The Value of Starting-Up The Power Plant

Avoiding perfect foresight with Least-Squares Monte Carlo: In this article we demonstrate the impact of various start-stop constraints and costs. This impact analysis is possible by applying advanced techniques for generating realistic Monte Carlo price simulations in combination with techniques for optimising the production pattern.

By Cyriel de Jong, Dirk van Abbema, Henk Sjoerd Los & Hans van Dijken

Gas-fired power plants provide the primary source of production flexibility in many power markets. An economically optimal use of the start-stop flexibility of gas plants is paramount to retrieving the maximum value from the asset. With the increasing penetration of wind power, this flexibility will become essential to balance the system. While starts and stops allow the owner to choose the production hours with the largest margin, they are also associated with various explicit and implicit costs.

An important insight that we gain is that different ways to limit starts lead to subtle differences in the actual use of the power plant and the corresponding value. We also find that the common modelling assumption of having perfect foresight about the future spark spreads may lead to a significant overstatement of plant value. This latter result contrasts our previous belief [Los *et al*, 2009], and statements of some other researchers [see Clewlow *et al*, 2009] who claim that perfect foresight is a reasonable assumption. In particular, when there is a fixed limit to the number of allowed starts, as is common in many Virtual Power Plant (VPP) contracts, uncertainty about future margins is definitely reducing plant value. We are able to show this result using the concept of Least-Squares Monte Carlo as applied to energy assets in e.g. Deng (2006, power plants) and De Jong and Boogert (2008; gas storage).

Building Blocks For Plant Valuation

The simplest way to assess the value of a power plant or a VPP deal is to discount the forward spark spread back to today, assuming production is shut down when forward spark spreads are negative. This means the plant is treated as a strip of European-style call options on the spark spread. This approach ignores the operational costs and constraints that tie production hours together, so overestimates true plant value. On the other hand, it also underestimates true plant value, because the variability in commodity prices generally leads to considerable real option value for flexible assets.

In fact, accurate valuation of thermal plants assets requires three fundamental elements:

1. A realistic model to describe how prices evolve over time; the model should be able to generate realistic Monte Carlo simulations.
2. A powerful methodology to find the optimal production pattern for the different price scenarios, incorporating all relevant costs and constraints.
3. A framework to analyse different trading and hedging strategies.

In a previous article in *WorldPower* (Los *et al*, 2009), we provided a description of all three building blocks. In particular, we highlighted the concept of cointegration: It links power prices to the fundamentals of the market (merit order) and thereby keeps spark spreads within reasonable bounds, while maintaining the stochastic nature of prices. In a case study for a 3-year VPP deal we found that the flexibility (extrinsic) value equaled 40% of intrinsic value if we included a realistic degree of cointegration. Without cointegration, the extrinsic value became far too large though.

Start Limitations

Last year's article[1] also showed the impact of various operational costs and constraints, ranging from minimum runtimes and start costs to maintenance and plant degradation. In this article, we introduce a new constraint, namely a hard limit to the number of times a plant operator may start in a year. We show how this compares to other start constraints and demonstrate that perfect foresight about future prices leads to over-optimistic assessments.

> ... we introduce a new constraint, namely a hard limit to the number of times a plant operator may start in a year

Some start costs are very clearly defined, such as the purchase of start fuel, which is the fuel consumed while firing-up to the minimum stable generation level. Other start costs are included to account for the detrimental effect on the condition of a unit caused by stopping and restarting the unit. A challenging aspect of start costs is that they may depend on the temperature of the unit, e.g. cold, warm, or hot. The temperature in turn depends on how long the unit has been offline. A plant operator will need to know for how long they have been offline, in order to make the correct cost calculation. In the solution framework of dynamic programming this creates additional states that the model needs to keep track of. For example, when a hot start is possible after 4 hours, a warm start after 8 hours and a cold start after 24 hours, we have 3 different start cost structures associated with 24 states.

