Valuation of Power Generation Assets: A Real Options Approach

Doug Gardner and Yiping Zhuang

Real options theory is an increasingly popular tool for valuing physical assets such as power generation plants. In this paper, we describe a model for power plant valuation that accounts for such important operating characteristics as minimum on- and off-times, ramp time, nonconstant heat rates, response rate and minimum electricity dispatch level. The power plant values and optimal operating policies are obtained by employing stochastic dynamic programming. Sample numerical results, using electricity price data from the New England power pool, show that operating constraints can have a significant impact on power plant values and optimal operating policies.

Deregulation of energy markets has dramatically changed the environment in which many power generation asset owners operate. In particular, utilities have become increasingly exposed to extremely volatile energy prices. Mismanagement of this risk exposure, even for an efficient power producer, may have a severe impact on its financial position.

The **real options approach** applies derivative pricing theory to the analysis of options opportunities in real assets (Dixit and Pindyck 1994). Unlike traditional discounted cash-flow analysis, real options theory explicitly accounts for flexibility in the manner in which an asset is developed and operated, often leading to higher asset values, as well as different optimal capacity planning and operating decisions. For example, accounting for different plant construction lead times in the face of demand uncertainty can lead to significantly different optimal capacity planning strategies (Gardner and Rogers 1999).

Valuing a power plant using real options theory has two main purposes in competitive markets. First, an investor who contemplates the purchase or sale of a power plant must accurately determine its value. The second purpose is to facilitate the use of risk management tools developed for financial markets in order to hedge both asset value and earnings. For example, a power plant can be hedged using forward electricity contracts (Eydeland and Geman 1998). In fact, much current trading activity in commodity derivatives is for precisely this purpose.

Ignoring non-fuel operating costs, the net profit per hour for a power plant is $q\left(P^E - HP^F\right)$, where q is the dispatch (or output) level (MW), P^E is the spot price of electricity (\$/MWh), P^F is the spot price of input fuel (\$/MMBtu), and H is the plant heat rate (MMBtu input fuel per MWh electricity). The quantity $(P^E - HP^F)$ is commonly referred to as the **spark spread** since it gives the difference between the price of electricity and the input fuel cost (expressed in $/MWh). If the spark spread is positive, the optimal dispatch level is full capacity; otherwise, the plant should not be run. The instantaneous plant pay-off per unit capacity is thus

$$max(P^E - HP^F, 0)$$

which is equivalent to an option to exchange one asset with price HP^F for another with price P^E. If it is assumed that the prices of electricity and fuel at some time in the future are lognormally distributed, then this option may be valued using Margrabe's exchange option formula (Margrabe 1978), an approach which is now widely used (Deng et al. 1998).

Given a set of available generating units, the lowest cost method to meet a power delivery commitment is via **merit order loading**: each unit is loaded to capacity in order of ascending operating cost until the required amount of power is made available. Plants with low-input costs ("baseload" generation) thus typically operate most of the year, while plants with high-input costs ("peakers") may be operated only a small fraction of time. In general, an optimal system design will include a mix of baseload, midload and peaking generation.

From an options perspective, **baseload generation** is normally an "in-the-money" option since the required electricity price at which it can be profitably operated is low. **Peaking generation**, on the other hand, is normally an "out-of-the-money" option since a high electricity price is required for profitable operation.

While the exchange option approach is useful from a conceptual perspective, it fails to account for some important plant operating characteristics that may affect plant value and the optimal operating policy:

- **Minimum on (up) and off (down) times**. These are imposed in order to limit the physical unit damage due to fatigue.
- **Minimum ramp (start-up) time**. Some time is required between the decision to turn on a unit and the time at which it is able to deliver power. In steam-powered generation units, for example, time is required to heat the boiler.
- **Minimum generation level**. Most units have some minimum level below which they cannot operate.
- **Response rate constraints**. Some time is required to effect a discrete change in the generation dispatch level.
- **Non-constant heat rate**. The heat rate of units normally varies with the generation level.
- **Variable start-up cost**. The cost of starting a unit may depend on the time spent off-line. For steam-powered units, for example, the boiler temperature declines with time spent off-line, increasing the cost to restart the unit.

When these operating characteristics are taken into account, the decision to turn on or off the plant depends not only on the market prices of electricity and fuel but also on the plant operating state, making the valuation problem path dependent. Johnson et al. (1999) provide results from a model that takes plant operating characteristics into account, but do not describe the underlying model. Tseng and Barz (1999) use Monte Carlo methods adapted to American options pricing. While their methodology is capable of describing most plant operating characteristics, it is computationally inefficient. Takriti et al. (2000) describe an efficient computational approach for determining the optimal dispatch of multiple plants under load and price uncertainty. However, the representation of uncertainty in the results they report is simplified for computational reasons.

