The UK carbon floor and power plant hedging

Due to the carbon floor, the price of carbon emissions has become a highly significant part of the generation costs for UK power producers. Vytautas Jurenas and Cyriel de Jong describe how the policy affects expected future carbon costs and optimal valuation and hedging decisions, by adjusting Monte Carlo simulations for the UK market

In April 2013, the UK government brought into force a tax on carbon emissions as a part of its Electricity Market Reform (EMR). The legislative target to promote decarbonisation is formulated as a carbon price floor (CPF), ie a minimum level of carbon cost to be achieved in the UK domestic market. In the 2014 Budget, it was announced that the carbon price support will be capped at £18 per tonne of carbon dioxide (tCO_2) until March 2020. This is lower than was planned initially, but is still a significant tax on top of the EUA (European Union Allowance) market prices.

The motivation for this type of measure was the collapse of EUA prices due to the oversupply of permits. Consequently, the European Union Emissions Trading Scheme (EU ETS) was deemed ineffective by the UK to achieve its decarbonisation goals, namely to draw investment into low-carbon generation facilities. The CPF mechanism was introduced around the same time as other mechanisms, such as feed-in tariffs that fix the selling price for power generated from renewable sources.

These mechanisms were considered more effective than the EU ETS mechanism. In particular, the CPF results in carbon costs that are not only higher, but also less volatile and hence easier to predict. At the same time, it should be noted that there is continuous opposition against the mechanism, which can (and has) lead to policy changes.

How to calculate the expected carbon costs?

The top-up rate or tax over the EUA market price in the UK is called the carbon price support (CPS). Every year it is calculated as the difference between the CPF and the average forward prices for EUA contracts during a certain pricing window. The CPS for a particular year is usually announced in a budget two years before the year in question. Hence, the CPS is the actual tax rate charged on fuel supplies with respect to the average carbon content of the fuel. To gain a better intuitive understanding of how the policy works, consider the following example.

Let's say that the CPF target level for 2025 is announced to be equal to £20 per tCO_2 in today's currency. Let's further assume positive inflation expectations, such that in nominal terms this level blows up to £25. The CPS for the fiscal year 2025 – spanning April 2025 to March 2026 – is going to be announced in the budget for 2023, so around March 2023.

The relevant calendar year forward EUA contracts for calculating CPS in this case are CY2025 and CY2026. The pricing window spans from March 2022 to February 2023. Let's say that the average prices for those two contracts in this pricing window turn out to be €9.60 and €9.70 respectively. Let's further assume that the forward EUR/GBP rate for that period turns out to be 0.80 (constant for simplicity). Hence, we obtain the expected EUA price in the fiscal year 2025 of £7.70 and a CPS of £17.30:

$$\text{expected EUA price} = €9.60 \times 0.80 \times \frac{9}{12} + €9.70 \times 0.80 \times \frac{3}{12} = £7.70$$

This far-away example was taken specifically to see the potential problem: the CPS for fiscal year 2025 is not known today (2016). Until we enter the pricing window period relevant for the fiscal year 2025 carbon costs, the best expectation of the carbon costs is probably the CPF, at least if we ignore the policy uncertainties and the possibility that the EUA price could rise above the CPF. During the pricing window, the EUA prices used to set the CPS level are known gradually. During this period of time, the expected carbon costs are a complex mix of the settled EUA prices and the current EUA forward price levels. More details are provided further on. After the pricing period, the expected carbon costs are relatively straightforward again, and equal the sum of the CPS and the current EUA forward (or spot) price. This shows that there are essentially three relevant time windows and the expected carbon costs behave differently in each. This should be taken into account when modelling power prices, when forecasting future power plant production and assessing the value and optimal hedges of power plants. In the following sections we provide further details about the modelling.

How to model the carbon costs?

Price modelling and Monte Carlo simulation

Kyos generates Monte Carlo spot and forward prices with its KySim Monte Carlo simulation model. This time-series model is in use by a wide variety of energy market players and contains a range of parameters, including time-varying volatility, correlation and cointegration. The Monte Carlo price simulations are based primarily on the statistical properties of historical market prices. To simulate power prices, the model takes into account the future merit order in the power market: fuel, foreign exchange and emission prices are simulated first in order to generate realistic power prices. With the introduction of the UK carbon floor, two additional variables are being modelled.

