Valuation of a Spark Spread: an LM6000 Power Plant

Mark Cassano and Gordon Sick*

May 14, 2009

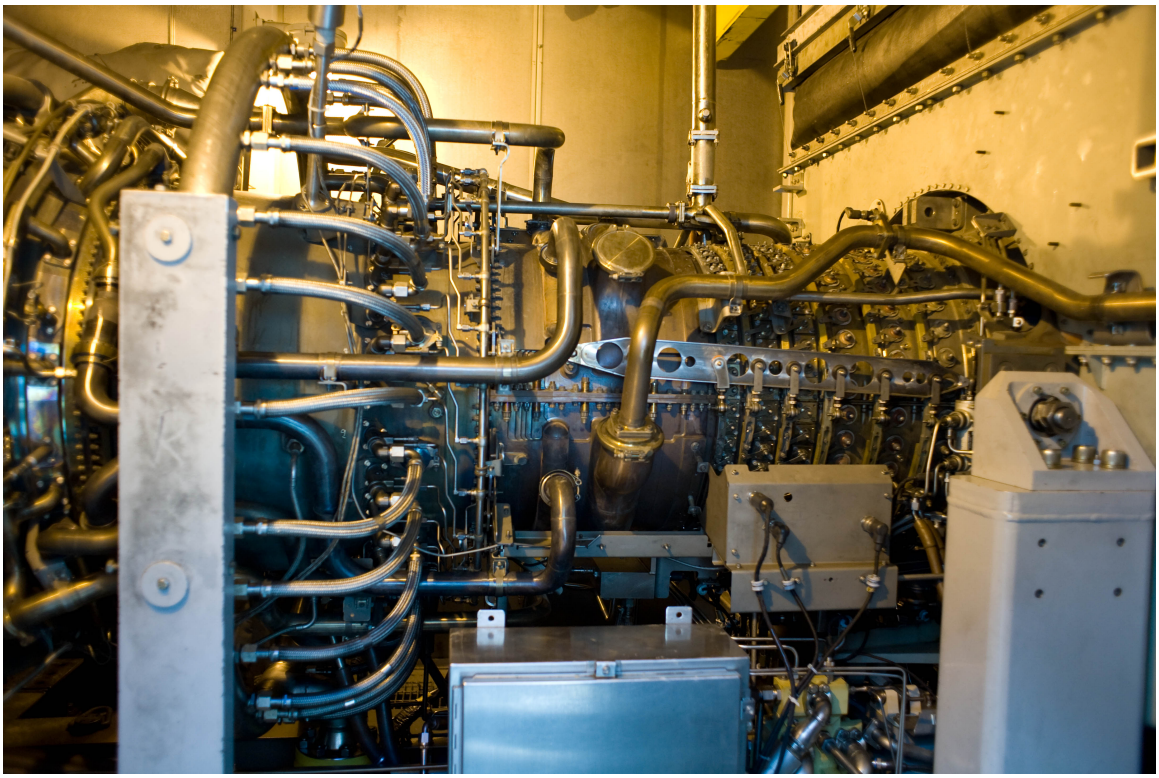

*Respectively, Independent Consultant and Professor of Finance at the Haskayne School of Business, University of Calgary. Please address all correspondence to gordon.sick@haskayne.ucalgary.ca . The analysis and data in this paper was prepared for a consulting client and we are grateful that they have allowed us to adapt our report in the form of this academic paper. We have modified the plant-specific results in Section 8 to protect the proprietary data and property of our client.

Contents

1 Executive Summary

This paper analyzes a power plant in Alberta, Canada, that is powered by two General Electric LM6000 gas turbines combined with a steam generator that allows combined cycle operations. The LM6000 is derived from a GE engine used on Boeing 767 and 747 airplanes, and is adapted for natural gas by General Electric. This power plant is popular in various power jurisdictions around the world as a turnkey power plant that can offer peaking capacity, and some baseload power delivery.

We consider 4 operating modes for the plant: cold metal (off), 15 MW idle in combined cycle, full simple-cycle power (95 MW) and combined cycle full power (120 MW).

It is common to refer to such a plant as generating a spark spread: converting natural gas to electricity by burning. A spark spread has two correlated stochastic variables: electricity price and natural gas price. To lower the dimensionality of the problem, we worked with heat rates. The market heat rate is the ratio of the electricity price to the natural gas price and the plant heat rate is the number of gigajoules (GJ) of gas needed to generate one MWh of electric power. This allows us to analyze the problem using the Margrabe approach, using natural gas as the numeraire commodity.

We estimate a stochastic model for market heat rates that incorporates time of day, day of week, month and the incidence or otherwise of a spike in heat rates. We use the model and its residuals in a bootstrap process simulating future market heat rates, and use a Least Squares Monte Carlo approach to determine the optimal operating policy.

We find that the annual average market heat rate is a good explanatory variable for the time-integral of the plant operating margin, denominated in the natural gas numeraire. This allows us to express plant values in terms of the numeraire and convert to dollars by multiplying this by the natural gas forward curve and a forward curve of riskless discount rates.

2 Option Valuation Based on the Market Heat Rate

The value of a natural gas based power plant centers on the spread or difference between the prices of electricity and gas (often termed the spark spread); in fact, one can look at a gas powerplant as being a call option on the spark spread (referred to in the finance literature as a “spread option"). The operator of the plant has the option to convert a given number of gigajoules (GJ) of natural gas, at the plant’s operational heat rate, into a number of megawatt-hours (MWh) of electricity. Being able to generate power strategically — only when it is profitable — is a significant source of value and the value of this flexibility can be found by applying modern financial derivative pricing techniques. Since the power plant converts gas to electricity, the “unit of account" for the exchange is GJ of natural gas. This is in contrast to the typical way of discussing options which is the right to exchange a quantity of dollars for a unit of the underlying

asset. For example, a standard call option on a stock is the right to buy the underlying stock for a fixed number of *dollars*. A spread call option on a stock could give the holder a right to buy a share of stock for a fixed number (e.g. say 2) shares of another stock. In fact, Margrabe (1978) exploited this framework by valuing spread options (or the option to exchange one asset for another) by first defining all values in units of the asset being acquired in the exchange. This leads us to modeling the price of electricity in units of natural gas: the Market Heat Rate.

The idea behind the Margrabe model is similar to that of the valuation of an international project. In many cases, an international project is analyzed by modelling all the cash flows in units of foreign currency and multiplying the expected free cash flow for a future period by the forward exchange rate for that period and a domestic discount factor. In the spark spread problem at hand, the foreign currency is GJ of gas, rather than US$ or Euros. The Margrabe model determines the value of each year of operations in units of GJ of gas. One then multiplies by the expected future price and discounts at an appropriate risk-adjusted discount rate.

3 Electricity Prices and Natural Gas Prices

In Alberta, Canada, there is "postage-stamp pricing" for electricity and natural gas. That is, the wholesale price of natural gas, in $ per gigajoule (GJ), is set on a daily basis (often called the AECO or NIT price) in a competitive market, and the same price prevails for all delivery points in Alberta. The wholesale price of electricity, in $ per megawatt-hour (MWh), is set on an hourly basis (often called the System Marginal Price or SMP) in a competitive market, but the same price prevails throughout the province.

Alberta is a Province with an abundant supply of natural gas and coal. Most electric power is generated by large baseload coal-fired plants, which are very efficient, but not very flexible in their operations. Thus, they have a relatively inelastic supply of electricity. The Province is east of the Rocky Mountains, and there is much less rainfall than on the western side of the mountains. Thus, there is very little hydro power available, which is a problem, since hydro power is a major source of price-responsive peaking power. The only other major source of peaking power is a gas-fired turbine, such as the GE LM6000 that we study in this paper.

The Alberta power grid is thus prone to price spikes, since there is so little price-responsive supply. One source of spikes is unplanned outages from the large baseload generators. Elliott et al. (2002) studied the stochastic process for electricity prices in Alberta and modeled the price spikes with a regime-switching model where the regimes were defined by the number of baseload generators in operation. They modelled the number of baseload generators in operation as a Markov chain.

3.1 Defining the Market Heat Rate

Let S_i^E (in \$ per MWh) and S_i^G (in \$ per GJ) be the electricity and gas prices, respectively, for observation hour i. The *market heat rate* is defined as the ratio of the two, $Y_i := S_i^E / S_i^G$. It is the natural gas value of electricity; specifically, how many GJ of natural gas to acquire one MWh of electricity. If the acquisition is done in the market, it is a market heat rate. As an example, assume the gas price is \$6.25/GJ and electricity is \$40/MWh. The market heat rate would be $40/6.25 = 6.40$ GJ/MWh. That is, one MWh or electricity is worth 6.40 GJ of natural gas.

We also speak of a *plant heat rate*. It is also the number of GJ of natural gas that is needed to produce one MW of power. In this case, the acquisition is done by burning the natural gas in the plant to produce electricity. Continuing our example above, a power plant that has an operational heat rate greater than 6.40 GJ would be "out of the money" — the amount gas required to generate a MWh is more than the quantity of gas that a MWh of electricity is worth in the market. A power plant that is inflexible may not be able to be turned off economically when the plant is out of the money.

As we will see, for modeling and valuation, it is necessary to define the natural logarithm of the market heat rate as $y_i := \ln(Y_i) = \ln(S_i^E / S_i^G)$.

Figures 1 and 2 show the heat rate and log heat rate[1] in the Alberta market for the years 2002-2006. The central feature of the graph, and the source of the option-like value of the power plant, are the spikes. These spikes typically do not last for long, as their causes (such as weather or generator outages) are likely temporary. In addition, at high heat rates, supply responds as less efficient generation becomes economical.

Figures 3 and 4 present the distribution of historical heat rates and log heat rates.[2] We find the majority of heat rates below 20, but there is positive skewness as there are a significant number of spikes (e.g. heat rates greater than 20). The log heat rate is more symmetrical .

Table 1 presents the averages and standard deviations for the market heat rate (Y) and the log of the market heat rate (y) for the years 2002 to 2006. As one can see, there is a "U" shape to the average heat rates; 2006 having the highest average heat rates, and 2004 and 2005 having the lowest market heat rates. As we will discuss in Section 4, spikes are present and these spikes make standard deviations harder to interpret. This is shown by the fact 2002's heat rate has the second largest standard deviation but the lowest log standard deviation.

[1] The values when the log market heat rate is less than -4 are observations when the electricity price was equal to zero. Since the log of zero is undefined, a value of 0.01 was used.

[2] The final bin of 100 in Figure 3 are heat rates greater than 100.