This paper describes how stochastic dynamic programming can be used to calculate plant values and optimal operating policies while considering plant operating characteristics. Specifically, we first develop a lattice for the underlying stochastic variables. We then use backward dynamic programming to compute the plant value and optimal operating policy. The methodology is very similar to that proposed by Hull and White (1993) for path-dependent options or, in the context of energy derivatives, to the method discussed by Thompson (1994) and Jaillet et al. (1999) for swing options. Although this paper focuses on power generation plants, the same methods may also be applied to the valuation of other real assets such as energy pipelines and storage facilities. Finally, we should note that, in some markets, generation asset owners may have additional means of generating revenue (such as providing ancillary services) that may be significant (see Griffes et al. (1999) for a discussion); we do not consider these here.

The remainder of the paper is organized as follows. The next section describes the mathematical model, followed by a discussion of the solution method. This is followed by an application of the model, using data from the New England power pool. In this section, we investigate the effect of the operating constraints, electricity price volatility and the expected spark spread on the value of a power plant and optimal operating policies. The final section concludes with some thoughts on the benefits and application of real options theory in practice.

Model description

We focus on valuing thermal power units over a short-term horizon (e.g., one week). Plant values over a longer time horizon may be estimated by appropriately prorating the plant value from a number of representative subperiods.

Without loss of generality, we assume that operating decisions are made at hourly intervals $t = 0, 1, \ldots, T$. To model the plant characteristics described above, we introduce the notation summarized in Table 1.

Parameter	Description	Units
t_{on}	minimum up time	hours
t_{off}	minimum down time	hours
t_{cold}	additional time over the minimum down time after which the unit start-up cost is constant	hours
t_{ramp}	time required to bring the unit on-line	hours
q_{min}	minimum dispatch level	MW
q_{max}	maximum dispatch level	MW
$H(q)$	heat rate as a function of plant output q	MMBtus/MWh

Table 1: Summary of notation

Operating state constraints

To model the constraints on minimum on-, off- and ramp times, we introduce a state variable s representing the operating state of the plant. A state is a combination of a plant's condition and the duration in that condition. The total number of possible plant states is equal to the sum of the minimum on-, cool-down, minimum off- and ramp times. Hence, s is a number between one and $t_{off} + t_{cold} + t_{ramp} + t_{on}$, which implies that the number of states depends on the plant operating characteristics. Within this range, the plant condition may be described as shown in Table 2.

Constraints on plant state transitions may be represented graphically via a state transition diagram. Figure 1 represents possible states (indicated by circles) and state transitions

Plant Condition	States
Off-line	$1 \le s \le t_{off} + t_{cold}$
Ramp (unable to sell power)	$t_{off} + t_{cold} < s \le t_{off} + t_{cold} + t_{ramp}$
On-line	$t_{off} + t_{cold} + t_{ramp} < s \le t_{off} + t_{cold} + t_{ramp} + t_{on}$

Table 2: Plant operating states

(arrows between circles) for a power plant with a minimum on-time of two hours ($t_{on} = 2$), minimum off-time of two hours ($t_{off} = 2$), extra cool-down time of two hours ($t_{cold} = 2$) and a ramp time of one hour ($t_{ramp} = 1$). Each state is defined by the plant condition (off-line, ramp or on-line) and the duration of time in that condition.

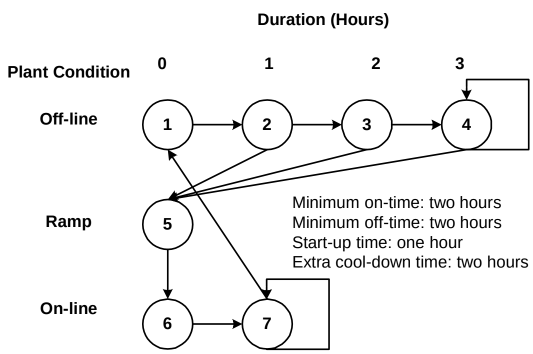

Figure 1: Feasible operating state transitions

State 1 represents a plant that has just gone off-line; given the minimum off-time restriction, the only possible transition one hour hence is to State 2, representing a plant that has been off-line for one hour. From State 2, a plant may either remain off-line (State 3) or start up (State 5), since one hour hence it will have been off-line for two hours. Similarly, from States 3 (off-line for two hours) and 4 (off-line for three or more hours), a plant may either remain off-line (State 4) or start up (State 5). States 3 and 4 are introduced only to model variable start-up costs, discussed later. Note that a plant cannot go directly from an off-line state to the on-line state due to the ramp time of one hour. Once started (State 5), a plant must go on-line (State 6); in this state, power may be produced for sale. A plant that has just gone on-line (State 6) must stay on-line, hence moving to State 7 (on-line for one or more hours), due to the minimum two-hour on-time restriction. Once in State 7, the plant can remain there or go off-line (State 1).

