Valuation and Optimal Interruption for Interruptible Electricity Contracts*

Ross Baldick† Sergey Kolos‡ Stathis Tompaidis§

*Acknowledgments: We would like to acknowledge the helpful comments and suggestions of seminar participants at the University of Texas at Austin, and the conference on "Modeling, Optimization and Risk Management in Finance" at the University of Florida. We also are thankful for financial support from NSF grant 0082506.

†Department of Electrical and Computer Engineering, University of Texas at Austin, Austin, TX 78712, baldick@ece.utexas.edu

‡Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX 78712, skolos@mail.utexas.edu

§MSIS Department, McCombs School of Business, University of Texas at Austin, Austin, TX 78712 stathis.tompaidis@bus.utexas.edu

Valuation and Optimal Interruption for Interruptible Electricity Contracts

ABSTRACT

We consider interruptible electricity contracts issued by a distributor of electricity, that allow for interruptions to electric service in exchange for either an overall reduction in the price of electricity delivered, or for financial compensation at the time of interruption. We provide an equilibrium model to determine electricity prices, based on stochastic models of supply and demand. In the context of this model we quantify the value of interruptible contracts and describe the optimal interruption strategy. Our numerical results indicate that, in a deregulated, competitive, market, interruptible contracts can alleviate supply problems associated with spikes of price and demand.

Introduction

The market for one of the most important commodities in today's economic environment, electricity, has recently undergone significant changes. For most of its North American history, a market with a few, heavily regulated, vertically integrated, participants, the electricity market is currently transitioning towards a restructured market with many more market players, most of which provide only a small part of the services provided by the original participants. While under the regulated environment risks to the market participants were mitigated by the mechanism of cost recovery, under deregulation, and facing competition, such cost recovery is unlikely, creating the need for the use of financial risk management tools and techniques. To mitigate financial risk, new financial products have been developed. Among them is the interruptible contract, which allows one party to renege on its obligation to provide electricity to the other party a certain number of times over a certain period of time. In this paper we study how interruptible contracts may help a large distributor reduce its exposure to fluctuations in the demand and supply of electricity.

Interruptible contracts have existed since at least the early 1990's, but have become a risk management tool only after the two California electricity crises, in the summer of 1998 and the winter of 2001. Prior to 1998, while interruptible contracts provided the right to interruption, these rights were almost never exercised, leading to a skewed perception of their risk among electric customers.[1] In the literature, interruptible contracts have been described in the paper by Kamat and Oren (2002). In that paper, a simple form of an interruptible contract is presented, in which one party can interrupt the other once over two possible interruption opportunities, and where it is assumed that interruption does not influence the spot electricity price.

[1]Indeed, since signing up for an interruptible contract provided a discount in the retail price of electricity, many customers that never intended to interrupt, such as hospitals, schools and nursing homes signed their electric load on interruptible contracts. Unsurprisingly, when called to interrupt, these customers refused to do so (see the report by the Energy Division of the Public Utilities Commision of California (2001)).

In our work, we extend and generalize the paper of Kamat and Oren (2002) in several directions. First, we allow for the possibility of multiple interruptions over many possible interruption dates, possibly with daily frequency. Second, we allow for different forms of payment for the right of interruption. Different forms of payment may generate differences in the optimal interruption policy, since in some cases the cost of interruption is already sunk. Finally, the most important difference involves the impact of the interruption on the spot price of electricity. While Kamat and Oren (2002) considered reduced form models for electricity prices (either geometric Brownian motion, or a mean-reverting process with jumps), we construct a structural model in which the spot price of electricity is determined by supply and demand. The importance of using a structural model that incorporates demand in determining electricity prices, is due to the fact that much of the benefit from interrupting a contract comes not from avoiding servicing the interrupted load, but from reducing the total load to the system, leading to system-wide lower prices. This feature is very valuable to a distributor of electricity that needs to resort to the spot market to cover spikes in demand (that are typically followed by spikes in price). By using its rights to interrupt, the distributor is able to effectively reduce demand, as well as the number of times that it resorts to the spot market and the number of spikes in electricity prices associated with high demand.

The structural model we develop is able to generate many of the observed characteristics of electricity prices. In particular, we can generate both mean reversion in electricity prices, as well as short-lived spikes. We attribute mean-reversion in prices to mean-reversion in temperature, which drives demand, while spikes are generated by a two-regime model for generation supply. Using this structural model we are able to numerically value interruptible contracts and determine the optimal interruption policy.

The rest of the paper is organized as follows: Section I describes the market setting as well as the different forms of interruptible contracts we consider. Section II describes a structural model for electricity prices that links electricity demand and generation supply. The demand is determined from average daily temperature, while supply can fluctuate due to outages or transmission constraints. The model is calibrated with data from ERCOT.[2] In Section III we formulate a stochastic control problem for the valuation of interruptible contracts from the point of view of a risk-neutral distributor, and describe the optimal interruption strategy as well as the value for the different forms of interruptible contracts. Section IV concludes. In Appendix A we collect all the notation used throughout the paper.

I. Model Setting

A. Market description

We consider the case of a large distributor of electricity that is obliged to serve all customers in a specified geographic area. To compensate for this obligation, the distributor is allowed to charge a fixed retail price per unit of electricity, p_{retail} to each of its customers. The distributor has available a certain generation capacity L_{internal}, either through the ownership of generators or through forward purchase agreements. We assume that the cost of the electricity available to the distributor is fixed in advance and does not depend on the spot price of electricity. When demand is higher than the energy available to the distributor, the distributor is obliged to serve the demand through purchases in the spot electricity market.[3,4]

Regarding the customers of the distributor, we assume that they can only purchase electricity from the distributor and can only use electricity for consumption; i.e., they cannot resell it. The customers belong to one of two categories: they are either "residential", with fluctuating

[2]ERCOT is the Electric Reliability Council Of Texas, and covers almost all of Texas; see `www.ercot.com`.

[3]In our model we assume that the cost for the generation available to the distributor is sunk; i.e, the distributor utilizes all power available from its own generators first, and then turns to the energy market. In the event that the generation available is greater than the load, this assumption implies that the utility company can sell the surplus in the spot market. In the examples we consider, we focus on situations where the distributor almost never has enough supply of electricity available to serve the entire demand without resorting to the spot market. In practice, capacity may be purchased in advance and be truly sunk, but generation of energy incurs additional costs that may be avoidable.

[4]This market setting is very similar to the one faced by Pacific Gas and Electric and Southern California Edison shortly after electricity deregulation in California.

demand $L_{\text{residential},t}$; or, they are "industrial" with constant demand $L_{\text{industrial}}$.[5] Total demand at time t is equal to $L_{\text{residential},t} + L_{\text{industrial}}$. We abstract from the daily variation in demand by assuming that $L_{\text{residential},t} + L_{\text{industrial}}$ represents peak demand, or the average demand during on-peak hours.

B. Interruptible contracts

There are several variants of interruptible contracts offered by distributors of electricity. In its most general form, an interruptible contract between a distributor and a customer allows the distributor to interrupt part or all of the supply of electricity to the customer over some period of time in exchange for some form of pecuniary compensation.[6]

We focus on two particular types of interruptible contracts, that appear to be among the most common. A detailed description of these contracts, as well as background on their use in California is available from the report of the Energy Division of the California Public Utilities Commision (see California (2001)). We assume that interruptible contracts are between the distributor and "industrial" customers only, and that upon request, the customer always curtails the requested load.

