KYOS white paper

Guidelines for valuation of real options in energy markets

1 Objectives: valuation, optimization, hedging

Energy markets contain a multitude of assets and contracts that offer the owner flexibility. Examples are transportation, gas storage, power production, gas production, heat storage and a range of flexibilities in contracts. Thanks to their flexibility, these assets can be seen as real (non-financial) options.

Companies can be more competitive if they know how to accurately value and manage their real options. They make better investment and trading decisions, capture more cash-flows from the available assets, and know better how to hedge the risks in the market.

In order to achieve these goals, a range of methodologies is available. This document describes the main methodologies and provides a few tests for their quality.

2 Main building blocks

In order to value assets and contracts in energy markets, the following three building blocks are of importance:

1. **Forward curves.** These should be as much as possible based on market prices. However, beyond the tradable horizon and in order to arrive at the required granularity (e.g. hourly for power) a process of forward curve building is required.
2. **Price dynamics.** Price dynamics can be introduced in a couple of ways. In energy markets there is generally the choice between (a) closed form formulas that require e.g. volatility as input, (b) approaches that model the price development as a tree, (c) Monte Carlo simulations. Monte Carlo simulations are the most flexible and general approach; the simulations can be based on any formulation of the price process, with multiple commodities and a complete specification of the whole forward curve (with spot, months-ahead, quarters-ahead, etc). The other two approaches are only possible with relatively simple formulations of price developments and generally assume there is only 1 relevant price.
3. **Optimization.** To capture the full value of an asset in the market, a trader or dispatcher will choose the best moments to operate. He will do so based on the latest information of spot and forward prices. Optimization software can provide support in making these decisions. Optimization software helps to find the optimal decisions not only today, but also for different

sets of possible future prices. This then leads to potential future cash-flows, the average of which (after discounting, and possibly after considering hedging) constitutes the fair value of the asset.

Each of the above three components is important. Small shifts in forward curves can cause large changes in value, extreme price scenarios or spread scenarios easily lead to an overvaluation of flexibility, and suboptimal operational decisions lead to a loss of value.

3 Optimization

The optimal use ("exercise") of an asset can be determined either with or without perfect foresight. Perfect foresight means that it is known in advance how prices will evolve in the future. If the valuation is based on simulations or scenarios, this means that in each scenario the optimization knows all future price levels. When valuing gas storage, whose value depends on price differences between periods, this clearly leads to an overvaluation. When valuing a power station, whose value mainly depends on the current spreads in the market, the degree of overvaluation may be very limited, e.g. 1-3% for an in-the-money plant and around 10% for a far out-of-the-money plant.

3.1 Optimization with perfect foresight

The perfect foresight optimization may lead to some overvaluation, but is also a bit easier to implement and understand. There are basically two methodologies that assume **perfect foresight of future price levels**:

1. **Mixed Integer Linear Programming (MILP).** This is a very popular approach in many fields. The optimization problem can be defined very flexibly and there are very powerful "solvers" available. The main disadvantage is that even the most powerful solvers (e.g. Gurobi, Cplex) need a long calculation time to find the (close-to) optimal solution for typical power plant optimization problems. Consequently, calculations may take several hours, or the problem needs to be simplified.

   Inputs:
   - 8760 hours
   - maximum of 100 starts per year
   - start cost of 50

   Definitions:
   - X1(t = 1 to 8760): 1 if produce at time t, 0 otherwise;
   - X2(t = 1 to 8760): 1 if start, 0 otherwise
   - Spread(t = 1 to 8760): margin when producing at time t

   Constraints:
   - Sum{X2(t)} <= 100
   - X1(t) – X1(t-1) – X2(t) <= 0
     When you run at time t, but not at t-1, then time t should be labelled a start. So, when X1(t) – X1(t-1) = 1, then X2(t) should be 1 too .

   Max revenues:
   - Sum{Spread(t) * X1(t) – 50 * X2(t)}

   Figure: example of MILP

2. **Dynamic Programming (DP).** This methodology breaks up the problem into various sub-problems. It starts on the last day or hour and decides what is the optimal decision. It then works backwards to find the optimal previous decision, knowing what is optimal further on. This methodology requires the definition of all possible "states" that the asset can be in. The main

challenge is to keep the number of states within reasonable bounds, though up to 10,000 or 20,000 states may still be feasible to solve. The main advantage of DP is that it can be many times faster with a greater level of detail than MILP.