The first is the expected carbon cost, which we denote by C. The second is the carbon price support, CPS. We furthermore assume that some CPS price levels are known, typically for the current and next fiscal year. For simplicity, we assume that the carbon price floor levels (CPF) are known with certainty and that there is no EUR/GBP forex price risk.

1 A single simulation of the expected carbon costs for September 2020

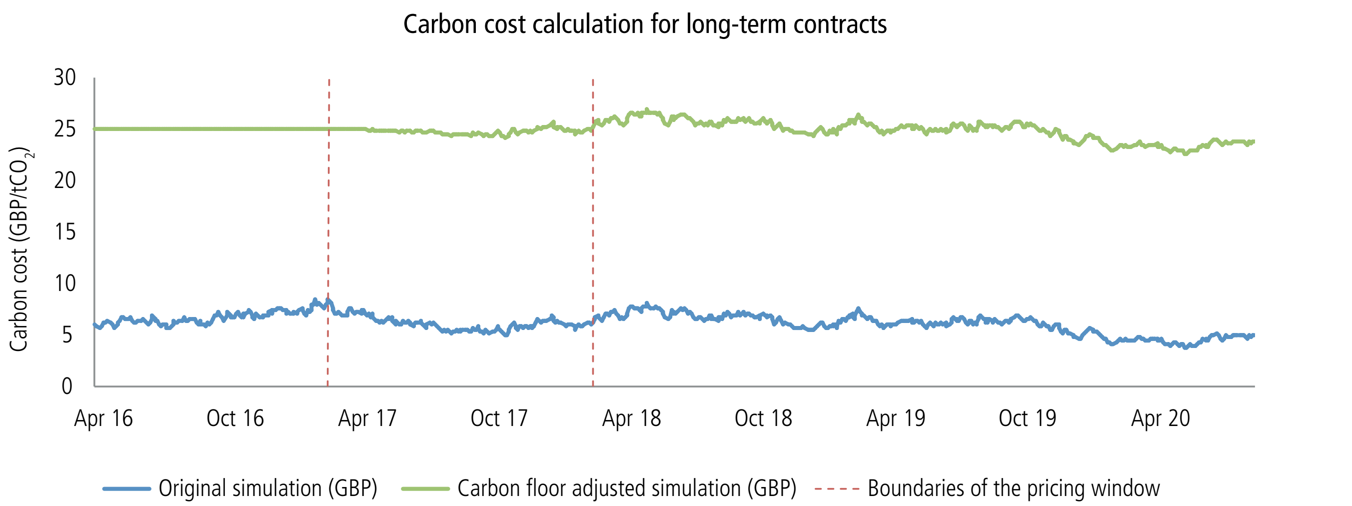

Source: Kyos modelling

In reality, in the KySim model, the EUR/GBP forex rate is included in the Monte Carlo simulations, but for the explanations it makes things unnecessarily complicated.

Although the tax based on the CPS is paid on the fuel used for generating power, for power producers it is easier to treat it as an add-on to the EUA emission price. We define the following variables, where *y* denotes the fiscal year, *t* the trading date and *s* the simulation:

$C_{y,s,t}$ is the expected carbon cost
CPF_y is the carbon price floor
$CPS_{y,s}$ is the carbon price support
$P_{y,s,t}$ is the forward market price of an EUA emission right
$\bar{P}_{y,s,t}$ is the average EUA settlement price over the whole pricing window
$T_{y,1}$ and $T_{y,2}$ are the first and last trading days of the pricing window
$\bar{P}_{y,s,u\leq t}$ is the average EUA settlement price over trading days, from the first day in the pricing period until trading date *t*.
$w_{y,t} = \dfrac{t - T_{y,1} + 1}{T_{y,2} - T_{y,1} + 1}$ is the weight of the prices which have already settled until trading date *t*.

For the modelling of the above variables, three time windows should be distinguished:

- before the pricing window;
- in the pricing window; and
- after the pricing window.

