Managing Market Risk in Energy

Michael Denton, Adrian Palmer, Ralph Masiello, *Fellow, IEEE*, and Petter Skantze

Invited Paper

***Abstract*—The market risks encountered by energy asset operators can be categorized as short term/operational, intermediate term/trading, and long term/valuation in nature. This paper describes how the market risks in operations can be measured and managed using real option models and stochastic optimization techniques. It then links these results to intermediate term value at risk and related risk metrics such as cash flow, earnings, and credit risk which can be used to measure trading risks over weeks to months; and how to optimize these portfolios for risk-return relationships. Finally, it then explores the risks in longer term energy portfolio management and how these can be simulated, measured, and optimized.**

***Index Terms*—Earnings at risk, energy risk management, portfolio optimization, potential credit exposure, real options, value at risk.**

I. Introduction

DEREGULATED energy markets and the emergence of centralized physical markets in electric power run by the ISO/RTO organizations lead to complexities in managing market risks in both operations and financial ways. This paper will cover the kinds of risks faced over different time frames by "asset operators" or generator/producers. These operators have to deal with *operational/earnings* risks over the short term (less than 1 month), *trading/financial* as well as operational risks over the intermediate term (1 month–1 year), and *asset valuation/equity* risks over long (>1 yr) time frames. The kinds of risks that fall into these time frames have strong analogies to classical power systems operations, planning, and economics but additionally, of course, involve market price and financial risks that require new techniques and mathematics to deal with them. Some of the methodologies have similarities or origins in the electric power world; some in the financial world, and some are truly new incorporating syntheses of both. We develop the "real option" model of a generator in detail as an exposition of financial modeling applied to physical asset operations and then discuss its application to longer term trading and valuation problems.

II. Operational Risk Assessment

In the regulated world the owner of a portfolio or fleet of generation plants had to solve the economic dispatch (minute by minute) and unit commitment (hourly) scheduling problem—how to most economically schedule the generating units considering the unit economics, physical constraints, and incremental transmission losses such that the operator's total commitment to deliver power was met [1]. A load forecast was assumed as given, and the revenues from the load were precisely known as rates were fixed. Incremental purchases and sales of generation from neighboring utilities were analyzed with the same tools.

The market price based unit commitment problem is central to risk assessment and management for the fleet owner/operator across all time horizons, and is our focus.

The "spark spread" model has generally been used to value a generation asset or a "trade" of power from that asset over a period of hours to years. The spark spread models the conversion of fuel (gas, oil, coal) into electricity at a specified conversion factor and incorporates the uncertainty (volatility) in the prices of the fuel and the output electricity. Classical financial modeling of a "spread option" can be used to "value" the option, and thus, price the unit's output—to make a trade or to decide whether or not to schedule production at a given price [9]. Spark spread models are used against long term forecasts of fuel and electricity prices and long term curves of daily and seasonal price variations to assist in making investment decisions around plants in different geographies and time frames.

Spark spread models are increasingly recognized as deficient, however, for two critical reasons: they fail to represent the nonlinear heat rate characteristics of different generating equipment and require an average heat rate representation; and they fail to consider the physical constraints on unit operations as well as associated costs—rate limits, cycle times, start up costs, and limits on total unit cycles in a maintenance period [5], [7]. As a result, spark spread models usually "overvalue" a plant by 10–30%. When used to value an asset for long term financing, they can produce over optimistic valuations—a fact which rating and financial institutions have recently realized.

A. Generation Model—Costs and Constraints

The problem is to maximize the net profit from operating an asset over a specified operating period spanned by the start and end dates, given a set of prices for fuel and electricity, the volatilities in those prices in a geometric Brownian motion price process, and the physical constraints on the operation of the asset.

The constraints include

- hourly minimum and maximum operating ranges;
- maximum hour to hour changes in output, or ramp rates;
- cycle time constraints, or minimum hours "off" and "on;"
- maximum number of cycles in a period.

Manuscript received November 18, 2002.
R. Masiello (e-mail: rmasiello@kemaconsulting.com).
Digital Object Identifier 10.1109/TPWRS.2003.810681

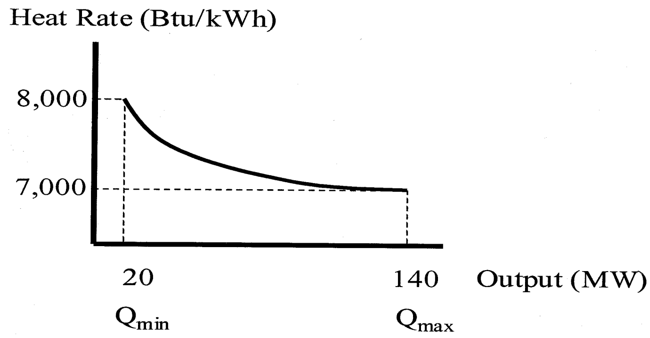

Fig. 1. Heat rate function describing generator efficiency.

The most critical cost component is the fuel usage of the plant, dictated by its efficiency as generally expressed in terms of the heat rate of the plant (Fig. 1).

The total fuel cost (TFC) in any given hour is a function of the heat rate (HR) at the given output level (Q), and the spot price of the fuel (S^f)

$$TFC = HR(Q) * S^f * Q.$$

In addition to the total cost of fuel, there is a component of the variable cost which is independent of the primary fuel type, used to model maintenance costs, taxes, labor costs, and so on. Start-up costs include the cost of heating the unit to operating temperatures. There may also be fixed run charges accrued whenever the unit is on.

