

Cleaning Up Spark Spreads

How Plant Owners Can Reduce Risk Through Carbon Markets

By Michael Cragg, Richard Goldberg, Varoujan Khatchatrian, and Jehan DeFonseka

Contents

Introduction

Section 1

Gas-Fired Power Plants as Clean Spark Spread Options.....3

Section 2

Valuation of Clean Spark Spread Options.....4

Section 3

Hedging the Value of a Power Plant: Should You Bother?.....7

Conclusion

Appendix

The Brattle Group provides consulting and expert testimony in economics, finance, and regulation to corporations, law firms, and governments around the world.

We have offices in Cambridge, Massachusetts; San Francisco, California; and Washington, DC. We also have offices in Brussels, London, and Madrid.

For more information, please visit www.brattle.com.

Experience with the European Union’s carbon dioxide Emission Trading Scheme has shown that carbon price volatility can significantly affect the value of power plants. In particular, neglecting to take into account carbon price uncertainty when designing hedging strategies can substantially increase the financial risk from owning and operating power plants. This is relevant to operators in countries introducing a carbon trading scheme as well as operators considering modification to their current hedging practices.

In this paper we use numerical simulations to demonstrate that simple “three-asset” hedging strategies that address power, fuel, and emissions price uncertainty can significantly mitigate these financial risks. Using three-asset hedging strategies, as opposed to “two-asset” hedging strategies that only address power and fuel, can result in a two-thirds reduction in overall cash flow variation over a three-year horizon. As a result, while being profit neutral on average, these more advanced hedging methods may result in a cash flow increase of over 50 percent in 10 percent of the outcomes. Given the magnitude of these results, we recommend that plant operators who are not currently doing so explicitly address the uncertainty arising from the pricing of carbon emissions in their current cash flow hedging practices. As we demonstrate, this is particularly relevant to operators of dispatchable resources.

We also report sensitivities of the hedging strategy outcomes to planning horizon as well as the choice of re-hedging frequency. As expected, in the absence of transaction costs the risk mitigation benefits are more pronounced for more frequently used three-asset hedging. Additionally, the relative impact of our hedging method that incorporates the carbon emission dimension is more sizeable for hedging the exposures that are farther into the future.



Introduction

Traditionally, uncertainty over power and fuel prices has been the primary source of cash flow uncertainty for power plant owners.¹ One way to reduce the financial uncertainty associated with a power plant investment is to enter into long-term power sale and fuel purchase agreements that lock in a minimum level of operating profits. In fact, merchant generators are sometimes required to have such contracts before they are able to obtain debt financing for their up-front capital investment.

Long-term contracting is just one of many possible hedging strategies that power plant owners and operators can use to reduce their financial risk from market uncertainty. For example, many power plant owners put in place near-term power and fuel hedges on a rolling basis to reduce the near-term uncertainty in operating profit without trying to hedge away the risks and opportunities from having long-term positions in those markets. Since wholesale power markets began evolving from capacity-sharing arrangements into trading markets during the 1980s, power generation owners have become familiar with and regular users of power and fuel hedging programs.

Power and fuel are not the only inputs and outputs to a power plant with uncertain prices. Power plants consume chemicals, require spare parts, and need labor to operate. They also provide installed capacity (which is a market product in some regions, not only a required level) and can provide other so-called “ancillary” services to the electricity grid. However, in most cases, the scale of cash flow uncertainty over near-term planning horizons from these other inputs and products is far less than that of fuel and power prices; hence the traditional focus on fuel and power hedging.

A challenge to the primacy of fuel and power price uncertainty in assessing market risk is the advent of emission markets in Europe and other locations. Fossil fuel-powered power plants can produce emissions of sulfur and nitrous oxides, particulate matter, and carbon dioxide. In recent years, there has been a trend away from traditional regulatory emission limits toward market-based programs for emission reductions. The basic idea of a market approach is to allow trading in emissions credits between companies that can reduce emissions at low cost and companies for which emissions reduction would be extremely costly. The low cost emission reducers can profit from further emissions, and the high cost reducers can save money by purchasing emission permits at a lower cost than required for actual emission reductions. With the advent of emission markets, there is suddenly another power plant input/product with the potential for significantly affecting the financial risk of operating a power plant.

The European Union’s (EU) Emission Trading Scheme (ETS) is a cap-and-trade system aimed at allocating EU emission allowances (EUAs) among originators of greenhouse gases via economically efficient market mechanisms. The ETS adds an additional uncertain cost to the operation of European power plants from the EUAs required for plant carbon dioxide emissions. For example, according to IntercontinentalExchange, Inc., the ICE-ECX emissions index for EUAs for December 2010 ranged from 8.5 €/tonne of CO₂ (tCO₂) to over 30 €/tCO₂ during the period from late 2007 (when the December 2010 index began to be quoted) up to its settlement in December 2010.²

According to a recent consultation paper available on the European Commission’s website,³ electric power generation was responsible for about 55 percent of total EU ETS emissions as of 2008. At the same time, in the current market environment, carbon dioxide emissions account for more than 10 percent of all variable costs for relatively new natural gas-fired combined cycle power plants and for more than 25 percent for coal-operated plants.⁴ These cost shares are likely to increase significantly in a “high carbon price” environment (i.e., in a situation where the costs associated with carbon dioxide emissions are higher in the future). Therefore, the cost of carbon dioxide emissions is an important issue for power plant owners and operators.