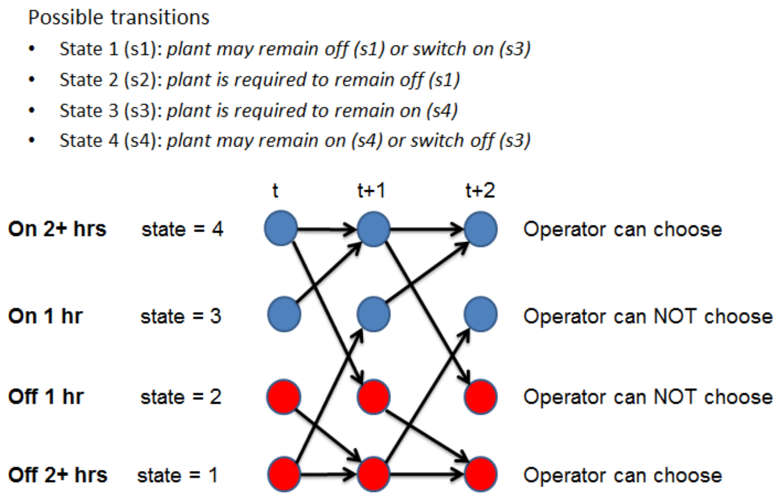

Figure: Example of dynamic programming for a power plant optimization. The plant has a minimum on-time of 2 hours and a minimum off-time of 2 hours, leading to a total of 4 states.

There are two main methodologies based on **non-perfect foresight of future price levels**:

1. **Tree approaches, including stochastic dynamic programming**. In these methodologies the actual price development is simplified. Prices can reach specific levels, denoted by nodes in the tree. Furthermore, there are specific probabilities determining how likely it is that the price evolves from one node at time t to another node at time t+1. The most popular specification is a recombining binomial tree. It assumes that prices can always go one movement up or one movement down, roughly with equal probability. With a binomial tree the number of nodes grows over time. In the SDP and SDDP methodologies the number of nodes is generally fixed, so from the most upper node the prices cannot go up further.
2. **Least-Squares Monte Carlo (LSMC).** This methodology has become popular since the publication by Longstaff and Schwartz (2001). It is very similar to dynamic programming (DP). However, the values for each state are not known with certainty as in DP. Instead, the values are estimated using a least-squares regression on the prices that are known at that time. This ensures that only the current day's information is used, and hence perfect foresight is avoided.

The LSMC method uses Monte Carlo simulations directly. The tree can be derived from Monte Carlo simulations as well, or directly from the price formulation (e.g. volatilities). The advantage of trees is that it can lead to very quick calculations. The real disadvantage, however, is the simplification of the price dynamics. For example, prices cannot only move up or down by a fixed amount. In particular, this becomes problematic when the value of the asset is not determined by a single price, but by multiple prices: multiple commodities and possibly multiple forward products. For example, for the valuation of a power station, a bundle point is a combination of a power price, fuel price and CO_2 price. The number of bundles is typically between 10 and maximum 40, which make it very hard to capture realistic dynamics of all relevant commodity prices and spreads.

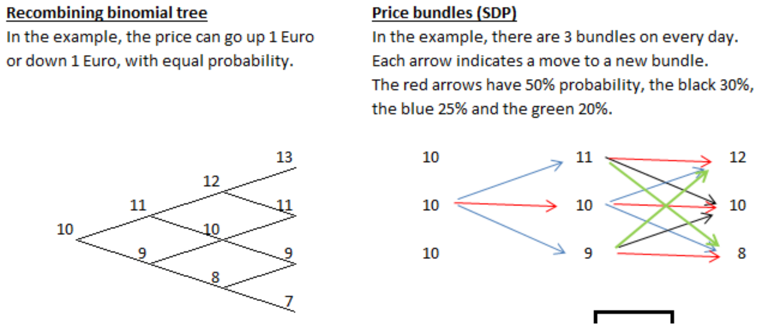

Monte Carlo scenarios
In the example, the next day's price is the previous day's price plus a random return.
The distribution can be any (e.g. Normal).

Sim 1:	10.00	9.59	8.85	9.47
Sim 2:	10.00	10.48	10.81	11.49
Sim 3:	10.00	9.42	9.88	10.23
Sim 4:	10.00	8.92	8.08	7.96
Sim 5:	10.00	11.59	11.83	11.67
Sim 6:	10.00	9.73	8.91	8.72

Figure: Simple examples of a biniomial tree, a price bundle (different type of tree) and Monte Carlo scenarios. In all cases, the user can change assumptions of probabilities, volatilities, price moves, etc.

4 Price dynamics

The forward curve can be interpreted as the average expectation of the market of future price levels. In reality the prices will be different though. A good price formulation considers both the dynamics in the forward market and the spot market.

In order to judge the quality of the price formulation, don't simply trust the mathematics. Instead, it is strongly advised to look at the scenarios, calculate a range of statistics and compare these with the historical statistics (mean, variance, correlation, etc).

