Pricing and Hedging Electricity Supply Contracts: a Case with Tolling Agreements

Shi-Jie Deng

Email: deng@isye.gatech.edu

Zhendong Xia

Email: dengie@isye.gatech.edu

School of Industrial and Systems Engineering

Georgia Institute of Technology

Atlanta, Georgia 30332–0205

January 31, 2005

Abstract

Customized electric power contracts catering to specific business and risk management needs have gained increasing popularity among large energy firms in the restructured electricity industry. A tolling agreement (or, tolling contact) is one such example in which a contract buyer reserves the right to take the output of an underlying electricity generation asset by paying a predetermined premium to the asset owner. We propose a real options approach to value a tolling contract incorporating operational characteristics of the generation asset and contractual constraints. Dynamic programming and value function approximation by Monte Carlo based least-squares regression are employed to solve the valuation problem. The effects of different electricity price assumptions on the valuation of tolling contracts are examined. Based on the valuation model, we also propose a heuristic scheme for hedging tolling contracts and demonstrate the validity of the hedging scheme through numerical examples.

Index Terms— Electricity options pricing, tolling agreement, spark spread, real options, dynamic hedging, risk management, Monte Carlo simulation.

1 Introduction

Electric power markets have been established worldwide due to the global restructuring of the electricity supply industry. With the power industry being restructured into three separate industrial segments of generation, transmission and distribution, firms in different segments possess distinct risk profiles. For instance, independent power producers in the generation segment face potential risks in both revenue and production cost since they are subject to the market price risk of both underlying commodities (e.g. electricity and input fuel). On the other hand, utility companies, which become more focused in the business of local transmission and distribution in the restructuring process, are mostly concerned with having ample electricity supply to serve their customers at a profitable margin.

In the early 2000's, the rise and fall of the several large U.S. electric power merchants created turmoils in the power markets and it consequently caused sizable financial losses to major financial institutions which offered loans to finance these power marketers' investment projects and business transactions. Basically, a large portion of the acquired power generation assets and the signed power purchasing contracts by the power merchants turned out to be far less profitable than what was expected due to optimistic valuations and insufficient risk management. It made the power marketers unable to pay back their loans in due time and put them under great financial distress. These adverse events have demonstrated the importance of an appropriate valuation and effective risk management methodology in power markets for both market participants and the financial institutions such as banks which have business dealings with these market participants.

The unique physical characteristics of electricity make its price the most volatile one among all commodity prices. Noting the extremely high price volatility, power market participants are especially wary of the price risk associated with business transactions and they resort to customized (most likely long-term) business transactions to hedge their respective unique risk profiles thus making the bilateral and multilateral power supply contracts ubiquitous. A market-based valuation approach is essential for pricing and risk managing these bilateral (sometimes multilateral) power transactions.

The valuation of electricity contracts differs from that of other financial contracts in that: a) the underlying electricity is not a traded asset, meaning that it cannot be bought and hold; b) electricity contracts often contain side constraints (e.g., various contract provisions) on how financial payouts are derived from the underlying electricity or a physical asset generating electricity. While a market-based valuation can be carried out by taking the price of electricity as a state variable and adopting a proper discounting factor, these side-constraints significantly increase the complexity of pricing electricity contracts. The goal of our paper is to propose a market-based approach for pricing and hedging electricity contracts with a complex contractual structure. We outline typical operational and contractual provisions in a structured electricity supply contract and incorporate them into a real options valuation framework. This approach is a valuable tool for both power market participants and financial institutions which are interested in exploring business opportunities in power markets.

The discounted cash flow method (DCF) was the norm for valuing power supply contracts and evaluating generation/transmission asset investments in the traditionally regulated electricity industry since power price was set by regulators based on cost of service. The basis of DCF valuation is a set of static (or, estimated) future cash flow. However, the electricity prices are no longer preset in the newly restructured power industry and they are driven by the ever-changing fundamental market supply and demand conditions. Under the new regime, a DCF valuation approach, which is based on static cash flow estimates rather than a dynamically evolving cash flow, undervalues power contracts and assets because it fails to capture the value associated with the inherent optionality for dynamically maximizing the cash flow of an underlying asset and takes little account of the extraordinary electricity price volatility into the valuation. Deng, Johnson, and Sogomonian (2001) propose a *real options* approach based on an analogy between the payoff of certain financial options and that of a physical asset for power asset valuation. They demonstrate that the option-pricing approach is the better alternative to the DCF method based on market information. Deng and Oren (2003) and Tseng and Barz (2002) advance the real options valuation of power plants further by incorporating operational constraints into the valuation framework.

Motivated by these works, we extend the real options approach for valuing power plants to the valuation of electricity contracts with embedded options. A complex electricity contract, such as a *tolling agreement* (or, *tolling contract*), is more challenging to value than a physical power asset (such as a power plant) since the contract can contain contractual constraints that are both operationally set and artificially designed. We formulate a tolling contract as a collection of multiple tolling options (introduced in section 3) with constraints on their exercising. We extend a Monte Carlo simulation approach with value function approximation, which is developed in Carriere (1996), Longstaff and Schwartz (2001), and Tsitsiklis and Van Roy (2001) for pricing American options with one single exercising decision to make, to the tolling agreement valuation problem with multiple exercising decisions and side-constraints under general assumptions on the electricity and fuel price dynamics. In particular, this approach is applicable to the models in Deng and Oren (2003) and Tseng and Barz (2002) with extensions to a wide range of electricity and fuel price assumptions. It can also be applied to other complex energy contract pricing problems as those in Thompson (1995), Jaillet, Ronn, and Tompaidis (2001), and Keppo (2004).

In the next two sections, we develop a real options based model for valuing a tolling agreement under both operational and contractual constraints.

2 Problem Description

Tolling contract is one of the most innovative structured transactions that has been embraced by the power industry. A tolling agreement is similar to a common electricity supply contract signed between a buyer (e.g. a power marketer) and an owner of a power plant (e.g. an independent power producer) but with notable differences. For an upfront premium[1] paid to the plant owner, it gives the buyer the right to either operate the power plant or simply take the output electricity during pre-specified time periods subject to certain constraints. In addition to inherent operational constraints of the underlying power plant, there are often other contractual limitations listed in the contract on how the buyer may control the power plant's operations or take the output electricity. For instance, a tolling contract almost always has a clause on the maximum allowable number of power plant restarts as frequent restarting of a generator increases the maintenance costs borne by the plant owner.

As a tolling agreement gives its buyer the right to take the electricity output of an underlying power plant subject to certain contractual constraints, holding a tolling contract is equivalent to owning the underlying plant but with operational flexibility constrained by additional contractual terms. By noting this analogy, we model a tolling agreement as a series of real options on operating a power plant coupled with contractual operating limits. We elaborate on the modeling details of the operational and contractual constraints involved in a tolling contract in this section.

[1] Woo, Olson, and Orans (2004) provide a statistical benchmark analysis on the reasonableness of the level of such premium based on historical price data of electricity and fuel.

Tolling agreements are written on fossil-fuel power plants. A fossil-fuel plant converts a generating fuel into electricity at certain conversion rate known as *heat rate*. In brief, heat rate measures the units of the fuel needed for producing one unit of electricity. The lower/higher is the heat rate, the more/less efficient is the power plant. The heat rate is measured in units of MMBtu/MWh where one MMBtu represents one million British thermal units and one MWh stands for one Megawatt (MW) hour of electric energy. In our model, we assume that there are only one power market and one gas market. The owner or any party who operates the power plant has the right but not obligation to generate electricity (e.g., an owner of a merchant power plant). This right-to-generate is known as an *operational option*, which falls into the category of *real options* (see Dixit and Pindyck (1994) for more examples of real options). By exercising the operational option, the plant operator receives the spot price of electricity less the heat rate adjusted input fuel cost by buying fuel and selling electricity in their respective spot markets. The "spread" between the electricity price and the heat rate adjusted fuel cost is called *spark spread*. Absent of operational constraints, a rational power plant operator turns on the plant to generate electricity whenever the spark spread (namely, the payoff of the operational option) is positive and shuts down the plant otherwise. Since a *spark spread call option* pays out the positive part of the price difference between the electricity and the generating fuel (namely, the spark spread), the payoff of a power plant at each time epoch t can be replicated by that of a properly defined spark spread call option (see Deng, Johnson, and Sogomonian (2001)). If ignoring both operational and contractual constraints, a tolling agreement is simply equivalent to a strip of spark spread call options with maturity time spanning through the duration of the contract.