In the general case, feasible state transitions from State s at time t to State s' at time t + 1 may be represented mathematically as follows:

$$s' \in \begin{cases} \{t_{off} + t_{cold} + 1, s + 1\} & t_{off} \leq s < t_{off} + t_{cold} \\ \{t_{off} + t_{cold} + 1, s\} & s = t_{off} + t_{cold} \\ \{1, s\} & s = t_{off} + t_{cold} + t_{ramp} + t_{on} \\ \{s + 1\} & \text{otherwise} \end{cases}$$

which may be denoted more compactly simply as $s' \in A(s)$. The first case in this expression (corresponding to States 2 and 3 in the above example) shows that a plant that has been off-line for longer than the minimum off-time may either turn on or stay off, in which case it proceeds to the next off-line state. The second case (corresponding to State 4 in the above example) says that a plant that is currently in the final off-line state (off-line three or more hours) may either startup or remain in that state. The third case (corresponding to State 7 in the above example) shows that a plant that has been on for more than the minimum on-time may either turn off or stay on. The fourth case indicates that, in any other state (States 1, 5 and 6 in the above example), the plant must proceed to the next operating state.

Price processes

A key input to the model is the description of the price processes for fuel and electricity. Discrete time, discrete state price processes may be obtained as an approximation to continuous price processes. Let $j \in J_t$ represent the set of energy price states possible at time step t, and P^E_{jt} and P^F_{jt} be the spot price of electricity ($/MWh) and the spot price of fuel ($/MMBtu), respectively, at time t in energy Price State j.

As an example of how a price process may be represented, suppose the spot price of electricity follows a mean-reverting geometric Brownian motion process:

$$d\ln P^E = \left(u^E - a_E \ln P^E\right)dt + \sigma_E dW_E$$

where u^E is a drift parameter, a_E is the mean reversion rate, σ_E is the volatility and dW_E is the increment of a Brownian motion. A trinomial lattice may be used to represent this process, following the approach described by Clewlow and

Strickland (1999), which involves the determination of the price states at each time step and the associated price-state transition probabilities. As part of this procedure, the drift term u^E is made time dependent, in order that the lattice may be calibrated to an observed forward price curve and, hence, describes the risk-adjusted price process required for pricing derivatives. This approach may be viewed as an extension of the Hull and White (1994a) methodology for creating lattices for short-rate interest rate models.

To allow fuel prices to be stochastic in addition to electricity prices, a number of techniques are available for construction of two-factor lattices (see, for example, Hull and White (1994b)). Alternatively, if the plant heat rate is assumed to be constant, it is possible to model the process followed by the spark spread directly.

Costs and revenues

Each operating state has an associated cost or revenue. While these may be quite general, we assume here that they take the form

$$f_{jt}(q, s) = \begin{cases} -K_{fix} & 1 \le s \le t_{off} + t_{cold} + t_{ramp} \\ q(P^E_{jt} - H(q)P^F_{jt}) - K_{fix} & otherwise \end{cases}$$

Thus, there is a fixed cost K_{fix} in all states. In the on-line states, revenue equal to the product of the plant dispatch q and the spark spread is received.

In addition to costs and revenues associated with different operating states, there may be a transition cost $g(s, s')$ associated with operating state transitions. We assume here the following functional form for the transition cost of moving from State s to State s':

$$g_{jt}(s, s') = \begin{cases} P^F_{jt}\left(c_1\left(1 - e^{-c_2(s - t_{off})}\right) + c_3\right) & s' = t_{off} + t_{cold} + 1 \\ 0 & otherwise \end{cases}$$

where c_1, c_2 and c_3 are non-negative constants. Thus, the cost of starting up a plant is an increasing function of the time spent off-line and the prevailing fuel price.

Dispatch and response rate constraints

Since the plant can only produce power in an on-line state and the output level is bounded, the following constraints are satisfied in the absence of response rate constraints:

$$q = 0 \quad if\ s \in \{1, \ldots, t_{off} + t_{cold} + t_{startup}\}$$
$$q_{min} \le q \le q_{max} \quad otherwise$$

Dispatch levels, q, satisfying these constraints are denoted $q \in B(s)$. Note that these constraints impose no restriction on how fast a plant can change its dispatch level: if it is on-line, it can be dispatched at any level between the minimum and maximum output levels.

To model response rate constraints, we discretize the possible plant dispatch levels and then add the plant dispatch level as a third dimension to the state descriptor (in addition to plant condition and duration). The state transition constraints and dispatch constraints must also be appropriately modified. To illustrate the steps required, we extend the previous example.

Suppose that changing the dispatch level from the minimum to the maximum dispatch level (and vice versa) may be achieved in a minimum of one hour. In this case, the modified state transition diagram may be represented as in Figure 2. In this figure, States 6 and 7 have been redefined as being on-line at the minimum dispatch level. An extra state (State 8) has also been added, defined as being on-line for one or more hours and being dispatched at the maximum dispatch level. From State 6, it is possible to stay on-line at either the minimum or maximum dispatch levels (States 7 and 8, respectively). From State 7, it is possible to go off-line, or stay on-line at either the minimum or maximum dispatch level. From State 8, it is possible to stay on-line at either the minimum or maximum dispatch level; we have assumed it is not possible to go directly to the off-line state from the maximum dispatch level. Similarly, we have assumed that it is not possible to go from the ramp condition (State 5) to the maximum dispatch level directly.

