Lacima GROUP

Innovations in risk analytics

How to accurately value thermal generation assets for greater profits

A white paper

by
Lacima Group

November 2009

1. Introduction

For many participants involved in power trading such as energy merchants, utilities, banks, and other investment companies, accurate valuation of generation assets and their associated risk calculations is crucial. Not only does rigorous valuation provide a more accurate view of a portfolio's current value, it also enhances the ability to properly manage and hedge the risks associated with generation assets. It is common to view the flexibility involved in the operation of generation assets as a "real option" from a valuation perspective. However, each plant also has certain operational constraints that affect how much flexibility the user has in operating the plant.

In this paper, we will detail the two most common ways in which industry practitioners typically value the flexibility of generation assets as real options while taking into account the constraints in operating the plant. As we will show, some of these methods can be relatively simple – for example, valuing the asset as a portfolio of spread options – while others are more complex, and take into account more of the asset's operational constraints. The level of complexity involved often depends on the type of generation asset that is being modelled, and so we begin by summarising many of the features of thermal generation assets and then progress to describe common techniques used to value them.

2. Properties of generation assets

Generation assets have many different properties. The following lists many of the most common properties that are routinely taken into account when valuing and measuring risk of generation assets:

- Maximum capacity
- Minimum stable generation
- Heat rate
- Variable operational and maintenance cost
- Start cost
- Ramp-up rate
- Ramp-down rate
- Minimum up time
- Minimum down time
- Emissions
- Outages
- Scheduled maintenance
- Fuel transportation costs
- Power transmission costs

Of the properties listed above, maximum capacity is probably most familiar to the reader, representing the maximum amount of power that can be produced in an hour. The maximum capacity can change from month to month, with the fluctuation in the capacity occurring because the capacity is dependent on the thermal gradient between the generation unit and the ambient air temperature. On the other hand, minimum stable generation is the lowest generation level that the unit can operate at, and still produce power that can be sold to the grid. The ability to run at the minimum stable generation allows the operator to run the unit at a minimal loss if the unit must be kept running for some reason.

The concept of the heat rate is also probably familiar to most readers, with the value representing the efficiency of the unit. As is the case with maximum capacity, the heat rate can also vary through time. A larger thermal gradient between the generator and the ambient air temperature improves the unit's efficiency. The efficiency can also be improved by running the unit at maximum capacity. In models that more accurately capture the operation of the generation unit, the unit can be dispatched at a capacity between the minimum stable generation and the maximum capacity, and in these cases a full heat-rate curve needs to be provided. The heat rate curve can be described via a step function or alternatively as a continuous function.

Variable operation and maintenance (VOM) costs represent the non-fuel costs associated with running the unit. Start costs are charges associated with starting the unit. Some of these charges are costs that actually occur such as the purchase of start fuel, which is the fuel consumed while getting the unit up to producing the minimum stable generation. Other start costs are included to account for the wear and tear on a unit caused by stopping and restarting the unit.

As we will see later, it can be challenging to account for these costs in any analytic valuation solution. One of the things that make start costs particularly challenging is the fact that they may be dependent on how long the unit has been off. Units that have just come off line are usually cheaper to bring back on line than similar units that have been off for a considerable amount of time. Ramp-up and ramp-down rates limit how quickly the unit can change its operating level of generation. Similar to start costs, ramp-up rates are also frequently dependent on how long the unit has been off line. Minimum up/down times control that if the unit is turned on/off that the unit remain in that state for a minimum number of hours. These constraints are put in place to minimise excessive wear and tear the unit would experience if it were constantly switched on and off.

Emissions costs associated with CO_2, NO_x and SO_2 are an important component to include in asset valuations. Since most emission markets are still relatively immature, it is often difficult to estimate reliable parameters for modelling their prices, and so these costs are often included as deterministic values and treated much like VOM costs. When parameters can be estimated, the incorporation of stochastic prices for emissions complicates the modelling process and typically requires Monte Carlo simulation methods to obtain valuations.

Every unit has the possibility of suffering a forced outage – random outages that reduce the operating capacity of the unit or take it off line completely for an extended period of time. In addition to forced outages there are scheduled maintenance outages. Although schedule maintenance outages are planned, they may take at least several days if they are for major issues. Both of these kinds of outages can have significant impact on the valuation of the plant. Finally, since, typically, generation units are not located at the gas supply or the load centre, we have to account for costs associated with getting the gas to the plant and the power to the grid. These costs may show up in the form of adders, multipliers, taxes or losses.

3. Real Options vs Financial Options

One of the benefits of treating an asset as a real option is that we can make use of the many techniques that have been developed for the valuation of financial options. As we have seen, there are a number of constraints that we would like to take into account when valuing a generation asset. Consequently, it is worth understanding the difference between financial and real options so that we understand the limitations these techniques impose on us when they are used to value generation assets.

