VALUING COMBINED CYCLE GAS TURBINE POWER PLANTS IN ENERGY TRADING

by Mark Houldsworth, PhD and Rich Pastore, CFA

I. Overview

The method by which energy trading operations calculate the value and measure the risk for combined cycle gas turbine (CCGT) power plants presents various challenges and pitfalls. In today's market, most practitioners utilize one of two dominant methods:

- Closed Form Approximations (Kirk) – Many choose this closed form spread option approach for its ease-of-use and speed, but this method has a number of weaknesses, especially relating to the incorporation of start-up costs and other state-transition costs.

- Monte Carlo Solutions – To compensate for the shortcomings of closed form approaches, other practitioners utilize Monte Carlo, a technique clearly capable of handling the multiple dimensions of the combined cycle valuation problem; but its computationally-intensive problem convergence often proves impractical, taking too much time to calculate all of the required hedge sensitivities.

In this article, we present a better valuation choice. Our method utilizes a three-dimensional Gaussian quadrature capable of generating the precision of Monte Carlo with the quick run-time of a closed form calculation. Our technique provides answers in seconds and the flexibility to integrate, either tightly or loosely, into most energy trading risk management (ETRM) systems.

II. Closed Form Solutions Oversimplify CCGT Plants

ETRMs most commonly deploy a closed form spread option method to value and capture the risk associated with generation assets. For example, at the end of the day, the ETRM will call the Kirk spread option model. All of the underlyings, their volatilities, the position, and other data will be fed into that spread option model, which will then generate and populate a long power delta, a short gas delta, and other hedge sensitivities for the relevant power price curves. These curves will generally include the 5x16, 5x8, 2x16, 2x8, and sometimes a super peak curve.

The closed form model provides the user with significant speed advantages, generating results – even for large portfolios – in minutes.

If the plant is a high heat rate peaker and one populates a super peak curve, then one might achieve satisfactory results with a spread option instrument struck at the elevated heat rate. Typically, these peaker plants have one operating mode and a small start cost that can easily be folded into the strike. Their off-peak risk profiles, for the most part, are negligible; however, imprecision still exists in this readily-dismissed area.

Where high heat rate peakers typically have a small degree of imprecision with regard to off-peak risk measurement, the CCGT power plant presents a far more substantive problem. The owner of the CCGT clearly has a number of spread-like options at expiry based on the realized day-ahead schedule. However, one cannot capture this array of possible expiry payoffs in a useful way with a 2-factor spread option instrument, like the Kirk model.

Consider the following three complicating aspects when analyzing a CCGT:

- ***Start Cost:*** Even if the start costs are fully loaded onto the 5x16 or 2x16 instrument strikes, the off-peak deltas typically vanish or are very low. Very frequently, however, CCGT operators run at a loss, either at HSL or LSL, to avoid the start cost. This may be the optimal solution, and it involves generating power and consuming natural gas off-peak; but the simple closed form solution shows sub optimal deltas and ignores the true terminal position possibilities.

- ***Multiple Operating Modes:*** Typical CCGT plants have multiple operating output states, each with its own average heat rate and sometimes its own variable operating and maintenance (VOM) charge. It is difficult to fold these conditions into a coherent spread option structure and maintain the logic.

- ***Hourly Optimization:*** Positing spread options on aggregated curves ignores the hourly optimization which plant operators perform in day-ahead markets. At expiry (day-ahead) plant operators will bid the solution to a dynamic programming problem. This will be a very precise solution that will obey Min Up, Min Down, and perhaps Max Start constraints imposed on the asset. It is impossible to capture this behavior in a spread option structure pointed at aggregated curves.

These three aforementioned features constitute the dominant problem when attempting to capture CCGT positions with spread options in an ETRM. They apply equally to physical positions, tolling positions, and revenue puts or calls. The illustration below demonstrates the potential magnitude of error that a closed form solution may suffer in comparison to the superior results of our Gaussian quadrature method.[1]

Inputs

CCGT Inputs	
Plant MW	500
Heat Rate	7.40
VOM	$2.00
Start Cost	$10,000
Market Inputs	
Gas	$3.04
5x16	$28.78
5x8	$18.63
Vol Gas	35%
Vol 5x16	40%
Vol 5x8	25%
Corr Gas/5x16	75%
Corr Gas/5x8	65%

Results for 4-Month Capacity Position

	Closed Form	Gaussian Quad	Difference
Margin	$13,658,705	$14,941,539	**$1,282,834**
Gas Delta	(22,952,872)	(24,220,059)	**(1,267,187)**
On-Peak Delta	3,167,961	3,086,811	**(81,150)**
Off-Peak Delta	100,939	396,917	**295,978**

[1] This is dominantly an illustrative result, there are a number of other assumptions behind the results here. The authors are happy to provide a remote demonstration with further details on request. Also, see Sample Output section on page 5 to see the full family of Greeks generated.

III. The Monte Carlo Alternative Presents a Different Set of Difficulties

If we leave the closed form methods that are native to ETRMs, we enter a realm of models developed outside the system. These external models are then integrated, either loosely or tightly, into the ETRM with the intent to more accurately capture risk for gas-fired power plants, especially CCGT plants, through the utilization of Monte Carlo techniques, thereby moving the problem to 3-dimensions, jointly realizing on-peak, off-peak, and natural gas prices

If one is willing to accept a GBM price process, and presume a very large number of samples, we can state the 3-dimensional Monte Carlo value solution as follows:

$$E[Value_0] = \frac{1}{n} * \sum_{i=1}^{n} Payoff_i * e^{-rT}$$

The term $Payoff_i$ is a very general term here. It contains the randomly realized and correlated underlyings for the i-th simulation.[2] Further, for every realized price triple, there will be a unique, optimal, algebraic payoff that can be calculated and stored.