B. State Space Model for Generation Asset

In classical unit commitment formulations, the generator can have a number of "on" or "off" state and for all combinations of portfolio generator states in each hour, an economic dispatch algorithm determines the best operating points and thus the "cost" of that combination. In order to build a real option model of a unit that can be integrated with a stochastic price process we define the states as being the output level Q and the number of hours the unit has been on or off, called RT. The state pair (Q_h, RT_h), fully characterizes the operation of the generator in any hour. In addition it contains all the information needed to determine which operating states are feasible for the asset in the next operating hour. The current output combined with the maximum ramp up and ramp down rates limit the feasible output levels in the next hour.

C. Discrete State Space Formulation

In order to be able optimize over the possible decisions in each hour, we first to create a limited number of possible decisions by discretizing the space of possible output levels of the generator.

This approach assumes that the operator of the asset will always take full advantage of the ramp capability of the plant. The concept of ruthless exercise is borrowed from swing option modeling in finance, where it can often be shown that a ruthless strategy is always optimal. It is possible to discretize the unit to accommodate this ruthless ramping philosophy or to specify discretization levels to capture desirable regimes of plant operation for nonconvex heat rates (Fig. 3). Fig. 2 is a decision tree for unit operations arising from a ruthless ramping model.

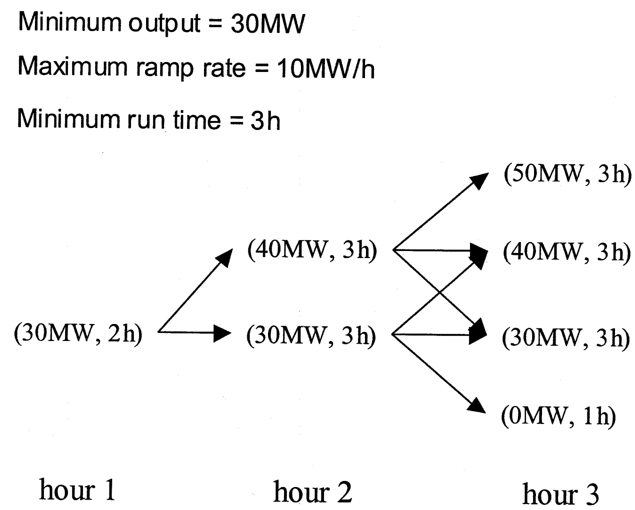

Fig. 2. Generator decision tree based on dynamic operating constraints.

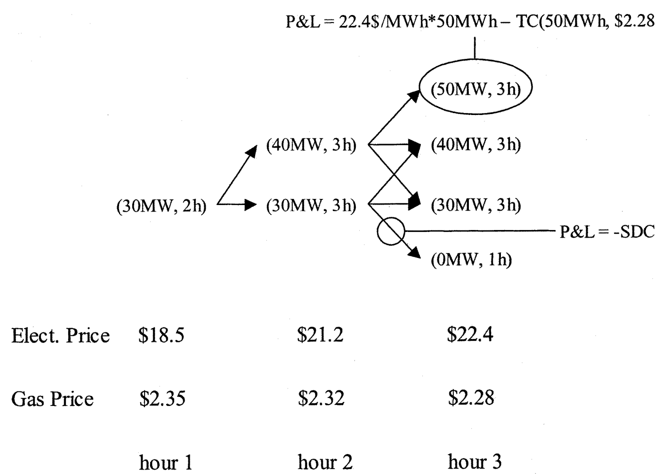

Fig. 3. Mapping operating costs and revenues to the decision tree.

It is worth noting that, anecdotally, plant operators in California in 1998 and 1999 tended to operate the units in exactly this fashion, ramping them at maximum rates based on hourly market awards and whether a plant was "in the money" or not. This result is not dissimilar to one that could be predicted from optimal control theory—a kind of "bang bang" control. A corollary to the bang-bang conclusion is that the number of discretization levels required to model the unit should be the number of time steps required for the unit to ramp to full load.

D. Optimization for the Deterministic Case

The first step in optimizing the dispatch schedule of the asset is to superimpose the known future fuel and power prices on the decision tree. Given these prices it is possible to map the applicable cost and revenue to each node in the decision tree.

To each node, we assign the cost equal to

$$TFC_h + (NFC_h * Q_h) + RC_h + FC_h.$$

And a revenue equal to

$$S_h^e * Q_h$$

where S_h^e is the spot price of power, TFC the total fuel cost from the heat rate, NFC the variable nonfuel cost, RC a running cost, and FC the fixed cost.

Finally, there are costs associated with the branches between the nodes. This is where startup and shutdown costs are captured. The revenue and the total running and transition costs combine to give the profit (or loss) at each node.

Having assigned a P&L value to each node, one can proceed to search for a profit maximizing path through the decision tree with a backward iterating dynamic programming algorithm. operation period. This algorithm will produce the optimal dispatch path and the optimal valuation (P&L) of the asset at every node along the path.

E. Optimization Under Uncertain Market Conditions

The scheduling optimization algorithm assumes that the operator of the asset has an exact knowledge of what the spot price of the fuel and power will be in each hour of the operating period. In reality, however, there is significant uncertainty associated with future price levels. Under these conditions, the operators objective is to maximize the expected profit from scheduling the asset.

As one approaches a given hour of operation, more information becomes available to help forecast the energy prices in this period (updated weather forecasts for instance). As a result the operator should be constantly updating his dispatching and spot trading decisions to reflect the best current knowledge of the future.

We assume that the spot price $S(t)$ of the energy commodity follows a risk-neutral stochastic process defined in terms of the logarithm of the price by

$$d\ln S(t) = [a(t) - k(t)\ln S(t)]\,dt + \sigma(t)\,dW(t)$$

with $a(t)$, $k(t)$ and $\sigma(t)$ the time dependent drift, mean-reversion rate, and volatility, respectively [4].