The first form of an interruptible contract, which we term a *pay-in-advance* contract, allows the distributor to interrupt a given percentage of an "industrial" customer's load a fixed number of times over the life of the contract. In exchange, the customer receives a discount on the retail price of electricity for the customer's entire load, $L_{\text{under\_contract}}$, and pays p_{reduced} per unit of electricity, rather than p_{retail}. Typical values for the parameters of this contract are a 15% discount on the retail price in exchange for 10 daily interruptions of 20% of the customer's load over the period of one year. Other constraints may also exist; e.g., the num-

[5]Under this specification, "industrial" customers may include both industrial and commercial users of electricity. In fact, industrial demand may also vary with time, complicating the design of the interruptible contracts.

[6]As a matter of fact, the distributor does not physically interrupt the customer, but rather gives the customer an advance notice, typically between 30 minutes and 24 hours, requesting curtailment of the customer's load. Failure of the customer to curtail the load to the specified level can lead to severe penalties, effectively resulting in the interruption of the customer's load.

ber of consecutive days of interruption may be limited, or no more than a certain number of interruptions may occur over a short period of time.

The second form of an interruptible contract, which we term the *pay-as-you-go* contract, allows the distributor to interrupt part of a customer's load a fixed number of times in exchange for compensation based on the load interrupted. Typical values for the parameters of this contract allow for 10 interruptions with compensation, p_{fine}, ranging from \$150 per MWh to \$600 per MWh of interrupted electricity, depending on whether notice of interruption is given 24 hours in advance or one hour in advance.

Assuming that the number of interruptible contracts signed between a distributor and "industrial" customers is large, and that the load interrupted under each contract is small, a distributor may neglect some of the constraints of individual contracts by pooling all the contracts together. For example, the number of times a particular customer may be interrupted is not relevant for the distributor, so long as the distributor is careful to rotate interruptions between all its customers. From the distributor's point of view, pooling simplifies the management of the portfolio of interruptible contracts. For each type of interruptible contract, the distributor need only keep track f the maximum amount available for daily interruption and of the total remaining amount of interruption until the end of the year.[7]

II. A Structural Model for Electricity Prices

While much of the literature on the stochastic process followed by electricity prices has focused on reduced form models that mimick the observed price behavior (see Pilipovic (1997), Deng (1999), Deng (2000), and Kholodnyi (2000)), such models are of limited value for the problem we consider. Intrinsically, in a reduced form model, the price process is not influenced by the actions of market participants, while, in the case of a large distributor with interruptible contracts, the interruption has the effect of lowering demand as well as the expected spot

[7]We assume that all interruptible contracts are effective over the same period.

price. Due to this interaction between distributor interruption and electricity price, we develop a structural model of the electricity market, where prices are determined by matching supply and demand.[8]

Due to the difficulty of relating real-time electricity prices to consumers, as well as the fact that almost all consumers of electricity have fixed price contracts, we assume that demand is inelastic, i.e. it does not depend on the spot electricity price. Given inelastic demand, it is important that electricity producers do not collude. We assume a competitive market for the generation of electricity, where each generator submits information about the amount of electricity it is willing to provide and the price associated with each unit of that electricity. The system operator in this model aggregates the information from each electricity provider, and creates the supply curve, where available generation is ranked according to price. The demand determines which of the generators are dispatched.[9] The price paid to each dispatched generator is the price of the marginal dispatched generator.[10] In the structural model we propose, we model demand and supply separately. We calibrate the model with data from the ERCOT area during weekdays in the summer months, since, in the case of ERCOT, that is the period when electric loads are very high and when interruption is likely to occur.

A. Demand

Stylistic facts concerning demand of electricity are that it is strongly seasonal (with daily, weekly, and annual patterns), strongly mean-reverting, and highly predictable. Demand is influenced by environmental factors, such as temperature and humidity, as well as population and industrial activity. In this paper we assume that demand has two components: one that is relatively stable, due to "industrial" customers, and one that is volatile and is due to "res-

[8]A similar structural model for the PJM area was proposed by Skantze, Gubina, and Ilic (2000).

[9]Sometimes there may be a violation of the order in the supply curve, due to congestion of the transmission system, transmission constraints, or transmission failures. We abstract from this problem by introducing random fluctuations to the supply curve.

[10]This compensation scheme is efficient, in the sense that it removes the incentive of generators to ask for prices above their costs, since they then run the risk of not being dispatched.

idential" customers.[11] We model demand fluctuations of the residential customers in terms of temperature fluctuations, which is the most important driving factor of demand in ERCOT during the summer.

A.1. Temperature Model

We use a model for forecasting temperature similar to the one introduced by Cao and Wei (2000 a), and Cao and Wei (2000 b) (see also Campbell and Diebold (2002)). In our model the deviation of the actual from the average temperature over the next day, $t+1$, is a function of the deviation of the actual from the historical average temperature today, t, and the deviation of the actual from the historical average temperature over the previous day, $t-1$. [12,13] The model allows for stochastic fluctuations around the historical average, with magnitudes that depend on the time of the year, and is described by the following equations:

$$\begin{aligned}
\Delta^T_{t+1} &= \rho^T_1 \Delta^T_t + \rho^T_2 \Delta^T_{t-1} + \sigma^T_t \varepsilon^T_t \\
\sigma^T_t &= \sigma^T_{(0)} - \sigma^T_{(1)} |\sin(\pi \frac{t+\phi}{365})| \\
\varepsilon^T_t &\sim \text{iid}(N(0,1)
\end{aligned} \tag{1}$$

where $\Delta_t = T_t - \bar{T}_t$, T_t is the actual temperature for day t, $\bar{T}_t$ is the average temperature for day t, and ρ^T_1, ρ^T_2 are the autocorrelation coefficients for average temperature. The magnitude of the random fluctuations is seasonal, with a fixed term $\sigma^T_{(0)}$ and a seasonal term of magnitude $\sigma^T_{(1)}$.

[11] In this context a customer with stable demand is considered an industrial customer, while a customer with volatile demand a residential one.

[12] Cao and Wei (2000 a) used a model in which future temperature deviations, at time $t+1$, depends on temperature deviations over three previous dates, $t, t-1, t-2$. We have found that for Texas dependence of temperature on temperature deviations three days ago is statistically insignificant and have not included this term in the model.

[13] By substituting temperature forecasts rather than historical averages, the model can also incorporate information from short- and long-term meteorological forecasts.

To calibrate the model for ERCOT we use data available at the National Climatic Data Center website (see `www.ncdc.gov`). We use daily data on average temperatures in Central Texas, from January 1948 through December 1999. Figure 1 presents the average daily temperatures. The variables $\bar{T}_t$ in equation 1 are set to these averages.

After obtaining the values for the average temperatures, we calibrate the temperature model in two steps: first, we construct the variable $\Delta_t^T = T_t - \bar{T}_t$ for each day in the data set. Since the model is heteroskedastic, we use an iterative procedure, in which we start with a guess for $\sigma_{(0)}^T, \sigma_{(1)}^T, \phi$. Using this guess for the heteroskedastic errors, we regress Δ_{t+1}^T on Δ_t^T and Δ_{t-1}^T to estimate the autocorrelation coefficients ρ_1^T, ρ_2^T. We then construct the deviations between the expected Δ^T and the actual Δ^T for each day, and use them to compute the deviations σ_t^T, from which we fit, using nonlinear regression (see Ratkowsky (1983)), the parameters $\sigma_{(0)}^T, \sigma_{(1)}^T, \phi$. We repeat the procedure until the values of the parameters $\sigma_{(0)}, \sigma_{(1)}, \phi$, converge. The estimated parameter values, and the standard errors are reported in Table I.

