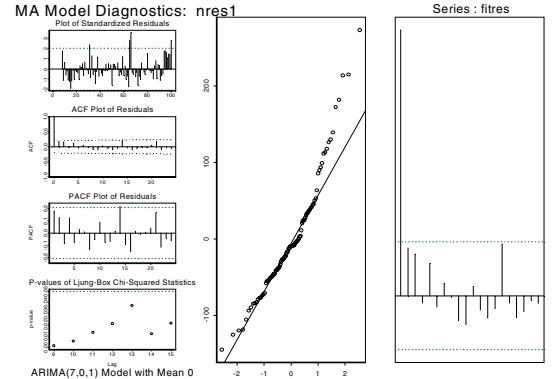


Figure 2: Diagnostic Checking of ARIMA(7,0,1) Fitting of De-trended Load Series, Quantile-quantile Plot, and ACF of Remaining Residuals

3.2 Modeling Electricity Forward Prices

This part of the DASH model forms the core of our scenario generation procedures. The inputs we use are the forward prices for the preceding year, together with recent trends in the market. Let us first focus on the forward prices for the preceding year. These are available as hourly quotes which we transform into “on-peak” and “off-peak” average prices. We have following format for the prices: $\pi_{\tau\kappa e}$, where π is the price, and τ, κ and e denote the contract week, delivery week and segment, respectively. Here, the range of indices are: $\tau = 1, 2, \dots, 52$, $\kappa = \{1, 2, \dots, N\}$, $e \in \{on, off\}$, N denotes the last week in which delivery will happen. For example, $\pi_{1,60,on}$ is the price (\$/Mwh) on

January 7th (i.e. end of week 1) for on-peak power delivered starting on March 1st (for the entire month of March). However we use “returns” to predict prices; that is, $r_{\tau,K,e} = (\pi_{\tau+4,K,e} - \pi_{\tau,K,e}) / \pi_{\tau,K,e}$. The quantity 4 above reflects the assumption of four weeks in a month.

There are two important reasons behind this choice (of using returns over prices). First, this approach allows us to treat different power contracts (associated with different months) with the same scenario tree, thus reducing the complexity of modeling the evolution of prices associated with each type of contract. We have empirically verified that it is the interval of time between contract and delivery that is important for modeling returns, and not the actual contract. Hence the same scenario tree remains a valid representation of returns for alternative contracts. Secondly, the econometrics literature recommends that “returns” are better for predictive purposes because empirical evidence suggests that they appear to have better properties (e.g. stationarity) from a computational point of view (Taylor 1986).

A discrete scenario tree may now be formed by grouping returns into subsets for each period (i.e. month), and modeling the return process as one that allows probabilistic transitions from one subset to another, over time. In order to maintain computational tractability, we consider only two subsets in each period: “High” and “Low” return states. Thus, the resulting scenario tree can be represented by a binary tree in which the returns can assume “High” or “Low” values over the course of the decision process.

To assign “High” and “Low” values for the return states, we adopt a sampling-based procedure that is guided by recent observations of the return series. The nominal value that we assign to each state (“High” or “Low”) is the median of the corresponding group for that period. However, without accommodating extreme values, the scenario tree (and consequently the decisions themselves) overlooks extreme events opening up the possibility for catastrophic losses. We will of course, include some loss constraints within the decision model, but in the absence of extreme scenarios, such constraints can only have limited impact. Accordingly, we use a combination of medians and extreme values (“Min” and “Max”) to assign values to the High and Low states. The precise manner in which we choose one or the other depends on a heuristic guided by market conditions prior to running the model.

Finally, the formation of the scenario tree requires a specification of transition probabilities between nodes which represent information states. Recall that our scenario tree is binary, and hence there are only two probabilities that need to be specified. In the event that our heuristic produced two nodes that are represented by medians (High and Low respectively), then we simply use equal conditional probabilities for these two transitions. On the other hand, if the heuristic produces an extreme value for one path, and a median for the other, then we associate a conditional probability of $\frac{1}{4}$ for the extreme value, and $\frac{3}{4}$ for the median value. Our

heuristic does not produce two extreme values from any node, and hence this possibility is not considered.

The above process creates a binary scenario tree for the return series, which is then used to create prices scenarios that are used within the stochastic programming model described in the following section.