The operational constraints in a tolling agreement are naturally tied to those in operating the underlying power plant. Among all aspects of operating a power plant, we consider three major operational characteristics (Wood and Wollenberg (1984) offer a good review on power plant operations). First of all, fixed costs are always incurred whenever a power generator is turned on from its "off" state (termed as *startup costs*). Startup costs are generally time dependent. Sometimes, there are costs associated with the turn-off process of a power plant as well which are called *shutdown costs*. The startup and shutdown costs are fixed costs borne by the tolling contract holder. Secondly, the tolling contract holder usually cannot get electricity output immediately after starting up a power plant. There is a ramp-up delay period D for a generating unit to reach certain operating output level starting from the "off" state. Costs incurred during the ramp-up period are also time dependent. Thirdly, a power plant may be operated at a continuum of output levels. At each output level, the generator has a different heat rate. A power plant is usually more efficient (consuming less fuel per unit of electricity generated) when operating in full capacity than running at a lower output level. Therefore, the heat rate of a power plant is a function of the output level.

On the contractual constraints, we use the maximum restart limit described above as one representative example. While the profit of generating electricity comes from the positive spark spread between generated electricity and the input fuel, it is clear that a power plant would only lose money when the spark spread becomes negative possibly due to too low an electricity price or too high a fuel cost. In times of the spark spread turning so negative that a temporary shutdown of the power generating unit is justified, the operator has to turn off the unit and restart it later when the profit of generating electricity becomes positive again. However, frequent restarts are detrimental to a generation unit since a restart reduces the unit's lifetime and increases the likelihood of a forced outage. Due to this fact, there is usually a provision specifying the maximum number of restarts allowed in a tolling contract. Sometimes this constraint is implemented through imposing an extremely high penalty charge on each restart beyond certain threshold on the cumulative number of restarts in the contract effectively capping the total number of restarts at the threshold level. As a result, a tolling contract holder cannot order to shut down the plant at will whenever the electricity spot price is lower than the heat rate adjusted generating fuel cost. Consequently, the value of a tolling agreement is affected by such a constraint.

Intuitively, the value of a tolling contract at any time depends on the *state* of the underlying power plant. The operational characteristics of a power plant provide natural guidelines for defining the state of the plant. The *state* of a tolling agreement encompasses both the operational state of the underlying plant and the operational status related to the contractual obligations. We elaborate on the definition of a power plant's state in a tolling agreement through an example. Suppose a power plant has 2 output levels: the minimum level and the maximum level, and it takes 2 phases to ramp up the production from the "off" state to the minimum output level. Then the power plant has 5 operational states: "*off*", "*ramp-up phase-1*", "*ramp-up phase-2*", "*operating at the minimum output level*" and "*operating at the maximum output level*". Consider a tolling agreement on this facility that allows the buyer to restart up to n times. In this example, the state of the contract at time t consists of the operational state and the number of allowable restarts left by t. Figure 1 illustrates all possible states with each circle representing one state and all feasible transitions between any two states of the contract subject to the number of restarts constraint. Each row in figure 1 corresponds to an operational state of the plant while every column is tied to the allowable number of restarts left. For instance, the circle at the intersection of the second row and the second column represents a state in which the plant is in the first phase of ramping up and there are $n - 1$ allowable restarts remaining. With all possible states of a tolling contract defined, we proceed with the problem formulation for valuing a tolling agreement.

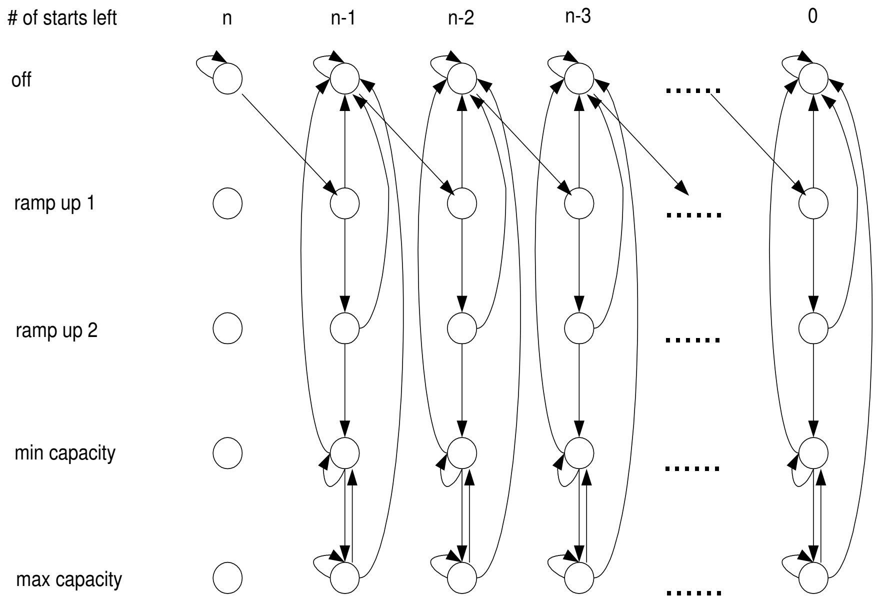

Figure 1: State Transition Diagram of a Tolling Contract with the Restart Constraint

3 A Stochastic Dynamic Programming Valuation Model

Consider a tolling contract written on an underlying power plant that has the three operational characteristics discussed in section 2. The contract allows no more than N re-starts of the power plant during its duration of T. Suppose the contract holder makes decisions on whether to take the output electricity at M discrete time points $t_1, t_2, \ldots, t_M$ over the horizon $[0, T]$ where $0 = t_1 < t_2 < \cdots < t_M = T$. N is very small comparing to M. The fact that the holder may take the electricity at any time t is modelled by letting M be an arbitrarily large integer. Since the electricity and the fuel are traded in the open markets, the holder elects to take the electricity whenever the spark spread is positive due to the no-arbitrage principle. As a result, the optimal take-or-not (and, quantity-to-take) decisions by the contract holder correspond exactly to the optimal produce-or-not (and, quantity-to-produce) decisions by the plant operator under the objective of maximizing the cumulative profit of the power plant subject to the tolling contract provisions.

Let us define a *tolling option* in a tolling contract to be the right of a contract holder to start taking the output electricity of the underlying plant at any time with self-supplied generating fuel, and the obligation, after exercising the right-to-take, to continuously take electricity (possibly in varying quantities) until she/he chooses to stop. The value of a tolling contract is therefore equal to the maximized total payoff associated with all exercised tolling options subject to a constraint that no more than N tolling options can be exercised during the life of the tolling contract. In exercising a tolling option, two sequential decisions need to be made: the first being when to start taking electricity and the second being when to stop. By no-arbitrage, the underlying plant is started (re-started) at the beginning of an exercised tolling option and shut down at the termination time. The contract holder is responsible for the corresponding startup and shutdown costs. If there were no startup or shutdown costs or other operational constraints of a power plant, then the optimal decisions in exercising a tolling option would be to start taking electricity whenever the spark spread turns positive and to stop doing so whenever the spark spread turns negative. In such an ideal case, a tolling option is simply a series of spark spread call options with the longest maturity time given by the first time of hitting zero by the spark spread with a positive initial value.

As explained in Dixit and Pindyck (1994), the real options (in this case, tolling options) can be valued by a stochastic dynamic programming (SDP) approach. When the payoff of the real options are perfectly replicated by traded financial instruments such as forward contracts on electricity and fuel, the correct discount rate used in the SDP approach needs to be the risk-free interest rate and the SDP approach becomes equivalent to the contingent claim analysis for option-pricing as developed in Black and Scholes (1973), Merton (1973), and Harrison and Kreps (1979). In the case where the available traded financial instruments cannot achieve perfect hedging (i.e., incomplete market), the discount rate is obtained by adding a risk premium to the risk-free rate.

3.1 Formulation of the Tolling Contract Valuation

The following notations are used throughout the paper.

V_t	:	value of the tolling contract at time t,
a_t	:	operational action taken at time t by a profit-maximizing power plant operator,
R_t	:	payoff of the operational option at time t,
n_t	:	number of power plant re-starts left at time t,
w_t	:	operational state of the power plant at time t,
Θ_t^P	:	state of the tolling contract, (w_t, n_t), at time t,

X_t, Y_t	:	natural logarithm of electricity and the fuel prices at time t, respectively,
Θ_t^S	:	log-price vector (X_t, Y_t),
Θ_t	:	vector (X_t, Y_t, w_t, n_t) representing the state of the world.