A.2. Demand vs. Temperature

To estimate the relationship between demand for electricity and temperature, we use a data set of power load for the summer 1999 period for ERCOT available at the ERCOT website (`www.ercot.com`). The data provide the average daily on-peak and off-peak load by region within ERCOT. We use average on-peak load, which includes load between 6 a.m. and 10 p.m. Monday through Friday. The reason for this choice is that night and weekend load is low enough that interruptions are not necessary. From Figure 2 it is clear that, for the range of temperatures encountered during the summer months, there is a close to linear relationship between average on-peak load and average temperature.

Based on Figure 2, we model the relationship between average temperature and average load by a linear function with additional random fluctuations.[14]

$$L_t = \alpha_L + \beta_L T_t + \sigma_L \varepsilon_{L_t}, \quad \varepsilon_{L_t} \sim N(0,1) \tag{2}$$

where L_t is the load at time t, T_t the temperature, α_L the load intercept, β_L the expected marginal increase in load for a unit increase in temperature, and σ_L the magnitude of the random fluctuations around the linear relationship between load and temperature. Table II presents that OLS regression results for the values of the parameters.

B. The Supply Curve

Most of the supply available in ERCOT is generated within ERCOT, due to limited transmission between ERCOT and surrounding areas. The generators that service the base load are coal-based or nuclear facilities, while intermediate and peaking plants include plants based on natural gas, oil or hydroelectric power. Since we do not have access to the marginal costs of the available generators, we calibrate our model of the supply curve through the observed relationship between spot electricity price and electric load. To justify this approach, we note that since all ERCOT customers pay a fixed retail price, we assume their demand to be inelastic with respect to the wholesale spot price. In addition, in 1999, there were no reports of electric customer interruption in ERCOT.

From Figure 3 we notice that there appear to be two regimes for the supply curve: the low demand regime, where load and prices are relatively low and price fluctuations are minor; and the high demand regime where load and prices are high and price fluctuations are large. Based on these observations, we propose a two regime model for the price/load relationship. To account for the fluctuations in price, we allow fluctuations in the load, corresponding fluc-

[14]Most variability of demand in Texas during the summer is driven by air-conditioning load that is dependent on temperature. In a colder climate one may need to include additional terms in the load-temperature relationship.

tuations in supply due to, for example, generator outages, transmission outages, transmission congestion, and possibly strategic behavior by market participants. For simplicity we assume that the magnitude of the fluctuations of supply is the same for both regimes.[15]

The model of the relationship between load and price is given by:

$$P_t = \begin{cases} \beta_{S,l}(L_t + \sigma_S \varepsilon_{S_t}) + \alpha_{S,l}, & \text{if} \quad L_t + \sigma_S \varepsilon_{S_t} \leq S_\mathrm{b} \\ \beta_{S,h}(L_t + \sigma_S \varepsilon_{S_t}) + \alpha_{S,h}, & \text{if} \quad L_t + \sigma_S \varepsilon_{S_t} > S_\mathrm{b} \end{cases} \tag{3}$$

where P_t is the wholesale price at time t, L_t the demand at time t, ε_{S_t} is a standard, normally distributed random variable, and S_b the supply level that determines the break between the high and low regimes.

To calibrate the supply curve model, we use the days from the data in Figure 3 with prices above \$60/MWh, assuming that they correspond to the high regime. From these days we estimate the parameters for the high regime, as well as the magnitude of supply fluctuations σ_S. We estimate the parameters for the low regime using days in which ERCOT load was below 39 GW. The break point S_b is calculated by requiring the expected price to be a continuous function of load; i.e.,

$$\beta_{S,l} S_b + \alpha_{S,l} = \beta_{S,h} S_b + \alpha_{S,h}$$

The OLS estimates for the parameter values are presented in Table III.

[15]This assumption is not critical for the valuation of interruptible contracts, since small errors in the calibration of the model parameters for the low regime have only a minor impact on the value of interruptible contracts.

III. Valuation and Optimal Interruption Policy for Interruptible Contracts

A. Stochastic Optimal Control Problem

The problem of determining the optimal interruption policy, as well as the value of the interruptible contract can be formulated as a problem of optimal stochastic control, with the objective of maximizing the utility of the distributor. We assume that the distributor is risk-neutral with respect to gains and losses and has intertemporal preferences that can be quantified through a constant discount factor.[16]

As we already discussed, the distributor can aggregate the information from all interruptible contracts into the load available for interruption the following day and the total load available for interruption during the remaining period. Therefore, the distributor may think of all its customers in terms of three representative customers: the first customer has not signed an interruptible contract and pays p_{retail} on its load; the second customer has signed a pay-in-advance interruptible contract and pays a reduced price on its load, p_{reduced}; the third customer has signed a pay-as-you-go contract and receives compensation p_{fine} per unit of interruption, upon interruption.

The net profit during peak hours on a day with load of L prior to interruption, load interrupted from the pay-in-advance contracts of l_{advance}, load interrupted from the pay-as-you-go

[16]Other choices for the risk-aversion of the distributor are possible. However, choosing a risk-neutral distributor is sufficient to capture the factors that are important in determining the optimal interruption policy, as well as the value of an interruptible contract.

contracts of l_{pago}, spot price as a function of expected load after interruption of $L - l_{\text{advance}} - l_{\text{pago}}$ and price fluctuations ε_S, $p_{\text{spot}(L-l_{\text{advance}}-l_{\text{pago}}),\varepsilon_S}$, is given by

$$\begin{aligned}\frac{\Delta V(L, p_{\text{spot}}, l_{\text{advance}}, l_{\text{pago}})}{16} =& (L - L_{\text{under\_contract}} - l_{\text{pago}}) p_{\text{retail}} \\ &+ (L_{\text{under\_contract}} - l_{\text{advance}}) p_{\text{reduced}} \\ &- L_{\text{generation}} p_{\text{generation}} \\ &- l_{\text{pago}} p_{\text{fine}} \\ &- (L - l_{\text{advance}} - l_{\text{pago}} - L_{\text{generation}}) p_{\text{spot}}(L - l_{\text{advance}} - l_{\text{pago}}, \varepsilon_S),\end{aligned} \tag{4}$$

where $L_{\text{under\_contract}}$ is the load under a pay-in-advance interruptible contract, $L_{\text{generation}}$ is the load available to the distributor at a fixed unit price $p_{\text{generation}}$, p_{retail} is the retail unit price paid by customers that have not signed an interruptible contract, p_{reduced} is the unit price paid by customers that have signed an interruptible pay-in-advance contract, and p_{fine} is the price paid by the distributor to customers under a pay-as-you-go contract per unit of interrupted load. The term $(L - L_{\text{under\_contract}} - l_{\text{pago}}) p_{\text{retail}}$ corresponds to the revenue to the distributor from the customers that have not signed a pay-in-advance interruptible contract, the term $(L_{\text{under\_contract}} - l_{\text{advance}}) p_{\text{reduced}}$ corresponds to the revenue from customers that have signed a pay-in-advance interruptible contract, the term $L_{\text{generation}} p_{\text{generation}}$ corresponds to the cost of procuring the load available to the distributor at a fixed price, the term $l_{\text{pago}} p_{\text{fine}}$ corresponds to the cost to the distributor for interrupting customers under a pay-as-you-go contract, and the term $(L - l_{\text{advance}} - l_{\text{pago}} - L_{\text{generation}}) p_{\text{spot}}(L - l_{\text{advance}} - l_{\text{pago}}, \varepsilon_S)$ corresponds to the cost of procuring the excess load in the spot electricity market.