Therefore, the following types of risk-neutral processes for the spot price are considered:

i) geometric Brownian motion process with constant volatility [$k(t) = 0$ and $\sigma(t) = \sigma =$ constant];

ii) mean-reverting process with constant volatility and constant mean-reversion rate [$k(t) = k =$ constant and $\sigma(t) = \sigma$ = constant];

iii) mean-reverting process with time varying volatility $\sigma(t)$ and time varying mean-reversion rate $k(t)$.

This model captures two crucial properties of electricity prices

- seasonality—average power prices vary predictably with respect to time of day, day of week, and month of year. This is captured by the time varying drift parameter $a(t)$;
- mean reversion—prices may suddenly spike, but tend to return to normal levels within a few hours or days. This property is captured by the mean reversion rate $k(t)$.

The parameters in this model are typically identified by a maximum likelihood estimation process from historical data. Unlike the classic closed formulations, this approach only requires a risk neutral assumption, not a no-arbitrage assumption.

The spot price model corresponds to a one-factor futures price model with a futures price stochastic process under the risk-neutral measure

$$dF(t, T) = F(t, T)\sigma(t, T)\,dW(t)$$

and with the futures price instantaneous volatility $\sigma(t, T)$ being a deterministic function of time t and maturity T of the form

$$\sigma(t, T) = \sigma(t)e^{-\int_t^T k(s)\,ds}.$$

This spot price model is used in many financial engineering models of commodities futures, derivatives, and, in particular, energy contracts.

As with the decision states, we need to discretize the space of possible price levels in order to perform the optimization. This is done utilizing a trinomial-tree based methodology. We approximate the stochastic process for $\ln S(t)$ over a time interval $[T_0, T]$, where $T_0 = t$ is the valuation date and T is the expiry (end of operation) date, with a trinomial tree process with N time steps. We assume that the initial forward price curve, the values of the mean-reversion rate function $k(t)$, and of the spot price volatility function $\sigma(t)$ are given.

The trinomial tree approximation of the price process is built in two stages such that the tree gives an exact match of the initial forward price term structure which is assumed to be known. More precisely, the final tree for the price $S(t)$ will be such that for each time slice T_i, the expectation of $S(T_i) = e^{X(Ti)}$ computed with the risk-neutral tree probabilities equals the forward price $F(T_0, T_i)$. This kind of a flexible trinomial tree is widely used in option valuation [3], [6], [8]. It produces a spot price distribution as in Fig. 4.

We will refer to a price state in the tree as a node and label the nodes by (i, j), where i refers to the time slice T_i (and equals the number of times steps from T_0) and j indicates the price level. Thus, $S_{i,j}$ will denote the price at node (i, j).

The generation option model is valued using a *multi-level trinomial tree—trinomial forest*—methodology: trinomial tree approximation of the stochastic evolution of the electricity spot price in a one-factor model as outlined above, followed by option valuation by "backward induction" through the trinomial forest (Fig. 5).

A node in the forest is characterized by time, price, and a pair (output level, run-time) characterizing the state of the generator, and thus, labeled by $(i, j, (Q, RT))$. One can regard the forest as containing a tree isomorphic to the spot price tree for each possible generator state (Q, RT).

Working backward from the final date in the exercise period to the valuation date through the trinomial forest, one computes the option value at each node in the forest, as the maximum of the following possibilities—(if these are permitted from the current: no change; ramp up, ramp down, turn off, turn on. In each case, the decision value is the discounted expectation of the option value at the next time step on the node to which the decision would branch, plus any transition costs of that decision).

For time steps i corresponding to dates T_i prior to the starting date T^S, one needs only to roll backward the option values from the nodes at time step i from the tree corresponding to the specified starting state of the generator in order to compute the final option value at the valuation date $T_0 = t$.

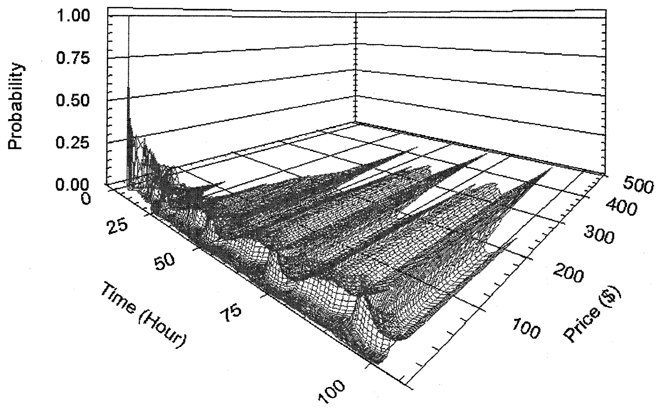

Fig. 4. Spot price probability distribution from tree-based model.

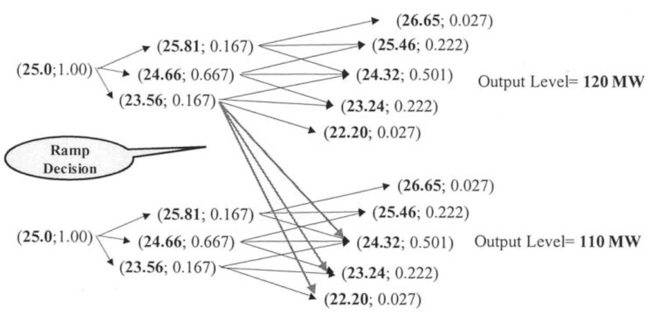

Fig. 5. Trinomial forest.

The stochastic optimization algorithm returns two sets of results

- value of asset over the operation period. This is equal to the expected P&L earned by optimally dispatching the unit over this period;
- decision rules for the asset at each time period and operating state and price. The algorithm stores the optimal dispatch rule at each node in the tree. This includes the optimal price issued at the current node in the tree.