Futures and option contracts for EUAs are now traded on centralized exchanges such as the European Climate Exchange (ECX). These and other carbon contracts are also traded in over-the-counter (OTC) markets. The emergence of forward markets for EUAs provides an opportunity for power plant operators to actively manage carbon costs as part of their normal risk management process (as opposed to being separately dealt with as an environmental compliance function). A natural question to ask is whether or not expanding hedging programs

Cleaning Up Spark Spreads

to address emission costs can significantly impact the resulting financial risk of operating a power plant. In other words, does it make more sense to pre-purchase emissions credits at fixed prices, or is it okay to wait and acquire them at spot market prices?

In this paper we compare the effectiveness of three-asset hedge programs (i.e., power, fuel, and emissions) versus two-asset programs (i.e., power and fuel only). Specifically, we use numerical simulations to demonstrate that by integrating forward emissions contracts into hedge programs using power and fuel forward contracts, companies can significantly reduce their financial risk. The next section introduces the notion of power plants as options, specifically focusing on the relationship between operating cash flows and carbon prices. Section 2 discusses several ways to value clean spark (or dark) spread options and highlights their relative merits. Section 3 addresses the key question of whether inclusion of forward carbon contracts into a hedge program can significantly further reduce the risks faced by power plant operators. This section also demonstrates the importance of hedging across a carbon dimension in a market environment with high carbon prices, such as the one predicted for Europe over the next decade. We conclude by summarizing our main findings.

Section 1 GAS-FIRED POWER PLANTS AS CLEAN SPARK SPREAD OPTIONS

The operating decisions for a dispatchable power plant have much in common with the exercise decision for a traditional financial spread option.⁵ In both cases, the decisions to transact are based on the relative prices of two market products: power and fuel for the power plant. In fact, it is very common to estimate the value and market price risk associated with the future operation of a dispatchable plant by treating the operating cash flows as if they came from financial spread options. The existence of the ETS market for carbon dioxide emissions introduces an additional layer of complexity into the valuation and risk management of gas-fired power plants operating in Europe.

In a world with no carbon markets, the payoff per unit of production from the operation of a dispatchable power plant over an operating cycle can be approximated with the following simple relationship:

$$\text{Payoff} = \max(E - HR \cdot F - K, 0),$$

where E is the price for electricity per MWh, HR is the power plant's heat rate (or the amount of fuel required to generate one MWh of electricity (measured in MMBtu/MWh)), F is the price of fuel per MMBtu, and K represents non-fuel variable operations and management (O&M) costs.^{6,7} The operating margin between the power price and the operating cost is sometimes called the "spark spread" for natural gas-fired units and the "dark spread" for coal-fired units. The "max" function in the payoff formula reflects the choice faced by the power plant operator: operate the plant when the electricity can be sold at profit but not otherwise. Such flexibility of choice gives rise to an "option value" of a power plant.

Like all other options, the option inherent in plant ownership can be valued as an expected value of future payoff discounted to the present. Because of the asymmetry between the high-margin case (where the plant operates at high levels to capture the benefit of high operating margins) and the low-margin case (where the plant operates at low levels, or shuts down, to limit the losses from cases with negative operating margins), the expected value of future payoffs will depend on the degree of uncertainty and correlations between future price outcomes. Hence, the value of future cash flows will not only depend on forward prices for power and fuel but also on their volatilities and the correlations between them.

In the presence of carbon markets and emission allowance price uncertainty, the payoff function needs to include an additional cost component:

$$\text{Payoff} = \max(E - HR \cdot F - ER \cdot C - K, 0),$$

Cleaning Up Spark Spreads

where ER is the power plant's emission rate measured in tCO_2/MWh and C is the price of carbon dioxide emissions per tCO_2 . We will refer to the operating margin modified to explicitly reflect emissions costs as a "green spark spread" for gas units and as a "green dark spread" for coal units. As in the case without internalized emission costs, the operator will keep the plant up and running only if it is profitable to do so, which is the case when the price of electricity exceeds the variable costs of fuel and carbon dioxide emissions by more than the variable O&M costs. As demonstrated in the examples below, the price of carbon emissions and emission rate can play an important role in deciding whether or not to operate the plant at a given point in time.

Example 1

Suppose a power plant operates at a heat rate of 7 MMBtu/MWh and has an emission rate of 0.36 tCO_2/MWh . Also assume that the price of electricity is 40 €/MWh, the price of gas is 4.7 €/MMBtu, and the variable O&M is 2.78 €/MWh. Finally, suppose that the operator faces an emission allowance price of 11 €/tCO₂. The profit per MWh of electricity produced by operating the power plant will therefore be equal to $(40 \text{ €/MWh} - 7 \text{ MMBtu/MWh} \cdot 4.7 \text{ €/MMBtu} - 0.36 \text{ tCO}_2/\text{MWh} \cdot 11 \text{ €/tCO}_2 - 2.78 \text{ €/MWh})$ or 36 cents per MWh. The operator is clearly better off running the plant than keeping it idle.

Example 2

Assume that conditions in the carbon market change so that the price of emissions jumps to 14 €/tCO₂. Also suppose this price change results in an increase in electricity prices from 40 €/MWh to 40.5 €/MWh, while all other parameters remain unchanged compared to the example above. If the operator kept running the plant, the payoff would now equal $(40.5 \text{ €/MWh} - 7 \text{ MMBtu/MWh} \cdot 4.7 \text{ €/MMBtu} - 0.36 \text{ tCO}_2/\text{MWh} \cdot 14 \text{ €/tCO}_2 - 2.78 \text{ €/MWh})$ or a loss of 22 cents per MWh. It is obvious that a rational operator would prefer to keep the plant idle rather than operate the plant, at least until the market conditions change such that plant operation is once again profitable.

Example 3

It is clear that the plant would be profitable even under the high carbon emission prices in Example 2 if the plant emission rate was lower than 0.36 tCO_2/MWh . Assume the same parameters as above but a "cleaner" plant with an emission rate of only 0.33 tCO_2/MWh . Such a plant is still profitable and will be operated in spite of a spike in carbon emission prices. The profit per MWh generated by such a plant is 20 cents.