Given our formulation of a structural model for electricity prices in Section II, the load and the spot price of electricity at time t depend on the temperature deviations from the temperature historical averages, or forecasted values, at time t and $t-1$. Given the values of the state variables $\Delta_t^T, \Delta_{t-1}^T$ and the remaining interruptible loads, $L_{\text{advance, remaining}}, L_{\text{pago, remaining}}$, the value function for the distributor is given by

$$
\begin{aligned}
V_t(\Delta_t^T, \Delta_{t-1}^T, L_{\text{advance, remaining}}, L_{\text{pago, remaining}}) = \\
\max_{l_{\text{advance}}, l_{\text{pago}}} \beta \{ \mathbb{E} [\Delta V(L_{t+1}, p_{\text{spot},t+1}, l_{\text{advance}}, l_{\text{pago}}) \\
+ V_{t+1}(\Delta_{t+1}^T, \Delta_t^T, L_{\text{advance, remaining}} - l_{\text{advance}}, L_{\text{pago, remaining}} - l_{\text{pago}}) | \mathcal{F}_t] \}
\end{aligned}
\tag{5}
$$

where the maximization is over

$$
\begin{aligned}
0 \leq l_{\text{advance}} \leq \min \left(L_{\text{advance, daily}}, L_{\text{advance, remaining}} \right), \\
0 \leq l_{\text{pago}} \leq \min \left(L_{\text{pago, daily}}, L_{\text{pago, remaining}} \right)
\end{aligned}
\tag{6}
$$

In Equation 5, β is the discount factor, and $\mathcal{F}_t$ denotes the information available at time t. Note that the interruption amounts $l_{\text{advance}}, l_{\text{pago}}$, are chosen at time t, but interruption occurs over the next day, at time $t+1$. The expectation in Equation 5 is taken over the random variables $\varepsilon_{L_{t+1}}, \varepsilon_{S_{t+1}}, \varepsilon_{t+1}^T$.

Assuming a teminal date t_f for the interruptible contracts, we set

$$
V_{t_f} = 0
$$

The maximization problem 5 can be solved using dynamic programming with state variables $\Delta_t^T, \Delta_{t+1}^T, L_{\text{advance, remaining}}, L_{\text{pago, remaining}}$, and choice variables $l_{\text{advance}}, l_{\text{pago}}$. Given the difficulty in solving a problem with many state and choice variables, we consider only one contract type at a time; i.e., the situation where the distributor has entered either pay-in-advance cotnracts, or pay-as-you-go contracts, but not both.

B. Base Case Interruptible Contracts

To further study the optimal interruption policy and the value of interruptible contracts, we specify base case contracts for the different types of interruptible contracts. The parameter values for these base case contracts have been chosen with ERCOT in mind. For both types of contracts we consider the possibility of interruption during weekdays of the summer months only, which is the time when interruption is most likely in ERCOT.

Pay-in-advance contract

In the base case pay-in-advance contract the distributor provides a 15% reduction to the retail price of electricity, $p_{\text{reduced}} = 0.85 p_{\text{retail}}$, to the entire load under contract, $L_{\text{under\_contract}}$. In exchange, the distributor may interrupt up to 20% of the load under contract daily, $L_{\text{advance, daily}} = 0.2 \times L_{\text{under\_contract}}$, up to ten times per year, $L_{\text{advance, yearly}} = 2 \times L_{\text{under\_contract}}$. Under this type of contract there is no additional fine paid by the distributor upon interruption, $p_{\text{fine}} = 0$.

Pay-as-you-go contract

In the base case of the pay-as-you-go contract, the distributor does not provide any reduction in the retail price — $L_{\text{under\_contract}} = 0, p_{\text{reduced}} = p_{\text{retail}}$ — but, in exchange for the right to interrupt customer load, pays a fine of either \$150/MWh or \$600/MWh of interrupted electricity.[17] In addition, the customer may be interrupted up to ten times per year, $L_{\text{pago, yearly}} = 10 \times L_{\text{pago, daily}}$.

C. Optimal Interruption Policy

The optimal interruption policy is determined by the first order condition that at the optimal policy, marginal benefit from additional interruption equals marginal cost. To describe the optimal interruption policy we first consider the case where there is no limit in the total yearly

[17]The fines of \$150/MWh and \$600/MWh, have been used in interruptible contracts in California depending on whether the interruption notification was overnight or 30 minutes prior to interruption, respectively.

amount available for interruption. Then, the value function in Equation 5 is independent of $L_{\text{advance, remaining}}, L_{\text{pago, remaining}}$, and the maximization is myopic; i.e., on each day t, the optimal interruption policy maximizes expected net profit on day $t+1$ only. In this case we can easily calculate the marginal cost and marginal benefit of interruption. For the case of pay-in-advance contracts, the marginal cost is p_{reduced}, while for pay-as-you-go contracts the marginal cost is $p_{\text{retail}} + p_{\text{fine}}$. The marginal benefit is the same for both contract types, and is a function of the expected load. For the base case contracts, the marginal benefit for different parameter values for the interruptible contracts is shown in Figure 4. The optimal policy can be directly determined from the figure in the following way: if the expected load is such that the marginal benefit is greater than the marginal cost, the distributor interrupts an amount that is the lesser of the maximum daily interruptible limit, or the amount for which the expected load is reduced to the point where the marginal benefit equals the marginal cost.

The discussion above is the proof of the following proposition.

Proposition 1. *In the case with no yearly limit for pay-in-advance and pay-as-you-go interruptible contracts, if $p_{reduced} \leq p_{retail} + p_{fine}$, then the optimal interruption policy is to interrupt the pay-in-advance contracts up to the daily limit, prior to interrupting any of the pay-as-you-go contracts.*

From Figure 4 we notice that the optimal interruption policy for interruptible contracts without yearly limits, depends on several factors. In particular: The pay-in-advance contract gets interrupted at lower expected loads than the pay-as-you-go contract, since the bulk of the cost of the pay-in-advance contract is sunk; The expected load at which interruption begins increases with the generation available to the distributor at a fixed price, since the marginal benefit of interruption at the same expected load is smaller because a reduction in the expected spot price only affects the excess demand; As the retail price of electricity paid to the distributor by its customers increases, the distributor interrupts at higher loads since the cost of interruption increases with the retail price; Finally, without yearly limits, the distributor interrupts aggressively, i.e., at expected loads significantly below the transition points between the two regimes in the supply curve. This aggressive behavior can be attributed to the large cost to the distributor of ending up in the high load regime, since the spot electricity price applied to the entire load procurred from the spot market.

In the case with no daily interruption limit, but with a yearly interruption limit, we can prove a proposition similar to Proposition 1.

Proposition 2. *In the case with no daily interruption limits for pay-in-advance and pay-as-you-go interruptible contracts, if $p_{reduced} \leq p_{retail} + p_{fine}$, then the optimal interruption policy is to interrupt the pay-in-advance contracts until the yearly limit for the pay-in-advance contracts is exhausted, prior to interrupting any of the pay-as-you-go contracts.*

Proof. By contradiction, assume that it is optimal to interrupt some amount from the pay-as-you-go contracts prior to exhausting the pay-in-advance contracts. Then, it is easy to see that the value function can be improved by following the strategy in which the interruption amounts from the pay-as-you-go contracts is transferred to the pay-in-advance contracts instead, if possible. If, later in the path, the pay-in-advance contract is exhausted, the pay-as-you-go contract is interrupted instead. For all scenarios the value function can not decrease, since the cost of interruption is at most equal to the original cost.