All of the possible output states for both on-peak and off-peak can be valued along with all of the transition costs between states. If hourly granularity and hourly constraints are desired, one can impose hourly shapes and embed dynamic programming solutions for each triple.

In the end, one can extract more accurate deltas, vegas, gammas, and cross gammas using finite difference methods.

The upside in using the Monte Carlo is that one achieves a far more accurate solution at expiry and the resulting deltas, therefore, will be far more accurate.

There is a big downside, however, with the Monte Carlo. The solution tends to be very slow – perhaps too slow to be practical in an end-of-day calculation for the ETRM. Even the calculation of a 2-factor spread option in Monte Carlo takes a very large number of samplings to achieve convergence. This problem compounds in three dimensions. Moreover, to generate a complete set of hedge sensitivities, the premium may need to be calculated 10 to 15 times for weekdays, weekends, and each calendar month.

[2] If one is confident that the correlation imposed is of full rank, then Cholesky can be used to generate correlated random samples in this. Otherwise one can use PCA or Singular Value Decomp methods.

IV. A Better Choice – Robust Valuation that Utilizes Gaussian Quadrature Techniques

"Closed form" generally means that the solution involves only mathematical functions – many of us take that term to mean "speedy". Keeping speed in mind, the Gaussian quadrature method loses very little of the speed advantage of the closed form while still maintaining the accuracy and robustness of Monte Carlo. The model achieves stable and fast calculations using Gaussian quadrature methods.[3]

Gaussian quadrature methods provide for specific abscissas and associated weights so that the function value can be specified as a weighted sum of function values where the function values are measured at the abscissas given. Given the weights one needn't use large samples with the weight 1/n. So long as the dimensions required are not large, say 2 to 5, the method can value very complex assets in seconds[4].

We can summarize the valuation method as follows:

$$E[Value_0] = \left[\sum_{i=1}^{ni} \sum_{j=1}^{nj} \sum_{k=1}^{nk} Payoff_{i,j,k} * f(x_i, x_j, x_k) \right] * e^{-rT}$$

And,

$$f(x_i, x_j, x_k) = \frac{1}{\sqrt{\pi}^3} (w_i * w_j * w_k)$$

Where *n* represents the number of evaluation points chosen for each dimension, *w* represents weights from the Gauss Hermite polynomials, and x represents abscissas from the Gauss Hermite polynomial.

Again, the $Payoff_{i,j,k}$ is a unique, optimal payoff for the weighted triple, and is driven by all of the items discussed above (start cost, ramp costs, state values, and transitions). It may accept a shape and deploy a dynamic programming solution over the hours for each triple.

---

[3] We use the same Gaussian quadrature method for our 2-factor spread options as the method avoids the Kirk problem where the non-stochastic strike is introduced into the d1 term, which can cause problems when the strike is large relative to the underlyings.

[4] One can generally get good answers with say 20 to 30 points per dimension, in contrast to several hundreds of thousands or more required in the Monte Carlo.

V. Sample Outputs

The following table provides the reader with a detailed view of the outputs of the model. The results were generated by modelling a plant with 3 operating modes, ramp fuel between states, and other details. The results of the model also include (not shown) probabilistic turbine starts and turbine hours. With these results the model quickly shows the margin per turbine start or the margin per turbine hour and can assist in optimizing the starts.

		1/1/18	2/1/18	3/1/18	4/1/18
1	BlendedPremium	$ 3.147	$ 2.462	$ 4.494	$ 4.570
2	WkDay Margin	$ 780,883	$ 546,915	$ 1,062,247	$ 1,031,582
3	Wknd Margin	$ 155,716	$ 114,886	$ 275,185	$ 284,649
4	Total Margin	$ 936,599	$ 661,801	$ 1,337,432	$ 1,316,231
5	GasDelta	(0.7071)	(0.6555)	(0.8124)	(0.7982)
6	5x16 Delta	0.9907	0.9317	0.9835	0.9742
7	2x16 Delta	0.8864	0.7607	0.9109	0.8905
8	OffDelta	0.1863	0.2323	0.5567	0.5598
9	GasVega	0.0104	0.0144	0.0230	0.0278
10	OnVega	(0.0128)	(0.0168)	(0.0256)	(0.0311)
11	OffVega	(0.0005)	(0.0010)	(0.0034)	(0.0042)
12	GasGamma	0.0666	0.1000	0.0674	0.0695
13	OnGamma	0.0333	0.0692	0.0241	0.0291
14	OffGamma	0.1462	0.1343	0.1345	0.1176
15	GasOnXGamma	(0.0384)	(0.0756)	(0.0289)	(0.0320)
16	GasOffXGamma	(0.1256)	(0.1189)	(0.1280)	(0.1158)
17	OnOffXGamma	(0.0870)	(0.1056)	(0.2322)	(0.2302)
18	WkDay Generation	152,728	134,094	176,155	167,054
19	Wknd Generation	56,324	44,663	69,290	67,122
20	WkDay FuelBurn	1,069,571	939,227	1,233,111	1,169,496
21	Wknd FuelBurn	394,362	312,664	485,034	469,860
22	Total Fuel Burn	1,463,934	1,251,892	1,718,145	1,639,356
23	Turbine Starts	23.37	17.45	12.95	11.87
24	5x16 Gen	139,473	118,937	137,905	130,459
25	2x16 Gen	51,022	38,783	53,425	51,253
26	7x8 Gen	18,557	21,037	54,115	52,464

VI. Contact Information

We hope that our readers have found this brief article interesting. Please feel free to contact us with any questions or comments.

Mark Houldsworth: 832-453-8319 Rich Pastore: 832-545-0243