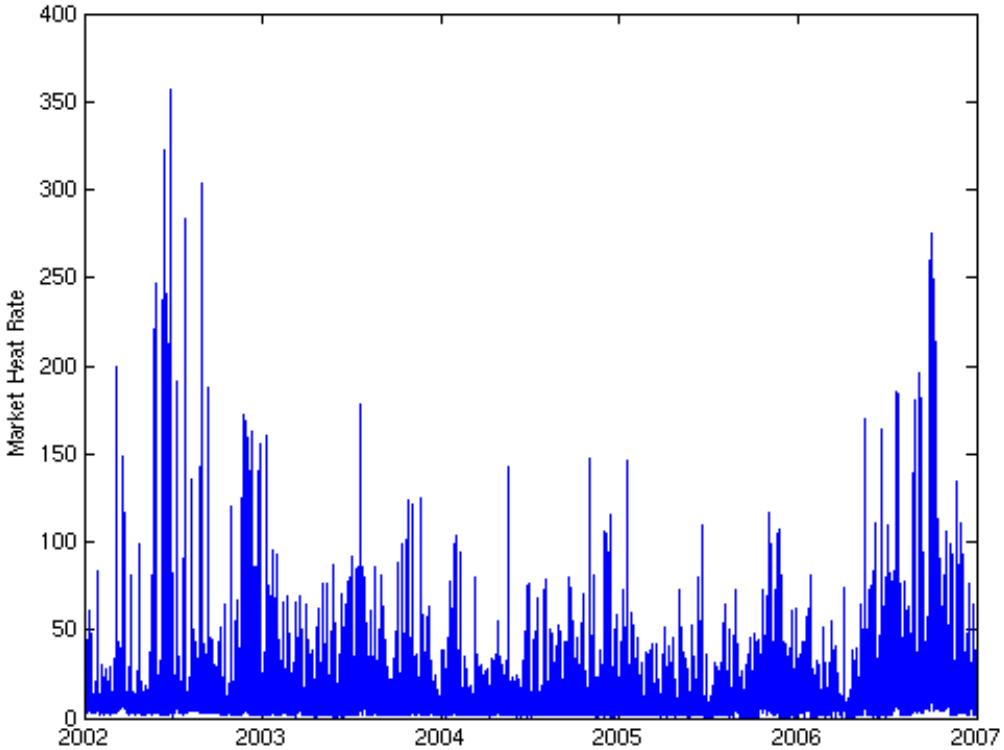

Figure 1: Historical Alberta Hourly Market Heat Rate: 2002-2006

Table 1: Summary Statistics of the Market Heat Rate

Year	Mean Y	Std Y	Mean y	Std y
2002	11.5830	17.7420	2.1013	0.7690
2003	10.0830	11.1920	1.9660	0.8165
2004	8.7909	8.4864	1.8766	0.7945
2005	8.2296	8.9049	1.7418	0.8685
2006	14.0070	23.4530	2.1460	0.8682
Overall	10.5380	15.2490	1.9663	0.8374

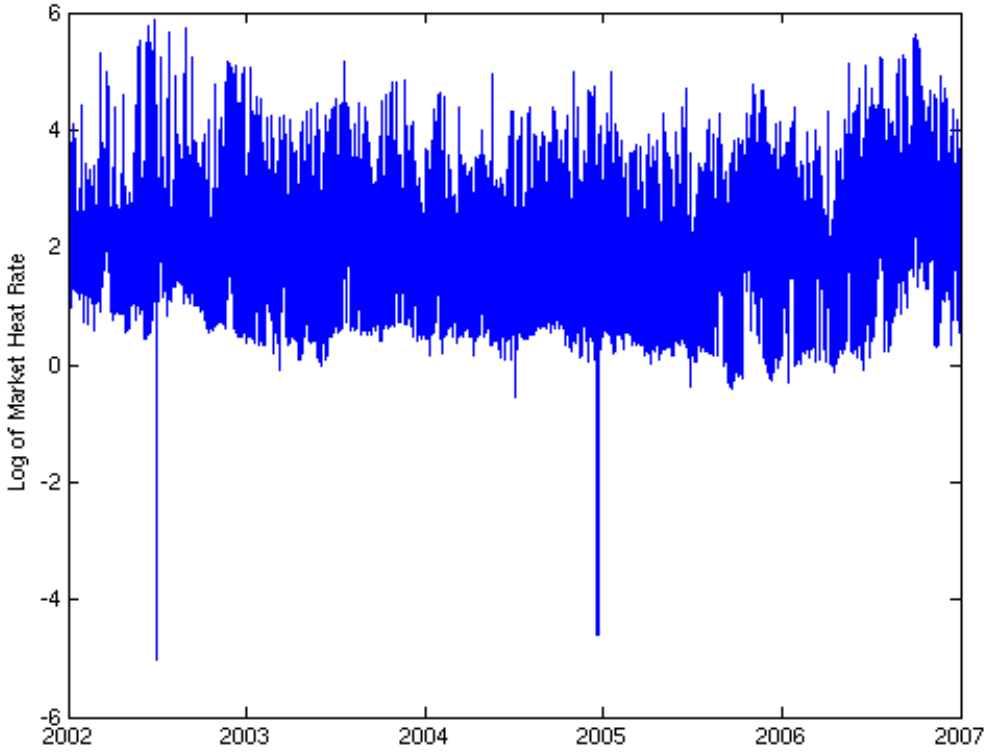

Figure 2: Historical Alberta Hourly Log Market Heat Rate: 2002-2006

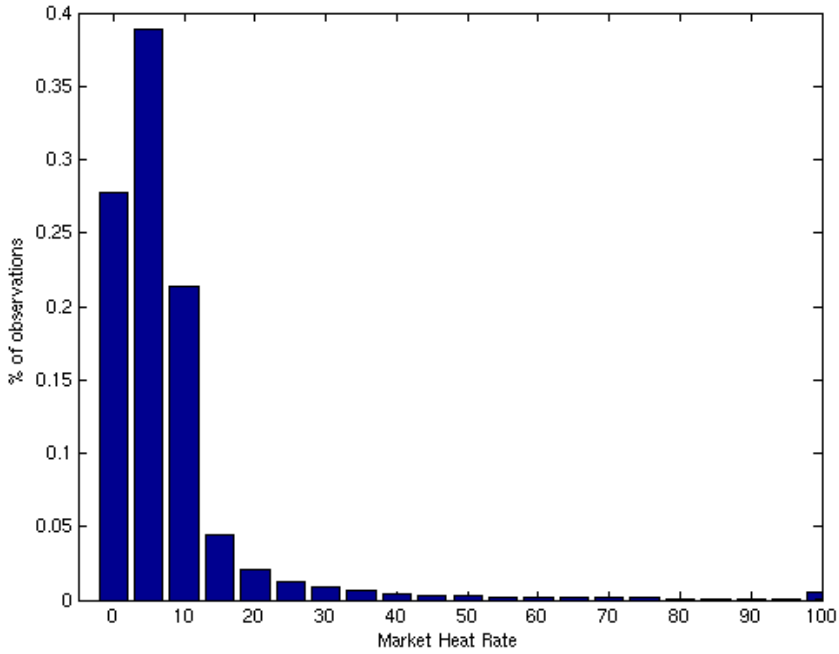

Figure 3: Histogram of Alberta Hourly Market Heat Rate: 2002-2006

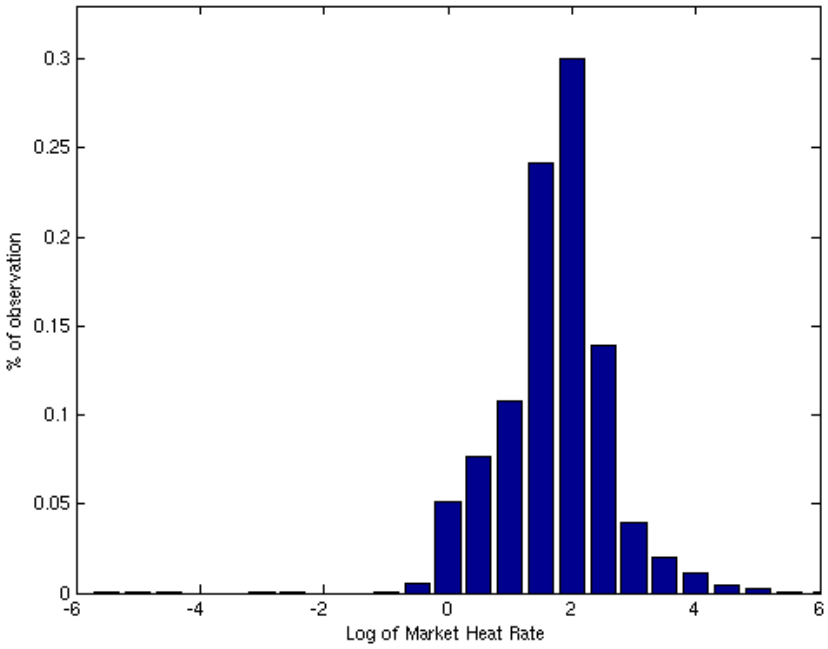

Figure 4: Histogram of Alberta Hourly Log Market Heat Rate: 2002-2006

3.2 Valuation Issues and the Margrabe Approach

There are two standard ways of valuing future cash flows (i.e. computing "Present Values" or PV):

1. The traditional approach using risk-adjusted discount rates (RADR) or cost of capital.

2. The approach used in the financial derivatives literature that uses certainty-equivalents (CE) and riskless discount rates.

In addition, valuation of business opportunities that exchange input goods for output goods that both have risky value requires the special technique developed by Margrabe. It can be coupled with the certainty-equivalent approach and other option-pricing methods.

Below, we discuss these three issues.

3.2.1 The Risk-Adjusted Discount Rate Method

The traditional capital budgeting and valuation models use risk-adjusted discount rates. The analyst calculates expected future cash flows and discounts these at a RADR that is obtained from a cost of capital model such as the Capital Asset Pricing Model (CAPM). Mathematically, a single future cash flow of CF_T received T years from today has a value of $\mathrm{PV} = e^{-(r+\lambda)T}\mathbb{E}[\mathrm{CF}_T]$ where $\mathbb{E}[\mathrm{CF}_T]$ is the expected value of the cash flow, r is the continuously compounded risk free rate (which would be the zero coupon continuous yield of a government bond with a maturity of T years), and λ is a "risk premium" for the cash flow.[3]

For our purposes, the future cash flow (operational profit or gross margin) is some function of the future market heat rate; i.e. $\mathrm{PV} = e^{-(r+\lambda)T}\mathbb{E}[f(y_T)]$ where $f(\cdot)$ is a function that comes from some value-optimizing decision rule (which will be determined).[4]

The difficulty with the RADR method is that high operating or financial leverage makes the appropriate risk-adjusted discount rate a random number because the leverage ratio changes rapidly with the risky revenue stream. Financial options achieve financial leverage from the fixed exercise cost. Operating leverage comes from fixed operating costs. This is discussed, for example in Hodder et al. (2001); Sick (2007).

[3] In the CAPM, $\lambda = \beta_{\mathrm{CF}}(\mathbb{E}(r_m) - r)$ where $\mathbb{E}(r_m)$ is the market's (e.g. S&P500) expected return and β is the cash flow's "Beta". Beta is the slope of a statistical regression of the cash flow on the market index. It can also be computed as

$$\beta_{\mathrm{CF}} = \frac{\rho_{\mathrm{CF},m}\sigma_{\mathrm{CF}}}{\sigma_m}$$

[4] We can use y_T without loss of generality, because any function of Y_T, say $g(Y_T)$, is a function of y_T since $f(y_t) = g(e^{y_T})$.

The option to exchange assets (gas for electricity) applies to risky cash-flow streams, but it generates leverage, as well. That is, leverage does not require fixed financial or operating costs — it arises as long as the costs are not perfectly correlated with the revenues.

Thus, the traditional RADR method puts the risk-adjustment in the denominator of the PV formula, but assumes that the RADR is constant. But, in fact, the RADR is continually changing because of the leverage. This leads to valuation errors which are excessive in highly levered situations, such as options and swap transactions.

3.2.2 The Certainty-Equivalent Method

Options and derivatives analysts value assets by incorporating risk in the expectation of the cash flow in the numerator of the PV formula.[5] This risk-adjusted expectation is a certainty-equivalent. The CE is discounted at the riskless discount rate. In other words, the risk premium is embedded in the expectation of the underlying asset and not the payoff function (which depends on the future value of the underlying asset). This is often termed *risk-neutral valuation* and is the key to option valuation such as Black-Scholes. Using our market heat rate payoff, risk-neutral valuation becomes $\text{PV} = e^{-rT}\mathbb{E}^*[f(y_T)]$ where $\mathbb{E}^*[f(y_T)]$ is the expected payoff after the market heat rate's expected value has been changed to incorporate risk. Notice, that this holds for all types of functions.

While certainty equivalents were merely a glimmer in finance professor's eyes 40 years ago, they are now commonplace in the derivatives industry. Futures and forward prices are certainty-equivalents of the uncertain spot price that will prevail at the maturity of the contract. As just discussed, option valuation is done by taking risk-adjusted (equivalently risk-neutral) expectations, which are also certainty equivalents.

3.2.3 The Margrabe Approach

The Margrabe approach is useful for valuing spread or exchange payoffs, where a good of uncertain value is exchanged or converted into another good of uncertain value. To understand the Margrabe approach, we first discuss a more common situation of capital budgeting for a foreign investment.

