Realistic Power Plant Valuations

How to Use Cointegrated Power & Fuel Prices.

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

The large investments in new power generation assets illustrate the need for proper financial plant evaluations. Traditional net present value (NPV) analysis disregards the flexibility to adjust production decisions to market developments, and thus underestimate true plant value. On the other hand, methods treating power plants as a series of spread options ignore technical and contractual restrictions, and thus overestimate true plant value. In this article we demonstrate the use of cointegration to incorporate market fundamentals and calculate dynamic, yet reasonable, spread levels and power plant values.

A practical case study demonstrates how various technical and market constraints impact plant value. It also demonstrates that plant value may contain considerable option value, but 64% less than with the usual real option approaches. We conclude with an analysis of static and dynamic hedges affecting risk and return profiles.

Power Plant Investment Needs

The combination of rising electricity demand with an ageing production park requires continuous investments in new production capacity. Although countries worldwide have ambitious targets for green energy consumption, fossil fired power plants will continue to play a key role in the coming years. RWE estimates that only 400,000 MW of existing capacity has to be renewed in Europe, 170,000 MW of which is from fossil-fired power plants. In the 2008 edition of WorldPower the authors investigated investments in wind production (De Jong and Van Dijken, 2008). Whereas investments in windmills will be massive, coal and gas-fired power plants will remain the backbone of the world's electricity production for the coming decades and are the subject of this article. This does not necessarily violate green energy targets, considering the possibilities of replacing fossil fuels with biofuels and possibilities of carbon capture and storage.

> ... coal & gas-fired power plants will remain the backbone of the world's electricity production

The need for investments may be clear, but each individual investment has to be justified before it can actually be made. If we assume that the price for a new gas plant equals around €700 per kW, it's easy to calculate that a 420 MW gas-fired plant costs almost €300 million. Investments in coal-fired plants easily involve a multiple of this number and these investments have to be earned back over a plant's lifetime.

The difficulty with estimating future income is the uncertainty about price levels combined with uncertainty about asset behaviour. Will prices remain at this level? Will the availability of the power plant meet expectations? Without doubt, today's expectations about future prices and plant performance will prove to be wrong. Therefore, it is essential to have a clear picture about potential price scenarios, likely plant behaviour and future hedging strategies. This combination of variables provides a range of outcomes, which gives valuable insight in the total value distribution and the optimal dispatch and hedging strategy to follow.

In this article we describe how to overcome the most common pitfalls in power plant valuation. We explain how a realistic Monte Carlo price simulation framework can be built in line with a market's merit order, using a cointegration approach. We also show how plant characteristics can be incorporated into this framework. This approach is especially relevant for assets that are relatively flexible and located in the back of the supply stack. We will demonstrate that the extrinsic value or the flexibility value for low efficient (gas) plants is relatively high but lower than traditional simulation and option pricing methods may suggest. Finally, we clarify the impact of asset-backed trading strategies on actual cash-flows.

Intrinsic Valuation

The gross margin of a power plant is determined by the difference between the power price and the production costs, consisting of costs for fuel, CO_2 emissions and variable operating costs. This margin is commonly denoted as (clean) spark spread for gas-fired units and (clean) dark spread for coal-fired units. Depending on the plant efficiency, the amount of fuel required to produce 1 MWh of electricity varies. A new Combined Cycle Gas Turbine (CCGT) with a 58%[1] efficiency requires 1.7 MWh of gas, whereas an older unit with a 50% efficiency requires 2 MWh of gas. We will refer to all spreads as 'spark spreads', not implying the discussion is limited to gas.

A traditional approach to determining plant value is to calculate the future spark spread levels and multiply this with a load factor of say 2,500 hours off-peak and 2,500 hours peakload. An NPV is obtained by discounting back all spark spreads to today, while deducting all cost components and the initial investment. This approach is often combined with a scenario analysis, where prices are assumed to be relatively high or low over the complete evaluation period.