An even more challenging operational constraint is a hard bound limitation to the number of starts, or alternatively, that start costs increase with the number of starts made before. This can be used to avoid an otherwise quite destructive number of starts, e.g. every working day. As a consequence of such a hard bound constraint, the owner of a plant will try to minimize the number of startups. In any case, the plant operator needs to keep track of two quite different statistics:

1. How many hours is the plant already offline, or how many hours is it already producing? This is a statistic ('state' in the dynamic programming language) which looks at the recent history of actions.

2. How many times did the plant start already this year? This is a statistic of a different nature, looking back until the beginning of the year.

Whatever the model implementation is, an optimisation tries to search for the optimal dispatch taking into account the current and future spark spread levels, the current 'short-history' state, and the current 'long-history state'.

Solution Approaches to Handle Start Constraints

In this paragraph we will demonstrate three different solution approaches, each using backward valuation dynamic programming. We show the pros and cons of each, but before doing so, a fourth alternative is worth mentioning. Mixed-integer linear programming (MILP) is a common alternative in which the objective (profit maximization) and constraints are written down in a set of linear equations. While this approach is relatively easy to implement, it is often a factor 5-20 slower than a dynamic programming approach due to the large number (at least 8,760 per year) of decision variables. Even more, MILP ignores any price uncertainty, which is part of our third approach below. These approaches are:

1. A dynamic programming approach for the 'short-history' constraint, and a 'shadow cost' for the 'long-history' constraint to limit the number of starts. This approach assumes perfect foresight about future price levels.

2. A dynamic programming approach for both types of constraints, meaning we have an additional dimension to keep track of the number of previously made starts. This approach assumes perfect foresight about future price levels.

3. As 2, but assuming no perfect foresight. This means the plant operator has to judge the likely level of spark spreads in the future and take decisions which are inevitably suboptimal in certain situations. We use Least-Squares Monte Carlo to establish this strategy.

The paper of Tseng and Barz (2002) is the primary academic reference for a description of a dynamic programming approach to handle short-history constraints of up/down times. The picture below shows the main logic. The model distinguishes between different states (y-axis) that the plant can reach at different hours (x-axis). In the shown example with 2 hours minimum on-time and 2 hours minimum off-time, if the plant is up ('on') for 2 hours or more it can stay 'on' or go 'off'; if it is off for 1 hour (idem if on for 1 hour), the operator has no choice other than to stay off for one more hour. At each node, the model compares the continuation values (F) for the alternatives it has and chooses the maximum. For example, when at the end of hour t the plant is already off for 2 hours or more, the choice is:

$$\max \begin{cases} F^{on\ 1hr} + \mathit{Margin} - \mathit{Start\ Costs}_{t+1} - \mathit{Shadow\ Costs} & \text{"Plant switched on in hour t+1"} \\ F^{off\ 2hr}_{t+1} & \text{"Plant switched on in hour t+1"} \end{cases}$$

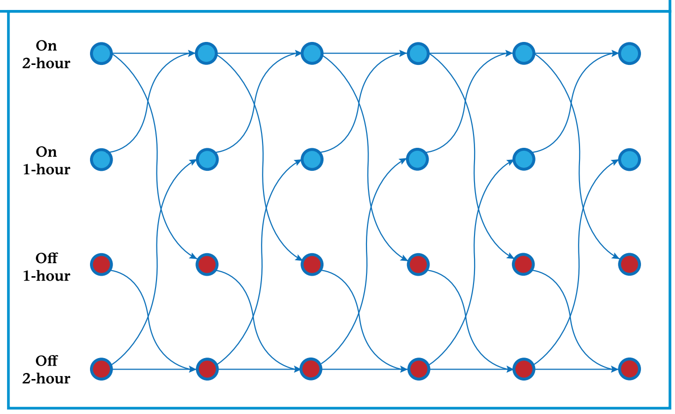

Solution approach 1 means we run through the dynamic program a couple of times. In round 1 we only include the real costs for starting up (Start Costs, e.g. fuel-related) and set the shadow costs for starts to zero (Shadow Costs = 0). If the number of starts is below the maximum, we are done. If the number of starts in the previous round is above the maximum, we modify the shadow cost. We make as many modifications to the shadow cost until the model reaches the maximum plant value within the start limitation per year. In our implementation typically 10-20 rounds are needed before reaching the optimum; the whole procedure generally takes less than 1 second per simulation over a 1-year horizon on an ordinary computer.