The discussion above has everything denoted in Canadian dollars (CAD). If the cash flows are in British Pounds (£), we must multiply all Pound values by appropriate exchange rates and also discount at the Canadian interest rate, in order to compute a value in CAD. For example, if the Pound cash flow ($\text{CF}_{\$}$) is in 2 years, we'd could convert this future Pound cash flow into CAD at the 2 year forward exchange rate, and then discount at the Canadian risk free interest rate. If F_T is the two year forward rate (in

[5] More technically, the probability distribution is altered by changing the "drift" of the underlying stochastic process.

CAD per £) and r_{CAD} is the canadian interest rate, the PV in CAD is $\text{CF}_{\$} \times F_T \times e^{r_{\text{CAD}}T}$. Alternatively, we could arrive at a Pound PV by discounting the Pound cash flow at the British interest rate and then convert it to CAD at the current (spot) exchange rate (call it $S_{\$}$). The PV in CAD would be $\text{CF}_{\$} \times e^{r_{\text{UK}}T} \times S_{\$}$. Both approaches arrive at the same answer because the interest rate parity arbitrage in foreign exchange markets implies $F_T = S_{\$}e^{(r_{\text{CAD}}-r_{\text{UK}})T}$. To summarize, the two ways to compute PV in CAD when the future cash flow is in British Pounds are:

- Use the T year forward exchange rate (CAD per £) and discount using the Canadian interest rate.
- Use the spot exchange rate (CAD per £) and discount using the British interest rate.

The previous discussion is relevant for our valuation in that there is nothing inherently special about currencies such as the British Pound. In fact, the key to our Margrabe valuation approach is putting values in terms of *natural gas* rather than British pounds or Canadian dollars. The future "cash flow" is in GJ of natural gas.[6] Substituting "£" with GJ of natural gas we have the two ways of converting the future value in GJ of gas to the PV in CAD of our power plant:

- Use the T year forward or futures price of natural gas (CAD per GJ) and discount using the Canadian interest rate.
- Use the spot price of gas (CAD per GJ) and discount using the natural gas "interest rate".[7]

For our valuation model, we chose the first method. If a plant operates optimally for a year (from year T to $T+1$), suppose the future gross margin is v_T GJ of gas.[8] Hence the PV in CAD is $V_0 = e^{-rT}F_T\mathbb{E}_0^*[v_T]$.[9]

[6] One could think of it as a "gas flow".

[7] This is termed the convenience yield, δ_G which is derived implicitly from the "cost of carry" formula for the natural gas futures curve $F_T = S^G e^{(r-\delta_G)T}$. We can also put the convenience yield in terms of the gas risk premium and expected growth rate in gas prices (denoted μ_G) by the CAPM relationship $\delta_G = r + \lambda_G - \mu_G$. Substituting this CAPM relation into the cost of carry formula and using some algebra gives

$$F_{0,T} = S_0 e^{(\mu_G-\lambda_G)T} = \mathbb{E}[S_T]e^{-\lambda_G T} \tag{1}$$

Thus, for example, if there is no risk premium, $\lambda_G = 0$, and the futures price equals the expected spot price.

[8] Technically, it is a function of the expected future plant heat rate; this will be determined by the user under optimal operation. This will be discussed in more detail shortly.

[9] Technically, v_T is an expectation of the operating margin over one year's time under an optimum operating rule. The law of iterated expectation implies that we can use the expected optimal future plant heat rate to arrive at $\mathbb{E}_0^*[v_T]$.

We assume that, for the market heat rate, risk-neutral expectations are equivalent to actual expectations. This will be correct if the market risk premium for electricity is approximately equal to that of gas, because the risk premia cancel each other in the exchange option.

Finally, we should add that instead of using $e^{-rT}F_T$ to discount gas futures prices to current dollars, one could also use gas price forecasts (i.e. expected actual future gas prices, or "price decks") and risk adjustments as in:

$$V_0 = \mathbb{E}[S_T]e^{-(r+\lambda_G)T}\mathbb{E}_0^*[v_T] \tag{2}$$

where λ_G is the annual market risk premium for gas.

In equation (2), the first factor $\mathbb{E}[S_T]$ is the expected gas price or price deck. The second factor $e^{-(r+\lambda_G)T}$ is the PV factor at the risk-adjusted discount rate or cost of capital of $r + \lambda_G$. The third factor $\mathbb{E}_0^*[v_T]$ is the risk-neutral expectation of the plant operating margin for year T, expressed in GJ of gas for delivery in year T rather than dollars for delivery in year T.[10]

4 Modeling and Estimation of the Market Heat Rate

In order to determine the risk-neutral expected natural gas payoff of a power plant, we must first model the random nature of the market heat rate. At first, it might seem appropriate to model electricity and gas prices separately; however, this is unnecessarily complicated and troublesome.[11] We chose to examine the historical behavior of the market heat rate and model the market heat rate directly, following the Margrabe approach.

The most common models used for valuation are based on Brownian motion, which makes the underlying processes normally and log-normally distributed. The spikes present in electricity prices[12] require a more complicated model; the empirical distributions do not come close to resembling a normal or log-normal distribution.[13] The most appropriate model class is that of "regime switching" models. These models have the underlying price switching randomly between two or more different processes. As an example, one regime could be considered "normal" and the other regime would be a "spike" regime that could occur due to a demand or supply shock such as abnormal weather or a base load generator outage. Each hour, the market heat rate is in one of

[10] Other authors have employed the Margrabe model to value generators as spark spreads. See, for example, Carmona and Durrleman (2003).

[11] Without a rich set of futures prices, natural gas prices are especially difficult to model accurately.

[12] Natural Gas does not have "spikey" behavior, hence spike in electricity leads to spikes in the market heat rate.

[13] Additionally, if one was to try to fit Geometric Brownian motion to the market heat rate, the volatility would be around 5000% which makes call options almost equal to the underlying price in the Black-Scholes model

the regimes and behaves according to the stochastic rules of the given regime. There must also be a stochastic “rule" to model which regime the market rate heat will belong to in a given hour. Hence we must estimate the following:

- The stochastic characteristics of the “normal" regime - we’ll name the normal regime Regime 1.
- The stochastic characteristics of the “spike" regime - we’ll name the spike regime Regime 2.
- The “switching" process which is the nature of how likely we are to switch from one regime to another.

The data set we used consisted of hourly Alberta Price pool electricity prices and daily AECO gas prices from 2002 to 2006.[14] We chose a simple and transparent way to identify regimes in the data . We define the spike regime as occurring when the market heat rate is greater than 20 (i.e. $Y > 20$ or $y > \ln(20) = 2.996$).[15] For each of Regime 1 and Regime 2, we use linear regression to estimate the following equations; for $j = 1, 2$,

$$y_i = \beta_j\,(\text{hour of day}, \text{day of week}, \text{month of year}, \text{year}) + \gamma_j y_{i-1} + \varepsilon_{i,j} \tag{3}$$

We’ve chosen, for ease of notation, to suppress the fact there are different β coefficients for each of the regime equations.[16] It is important to note the lagged observation, y_{i-1}; this will be able to capture any temporary non-regime-switching shocks. Additionally, the lagged variable also eliminates serial correlation in the regression residuals, ε_j. This is an important quality when using these residuals for bootstrapping in the Monte Carlo study.

For the regime switching process, we specified a logistic regression. With the definition of a spike given, define the binary variable $R_i = 1$ if a spike occurs at hour i and $R_i = 0$ otherwise. Then use a logistic regression to find the best fit minimizing the least squares distance between the observed R_i and the logistic function $\hat{R}$ where $\hat{R} = e^{z_i}/(1 + e^{z_i})$ and

$$z_i = \beta_r\,(\text{hour of day}, \text{day of week}, \text{month of year}, \text{year}) + \psi R_{i-1} \tag{4}$$

Logistic regression is a numerical way of estimating the probability of each regime in the data observed; for example, if spikes occur more in October, R_i will be more likely to be a 1 and the logistic regression coefficient for an October dummy variable will be positive and significant. The results of our analysis are given in Table 2.

[14] Empirically, one could estimate which regime the process is in by using generator data (i.e. aggregate wind and load based coal generation), as in Elliott et al. (2002). However, this strategy fails to capture a significant proportion of spikes, making additional regime modeling necessary.

[15] We looked at other spike definitions and determined this was the most reasonable specification in that the simulations from this specification resembled the actual historical behavior best.

[16] The estimation was done with dummy variables so there are no regression coefficients for Hour 1, Sunday, January, or the Year 2002.)

Table 2: Estimated Regime and Regime Switching Probability models

Variable	Regime 1	Regime 2	Logistic Regression
Constant	0.3685	2.6246	-4.6182
Hour 2	0.1679	0.0000	0.0000
Hour 3	0.0845	0.0000	-0.5733
Hour 4	0.0806	0.0000	-0.8620
Hour 5	0.1342	0.0000	-0.9841
Hour 6	0.2608	0.0000	0.0000
Hour 7	0.4362	0.2554	1.4019
Hour 8	0.2578	0.0000	0.6749
Hour 9	0.5941	0.0000	1.2641
Hour 10	0.5478	0.0000	1.2453
Hour 11	0.5640	0.1074	1.6926
Hour 12	0.5097	0.0000	1.7470
Hour 13	0.4363	0.0000	1.2879
Hour 14	0.4486	0.0984	1.6905
Hour 15	0.4372	0.1103	1.5376
Hour 16	0.4505	0.0843	1.6950
Hour 17	0.5530	0.1437	2.2390
Hour 18	0.4877	0.1267	2.1675
Hour 19	0.3023	0.0000	1.1408
Hour 20	0.4112	0.0000	1.6318
Hour 21	0.4715	0.0000	1.4570
Hour 22	0.4953	0.0000	1.7075
Hour 23	0.0748	0.0000	0.0000
Hour 24	0.3056	0.0000	1.6465
Mon	0.0728	0.0000	0.3061
Tues	0.0701	0.0000	0.1144
Wed	0.0694	0.0000	0.0000
Thurs	0.0830	0.0000	0.0000
Fri	0.0970	0.0000	0.0000
Sat	0.0422	0.0000	0.0000
Feb	-0.0238	0.0000	-0.4007
Mar	0.0000	0.0000	0.0000
Apr	-0.0350	0.0000	-0.3650
May	-0.0212	0.0000	0.0000
Jun	-0.0557	0.0000	0.0000
Jul	0.0221	0.0000	0.4321
Aug	0.0522	0.0000	0.0000
Sep	0.0264	0.0000	0.0000
Oct	0.0377	0.1020	0.5759
Nov	0.0551	0.0000	0.2553
Dec	-0.0618	0.0992	0.0000
2003	-0.0829	0.0000	0.2633
2004	-0.0829	-0.0647	-0.1657
2005	-0.1488	-0.0796	0.0000
2006	-0.0424	0.1339	0.5156
Lag[17]	0.5978	0.2984	2.6853

4.1 Example using the estimation results

As an example of how to use Table 2, assume it is Hour 11 (HE 11) on a Wednesday in October 2007 and the current market heat rate is 8. Note that $\ln(8) = 2.0794$. Since $8 < 20$, we are in Regime 1. If we use the stylized year 2002[18] to determine the the expected value of the log market heat rate in the next hour (Hour 12), we add the appropriate coefficients to the constant term. For example, the expected log market heat rate for regime 1 is $.3685 + .5097 + .0377 + .5978 \times \ln(8) = 2.1591$. Given that there is volatility, the expected price is not $e^{2.1591} = 8.66$. If the heat rate was normally distributed, the expected heat rate would be $e^{2.1591+\sigma^2/2}$ where σ is the standard deviation of the log heat rate.[19] Given that the historical standard deviations are approximately 0.8, this would imply an approximate mean heat rate of $e^{2.1591+.8^2/2} = 11.93$. A similar exercise for regime 2 has the expected log market heat rate being 3.3471 and the expected heat rate being approximately 39.13.

Given that we are in regime 1, the probability of being in a spike regime in the next hour (Hour 12) is determined by the logistic equation. Using the logistic regression, we repeat a similar coefficient as above, arriving at a total of $z = -4.6182 + 1.747 + .5759 = -2.2953$. This is converted to a probability using the logistic function $e^z/(1+e^z)$, hence the probability of being in a spike regime is $e^{-2.2953}/(1+e^{-2.2953}) = 9.51\%$. This implies that the expected market heat rate would be equal to the expected heat rate across the two regimes: $11.93 \times .905 + 39.13 \times .095 = 14.42$.