3.1.1 Value Function

From here on, we refer to X_t and Y_t as prices with the understanding that they represent log-prices. Suppose that the price vector $(X_t, Y_t) \in R^2$ evolves according to a Markov process defined in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with initial value (X_0, Y_0) at time 0. The σ-field generated by the stochastic process $\{(X_s, Y_s) : 0 \leq s \leq t\}$, denoted by $\mathcal{F}_t \subset \mathcal{F}$, forms a filtration $\mathbb{F}$ over time interval $[0, T]$. Let A_t denote the set of admissible operations available to the plant operator at time t and R_t be the payoff of the time-t operational option. Based on the previous analysis, the value of a tolling contract at time $t \in \{t_1, t_2, \ldots, t_M\}$ is given by

$$V_t = \max_{\{a_{t_i} \in A_{t_i}, \cdots, a_{t_M} \in A_{t_M}\}} E[\sum_{I=i}^{T} e^{-r(t_I - t)} R_{t_I} | \mathcal{F}_t] \tag{1}$$

where $t = t_i$ and r is a discount rate. R_t can be interpreted as the operating profit at time t. It depends on both the operational action and the state of the world, namely, $R_t \equiv R(a_t, \Theta_t) \equiv R(a_t, X_t, Y_t, w_t, n_t)$. While the range of (X_t, Y_t) is R^2 and $n_t \in W_R \equiv \{0, 1, 2, \ldots, N\}$, we need to introduce w_t and a_t before defining $R(a_t, \Theta_t)$.

For the ease of exposition, we make further simplifying assumptions on the operating characteristics which can be readily generalized. Specifically, the underlying power plant has one "off" state and two output states: the minimum output state of generating $\underline{Q}$ MW per time unit (with heat rate $\underline{Hr}$) and the maximum output state of generating $\overline{Q}$ MW per time unit (with heat rate $\overline{Hr}$). Since a power plant works more efficiently at the maximum output level than at the minimum output level, $\overline{Hr}$ is smaller than $\underline{Hr}$. We also assume no-delay and no-cost in switching between the minimum and the maximum output levels. A start-up cost c_{up} is incurred whenever the power plant is turned on from the "off" state. Recall there is a delay (or, ramp-up) period of D (called the ramp-up time) before a power plant can output electricity after the plant being turned on from the "off" state. Without loss of generality, we assume that D is a multiple of Δt where Δt is the length of the small time intervals over which the operating decisions are made. Let K_D denote $\frac{D}{\Delta t}$. Then there are $K_D - 1$ ramp-up states. During the ramp-up period, the cost is $c_r(Y_t)$ per time unit at time t, which is a positive increasing function of the generating fuel price Y_t. A shut-down cost c_{down} is incurred whenever the power plant is turned off. To summarize, we use a set $W_D \equiv \{0, 1, \cdots, K_D - 1, K_D\}$ to represent the $(K_D + 1)$ possible operational states. Specifically, w_t takes on $(K_D + 1)$ possible values at time t.

- $w_t = 0$: The power plant is in *off* state at time t.
- $w_t = i$: The power plant is *on* but in the i^{th} stage of the *ramp-up* period D at time t for $i \in \{1, 2, \cdots, K_D - 1\}$.
- $w_t = K_D$: The power plant is *on and ready* to generate electricity outputs at time t.

The operational action of the plant operator, a_t, has three possible choices a^i $(i = I, II, III)$ corresponding to the operational states. The admissible action set A_t in (1) is a subset of $A \equiv \{a^I, a^{II}, a^{III}\}$ for all time t.

- a^I: The operator operates the power plant at the maximum capacity level. The plant generates $\overline{Q} \cdot \Delta t$ units of electricity in time Δt with an operating heat rate of $\overline{Hr}$ if it is not in a ramp-up stage; otherwise it generates 0 units of electricity.
- a^{II}: The operator keeps the power plant running at the minimum capacity level. The plant generates $\underline{Q} \cdot \Delta t$ units of electricity in time Δt with an operating heat rate of $\underline{Hr}$ if it is not in a ramp-up stage; otherwise it generates 0 units of electricity.
- a^{III}: The operator turns the power plant off from any non-"*off*" state.

The operating profit of the power plant at any time t, $R(a, x, y, w, n) : A \times R^2 \times W_D \times W_R \to R^1$, is defined as follows. The operational characteristics described above are reflected in the definition of $R(a, x, y, w, n)$.

- When $w = 0$,

$$R(a, x, y, 0, n) = \begin{cases} -\infty, & \text{if } a_t = a^I \text{ or } a^{II},\ n = 0,\ \forall (x, y). \\ -c_{start} - c_r(y) \cdot \Delta t, & \text{if } a = a^I \text{ or } a^{II},\ n \geq 1,\ \forall (x, y). \\ 0, & \text{if } a_t = a^{III},\ \forall n,\ \forall (x, y). \end{cases} \tag{2}$$

- When $w = 1, 2, \cdots, (K_D - 1)$,

$$R(a, x, y, w, n) = \begin{cases} -c_r(y) \cdot \Delta t, & \text{if } a = a^I, \forall\ n, \forall\ (x, y). \\ -c_r(y) \cdot \Delta t, & \text{if } a = a^{II}, \forall\ n, \forall\ (x, y). \\ -c_{down}, & \text{if } a = a^{III}, \forall\ n, \forall\ (x, y). \end{cases} \tag{3}$$

- When $w = K_D$,

$$R(a, x, y, K_D, n) = \begin{cases} \overline{Q} \cdot \Delta t \cdot [e^x - \overline{Hr} \cdot e^y], & \text{if } a = a^I, \forall\ n, \forall\ (x, y). \\ \underline{Q} \cdot \Delta t \cdot [e^x - \underline{Hr} \cdot e^y], & \text{if } a = a^{II}, \forall\ n, \forall\ (x, y). \\ -c_{down}, & \text{if } a = a^{III}, \forall\ n, \forall\ (x, y). \end{cases} \tag{4}$$

3.1.2 Hamilton-Jacobi-Bellman Equations for the Value Function

With $\{(X_t, Y_t) : t \geq 0\}$ being a Markov process and R_t defined by (2)-(4), the contract value function V_t in (1) simplifies to a function of the state variables (X_t, Y_t, w_t, n_t) at time t, namely, $V_t = V_t(X_t, Y_t, w_t, n_t) \equiv V_t(\Theta^S_{t_{i+1}}, \Theta^P_{t_{i+1}})$. Moreover, the value function $V_t(X_t, Y_t, w_t, n_t)$ satisfies the following Hamilton-Jacobi-Bellman equations in all possible states of the contract Θ^P_t at every $t_i \in \{t_1, t_2, \ldots, t_{M-1}\}$. Let $E_t[\cdot]$ denote the conditional expectation operator $E[\cdot|\mathcal{F}_t]$.

- For every price vector $\Theta^S_t = (X_t, Y_t) \in R^2$ and the state of the contract Θ^P_t being $(w_t, n_t) =$

$(0, n_t)$ for every $n_t \geq 1$ at $t = t_i$,

$$V_t(\Theta_t^S, \Theta_t^P) = \max_{a_t \in A_t} \begin{cases} a_t = a^I : -c_{up} - c_r(Y_t)\Delta t + e^{-r\Delta t} E_t[V_{t_{i+1}}(\Theta_{t_{i+1}}^S, \Theta_{t_{i+1}}^P)] \\ a_t = a^{II} : -c_{up} - c_r(Y_t)\Delta t + e^{-r\Delta t} E_t[V_{t_{i+1}}(\Theta_{t_{i+1}}^S, \Theta_{t_{i+1}}^P)] \\ a_t = a^{III} : e^{-r\Delta t} E_t[V_{t_{i+1}}(\Theta_{t_{i+1}}^S, \Theta_t^P)] \end{cases} \quad (5)$$

where $\Delta t = (t_{i+1} - t_i)$ and $\Theta_{t_{i+1}}^P = (1, n_t - 1)$.

- For every price vector $\Theta_t^S = (X_t, Y_t) \in R^2$ and every state of the contract $\Theta_t^P = (w_t, n_t) \in \{(W_D \setminus \{0, K_D\}) \times W_R\}$ at $t = t_i$,

$$V_t(\Theta_t^S, \Theta_t^P) = \max_{a_t \in A_t} \begin{cases} a_t = a^I : -c_r(Y_t)\Delta t + e^{-r\Delta t} E_t[V_{t_{i+1}}(\Theta_{t_{i+1}}^S, \Theta_{t_{i+1}}^P)] \\ a_t = a^{II} : -c_r(Y_t)\Delta t + e^{-r\Delta t} E_t[V_{t_{i+1}}(\Theta_{t_{i+1}}^S, \Theta_{t_{i+1}}^P)] \\ a_t = a^{III} : -c_{down} + e^{-r\Delta t} E_t[V_{t_{i+1}}(X_{t_{i+1}}, Y_{t_{i+1}}, 0, n_t)] \end{cases} \quad (6)$$

where $\Delta t = (t_{i+1} - t_i)$ and $\Theta_{t_{i+1}}^P = (w_t + 1, n_t)$.