These examples make it clear that the carbon emission rates and the market price of emission allowances play an important role in plant operating decisions and their profitability.

Section 2 VALUATION OF CLEAN SPARK SPREAD OPTIONS

In the previous section we introduced the notion of option value of a power plant and clean spark spreads, and described how cash flows from clean spark spread options are a function of electricity, gas, and carbon prices. In this section, we will discuss valuation techniques and computational specifics of these options.

As discussed in Carmona and Durrleman,⁸ simple two-asset spark spread options can be valued via a closed-form formula under certain restrictive assumptions and via closed-form approximations in the absence thereof. Things become somewhat more complicated once one adds the third dimension — carbon emission markets. The analytic calculation of the option value amounts to a triple integration over possible outcomes of each of the three uncertain prices. Instead, one can rely on a Monte Carlo simulation of the three prices to determine the distribution of option payoffs at maturity. In order to do so, one needs to know the probability distributions for the random variables of interest.

In this paper, we simulate these price processes using the Schwartz-Smith two-factor volatility model.⁹ According to this specification, forward prices for each maturity date evolve based on two sources of uncertainty — one that has a permanent impact on price (the "long" factor) and one that is transitory in nature and has an impact

Cleaning Up Spark Spreads

decaying over time (the “short” factor). The use of such two-factor models for electricity and gas prices is quite common among energy market practitioners and academics alike. We use this specification to simulate the electricity, gas, and carbon forward prices for one-month delivery futures contracts for each week between the valuation day and maturity. In the simulation examples we present below, we rely on the distributional and plant-specific assumptions outlined in Table 1.

The current price levels and plant operating parameters (emission rates, heat rates, and the variable O&M costs) are chosen to be consistent with a relatively new combined cycle gas-fired power plant. For example, the assumed plant heat rate of 7 MMBtu/MWh and the variable O&M of 2.78 €/MWh match typical mechanical and operating parameters for such a facility. Additionally, the assumed plant emission rate of 0.36 tCO₂/MWh is largely consistent with the heat rate assumption above and carbon dioxide emission factors for typical combined cycle gas-fired plants.¹⁰

The price volatility parameters were selected to be consistent with market data and those used in other studies. The volatility parameters for the electricity market are based on Kiesel et al¹¹ while the gas price parameters are based on our own estimates,¹² which are consistent with those reported by Treeck.¹³ Selection of “typical” levels for carbon volatility is more difficult, and in any implementation will require careful modeling as well as modifications for parameter uncertainty. Implementation challenges arise from the fact that carbon emission markets are relatively new and have been rapidly developing during the last several years. The observed volatilities of futures prices for December deliveries in 2010 through 2013 have been fluctuating between 20 percent and 40 percent during the January to September 2010 time interval. The carbon price volatilities inferred from the observed option prices (the option-implied volatilities) are generally higher and have been fluctuating in the 30 to 45 percent range. Finally, during September 2010 we observe a correlation between carbon and gas prices of 48 percent for December 2010 deliveries.

Based on what we observe in carbon markets, our choice of 25 percent volatility parameters and a 50 percent carbon and gas price correlation in Table 1 is rather conservative.¹⁴ We emphasize that in actual applications, careful analysis of the carbon and gas price processes for the given market is necessary. We have chosen to

Table 1 Model Parameters

Parameter	Value
Current Electricity Forward Price (<i>E</i>)	40 €/MWh
Current Fuel Forward Price (<i>F</i>)	4.7 €/MMBtu
Plant Heat Rate (<i>HR</i>)	7 MMBtu/MWh
Current Carbon Forward Price (<i>C</i>)	12 €/tCO ₂
Plant Emission Rate (<i>ER</i>)	0.36 tCO ₂ /MWh
Deterministic Costs (<i>K</i>)	2.78 €/MWh
Discount Rate (<i>r</i>)	5%
Electricity “long factor” volatility	15%
Electricity “short factor” volatility	40%
Electricity “short factor” mean reversion	1.40
Gas “long factor” volatility	45%
Gas “short factor” volatility	60%
Gas “short factor” mean reversion	2.75
Carbon “long factor” volatility	0%
Carbon “short factor” volatility	25%
Carbon “short factor” mean reversion	0.04
“Long” vs. “short” factor correlation (for each price process)	0%
“Long factor” correlations (between the three price processes)	50%
“Short factor” correlations (between the three price processes)	50%

Cleaning Up Spark Spreads

be conservative in our parameter selection because our objective is to demonstrate that even if carbon price volatility is small, there are large benefits to hedging carbon price exposure.

If we simulate the uncertain changes in the forward prices for all three markets over time between now and delivery, we can compute the cash flows from the plant during the delivery month. Those cash flows will be the same as the ones from spread options on the three prices. Repeating the process for many simulation runs provides us with a distribution of potential option payoffs. We can calculate the average of these payoffs and discount this average payoff back to the valuation day to obtain a Monte Carlo-based estimate of the clean spark spread option value. The results, based on 50, 100, and 500 thousand iterations, for three different maturity dates (one, two, and three years) are shown in the three Monte Carlo columns of Table 2.

Table 2 Value of Spread Option Using Different Pricing Techniques

Time to Maturity	Monte Carlo (50,000 iterations)	Monte Carlo (100,000 iterations)	Monte Carlo (500,000 iterations)	Second Order Boundary Approx.	Binomial Sampling Distribution (N = 20)	Binomial Sampling Distribution (N = 250)
1 year	5.9	5.9	5.9	5.8	5.9	5.9
2 year	7.3	7.3	7.3	7.0	7.3	7.3
3 year	8.2	8.2	8.1	7.7	8.1	8.1

As with any other Monte Carlo simulation, the higher the number of simulation runs, the more accurate the results will be. However, this accuracy comes at a cost — a simulation may take a long time to run, especially for long horizons. For this reason, other approximations are useful in cases where multiple valuation calculations are needed. One such approximate valuation method is the second order boundary approximation of Deng et al.¹⁵ Valuation using this method is shown in the fourth column of Table 2.