The decision rule, in effect, enables the scheduler to know the expected profitability of turning a plant on, ramping, etc. Since the model is stochastic and incorporates the effect of spot price volatility, it will produce the correct expectation. A deterministic solution taking the expected prices as given will not in general produce the correct expectation as it ignores volatility. One benefit of the model is to show the effect of price volatility on plant valuation; it also will produce different commitment decisions which take advantage or avoid the consequences of possible price departures from the expected path.

The tree can also be used to produce a set of decision rules—what to do at a given node given the state of the unit and the price, typically generated for each month. A Monte Carlo simulation of prices can be run through the decision rules to produce high-speed simulations of optimal unit operations for a variety of purposes. Using this technique, it is also possible to simulate random outages and deratings which will be important in portfolio operations.

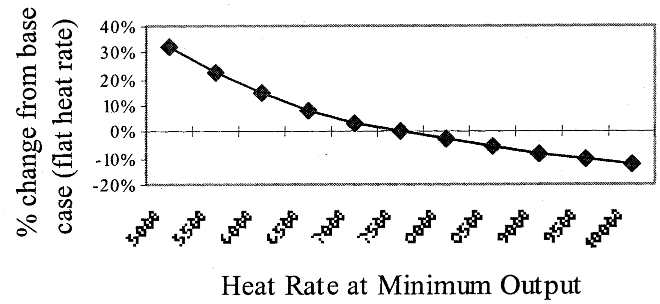

Fig. 6. Heat rate impact on spark spreads.

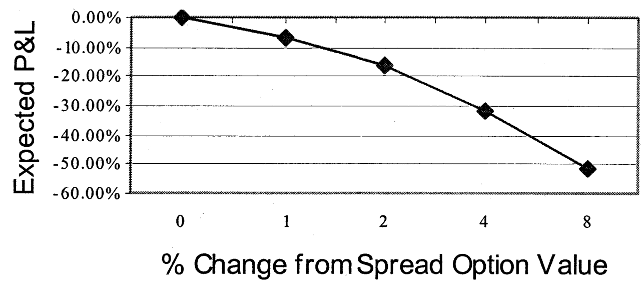

Fig. 7. Ramp rate impact on spark spreads.

F. Comparison With Spark Spread Model

Fig. 6 shows a comparison of the spark spread model and the real option model of a generator across a range of heat rate variability—as expected, as the heat rate becomes less constant the two valuations differ. Fig. 7 shows the effect of ramp rates on unit valuations.

The effect of ramp rates on valuations is always unfavorable; a unit with a 4-h ramp time from "idle" to full load is worth 30% less than the spark spread model would indicate.

Similar variations occur when start-up costs and minimum run times are considered, and additionally when these are coupled with ramp rates and nonlinear heat rates, the effects in combination produce larger errors. For instance, one of the adverse effects of a slow ramp rate is to force the unit to spend more time at less efficient heat rate points.

G. Generation Portfolio Optimization

In a perfectly liquid environment, where all the output of the plants can be sold without limit, and the total portfolio output has no effect on the market price, each of the units is an independent price taker and a fleet of units can be optimized by individually optimizing models of each unit as described before.

However, generation fleet operators rarely operate in such a simple environment. In general, they operate in a complex world with some of the flavor of the regulated load fulfillment/cost minimization world and with some of the flavor of a trading and option selling/exercising world. The typical fleet operator has an existing portfolio of committed load to serve, regulated and competitive, plus a portfolio of traded sales/purchases to fulfill. Typically, the portfolio operator exists in a semi-liquid market environment where spot purchase/sales can be made to improve the profitability, indeed to meet load and traded requirements above capacity, but also where the total spot purchases may and probably will impact the market price.

A conceptually simple algorithm for handling the portfolio problem in an illiquid market with a portfolio of supply commitments is to first model the units as price takers in a liquid market using real option models of each unit. By running a price forecast for the scheduling period through all the models, a profit optimizing production schedule for each unit will be obtained. However, the production schedules of the units, when added up, will sum to a long or short (excess or deficient) position vs. the total committed supply.

In a manner analogous to Lagrangian relaxation solutions, the hourly price forecast can be adjusted iteratively until the total production output matches the supply commitment. In practice, however, physical long or short positions are balanced in the spot market. Bid ask data can be used to maintain a spot market demand/supply curve which becomes a market or slack unit in the portfolio. Now, as the market price is adjusted in each hour to eliminate portfolio imbalance, the "market" unit will contribute more or less as well until the imbalance is zeroed out. At this solution point, the solved for market price reflects the market reaction to additional purchases/sales in the spot market, and the units are profit optimized against the resulting price curve and balanced against supply commitments at that price in the spot market.

Physical options in the trading portfolio can similarly be modeled and this methodology will indicate when they should be exercised to improve expected P&L, as well as to anticipate when sold options will likely be called.

The trinomial forest grows linearly with the number of time steps in the scheduling horizon as the price process saturates due to mean reversion. Even so, the models are computationally intensive for periods of more than a week or two.

The real option model can be used to produce a value surface which provides valuations and greeks (derivatives of value wrt to prices) as a function of fuel and electricity prices. This value surface will be representative of the unit value with sufficient accuracy for a longer term solution—it captures the intra-day and intra-week optionality value of the unit considering all the physical characteristics. The value surface can then be used directly in the Monte Carlo VAR process very efficiently.

For longer term hour by hour simulations where hourly behavior is modeled the model decision rule which looks at fuel, electricity prices, and the "state" of the unit can be used very efficiently in random price path simulations for long term valuations.