- For every price vector $\Theta_t^S = (X_t, Y_t) \in R^2$ and the state of the contract $\Theta_t^P = (K_D, n_t)$ for every $n_t \in W_R$ at $t = t_i$,

$$V_t(\Theta_t^S, \Theta_t^P) = \max_{a_t \in A_t} \begin{cases} a_t = a^I : \overline{Q}\Delta t[e^{X_t} - \overline{Hr} \cdot e^{Y_t}] + e^{-r\Delta t} E_t[V_{t_{i+1}}(\Theta_{t_{i+1}}^S, \Theta_{t_{i+1}}^P)] \\ a_t = a^{II} : \underline{Q}\Delta t[e^{X_t} - \underline{Hr} \cdot e^{Y_t}] + e^{-r\Delta t} E_t[V_{t_{i+1}}(\Theta_{t_{i+1}}^S, \Theta_{t_{i+1}}^P)] \\ a_t = a^{III} : -c_{down} + e^{-r\Delta t} E_t[V_{t_{i+1}}(X_{t_{i+1}}, Y_{t_{i+1}}, 0, n_t)] \end{cases} \quad (7)$$

where $\Delta t = (t_{i+1} - t_i)$ and $\Theta_{t_{i+1}}^P = (K_D, n_t)$.

Since the contract has no value to the holder in two scenarios: a) when the contract reaches its expiration time T, regardless of the power plant being *off* or not; b) when the power plant is *off* and the number of allowable re-starts is 0, the boundary conditions are

$$
\begin{aligned}
&V_t(X_t, Y_t, 0, 0) \equiv 0,\ \forall\ (X_t, Y_t) \in R^2 \text{ and } \forall\ t \in \{t_1, t_2, \ldots, t_M\}, \\
&V_T(X_T, Y_T, w_T, n_T) \equiv 0,\ \forall\ (X_T, Y_T) \in R^2 \text{ and } \forall\ (w_T, n_T) \in W_D \times W_R, \\
&V_t(X_t, Y_t, w_t, -1) \equiv -\infty,\ \forall\ (X_t, Y_t) \in R^2 \text{ and } \forall\ w_t \in W_D.
\end{aligned} \tag{8}
$$

3.2 Specification of the Underlying Commodity Price Processes

In the SDP formulation, price processes $\{(X_t, Y_t) : t \geq 0\}$ are key components. To illustrate the impacts of the commodity price assumptions on the contract valuation, we specify two alternative underlying commodity price models and investigate the effects of different price models on the valuation of a tolling contract. We shall see from the numerical examples in section 5 that different commodity price models systematically bias the value of a tolling agreement against one another.

For modelling the electricity price X_t, we consider both a simple mean-reverting process and a mean-reverting jump-diffusion process which are adapted to the filtration $\mathbb{F}$. The generating fuel price Y_t is modelled by an adapted simple mean-reverting process.

Case 1: both $\{X_t : t \geq 0\}$ and $\{Y_t : t \geq 0\}$ are mean-reverting processes.

$$
\begin{aligned}
dX(t) &= \alpha_1(\mu_1 - X(t))dt + \sigma_1 dW_1(t) \\
dY(t) &= \alpha_2(\mu_2 - Y(t))dt + \rho\sigma_2 dW_1(t) + \sqrt{1-\rho^2}\sigma_2 dW_2(t)
\end{aligned} \tag{9}
$$

where α_i $(i = 1, 2)$ are mean-reverting speeds, μ_i $(i = 1, 2)$ are long-term means of log-prices, σ_i $(i = 1, 2)$ are price volatilities, ρ is the instantaneous correlation coefficient between the two price processes, and $W_i(t)$ $(i = 1, 2)$ are independent standard Brownian motions. Regarding parameter estimation, we observe that, over a small time increment Δt, the diffusion terms $\sigma_i dW_i(t)$ $(i = 1, 2)$

of $dX(t)$ and $dY(t)$ in (9) are simply Normal random variables with mean zero and variance $\sigma_i^2\Delta t$ $(i = 1, 2)$. Thus, for price $X(t)$ sampled at time $0, \Delta t, 2\Delta t, \ldots, n\Delta t$, we can regress the log return $\Delta X \equiv X(t + \Delta t) - X(t)$ over Δt onto $X(t)$ scaled by Δt to estimate μ_1 and α_1. The standard error of the residuals provides an estimate of σ_1 (Clewlow and Strickland 2000). The parameters of $Y(t)$ can be estimated in the same fashion.

Case 2: $X(t)$ is a mean-reverting jump-diffusion process and $Y(t)$ is a mean-reverting process:

$$dX(t) = \alpha_1(\mu_1 - X(t))dt + \sigma_1 dW_1(t) + \kappa dq(t) \tag{10}$$

$$dY(t) = \alpha_2(\mu_2 - Y(t))dt + \rho\sigma_2 dW_1(t) + \sqrt{1-\rho^2}\sigma_2 dW_2(t)$$

where α_i, μ_i, σ_i, ρ and $W_i(t)$ have the same interpretations as those in (9), $q(t)$ is a Poisson process with intensity ϕ independent of everything else, and κ denotes a Normal random variable $N(\bar{\kappa}, \gamma)$. $q(t)$ and κ model the occurrence and size of jumps, respectively. To estimate the parameters for $X(t)$ in (10), we adopt a heuristic method outlined in Clewlow and Strickland (2000). The first step is to determine a reasonable threshold level of return beyond which a log-return is considered to be due to a price jump. For instance in the numerical examples in section 5, if a log-return is at least three times of the standard deviation away from the mean of the log-returns, we determine that it is resulted from a price jump. With this threshold, we examine all return data and take out those associated with jumps from the entire data set. We then repeat the previous step with a new threshold based on the mean of the remaining data until no more jump returns can be removed. The mean-reverting and diffusion parameters α_1, μ_1 and σ_1 are estimated from the resulting data set by applying the same procedure as the one used for the simple mean-reverting model (9). We next count the total number of jumps that we remove from the original data set and divide it by the length of the sample period to obtain the jump frequency ϕ. The mean and the standard deviation of the jump size κ are also computed from these jumps.

3.3 Value Function Approximation

For pricing a financial option, there are two classical numerical approaches: pricing by a lattice (Cox, Ross, and Rubinstein 1979) and pricing by Monte Carlo simulation (Boyle 1988). As the valuation problem is framed as pricing a series of real options, both the lattice method and the simulation method can be applied in theory. However, when the underlying price process has a jump component as in the case of (10), it is difficult to construct a end-recombining lattice which converges to the continuous-time model in distribution. Moreover, if the valuation horizon is long, the computational effort of the lattice approach can be huge thus making the computational time prohibitively long. As an alternative, we propose a Monte Carlo simulation based method using value function approximation by a least-squares regression technique, which is developed in Carriere (1996), Longstaff and Schwartz (2001), and Tsitsiklis and Van Roy (2001) for pricing American-style options, to solve the dynamic programming problem of tolling contract valuation.

3.3.1 Commodity Price Process Simulation

To simulate the electricity and the generating fuel prices, we apply the Euler method (Higham 2001) to discretize the continuous-time models defined in (9) and (10). The two correlated price processes in the two cases are simulated through the stochastic difference equations (11) and (12).

Case 1: $X(t)$ and $Y(t)$ are mean-reverting processes:

$$\begin{aligned} X(t+\Delta t) &= X(t) + \alpha_1(\mu_1 - X(t))\Delta t + \sigma_1\epsilon_1\sqrt{\Delta t} \\ Y(t+\Delta t) &= Y(t) + \alpha_2(\mu_2 - Y(t))\Delta t + \sigma_2\rho\epsilon_1\sqrt{\Delta t} + \sqrt{1-\rho^2}\sigma_2\epsilon_2\sqrt{\Delta t} \end{aligned} \qquad (11)$$

Case 2: $X(t)$ is an mean-reverting jump-diffusion process and $Y(t)$ is an mean-reverting process:

$$
\begin{aligned}
X(t+\Delta t) &= X(t) + \alpha_1(\mu_1 - X(t))\Delta t + \sigma_1 \epsilon_1 \sqrt{\Delta t} + (U < \phi \Delta t)(\bar{\kappa} + \gamma \epsilon_3) \\
Y(t+\Delta t) &= Y(t) + \alpha_2(\mu_2 - Y(t))\Delta t + \sigma_2 \rho \epsilon_1 \sqrt{\Delta t} + \sqrt{1-\rho^2}\sigma_2 \epsilon_2 \sqrt{\Delta t}
\end{aligned}
\tag{12}
$$

where ϵ_1, ϵ_2 and ϵ_3 are independent standard Normal random variables, and U is a uniform random variable in [0, 1].