As a first improvement, more detailed forward curves for the relevant commodities should be constructed. Initially, the curves typically have a monthly granularity. Further out in time, the curve inevitably involves some (solid) guesswork. The monthly forward curves for the peak and offpeak spark spreads form the basis for the expected operation and the intrinsic valuation of a plant. Refining the power and gas curves with daily and hourly profiles improves the valuation further. Ultimately, the largest part of the power plant's capacity will be dispatched on an hourly basis. Consequently, hourly price curves are required to make the dispatch decision.

Figure 1: Comparison Spark Spreads

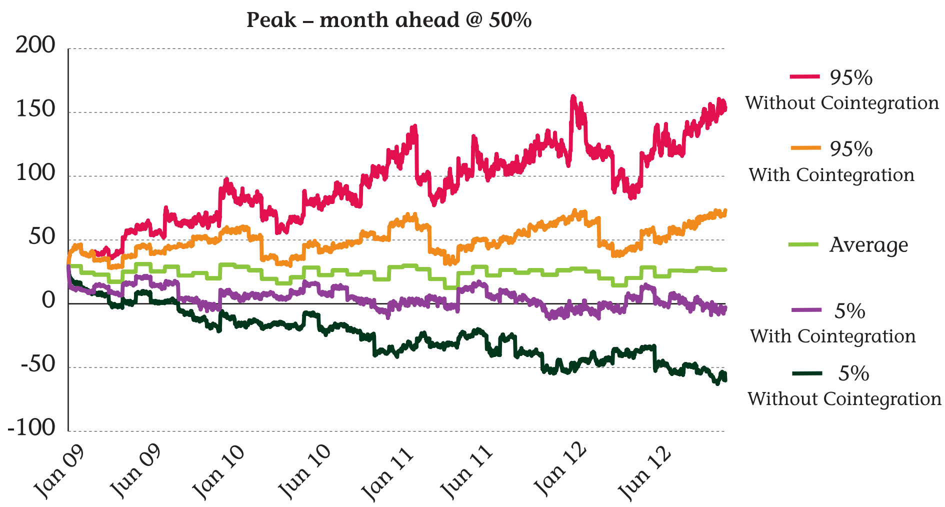

Source: KYOS Energy Consulting

Price Uncertainty & Real Option Valuation

The hourly and daily forward curves may be treated as the best forecast of future spot price levels (if we leave aside risk premia). However, actual spot price levels will surely be different. On the one hand, this creates a risk, which may be reflected in a high discount rate. On the other hand, price variations offer opportunities for extra margin if the plant's dispatch and trading decisions can respond to them. To capture this uncertainty, it is not sufficient to create high/medium/low price or spread scenarios. Actual market dynamics are far more diverse than that. For example, a period of low margins may be followed by a period of high margins in the same day, week, month or year. A plant operator will respond by reducing the production in the low spread period to minimise losses. At the same time, they will maximise production in the high spread periods. In fact, a flexible plant offers the ability to limit the downside and take full advantage of the upside. This is the basis for any real option approach and is actually the way plant owners make a large part of their asset-backed trading profits in the market place.

Models based on correlated returns lead to unrealistic spreads

Still, to many in the power industry, this seems a non-real financial trick. Indeed, such an approach is sensitive to 'model error' or 'analyst bias'. It easily leads to an over-estimation of true plant value. First, approaches which treat the plant as a strip of spark spread call options ignore the real-life restrictions on plant flexibility; restrictions may have either a technical or contractual nature. Second, approaches which are directly or indirectly based on unrealistic spark spread levels suffer from the same overestimation bias.

Correlated Returns: Unrealistic Spreads

To capture the dynamics between commodities over time, analysts rely on Monte Carlo price simulations. This covers a wide range of model implementations and we will demonstrate that the usual approaches exaggerate actual variations in spark spread levels.

The most common approach to combine multiple commodities in a Monte Carlo simulation model is applying a correlation matrix between the different commodities. This includes Principal Component Analysis (PCA). A correlation matrix captures the degree to which prices move together from one day to the next; it is derived from daily (or weekly) price returns. A correlation matrix, in combination with market volatilities, describes actual price behaviour quite well for relatively short horizons, for example in Value-at-Risk models. However, extensive research and practical experience lead to the insight that a correlation matrix is too weak to maintain the fundamental relationships between commodities over a longer period. Consequently, very large or negative spark spreads will be the result. These extreme scenarios are not possible in reality though, as they would mean that either no power plant makes money or all power plants make huge amounts of money. So, whereas an intrinsic valuation disregards the value of plant flexibility, the usual Monte Carlo simulation approach of correlated returns results in an overestimation of plant flexibility.