If the market heat rate in Hour 11 was much higher we alter the previous analysis by changing the lagged variable. For example if the market heat rate was 50 (where $\ln(50) = 3.9120$), we would expect log market heat rates of 3.2547 and 3.8940 for regimes 1 and 2, respectively. This gives approximate heat rates of 35.68 and 67.62 for the two respective regimes. The likelihood of staying in regime 2 would be 59.63% implying an expected heat rate of 54.73.

To show how the stylized year affects these results, assume the stylized year was 2006 (the "spikiest" year). The market heat rate of 8 results in expected log heat rates of 2.1167 and 3.481, and expected heat rates of 11.44 and 44.75, in the Regimes 1 and 2, respectively. Also, there is a 14% chance of Regime 2 (spike) in the next hour. For a market heat rate of 50, we find expected log heat rates of 3.2123 and 4.078 and expected heat rates of 34.20 and 77.32 for the two regimes, with a 71% chance of continuing the spike.

[17] Lagged log heat rates for regime regressions; lagged regime state for logistic regression.

[18] The difference in the years is the average heat rates and spike probability.

[19] The extra adjustment $\sigma^2/2$ is sometimes called the Itô adjustment. It arises because we are taking a non-linear exponential transform of an uncertain variable. Another way to notice this is that with the non-linear convex exponential transformation, the expected value of the exponential transform is greater than the exponential transform of the expected value, given Jensen's inequality.

5 Monte Carlo Simulations

We use the *Least Squares Monte Carlo* (LSM) method,[20] which relies on simulating hourly heat rates of the power plants; this section describes the simulation method used. One often uses a pre-specified distribution such as normal or log-normal and matches the distribution's parameters to the data. However, for market heat rates, there is no obvious distribution that adequately describes the salient features of the data.

To simulate heat rates, one could try to search for a more appropriate parameterized distribution from which to draw random numbers from; however, a popular and effective simulation method used, *bootstrapping*, uses the actual residuals from the regressions to create the simulations. Bootstrapping uses the residuals from the two regime regressions, $\varepsilon_{i,j}$ from equation (3) to generate the random characteristics of the market heat rate. The specifics of a one-year simulation are as follows:

- An initial market heat rate, hour, day of week, month, and stylized year are specified. We used an initial market heat rate of 10, Hour 1, Monday, January for all of the simulations.[21]
- The logistic regression is used to simulate the regime of the first hour (Hour 1). This is done by drawing a uniform random number (i.e. a random number between 0 and 1, which is the typical way a computer generates a random number) and comparing this number with the probability of being in the spike regime which is determined as described in Section 2. If the simulated uniform number is less than the spike probability, then there will be a spike (i.e. hour 1 will be in regime 2). If the uniform number is greater than the spike probability, then the regime will be normal (i.e. regime 1).
- Once the regime is determined, the market heat rate comes from the expected log price regression and drawing randomly from the residuals of the relevant regime regression.
- This process is then repeated. The hour, day of week, month is adjusted appropriately; the new regime and heat rate are now used in the 3 estimated equations.

Figures 5 and 6 shows simulated runs using this bootstrapping method for 4 stylized years.[22] Comparing these figures with Figures 1 and 2 the random nature of the simulated heat rates bear a strong resemblance that of the historical market heat rates. This is important for the option value of the power plants depend crucially on the dynamic behavior of the market heat rate, the most important being the spike levels and durations.

[20] See, for example Longstaff and Schwartz (2001).

[21] The nature of the simulations (and the subsequent valuations) is not affected by the the initial market heat rate and day of week.

[22] 2004 simulations looked similar.

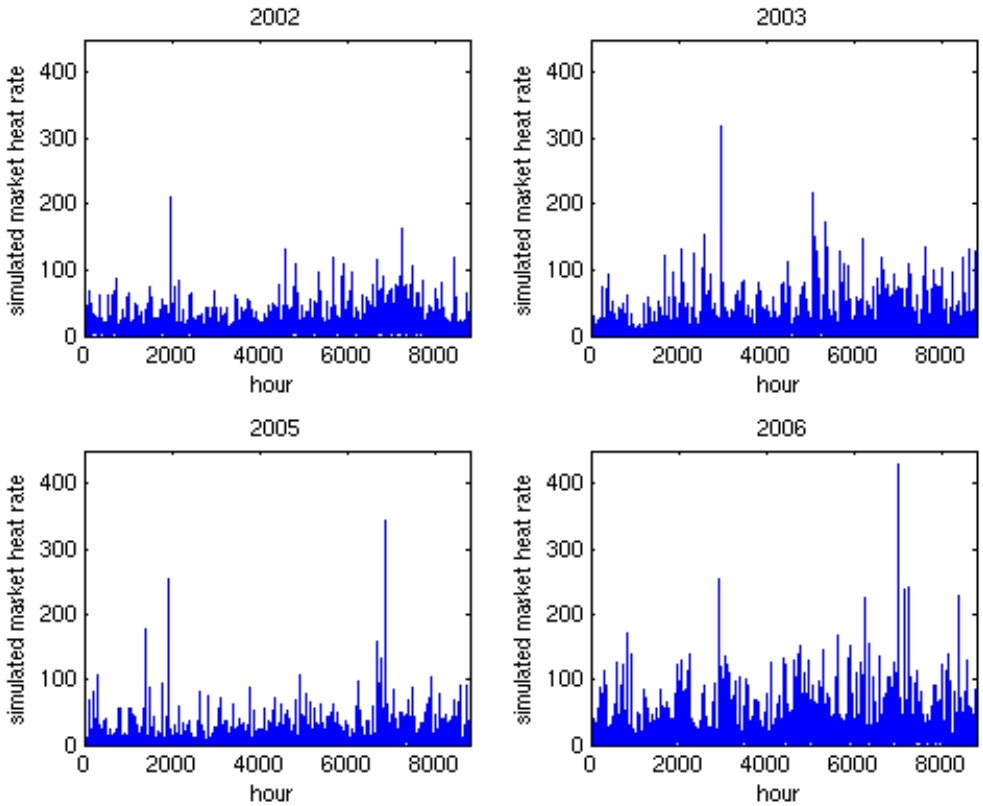

Figure 5: Simulated heat rates for four stylized years.

These figures were a single representative simulation for each year. We ran 100 simulations for each stylized year and looked at the distribution of heat rates and log heat rates. Figures 7 and 8 present the results. We see that 2006, besides having a higher average heat rate, has a much more dispersed distribution. This is due to the higher spike probability. The stylized year 2005 distribution is the most “bearish" distribution. When we perform our valuation exercise, it is these 5 possible stylized years that we use. Finally, Figures 9 and 10 look at the distribution of the simulated heat rates and log year rates of the combined 5 stylized years. Once again, they do a good job of approximating the historical distribution of heat rates in Figures 3 and 4.

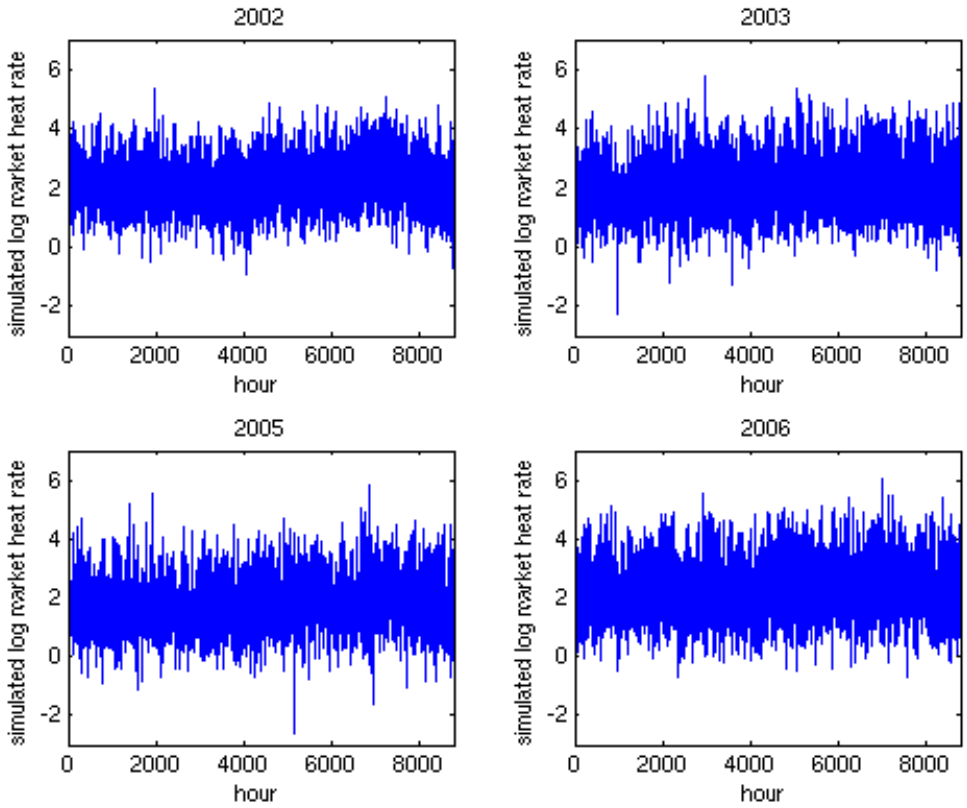

Figure 6: Simulated log heat rates for four stylized years.

6 Modeling the Operating Characteristics

An earlier version of this paper was developed for a consulting client that wishes to remain anonymous. We are grateful that they have permitted us to create a modified form of the report as a research paper. To protect their proprietary interests, we have modified the plant characteristics slightly with the effect that a precise valuation is not possible. However, the broad issues in the analysis are still correct. We have not modified the analysis in the previous sections, since they are based on publicly available data on market prices, although we are grateful to the client for having gathered it for us.

We obtained operating performance characteristics from a visit to a similar plant,[23]

[23] The plant we model here is not identical to the one we visited, but we were able to infer how the plant modeled here would operate, given our observations, since the differentiating characteristic can

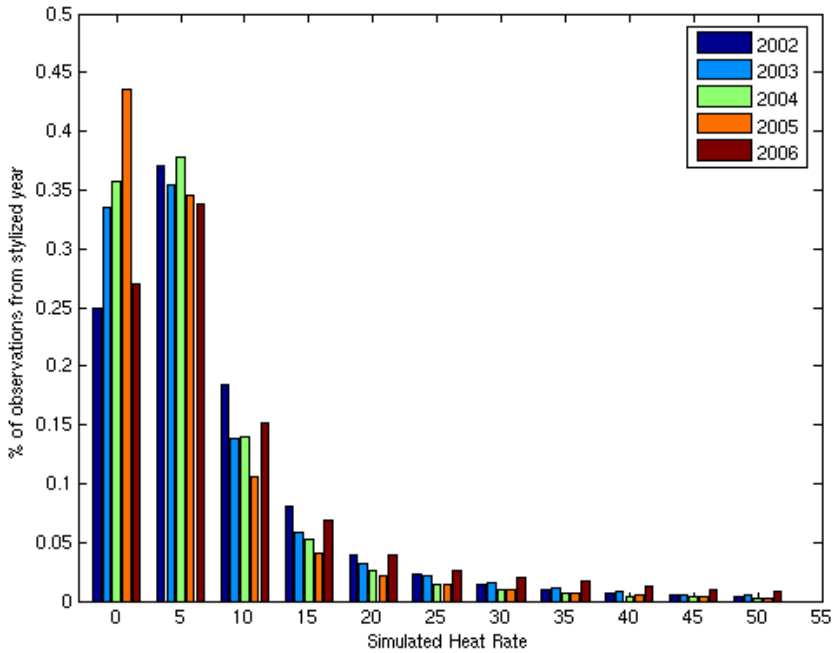

Figure 7: Distribution of heat rate from 100 simulations in each stylized year.

to view the operations and collect actual operating data (heat rates and output capacity on a minute-by-minute basis) for a ramp up and ramp down of one gas turbine in combined cycle mode. We understand that the ramp up and ramp down can be done more rapidly than in our sample, and it is often limited by the Power Pool System Operator, which does not like to see the plants increasing or decreasing output more rapidly than 5 MW/minute because they need time to adapt the system to the changes in generation. We modified the data to be consistent with ramp up and ramp down at 5 MW/minute, while still allowing for pauses for heat soak and other characteristics specific to the gas and steam turbines.