The performance of the model and of the stochastic portfolio optimization is competitive with alternatives—both spark spread models and deterministic unit commitment. Because the models in a portfolio are valued and simulated independently of each other, and because the dispersion of nodes in the forest is bounded by mean reversion, the run times increase linearly with both the time span of the problem and the number of units in the portfolio.

III. Intermediate Term Risk Management

In the period ranging from weeks to a year or more, the market risk is dominated by fluctuations in the forward price curves and their correlation. Operational issues are less important than trading and hedging decision making. It remains important to represent the hourly granularity of price, demand, and best practice operational behavior in the intermediate term analyzes and simulations, but the decision support questions are around monthly forward contracts. Hourly forward prices are generally not available for the out months and have to be simulated and related to the monthly prices that are available.

A. Value at Risk

The primary risk assessment metric in this time frame is value at risk, VAR, and related applications of the same methodology used to assess cash flow at risk, earnings at risk, and credit exposure. VAR is the classic risk management tool widely used by financial institutions and corporate treasury functions in many industries. It measures the minimum occasional loss expected in a given portfolio within a stated time period. Thus a one-day 95% VAR of U.S.$ 2 000 000 means that the single day loss will be less than U.S.$ 2 000 000 95% of the time and 5% of the time the loss should be greater. There are many excellent references on VAR in the literature, and a good number deal with VAR applied to energy trading.

VAR measures the change in the portfolio value through the end of the specified time horizon. Changes in portfolio value are marked to market for all forward periods, so VAR measures risk in the portfolio value for the total period.

For an operational business such as power generation, VAR may not be the single most useful metric available. For such a business, the bulk of trades are physical trades for delivery or options on physical energy sold as hedges on unsold generation or peak load commitments. Per FAS 133, hedges on the operational business which are deemed effective, that is, sufficiently correlated with the profitability of specific physical trades/schedules, need not be marked to market for periodic P&L reporting and instead the cost of the hedge (the premium to buy the option) is expensed as an operating cost. 90% of the trades in large power producers' books may be physical trades or hedges on them; with only 10% "derivative" trades that are deemed speculative and have to be marked to market.

B. Earnings and Cash Flow at Risk

In these cases alternative risk metrics such as earnings at risk (EAR) and cash flow at risk (CFAR) are more important. EAR measures the variability in accrued earning from physical deliveries made and financial positions that settle during the period. It does not include any change in the ongoing portfolio value after the risk time horizon as VAR does. CFAR differs from EAR only in that payment dates on settled trades/deliveries have to be considered for cash flow purposes; and these can vary anywhere from end of day in some pools to 30 days after the end of the month for some over the counter bilateral contracts. EAR and CFAR require hourly representations of price processes and production, and detailed analyzes of each contract in order to be accurate. Incorporation of decision rules and value surfaces from real option models is essential to EAR and CFAR validity.

C. VAR Estimation

VAR, EAR, and CFAR are all estimated using the same basic methodology. Given the volatilities and the correlations of the prices of the commodities in the portfolio, the distribution of portfolio returns is analyzed to identify a confidence level for the left tail. If an instrument's returns are normally distributed

$$\text{VAR} = \alpha\sigma V_o\sqrt{\Delta t}$$

where α is the confidence level (1.65 for a 95% level), σ is the annualized standard deviation of returns, V_o is the initial mark to market value of the instrument, and Δt is the time horizon in years. This approach is complicated a bit for total portfolio returns by the covariance of the instruments. Typically, a variance-covariance approach is used in which

$$\text{VAR} = \alpha\sqrt{x'\sum x}$$

and x is diagonal matrix of instrument returns and Σ is their within-period covariance matrix [10]. This is sometimes called the "delta-normal" approach, since the x matrices may contain the delta (first order) equivalents of positions that are less liquid. Obviously, the accuracy of this closed form calculation is limited by the normality of the expected returns and the extent to which the positions have well behaved, linear returns. This is rarely the case for energy contracts due to the decidedly nonlog normal behavior of energy prices and the nonlinear behavior of many energy instruments such as swing contracts, interruptible supply contracts, and more complex derivatives.

Higher order approximations are sometime used, but simulation approaches dominate in the tools used for larger and more complex portfolios. Here, correlated prices at multiple locations and times are simulated in a Monte Carlo process where a variety of techniques can be used to model nonlog energy price behavior. Valuing each position for each price simulation is straightforward, even for path dependent instruments.

Detailed MC VAR calculations can be cumbersome to run. For a portfolio with 100 000 to 500 000 open positions in trades and contracts that involve 500–1000 different market prices it may require several hours to perform a detailed analysis. Energy firms typically have to consider so many market prices due to the locational variability of energy prices. Normally locational prices are tightly correlated to major hub prices, but in times of high demand and transportation bottlenecks the locational or basis prices can deviate significantly. In order to reduce the number of prices modeled, principle component analysis is often used to reduce the dimensionality of the problem; additionally quasi-Monte Carlo techniques can be used to achieve a linear rather than square root growth in the number of simulations required to increase the precision of the results.

An important variation of VAR is the Conditional VAR or CVAR. Where VAR establishes the loss at the 95% level, it does not measure the expected loss. CVAR is the integral of the loss distribution and measures, instead, the expected amount of loss given that it is greater than the 95% level. As such it is a more meaningful metric. Also, because it is an integral it is inherently better behaved mathematically and more tractable in, for instance, portfolio optimization.

Such calculations also must consider the volumetric risk in full requirements contracts—peak load demand varies with weather, and peak demand correlates with price spikes. So these contracts require not only price simulation but correlated demand simulation for valuation. It is not sufficient to only vary a fixed daily load shape of hourly values, since prices are sensitive to regional loads in a highly nonlinear fashion. Hourly modeling is typically used for small portfolios, and nonlinear correlation methods for larger portfolios.