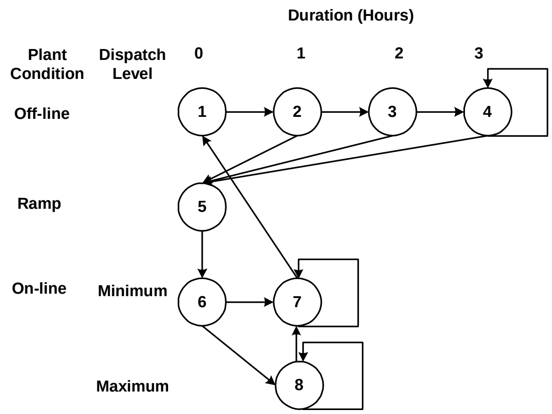

Figure 2: Feasible operating state transitions with response rate constraints

Solution method

To obtain the plant value together with the optimal operating policy, we employ stochastic (or probabilistic) dynamic programming (see Wagner (1975) for an introduction). Dynamic programming is a standard technique for solving optimization problems that may be formulated in a set of stages or time periods. An optimal policy with n stages (or periods) remaining may be determined by selecting the policy that maximizes the sum of net revenue in stage n plus the expected net revenue in the subsequent $n - 1$ remaining stages.

The optimal policy for this problem is found by solving

$$F_{jt}(s) = \max_{q \in B(s)} f_{jt}(q, s) + \sum_{j'} p_{jt}^{j'} \left\{ \max_{s' \in A(s)} [F_{j', t+1}(s') - g_{jt}(s, s')] \right\} \quad (1)$$

where $F_{jt}(s)$ denotes the value of the power plant over the period t to T, conditional on being in energy Price State j at time t and operating State s; $p_{jt}^{j'}$ represents the probability of moving from Price State j at time t to Price State j' at time $t + 1$.

Equation 1 states that the value of the power plant over the remaining stages (i.e., from time t to T) is the sum of two terms. The first term is the net revenue in period t. We choose the optimal plant output level q to maximize the net revenue, subject to the plant operating constraints $B(s)$. The second term is the expected value of the power plant from time $t + 1$ to T, which is conditional on the plant operating state at time $t + 1$. We select the operating State s' that results in the maximum plant value (net the state transition cost), conditional on the requirement that it is a feasible transition from State s. This maximization determines the optimal operating state transition policy for the plant. This optimal state transition policy is a generalized version of the optimal exercise boundary that is obtained as part of the valuation of American-style options.

The plant value at time 0 is obtained by solving Equation 1 recursively, working backward from time T for all possible Price States $(j \in J_t)$ and operating States $s \in S$, to time 0 (which has only a single known price state). $F_{0,0}(s)$ then represents the value of the plant over the entire period, conditional on being in State s at time $t = 0$.

In addition to plant value, a key output of the solution procedure is the optimal operating policy which consists of the optimal plant output in each on-line state as a function of price state and time, and the optimal state transition strategy as a function of the current operating state, price state and time. The optimal operating policy should be used by plant operators to maximize the plant value—operating the plant using a different operating policy is, by definition, suboptimal and, hence, will result in a lower plant value.

Typically, the optimal state transition strategy may be expressed in terms of a set of exercise boundaries. For example, if the current plant state is on-line, the optimal transition in the next period will be to remain on-line for all values of the spark spread greater than a certain critical value and to go off-line (assuming this is a feasible transition) for all spark spreads that are less. Given the cyclical variation in electricity prices over time, this critical value will also vary through time.

To illustrate the methodology, we consider the following simple example. Consider the valuation of a one-MW capacity plant that has a minimum on-time of two hours, a minimum off-time of one hour, and no start-up or cool-down time. We also assume there are no fixed or start-up costs. The state transition diagram for this plant is shown in Figure 3. We assume the plant's heat rate is constant for all levels of output and that the minimum generation level is 0.5 MW. Since the plant heat rate is constant, we may represent the price state by the spark spread. We assume the spark spread evolves according to the three-period binomial lattice shown in Figure 4. In hour 0, the spark spread is 0 and may go to either 8 or –2 with equal probability in hour 1. In hour 2, possible Price States are –4, 6 and 16.

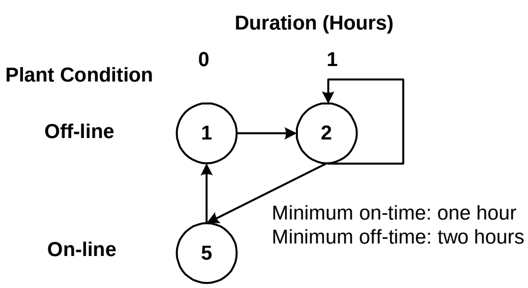

Figure 3: Feasible operating state transition

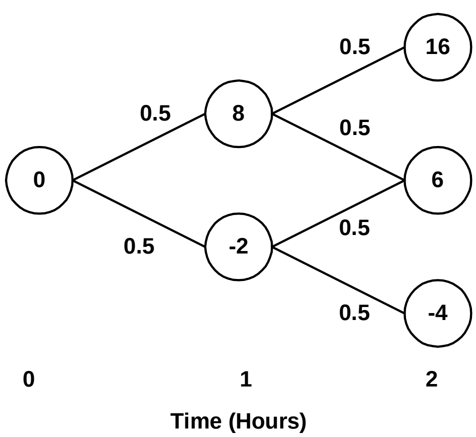

Figure 4: Spark spread lattice

The values of $F_{jt}(s)$ and the optimal operating policy are shown in Figure 5. Each table in this figure shows the optimal dispatch level q, the optimal state transitions s' for "up" and "down" price moves, respectively, and the plant value F_{jt} as a function of the time and price state. The plant value over the three periods is 7, 11 or 10.5, depending on whether the plant state is initially in State 1, 2 or 3.