Our analysis only considered a restricted set of operating modes for the plant, which were consistent with optimal heat rates for a given range of outputs. We analyzed the plant as a merchant power plant that did not have any contract or obligation to provide Ancillary Services to the Power Pool. The actual operating data had the plants running in intermediate modes that we did not model. While these intermediate modes could represent suboptimal operating decisions, it seemed more likely that the intermediate operating modes arose from one of the following considerations:

be modeled out.

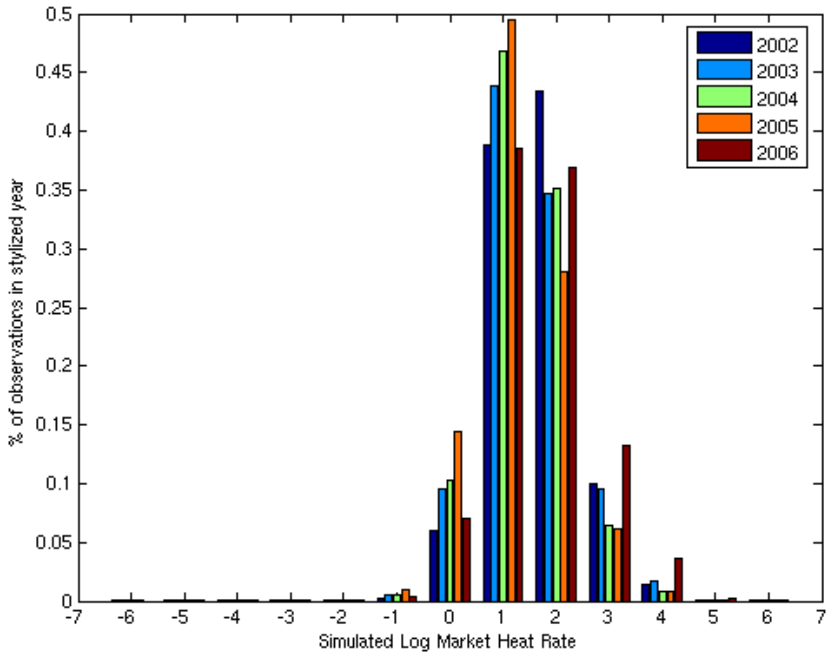

Figure 8: Distribution of log heat rate from 100 simulations in each stylized year.

- The plants could have been contracted to provide Ancillary Services (spinning reserve, for example) with the Power Pool. This generates revenue, and we did not collect data on this, nor did we examine the types of Ancillary Services contracts that were available. The decision to contract for Ancillary Services could generate more value than our merchant power model, but we were not in a position to assess this.

- Our visit to the plant led us to understand that the plant operators use more data to make their decisions than we used in our model. In particular, they pay attention to the whole merit order and the demand curve for electricity in the system. In many cases the plant is running at a heat rate that makes it the marginal supplier of power, so bidding in a second generator could depress the system marginal price (SMP) so much that the first generator would start losing money. Our model assumed a perfectly elastic demand curve, and that no other power suppliers would adjust their production in response to changes by our plant. In other words, we assumed a perfectly competitive electricity market. But, when the plant is at the margin (its operating heat rate is the market heat rate), optimal decisions by the plant operator should include consideration of the effect on SMP

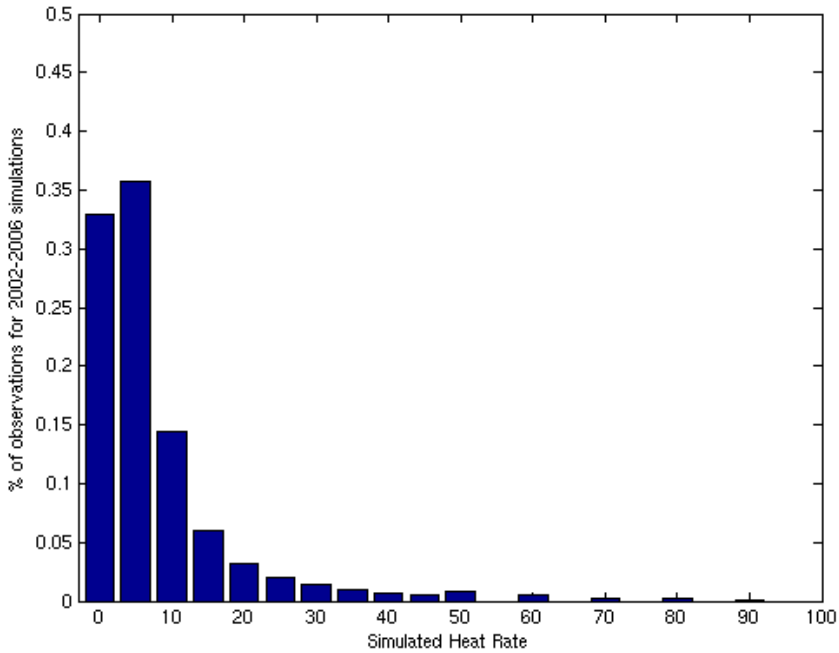

Figure 9: Distribution of heat rate from 500 simulations, 100 from each stylized year.

and this can result in intermediate operating modes.

6.1 Plant Operating Modes

We considered 4 operating modes:

Cold Metal was a mode where no fuel was burned and no power was generated.

15 MW Idle Combined Cycle was a mode where 1 gas turbine was synchronized and idling in combined cycle mode. In this mode, it burned fuel at 168 GJ/hour for a heat rate of 11.2 GJ/MWh. This mode allowed a fairly rapid ramp-up to full combined cycle output, but power produced in this idle mode is quite expensive, compared to typical market heat rates.

Full power from two simple cycle gas turbines generated 95 MW of power with a fuel burn of 931 GJ/hr and a heat rate of 9.8 GJ/MWh. The gas turbines ramp up more quickly than the steam turbine, so this mode capture price spikes before steam generation can start up.

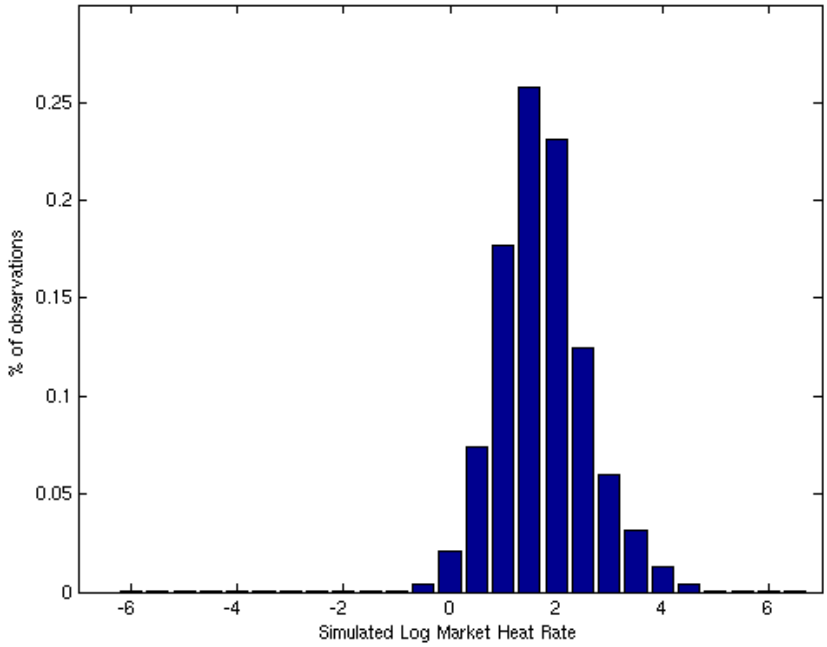

Figure 10: Distribution of log heat rate from 500 simulations, 100 from each stylized year.

120 MW Full Combined Cycle Output burned fuel at 938 GJ/hr for a heat rate of 7.90.

We also considered all possible transitions amongst these four modes, using the performance data from the plant visit. To ramp up steam production and enter combined cycle mode required 4 hours from cold metal or a gas turbine mode. Ramp down took 0.483 hours to 1.4 hours, depending on the transition. The ramp up from the 15 MW Idle Combined Cycle to full Combined Cycle output took 0.95 hours.

6.2 Sequential Ramp Up and Ramp Down

All of our transitions assumed that the operator would start up one gas turbine at a time. This sequential operation allows the operator to study the generator closely, to check for problems. In theory, it might be possible to start up to gas turbines in tandem, but the System Controller would not like such a sudden surge of power anyway.

7 Least Squares Monte Carlo Valuation

In order to value a power plant with flexible operating characteristics, one must determine the optimal operating policy of the plant. Entering into a given hour, the plant has a current operating mode. The manager uses all available information[24] to determine whether to have the plant remain in the current mode or switch to another mode. The manager must consider the cash flow for the next hour, plus the remaining value of the plant, given the operating mode in which the manager's decision leaves the plant. For example, the manager may be tempted to capture a spike at 9 pm, just as Hour 22 starts. This is an hour with generally high cash flows and a high probability of a spike, given the data in Table 2. The payoff could be lucrative for the next hour, but turning on the generator at this time would leave it in operation in the next hour, when value is likely to be low. Thus, the manager must consider not only the cash flow for the next hour, but also the "residual value" of the state that the decision gives the plant at the end of the hour. This residual value is the value of the cash flow for the remainder of the year, given the decision.

Each decision implies a "switching cash flow"[25] which is the value (in GJ of natural gas) of the electricity produced minus the amount of natural gas used during the switch, given the plant heat rate. This also includes the fractional part of the hour (if there is one) after the switch is completed. As an example, if the switch takes 40 minutes, the switching cash flow is the electricity produced minus the natural gas used during the 40 minutes of the switch plus the 20 minutes at which the plant is operating in the new mode. When there is no switch, the "switching" cash flow is simply the value of producing electricity at the current mode. The remaining value is the total value of operations for the remainder of the year from the hour after the switch has been made, given the mode that the manager into which the manager has put the plant and the market data on gas prices and electricity prices. This is a dynamic programming problem, which is solved recursively in the next subsection.

7.1 The Dynamic Programming Problem

Assume the plant operates for one year. The value of the plant in operational model m is a function of the hour h of the year and the current market heat rate. We denote this as $V_h(m, \mathcal{Y}_h)$. The solution method uses backward induction; we start at the final hour H in the year, and determine the switching cash flows going from mode m to all other possible modes. Let $\mathcal{M}$ be the set of operating modes and denote the switching cash

[24] Besides the hour, day of week, month, and stylized year, the relevant information is the current operating mode and market heat rate. We also assume the market heat rate describes the current regime.

[25] Although we use term "cash flow", it is actually a natural gas flow; i.e. the net amount of natural gas used in the switch.

flows of going from mode m to m' as $\pi_h(m, m')$.[26] As mentioned, if we remain in the current operating mode, the switching cash flow is the operating cash flow $\pi(m, m) = y_h - k_m$. For the final hour the value function is $V_H(m, y_H) = \max_{m' \in \mathcal{M}} \pi_h(m, m')$. We need the following notation; let $\tau(m, m')$ be the hours of time required to transition from operating mode m to m'.