3.3 Modeling Gas Forward Prices

The process used to model gas forward prices is similar to the process described in the previous subsection (on electricity forward prices). We will also assume that the returns for gas and electricity are perfectly correlated so that a scenario obtained from the electricity forwards returns tree generates a similar scenario from the gas forwards return tree.

3.4 Modeling Electricity Spot Prices

Recall that the forward price process is discretized on a monthly basis. However, spot prices must be modeled on a different time scale. As discussed earlier, on-peak and off-peak spot prices will be modeled on a daily basis, with the understanding that they will be correlated with an appropriate forward price scenario. As with the forwards, we resort to modeling the return series of spot prices.

The spot prices during a delivery month are generated from the following formulation (of spot returns): $r_{p_{e,d,t,\omega}} = r_{f_{e,t,\omega}} + \sigma_{e,t} z_{d,t,\omega}$, where ω is the node number of the forward scenario tree, $\sigma_{e,t}$ is the standard deviation of spot returns which changes from delivery month to delivery month, and $r_{f_{e,t,\omega}}$ is daily equivalent of the forward return on node ω for month t . The quantity z represents a standard normal random variate. Here $\sigma_{e,t}$ may be interpreted as the volatility associated with on-peak and off-peak returns during month t and are estimated using a GARCH (Generalized Auto Regressive Conditional Heteroskedasticity) model (Hamilton (1994)). Because the expectation of spot market prices may be assumed to equal the expected forward prices (Hull 1997), the above relationship between spot and forward returns captures both the first as well as second moments of the spot price process.

4 THE DECISION MODEL

The DASH model may be classified as a multi-stage stochastic integer program which recommends forward decisions on a here-and-now basis, whereas, the operational decisions (generation, spot market activity etc.) are used to evaluate the viability of the portfolio. In this sense, the generation and spot market decisions are adaptive (i.e. wait-and-see), and allow us to compute medium (six months to a year) decisions without being mired in daily (here-and-now) details.

In formulating the stochastic program, all decision variables and parameters are dependent on the scenario. However, in the interest of simplifying the notation we have suppressed this dependence below. We remind the reader that all forwards variables will be required to satisfy the non-anticipativity requirements of stochastic programming (Birge and Louveaux 1997). The formulation is presented in two parts: the financial problem and generation costing problem.

4.1 The Financial Problem

Scenario Independent Parameters

- α : Max liquidity limit coefficient;
 T : Number of periods;

Scenario Dependent Parameters

- $PP_{\tau e}$: Price of power forward for delivery period t , peak e (on/off peak) at contract period τ ;
 PG_{τ} : Price of gas forward for delivery period t , at contract period τ ;

Scenario Dependent Decision Variables

- $FP_{\tau e}$: Power forward for delivery period t , peak e (on/off peak), signed at contract period τ (positive for long position, negative for short position);
 $FP_{\tau e}^+$: Power forward in long position for delivery period t , peak e (on/off peak), signed at contract period τ ;
 $FP_{\tau e}^-$: Power forward in short position for delivery period t , peak e (on/off peak), signed at contract period τ ;
 FG_{τ} : Gas forward in long position for delivery period t , signed at contract period τ ;
 $YP_{\tau e}$: Power forward for delivery period t , peak e held at period τ (positive for long position, negative for short position);
 $YL_{\tau e}$: Power forward in long position for delivery period t , peak e held at period τ ;
 $YS_{\tau e}$: Power forward in short position for delivery period t , peak e held at contract period τ ;
 YG_{τ} : Gas forward for delivery period t held at contract period τ ;

Scenario Dependent Constraints

$$FP_{\tau e} = FP_{\tau e}^+ + FP_{\tau e}^- \quad \forall \tau, e, t \geq \tau;$$

$$YP_{\tau e} = YL_{\tau e} + YS_{\tau e} \quad \forall \tau, e, t \geq \tau;$$

$$YL_{\tau e} = YL_{(\tau-1)e} + FP_{\tau e}^+ \quad \forall \tau, e, t \geq \tau;$$

(Power forward balance in long position at period τ ; for $\tau = 1$, initial position YL_{0te} is assumed given);

$$YS_{\tau e} = YS_{(\tau-1)e} + FP_{\tau e}^- \quad \forall \tau, e, t \geq \tau$$