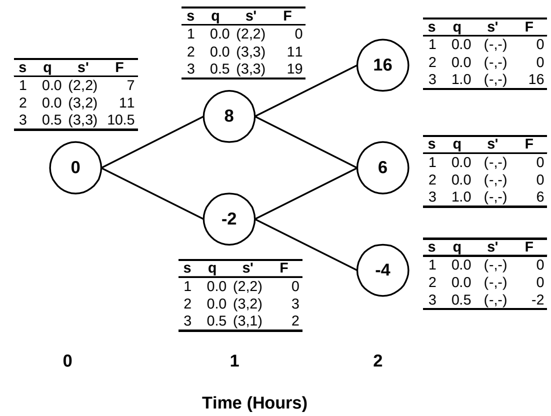

Figure 5: Plant value and optimal operating policy

A sample calculation is as follows. Beginning at the terminal period (hour 2), consider the plant value in Price State 16 assuming the plant is on-line (State 3). Since the spark spread is positive, the optimal dispatch level is one MW (i.e., full capacity) for current period net revenue of $1 \times 16 = 16$. As this is the final period, this is also the plant value. In Price State –4, the spark spread is negative, so the optimal dispatch level in State 3 is only 0.5 MW (i.e., the minimum generation level) for a plant value of $0.5 \times (-4) = -2$.

Now consider the plant value in the preceding hour (hour 1) in operating State 3, Price State –2. Once again the spark spread is negative so the optimal dispatch level is 0.5 in the current period for net revenue of $0.5 \times (-2) = -1$. To this, we must add the value over the remaining stage. From State 3, the feasible state transitions are to State 1 or State 3. We must choose the optimal transition for each possible price state in hour 2 (this is the second maximization in Equation 1). If the price state in hour 2 is 6, the realized values for States 1 and 3 are 0 and 6, respectively; hence the optimal state transition is to State 3. On the other hand, if the price state in hour 2 is –4, the realized values are 0 and –2; hence the optimal state transition is to State 1. The probability-weighted sum for these two possible price states is then $0.5 \times (6) + 0.5 \times (0) = 3$. When added to the net revenue of –1, we obtain a value of 2, as reported in the associated table in Figure 5.

Of note is the fact that state transition decisions take into account not just immediate net revenue but also the opportunity cost in terms of future decision-making flexibility; the simple exchange option approach does not consider this. Consider the optimal state transition from hour 0, operating State 3, for example. In hour 1, the possible Price States are 8 and –2. It is no surprise that the optimal decision is to stay turned on (State 3) in the former. What is perhaps surprising is that this is also the optimal decision in the latter, since the net revenue will necessarily be negative in this case. The reason for this behavior may be explained by the fact that moving to State 1 (off for zero duration), rather than State 3, would prevent the plant from taking advantage should the spark spread become positive in hour 2, due to the minimum off-time constraint. This phenomenon, in fact, explains why electricity prices have gone to zero or even have become negative for short time periods in some markets.

Also of interest is the fact that the optimal state transition may vary depending on the current operating state, even when the possible state transitions appear similar. For example, in hour 0, should the next hour Price State be –2, the optimal decision is to stay off (State 2), if the current state is State 2, and to stay on (State 3) if the current state is State 3; this, despite the fact that turning on the plant is possible in the first case and turning off the plant is possible in the second. The explanation for this lies in the fact that the possible choices in the two cases, though similar, are not identical: the off option in the first case is to State 2, while the off option in the second case is to State 1.

Numerical results

In this section, we use the methodology developed above to value a power plant over a time horizon of five days. We have assumed that the price of input fuel is constant over the operating period. Given the much greater volatility of electricity prices, the results are unlikely to differ significantly over the time horizon we consider.

We estimated the parameters for this model using hourly electricity spot price data for the New England power pool (ISO New England 2000). Figure 6 shows, for each hour of the day, the mean electricity price in the months of March and June 2000. Of note is the fact that electricity prices are on average lowest during the "off-peak" hours from 11:00 p.m to 7:00 a.m. For each hour of the day, prices in June are higher than those in March, although the pattern is different. The highest average prices in March are recorded in the evening (7:00–9:00 p.m.) driven by domestic lighting and appliance requirements, the highest average prices in June are around 12:00 p.m., driven by demand for air-conditioning.