Carbon cost before the pricing window

Let's denote the starting day of the pricing window for CPS_y as $T_{y,1}$. Before this date, the forward carbon costs should be equal to the expected future carbon costs across all scenarios. In most scenarios, the EUA emission prices are likely to stay below the carbon floor price, implying a positive carbon price support. To be precise, we should take into account the probability that the EUA prices rise above the floor. If we do that, then the expected carbon costs are derived in a similar way as a far out-of-the-money call option, where the floor is the strike price. The expected pay-off of this option should be added to the floor price. This means we set *C* for far-away fiscal years equal to the expected value (E) of the maximum of price floor target for that year (CPF_y) and simulated EUA price ($P_{y,s}$):

$$C_{y,s,t} = \mathrm{E}\left[\max\left\{CPF_y; P_{y,s}\right\}\right] \; if\ t < T_{y,1}$$

If we assume that the market expects the EUA price to stay below the carbon floor as long as it is below the carbon floor today, this can be simplified to:

$$C_{y,s,t} = \max\left\{CPF_y; P_{y,s,t}\right\} \; if\ t < T_{y,1}$$

Given the large difference between the carbon floor and the EUA price in today's markets, it is very unlikely that the EUA price rises above the floor in any of the coming years. Furthermore, the whole political intention of the floor is to set a level for carbon costs that is well above the EUA market price level. For this reason, we argue that it is a minor simplification to simplify the full calculation of optionality.

Carbon cost after the pricing window

The carbon price support for fiscal year *y*, CPS_y, is calculated in the period between $T_{y,1}$ and $T_{y,2}$ as the average difference between CPF_y and forward prices for relevant contracts. It is now easier to assume that we have a specific fiscal year EUA contract, while in reality we take the weighted average of two EUA calendar year contracts. Hence, the carbon price support rate is obtained by:

$$CPS_{y,s} = \max\left\{0; CPF_y - \bar{P}_{y,s}\right\}, where\ \bar{P}_{y,s} = \frac{1}{T_{y,2} - T_{y,1} + 1}\Sigma_{t=T_{y,1}}^{T_{y,2}} P_{y,s,t}$$

Here $\bar{P}_{y,s}$ is the average settlement price for this contract in its pricing window. Thus, it follows that the expected carbon costs can be derived by adding the inferred CPS rate to the simulated price of the EUA forward contract:

$$C_{y,s,t} = P_{y,s,t} + CPS_{y,s} \; if\ t \geq T_{y,2}$$

Note that we have a different CPS per simulation.

Carbon cost in the pricing window

The expectation of carbon costs inside the pricing window is slightly different from the previous section. Because we do not have all settlement EUA prices, the actual $\bar{P}_{y,s}$ cannot be calculated. Nevertheless, we can infer it with increasing accuracy, as we are approaching the end of the pricing window.

Let's denote the expected average settlement price at trading date t as $E[\bar{P}_{y,s,t}]$. This expectation comprises two parts:

- the realised forward settlement prices until day t; and
- the expected settlement price after day t, which is equal to the forward settlement price on the day t.

For example, suppose that we are three months into the 12 months of the pricing period and the average realised EUA prices in the first three months are equal to eight, whereas the current forward price level is £12/tCO_2. Then the expected average settlement price equals 3/12x8 + 9/12x12 = £11/tCO_2. With a carbon floor price of £18, the expected CPS is 18-11 = £7/tCO_2 and the expected carbon costs 12+7 = £19/tCO_2.

The expected EUA settlement price equals the weighted average of the settled EUA prices and the current EUA forward price. The expected CPS equals the difference between the carbon floor and the expected EUA settlement price. The expected carbon costs equal the EUA forward price at trading date t plus the expected CPS.

Mathematically, we have:

$$E\left[\bar{P}_{y,s,t}\right] = w_{y,t} \cdot \bar{P}_{y,s,u\le t} + \left(1 - w_{y,t}\right) \cdot P_{y,s,t}$$

$$C_{y,s,t} = P_{y,s,t} + E\left[CPS_{y,s}\right] = P_{y,s,t} + \max\left\{0; CPF_y - E\left[\bar{P}_{y,t,s}\right]\right\} \text{ if } T_{y,1} \le t \le T_{y,2}$$

To sum up, in this case we not only have a different CPS per simulation, but CPS slowly changes depending on the trading day in the pricing window. At the end of the pricing window, all information becomes available, so when $t=T_{y,2}$, we obtain the actual CPS for that simulation.