In the last two columns of Table 2 we report spread option values calculated using binomial sampling of the joint price outcomes at the delivery date (using two different numbers of samples for each price outcome). Although the binomial sampling method is generally quite accurate and far less computationally intensive than the Monte Carlo approach, it is still a bit more computationally costly than the second order boundary approximation when valuing a three-asset spread option.

In spite of the second order boundary method being less accurate than the other two alternatives, it is well-suited for the purposes of our exercise. As will become clear in the next section, we rely on the approximation methods only to the extent that we need to calculate the deltas — the sensitivities of the spread option prices to changes in each of the three underlying price components. As Table 3 demonstrates, option deltas with respect to power, fuel, and carbon prices are nearly identical across the two approximation methods. This holds true both for the “at the money” case (the starting point for our simulation) and for “in the money” and “out of the money” scenarios.

The results shown in Table 3 confirm that the second order boundary approximation is reasonably accurate as a method for assessing price sensitivities of the clean spark spread options. Given the computational speed of the second order boundary approximation, in our subsequent numerical experiments we use it to examine the relative merits of various hedging strategies.

Table 3 Spread Option Deltas Using Different Pricing Techniques for Different Plant Profitability Scenarios

Deltas with respect to	At the money ($E = 40$ €/MWh)		In the money ($E = 48$ €/MWh)		Out of the money ($E = 32$ €/MWh)	
	Second Order Boundary Approx.	Binomial Sampling Distribution (N = 250)	Second Order Boundary Approx.	Binomial Sampling Distribution (N = 250)	Second Order Boundary Approx.	Binomial Sampling Distribution (N = 250)
Power	0.5	0.6	0.6	0.6	0.4	0.4
Fuel	-2.2	-2.1	-2.8	-2.8	-1.5	-1.4
Carbon	-0.2	-0.2	-0.2	-0.2	-0.1	-0.1

Cleaning Up Spark Spreads

Section 3 HEDGING THE VALUE OF A POWER PLANT: SHOULD YOU BOTHER?

In a world with well-developed and highly liquid markets for clean spark spread options, a plant operator would be able to lock in a fixed payoff by simply writing a call option with heat rate (HR), emission rate (ER), and variable O&M cost (K) parameters that match the plant attributes. This would effectively create a deterministic stream of future income under any possible combination of the market prices. Unfortunately, such markets are highly illiquid and the existing clean spark spread products may offer only a limited menu of the aforementioned parameters and option exercise dates from which to choose. On the other hand, individual futures contracts on electricity, fuel, and (presently) carbon are more liquid and further developed. Such trades routinely take place in the centralized marketplace, such as ECX, and are also prevalent in the OTC market. Companies can therefore use these derivative instruments to hedge their cash flows, taking into account the variation in gas, electricity, and carbon prices. But the question remains: should they bother?¹⁶

Following, we measure the relative benefits of two hedging strategies. In one of the strategies, the power plant operator hedges its exposure with electricity, gas, and carbon futures based on its anticipated production levels. Under the second strategy, the power plant operator periodically rebalances only its portfolio of electricity and gas futures, while not hedging its carbon exposure. These strategies are often referred to as “delta hedging.”

In delta hedging, the operator takes long and/or short positions in a combination of futures contracts on electricity, gas, and carbon. It calculates its exposure to movements in each of the three prices — calculated as the sensitivity of the spread option value to a movement in the price (also known as “deltas”). These deltas dictate the number of futures contracts that should be held in order to hedge the cash flow exposure of the power plant. In a delta hedging strategy, the operator would acquire futures contracts that fully offset the deltas of its spread option (i.e., power plant). Then, at each portfolio rebalancing point, the operator would adjust its hedging positions to match the current deltas computed for the power plant.

In our numerical experiment, we carry out Monte Carlo simulations where we calculate for each simulation run the net present value of profits and losses from the hedge transactions combined with cash flows from the power plant. Unlike the previous case, where we used Monte Carlo simulation to calculate only the value of the spread option, in this simulation we also need to calculate the price sensitivities of the option value (i.e., the deltas) at intermediate time rebalancing points.

Multiple simulation runs produce a range of payoffs¹⁷ (as well as various summary metrics of risk) under two-asset and three-asset hedging strategies that can be compared. The key steps of our numerical experiment are outlined in the Appendix on page 13.¹⁸ In our exercise, we experiment with rebalancing frequency (weekly, monthly, quarterly, and yearly) and maturity time T (one, two, and three years).¹⁹ For maturities of two and three years, we also report results of the static hedging strategy, which amounts to hedging with futures at time $t = 0$ with no subsequent rebalancing of the positions.

In Table 4, we compare the results of our three-asset dynamic hedging strategy with the present value of the exercise value of the power plant option. The latter corresponds to a situation in which the power plant operator decides against any kind of hedging. These results are based on the parameter specifications from Table 1. The results for the one-year and three-year horizons are also shown in Figure 1 and Figure 2.