(Power forward balance in short position at period τ ; for $\tau = 1$, initial position YS_{0te} is assumed given);

$$YG_{\tau} = YG_{(\tau-1)t} + FG_{\tau} \quad \forall \tau, e, t \geq \tau$$

(Gas forward balance at period τ ; for $\tau = 1$; initial position YG_{0t} is assumed given);

$$\sum_{t \in [\tau, T]} FP_{\tau e}^+ \leq \alpha \sum_{t \in [\tau, T]} YL_{(\tau-1)e} \quad \forall \tau, e$$

(Max liquidity limit for long position);

$$\sum_{t \in [\tau, T]} FP_{\tau e}^- \leq \alpha \sum_{t \in [\tau, T]} YS_{(\tau-1)e} \quad \forall \tau, e$$

(Max liquidity limit for short position);

The last two constraints above provide a way control the extent to which a portfolio is allowed to change. These constraints help avoid speculation, thus limiting risk exposure. The less the value of α , the tighter the control on the trajectory allowed by the model is. Finally, there are two important factors required in specifying the financial problem.

- Non-anticipativity constraints require that scenarios which share the same history until period t should be associated with decisions which have the same values until period t . These linear constraints couple decisions from different scenarios, thus allowing a well hedged plan.
- The objective function for the financial problem maximizes discounted expected profits associated with the portfolio. In calculating the profits, we accommodate the generation cost, which is computed via the model discussed next.

4.2 The Generation Problem

With each scenario we associate a generation problem that models power production. This is accomplished via a sequence of biweekly unit-commitment problems. Thus for any scenario, there will be twice as many unit-commitment problems as there are months in the financial model. In this formulation, the generation and spot market variables are allowed to be adaptive. As before, the notation suppresses the dependence on scenarios.

Scenario Independent Parameters

- J : The number of segments in a period;
 Gas : The set of gas generators;
 p(j) : Peak status(on/off) of segment j;
 j(d) : The set of segments associated with day d;
 P : Regulated power price;
 L : Maximum loss limitation for any day;
 H_e : Hours of one on/off peak, H_e =16h for e=on peak, and 8h for e=off peak;
 Q_i : Maximum generation capacity of generator i;
 q_i : Minimum generation capacity of generator i;
 L_i : Minimum up time requirement for generator i;
 l_i : Minimum down time requirement for generator i;
 F(x) : Conversion formula (based on heat rate) specifying the amount of gas needed to produce electricity;

Scenario Dependent Parameters

- W_{ij} : Scheduled outage (W_{ij} =0, if outage is scheduled in period t, segment j for generator i; 1, otherwise);
 ω_{ij} : Forced outage (ω_{ij} =0, if outage is forced in period t, segment j for generator i; 1, otherwise);
 PS_{ij} : Price of power in spot market in period t, segment j;
 D_{ij} : Electricity demand in period t, segment j;

Scenario Dependent Decision Variables

- G_{ij} : Total generated power in period t, segment j;
 C_{ij} : Total generation cost in period t, segment j;
 G_{ij} : Power generated by generator i in period t, segment j;
 U_{ij} : Decisions about turning on/off generator i in period t, segment j (binary variables);
 SP_{ij} : Power exchanged with spot market in period t, segment j (positive for purchase, negative for sale);

Scenario Dependent Constraints

$$YP_{te} + SP_{ij} + G_{ij} = D_{ij} \quad \forall t, j, e=p(j)$$

(Demand -generation- forward - spot relationship);

$$YG_{it} = F\left(\sum_{i \in Gas, j} G_{ij}\right) \quad \forall t$$

(Conversion from gas to power for period t);

$$G_{ij} = \sum_i G_{ij} \quad \forall t, j$$

(Total generated power at period t, segment j);

$$q_i U_{ij} \leq G_{ij} \leq Q_i U_{ij} \quad \forall i, t, j$$

(Operating range for each generator);

$$U_{ij} - U_{i, j-1} \leq U_{it} \quad \forall t, i, \tau=j+1, \dots, \min\{j+L_i-1, J\}, j=2, \dots, J,$$

(Minimum up-time requirement);

$$U_{i, j-1} - U_{ij} \leq 1 - U_{it} \quad \forall t, i, \tau=j+1, \dots, \min\{j+1-l_i-1, J\}, j=2, \dots, J,$$

(Minimum down-time requirement);

$$U_{ij} \leq W_{ij} \quad \forall i, t, j \text{ (Scheduled outage);}$$

$$U_{ij} \leq \omega_{ij} \quad \forall i, t, j \text{ (Forced outage);}$$

$$\sum_{\tau \leq t, e} FP_{ae} PP_{ae} H_e + \sum_{j \in j(d)} [(D_{ij} P - SP_{ij} PS_{ij} - C_{ij}] H_{p(j)} + L \geq 0$$

∀d, t is the period associated with the day d (Max daily loss constraint);

Finally, there is an important factor required in specifying the generation problem. The objective function for the generation costing problem minimizes the generation cost (including the cost in spot market activities) associated with each segment in the planning period.