The hourly power prices from 06:00 to 22:00 in every day are called on-peak prices and the rest are called off-peak prices. One important fact about power prices is the time-of-day effect: on-peak prices are much higher than off-peak prices. To reflect this fact, we discretize the time horizon into alternating small intervals of two different lengths. One length is 16-hour corresponding to the peak-hour interval in one day and the other is 8-hour which corresponds to the off-peak interval. Each peak-hour interval is followed by an off-peak interval and vice versa. Finally, we simulate the "daily" prices over these small time intervals according to (11) and (12). The on-peak and off-peak prices are obtained by multiplying the daily prices over on-peak and off-peak intervals with different scaling factors k_{on} and k_{off}, respectively. In section 5, we choose $k_{on} = 1.2$ for the on-peak price and $k_{off} = 0.6$ for the off-peak price so that the average price over 24 hours in one day equals the daily price simulated according to (11) and (12).

3.3.2 Value Function Approximation by Least-squares Regression

In solving V_t through (5)-(8) in either of the two cases, it is a difficult task to compute the exact conditional expectation of the value function. To overcome this difficulty, we propose to approximate the conditional expectation of the value function in our problem by expanding it with respect to a set of complete basis functions and obtaining the expansion coefficients through least-squares regression. Such an approach is developed in Carriere (1996), Longstaff and Schwartz (2001), and Tsitsiklis and Van Roy (2001) for pricing American options with one single exercising decision to make. We show that it is applicable to the tolling agreement valuation problem which involves multiple exercising decisions and a constraint on the total number of allowable exercises as well. The key is to use a truncated expansion of the conditional expected value function with respect to a set of basis functions for the approximations in (5)-(8). As the approximation is a linear combination of the basis functions, a least-squares regression can be applied to obtain the expansion coefficients. Specifically, we regress a conditional expected value function onto basis functions which are functions of the underlying price variables. Polynomial functions up to the third order are chosen as the basis functions for the implementation specified in (13).

Let $\widetilde{V}_t(X_t, Y_t, w, n)$ denote $e^{-r\Delta t}E_t[V_{t+1}(X_{t+1}, Y_{t+1}, w, n)]$. Given a set of simulated sample paths $\Lambda \equiv \{(X_t(\omega_i), Y_t(\omega_i)) : \omega_i \in \Omega \ (i = 1, 2, \cdots, J),\ 0 \leq t \leq T\}$, we regress the vector consisting of $e^{-r\Delta t}V_{t+1}(X_{t+1}(\omega_i), Y_{t+1}(\omega_i), w, n)$ $(i = 1, 2, \cdots, J)$ onto a set of polynomial functions of $(X_t(\omega_i), Y_t(\omega_i))$ as given in the right hand side (RHS) of (13) and obtain the coefficients $\{b_i^t(w, n) : i = 0, 1, \cdots, 9;\ (w, n) \in W_D \times W_R\}$ by a least-squared regression.

$$\begin{aligned} e^{-r\Delta t}E_t[V_{t+1}(X_{t+1}, Y_{t+1}, w, n)] \ \cong\ & b_0^t(w,n) + b_1^t(w,n)e^{X_t} + b_2^t(w,n)e^{Y_t} \\ & + b_3^t(w,n)e^{2X_t} + b_4^t(w,n)e^{2Y_t} + b_5^t(w,n)e^{X_t}e^{Y_t} \\ & + b_6^t(w,n)e^{3X_t} + b_7^t(w,n)e^{3Y_t} + b_8^t(w,n)e^{2X_t}e^{Y_t} \\ & + b_9^t(w,n)e^{X_t}e^{2Y_t} \qquad \forall (w,n) \in W_D \times W_R \end{aligned} \tag{13}$$

We then use the RHS of (13) to approximate $\widetilde{V}_t(X_t, Y_t, w, n)$ and substitute it into the term

$e^{-r\Delta t}E_t[V_{t_{i+1}}(\Theta^S_{t_{i+1}}, \Theta^P_{t_{i+1}})]$ in (5)-(7) to solve for $V_t(X_t, Y_t, w_t, n_t)$ by backwards induction starting with terminal condition (8). Working backwards until reaching the starting time, we obtain the initial value of a tolling agreement.

4 Hedging Tolling Contracts

If markets were complete, then the value function of a tolling agreement plus the cumulative payouts deposited into a bank should have a martingale representation with respect to the filtration generated by the Browian motions and the compound Poisson process in (9) and (10) thus a perfect hedging strategy with continuous trading would exist. The electricity markets are inherently incomplete and the continuous trading is only an ideal assumption. Nevertheless, delta-hedging strategies derived with the continuous-trading assumption and implemented through discrete-trading still provide greate practical value. In this section, we present a heuristic delta-hedging strategy for hedging a tolling contract.

As explained in section 3, tolling options in a tolling agreement closely resembles a strip of spark spread call options with maturity time spanning through the contract period. This observation prompts the idea of using delta positions derived for the European-style spark spread options to construct a hedging portfolio of a tolling contract.

To illustrate how the proposed hedging strategy works, we present an example of hedging an one-year tolling contract. First, we replicate the one-year tolling agreement by a portfolio of spark spread call options with maturity times ranging from 1-month to 12-month and strike heat rate equaling the plant's average heat rate. In the spark spread option portfolio, options of different maturities all have the same number of shares which equals the capacity of the underlying plant.

Also, we adjust the option positions every day to match the value of the option portfolio to that of the tolling agreement.

Under a typical delta-hedging scheme, to hedge an option on a generic underlying asset S_t with maturing time T, an instrument with payoff S_T at time T is needed. The instrument is usually the underlying asset itself if it is traded (e.g., in the case of hedging stock options). However, for an option on electricity, we cannot hold electricity to get S_T at time T since electricity is non-storable. Instead, an electricity futures contract with a payoff of $S_T - F_E$ at time T is needed for hedging electricity options, where F_E is the electricity futures price. We use a futures contract and a bond to construct a synthetic security that is traded and pays out the electricity price S_T at time T for implementing delta-hedging for a tolling agreement. Specifically, each share of the synthetic security consists one share of futures contract and F_E shares of a riskless bond paying \$1 at time T. As the time-T payoffs of the futures contract and the bond are $S_T - F_E$ and F_E, respectively, they yield the synthetic security a combined payoff of S_T. The price of the synthetic security at time 0 is $F_E \cdot e^{-rT}$ since the cost of entering into futures contracts is zero, where r is the risk-free rate.

We next construct a hedging portfolio out of a bank account and 24 synthetic securities consisting of 12 monthly electricity futures and 12 monthly nature gas futures over a time period of one year. For $t \in [0, 1]$, let V_t denote the remaining value of a tolling contract and M_t denote the month in which t falls. The hedging portfolio of V_t is denoted by $(\overrightarrow{\delta}_E(t), \overrightarrow{\delta}_G(t), B_t)$ where $\overrightarrow{\delta}_E(t) \equiv (\delta_E^{M_t}(t), \delta_E^{M_t+1}(t), \cdots, \delta_E^{12}(t))$ represent the shares of synthetic securities based on electricity futures with maturity date $T_{M_t}, T_{M_t+1}, \ldots, T_{12}$ (while time-to-maturity ranges from 1-month to $(13 - M_t)$-month), $(\delta_G^{M_t}(t), \cdots, \delta_G^{12}(t))$ represent the shares of the corresponding synthetic natural gas contracts, and B_t represents the balance of a bank account. Let $C_{t,T}$ denote the time-t value of a spark spread call with a payoff of $\max(S_E(T) - H \cdot S_G(T), 0)$ at maturity T with H being the average heat rate of the power plant and T ranging from T_{M_t} to T_{12}. The shares of $C_{t,T}$ $(T = T_{M_t}, T_{M_t+1}, \cdots, T_{12})$ to hold at time t for hedging V_t is obtained by solving for Q_t in (14).

$$V_t = Q_t \cdot \sum_{i=M_t}^{12} C_{t,T_i} \tag{14}$$

Therefore $\delta_E^i(t)$ is given by

$$\delta_E^i(t) = Q_t \cdot \frac{\partial C_{t,T_i}}{\partial F_E^i} \qquad \forall i = M_t, M_t+1, \cdots, 12 \tag{15}$$

where F_E^i is the electricity futures price of maturity T_i and C_{t,T_i} is calculated by the transform method in Deng (1999). The delta positions of the gas synthetic securities are obtained by solving (16) utilizing (15). That is, letting the total value of synthetic securities equal to the value of the tolling agreement less operating profit (loss) for each day.

$$Q_t \cdot C_{t,T_i} = e^{-r(T_i - t)}(\delta_E^i(t)F_E^i(t) + \delta_G^i(t)F_G^i(t)) \quad \forall i = M_t, M_t+1, \cdots, 12. \tag{16}$$

We re-balance the futures positions at the current futures prices at the end of each trading period Δt. Trading gains/losses get deposited or withdrawn from the bank account B_t (assuming no limit on the size of deposit and withdrawal). The balance of the bank account at the end of each trading period t is given by:

$$
\begin{aligned}
B_t \;=\; & \sum_{i=M_{t-1}}^{12} \{\delta_E^i(t-\Delta t)(F_E^i(t) - F_E^{i-1}(t-\Delta t)) + \delta_G^i(t-\Delta t)(F_G^i(t) - F_G^{i-1}(t-\Delta t))\} \\
& + \sum_{i=M_{t-1}}^{12} \{\delta_E^i(t-\Delta t)F_E^i(t-\Delta t)e^{-r(T-t)} + \delta_G^i(t-\Delta t)F_G^i(t-\Delta t)e^{-r(T-t)}\} \\
& - \sum_{i=M_t}^{12} \{\delta_E^i(t)F_E^i(t)e^{-r(T-t)} + \delta_G^i(t)F_G^i(t)e^{-r(T-t)}\} + B_{t-\Delta t} \cdot e^{r\Delta t}. \qquad (17)
\end{aligned}
$$

By the end of the contract period, the balance $B_{T_{12}}$ represents the cumulative hedging error.