The table and the two graphs reveal a few important results. First and foremost, while unhedged positions may result in a rather significant upside,²⁰ they also produce zero payoff in more than 35 percent of the cases (corresponding to the decision to keep the plant idle). By contrast, hedging positions produce profitable outcomes in more than 95 percent of the cases. Also note that dynamic hedging ensures that companies achieve at least 50 percent of average discounted total value with 30 to 45 percent higher probability than under no hedging.²¹

Cleaning Up Spark Spreads

Table 4 Descriptive Statistics of the Discounted Total Value with Three-Asset Hedging

Hedging Frequency	Mean	Standard Deviation	1%	5%	10%	25%	50%	75%	90%	95%	99%
<i>Time-to-Maturity: 1 year</i>											
Weekly	5.9	0.8	3.7	4.6	4.9	5.5	5.9	6.4	6.8	7.2	8.2
Monthly	5.9	1.6	2.2	3.4	4.1	5.1	5.9	6.7	7.6	8.3	10.5
Quarterly	5.9	2.6	0.5	2.0	2.8	4.3	5.9	7.3	8.8	10.1	13.9
Yearly	5.9	4.7	-1.8	0.1	1.0	2.7	5.1	8.3	11.7	14.1	20.9
No hedging	5.9	7.8	0.0	0.0	0.0	0.0	2.2	10.1	17.2	21.6	30.7
<i>Time-to-Maturity: 2 years</i>											
Weekly	7.3	0.8	5.1	5.9	6.3	6.8	7.3	7.8	8.2	8.5	9.4
Monthly	7.3	1.4	3.5	4.9	5.6	6.5	7.3	8.0	8.8	9.4	11.3
Quarterly	7.3	2.4	1.7	3.5	4.4	5.9	7.3	8.5	9.8	10.9	14.4
Yearly	7.3	4.3	-1.1	1.2	2.4	4.4	7.0	9.5	12.0	14.1	20.8
Static	7.3	5.9	-2.7	0.0	1.3	3.5	6.4	10.1	14.0	16.9	26.1
No hedging	7.3	9.0	0.0	0.0	0.0	0.0	3.9	12.4	20.0	25.1	36.0
<i>Time-to-Maturity: 3 years</i>											
Weekly	8.1	0.9	6.0	6.7	7.1	7.5	8.1	8.7	9.3	9.7	10.5
Monthly	8.1	1.4	4.4	5.8	6.5	7.3	8.1	9.0	9.7	10.3	11.9
Quarterly	8.1	2.3	2.6	4.5	5.4	6.8	8.2	9.4	10.6	11.5	14.7
Yearly	8.1	4.1	-0.4	2.2	3.4	5.6	8.0	10.3	12.5	14.4	20.9
Static	8.1	6.8	-3.1	0.0	1.5	3.9	7.3	11.2	15.3	18.5	29.6
No hedging	8.1	9.6	0.0	0.0	0.0	0.0	5.1	13.6	21.5	26.8	38.8

Second, a more active involvement in derivatives markets generally results in better protection from the downside risk. For example, for each of the hedging horizons, weekly re-hedging yields better results at the bottom 1st, 5th, and 10th percentiles. Also note that the standard deviations of the hedged positions are smaller, which implies that weekly re-hedging results in a more predictable stream of profits for the power plant operator. Essentially the dynamic hedging allows the operator to make the payoff profile significantly smoother by “trading-in” the upside in exchange for protection against unfavorable realizations. In some situations, such as a yearly re-hedging for a two- or three-year horizon, such risk mitigation may be achieved at the expense of taking an exposure to a loss at a very low probability (far less than 5 percent).

Finally, the farther away the plant operation decision date, the higher the cumulative uncertainty, and therefore the more valuable the dynamic hedging. For example, while a “no hedging” scenario produces zero payoff at the 5th, 10th, and 25th percentiles, irrespective of the time-to-maturity, the discounted to total value with weekly, monthly, and quarterly frequencies grows uniformly for the same probability cutoff levels as the time-to-maturity grows.

By now it should be clear that a power plant operator can protect itself, to a large degree, from unfavorable realizations of market prices by engaging in active hedging. But how important is it to adjust the traditional spread option approach for gas and electricity pricing to include hedging across the third carbon dimension? Would the power plant operators be significantly worse off if they ignored the uncertainty in the carbon markets and simply focused on hedging the electricity price and fuel cost risks?

In order to answer the question posed above, we compare the distribution of losses under the weekly hedging strategy across all three assets with a weekly hedging strategy in which only electricity and gas price risks are hedged away. The simulations are run under the assumption that carbon price levels are three times as high as today.²² A complete set of summary statistics for weekly dynamic hedging strategies involving two- and three-assets is presented in Table 5.

The comparative results for various hedging horizons are also plotted in Figures 3 through 5. A couple of observations are evident from these analyses. Unlike the case with no hedging, both two- and three-asset hedging strategies yield positive total discounted values at the lower spectrum of outcomes, such as the bottom 5th, 10th,

Cleaning Up Spark Spreads

Figure 1 Discounted Total Value with Three-Asset Hedging (1-Year Horizon)





Figure 2 Discounted Total Value with Three-Asset Hedging (3-Year Horizon)



Cleaning Up Spark Spreads

and 25th percentiles. Upon closer examination, however, the benefits of a more comprehensive hedging strategy involving carbon become obvious. More specifically, for the one-year horizon case, the bottom 5th percentile of a two-asset strategy is only 2.8 €/MWh, while it is 4.8 €/MWh for a three-asset hedging strategy. This difference constitutes approximately 32 percent of the average (or expected) discounted total value of 6.2 €/MWh. The difference is more pronounced at the 1st percentile, equaling about 51 percent of the average discounted total value.

We can also see the benefits of a three-asset hedging at the bottom 10th and 25th percentiles. Hedging involving carbon also allows the operator to significantly reduce the uncertainty of the outcomes. For example, the range between the bottom and top 5th percentiles in the one-year horizon case reduces from 6.5 to 2.7 €/MWh (i.e., by more than 60 percent of the average discounted total value) when substituting a two-asset strategy with a more comprehensive one.