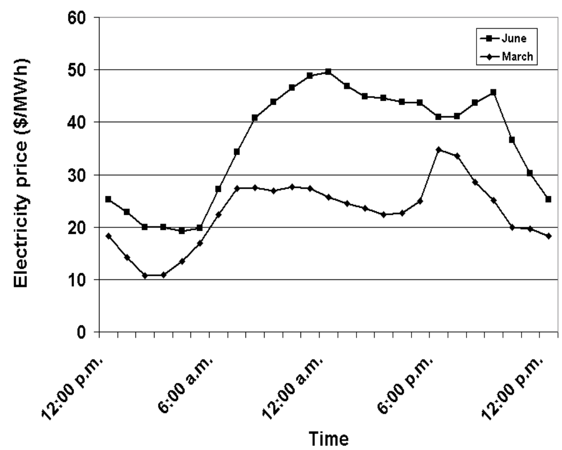

Figure 6: Mean price of electricity in New England power pool, 2000, by hour of day

The estimated model parameters for March are $\sigma_E = 3{,}023\%$ and $a_E = 2{,}899$, for June, $\sigma_E = 2{,}733\%$ and $a_E = 2{,}286$; all figures are expressed on an annualized basis. Note that the extreme volatility of the New England market is common to deregulated electricity markets. Volatility in equity markets is typically one or two orders of magnitude less. The mean reversion of prices is also very strong. For this model, the "half-life" of deviations, or the expected time for a deviation from the mean to halve, is $ln(2)/a_E$. Thus the half-lives of deviations in March and June are only 2.1 and 2.7 hours, respectively.

For the purpose of this example, the price process was calibrated to the mean hourly electricity curve; in practice, market-quoted forward prices should be used where available. Except as noted elsewhere, all other parameter values used are listed in Table 3.

Parameter	Value
q_{min} (MW)	100
q_{max}(MW)	500
H (MMBtu/MWh)	10
K_{fix} (\$/h)	0
P^F (\$/MMBtu)	2
c_1 (\$/h)	0
c_2 (\$/h)	0
c_3 (h)	0

Table 3: Base case parameter values

In order to determine the importance of different operating constraints on plant value, we consider a variety of cases summarized in Table 4. Given the assumption that operating decisions are made at hourly intervals, Case 1 corresponds to having no operating constraints. The plant values obtained for this case are thus equal to those obtained using the exchange option approach. Cases 2 through 4 consider the impact of adding a minimum off-time, minimum on-time and ramp time, respectively.

Case 5 considers the impact of adding a minimum on-time and a response time of one hour. Case 6 considers the impact of increasing the minimum dispatch level from 100 to 250 MW in combination with an increased minimum on-time. Case 7 considers the impact of increasing the minimum off-time, ramp time, minimum on-time and response rate together. Case 8 is identical to Case 7 except that the minimum dispatch level is greater.

Figures 7 and 8 illustrate the optimal operating policy boundaries for Cases 4 and 7, respectively, based on the electricity process parameters for March. The "turn-on" boundary is the spark spread above which the plant should be turned on, if it is currently off and may be turned on (i.e., on-line must be a feasible state transition). The "turn-off" boundary is the spark spread below which the plant should be turned off, if it is currently on and may be turned off.

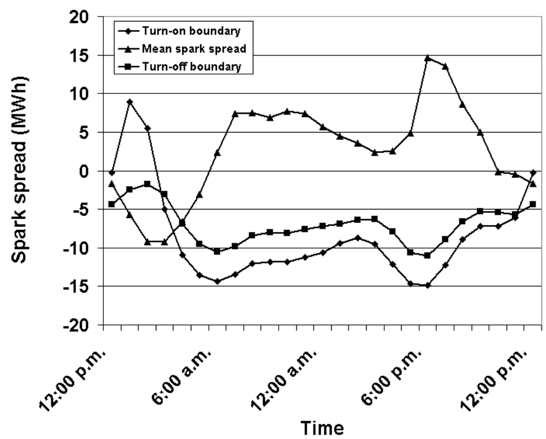

Figure 7: Optimal dispatch boundaries (\$/MWh), Case 4

	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6	Case 7	Case 8
Minimum off-time (h)	1	4	1	1	1	1	4	4
Ramp time (h)	0	0	0	2	0	0	2	2
Minimum on-time (h)	1	1	8	1	8	8	8	8
Response time (h)	0	0	0	0	1	0	1	1
Minimum dispatch level (MW)	100	100	100	100	100	250	100	250