5 Numerical Examples and Managerial Insights

To validate the valuation model, we apply it to a one-year tolling agreement on a hypothetical natural gas fired power plant with a capacity of 150 MW subject to the operational constraints and the maximum restart constraint discussed in section 3. Suppose the underlying plant sells wholesale power to the Electric Reliability Council of Texas (ERCOT) region of the U.S. and purchases its fuel supply from the Henry Hub natural gas market. Thus we use ERCOT power price data and Henry Hub gas price data to estimate parameters for the stochastic price processes in (9) and (10).

We assume that the ramp-up cost rate function $c(y)$ has the following form (Deng and Oren 2003).

$$c(y) = \underline{Q} * \underline{Hr} * e^y + M \qquad (18)$$

where M is a constant. Table 1 summarizes the set of parameters characterizing the underlying power plant. In addition, we test our model for the maximal allowable number of restarts being 3 and 6 per year, respectively.

C_{start}	C_{down}	$\underline{Q} : \overline{Q}$	$\underline{Hr} : \overline{Hr}$	M	r
\$2000	\$1000	0.2 : 1	1.38 : 1	1	5%

Table 1: Parameters for a Hypothetical Natural Gas Fired Power Plant

We examine the effects of different electricity price modeling assumptions on the valuation of a tolling agreement using the two cases specified in section 3.2.

- Case 1: both power and gas price processes are mean-reverting (MR) processes.
- Case 2: the power price process is a mean-reverting jump-diffusion (MRJD) process and the gas price process is a MR process.

While both power and gas prices have the mean-reverting feature, power prices are much more volatile than other commodity prices (including gas price) as demonstrated by the salient spikes and jumps in the historical power prices. A MRJD process is a more realistic assumption for modeling power prices than a simple MR process. Case 1 serves as a benchmark case in which the price jumps are not explicitly modelled and case 2 offers the contrasting results with explicit modeling of price jumps. Tables 2 and 3 report parameters for the power and the gas price processes, respectively, which are estimated from the historical ECORT daily electricity prices and Henry Hub daily gas prices. The correlation coefficient ρ between the power price and the gas price is set to be the sample correlation 0.177 in both cases. The initial prices of electricity and natural gas are sampled from the historical data as \$34.7 per MWh and \$3 per MMBtu, respectively. Figure 2 plots three simulated sample paths for power and gas price models (9) and (10) with the estimated parameters in tables 2 and 3.

We compute the value of the one-year tolling agreement for different levels of heat rate $\overline{Hr}$: 7.5, 10.5

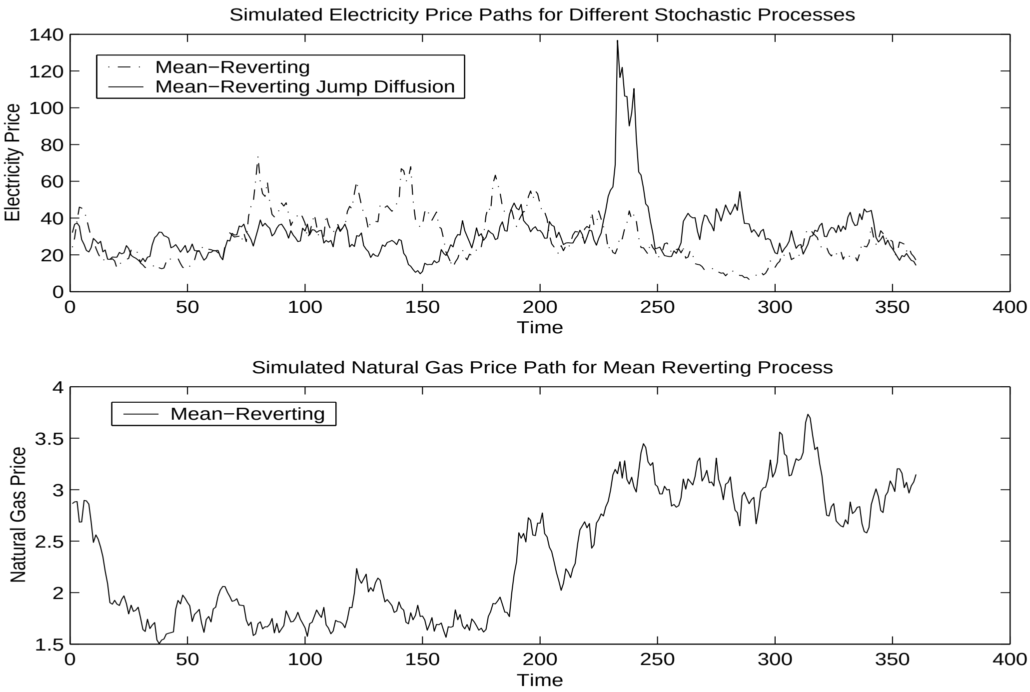

Figure 2: Simulated Sample Paths of Different Electricity and Gas Price Models

	μ	σ	α	ϕ	$\bar{\kappa}$	γ
MR	3.5527	0.1507	0.0651	–	–	–
$MRJD$	3.5304	0.1299	0.0584	0.0281	0.0483	0.2566

Table 2: Parameters for the Power Price Process

and 13.5 MMBtu/MWh. Table 4 reports the estimates of the time-0 value of the tolling agreement and the corresponding standard errors based on 2000 Monte Carlo sample paths of the power and gas price processes. The contract values in case 1 are slightly lower than those in case 2. A MRJD electricity price model is more realistic than a simple MR model, given the empirical features of power prices such as mean-reversion, jumps, and spikes (Deng 1999). Therefore, we argue that the contract value obtained with the MRJD power price model is a better approximation to the fair value than the value obtained with the MR model.

To put our valuation model into perspective, a couple of large energy merchant firms in the elec-

	μ	σ	α
MR	1.3638	0.0468	0.0087

Table 3: Parameters for the Gas Price Process

tricity industry in the U.S. paid about 15 million dollars per year in the early 2000's for tolling contracts in the ECORT region written on a 150 MW power plant with an average heat rate of around 8.0 MMBtu/MWh, a similar cost structure to our hypothetical power plant, and about 20 restarts allowed. Such a tolling premium is quite close to the estimated values in table 4. However, as we incorporate only the major, not all, operational characteristics and constraints into our valuation model, the true value of a tolling agreement should be lower than our model-predicted value. Another important point is that the values in table 4 are computed under the assumption of underlying prices processes being stationary which means that the electricity price (or, fuel price) regime remains the same over time. However, if either of the price regimes were to change from the present state to a fundamentally different one in the future due to dramatic increases in energy demand (or, a flood of newly built generation capacity coming online), then the value of a tolling agreement would be significantly different from the value obtained under the stationary price assumption (this is further discussed in the conclusion section).

$N = 3:\ \overline{Hr}$	7.5	8.0	10.5	13.5	$N = 6:\ \overline{Hr}$	7.5	8.0	10.5	13.5
MR	15.02	14.94	8.09	4.06	MR	16.29	15.08	8.91	4.87
(STD-1)	0.28	0.33	0.27	0.18	(STD-1)	0.32	0.32	0.29	0.20
MRJD	15.40	15.18	8.33	4.11	MRJD	16.79	15.31	9.48	4.79
(STD-2)	0.32	0.34	0.28	0.17	(STD-2)	0.34	0.34	0.31	0.21

Table 4: Estimated Values (in $ millions) of the Tolling Agreement for Different Price Models and Heat Rates (Max_Restarts N: 3 and 6; STD: standard error).