The benefits of three-asset hedging are even more pronounced for longer horizons. For example, with a three-year horizon the weekly hedging strategy involving carbon results in the bottom 5th percentile value of 7.0 €/MWh, whereas a less sophisticated two-asset hedging produces only 3.3 €/MWh, a difference of 44 percent of the average discounted total value of 8.4 €/MWh. The difference between the two hedging alternatives at the 1st percentile equals 75 percent of the average discounted total value. Also, the range between the bottom and top 5th percentiles reduces from 9.8 to 3.1 when shifting from a two-asset strategy to the three-asset hedging — a reduction of nearly 80 percent of the average discounted total value. Thus, hedging the clean spark spread instead of the just the spark spread (i.e., revamping the two-asset hedging approach to a three-asset hedging approach) can raise cash flow by over 50 percent in 10 percent of the outcomes (see Table 5) and reduce overall cash flow variation by two-thirds, from 36 percent of the mean to 12 percent of the mean.²³

Hedging across the carbon dimension (along with electricity and gas) provides much better protection at the lower end of the distribution and reduces the overall riskiness of the payoff per unit of production compared to a strategy that ignores carbon price uncertainty.²⁴ Even if one expected that at the current levels of carbon prices the relative benefit from dynamic hedging across all three dimensions (i.e., electricity, gas, and carbon) might have been somewhat muted and provide only marginally better downside protection compared to hedging that ignores carbon volatility, this is not true in an environment with higher carbon prices.

As recently as May 2010, the European Union's climate commissioner announced that she was planning to make the case for setting a much higher CO₂ reduction target than originally planned.²⁵ Even prior to an announcement of such plans, it was argued in a June 2009 Point Carbon press release²⁶ that carbon prices were expected to double by 2013 and nearly triple by 2016, reaching 40 €/tCO₂. Thus, going forward, the benefits from actively hedging carbon exposure will become much more sizable and should not be ignored by plant operators.

Table 5 Descriptive Statistics of the Discounted Total Value for Weekly Hedging (Higher Carbon Prices)

Hedging Frequency	Mean	Standard Deviation	1%	5%	10%	25%	50%	75%	90%	95%	99%
<i>Time-to-Maturity: 1 year</i>											
No hedging	6.2	8.6	0.0	0.0	0.0	0.0	1.6	10.4	18.5	23.8	34.5
2-Asset	6.2	2.0	0.7	2.8	3.7	5.0	6.3	7.4	8.6	9.3	10.6
3-Asset	6.2	0.9	3.9	4.8	5.2	5.7	6.2	6.7	7.1	7.5	8.7
<i>Time-to-Maturity: 2 years</i>											
No hedging	7.5	9.9	0.0	0.0	0.0	0.0	3.1	12.6	21.6	27.5	40.7
2-Asset	7.5	2.6	0.2	3.1	4.3	6.1	7.6	9.2	10.7	11.5	13.0
3-Asset	7.5	0.9	5.2	6.1	6.5	7.0	7.5	8.0	8.5	8.9	9.8
<i>Time-to-Maturity: 3 years</i>											
No hedging	8.4	10.8	0.0	0.0	0.0	0.0	4.2	14.0	23.4	29.9	44.2
2-Asset	8.4	3.0	-0.2	3.3	4.7	6.7	8.5	10.4	12.1	13.1	14.8
3-Asset	8.4	1.0	6.1	7.0	7.3	7.8	8.3	9.0	9.7	10.1	11.0

Cleaning Up Spark Spreads

Figure 3 Three-Asset Versus Two-Asset Weekly Hedging (1-Year Horizon and Higher Carbon Prices)




Figure 4 Three-Asset Versus Two-Asset Weekly Hedging (2-Year Horizon and Higher Carbon Prices)





Figure 5 Three-Asset Versus Two-Asset Weekly Hedging (3-Year Horizon and Higher Carbon Prices)



Conclusion

The results of our hedging experiments are rather striking and directly relevant to any power plant operator who questions the potential benefits of carrying out a three-asset hedging strategy. The regulatory push to internalize emissions costs by emitters creates additional financial risk to power plant owners and operators. Our paper demonstrates that hedging strategies that explicitly address emissions cost uncertainty can significantly mitigate the overall market price risk associated with power generating assets. Whether falling under the European Union's carbon dioxide Emission Trading Scheme or in other jurisdictions where regulators are planning to address emissions costs through market means, power generators will increasingly need to consider emissions as another commodity associated with plant operations and become accustomed to treating emissions on equal footing with the more traditional commodities of power and fuel.

By taking a combination of short and long positions in the electricity, fuel, and carbon futures contracts, and by periodically reevaluating and rebalancing those positions in response to changes in market conditions, companies can enjoy much more predictable streams of profit per unit of electricity generated than under a passive wait-and-see approach, especially as hedging is done more frequently.²⁷ These results are hardly surprising given the extensive use of fuel and power hedging instruments by plant operators to make cash flows predictable. What may be surprising is the size of the reduction in cash flow volatility that can be achieved from directly hedging for carbon price risk. Our simulations show that in high carbon price environments, overall cash flow variation can be reduced by two-thirds, from 36 percent to 12 percent of average cash flow, over a three-year horizon. Consequently, while being overall profit neutral in expectation, cash flows can be raised by over 50 percent 10 percent of the time. The additional reduction in cash flow risks should reduce financing and management costs that arise from extreme outcomes. In other words, "cleaning up" the hedging of spark spreads (and dark spreads) by including emissions costs can add significant value to generators.