4.3 The Solution Approach

The stochastic programming model presented above is a very large scale optimization problem. Fortunately, the model is amenable to solution using decomposition techniques. Our approach decomposes the stochastic program into three interrelated optimization problems which are motivated by a nested column generation (i.e. Dantzig-Wolfe) type method. The three problems may be summarized as follows.

1. We use a master problem to enforce non-anticipativity restrictions. Each scenario is represented by a collection of columns in this problem, and its goal is to find a convex combination of columns of each scenario that also satisfy non-anticipativity restrictions. Initially, a Phase I problem is solved to obtain a feasible solution to this problem. We note that the objective function coefficient for each column in this problem represents the total profit under a particular scenario of forwards prices, spot prices, and electricity demand
2. Given the price for achieving non-anticipativity from the master problem described above, a mid-level coordinating problem is formulated to make

the best forward decisions for each scenario. This is essentially the same formulation as the financial problem described in section 4.1. However, the summation of forward decisions for a certain delivery period is once again represented via a convex combination of forward columns that are generated by generation costing model where the summation of forward decisions appear in the demand constraint.

3. Finally, the lowest level problem, which generates the aggregation (i.e. summation) of forward columns for the higher levels, consists of a series of biweekly unit-commitment problems. As with the level 2 coordinator, this problem assumes that the scenario is given, and a series of deterministic instances of the unit commitment problem are solved. The prices of forwards in this model are modified by the dual prices from forward balance constraints in the mid-level coordinator.

The details associated with the above mathematical programming procedure will be the subject of a forthcoming paper, and providing further details here would not be germane to the experimental study provided in this paper.

5 EXPERIMENTAL RESULTS

In this section we present experimental evidence that the stochastic programming approach provides significant advantages over other commonly used policies for hedging. On such policy is a “fixed-mix” policy, which in this industry may be described as follows.

On any contract date, an appropriate hedging position for a future delivery date (month) is one that is determined according to the following strategy. Make a prediction of expected demand and expected capacity for the delivery month. If expected demand exceeds expected capacity, then assume a long position for forwards in that delivery month, and the quantity of this transaction should be a fraction “ f ” of the difference. On the other hand, if expected capacity exceeds expected demand, then one should assume a short position for forwards in the delivery month being considered. Once again, the quantity associated with this transaction should be a fraction “ f ” of the difference.

One can devise several variations on this scheme. For instance, instead of using expected demands and capacities, one may use scenarios to determine scenario-dependent strategies, and then use some weighted averaging to determine the exact mix. For our experiments we only tested the basic scheme outlined in the previous paragraph. However, we ran our simulations using several values of the fixed-mix fraction f , including 0, 0.1, 0.2, 0.3 and 0.4.

5.1 The Backcasting Experiment

As outlined in the introduction, this experiment covers a five month operating period from January 2001 through May 2001, with hedging decisions being made once each month. The decisions at the beginning of each month are, of course, made prior to observing the markets. Once the transactions are carried out, no portfolio changes are allowed for the rest of the month. During this period, we do run a generation costing simulation based on biweekly unit-commitment. At the start of the next month, we once again use the fixed-mix policy to obtain the newly rebalanced positions, and the process resumes again. For all runs reported here, we used an initial position of forwards amounting to 15% of the averaged electricity load for a certain period. The electricity market data for our study reflects prices at Palo Verde, AZ, whereas, the gas market prices reflect data from Henry Hub, LA. The hedging decisions made in this study allowed delivery dates up to six months in the future. For the sake of this study, transaction costs were not included, although such calculations are easily accommodated within a simulation. Moreover, since all rules carry out the same number of transactions, the difference in transactions costs between the different policies can be ignored. Finally, a word is about costs and revenue calculations. Costs/revenues are calculated using the unit commitment (generation) model which includes spot market and forwards activity. Thus revenues are accounted for in a delivery month only.