A carbon cost calculation example in all three time windows can be seen in Figure 1. It displays a single simulation, as of April 2016, for the contract of September 2020. Hence, the pricing window is from March 2017 until February 2018. Assumptions: CPF_{2020} = £25, last EUA price = £6. Until March 2017, the expected carbon cost is equal to CPF_{2020}. The CPS_{2020} obtained in the pricing window equals CPF_{2020} minus average simulated $EUA_{SEP/2020}$ price in the pricing window (25 6.13 = £18.87). This CPS_{2020} is applied in March 2018 and onwards. In the pricing window, on the day t, the expected CPS_t is calculated based on realised settlement prices until this day and forward settlement price at the day t for the remaining days.

Implications for power plant hedging

Power plant hedging using Monte Carlo price simulations

In certain market conditions, a power plant will produce considerably more megawatt hours and at a higher margin than in other market conditions. Due to the timing element in the UK's carbon floor policy, power plant owners cannot assume that their carbon costs in a certain year will be equal to the CPF for that year. Because the CPS is calculated two years before the actual costs are incurred, and EUA prices can change quite a lot in two years, the actual costs can be significantly lower or higher than the CPF. Consequently, the assumption of a fixed carbon cost would result in a biased valuation and suboptimal hedging decisions. Therefore, it is best to repeatedly simulate the whole distribution of expected carbon costs, including the newest information about the UK carbon tax policy.

To stabilise the future income of a power generator, price exposures have to be hedged. The main price exposure is typically towards the power prices, but changing production costs can also have a significant impact on profitability. Therefore, forward trading of fuel and carbon are part of a sound hedging strategy. At the currently low level of carbon prices, it may seem that carbon price exposures are virtually non-existent. However, from 2010 to the beginning of 2016, EUA prices exhibited a relatively high volatility (around 50%), and the average year-on-year change in price was around €3/tCO_2, mostly downward. Given the global pressure to reduce future emission caps more sharply, the possibility of a price hike exists, and hedging this exposure should be part of an overall power plant hedging strategy.

The valid simulation of expected carbon costs is crucial to assess correctly the future value of a power plant, to assess the price exposures, and to take the right hedging decisions for managing these exposures. The value of a power plant equals the expected power income minus the costs for fuel, carbon and all relevant variable and fixed costs. A large number of price simulations has to be generated in order to assess the impact on this value of uncertain fuel, carbon and power prices. And for each simulation the optimal

2 The same example as in Figure 1. The hedge ratio in the pricing window gradually increases from 0% to 100%

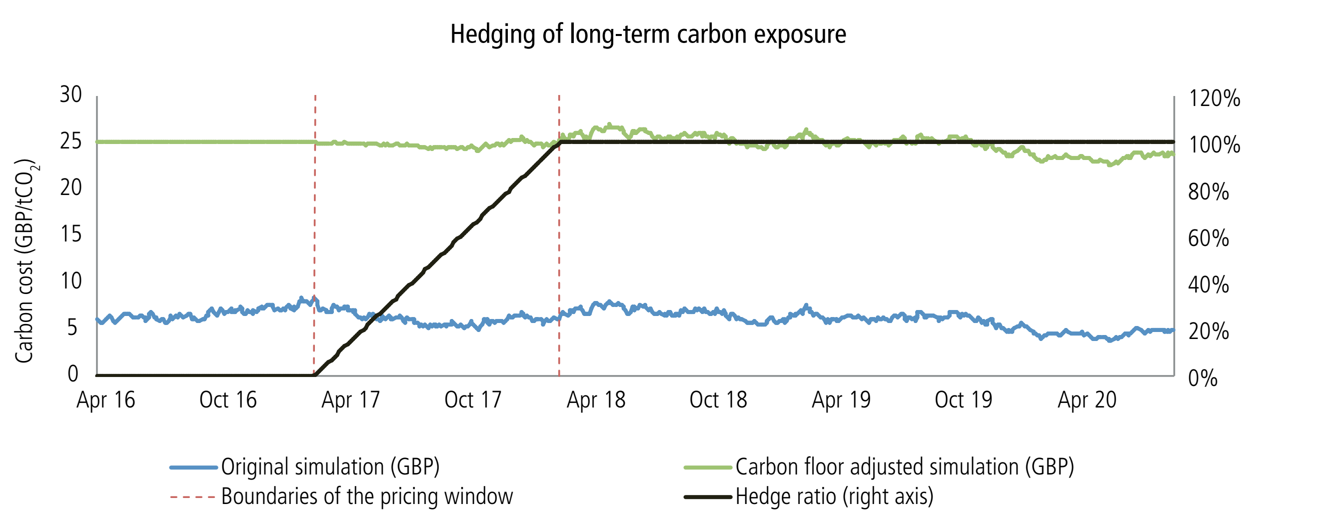

Source: Kyos modelling

3 Carbon cost simulations for power plant valuation in UK power market, adjusted for CPF policy. $CPS_{2016\text{–}2018}$ = £18