Another approach is not simulating the individual commodities, but simulating the spark spread directly. There are clear benefits to this approach. The spread can fluctuate between certain 'logical' boundaries, with the result that the (undesired) extreme outcomes are avoided. However, information is lost about the movement of underlying power and fuels prices, e.g. relevant for hedging decisions. This will lead to various practical problems, for instance when combining a dedicated gas contract to the power plant.

In short, we feel that the most common approaches are inadequate solutions when they are applied to power plant evaluation projects. The alternative is the explicit incorporation of fundamental price relationships. This approach has the benefit that spark spreads remain at logical levels, but that information about underlying prices is not lost.

Power prices are the result of the movement in underlying fuel and carbon prices. This relationship can be captured with cointegration: power prices are cointegrated with price movement of fuels (mainly coal and gas) and carbon. With cointegration, power prices are fundamentally driven by dynamic market marginal costs in peak and off-peak and will react properly when commodities are substituted (for instance, change from coal to natural gas in summer periods). Actual commodity prices may temporarily deviate from the fundamental relationships, but not for too long and not by too much.

Cointegrated Forward & Spot Price Simulations

KYOS started the development of a proprietary price simulation model for energy commodities several years ago, part of which has been published in the literature (see e.g. De Jong, 2007). It is now in use by several leading commodity trading companies. Fuel and CO_2 prices are simulated first, with power prices following. The model captures the many shapes that forward curves display over their lifetime. They may, for example, turn from contango (future price higher than today) in backwardation (future prices lower than today). To capture these dynamics we use a multi-factor model whose parameters can be accurately estimated on the basis of a limited set of historical data. The major parameters capture general level shifts, shifts from contango into backwardation and shifts in the size of the winter-summer spread (for power and gas). The volatilities and correlations of the different maturities along the curve can be calibrated to properly match the historical price data, both between different maturities and between different commodities. This is especially important when hedging strategies are evaluated. The model also contains spiky ('regime-switching') power and gas spot prices, mean-reverting to forward price levels, with appropriate random hourly profiles.

The model as described produces realistic price simulations for individual commodities. At first sight, it also nicely ties commodities together through correlations. Still, we experienced that it does not produce realistic spreads between commodities, whether it be oil-gas spreads, regional gas spreads or power-fuel spreads. Yet spreads are actually the most important input to most valuations, including power plant valuations.

We solved the issue through cointegration, a Nobel Prize winning econometric innovation (Engle and Granger, 1987). For spark and dark spreads it is complemented with the explicit incorporation of the merit order. Essentially, the cointegration approach captures the correlation between price *levels* rather than (only) price *returns*. Intuitively, it uses a regression to find the 'stable' relationship between commodity prices and then assumes that 'actual' commodity prices move around this stable level. The concept is very similar to a spot price mean-reverting around a forward price level. The primary challenge is to align the approach with the return-driven movements of the forward curve, something we managed to resolve over time.

Case Study: In order to bring this theoretical explanation to a practical level, we next consider a case study involving a power plant over a three year period.

Case Study

We consider a new gas-fired power plant in Germany. With a 58.5% efficiency the plant produces a maximum of 420 MW; at the minimum stable level it produces 170 MW (47% efficiency). The plant has fixed annual operation and maintenance (O&M) costs of €6.3 million. We disregard discounting for simplicity.

If a plant is dispatched economically, it produces when its spark spread – the gross marginal revenue – is positive and does not produce when the spark spread is negative. Simple as it seems, technical, contractual and market restrictions hinder plant owners aiming to exactly dispatch along this principle. Actual dispatching is an optimization challenge, involving issues such as ramp rates, minimum run-times, plant trips, maintenance and production-dependent heat rates. Optimal dispatch decisions can be derived with various mathematical techniques. KYOS generally works with dynamic programming techniques.