Figure 3: Comparisons between fixed-mix strategies

The experiments reported in this study are based on data obtained from Pinnacle West Capital. In order to maintain confidentiality of the data, we will report performance in terms of fractions, with the best policy assuming the value of 1. Figure 3 is based on outputs that showed that using $f = 0.1$ provided the most profitable fixed-mix strategy. Note that although some other fractions appear to be competitive during certain months, using $f = 0.1$ provides the overall winner among the fixed-mix strategies.

Next we proceed to experiments with the Stochastic Programming approach. These experiments were run with the same data as above, except that the fixed-mix hedging rule was replaced by decisions from the stochastic pro-

gramming model. During each month (January 2001 through May 2001), we run the stochastic programming model once. As before, decisions are made before observing market prices at Palo Verde, AZ and Henry Hub, LA. The planning period used within the decision model was five months long (i.e. $T = 5$). Hence as in the previous experiments, delivery dates of six months in the future were permitted in the model. Thus, the experimental setup, and data are exactly the same as in the previous study, and this permits comparisons between hedging decisions from stochastic programming and those from the fixed-mix rule.




Figure 4: Comparing DASH with Fixed-mix Strategy (Backcasting Data)

In Figure 4 we use the best fixed-mix strategy ($f = 0.1$) as the basis for our comparisons. We made two series of runs with the DASH decisions: one using $\alpha = 0.3$ (i.e. 30% change allowed in the portfolio from one month to the next), and another series of runs using $\alpha = 0.5$ (i.e. 50% change allowed in the portfolio from month to month). The motivation for such controls was discussed earlier in the paper (see section 4.2). In any event, both series of DASH runs perform significantly better than the best fixed-mix strategy. It turns out that the series of revenues for $\alpha = 0.5$ exceeds that for the best fixed-mix strategy by approximately 7% per month, on average. This is a significant advantage in favor of the DASH model.

Before closing this subsection, we should comment on a certain initialization bias that results from restrictions imposed by the initial portfolio. Recall that when we allow a 50% change allowed in the portfolio from month to month ($\alpha = 0.5$), it takes about 2 months for the effect of the initial portfolio to wear off. It is therefore appropriate to focus our attention on the performance of DASH (with $\alpha = 0.5$) for months 3, 4 and 5. Similarly, when $\alpha = 0.3$, the output for months 4 and 5 are critical. Thus if we set aside the initialization bias, the performance of the DASH model for the period March – May 2001 is clearly superior to all tested fixed-mix strategies.

5.2 Experiments with Synthetic Scenarios

In order to test the robustness of the decisions provided by the DASH model, we created synthetic scenarios and

tested the decisions provided by the model against these scenarios. In conducting this phase of our experiments, we did not re-optimize to allow DASH to adapt to the observed (synthetic) scenario; instead, we used the decisions obtained from the backcasting experiment, and applied those to the synthetic scenarios. Hence the gains reported here are lower bounds on potential improvements.

The synthetic scenarios were created in two steps. First, we create a series of forward prices from a discrete-time stochastic process with each time step reflecting the passage of a month. During each month, we draw a random number representing a particular outcome of forward prices. We allow four such outcomes in any month: {Max, High-Median, Low-Median, Min}. The values for these quantities are obtained from historical data as described in section 3.2, and the probability of these outcomes is assumed to be $\{1/8, 3/8, 3/8, 1/8\}$. Note that over a five month period, we can create a total of 1024 scenarios. For the purposes of our tests, we generate 30 scenarios, against which the model is tested. For each of these scenarios, we also generate spot market prices, and loads. The latter are created in the same manner as described in section 3.

Due to the initialization bias in the first two months (see the last paragraph of section 5.1), the comparison we report pertains to months 3, 4 and 5. This comparison involves the DASH model ($\alpha = 0.5$) and the fixed-mix strategy using $f = 0.1$. Figure 5 depicts the fraction of differences (i.e. $(\text{DASH} - \text{Fixed-Mix})/\text{Fixed-Mix}$) over all 30 scenarios, for months 3, 4 and 5. Upon examining this figure, it is clear that DASH is the winner over most scenarios, with the magnitude of wins being significantly higher than the magnitude of losses. A summary of Figure 5 in terms of win-loss statistics is provided in Table 1.