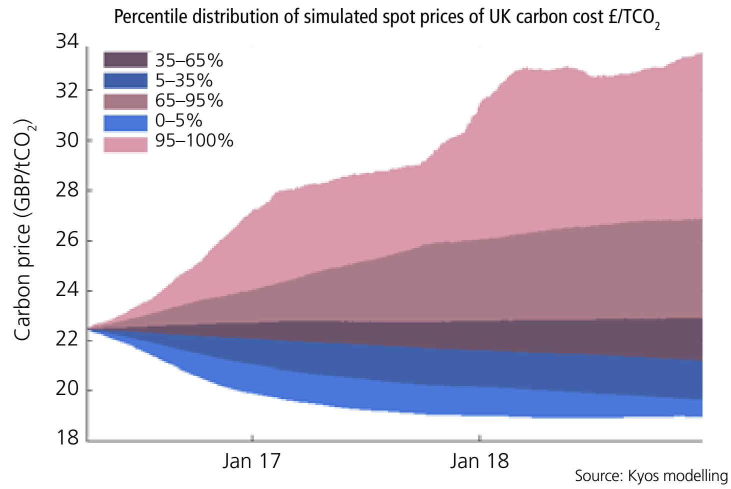

Source: Kyos modelling

dispatch and corresponding cashflows provide an assessment of the expected plant value and the price exposures. This is the essence of the Kyos KyPlant modelling software, which is used to optimally value power plants and optimally hedge the price exposures. The model derives the optimal hourly or half-hourly dispatch for a large number of price simulations using the least-squares Monte Carlo method. The KyPlant model takes into account all technical characteristics and costs of a power station, including for example detailed start curves, efficiency curves, emission caps, maintenance periods, minimum run-time and must-run obligations.

The concept of delta hedging

We assume that optimal hedging decisions are calculated over the *delta exposures*, which determine the optimal *delta hedges*. The optimal delta hedge is the volume that should be hedged in the market to minimise the exposure to commodity price changes. For a fuel-fired power plant, the delta exposure can be calculated for baseload power, peakload power, fuel and carbon (EUA) prices. There are delta exposures for different horizons, such as one month ahead, two months ahead and two years ahead.

The delta exposure of the expected power plant value V_t to changes in the commodity forward price $F_{y,t}$ at trading date t and for delivery period y, equals:

$$\delta_{y,t} = \frac{\partial V_t}{\partial F_{y,t}}$$

The optimal delta hedge is a market trade with a volume equal to minus the delta exposure. For example, if a power plant of 400MW capacity has a 100MW delta exposure to forward baseload power 2017, then a plant owner should optimally sell 100MW in the market. This hedge will ensure that when there is a (small) decrease in the 2017 baseload forward price, the loss in plant value is compensated by a profit on the hedge transaction of almost equal size. For example, suppose the forward price goes down by £1/MWh, and we have 8,760 hours in 2017. Then the plant value is expected to go down by 1 x 100 x 8,760 = £87,600, offset by a hedge profit of the same amount.

It can be shown that the value of the delta hedge is equal to the expected value of the commodity price exposure. In particular, the delta hedge of the 2017 power price can be calculated as the average value of the power production in 2017. For example, suppose that we have two simulations. In simulation one the average 2017 power production is 50MW, at a (weighted) average spot price of £30/MWh, while in simulation two it is respectively 150MW and £40/MWh. Then the expected value of the production is 8,760 x ½ (30 x 50 + 40 x 150) = 8,760 x 3,750 = £32,850,000. In order to hedge this value against today's forward price of £35/MWh (in between 30 and 40), the plant owner should sell 107MW in the forward market. This is the price-weighted average production: ½ (30 x 50 + 40 x 150)/ 35 = 107.