Firstly, typically financial options are paid for 'up front' and there is no significant cost to exercising the option. As we have seen, there usually is a start cost associated with the generation asset. Since the start charge is accounted per start and not per hour run, it is more complicated to implement start costs in a closed-form solution than in a Monte Carlo solution.

The second major difference between financial options and generation assets is that once the financial option matures, we can immediately exercise it. Generation assets, on the other hand, have a ramp rate, which implies that we can't instantly go from having the unit off to running at maximum capacity. In other words, we need to decide to exercise the real option of the generation asset before its 'expiry'.

Thirdly, most financial options can have the payoff described in a single payoff function that can be easily written down. The operational constraints of a generation asset such as start costs, ramp rates, and minimum up/down times, require us to keep track of prior states of the unit. This requirement makes it difficult to write out a simple payoff function for a generation asset with all the operational constraints. In order to make use of many of the standard techniques from financial options, many of the constraints of the generation asset are typically ignored or modelled in a less than ideal way.

There are two main methodologies of valuing a generation asset as a real option; as an analytic approximation, or via Monte Carlo techniques – we now go on to describe the advantages and disadvantages to each of the techniques.

4. Real Option Valuation – Analytic Spark-Spread Option

There is a long history of using spread options to value many different kinds of energy assets as real options; for example, refineries, (crack-spread options); storage assets (calendar-spread options), and transportation/transmission systems (geographic-spread options). It is a natural step to treat a generation asset as a spark-spread option. A spark-spread option is an option on the spread between the power price and the input fuel price used to generate it. The advantage of this approach is that it is very simple and easy to obtain a quick evaluation of the asset. The payoff function for a spark-spread option maturing at time *T* is given by

$$Payoff = Q \times \Delta t \times Max(P_T - HR \times G_T - K, 0) \ ,$$

where Q is the maximum capacity, Δt is the time the unit is generating power, P_T is the power price at the maturity of the option, HR is the heat rate, G_T is the natural gas price, and *K* is a fixed strike. The fixed strike will be composed of the VOM as well as other costs from the other operational constraints. A transformation of the well-known Black-Scholes model can be used to value options with this payoff.

This valuation method is easy to implement, straightforward to understand, and provides a good basis for a rough estimate. However, it doesn't account for many of the constraints that we listed earlier in this paper. For those that can be handled, they are only handled in a very crude fashion. For example, emissions costs can be treated as a fixed cost that is incorporated in the strike term. The cost per start can be incorporated similarly. However, the start gas needs to be handled in an alternative way because the gas price is stochastic. Typically, start gas is incorporated by adjusting the heat rate input.

The only way to handle outages, forced or scheduled, in this analytic framework is by de-rating the volume. If the unit is expected to be forced out for 10% of the hours run then the hours run should be scaled down by a similar percentage. Unfortunately, this does not capture the real risk of a forced outage. We are also unable to capture the effects of ramp rates and minimum up/downtime constraints in this methodology.

One of the main limitations of the analytic approach is that there is no practical way to make the dispatch decision in the current period dependent on the dispatch decisions in prior periods, making it impossible to handle constraints such as ramp rates, and minimum up/down times. A further limitation is that the blocks of time typically quoted for forward power prices do not describe the hourly variability in observed prices – this hourly variability is crucial to the valuation of peaking units.

5. Real Option Valuation – Monte Carlo simulation

A solution to a number of the issues of the analytic technique is to implement hourly simulations of the spot power price, which gives us the granularity in prices to determine how the plant should be run each hour. In this section we describe the use of Monte Carlo methods for valuing generation assets in more detail. In particular, we will discuss the appropriate price models to use and how to develop path-dependent dispatch algorithms to handle the operational constraints of the assets.

i) Appropriate price models

The choice and use of an appropriate power price process is critical for accurate valuation because the resulting prices affect how we run the generation asset within the confines of the operational constraints. If we use an unrealistic power price process the model may tell us to operate the generation asset differently than we would expect to run it in the real world. This can be especially true for peaking units that substantially realise their profits in a small number of hours.

Realistic simulation-based methodologies for valuing generation assets require simulated power prices that are consistent with the behaviour seen in real-world power markets. These prices tend to exhibit spikes and very strong mean-reversion effects, and these can be seen in the hourly spot prices for a location within the New York Independent System Operator (NYISO) from 2005 through 2008, as shown in figure 1.