Since our transition times are either less than one hour or equal to four hours, we can find the value functions for the penultimate hour, $H-1$ as follows. Given a current operating mode m and market heat rate y_{H-1} the value of changing the operating mode to m' is

$$v_{H-1}(m, m', y_{H-1}) = \begin{cases} \mathbb{E}\left[\pi_h(m, m') + V_H(m', y_H) \mid y_{H-1}\right] & : \quad \tau(m, m') \leqslant 1 \\ \mathbb{E}\left[\pi_h(m, m') \mid y_{H-1}\right] & : \quad \tau(m, m') > 1 \end{cases} \tag{5}$$

Note than if $m' = m$, the plant stays at the current operating mode. $\mathbb{E}\left[\bullet \mid y_{H-1}\right]$ is the expected value operator conditional on information available at hour $H-1$. The value function in hour $H-1$ is then simply

$$V_{H-1}(m, y_{H-1}) = \max_{m' \in \mathcal{M}} v_{H-1}(m, m', y_{H-1}). \tag{6}$$

The same thing is done for hours $H-2$ and $H-3$. For hour $H-4$ the values of the alternative operating modes are different, because a final 4-hour transition is feasible:

$$v_{H-4}(m, m', y_{H-4}) = \begin{cases} \mathbb{E}\left[\pi_h(m, m') + V_{H-3}(m', y_{H-3}) \mid y_{H-4}\right] & : \quad \tau(m, m') \leqslant 1 \\ \mathbb{E}\left[\pi_h(m, m') + V_H(m', y_H) \mid y_{H-4}\right] & : \quad \tau(m, m') > 1 \end{cases}$$

The value function for hour $H-4$ is $V_{H-4}(m, y_{H-4}) = \max_{m' \in \mathcal{M}} v_{H-4}(m, m', y_{H-4})$. This is procedure done for earlier periods $H-i$, where $i = 5, ..., H$. Specifically

$$v_{H-i}(m, m', y_{H-i}) = \begin{cases} \mathbb{E}\left[\pi_h(m, m') + V_{H-i+1}(m', y_{H-i+1}) \mid y_{H-i}\right] & : \quad \tau(m, m') \leqslant 1 \\ \mathbb{E}\left[\pi_h(m, m') + V_H(m', y_H - i + 4) \mid y_{H-i}\right] & : \quad \tau(m, m') > 1 \end{cases} \tag{7}$$

and

$$V_{H-i}(m, y_{H-i}) = \max_{m' \in \mathcal{M}} v_{H-i}(m, m', y_{H-i}). \tag{8}$$

The only remaining item for valuation is to determine the conditional expectations of the relevant values. Due to the complicated nature of the problem as well as the stochastic behavior of the market heat rate, an analytical or closed form solution is not feasible. Thus, we turn to the power of using Least Squares Monte Carlo (LSM) next.

[26] π_h will also depend on the current and, for transitions longer than an hour, expected future market heat rates. We suppress these heat rate arguments in this discussion.

7.2 Least Squares Monte Carlo

In order to compute the conditional expectations needed to solve the dynamic programming problem in the previous section, one can use simulation techniques. The idea is that a condition expectation such as $\mathbb{E}\left[\pi_h(m, m') \mid y_{H-1}\right]$ from equation (5), is a function of the current market heat rate (e.g. y_{H-1}) which can be approximated by a simple polynomial of the variable. Specifically, for some random variable X that is only revealed after time h,

$$\mathbb{E}\left[X \mid y_h\right] \cong \beta_0 + \beta_1 y_h + \beta_2 y_h^2 + \beta_3 y_h^3 + \beta_4 y_h^4 \tag{9}$$

The regression coefficients, β_i, depend on variables known at time h, including as the information about seasonal factors (hour, day of week, and weekday). Thus, these regression coefficients are different for each hour of the year and for the regime (spike or no spike). We found that a fourth order polynomial is adequate for our purposes. Since we are using these approximations to arrive at an optimal decision, any error in the approximation leads to a sub-optimal decision. Hence this will undervalue the power plant. We found very little difference in 3rd and 4th order approximations, hence the degree of under-valuation is unlikely to be significant.

Equation (9) can be estimated by simulations using ordinary least squares. For example, let ω be a given simulation. For equation (5), we can regress the observed switching cash flows, $\pi_h(m, m')_\omega$ on powers of $y_{h,\omega}$. This procedure gives an estimate of the switching alternatives, $\hat{v}_{H-1}(m, m', y_{H-1})$ in equation (7). Note this depends only on current, i.e. hour $H-1$, information. For a given initial mode, we choose the m' that maximizes $\hat{v}_{H-1}(m, m', y_{H-1})$ and use that decision to evaluate equation (8).[27] This least squares approach is used when the recursion discussed in section 7.1 is done. The only step where least squares is not used is in the initial value; this is due to there being only a single initial heat rate, y_0. In this case, the initial switching values are simple averages.

8 Valuation Results

In this section, we apply the model to compute values of the plant, and assess the characteristics of optimal operations. To protect the proprietary data of the owner of this plant, we have applied a positive linear transformation to the valuation data and cash flows (denominated in GJ). This does not change the nature of the optionality and optimal operation, so the results here are consistent with a plant that has slightly different ramp rates and output characteristics than the actual plant we studied.

[27] Note the actual v is used; hence the least squared technique is only involved in the decisions. A common error is to use $\hat{v}_{H-1}$ which will lead to additional estimation error in the value function.

First, in subsection 8.1, we will study the operations and valuation of the plant under the model decisions when applied to the actual market heat rate data for 2002-2006. We will perform a sensitivity analysis to show that the results are robust to various changes in the model. Then, in subsection 8.2, we revisit the market heat rate data for that year, and see what decisions and cash flows result from those decisions.

8.1 Applying the valuation to the market heat rates of 2002-2006

Figure 11 shows the number of model transitions (changes in plant operating mode) per year under four different optimization scenarios. In the first model, we average the LSM beta coefficients over one year to calculate conditional expectations of next period value, depending on the current operating mode and exogenous market parameters (market heat rate, time of day, day of week and month of year). This is used on each simulation path to determine an optimal sequence of operation decisions. This corresponds to "Beta Reset Each Year, No Penalty". This is how the model would work if the analyst recalibrates the LSM model annually.

The first model has two deficiencies. First, it generated many more operating mode transitions than were customarily used for this type of plant, and there was concern that it might impose excessive wear and tear on the equipment.Second, the company was interested in a model that did not require annual recalibration.

To address the first problem, we considered a variation of the model where a 1000 GJ penalty was imposed on each operating mode transition, to reduce the number of mode transitions.[28] To address the second problem, we used the average LSM conditional expectation betas estimated for 2002 and used them for determining optimal decisions in the subsequent years, without re-estimating the LSM model.

Incorporating the first variation only yields the model "Beta Reset Each Year, with Penalty". This operating model reduces the number of transitions per year from a range of 2500 to 3300 down to a range of approximately 1000, which is slightly less than 3 transitions per day.

Incorporating the second variation only yields the model "Constant 2002 Beta, No Penalty", and this gives almost the same number of transitions per year as the base model. Incorporating both variations together yields the fourth model "Constant 2002 Beta, With Penalty". This latter model gives almost the same number of transitions as the penalty model where the betas are re-estimated annually.

Thus, introducing the penalty for transitions does reduce the number of transitions, but there is little change in the number of transitions when we use the betas from 2002, rather than re-estimating each year.

Generally, there are more transitions in the middle years, which also had the lowest heat rates. The no-penalty model generated the highest number of transitions in 2005,

[28] With typical gas prices of $6/GJ, this corresponds to a $6000 penalty for changing the operating mode, which the company regarded as a reasonable estimate of incremental costs from wear and tear.

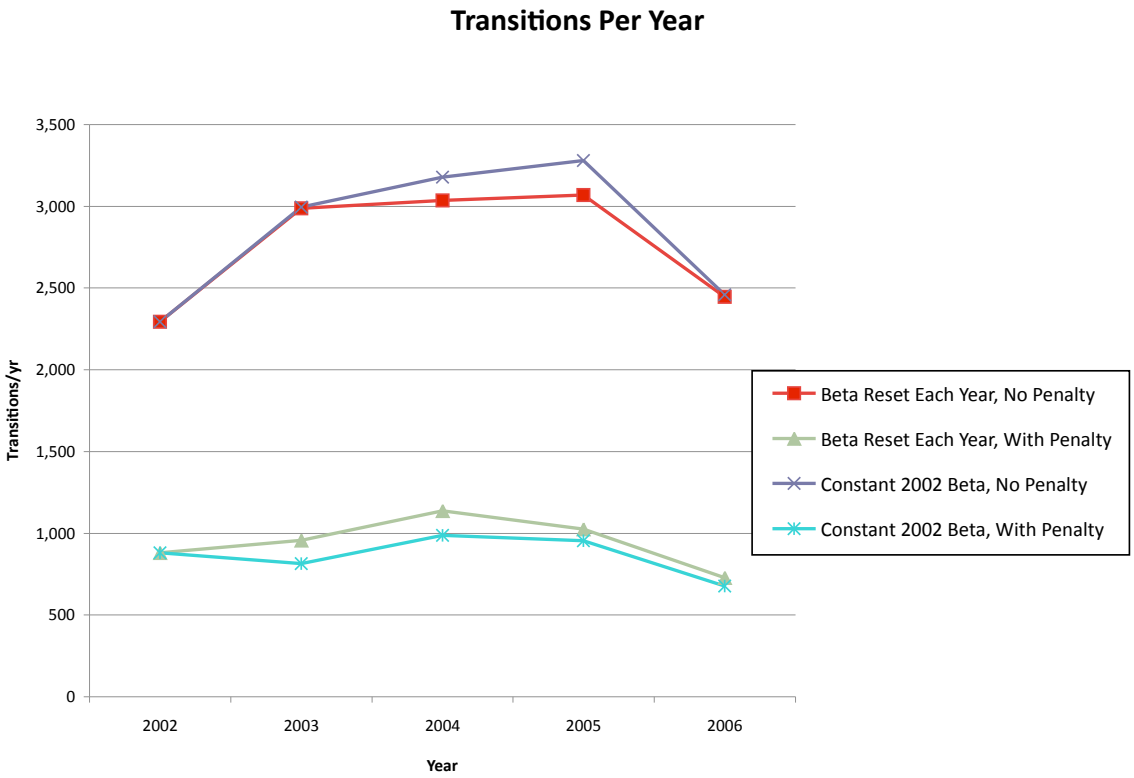

Figure 11: The number of transitions per year under 4 scenarios of optimal operation. Two of the scenarios are based on resetting the LSM projection betas with each year of LSM optimization, and two of the scenarios are based on keeping the LSM betas at the same as for the first year of LSM optimization (2002). The other criterion is whether the optimal transitions were chosen for with costs equal to the actual fuel burn rates, or whether transitions were discouraged by imposing a 1000GJ cost for each transition.

which is a low heat rate year. Interestingly, this had more transitions than 2006, which is the year with the most power price spikes.[29] Thus, volatility (or spikes) do not induce as many transitions (flexible behaviour), as we might initially expect for a real options model. This is likely because the year 2006 had a generally high heat rate, so the volatility apparently did not result in a large number of low market heat rates, where it would be optimal to turn the power plant off. Indeed, Figure 17 shows that in 2006, the power plant would be at full output 65% to 75% of the time, depending on the model used, so there would have been little reason to transition the plant to a lower production mode.

In contrast, there were a lot of transitions in 2005, which had a low market heat rate, because the real option flexibility can be used to ramp up the plant and capture the few periods where running the plant is optimal.

Figure 12 shows the amount of power generated each year under the 4 different models. The first broad observation to make is that the plant would generate more power in the early and late years than it would in the middle years. This is consistent with the U-shaped graph of average market heat rates Y that is shown in Table 1. In the middle years, the market heat rates were low and optimal operation had the plant shut down much of the time. Generally speaking, the penalty models tended to have the plant generating more power than the non-penalty models.[30]

Figure 13 shows the annual gross margin associated with each of the four strategies. To make the penalty models comparable to the no-penalty models, we added back the 1000GJ per transition penalty to the gross margin of the penalty models. That is, the penalty was used only to change behaviour, rather than reported profit.[31]

Two things stand out with this graph. First, the gross margin is U-shaped, just like the market heat rates, so the power plant generates more profit when heat rates are high, which is no surprise. What may seem more surprising, however, is that the gross margin is quite insensitive to the precise operating model used.

The fact that the profit margin is not sensitive to changes in the operating strategy, under optimal operations, arises in this case because the transitions that are avoided with the penalty model are those where the market heat rate was very close to the plant heat rate. Thus, the number of transitions can be reduced significantly, if the operator allows a plant to stay on even if the profit margin (spark spread) has temporarily turned negative. The losses in this case are not significant. Similarly, the penalty model keeps the plant from turning on to capture very small positive spark spreads, but missing

[29] Table 1 shows the highest standard deviation for 2006, for example.