Case Study Results

We evaluate the plant over the period 2010-2012 based on forward prices at the end of March 2009.

- ***Traditional approach, no constraints***

With the traditional approach, power is constantly produced during 2,500 peak and 2,500 offpeak hours. Taking fixed cost components of €6.3 million/year into consideration, this leads to an average annual value of €20.4 million.

- ***Monthly intrinsic valuation, no constraints***

A more detailed monthly curve shows that the winter periods have the highest spark spreads, where the high power forward prices compensate for the also high gas prices. In the 36 months, the plant produces only peakload, with an average spark spread of €31.70/MWh. This generates an annual value of €35.2 million. If the company could trade all monthly periods individually, this would ensure a minimum value the company can lock in on the forward market.

- ***Hourly intrinsic valuation, no constraints***

In most of the months, the expected hourly spark spread is negative in some hours, but positive in other hours. Assuming the plant has maximum ramping flexibility and is fully traded on the spot market, the average expected value totals €43.0 million. This is more than the monthly intrinsic value, because of the larger expected variations in the spot than in the forward market. However, prices will not follow the current curve for sure. This creates risk, part of which can be hedged on the forward market, but also additional profit opportunities.

- ***Simulations with cointegrations, no constraints***

Based on our price simulation model we calculate an optimal dispatch schedule per simulation path. This yields a value per simulation, with an average of €53.8 million, but with a large standard deviation of €8.0 million. As we will analyse later, the uncertainty in outcomes may be partially hedged on the forward market, but some risk certainly remains. The €10.7 million difference with the hourly intrinsic

Table 1: Breakdown of Power Plant Value

	Total Value [m€/yr]	Intrinsic Value [m€/yr]	Flexibility Value [m€/yr]	Power [GWh/yr]	Gas [GWh/yr]	Carbon [kton/yr]	Starts [#/yr]	OH [#/yr]
Fixed Price Curves								
2500 peak + 2500 offpeak	20.4	20.4	0.0	2,100	3,621	740	N/A	5,000
Monthly forward curve	35.2	35.2	0.0	1,379	2,377	486	250	3,283
Hourly forward curve	43.0	43.0	0.0	2,059	3,520	720	417	4,902
Price Simulations,								
With Cointegration	53.8	43.0	10.7	2,003	3,423	700	343	4,769
Variable O&M	50.8	40.0	10.8	1,919	3,280	671	345	4,569
Minimum runtimes	44.3	32.6	11.8	2,008	3,488	713	52	5,232
Start costs	42.3	30.2	12.1	2,025	3,545	725	44	5,294
Maintenance	39.6	28.1	11.5	1,907	3,339	683	42	4,987
Plant trips	36.9	25.8	11.0	1,800	3,153	645	44	4,706
Seasonality	37.1	26.0	11.1	1,803	3,159	646	44	4,703
Degradation	36.5	25.2	11.4	1,782	3,148	644	44	4,661
ALL, incl Take-or-Pay	**35.9**	**25.9**	**10.0**	**1,894**	**3,341**	**683**	**43**	**4,935**

Source: KYOS Energy Consulting. Note: the values exclude investment or financing costs; total value = intrinsic + flexibility value

Table 2: The Impact of Cointegration on Power Plant Value

	Total Value [m€/yr]	Flexibility Value [m€/yr]
Without Cointegration	53.9	28.0
With Cointegration	35.9	10.0
Difference	-33%	-64%

Source: KYOS Energy Consulting

is labelled the option value, extrinsic value or flexibility value.

- ***Variable O&M and start costs***

Now we make the case gradually more realistic by adding variable costs. They depend on the number of operating hours (inspections, overhauls) or on the number of starts (extra fuel, extra maintenance). With variable costs per production hour of €1.50/MWh, the plant value reduces by €2.9 million.

- ***Minimum runtimes and start costs***

In practice, there are no fossil-fired plants that are switched on and off from one hour to the next. Actual plant operation is constrained by minimum times to be on or off, which we set at 24 hours each. The impact on plant value is €6.5 million. Taking into account costs per starts of €12,600 plus 2,000 GJ of gas, the plant value reduces further by €2.0 million.