□

When there are both daily and yearly limits on both types of contracts, there is no generalization of Propositions 1 and 2, since one may want to avoid exhausting the pay-in-advance contracts in order to be able to interrupt larger amounts in a daily basis. In this case, we consider the situation with only one contract type present. In Figure 5, we provide the optimal interruption strategy for pay-in-advance contracts with yearly limits. The figure provides the optimal interruptible strategy at the beginning of July (60 days before summer end) and the beginning of August (30 days before summer end), and different amounts of interruption available, for a distributor with 35 GWs of generation available at a fixed price, who charges a retail price of $60/MWh to its customers. From the figure we notice that the most significant difference between the contract with yearly limits and the contract without yearly limits is that with yearly limits interruption occurs at higher expected loads. This behavior is more pronounced at low levels of remaining interruptible load, and at times that are far from the end of the summer. In addition, at low levels of remaining interruptible load and at times that are far from the end of the summer, the interruption policy is "fuzzier", since it depends on two state variables, rather than just the expected load (these state variables are the deviation from historical temperatures at times t and $t-1$). Moreover, the slope of the interruption policy, with respect to the expected load, decreases for lower remaining interruption loads, and far from expiration, in line with the intuition that in those circumstances waiting before interrupting is valuable.

In calculations we do not report, we verified that the differences between the different contract types are relatively minor, and that the interruption policy for reasonable parameter ranges for the pay-as-you-go contract is very similar to the interruption policy for the pay-in-advance contract.

D. Value of Interruptible Contracts

The value of an interruptible contract is defined as the difference in the value function of the distributor between having the interruptible contract and not having the interruptible contract. The value function for a distributor with no interruptible contracts can be easily calculated using Monte Carlo simulation, since no choice variables are involved in that case.

In Figure 6, we provide contour plots for different interruptible contracts as the load under contract and the retail price change. From the contour plots we notice that, for the pay-in-advance contract, the value of the interruption may be positive or negative, since a discount is provided in advance to the customer. The figure demonstrates the importance of having generation available at a fixed price, since if the amount of generation available to the distributor at a fixed price is low, the distributor interrupts often in order to avoid paying the spot electricity price on a large portion of the customer load. On the other hand, when the amount of generation available increases, the value of the interruption drops and may become negative for large amounts under contract, or for high retail prices. Under these situations the reduction in the income of the distributor, due to the interruptible contract, is higher than the value added by the reduction of the spot price due to interruption, when the load is high.

For the pay-as-you-go contract, the value of the interruption is always positive, since payment is made only after it is optimal to interrupt. The interruption value increases with the load available for interruption, and decreases with the retail price. In results not reported here, we verified that, similar to the case of the pay-in-advance interruptible contract, the interruption value is higher for low generation load available to the distributor at fixed prices. We note that for both contract types the marginal value of the interruptible contract decreases with increasing amount of load under contract.

IV. Conclusions

We have presented a model that quantifies the value of interruptible contracts. The model builds on a structural model of electricity prices, where price is determined from stochastic models for supply and demand for electricity. The model accounts for fluctuations to demand due to temperature changes, and fluctuations to supply due to outages and transmission congestion. The interruptible contracts are priced from the point of view of a distributor of electricity that has the obligation to provide electricity to all its customers. The distributor needs to rely on the spot market to satisfy at least part of the electricity demand, and we have shown that the interruptible value and the optimal interruption policy depend critically on the amount of generation available to the distributor at a fixed price. In the absence of forward contracts or ownership of generating assets, the interruptible contracts are the most valuable, and the distributor interrupts aggressively. As more generation is available for a fixed price, the value of interruptible contracts diminishes, and interruption occurs at higher expected loads.

We have studied two types of contracts: the pay-in-advance contract, in which the distributor agrees to a discount for the entire load of a customer, in exchange for the right to interrupt part of the load a certain number of times; and the pay-as-you-go contract where the distributor compensates the customer for the interrupted load upon interruption. Given a choice between different types of interruptible contracts, pay-as-you-go contracts are preferable to the distributor, since, due to the advance payment of the pay-in-advance contracts, it is possible, in cases where the distributor signs up too large a load, that the value of the interruptible pay-in-advance contract is negative, while, on the other hand, the value is always positive for the pay-as-you-go contracts. Alternatively, our methodology can be combined with information on customer preferences regarding types of interruptible contracts, to decide the optimal design and mix of different contract types.

To price interruptible contracts, we have presented and calibrated a structural model for electricity prices based on publicly available information. Other than pricing interruptible contracts, an accurate structural model can be useful in the optimal asset allocation problem for a generator that can choose among generation plants, forward contracts, options, and interruptible contracts, as well as the optimal design of new types of contracts.

A. Notation

Temperature model

T_t: actual temperature on day t

$\bar{T}_t$: average temperature for the t-th day in the year

Δ_t^T: difference between actual and average temperatures on day t, $\Delta_t^T = T_t - \bar{T}_t$

ρ_1^T: first order autocorrelation for temperature differences from the average temperature

ρ_2^T: second order autocorrelation for temperature differences from the average temperature

σ_t^T: magnitude of temperature fluctuations on day t

$\sigma_{(0)}^T$: fixed term of temperature fluctuations

$\sigma_{(1)}^T$: magnitude of seasonal term of temperature fluctuations

ϕ: day during the year on which temperature fluctuations are greatest

Load vs. Temperature Model

L_t: load at time t

α_L: load intercept

β_L: marginal expected increase of load per one degree Fahrenheit increase in temperature

σ_L: magnitude of fluctuations in the load-temperature model

Load vs. Price Model

$\beta_{S,l}$: Marginal increase in expected spot price per unit increase in load in the low demand regime

$\alpha_{S,l}$: intercept for the load-price relationship in the low demand regime

$\beta_{S,h}$: Marginal increase in expected spot price per unit increase in load in the high demand regime

$\alpha_{S,h}$: intercept for the load-price relationship in the high demand regime

σ_S: magnitude of fluctuations in the load-price relationship

S_b: supply level that marks the boundary between the low demand and the high demand regimes.

Prices

P_t: spot electricity price at time t

p_{retail}: fixed retail price, charged by the distributor to its retail customers.

p_{reduced}: fixed retail price paid by customers that have signed a pay-in-advance interruptible contract

p_{fine}: fine per unit of interrupted load paid to the customers that have signed a pay-as-you-go interruptible contract

$p_{\text{generation}}$: unit cost per unit load available to the distributor at a fixed price

Interruptible contracts

$L_{\text{advance, daily}}$: maximum amount available for interruption under a pay-in-advance contract, for one day

$L_{\text{advance, yearly}}$: total amount available for interruption under a pay-in-advance contract for one year

$L_{\text{advance, remaining}}$: total amount available for interruption, under a pay-in-advance contract, for the remaining period, $L_{\text{advance, remaining}} \leq L_{\text{advance, yearly}}$

l_{advance}: interrupted load under a pay-in-advance contract, in a particular day, $l_{\text{advance}} \leq L_{\text{advance, daily}}$

$L_{\text{under\_contract}}$: Total daily load of customers under a pay-in-advance interruptible contract

$L_{\text{pago, daily}}$: maximum amount available for interruption under a pay-as-you-go contract for one day

$L_{\text{pago, yearly}}$: total amount available for interruption under a pay-as-you-go contract for one year

$L_{\text{pago, remaining}}$: total amount available for interruption, under a pay-as-you-go contract, for the remaining period, $L_{\text{pago, remaining}} \leq L_{\text{pago, yearly}}$

l_{pago}: interrupted load under a pay-as-you-go contract, in a particular day, $l_{\text{pago}} \leq L_{\text{pago, daily}}$

$L_{\text{generation}}$: power available to the distributor at a fixed price.