Early attempts to solve the EAR and related problems would typically subject a single, deterministically obtained, portfolio schedule to a Monte Carlo analysis of different prices after the schedule determination. This methodology ignores the operational flexibility and ability to react to new operation. Another approach was to schedule against each simulated price path; this builds in an assumption of perfect foreknowledge in operations. Both of these approaches are invalid and unsatisfactory; thus the need for the stochastic dynamic programming real option approach.

By running Monte Carlo simulations through the tree-derived decision rules described in Section II-E it is possible develop tractable value surfaces which can be used to quickly value a generation plant or other real option at each step in the Monte Carlo price path. This enables the effective incorporation of all the physical constraints and nonlinearities of the real asset into the VAR and other metrics. The ultimate performance constraint on VAR calculations is that they must be produced at least daily; the use of this technique makes detailed modeling possible in daily VAR runs.

D. Credit Exposure

In the financial world, credit exposure has two components—a change in the credit rating of a corporation affects the value of bonds that it has issued, and thus the value of a portfolio containing those bonds. Also, depending upon the credit rating of the issuer, there are varying probabilities of default that can be used to assess the risk of the bond's value disappearing in a default/bankruptcy scenario. This is traditional credit risk.

In the energy industry, the primary credit risk has been the default risk, either of supply of contracted deliveries or of making payment on settled trades. Added to this is the replacement risk of, for instance, acquiring power on the spot market to replace the defaulted delivery. Given that a delivery default may correlate strongly with price peaks, the replacement power can be at a real premium over the original contract. This potential credit exposure has not been routinely measured by energy firms but is being focused on in today's environment as equal in priority to VAR and EAR. Ratings become a major concern in energy credit risk assessment due to the linkage of ratings declines to debt triggers and consequent liquidity crises. Potential credit exposure is analyzed with Monte Carlo techniques similar to VAR, although default probabilities are difficult to establish.

In June 2002 Standard & Poor identified liquidity risk due to margin calls as a major factor in the creditworthiness of energy trading firms. S&P announced that henceforth it would require firms to report the cash exposure they had to margin calls under stress tested conditions; in other words, to add to the calculations of VAR, EAR, etc., a calculation of cash position risk due to margin calls. Computationally simple, this calculation is programmatically tedious since each contract may have different applicable margin requirements.

E. Transmission Congestion

Transmission congestion can be hedged by owning transmission congestion contracts (TCCs) or firm transmission rights (FTRs) in transmission markets where they are available. TCCs pay off based on the price difference between their "from and to" nodal locations in an locational-based marginal pricing (LBMP) environment and can be represented as a swap—trading one commodity for another, and settled financially. Thus, they perfectly hedge the market exposure to LBMP. A producer who uses TCCs to hedge a physical delivery position will not only hedge against LBMP incremental costs, he will also concede any profit opportunities from reverse congestion. FTRs only pay off in one direction; therefore, they are represented as swaptions or options on a swap. Modeling and valuing transmission contracts as swaps assumes that they are perfectly correlated with underlying locational prices; an area of current interest is how to incorporate factors such as transmission and unit outages into simulations used for valuation purposes.

F. Portfolio Optimization

In classic portfolio theory, optimizing the expected return for a specified level of risk is a well-known problem [11]. There are three dimensions to the problem—the expected return (or P&L) on each instrument in the portfolio; the risk in that return (as measured by standard deviation in the expected return, or alternatively by the VAR); and the quantity of each instrument held. The latter of these is the choice variable. Simplistically, portfolio optimization is the search for a vector of quantities that satisfies a number of constraints and provides minimum total variance with maximum return. There are typically many such vectors, and their risk/return metrics comprise the "efficient frontier" of this problem space as shown in Fig. 8. Any imperfect correlation, positive or negative, between two different assets or positions allows the creation of optimal portfolios. As prices, volatilities and correlations of the component positions change, the set of optimal quantity vectors also moves. At this point, the transactions costs of maintaining a desirable portfolio become a complicating factor in dynamically hedging (optimizing) a portfolio over time.

This problem is of great interest in the energy world to unregulated market participants, and is becoming more interesting to many regulated local distribution companies (LDC) as well, and, potentially, to ISOs that act as single buyers. In many states the cost of hedging volumetric risk (load peak) has been recognized as a legitimate cost of service for LDCs. In the past, this was handled physically, by maintaining reserve margins. Today it is often handled by purchasing capacity options, swing contracts, weather derivatives, and other physical and financial instruments that provide a hedge to the LDC against high loads. In order to recover all the costs of that hedging, the LDC needs to show that it is following prudent practice—both in terms of its assessment of relevant risks and its expenditures to hedging against them. Portfolio optimization provides the vehicle to do this. The generation operator must consider how to price physical options, how many to hold to hedge against outage and volumetric risk, and how many transmission hedges to hold against specific locations and times.

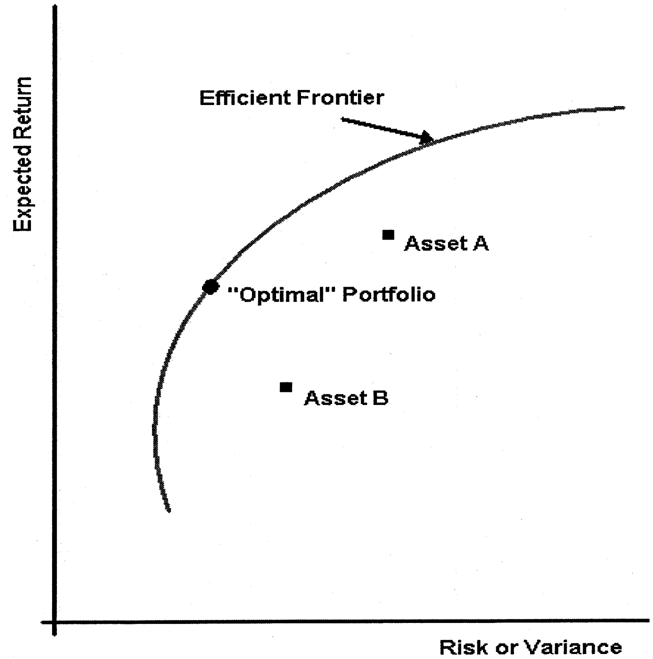

Fig. 8. Portfolio optimization space.