- ***Maintenance and trips***

Planned maintenance is the time required for inspections and planned repairs. For a longer-term analysis it is worth incorporating an inspection scheme with both smaller inspections and major overhauls. Assuming the plant will be in maintenance for 20 days per year, the plant value is reduced by €2.7 million. Unplanned outages (trips) have more effects than simply reducing the generated power production by a single percentage. A trip can occur at the start of a production period, but also at the end where the financial consequences are limited. Furthermore, after a trip, a decision needs to be made if the plant can and should start again. With an outage rate of 6%, in our example the plant value is reduced by €2.8 million.

Figure 2: Build-Up Plant Value

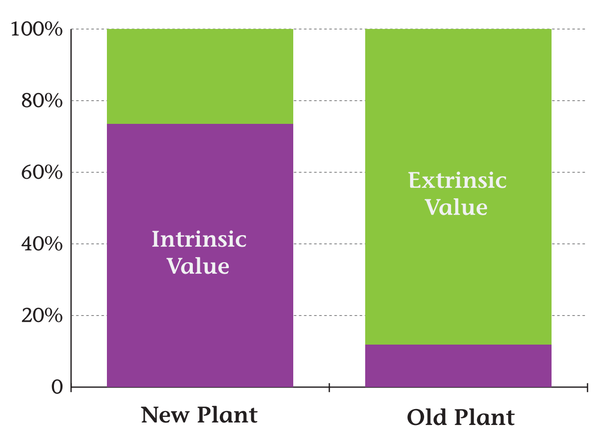

Source: KYOS Energy Consulting

- ***Seasonal effects and plant degradation***

The outside temperature influences the capacity of gas-fired power plants. In the winter, with colder temperatures, more oxygen results in higher capacities than in the summer. The impact of 5% more capacity in favourable periods (winters tend to have larger spark spreads) and 5% less capacity in less favourable periods (summer) leads to a small increase of €0.2 million. During the lifetime a power plant will lose some of its efficiency. Although maintenance reduces the consequences, degradation may be expected especially in the first period after commissioning. An average efficiency of 58% leads to a decrease of €0.5 million.

- ***Contractual: take-or-pay obligation for natural gas***

Besides the physical constraints there can also be contractual limitations to fully exploit the plant flexibility. A dedicated gas contract with a take-or-pay clause restricts the flexibility of the power plant, as the gas cannot be transported elsewhere. In our case, a take-or-pay obligation is translated in a minimum number of operating hours of 5,000 in the first year. As a take-or-pay contract is usually aligned with the expected consumption, the impact is limited to a decrease of €0.7 million.

Besides the described limitations, more constraints could be applied, for example, the ramp rate, although this is more a limitation for coal plants. Also, the delivery of heat could impose must-run obligations for specific plants. Environmental constraints like maximum NOx emissions would also limit the flexibility, similar as for take-or-pay contracts. To highlight the effect of cointegration, a comparison is made with the full simulation model, but the cointegration switched off. The lack of cointegration causes a value increase from €35.9 to €53.9 million. So, cointegration reduces the plant value to 67% of the 'normal' Monte Carlo approach. This reduction is solely attributable to the price scenarios, where spark spreads become more extreme. This becomes an even larger problem when the valuation horizon increases.

Comparing Option Values

The option or flexibility value of a power plant is the difference between the intrinsic value, derived from a static curve (hourly, monthly or something else), and the average value over the simulations. This value is realised by adapting the production profile to changed price scenarios: If spreads turn positive, the plant is switched on. If spreads turn negative, the plant is switched off. With this behaviour profits are added in positive market circumstances, while losses are avoided by stopping the production in negative market circumstances. New-build plants with relatively high efficiency produce in more hours than older, less efficient plants. This impacts the option value: If a plant is already operating, there is the possibility to reduce output or stop producing, while if a plant is not yet running, there is the possibility to switch on.