[30] Except for the penalty model with the beta being reset each year, and only only for the year 2005, where the power generated is quite low. Clearly, the plant is being shut down completely for much of that year.

[31] Note that the pretax gross margin is in units of GJ of gas, given our Margrabe methodology. To convert the gross margin to a dollar figure, one can multiply by the price of gas in $/GJ. This can be done by taking a forward curve for gas prices.

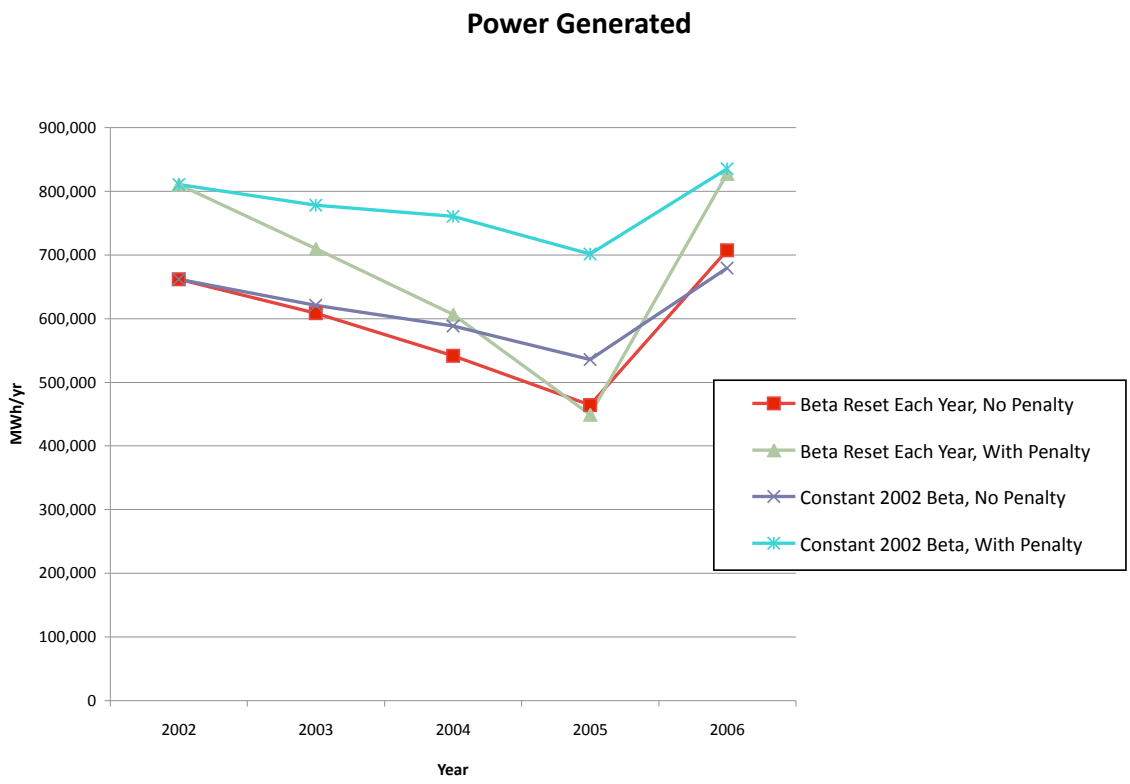

Figure 12: The amount of power generated per year under 4 scenarios of optimal operation. Two of the scenarios are based on resetting the LSM projection betas with each year of LSM optimization, and two of the scenarios are based on keeping the LSM betas at the same as for the first year of LSM optimization (2002). The other criterion is whether the optimal transitions were chosen for with costs equal to the actual fuel burn rates, or whether transitions were discouraged by imposing a 1000GJ cost for each transition.

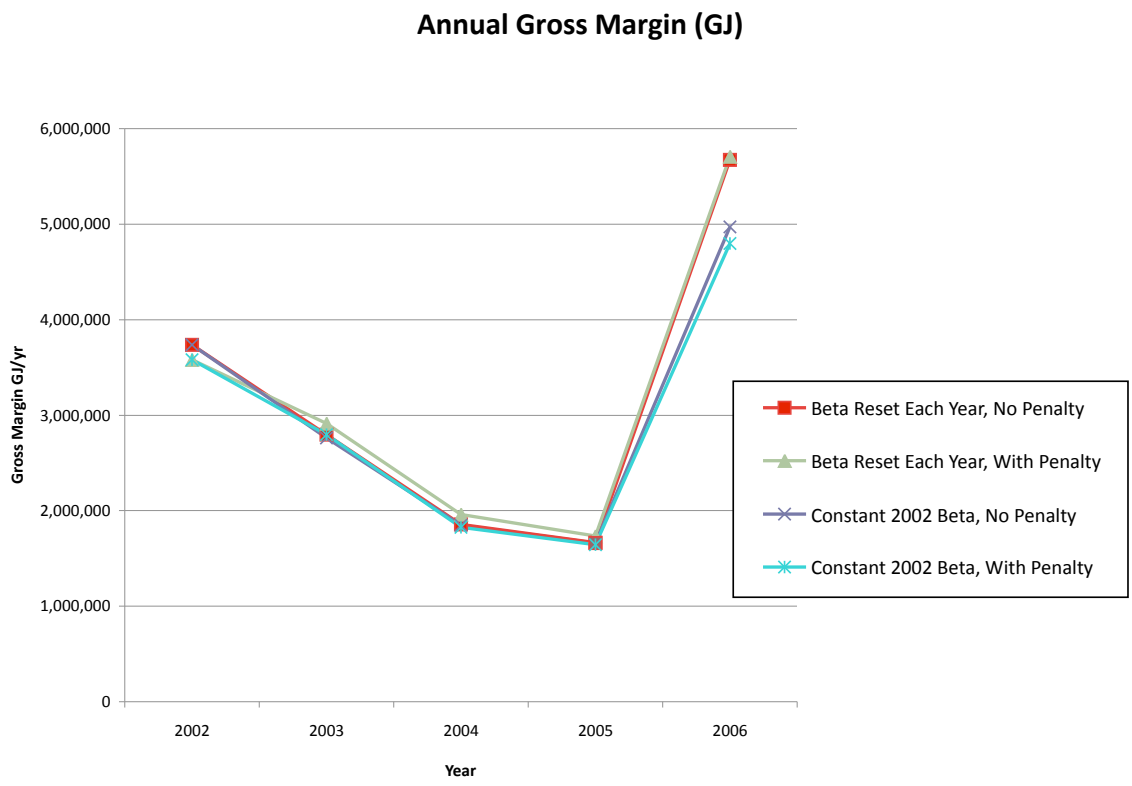

Figure 13: The annual gross margin (in GJ) under 4 scenarios of optimal operation. Two of the scenarios are based on resetting the LSM projection betas with each year of LSM optimization, and two of the scenarios are based on keeping the LSM betas at the same as for the first year of LSM optimization (2002). The other criterion is whether the optimal transitions were chosen for with costs equal to the actual fuel burn rates, or whether transitions were discouraged by imposing a 1000GJ cost for each transition.

these opportunities is not very costly in terms of overall profit.[32]

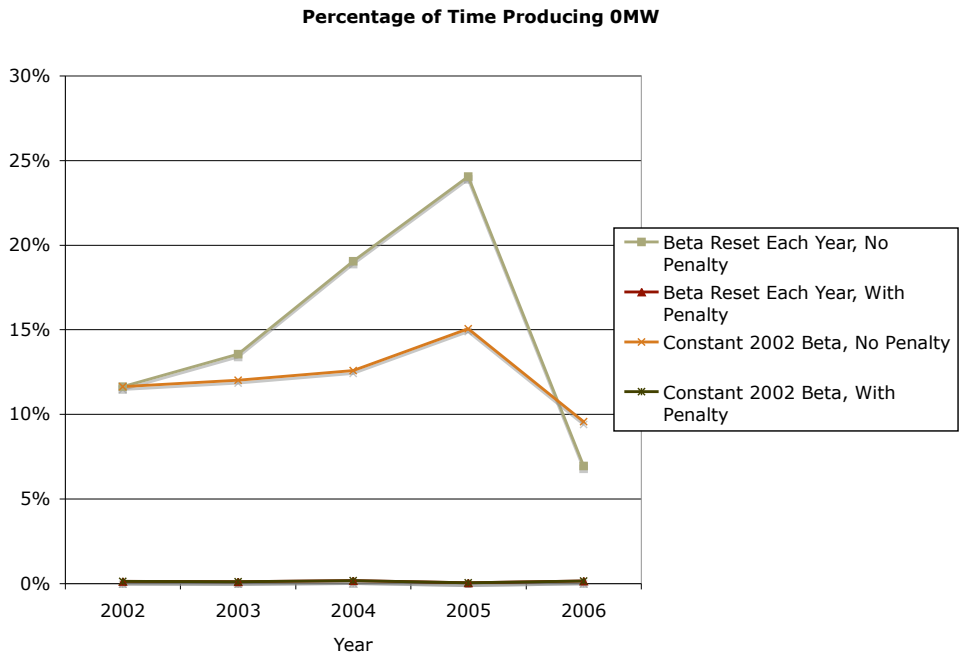

Figure 14: Proportion of time spent in the 0MW "Cold Metal" operating mode for each of the 4 operating models and operating years.

Figure 14 shows how often the model has the plant completely shut down ("Cold Metal") for each of the 4 years and for each of the 4 operating models. Note that the penalty models do not let the plant shut down, but the no-penalty models will allow it to shut down between 10% and 25% of the time. The no-penalty models would have been more likely to shut the plant down in the middle years when the market heat rate was low. The times when these models shut down the plant must have been mainly at the margin of indifference between running and shutting down, because the two different strategies led to very similar valuations in Figure 13.

Figure 15 shows that all of the models would have had the plant idling at 15MW in all years. In this case the stragegies led to essentially the same proportion of time idling, with greater variation occurring for different years with different heat rates. The plant

[32] The fact that the optimal value still stays near the optimum even when the decision rules are changed from the optimum is a result of the "smooth pasting condition" in this real options model of operating a plant. The smooth pasting condition is a condition of optimality and it says, that, when the plant is operated near the optimum, infinitesimal changes in policy do not affect value (gross margin). Significant changes in policy can affect value, and the only way to be sure is to compute the value under the optimal and adjusted policy, as we do here.

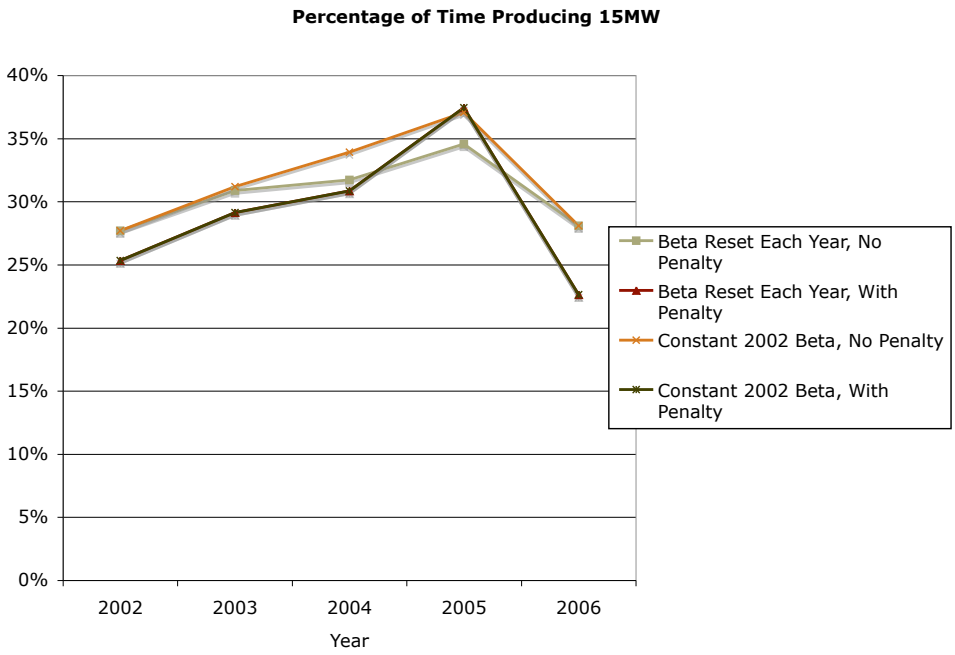

Figure 15: Proportion of time spent in the 15MW "Idling and Generating Steam" operating mode for each of the 4 operating models and operating years.

was more likely to be running at idle in 2005 when market heat rates were particularly low, and less likely in the other years with higher heat rates.