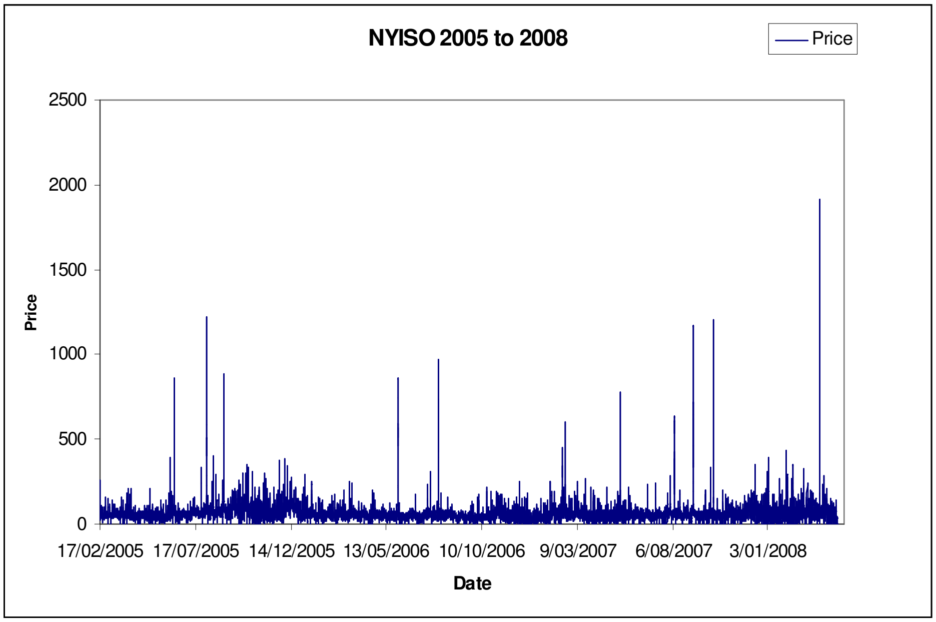

Figure 1: NYISO hourly spot prices from 2005 through 2008

The very large spikes in the data make it difficult to see and understand the behaviour of prices that are not spikes. By looking at smaller segments of the data, we can see that power prices are very volatile even without the large price spikes and that they exhibit strong mean reversion and that the price level changes constantly through out the day. Using a realistic price process is crucial – a price process that does not correctly account for jumps, mean reversion, and other properties of the data can incorrectly value the generation asset and will provide inaccurate risk metrics such as earnings-at risk or gross margin-at-risk.

There are two main types of models that practitioners use to implement power-price models. The first set of models are pure stochastic models that attempt to simulate the price directly as a stochastic process, with process parameters typically derived from both market data and historical data. Alternatively, we could use a hybrid model – hybrid models make use of functional relationships between power prices and the fundamental drivers of those prices. Fundamentals are used to represent supply and demand relationships, while stochastic processes are used to represent the evolution of the underlying drivers.

ii) Incorporating Physical Constraints

With the consideration of the physical constraints, the evaluation of a generation plant's profitability becomes far more complicated. The decision to operate the plant in a given hour is no longer simply based on the cashflow for that hour, but must also take into consideration how the unit was operated in prior hours. In order to determine when the plant will run, we need to develop a dispatch algorithm, which not only takes into account the economic situation but also meets the physical constraints. As part of these constraints we are able to simulate emissions prices and outages, greatly improving the accuracy of the valuation.

iii) Incorporating Dispatch Strategies

There are many ways to determine how to dispatch the generation asset. One can develop sophisticated algorithms like linear or dynamic programming. However, both these approaches require a substantial amount of time to design the algorithm. In the case of a linear programming approach, we need to determine all the equations that will be used to model our constraints. The dynamic programming approach requires states and transitions to be defined, with the number increasing as the number of constraints increases. Both the linear and dynamic programming approaches are guaranteed to determine an optimal solution, but can be difficult to implement.

Although they are not guaranteed to determine the optimal solution, heuristic algorithms are often satisfactory. Because heuristic algorithms are a collection of simple rules, they are often easier to design and implement. Additionally, if new constraints are added, it is fairly simple to add extra rules for the inclusion of the new constraint.

The differences between an optimal solution and a suboptimal solution are usually insignificant when compared to the difference between implementing a realistic power price model and an unrealistic power price model, on the valuation of a generation asset. A solution from a good heuristic model may turn the generator on an hour too late or too early from time to time. This will have less effect on the value of a plant than an unrealistic model for hourly power prices. The unrealistic model will ultimately affect the valuation of the generator on all hours.

One drawback of using Monte Carlo simulation is that many implementations of the technique suffer from an effect sometimes termed, 'the perfect foresight problem'. This occurs when the dispatch of the generation asset is optimised over the complete simulated path of spot prices. This is effectively assuming that the plant operator can choose the best spot prices in the future at which to dispatch the plant. In reality a plant operator would never know exactly what the spot prices will be in the future – instead relying on experience and expectations of future prices. This means that if there is an unexpected spot price realisation the generator may lose out on extra profit or run at a loss.