4.1 Transform trees and bundles into scenarios

For a fair comparison, the trees and bundles should be transformed into scenarios. Taking the price bundles as an example, this can be achieved as follows:

1. For all scenarios time 1 price is 10
2. For time 2, draw a random number between 0 and 1 (uniform distribution, Excel function rand())
    a. If it is above 0.75, go up to 11
    b. If it is between 0.25 and 0.75 stay at 10
    c. If it is below 0.25 go down to 9
3. For time 3, draw another random number between 0 and 1 and apply the same logic as for day 3
4. This leads to a first scenario. Repeat this e.g. 1,000 times to get 1,000 scenarios

Importantly, even if Monte Carlo simulations are used as input to create a tree, the actual valuation is based on the simplified tree. Therefore, the original Monte Carlo simulations should not be tested, but instead the price scenarios resulting from the above process.

4.2 Recommended test statistics for power plant valuation

We assume there are hourly and daily scenarios for power, fuel prices and CO_2. The below suggested statistics can be regarded as a minimum and can always be extended with further analysis. The calculated statistics should be compared with the historical statistics, except for the absolute price levels. There will never be a 100% match between e.g. historical and simulated volatility, but the average of the simulations shouldn't be too much different.

All prices and relevant spreads (e.g. the spark spread)

a) Calculate the mean, the 5% lowest, 95% highest and possibly extra percentiles for:
   i) Each hour
   ii) Each day (average of the 24 hours)
   iii) Each month (average of all prices in that month)
b) Calculate the correlation of i, ii and iii between the different time periods. In particular, the correlation of average prices and spreads per month, is a useful statistic to judge the whether the scenarios are realistic or not.

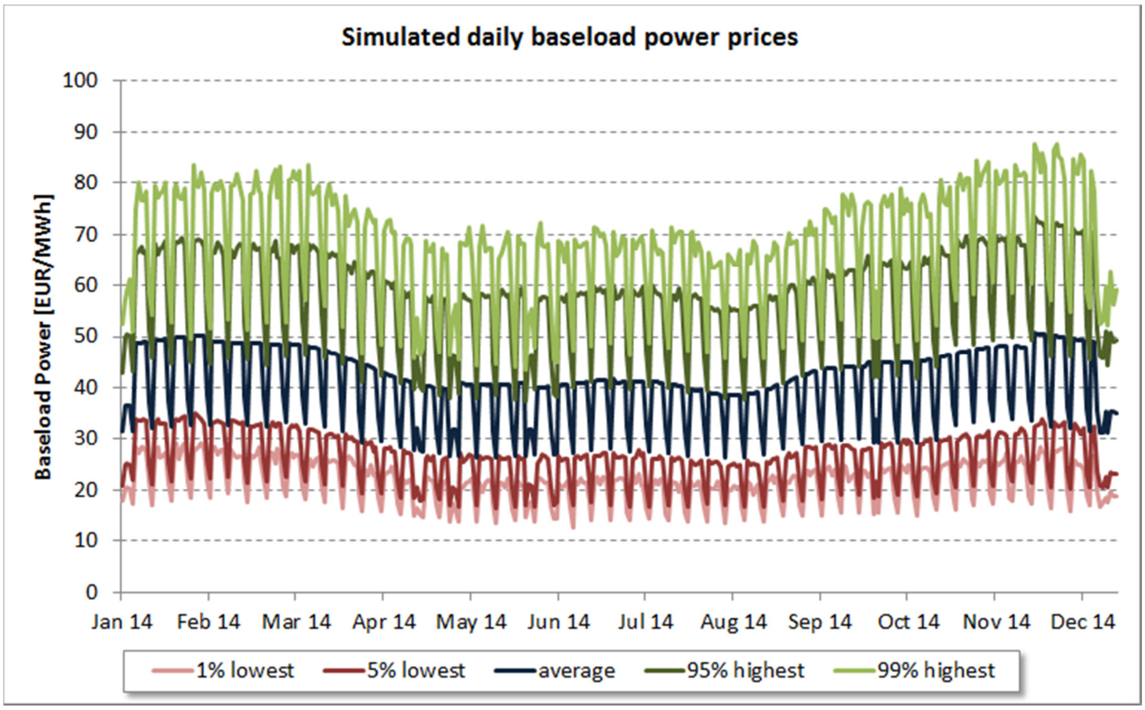

Figure: example analysis of the daily average price distribution for power baseload

Percentage price returns

c) Calculate the standard deviation of the percentage price changes (= returns) per commodity for:
   i) Each hour, so the return from hour 1 to hour 2
   ii) Each day
   iii) Each month
d) Calculate the correlation of i, ii and iii between the different commodities. For example, calculate the correlation between power and gas.

Absolute spread changes

Because spreads can be negative and close to 0, percentage changes are not good for the evaluation of spreads. Instead, the absolute price change is generally more insightful to analyze.

e) Calculate the standard deviation of the spread changes for:
   i) Each hour, so the return from hour 1 to hour 2
   ii) Each day
   iii) Each month