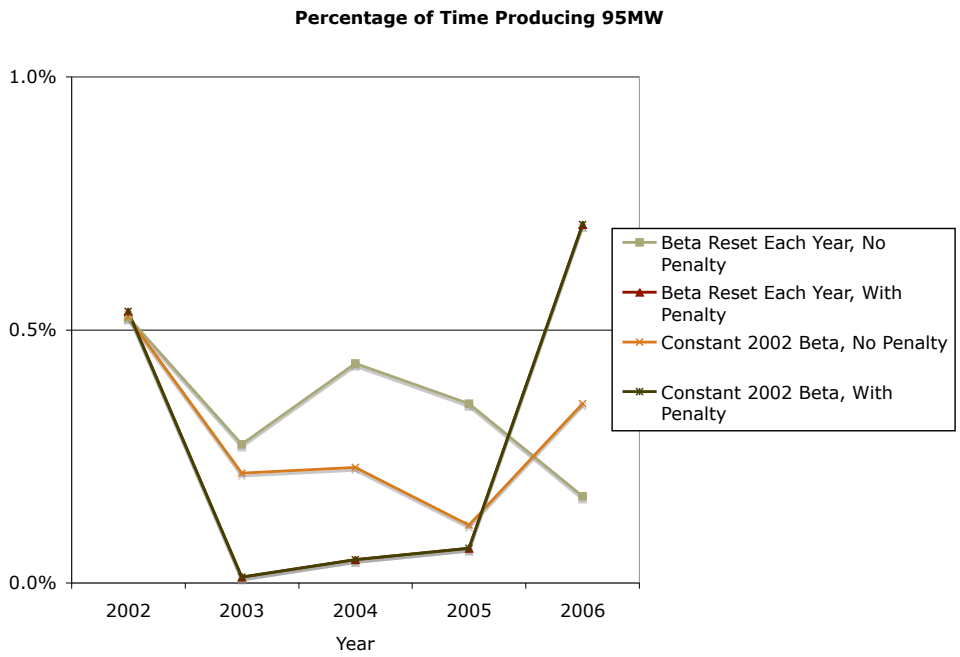

Figure 16: Proportion of time spent in the 95MW "Full Gas Turbine" operating mode for each of the 4 operating models and operating years.

Figure 16 shows that none of the models tended to let the plant run in full gas turbine mode, since this mode occurred in all years for all models less than 1% of the time. This is particularly interesting, since the 95MW gas turbine mode is a peaking mode that makes great use of the plant's ability to ramp up quickly. The operator would enter this mode if there is a surprising spike in the heat rate — the surprise is needed, since if the spike is predictable (based on time of day, week or year), the operator would also put it into combined cycle mode beforehand, and we would see the plant switching between the 15MW combined cycle mode and the 120MW full output combined cycle mode. We have already seen in Figure 15 that the 15MW combined cycle mode is, indeed used quite often.

Figure 17 shows that the plant is most often run in the full-output combined cycle mode, producing 120MW of power. The penalty models use this mode more often than the non-penalty models, which is the mirror image of the differential for their use of the 0MW Cold Metal mode. The plants were least likely to be in this full output mode in the middle years when the market heat rate was low.

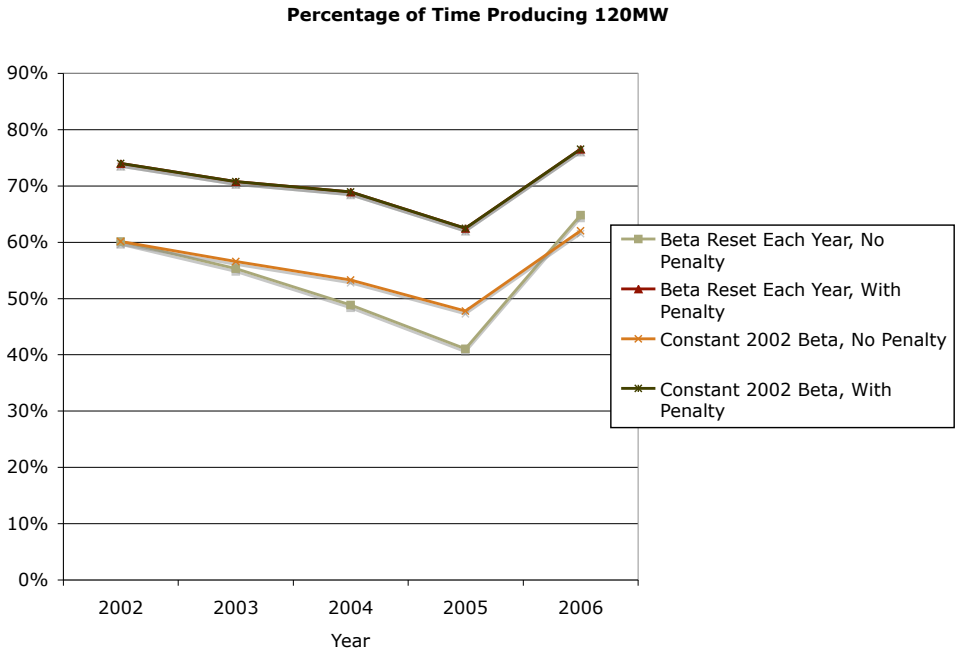

Figure 17: Proportion of time spent in the 120MW "Full Combined Cycle" operating mode for each of the 4 operating models and operating years.

8.2 Extending the valuation to other scenarios with simulation

The previous analysis only studied the 5 market scenarios that occurred in 2002 to 2006. It is important to consider the results for other plausible market scenarios. In this section, we compute other scenarios by simulation, varying model parameters to get different average market heat rates for various years. We will see that the overall value is mainly dependent on the average heat rate for the year. Other market issues such as the frequency of spikes are important, but not as much as the heat rate.

To value a yearly operation of power plants, we must run a given number of simulations of one year of hourly market heat rates. Note this gives 8760 hourly prices per simulation. The number of simulations we chose is 500; we found this was a reasonable tradeoff in computation time and accuracy. Once the 500 simulations are done, we can perform LSM to find a value for each power plant. There are a few items/questions that must still be addressed:

- We must generate the present value for cash flows greater than a year.
- How do we choose a stylized year? The years 2002-2006 were widely varied in average market heat rates and spike occurrences.
- How can the modeler compute the year's plant value for heat rate forecasts that are different than the average heat rate for the stylized years?

It turns out that there is a simple solution to all three of these concerns. We chose to run the valuations for all 5 stylized years and then altered the mean regime heat rates by adding and subtracting constants to the intercept of the estimated regime regressions. Specifically, we add and subtract 1%, 2%, 5%, and 10% to the intercept term in equation (3). This implies the mean heat rates of the simulations will be different which will lead to different simulated valuations. This leads to a sequence of expected heat rates and associated power plant values. These are presented in Figure 18, which plots the observed valuation and the approximated valuation functions for one year of operations, again denominated in GJ rather than $. The lowest value is obtained when the plant is always generating power, which is to say that it is a baseload plant. It represents the value of the plant without any option value. The upper values correspond to the plant being operated optimally. As the mean heat rates become very high, the two values converge to each other because there is little to be gained by shutting down the plant when the heat rate is high. We also fit a fourth order polynomial to fit the points of the optimal valuation. When the plant is operated as a baseload plant, we can see that the gross margin is linear in the mean market heat rate, much like the intrinsic value of an option. But, when it is allowed to operate optimally, the downside losses at low heat rates are mitigated and the overall value rises. The polynomial fit gives an increasing convex function over the relevant range, much like an option value lies about the underlying intrinsic value of the option.

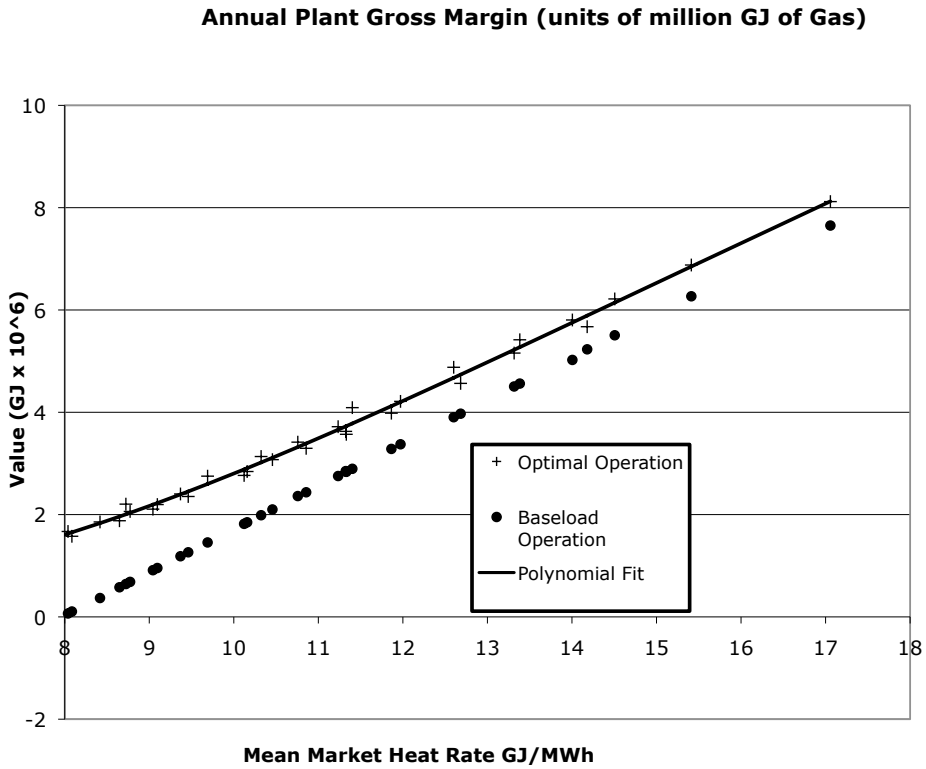

Figure 18: Valuation of the two power plant as a function of average yearly heat rates. The upper curve is a polynomial fit to the simulated value when the plant is run optimally. The lower circles are simulated values when the plant is always on, as in baseload operation. The option value lifts the value above the baseload value, but the value gain is reduced when the market heat rate is high.

9 Conclusion

This paper has analyzed power plant operating in Alberta that is based on two General Electric LM6000 gas turbines, that can also generate steam power from the exhaust gases.

We have found that the gross margins for the plants are sensitive to market conditions, such as the average heat rate for the year. However, they are not very sensitive to varying the strategy to reduce the amount of plant cycling. We have shown that the operating rules can be modified to cut the number of plant transition cycles in half without significantly impairing plant financial performance.

References

Carmona, R. and Durrleman, V. (2003). Pricing and hedging spread options. *SIAM Review*, 45(4):627–685.

Elliott, R., Sick, G., and Stein, M. (2002). *Real Options and Energy Management*, chapter Pricing Interactions of Baseload Supply Changes and Electricity Demand, pages 371–391. Risk Waters Group.

Hodder, J., Mello, A., and Sick, G. A. (2001). Valuing real options: Can risk adjusted discounting be made to work? *Journal of Applied Corporate Finance*, 14(2):90–101.

Longstaff, F. A. and Schwartz, E. S. (2001). Valuing american options by simulation: a simple least-squares approach. *The Review of Financial Studies*, 14(1):113–147.

Margrabe, W. (1978). The value of an option to exchange one asset for another. *Journal of Finance*, 33(1):177–186.

Sick, G. (2007). An overview of real options. Presented at NACVA Conference June 2007.