The example shows that a delta hedge is generally close to the expected volume of the underlying commodity (in this case 100 x 8,760MWh), in particular when there is limited correlation between price and volume. This expected volume $E[L_{y,t}]$ is therefore often used by market players as a proxy for the 'true' delta hedge. For power, this volume is the expected power production, for fuel the expected fuel consumption and for EUA the expected CO_2 emissions, all measured over a large number of simulations. Using the same set of Monte Carlo simulations, it is actually quite straightforward to derive the more accurate value-based hedges, as we did in the example with two price simulations. Mathematically, it equals:

$$\delta_{y,t} = E\left[P_{y,t} \cdot L_{y,t}\right] / F_{y,t} \quad \text{“delta hedge, value based”}$$

Using a large number of simulations, the expected value of the commodity, $E[P_{y,t} \cdot L_{y,t}]$, can be approximated by the average over the simulations.

Delta hedging carbon price exposure with a carbon floor

Without the carbon floor, the EUA delta hedge is roughly equal to the expected volume of CO_2 emissions across a large number of scenarios. Or, more accurately, the value of the delta hedge equals the expected value of the emissions. With the carbon floor, this is no longer correct. Instead, the optimal delta hedge volumes depend on the period we are trying to hedge. Hedging of the EUA price *exposure* should follow a similar logic as the modelling of the expected carbon *costs* described previously. There are three periods to consider:

1. After the pricing window, ie relatively nearby periods. The CPS price is known and is a tax on carbon emissions. Variations in the EUA market price affect the expected carbon costs directly. This means the EUA hedge is equal to the delta hedge without the carbon floor. For example, if the average value of the emissions corresponds to 400,000 tCO_2 in 2017, the EUA delta hedge should equal this volume. In Figure 2 this is displayed as 100% hedge ratio.
2. Before the pricing window, ie relatively far-ahead periods. The expected carbon costs equal the carbon floor, independent from (smaller) variations in the EUA price. This is at least true as long as the EUA market price is well below the carbon floor, and there is a negligible probability that it will rise above it. In these conditions, the optimal hedge ratio is 0%. This means that UK power producers should not procure any EUA emission rights for further-ahead time periods. That would be a speculative position.
3. In the pricing window, ie intermediate periods of about one-to-three years ahead, UK power producers should gradually buy emission rights. In the pricing window, the CPS is gradually determined based on the settlement prices in the EUA market. This means the hedge ratio linearly increases from 0% up to 100%.

If we call the optimal EUA hedge under the carbon floor policy $\delta^{CF}_{y,t}$ for emissions in fiscal year y and calculated at trading day t, we can write mathematically:

4 Distribution of power plant earnings with no hedges, hedging only power and gas risk, and hedging all exposures including carbon costs. When all exposures are hedged, the earnings distribution is the most narrow