References

California, Public Utilities Commision, 2001, Energy Division's report on interruptible programs and rotating outages, available at `http://www.cpuc.ca.gov/PUBLISHED/REPORT/5119.htm`.

Campbell, Sean D., and Francis X. Diebold, 2002, Weather forecasting for weather derivatives, working paper, available at `www.ssrn.com`.

Cao, Melanie, and Jason Wei, 2000 a, Pricing the weather: An intuitive and practical approach, *Risk* pp. 67–70.

Cao, Melanie, and Jason Wei, 2000 b, Equilibrium valuation of weather derivatives, Working paper, University of Toronto.

Deng, S., 1999, Stochastic models of energy commodity prices: Mean-reversion with jumps and spikes, Working paper, Georgia Institute of Technology.

Deng, S., 2000, Pricing Electricity Derivatives Under Alternative Spot Price Models, in *Proceedings of the* 33$^{\text{rd}}$ *Hawaii International Conference on System Sciences*.

Kamat, Rajnish, and Shmuel S. Oren, 2002, Exotic options for interruptible electricity supply contracts, *Operations Research* 50, 835–850.

Kholodnyi, Valery A., 2000, The stochastic process for power prices with spikes and valuation of European contingent claims on power, working paper.

Pilipovic, D., 1997, *Energy Risk: Valuing and Managing Energy Derivatives*. (McGraw-Hill).

Ratkowsky, David A., 1983, *Nonlinear regression modeling*. (Marcel Dekker, Inc. New York and Basel).

Skantze, Petter, Andrej Gubina, and Marija Ilic, 2000, Bid-based stochastic model for electricity prices: The impact of fundamental drivers on market dynamics, MIT Energy Laboratory Report EL 00-004, MIT.

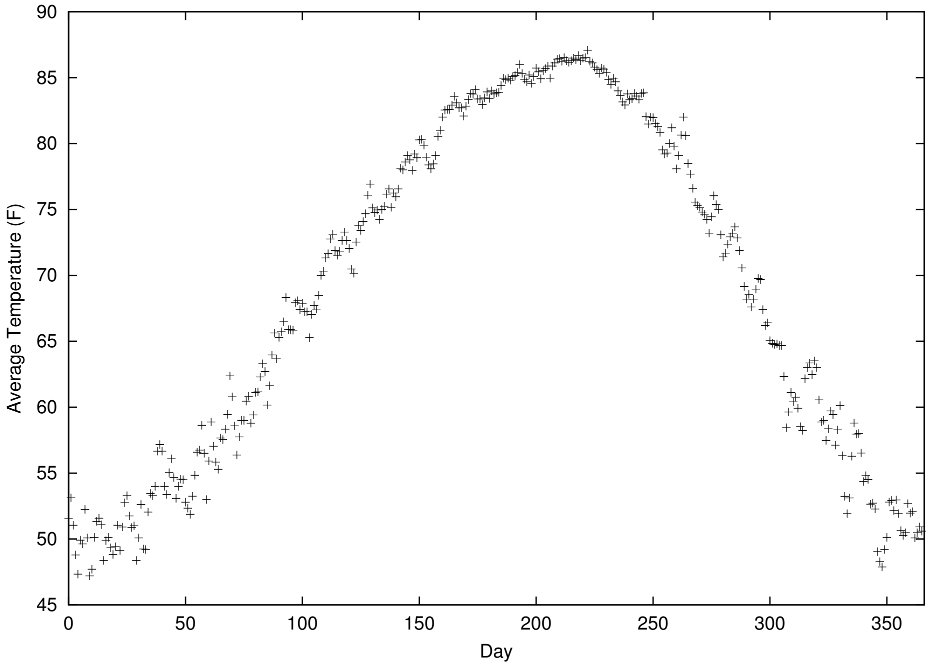

Figure 1. Average daily temperatures for central Texas, averaged over 1948-1999.

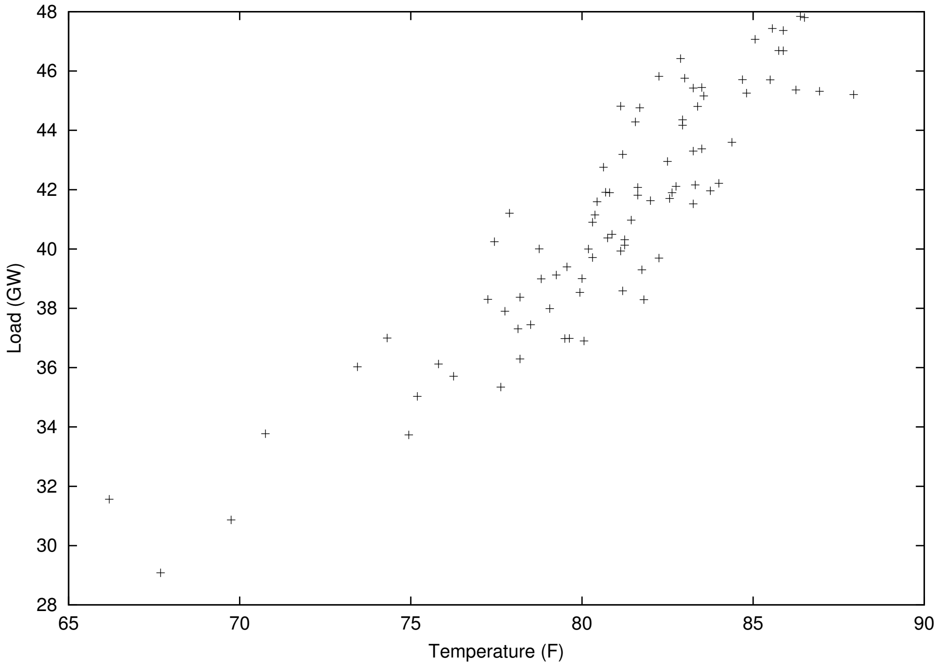

Figure 2. Average on-peak Load vs. Average Temperature during weekdays for the period June 1st to August 31st, 1999 in ERCOT.

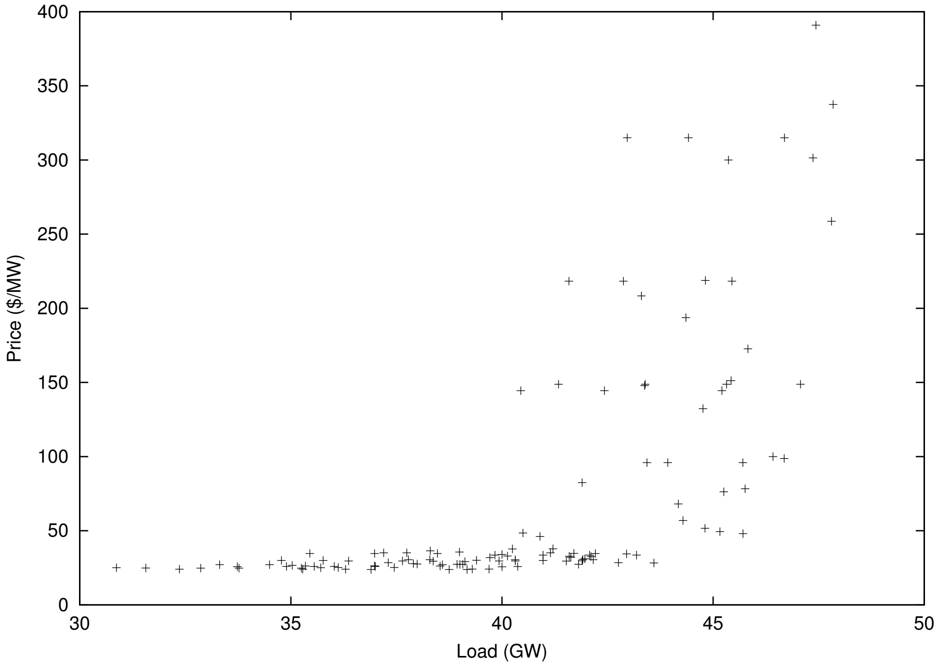

Figure 3. On-peak price per MW of electricity vs. average daily load during weekdays for the period June 1st to August 31st, 1999 in ERCOT.