Solution approach 2 can be carried out with a single model run, but overall it is generally more time consuming. The reason is that we expand the state-space by a factor of NMax, where NMax equals the number of allowed starts per year: whereas the original dimensions in the previous example are 8,760 x 4, they now become 8,760 x 4 x 50 when, for example, 50 starts are allowed. Multiplying the dimensions gives the number of 'nodes' in the state-space: from each node the model decides which other node in the next hour is optimal. Continuing the previous example, and assuming we already started 20 times before (N=20), the choice is:

$$max \begin{cases} F_{t+1}^{on\ 1hr,\ N=21} + Margin_{t+1} - Start\ Costs_{t+1} - Shadow\ Costs & \text{“Plant switched on in hour t+1”} \\ F_{t+1}^{off\ 2h,\ N=20} & \text{“Plant remains off in hour t+1”} \end{cases}$$

The time-consuming part is that this type of calculation needs to be performed for each N = 0 to 49.

Solution approach 3 is Least-Squares Monte Carlo. The general idea is to estimate the values at the different nodes, instead of assuming perfect knowledge about the true value. Longstaff and Schwartz (2001) made this approach popular and it is a key algorithm in many KYOS applications. In the above example and a maximum of 50 starts, the model carries out 8,760 x 4 x 50 individual linear regressions ('least squares') to estimate the values at the node.

For example, in order to estimate $F^{on\ 1hr\ \ N=21}$ at $(t+1)$ the continuation value it is based on a regression estimate, which may take the following form:

$$F_{t+1}^{on\ 1hr\ \ N=21} = \alpha + \beta_1 \cdot St + \beta_2 \cdot S_t^2 + \gamma_1 \cdot P_t + \gamma_2 \cdot P_t^2$$

This equation is estimated at each node, so in our example we get different parameters for each of the 8,760 x 4 x 50 nodes. The regression is carried out on simulated price levels at time t, which are here denoted by $S(t)$ and $P(t)$. In our implementation we use the current spark spread as $S(t)$ and the average spark spread over the past 24 hours as $P(t)$. Additional price information, such as the forward spark spread, may be incorporated to improve the estimate, but our analysis suggests this hardly improves the performance. However, this whole process takes quite some time, making the model roughly 10 times slower than with approach 2 and roughly 50 times slower than with approach 1. This is the price we pay for avoiding perfect foresight.

Case Description

Developing the above solutions is a considerable effort. The question is whether this pays off in better insights when comparing different investment options. Furthermore, the question is whether the additional calculation time of the Least-Squares Monte Carlo method leads to fundamentally new outcomes.

In order to verify this, we consider a power plant over the year 2011 in the UK market. As a first step, we generate Monte Carlo simulations over 2011, using KYOS' model KySim as we did for our study last year. The model incorporates cointegration and a multi-commodity and multi-factor model for the forward curve movements. It is also based on a regime-switches and cointegration for spot prices. The CCGT plant under consideration has a maximum output of 860 MW at an efficiency of 57.5%.

When spark spreads are negative, the plant can reduce its output to the minimum stable level of 510 MW at an efficiency of 54.5%. It can also turn off the production and restart at a later point in time. In the base case the plant can start and stop without any limitation and without any start costs. Based on a single scenario (the hourly forward curves), the optimal


Siemans Press Picture

Figure 1: Costs Per Plant Start

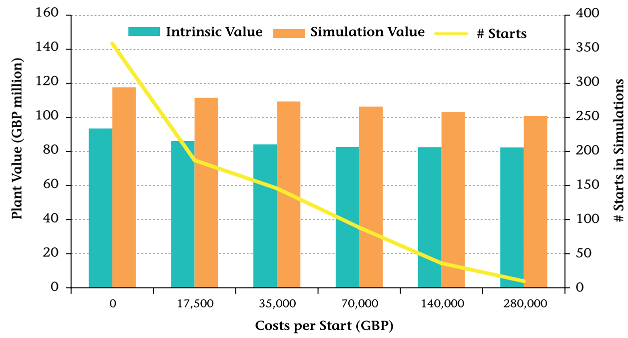

Figure 2:

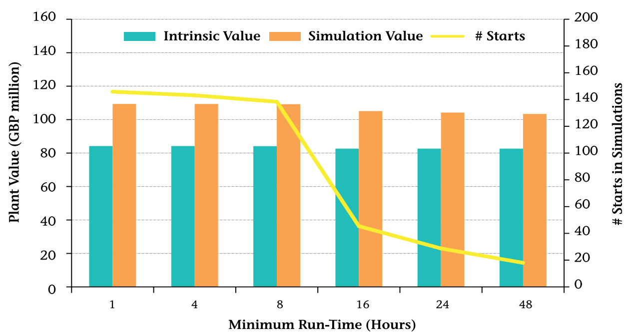

Figure 3:

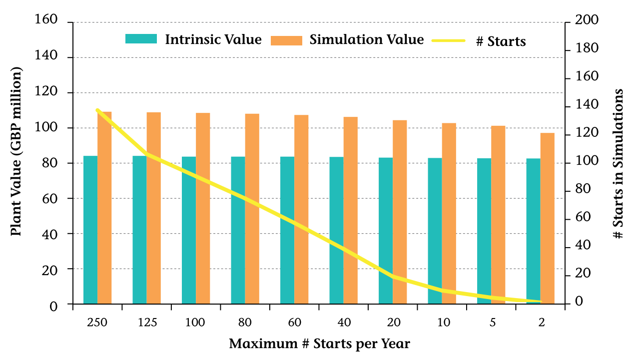

Source: KYOS Energy Consulting

production pattern produces an intrinsic value of GBP 93.5 million. Volatility in the market prices create extrinsic value on top, leading to a total average value across the simulations of GBP 117.7 million.

When Starts Are Costly or Limited

There are no CCGT plants that can actually run at any chosen hour. A start requires additional fuel consumption and reduces maintenance cycles. Restarting may only be possible after a few cool-down hours, but too many cooling hours may also require additional fuel to restart. We distinguish a hot start (within 4 hours), a warm start (within 16 hours) and a cold start (after more than 16 hours off). Hot and warm starts are cheaper than a cold start by a factor of 3/7 and 5/7 respectively.

Increasing the costs for a cold start strongly reduces the number of starts. In the base case scenario the plant made on average about 1 start per day. A relatively low impediment of GBP 17,500 (about GBP 20/MW capacity) for a cold start roughly halves the number of starts to 187 per year (Figure 1). A doubling of start costs reduces the average number of starts further to 145. And with costs of GBP 280,000 only 10 starts per year are made on average. With start costs at such a high level, plant value obviously goes down – from GBP 117.7m to 100.9m.

There are at least two other ways to reduce the start-stop frequency. First, one can impose a minimum number of production hours after which a stop is allowed. Our analysis combines this with a start cost of GBP 35,000. Figure 2 shows that a minimum run-time of up to 8 hours is not a problem, but at 16 hours there is a sudden drop in number of starts and in plant value. Second, one can explicitly impose an upper bound to the number of allowed starts in a year, a common element in VPP contracts. Again, we combine this with a (cold) start cost of GBP 35,000. In Figure 3 it is shown that the plant value with up to 60 allowed starts hardly declines relative to the unconstrained case: the value goes

Siemans Press Picture

Figure 4: Optimal Moments to Start – Valuation

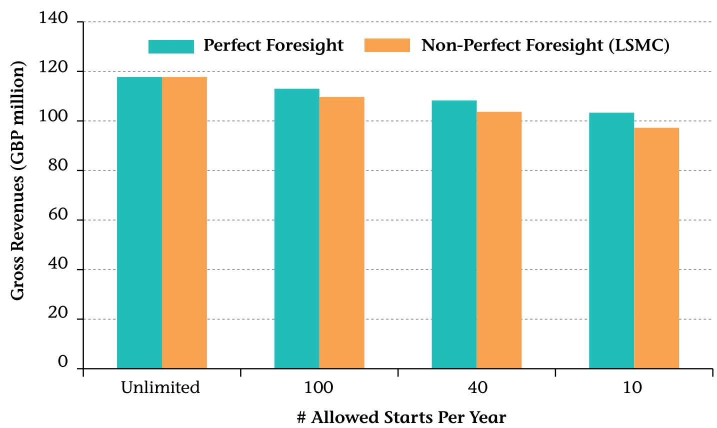

Source: KYOS Energy Consulting

down from GBP 109.2 to 107.4 million. However, with fewer starts the plant is eventually affected with a value going below GBP 100 million for 2 allowed starts.