Cleaning Up Spark Spreads

Appendix: Simulation Models

Our numerical experiment proceeds in the following steps:

1. At $t = 0$, calculate the spread option deltas with respect to electricity, gas, and carbon forward prices.²⁸
2. Use the calculated deltas to establish the delta hedge position in the three underlying markets. Put in place these hedge positions in either two or three markets (depending on the hedging strategy being simulated).
3. Move forward one time step and calculate the current time realizations of forward prices using the two-factor price processes.
4. If the current time is a hedge rebalance time, compute updated deltas based on the current realization of forward prices. Rebalance the hedge position and save the associated profit and loss from liquidating existing contracts as well as the prices of any new hedge contracts being held. (Since transaction costs are not being modeled in the simulation, rebalancing can be carried out by liquidating all previously-held hedges at current market prices and then acquiring new delta hedges at the current prices.)
5. Repeat steps 3 and 4 above for time $t = 2, \dots, T - 1, T$ (Steps 3 and 4 amount to simulating different price paths for the underlying assets as shown in Figure A below. In all of our experiments, we choose a weekly time step for simulation.)
6. At time T , calculate the cash flows from the operation of the plant if such operation is “in the money” and from the liquidation of any hedge contracts still being held.
7. Repeat steps 1 through 6 above 100,000 times.
8. Calculate the present value of hedging strategy profit/loss by discounting quantities calculated above to time $t = 0$. This measure is referred to as a “discounted total value” or simply a “total value.” We also plot its probability distribution function and calculate the distribution’s mean, median, standard deviation, and 1st, 5th, 10th, 25th, 75th, 90th, 95th and 99th percentile values.

Figure A Sample Paths for Electricity Forward Price



► To learn more about our expertise and our consulting staff, including our affiliated senior advisors, please visit www.brattle.com.

Endnotes

¹ In this paper we are focused solely on market price risk. There are, of course, other important risks associated with owning and operating power plants including operating safety risk, counterparty credit risk, and regulatory risk, among others. Although we will not discuss these risks here, we do not intend to imply that they are not important and interesting topics of study and attention.

² From "ICE-ECX European Emissions - Emissions Index," *IntercontinentalExchange, Inc.*, data retrieved January 10, 2011. Available at: <http://www.theice.com/marketdata/reports/ReportCenter.shtml?reportId=82&productId=390&hubId=564>.

³ "Technical Aspects of EU Emission Allowances Auctions," European Commission, August 3, 2009, page 11. Available at: http://ec.europa.eu/clima/consultations/0002/cons_paper.pdf.

⁴ These calculations are based on typical heat rates and emission factors for coal- and gas-fired plants in a reasonable higher heating value (HHV) range reported on the U.S. Department of Energy's website. Available at: <http://www.eia.doe.gov/oiaf/1605/coefficients.html>. See Section 3 for further discussion of reasonable parameters for a relatively new gas-fired power plant.

⁵ A spread option is a type of option that derives its value from the difference between the prices of two or more assets. In the traditional setup, the option is exercised when the difference between the two prices is positive (i.e., when the difference between the revenue derived from a power plant is greater than the costs of operating it). This was traditionally modeled as the spark spread, or the spread between electricity revenue from a plant and the fuel costs to produce this electricity less the fixed production costs. The clean spark spread is a three variable spread because in addition to the spark spread, the costs of carbon emissions (which vary with the price of carbon) are taken into consideration.

⁶ We refer to O&M costs as variable costs because they vary linearly with the changes in output. Throughout this paper, we treat this cost component as non-stochastic and known in advance. In other words, the variable O&M is treated as fixed (or pre-determined) per one MWh of electricity produced.

⁷ In this approximation, estimated fuel and non-fuel start costs over the operating cycle are often allocated on an estimated per-unit-output basis and included in K when the unit is typically not operating outside of the operating cycle.

⁸ Carmona and Durrleman, "Pricing and Hedging Spread Options," *SIAM Review*, October 2003.

⁹ See Schwartz and Smith, "Short-Term Variations and Long-Term Dynamics in Commodity Prices," *Management Science*, July 2000, or "Forward Curve Dynamics and Asset Valuation," EPRI Technical Brief TB-108991, October 1997.

¹⁰ See <http://www.eia.doe.gov/oiaf/1605/coefficients.html>. For example, for a gas-fired power plant with HHV in the range of 1000-1025 Btu/scf, the emission factor of 52.9 kgCO₂/MMBtu would imply a carbon emission rate of 0.37 tCO₂/MWh (= 52.9 kgCO₂/MMBtu × 7 MMBtu/MWh multiplied by a factor of 0.001), which is reasonably close to the assumption used in our analysis.

¹¹ Kiesel, Schindlmayr, and Borger, "A Two-Factor Model for the Electricity Forward Market," Ulm University Working Paper, August 20, 2007. Available at: http://www.uni-ulm.de/fileadmin/website_uni_ulm/mawi.inst.050/people/kiesel/publications/articleresubmitv1008.pdf.

¹² We analyzed the daily National Balancing Point gas price data for August 1996 through October 2009 for one-, two-, and three-month ahead futures contracts with monthly delivery.

¹³ Treeck, "The Hedge Effectiveness of European Natural Gas Futures," University of St. Gallen Doctorate Dissertation, November 23, 2009.

¹⁴ Note that for the carbon prices we are using a model with only one factor with a very slow decay. This is consistent with our observation that future price volatilities of carbon decline only slightly for more distant contract maturity dates.

¹⁵ Deng, Li, and Zhou, "Multi-Asset Spread Option Pricing and Hedging," SSRN Working Paper, October 29, 2007. This approximation essentially amounts to reducing the three-dimensional integration in the analytical expression of option value to a combination of four two-dimensional integrals. It works under certain assumptions such as jointly-normal returns, non-negative strike price, and non-zero determinant of the correlation matrix.