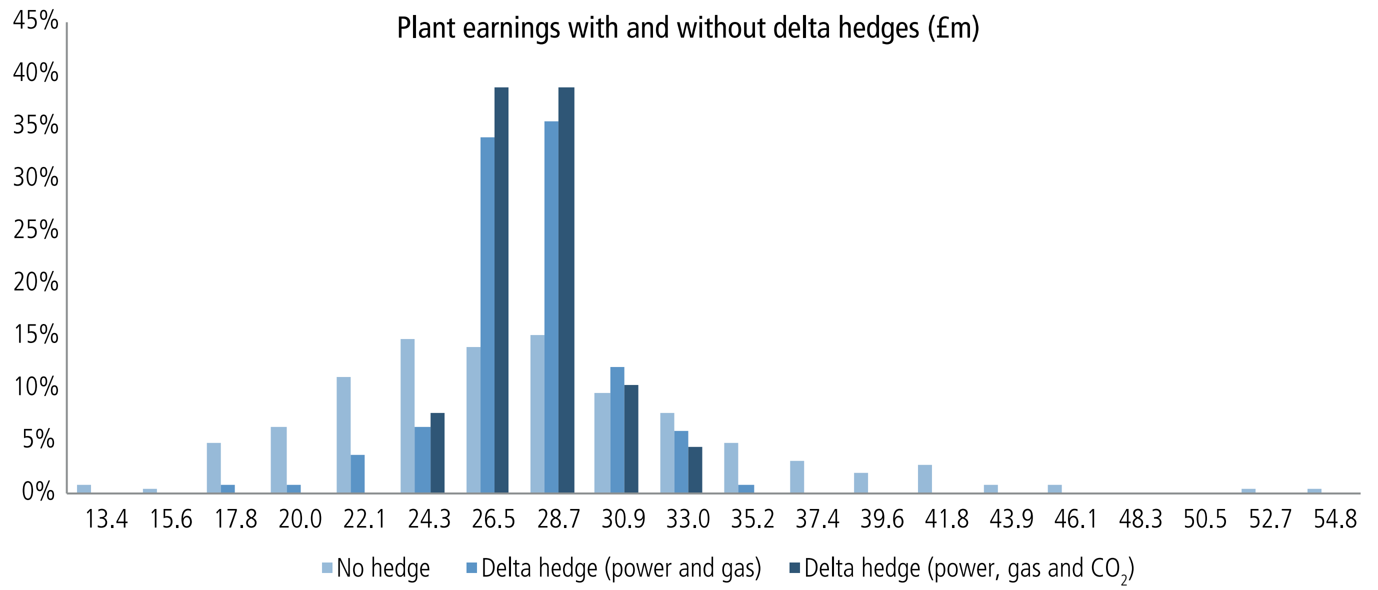

Source: Kyos modelling

$$\delta^{CF}_{y,t} = \omega_{y,t} \cdot \delta_{y,t}$$

$$\omega_{y,t} = \begin{cases} 0 & \text{if } t \le T_{y,1} \\ (t - T_{y,1} + 1) / (T_{y,2} - T_{y,1} + 1) & \text{if } T_{y,1} < t < T_{y,2} \\ 1 & \text{if } t \ge T_{y,2} \end{cases}$$

The optimal hedging of the EUA carbon price exposure is summarised in Figure 2. The black line corresponds to the hedge ratio (right y-axis). This hedge ratio should be multiplied with the 'normal' EUA delta exposure to know which volumes should be optimally hedged in the market to minimise the impact of EUA market price changes. This 'normal' delta exposure can be approximated by the expected volume of carbon emissions or, more accurately, by the value of those emissions. In any case, both should be calculated over a large number of Monte Carlo price simulations that capture the uncertain nature of commodity markets.

To assess the benefits of hedging the carbon price exposure, we valued a combined cycle gas turbine in the UK market. The 250MW plant has been valued over the period April 2016 to December 2018, assuming a maximum efficiency of 55%. The expected cumulative earnings are equal to £27.9 million. Over this time window the CPS level is already set by the UK government at £18/tCO_2. Figure 3 illustrates the distribution of carbon cost, which is one of the risk factors for the plant owner: the cost cannot go below £18/tCO_2, the CPS level, but due to the high EUA price volatility it may reach a level of more than £30/tCO_2 by the end of 2018. Without any forward hedges the earnings have a wide distribution with a 95% earnings-at-risk (EAR) of £9.22 million. Hedging the power and gas exposures reduces this risk considerably to an EAR of £4.75 million. Figure 4 illustrates that a considerable further reduction is achieved when the carbon price risk is hedged as well. In that case the EAR is £2.96 million.

Conclusion

Although the UK carbon floor is paid as a taxation on fuels, the implications for pricing and hedging are more complicated. For long-term periods, the relatively high level of the carbon floor determines the expected carbon costs. The costs are essentially fixed and there is virtually no exposure to EUA market price changes. Consequently, UK power producers should not buy any EUA emission credits for longer-ahead periods. For shorter-term time windows, the situation is entirely different: once the tax rate (CPS) has been determined, the carbon costs are a mix of the CPS and the EUA prices in the market, and there is full exposure to EUA market price movements. In between short- and long-term horizons, during the pricing window the CPS is set gradually and the EUA price exposure built up. Power producers should respond by gradually buying EUA credits in the forward market. The optimal hedge volume in the pricing windows is therefore derived from the expected future emissions and the number of settlement days that have already passed. ■

Vytautas Jurenas is quantitative analyst and
Cyriel de Jong is director at Kyos Energy Consulting.
E-mails: jurenas@kyos.com; dejong@kyos.com

REFERENCES

Energy Bill, 2013
The Parliament of Great Britain

Longstaff, FA and Schwartz, ES, 2001
Valuing American options by simulation: a simple least-squares approach
The Review of Financial Studies, p113–147

De Jong, C, Van Dijken, H and A Bundalova, 2011
Power plant hedging strategies: keeping your options open
Energy Risk, November, p63–67