It is therefore important to realise that the flexibility value is highly dependent on the power plant characteristics and the degree to which the plant is already 'in-the-money' (i.e. profitable to run). For a new plant, the flexibility value is limited compared to the relatively high intrinsic value. However, for an older plant the flexibility value has a larger influence on the total plant value. This is illustrated with comparing our reference plant (58% efficiency) with a 10 year old power plant (54% efficiency). Note that the flexibility value of the plants is relatively high, as result of the chosen forward curves with low spreads.

Figure 3: Distribution P&L

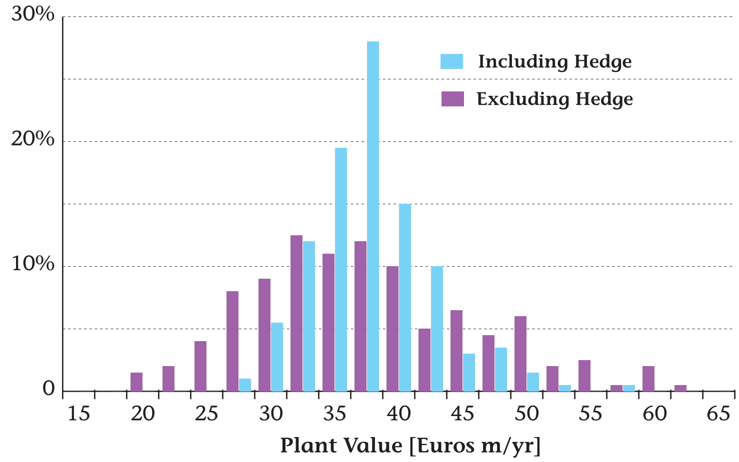

Source: KYOS Energy Consulting

Hedging Strategies

It is common to sell a majority of the expected plant production in the forward market, while at the same time purchasing forward the required fuels and CO_2 credits. This is called hedging. Hedging a power plant serves two main purposes:

1. Risk Reduction. First, with hedging the dependency on price levels of highly volatile spot markets decreases. In relation to this, hedging reduces potential liquidity issues on spot markets.

2. Profit Optimization. Forward spark and dark spreads vary over time. Dynamic trading strategies can increase value by selling more power against high spreads and selling less power (or buying it back) against low spreads.

In this case study the hedge volume is defined as the expected production over the evaluation period, i.e. the average volume over all scenarios. It can be verified that this volume hedge is very close to the concept of a delta hedge. To begin with, the spark spreads are sold forward in March 2009 using calendar forward contracts for delivery in 2010, 2011 and 2012 for peakload power, natural gas and CO_2. We assume no transaction costs. If the hedge is not adapted over the lifetime, this is defined as a static hedge. The result of static hedge is illustrated in Figure 3. Where a spot strategy leads to a wide value distribution, hedging reduces the bandwidth. Scenarios with high spot spreads yield a loss on the hedge, whereas scenarios with low spot spreads yield a profit on the hedge. This dampens the total profit and loss on the spot market and clarifies that hedging reduces the risk profile.

In reality, the expected production volume, which drives our hedge volume, varies with a change in spark spreads. Re-hedging on the basis of this information is called dynamic hedging. Dynamic hedging leads to a further narrowing of the value distribution. And more importantly, a higher profit is expected as more production is sold against higher spark spreads.

Different plant types offer different degrees of flexibility to respond to future price developments

Conclusion

The energy industry is facing important investment decisions, shaping the power production portfolio for the next decades. Different plant types offer different degrees of flexibility to respond to future price developments. An important consideration in the decision process is therefore the accurate assessment of the value to assign to this flexibility. This article demonstrates how the concepts of cointegration and dynamic programming can help to avoid a bias towards either very flexible, yet expensive, or very inflexible power plants. ●

Henk Sjoerd Los, Cyriel de Jong & Hans van Dijken – KYOS Energy Consulting, 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. Lower heating value.

References:

Effective Pricing of Wind Power, De Jong, C. and H. van Dijken, World Power, 2008.

The nature of power spikes: a regime-switch approach, De Jong, C., 2007, Studies in non-linear dynamics and econometrics.

Co-integration and error correction: Representation, estimation and testing, Engle, R. and C. Granger, Econometrica, p. 251-276, 1987.