¹⁶ Carbon derivatives traded on the centralized exchanges are still somewhat limited in terms of the exercise dates (i.e., carbon options traded on ECX are concentrated in December exercise dates). In our hedging example, we assume that the decision date is on the same day in December as the futures exercise date.

¹⁷ This hedging strategy may also be subject to transaction costs associated with periodic rebalancing of the portfolio of futures contracts, such as fixed trading costs or variable costs arising from the bid-ask spread. In this paper, we disregard such costs. According to some recent studies, bid-ask spreads have declined dramatically as the trading activity in the carbon markets intensified (see, for example, Kruk, "Execution costs in money and futures markets," University of Sydney Doctorate Dissertation, December 2009).

¹⁸ Note that under the hedging strategy that involves only electricity and gas futures contracts, the deltas are calculated only with respect to electricity and gas prices, while the simulated plant payoffs will depend on all relevant components including, of course, the emission rate and the carbon emission price.

¹⁹ Since our simulation time step is weekly, the "monthly" re-hedging frequency means a hedging that takes place every four weeks (i.e., slightly more frequently than once in every calendar month), while the "quarterly" re-hedging happens every thirteen weeks.

²⁰ See 90th, 95th, and 99th percentiles for the "No hedging" rows in Table 4.

²¹ Note from Figure 2 and Table 4 that at least half of the average discounted total value (i.e., 4.05 = 8.1 × 50%) is attained with only about 55 percent probability in the no hedging case, but with approximately 85 percent probability with annual re-hedging and with almost certainty in the weekly re-hedging case.

²² We increase electricity cost per MWh dollar-for-dollar (i.e., in the amount equal to increment in carbon prices times the plant's emission rate). Thus the plant's option value is still "at the money".

²³ In Table 5 we calculated the coefficient of variation (ratio of standard deviation to mean) for the two-asset and three-asset weekly hedging over three years to be 36 percent and 12 percent. Thus, the three-asset case is two-thirds of the two-asset case.

²⁴ Some industry professionals note that it is not unlikely that European policy-makers will institute some type of carbon price caps and floors if emission prices become extremely high or if they drop to extremely low. In order to address the concern that the comparative results for the two- and three-asset hedging strategies are driven by unrealistically high/low realizations of carbon prices, we performed a separate exercise and discarded iterations in which the simulated carbon price was very low or very high at any point during the three-year simulation horizon. For example, limiting our analysis only to iterations with carbon prices inside the 16 €/tCO₂ to 90 €/tCO₂ range, we end up with more than 92,000 (out of 100,000) iterations without impacting the conclusions above in any significant way.

²⁵ See "Europe to examine case for bigger CO₂ cuts," BBC News, May 11, 2010. Available at: http://news.bbc.co.uk/2/hi/science_and_environment/10109088.stm.

Cleaning Up Spark Spreads

²⁶ See "CO₂ trading at €30/t in 2013, €40/t in 2016, predicts Point Carbon," Point Carbon, June 3, 2009. Available at: <http://www.pointcarbon.com/aboutus/pressroom/pressreleases/1.1130832>. Also see "EU carbon price forecast to nearly triple by 2016," ENDS Europe, June 3, 2009. Available at: <http://www.ends europe.com/21466>.

²⁷ Real-world implementation of these strategies must take into account a variety of factors, not the least of which are transaction costs. Although doing so will limit the effectiveness of very high frequency rebalancing, the overall conclusions of the paper will still apply.

²⁸ Deng, et al, derive an approximation for deltas based on the second order boundary method. Alternatively, such deltas could have been calculated numerically. See Deng, Li, and Zhou, "Multi-Asset Spread Option Pricing and Hedging," SSRN Working Paper, October 29, 2007.

About the Authors



Michael Cragg
Principal

Phone: +1.617.864.7900
Email: Mike.Cragg@brattle.com

Dr. Cragg has extensive research, teaching, consulting, and expert witness experience in financial services, corporate finance, public finance, and valuation. He has assisted numerous corporations, the U.S. Department of Justice, and the U.S. Internal Revenue Service in developing economic and financial testimony in complex financial litigation. In addition, he has worked for senior management at energy companies to understand and optimize their hedging and marketing operations, and advised on issues related to capital structure.

Dr. Cragg holds a Ph.D. in Economics from Stanford University.



Richard Goldberg
Principal

Phone: +1.415.217.1000
Email: Rich.Goldberg@brattle.com

Dr. Goldberg helps clients make better asset management, risk management, resource planning, and pricing decisions in a complex, uncertain, and competitive business environment. He provides consulting services on calculating the value and managing the risks associated with physical assets, contract commitments, and retail customers. He also performs statistical analyses of historical price and trading data to develop probabilistic price forecasts for use in valuation and risk management analyses.

Dr. Goldberg holds a Ph.D. in Physics from Stanford University.



Varoujan Khatchatrian
Associate

Phone: +1.617.864.7900
Email: Varoujan.Khatchatrian@brattle.com

Dr. Khatchatrian's areas of expertise include risk management and asset pricing. His cost of capital work focuses on engagements in the context of regulatory policy and ratemaking in water, electricity, and natural gas sectors in the United States and Canada. Prior to joining *The Brattle Group* he spent several years in the risk modeling industry.

Dr. Khatchatrian holds a Ph.D. in Economics from Duke University.

Jehan DeFonseka

Mr. DeFonseka is a former Research Analyst at *The Brattle Group*. He is currently pursuing his MBA at Harvard Business School.

The Brattle Group

ECONOMIC AND FINANCIAL EXPERTS

CAMBRIDGE

SAN FRANCISCO

WASHINGTON

BRUSSELS

LONDON

MADRID

www.brattle.com