The way around this shortcoming is to dispatch the plant at each point in time based only on the simulated spot and forward prices at that point in time and then use the future simulated spot prices to calculate the profit and loss.

Although this is not that much more complicated than the normal Monte Carlo method it does require additional time due to the extra calculations.

6. Types of generation assets

We can group generation assets in to three broad categories based on how they are expected to serve load:

- Base load
- Mid-merit
- Peaking

Base-load units, as the name implies, handle the base load of the grid. They tend to be efficient units with low heat rates and consequently, typically operate in-the-money with most of their value being intrinsic value rather than extrinsic. On the negative side, they have high start costs and have long minimum up and down times. It also takes a long time to ramp the unit up to the maximum capacity. These constraints tend to minimise the amount of optionality that we have for base-load units. Since base-load units tend to be in-the-money and have some significant operational constraints, the use of an analytic spark-spread option to value them is often seen as a good approximation.

Mid-merit units, on the other hand, tend to be less efficient than base-load units and/or use a costlier fuel, which often puts mid-merit units at the cusp of being used. In terms of moneyness, these units are considered to be at-the-money and as such have the largest amount of extrinsic value and, consequently, it is very important to properly value the optionality of a mid-merit unit. Mid-merit units typically have lower start costs, shorter minimum up/down times, and less time to ramp up to maximum capacity than baseload units. This implies that the constraints do not constrict the optionality of the mid-merit unit as much as they do for baseload units. Consequently, using an analytic spark-spread option model for a mid-merit unit would be less than ideal, and so many users prefer to value mid-merit units using the Monte Carlo techniques that we discussed earlier.

Peaking units tend to have very high heat rates and as such are seen as out-of-the-money options with all of their value seen as extrinsic value. Peaking units are designed to be able to meet sudden peaks in demand and consequently have low start costs, short minimum up/ down times, and can quickly ramp up to maximum generation. Since peaking units are very flexible and have only extrinsic value, the ideal way to value them is to use Monte Carlo techniques.

7. Summary

The question of how to best value generation assets is an important issue for many participants involved in power trading such as energy merchants, utilities, banks and other investment companies. Not only are the right techniques important for the accurate valuation of a generation-asset portfolio, they also enhance the ability to properly manage and hedge the risks associated with the assets.

In this paper, we firstly demonstrated a quick and simple analytic approach that can be used to value generation assets. Typically this approach is insufficient for most real generation assets because it does not account for path-dependent dispatch decisions and also uses a price process that does not accurately represent real world power-price dynamics. Our second, Monte Carlo approach, is much better able to capture the real world price dynamics of both the fuel and power prices, and is more suited to incorporation of the physical assets associated with the assets.

8. Lacima's expertise in generation asset valuation

Whether contemplating the valuation of a virtual power plant, purchase or sale of a generation asset, upgrading a facility, or needing to make optimal operational decisions, Lacima's software incorporates the latest valuation methods to help you to value and optimise generation asset portfolios. Available as a module within our energy risk management, valuation and optimisation system – "Lacima Analytics", Lacima's generation assets solution helps you to:

- Value portfolios of power contracts and generation assets at once with a holistic view of optimal dispatch strategies

- Choose from a range of generation dispatch algorithms including ramping, price based, energy limited, wind and multi-unit generation

- Capture wide ranging constraints such as:
    - Capacity
    - Contracted minimum run
    - Minimum stable generation
    - VOM costs
    - Start-up cost
    - Ramp-up/ramp-down rates
    - Minimum up/down time
    - Fixed costs

- Heat rates
- Fuel switching capability
- Tax rates
- Emissions costs
- Outages

- Obtain a comprehensive range of results/outputs including:
  - $ value
  - Maximum/total MWh
  - Average peak/off-peak forward
  - Average off-peak forward
  - Face value
  - Distributions of profit/loss, earnings, and costs

- Integrates effectively with ETRM and other operational systems

9. About Lacima Group

Lacima Group is a specialist provider of software and advisory services dedicated to valuation, optimisation and risk management for global energy markets. We help you to maximise your profit potential and make more informed decisions by providing tools that yield more accurate valuations, hedging analysis and risk exposure analysis for portfolios of financial contracts and physical assets.

Clients of our software and services include structuring, valuation and risk teams in vertically integrated energy companies, energy retailers, financial institutions and large energy consumers in Europe, North America and Australasia.

Our software solutions have been developed and implemented by peer-recognised experts in energy analytics, offering an unparalleled level of expertise and personalised support.

Contact us:
London: +44 (0)20 7036 0360
Houston: +1 713 353 3949
Sydney: +61 (2) 9275 8818
Email: info@lacimagroup.com
Website: www.lacimagroup.com