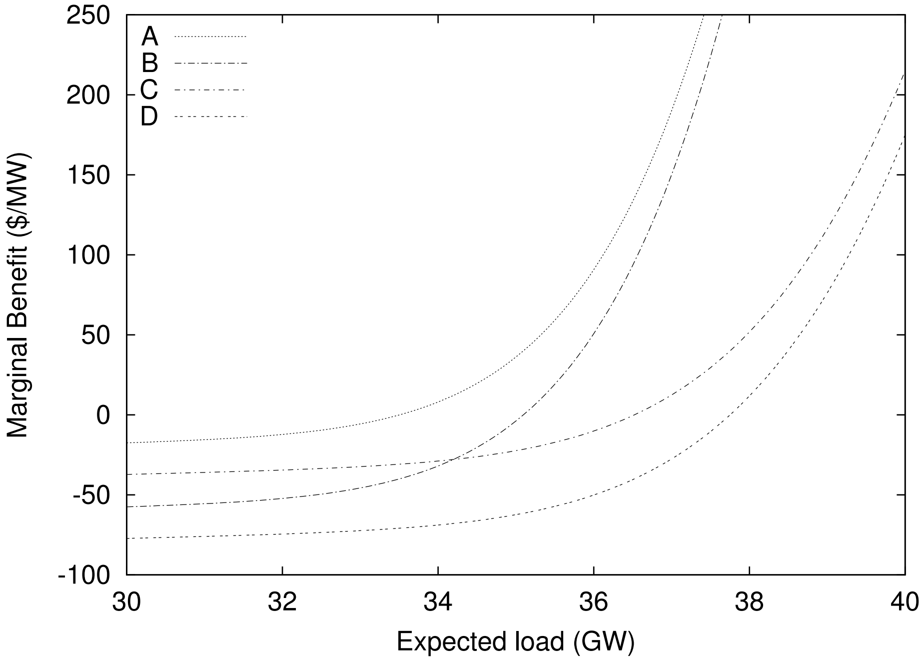

Figure 4. Marginal benefit from interrupting a MW of electricity on August 31st in ERCOT, when there is an unlimited amount of interruption available. Curve A corresponds to a distributor that provides electricity at a retail price of $60/MW and has zero generation available at a fixed cost; curve B corresponds to a distributor that provides electricity at a retail price of $100/MW and has zero generation available at a fixed cost; curve C corresponds to a distributor that provides electricity at a retail price of $60/MW and has 35 GWs of generation available at a fixed price; curve D corresponds to a distributor that provides electricity at a retail price of $100/MW and has 35 GWs of generation available at a fixed price.

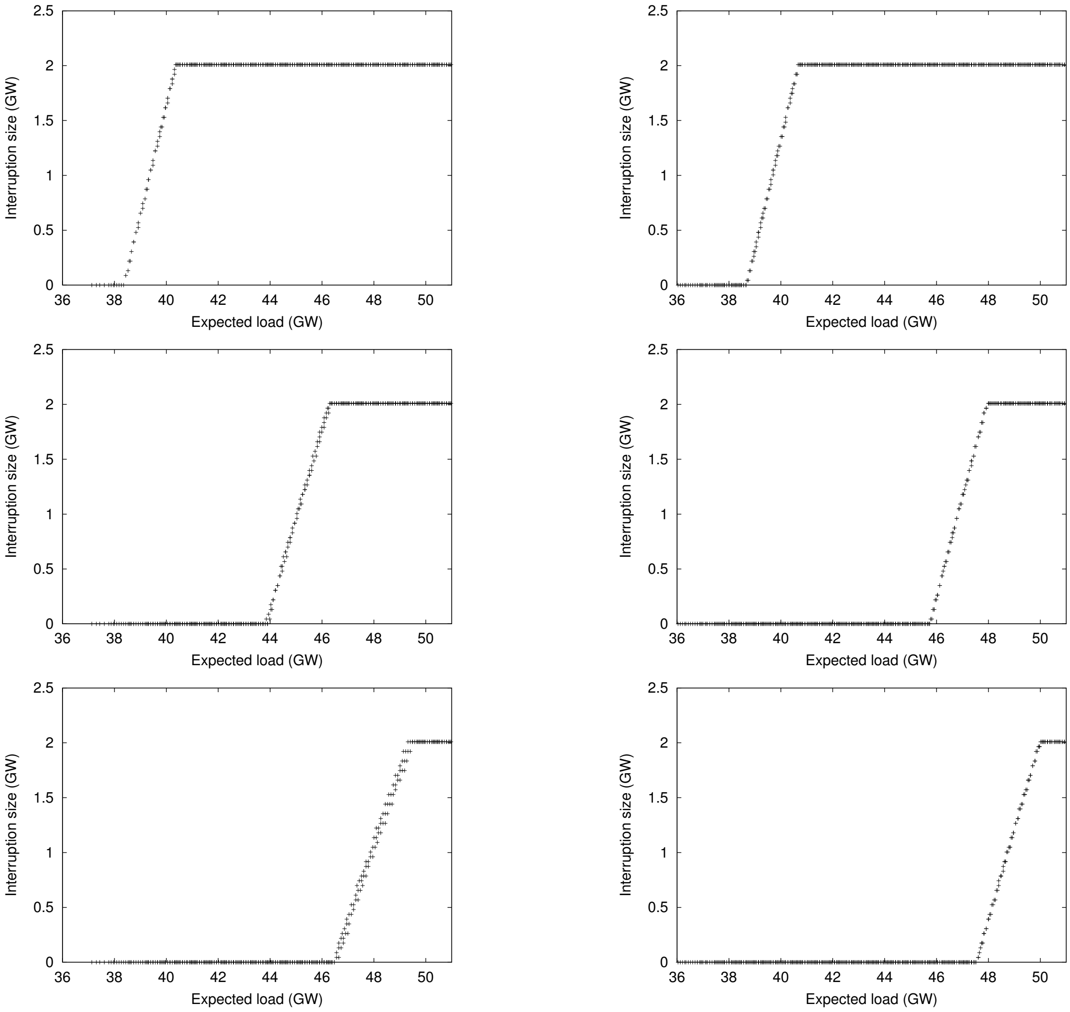

Figure 5. Plots describing the interruption strategy as a function of the expected load for the following day. The figure on the left correspond to a pay-in-advance contract 60 days prior to the end of the contract, while the figures on the right correspond to the same contract 30 days prior to the end of the contract. The figures on the top row correspond to an unlimited amount of interruption remaining, those in the middle row to 20 GW of interruption remaining, and those on the bottow row to 5 GW of interruption remaining. For all contracts, the distributor has 35 GWs of generation available, and the retail price is $60/MWh. The daily amount that can be interrupted is 2 GWs. The pay-in-advance contract provides a discount of 15% to the entire load under contract. The fuzziness in the figures in the middle and bottow rows is due to the fact that the optimal interruption strategy depends on two, rather than one, state variables.