In the intermediate term, the portfolio optimization problem is largely a trading and hedging problem and the major challenges in implementation are focused on simulations that effectively span the complete space of instrument returns—and quickly enough to be useful—and on robust optimization or searching that can solve the problem in usable time frames. Algorithms in use for the energy portfolio optimization problem include iterative quadratic programming techniques and TABU search engines.

Problems in portfolio optimization include how to realistically group many detailed physical forward positions into synthetic aggregate positions for use in optimization; and how to incorporate constraints on other risk metrics such as EAR and CFAR in addition to VAR.

IV. Long-Term Asset Valuation

Beyond the time frame of the markets for liquid traded energy instruments, the key focus is upon understanding the market risks that affect the value of energy assets whose operational lifetimes typically extend 20 to 30 years or more. Forward price curves going out several years are readily available in the oil and natural gas markets, underpinned by the key futures contracts for these commodities listed on the NYMEX exchange. In electricity, however, price discovery has yet to extend much beyond a time frame of 18 months to two years.

Energy deregulation has been accompanied by significant investment and divestiture activity, involving both the transfer of existing assets and the construction of new facilities. Motivations include regulatory edict, business refocusing, and balance sheet improvement efforts. In the U.S., the level of generation capacity transferred from utility to nonutility ownership has ranged from 23 to 51 GW annually over the past four years. In terms of greenfield projects, it has been reported that more power plants will have been commissioned in the U.S. during 2001 and 2002 than were added in all of the 1990s. All this investment activity places a considerable emphasis upon asset valuation.

The fundamental challenges of asset valuation—modeling the evolution of forward price curves and the operational P&L of the assets against them—is essentially the same as that involved in intermediate term risk management. However, additional risk drivers need to be considered, drivers that can usually be safely ignored over shorter time scales. These include such factors as technology risk and regulatory risk.

Natural gas-fired combined cycle gas turbines (CCGTs) have long been the new generation technology of choice in many regions. Enhancements in turbine technology have produced steady improvements in plant heat rates over the past decade, undermining the market competitiveness of early CCGT stations. As a result, long term asset valuation studies often include assumptions on the rate of technological innovation (reductions in heat rate and/or capital costs) for benchmark new entrant generation technologies. These assumptions determine the market prices at which new entrants will break even—the long run marginal cost of electricity generation in classical economics and a key consideration in many market analyzes over this time frame. Harder to model is the possibility of a breakthrough in alternative generation technologies, such as clean coal or fuel cells. From a wider perspective, large-scale generation plants face competitive threats through the development and application of distributed generation and demand-side management technologies.

California's electricity restructuring experience has dramatically highlighted the regulatory risks inherent in the energy sector. Price caps and fundamental market rule changes can significantly influence the economics of operating and investing in deregulated electricity markets. Even relatively successful market designs have witnessed a steady stream of rule changes and enhancements aimed at correcting perceived design flaws and improving operational efficiency. Only a handful of deregulated power markets around the world, such as England, Wales, and Norway, have been operational for more than ten years: none of these markets could yet be characterized as mature. It is reasonable to expect that the market and regulatory frameworks will continue to evolve in the future, although it is far from straightforward to incorporate the possibility of such changes in an asset valuation model. Environmental policy presents another area of considerable regulatory risk for the power sector, with the potential to significantly impact the relative competitiveness of different generation asset types.

A. Scenario Analysis and Production Cost Models

Scenario analysis utilizing production cost models has probably featured in the majority of asset valuation studies over the past decade. This approach typically involves examining the fundamental market drivers and then developing scenario-based price projections using detailed models of the production facilities in the region of interest. Aggregate industry supply and demand curves are built up for the regional marketplace under each set of scenario assumptions. After the characteristics of all existing and planned generation facilities in the region are compiled, plants are stacked up in order of increasing expected production cost. The model then determines the plant output levels and market clearing prices at which projected aggregate demand is met at minimum cost. Key scenario variables might include fuel prices, new plant build costs, and environmental restrictions.

A production cost model can be as simple as a spreadsheet, providing a static snapshot of supply in the region. Most models, however, apply linear or dynamic programming techniques to determine equilibrium prices and output levels, enabling more complex factors such as plant dynamic characteristics, transmission congestion, and emission allowances to be accounted for. Some of these models are essentially stripped-down variants of the unit commitment algorithms referred to in the first part of this paper, suitably modified to run multiyear studies utilizing several characteristic days per year. Plant build and retirement decisions can be incorporated in the model optimization process, given assumptions on the range of available new entrant technologies. In other formulations, these decisions are treated as scenario variables, subject to review and subsequent iteration.

Experience has exposed drawbacks to relying on production cost based models for long term valuation purposes. First, cost-based production models rarely reproduce all the features of observed market prices. Electricity spot prices are generally well behaved at low demand levels and often show a strong correlation to underlying cost drivers at such times, but this relationship tends to break down at higher load levels. Compared to the results of cost-based models, actual market prices typically exhibit much greater volatility and more frequent price spikes. As a result, cost-based models can significantly undervalue generation assets.