Non-Perfect Foresight: Least-Squares Monte Carlo

The above results may not be a fair assessment of the value that can be extracted from this power plant. The dispatch optimisation has been carried out using perfect foresight about the price levels in the rest of the year. While operators may be able to have a quite precise assessment of prices in the next 24-48 hours, assessing spread levels and start-stop frequency a couple of weeks or months out, is more challenging.

The Least-Squares Monte Carlo approach provides a methodology to introduce a degree of non-perfect foresight. Its impact is most pronounced when the number of allowed starts per year is the only limiting factor. So, there are no explicit start costs or minimum run-times. Figure 4 shows that there is a decline in value when the optimal moments to start cannot be selected with perfect foresight. The more stringent the start limitation, the more difficult it is to operate optimally, ultimately leading to a value decline of around GBP 6 million if only 10 starts per year are allowed.

The implemented Least-Squares Monte Carlo approach provides an estimate of the plant's continuation value using only spot price information. Alternative specifications, incorporating forward price data, reduces the GBP 6 million slightly, but by never more than GBP 1 million in our test setups.

Conclusion

Gas-fired power plants have generally higher marginal production costs than other production units. Yet, what motivates their investment is the combination of lower investment costs (per MW capacity) and their flexibility to vary production. With a rising penetration of wind power, and especially offshore wind, the price fluctuations in the hourly markets are expected to become more extreme. Nevertheless, the flexibility of gas-fired plants is not unlimited as each start is associated with explicit and implicit costs. In this article we showed how such costs can be introduced in an investment analysis. We also showed that there is considerable real option value in the hourly production flexibility.

> Using LSMC methodology ... the value loss due to non-perfect foresight may amount to 5% of the total plant value

However, this flexibility is likely to be overstated when assuming perfect foresight in the optimal dispatch decision. Using the Least-Squares Monte Carlo methodology, our calculations demonstrate that the value loss due to non-perfect foresight may amount to 5% of the total plant value. This is a fundamentally new outcome and worth the extra modelling effort. ■

By Cyriel de Jong, Dirk van Abbema, Henk Sjoerd Los & Hans van Dijken:
KYOS Energy Consulting and Erasmus University Rotterdam.
KYOS Energy Consulting is an independent consultancy firm offering specialised advice on trading and risk management in energy markets. KYOS advises energy companies, end-users, financial institutions, policy makers and regulators.
www.kyos.com

Footnote:
1. Realistic Power Plant Valuations, *WorldPower 2009*, pp. 48-53. Download from www.commodities-now.com/worldpower

References:

- Boogert, A., and C. de Jong, 2008, "Gas Storage Valuation Using a Monte Carlo Method", *Journal of Derivatives*, Vol. 15, No. 3, pp. 81-98.
- Tseng, C.L., and G. Barz, 2002, "Short-term generation asset valuation: a real options approach." *Operations Research*, Vol. 50, No. 2, pp. 297-310.
- Deng, S.J., and Z. Xia, 2006, "A Real Options Approach for Pricing Electricity Tolling Agreements", *International Journal of Information Technology & Decision Making*, Vol. 5, No. 3, pp. 421-436.
- Los, H.S., C. de Jong and H. van Dijken, 2009, "Realistic power plant valuations: How to use cointegrated power and fuel prices", *WorldPower 2009*, pp. 48-53.
- Clewlow, L. J. Lujun Liu, D. Meador, R. Sobey and C. Strickland, 2009, "Valuing generation assets using Monte Carlo Simulations", *Energy Risk*, September, pp. 56-61.