Figures 3 and 4 plot the optimal action regions for executing this tolling contract at the end of the third month and the sixth month over the 12-month contract period. Specifically, it is optimal for the contract holder to start taking electricity (i.e., exercising a tolling option, or turning on the plant) in the region to the south-east of the boundary line formed by the $\times$'s and stop taking electricity (i.e., terminating an exercised tolling option, or shutting down the plant) in the region to the north-west of the boundary line formed by the circles. We term these two regions as the "tolling" and the "no-tolling" regions. We choose 4 contract states (shown in the figures) to plot the corresponding "tolling" and "no-tolling" regions for the heat rates being 7.5 and 13.5 MMBtu/MWh. Turn-on and turn-off boundaries are clearly shaped in each plot.

These computational results shed lights on the structure of the optimal execution strategies for a tolling agreement and offer valuable insights for managing and operating such contracts. First of all, the optimal actions for operating and managing a tolling contract are separated into "tolling" and "no-tolling" regions in the plane of all possible price pairs of electricity and the fuel by some curves (or, boundaries). With this insight, a tolling contract holder knows that the optimal execution strategy of the contract is governed by some threshold curves in the form of certain functional relationships between the electricity price and the fuel price. Thus she or he can efficiently identify the optimal action regions for operational guidance by testing various relationships between the power price and the fuel price. We also observe that, as the remaining time of the contract gets shorter, both the "no-tolling" and the "tolling" regions get larger meaning that it is optimal for the contract holder to exercise the "tolling" and "shutting down" options more frequently as the time approaches contract expiration. On the other hand, the contract holder shall be patient in determining whether to start taking electricity or to terminate an exercised tolling option immediately in the early stage of the tolling contract since there is only a limited number of restart opportunities available. The two regions also get larger as the remaining number of re-starts gets larger since a larger number of re-starts reduces the opportunity cost of exercising a tolling option/shutdown option and it leaves more flexibility with the contract holder in determining the best time to exercise a tolling or shutdown option.

All price pairs falling in between the "tolling" region and the "no-tolling" region constitutes a "no action" band in the sense that, if a pair of the electricity price and the fuel price belongs to this band, the optimal action for the contract holder is to maintain status quo under current market conditions: either to continue taking the electricity output if under the obligation of an exercised tolling option, or to keep putting off the exercising of a tolling option if not under any tolling option's obligation. Based on numerical experiments, we observe that the area of the no-action band shrinks as the remaining contract time gets shorter and it expands as the number of remaining re-starts gets smaller. This implies that, as the tolling contract approaches expiration, the holder should be more active in determining which best action to take rather than passively sticking to its current operating state. On the other hand, when the number of remaining restarts get smaller, the holder should be patient in determining the optimal action for now so as to leave the limited optionality for the best time to capture the most economic benefits.

The mean, standard deviation, and 90% confident interval of the cumulative hedging errors for case 2 are reported in table 5. We observe that while the mean and the standard deviation of hedging errors decrease as the heat rate increases, the percentage of the hedging error with respect to the tolling contract value decreases as the heat rate decreased. Namely, the percentage hedging error of the delta-hedging strategy for a tolling contract written on an efficient power plant (i.e., with low heat rate) is smaller than that for a contract written on an inefficient plant. This is as expected because the hedging strategy for a tolling contract is designed based on the hedging portfolio of a

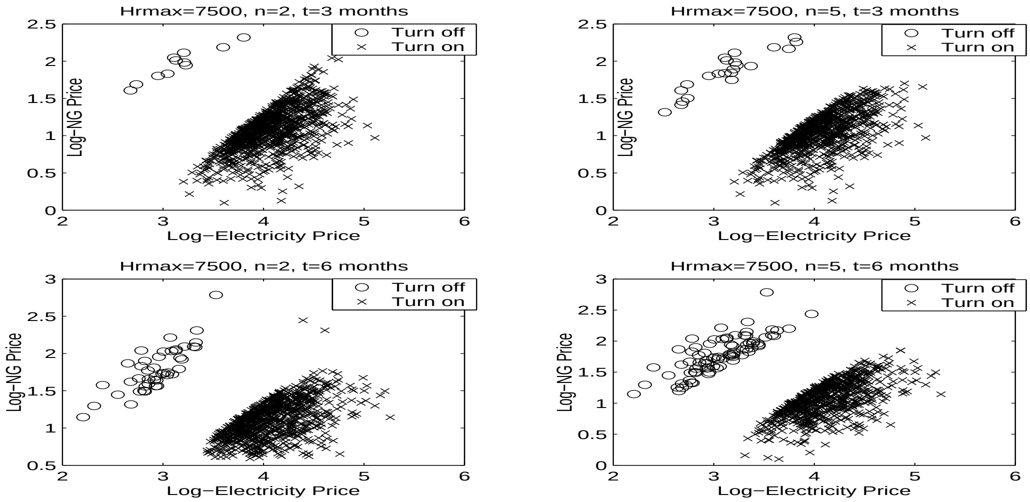

Figure 3: Optimal boundaries for heat rate 7.5

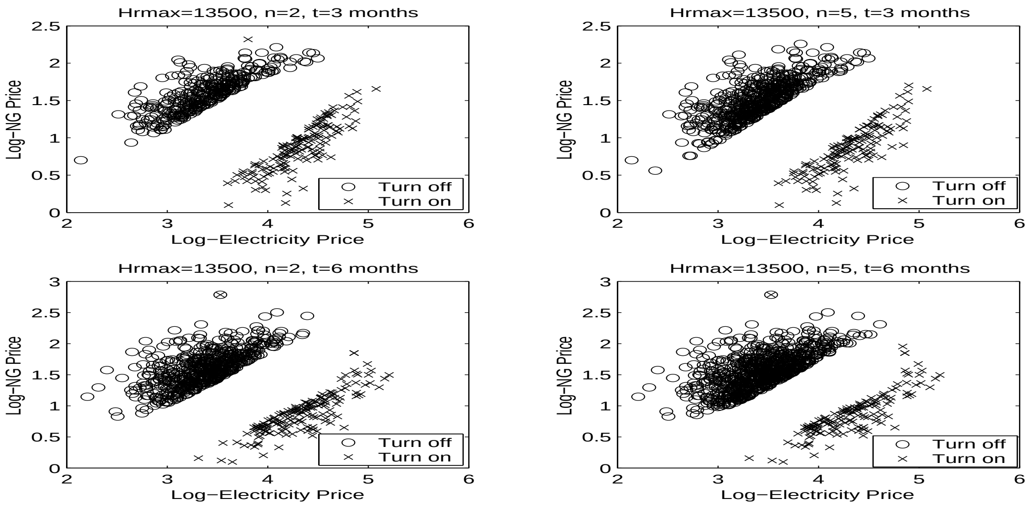

Figure 4: Optimal boundaries for heat rate 13.5

series of spark spread call options. For a tolling agreement on an efficient power plant, it can be well approximated by a series of spark spread call options.

The two panels in figure 5 show the histograms of cumulative hedging errors with x-axis indicating the dollar amount for different heat rates with the maximal number of restarts being 6. For heat rate being 7.5 and 13.5 MMBtu/MWh, the distributions of hedging errors skew to the right.

$N = 3 : \overline{Hr}$	7.5	10.5	13.5
$Mean$	1.03	1.02	0.66
STD	2.01	1.56	1.02
$Conf_{5\%-95\%}$	(-1.79, 4.67)	(-0.83, 3.98)	(-0.42, 2.60)
$N = 6 : \overline{Hr}$	7.5	10.5	13.5
$Mean$	1.03	0.95	0.69
STD	2.06	1.48	1.05
$Conf_{5\%-95\%}$	(-1.86, 4.90)	(-0.88, 3.72)	(-0.46, 2.65)

Table 5: Statistics of Hedging Errors for Case 2 (Max_Restarts N: 3 and 6)