The strategic bidding behavior of market participants often contributes to this discrepancy because players exerting market power often elevate prices above cost-based levels, particularly as the margin of supply over demand tightens. By drawing upon game-theory approaches such as Cournot pricing, it is possible to extend the cost-based modeling framework to incorporate strategic bidding considerations. However, even these extended production cost models still tend to be geared toward analyzing equilibrium market conditions. As such, they fail to replicate the dynamic nature of the price setting process in deregulated electricity markets.

This highlights a second significant drawback of traditional production cost models—their inability to adequately address market price uncertainty. Traditional models are generally deterministic in nature, with plant output choices made in full knowledge of the paths of key input variables such as fuel costs and demand. As discussed earlier in this paper, operational decisions in the real world have to be made on the basis of uncertain views of the future. Even if multiple scenarios are run by sampling a range of possible values for key parameters, most models will not overcome their perfect foresight and will therefore tend to produce over-optimistic valuations.

The third disadvantage of production cost modeling relates to changes in the level of information transparency resulting from electricity deregulation. The traditional models tend to require extensive knowledge of generation production costs and transmission system conditions for all facilities in a region, regardless of who owns them. While this information was historically available in the regulated world, it is now increasingly regarded as commercially confidential in the competitive marketplace. Conversely, conventional generation models often fail to capitalize on the rich pricing information now obtainable in spot and forward markets.

These shortcomings have resulted in the production cost-based scenario approach to asset valuation being supplanted or augmented by techniques more adequately suited to supporting trading and investment decisions in the face of highly volatile prices and dynamic, uncertain market conditions. The stochastic real option models described in this paper, used via value surfaces, lend themselves to longer term scenario valuation exercises. With current technology, it is feasible to obtain multiple year simulations in compute times on the order of an hour or less.

B. Real Options and Dynamic Portfolio Optimization

Leading merchant generators have sought to reduce their exposure to specific regional markets by building a broad geographic portfolio of assets. This portfolio approach enables them to mitigate the impact of an adverse change in market prices or regulations in a given region. Furthermore, at any one time, a merchant generating company typically has a range of projects at various stages of operation or development. This provides a valuable portfolio of real options since the merchant generator may choose to accelerate, expand, postpone, scale back or sell a given project. Having advanced a project through the siting and permitting phases, a developer may decide to defer construction until the implied spark spread in the forward markets is sufficiently attractive to begin hedging the plant's expected output. By employing a portfolio risk management approach, merchant generators can ensure that their capital and construction resources are directed to the most attractive development opportunities.

In financial terms, real strategic options such as the ability to expand capacity or mothball a generation plant can be considered as compound options [2]. The fixed costs incurred or avoided represent the option fee on a spark spread option between fuel and power.

The portfolio optimization techniques outlined in the intermediate term section can be extended to help value these and other real options. The intermediate term problem was limited to portfolios of traded energy instruments. Over the long term, assets such as power plants can also be bought and sold. The objective is to dynamically optimize the mix of generation assets in the overall portfolio with respect to the corporate risk profile, taking into account the time lags in the investment process and the discrete nature of generation assets. The results provide insights into the composition of asset portfolios that maximize expected returns over a given time horizon while satisfying key risk metrics such as EAR or CFAR. Portfolio optimization as understood today is a static problem considering a portfolio at one point in time—there is need for theoretical and algorithmic development in dynamic, decision contingent optimization over a time horizon.

V. Conclusions

We have demonstrated a real option model for power generation that can be used in valuing the asset and in determining optimal stochastic schedules, and in optimizing the schedules of a portfolio of assets in an illiquid market. Then we have discussed the intermediate and longer term risk management problems and how these real option models factor into trading and investment risk management, and portfolio optimization.

References

[1] A. J. Wood and B. F. Wollenberg, *Power Generation, Operation and Control*. New York: Wiley, 1996.
[2] A. Dixit and R. Pindyck, *Investment Under Uncertainty*. Princeton, NJ: Princeton Univ. Press, 1994.
[3] J. Hull, *Options, Futures and Other Derivatives*, 4 ed. Englewood Cliffs, NJ: Prentice-Hall, 1999.
[4] E. Schwartz, "The stochastic behavior of commodity prices: Implications for valuation and hedging," *The J. of Finance*, vol. 7, no. 3, July 1997.
[5] M. D. Ilic and P. Skantze, "Electric power systems operation by decision and control: The case revisited," *IEEE Control Syst. Mag.*, vol. 20, pp. 25–39, Aug. 2000.
[6] P. Skantze and M. Ilic, *Valuation, Hedging and Speculation in Competitive Electricity Markets: A Fundamental Approach*. Norwell, MA: Kluwer, Oct. 2001.
[7] E. Allen and M. Ilic, *Price-Based Commitment Decisions in the Electricity Market*. New York: Springer-Verlag, Apr. 1999.
[8] L. Clewlow and C. Strickland, *Energy Derivatives: Pricing and Risk Management*: LaCima Publications, 2000.
[9] S. Deng, B. Johnson, and A. Sogomonian, "Exotic electricity options and the valuation of electricity generation and transmission assets," in Chicago Risk Manage. Conf., Chicago, IL, May 1998.
[10] P. Jorian, *Value at Risk: The New Benchmark for Managing Financial Risk*, 2nd ed. New York: McGraw-Hill, 2001, p. 112 and 150.
[11] H. Markowicz, *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley, 1959.

Michael Denton, photograph and biography not available at the time of publication.

Adrian Palmer, photograph and biography not available at the time of publication.

Ralph Masiello (F'94), photograph and biography not available at the time of publication.

Petter Skantze, photograph and biography not available at the time of publication.