Figure 5: Comparing DASH Runs with Fixed-mix (Synthetic Scenarios)

Table 1: Win-Loss Statistics

Month\Statistics	Wins-Losses for DASH	Average Size of Wins	Average Size of Losses
3	22-8	3.2%	4.7%
4	19-11	11.03%	6.7%
5	19-11	21.81%	12.8%

The win-loss advantages in favor of DASH are unmistakable. Moreover, these results may underestimate the gains because the DASH model was not re-optimized based on observations of the evolving (synthetic) scenario.

6 CONCLUSIONS

In this paper we have provided a summary of the DASH model, and reported experimental evidence that the stochastic programming approach provides a powerful tool for scheduling and hedging in deregulated electricity markets. There are several additional features (e.g. options, swaps etc.) that are being incorporated into the DASH model, and future papers will report on these extensions. The integration of market and production data with statistical models, optimization models, and simulation within one software framework requires fairly heavy investments in modeling and simulation technology. However, as demonstrated by our experiments, such an investment is very likely to bear fruit.

ACKNOWLEDGMENTS

This research has been partially supported by a grant from the National Science Foundation. We thank several of our colleagues at Pinnacle West Capital, the holding company for APS. In particular, we thank Ned Lesnick (formerly at Pinnacle West Capital), Wen Lin, Bill Maese, Steve Murphy, and Gary Zhu for their insights and support.

REFERENCES

- Bertsekas, D.P., G.S. Lauer, N.R. Sandell Jr., and T.A. Posbergh. 1983. "Optimal short-term scheduling of power systems," *IEEE Transactions on Automatic Control*, AC-28, pp. 1-11.
- Birge, J.R. and F. Louveaux. 1997. *Introduction to Stochastic Programming*, Springer
- Dupacova, J., N. Grawe-Kuska and W. Romisch. 2000. "Scenario reduction in stochastic programming: an approach using probability metrics," Working paper, Institute for Mathematics, Humboldt University, Berlin, Germany.
- Feinberg, E. 2002. "Electricity Power Load Pocket Modeling and Forecasting," University of Florida Seminar, March.
- Georgantzas, N. C. and W. Acar. 1995. *Scenario-Driven Planning*, Quorum Books, Westport, CT.
- Hamilton, J. D. 1994. *Time Series Analysis*, Princeton University Press.
- Hobbs, B.F, M.H. Rothkopf, R.P. O'Neill, and H-P. Chao (eds.). 2001. *The Next Generation of Electric Power Unit Commitment Models*, Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Hull, J.C. 1997. *Options, Futures, and Other Derivatives*, Prentice-Hall, NJ.
- Nowak, M.P. and W. Romisch. 2000. "Lagrangian relaxation applied to power scheduling in a hydro-thermal system under uncertainty," *Annals of Operations Research*, 100, pp. 251-272.
- Takriti, S., J.R. Birge and E. Long. 1996. "A stochastic model for the unit commitment problem," *IEEE Trans. of Power Systems*, 11, pp. 1497-1508.
- Takriti, S., B. Krasenbrink and L. S-Y. Wu. 2000. "Incorporating fuel constraints and electricity spot prices into the stochastic unit commitment problem," *Operations Research*, 48, pp. 268-280.
- Taylor, S. 1986. *Modeling Financial Time Series*, John Wiley & Sons.

AUTHOR BIOGRAPHIES

SUVRAJEET SEN is Professor of Systems and Industrial Engineering at the University of Arizona. His interests are in optimization algorithms, especially those arising in stochastic optimization problems. He has published over 60 papers, many of which deal applications arising in electric power systems, communications networks and transportation. He has been active in INFORMS for over 20 years, and led the formation of the very successful Optimization Section. He is also a past Optimization (Area) Editor for Operations Research, and a past President of the INFORMS Telecommunications Section. He can be reached at <sen@sie.arizona.edu>.

LIHUA YU is a PH.D. student of Systems and Industrial Engineering at University of Arizona. He received his Master's degree from University of Arizona in May, 2002. His research interests are mainly in stochastic programming and its application to power portfolio optimization. His email address is <lihua@email.arizona.edu>.

TALAT GENÇ is a Ph.D. student of Economics and M.S. student of Industrial Engineering at the University of Arizona. His research interests are industrial organization and applied econometrics. He is working on stochastic general equilibrium and power portfolio optimization. He can be reached at <tgenc@u.arizona.edu>.