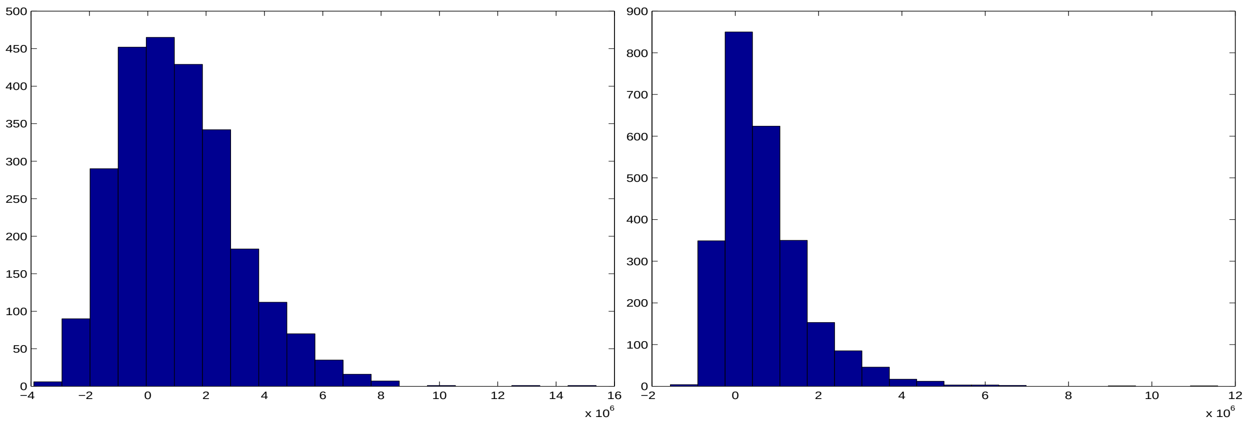

Figure 5: Histogram of Hedge Errors: Hr = 7.5 (left) vs. Hr = 13.5 (Right)

6 Conclusion

Electricity tolling agreements, as well as other structured transactions, have played important roles in facilitating risk-sharing and risk-mitigation among independent power producers, utility companies and unregulated power marketers in the restructured power industry. As power markets continue to evolve, market information and variables needed for pricing these complex structured transactions shall become more and more transparent. However, at the present time, we need to work with assumptions, which are plausible but yet to be empirically justified, in investigating the problem of pricing and hedging customized electricity contracts such as tolling agreements. Under these assumptions, we formulate a real options based valuation model and solve it through a dynamic programming algorithm with value function approximation by least-squares regression. The valuation model yields a fairly accurate approximation to the market value of a tolling agreement as gauged by the limitedly available market transaction data.

Through numerical examples, we also examine the effects of different power price assumptions on the tolling contract valuation. The simple mean-reverting power price model results in a slightly lower tolling premium than does the mean-revering jump-diffusion power price model. The well-known empirical features of electricity prices make the jump-diffusion assumption a more realistic choice for pricing tolling agreements. One other crucial factor affecting the valuation is the stationarity assumption on the underlying price processes. Electricity prices are indeed non-stationary due to seasonality effects, fundamental changes in electricity supply and demand, evolutions in market designs, and other modifications to regulatory policies. Tolling contracts signed in the late 1990's turn out to be significantly overvalued as they were priced under the assumption that electricity prices would fluctuate in a high-price regime which was made based on the limited market price data by that time. The risk of a regime shift in electricity prices from a high-price regime to a low-price one was not properly incorporated thus the overvaluation. To mitigate such over- or under-valuation risks, one needs to account for non-stationarity in modeling the underlying prices, particularly when valuing a contract with a long horizon such as 10 to 15 years.

One point worth mentioning is that the real options valuation formulation is applicable under quite general price modeling assumptions. If the underlying prices are non-stationary with jumps, most numerical schemes such as the lattice approach and the partial differential equation approach would encounter difficulties in solving real options valuation problems. In such cases, the value function approximation scheme with Monte Carlo simulation is the only viable approach for solving the valuation problem.

Regarding the other important issue on hedging a tolling agreement, we propose a heuristic delta-hedging scheme by utilizing futures contracts. Starting with the premium of a tolling contract, we construct a portfolio for hedging (or, replicating) the contract using electricity and generating fuel futures and a riskless bank account. The hedging position in electricity futures is calculated through the closed-form pricing formula of a spark spread call option and the position in generating fuel futures is obtained by matching the value of the hedging portfolio to that of the tolling contract. We then continuously re-balance the hedging portfolio by adjusting the number of shares of each security held. In the numerical examples, this straightforward hedging scheme results in an around 6% hedging error with respect to the tolling contract value for an efficient power plant (i.e. heat rate is 7.5) and a 14% $\sim$ 16% hedging error for an inefficient power plant (i.e. heat rate is 13.5). The less/more efficient a power plant is or the less/more frequent the restarts are allowed, the larger/smaller percentage error the hedging scheme yields.

As for future work, we plan to carry out rigorous empirical tests on one major assumption of the tolling contract valuation model: the specification of a power price model. Another fruitful direction deserving further pursuit is the search for a more efficient and accurate hedging scheme for a tolling contract, especially when the underlying power plant is inefficient meaning that it has a very large operating heat rate.

Acknowledgment

This research is supported by the NSF Grant ECS-0134210. The authors would like to thank Kenneth Deng for his valuable comments and assistance.

Appendix

A Notations

In addition to those defined in section 3.2, the following notation are used in the paper.

a_t	:	operational action taken by a plant operator at time t,
$b_i^t(w,n)$	:	regression coefficients for the SDP value function approximation at time t,
c_{start}/c_{down}	:	start-up/shut-down costs,
D	:	ramp up delay period,
K_D	:	number of Δt in the ramp up delay period D,
H	:	strike heat rate,
$\overline{Hr}/\underline{Hr}$	:	the heat rate associated with $\overline{Q}$ and $\underline{Q}$, respectively,
k_{up}/k_{down}	:	the scale factor for simulating on-peak/off-peak prices,
N	:	the total number of allowed restarts,
n_t	:	the remaining number of allowed restarts at time t,
$\overline{Q}/\underline{Q}$	:	maximum/minimum output levels of a power plant,

r	:	discount rate,
T	:	contract expiration time,
w_t	:	the operational state of the power plant,

α_i : rate of mean-reversion in the price of commodity i,

μ_i : long-run mean price for commodity i,

σ_i : volatility parameter of price process,

ρ : correlation coefficient between electricity and NG prices,

κ : size of jumps in a Poisson process,

ϕ : arrival intensity of a Poisson process,

$\bar{\kappa}$: mean size of jumps,

γ : standard deviation of jump size,

References

Black, F., and M. S. Scholes, 1973, "The pricing of options and corporate liabilities," *Journal of Political Economy*, 81, 637–654.

Boyle, P. P., 1988, "A Lattice Framework for Option Pricing with Two State Variables," *Journal of Financial & Quantitative Analysis*, pp. 1–12.

Carriere, J. F., 1996, "Valuation of the Early-Exercise Price for Derivative Securities using Simulations and Nonparametric Regression," *Insurance: Mathematics and Economics*, 19, 19–30.

Clewlow, L., and C. Strickland, 2000, *Energy Derivatives - Pricing and Risk Management*, Lacima.

Cox, J., S. Ross, and M. Rubinstein, 1979, "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7, 229–263.

Deng, S. J., 1999, "Stochastic Models of Energy Commodity Prices and Their Applications: Mean-reversion with Jumps and Spikes," *POWER Working Paper*, University of California Energy Institute.

Deng, S. J., B. Johnson, and A. Sogomonian, 2001, "Exotic electricity options and the valuation of electricity generation and transmission assets," *Decision Support Systems*, 30, 383–392.

Deng, S. J., and S. S. Oren, 2003, "Incorporating Operational Characteristics and Start-up Costs in Option-Based Valuation of Power Generation Capacity," *Probability in the Engineering and Informational Sciences*, 17, 155–181.

Dixit, A. K., and R. S. Pindyck, 1994, *Investment under Uncertainty*, Princeton University Press, New Jersey.

Harrison, J. M., and D. Kreps, 1979, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381–408.

Higham, D. J., 2001, "An Algorithmic Introduction to Numerical Simulation of Stochastic Differential Equations," *SIAM Review*, 43, 525–546.

Jaillet, P., E. Ronn, and S. Tompaidis, 2001, "Valuation Commodity Based Swing Options," *Working Paper*, University of Texas, Austion.

Keppo, J., 2004, "Pricing of Electricity Swing Options," *Working Paper*, University of Michigan.

Longstaff, F. A., and E. S. Schwartz, 2001, "Valuing American Options by Simulation: A Simple Least-squares Approach," *The Review of Financial Studies*, 14, 113–147.

Merton, R. C., 1973, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.

Thompson, A., 1995, “Valuation of path-dependent contingent claims with multiple exercise decisions over time: the case of take-or-pay,” *Journal of Quantitative and Computative Analysis*, 30, 271–293.

Tseng, C.-L., and G. Barz, 2002, “Short-term Generation Asset Valuation: a Real Options Approach,” *Operations Research*, 50, 297–310.

Tsitsiklis, J. N., and B. Van Roy, 2001, “Regression methods for pricing complex American-style options,” *IEEE Transactions on Neural Networks, forthcoming.*

Woo, C.-K., A. Olson, and R. Orans, 2004, “Benchmarking the Price Reasonableness of an Electricity Tolling Agreement,” *Electricity Journal*, 17, 65–75.

Wood, A., and B. Wollenberg, 1984, *Power Generation, Operation and Control*, John Wiley & Sons, New York.