Table 4: Summary of test case characteristics

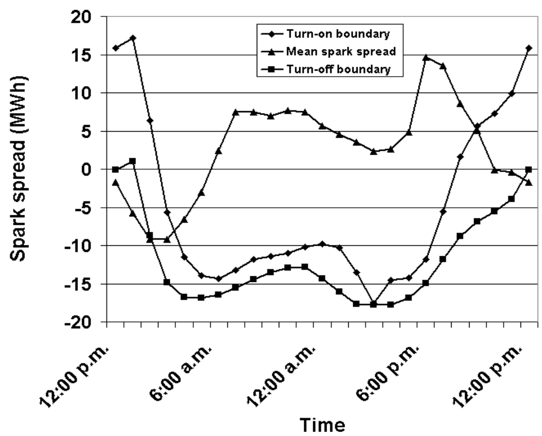

Figure 8: Optimal dispatch boundaries ($/MWh), Case 7

If the plant is currently on, it will tend to be uneconomic to turn it off if the current spark spread is positive: immediate revenue is foregone and a decision to turn off the plant in the next hour can always be made if the spark spread becomes negative. Thus, the turn-off boundary is roughly bounded above by zero. The smaller the spark spread is expected to be, the more willing an operator should be to turn off the plant than otherwise. Thus, the turn-off boundary peaks around midnight, since the expected spark spread over the next several hours is negative. If the spark spread is expected to become positive, the operator should be less willing to turn off the plant, given the time required to restart it. Hence, the troughs in the turn-off boundary occur in the hours in which electricity prices are highest. This effect will be more pronounced the longer it takes to come back on-line once the plant shuts down (i.e., the greater the minimum off and ramp times): this explains the lower off-on boundary in Figure 8 (around –$17) than in Figure 7 (around –$10).

If the plant is currently off, the decision to turn it on is determined by the expected spark spread at the time the plant would come on-line, accounting for the time required for start-up. Hence, if the spark spread is expected to be positive, an operator should be willing to turn on the plant even if the current spark spread is unfavorable. This explains why the turn-on boundary is low during peak hours. Conversely, if the spark spread is expected to be negative, the operator should be less willing to turn on the plant even if current prices are high: this explains the high turn-on boundary around midnight. Furthermore, in this situation, the greater the minimum on-time, the greater the turn-on boundary peak, since the period in which the plant would be forced to be on is longer should the spark spread actually become negative. This explains why the turn-off boundary peaks at a higher value in Figure 8 (around $17) than in Figure 7 (around $9).

A key parameter in the analysis is the expected spark spread. If one assumes that the hourly electricity price pattern is roughly the same in each month (only shifted up or down), then, by calculating the value of a plant as a function of the expected spark spread, one may estimate the plant value for any month and for any fuel price and heat rate. Figure 9 shows the results for each case.

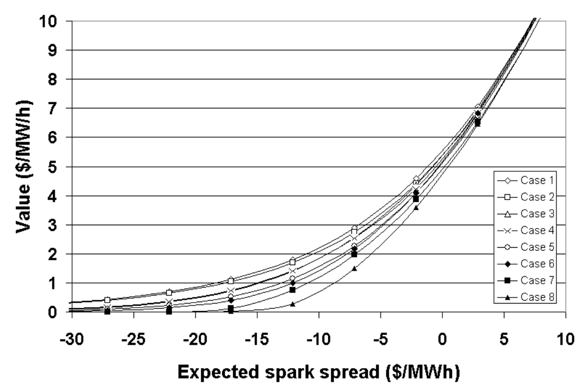

Figure 9: Plant value versus expected spark spread

As expected, the greater the expected spark spread, the higher the plant value. In fact, the shape of this function is much the same as that of a call option as a function of the strike. An expected spread of zero corresponds to an option that is "at the money." Of particular note is the fact that a plant with zero "intrinsic" value (i.e., the expected spark spread is zero) has a significant value.

Figure 10 illustrates the difference between the plant value in Case 1 and the other cases. Depending on the expected spark spread, the differences may be significant. For example, with an expected spark spread of –12 $/MWh, the plant value in Case 8 is 1.5 $/MWh less than the plant value in Case 1—a reduction of 85%.

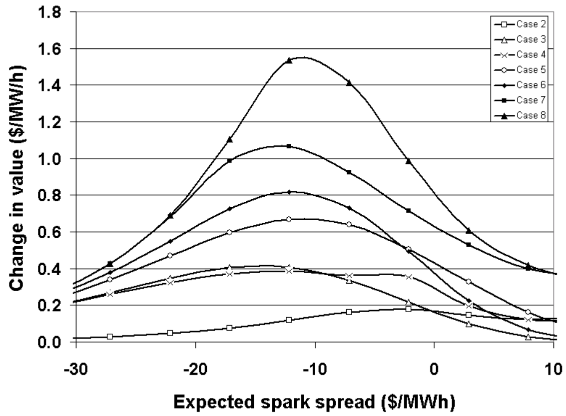

Figure 10: Decrease in plant value (difference from Case 1)

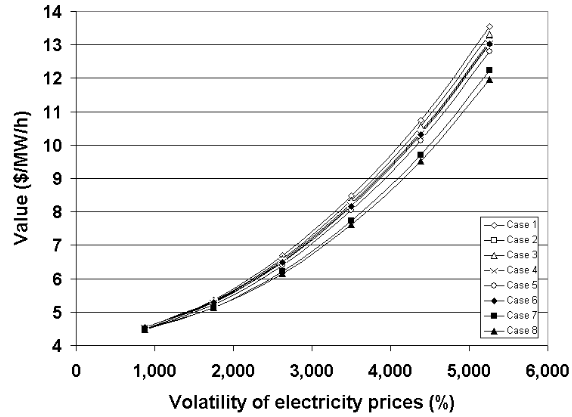

Figure 11: Plant value versus volatility

In almost all cases, the difference in value is low for very low expected spark spreads (for which the plant value is low in any case), rises for moderately negative spark spreads and declines as the spark spread increases. The decline in importance of operating constraints for high spark spreads may be attributed to the fact that the plant is expected to run a greater fraction of time. For very high spark spreads, the plant will almost surely be run continuously, implying that operating constraints assume almost no importance. Operating constraints are most important for plants that are expected to be turned on and off frequently.

Figure 11 shows the plant value for each case as a function of the volatility of electricity prices. In all cases, increasing volatility leads to higher plant values. This is expected: the price of any vanilla option increases with higher volatility. Comparing Case 1 to the other cases, the impact of the operating constraints is also seen to be an increasing function of volatility, both in absolute and percentage terms. The greatest absolute and percentage differences (1.6 $/MWh and 12%, respectively) are obtained with Case 8 when volatility is in the 5,000% range. Intuitively, since operating constraints reduce flexibility to respond to price changes, their impact will be greater the higher the level of uncertainty regarding those prices.

Conclusion

In this paper, we describe how real options theory may be applied to value power generation assets. In particular, the model we develop is capable of handling constraints related to minimum on- and off-times, ramp times, minimum dispatch levels and response rates. Numerical results illustrate that these constraints may have a significant impact on the power plant's value, particularly for plants that are just slightly "out of the money." The optimal operating policy also may be significantly affected.

Real options theory provides a methodology for quantifying the value of the operating flexibility of real assets and for determining optimal operating policies. It offers the potential to improve greatly the effectiveness of operating decisions and to unlock "hidden" asset value. Understanding the sources of asset value and its sensitivity to fuel and electricity prices is also critical for companies seeking to determine a suitable hedging policy through either forward sales or other derivatives contracts. As with any theory, effective application of the insights provided by real options theory requires that managers become familiar with its underlying assumptions in order to understand both its strengths and weaknesses. The pay-off for companies that are able to do so is the ability to effectively leverage a company's assets to achieve an optimal trade-off between risk and reward.

References

Clewlow, L. and C. Strickland, 1999, “Valuing energy options in a one-factor model fitted to forward prices,” University of Technology, Sydney, Working Paper.

Deng, S., B. Johnson and A. Sogomonian, 1998, “Exotic electricity options and the valuation of electricity generation and transmission assets,” Proceedings of the Chicago Risk Management Conference.

Dixit, A. and R. Pindyck, 1994, *Investment Under Uncertainty*, Princeton, NJ: Princeton University Press.

Eydeland, A. and H. Geman, 1998, “Fundamentals of electricity derivatives,” *Energy Modelling and the Management of Uncertainty*, London, UK: Risk Books.

Gardner, D. and J. Rogers, 1999, “Planning electric power systems under demand uncertainty with different technology lead times,” *Management Science* 45(10): 1289–1306

Griffes, P., M. Hsu and E. Kahn, 1999, “Power asset valuation: real options, ancillary services and environmental risks,” *The New Power Markets*, London, UK: Risk Books.

Hull, J. and A. White, 1993. “Efficient procedures for valuing European and American path-dependent options,” *Journal of Derivatives*, 1(1): 21–31.

Hull, J. and A. White, 1994a, “Numerical procedures for implementing term structure models I: single-factor models,” *Journal of Derivatives* 2(1): 7–16.

Hull, J. and A. White, 1994b, “Numerical procedures for implementing term structure models II: two-factor models,” *Journal of Derivatives* 2(2): 37–48.

ISO New England Web site: http://www.iso-ne.com/Historical_Data/hourly_data/2000_hourly_data.txt, accessed November, 2000.

Jaillet, P., E. Ronn and S. Tompaidis, 1999, “Modeling energy prices and pricing and hedging energy derivatives,” University of Texas, Austin, Working Paper.

Johnson, B., V. Nagali and R. Romine, 1999, “Real options theory and the valuation of generating assets: a discussion for senior managers” *The New Power Markets*, London, UK: Risk Books.

Margrabe, W., 1978, “The value of an option to exchange one asset for another,” *Journal of Finance* 33(1): 177–186.

Takriti, S., B. Krasenbrink and L. Wu, 2000, “Incorporating fuel constraints and electricity spot prices into the stochastic unit commitment problem,” *Operations Research*, 48(2): 268–280.

Thompson, A., 1994, “Valuation of path-dependent contingent claims with multiple exercise decisions over time: the case of take-or-pay,” *Journal of Financial and Quantitative Analysis*, 30(2): 271–293.

Tseng, C. and G. Barz, 1999, “Short-term generation asset valuation,” Proceedings of the 32nd Hawaii International Conference on System Sciences.

Wagner, H., 1975, *Principles of Operations Research*, 2nd edition, Englewood Cliffs, NJ: Prentice-Hall.