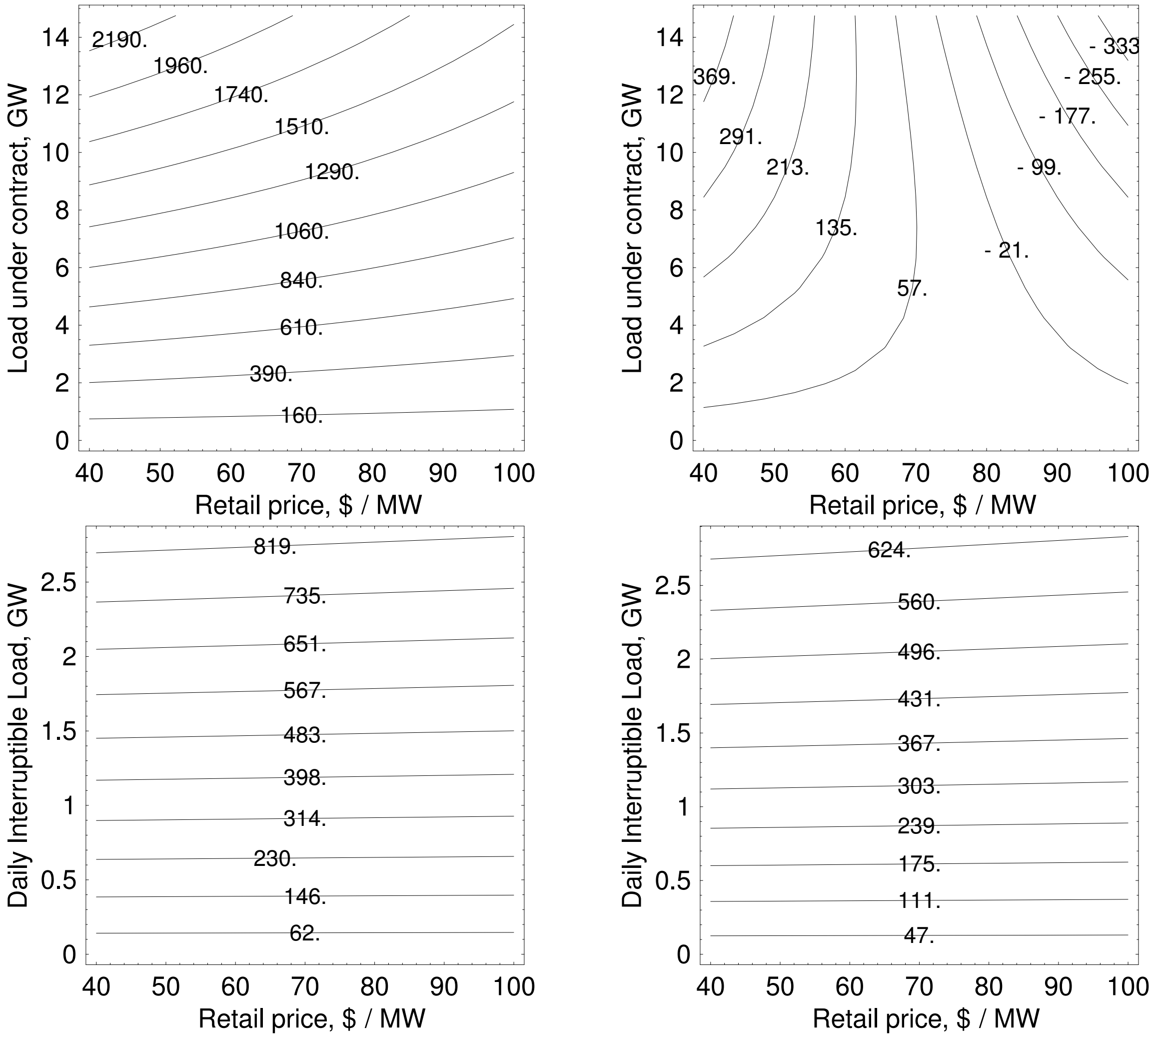

Figure 6. Contour plots of the value of interruptible contracts. The figures on the top row correspond to pay-in-advance interruptible contracts as the retail price that a distributor charges and the total load under contract change. The discount provided to the entire load under contract is 15% from the retail price. The load available for interruption is equal to 20% of the load under contract, and interruption can occur up to ten times. The figure on the top left corresponds to a distributor that has no generation available at a fixed price and is forced to serve all the load from the spot market. The figure on the top right corresponds to a distributor that has 35 GWs of generation available. The bottom row figures correspond to pay-as-you-go interruptible contracts. Interruption can occur up to ten times. The distributor has 35 GWs of generation available. The bottom left figure corresponds to a contract with a penalty of \$150/MW of interrupted load, and the bottom right figure on to a contract with a penalty of \$600/MW. The value of the interruptible contracts is in millions of dollars.

Table I
Temperature Model

This table presents the estimated values and the standard errors of the parameters for the temperature model. The parameter ρ_1^T is the one-day autocorrelation, ρ_2^T the two-day autocorrelation, $\sigma_{(0)}^T, \sigma_{(1)}^T$ parameters that determine the magnitude of the fluctuactions, and ϕ corresponds to the date during the year when the fluctuations are the largest.

	Estimate	Standard Error
Intercept	-0.0002	0.010
ρ_1^T	0.837	0.010
ρ_2^T	-0.188	0.010
$\sigma_{(0)}^T$ (Fahrenheit)	8.316	0.131
$\sigma_{(1)}^T$ (Fahrenheit)	5.747	0.185
ϕ (days)	-14.5	1.6

Table II
Load vs. Temperature Model

This table presents the estimated values and the standard errors of the parameters for the model of the dependence of load on temperature. The model is given by

$$L = \alpha_L + \beta_L T + \sigma_L \varepsilon_L$$

where L is the load, T is the temperature, and ε_L is a standard, normally distributed random variable.

	Estimate	Standard Error
Intercept α_L (GW)	-29.5	3.5
Slope β_L (GW/Fahrenheit)	0.874	0.044
σ_L (GW)	1.80	0.16

Table III
Supply Curve Model

This table presents the estimated values and the standard errors of the parameters for the model of the supply curve. There are two regimes for the supply curve, high and low. The model for prices is described by

$$P = \begin{cases} \beta_{S,l}(L+\sigma_S\varepsilon_S)+\alpha_{S,l}, & \text{if } \; L+\sigma_S\varepsilon_S \leq S_b \\ \beta_{S,h}(L+\sigma_S\varepsilon_S)+\alpha_{S,h}, & \text{if } \; L+\sigma_S\varepsilon_S > S_b \end{cases}$$

where P is the price, L the load, ε_S is a standard, normally distributed random variable, and S_b the supply level that determines the break between the high and low regimes.

	Estimate	Standard Error
$\beta_{S,l}$ (\$/GW)	0.554	0.281
$\beta_{S,h}$ (\$/GW)	146.0	78.6
$\alpha_{S,l}$ (\$)	8.86	10.41
$\alpha_{S,h}$ (\$)	-6344.5	3418.9
σ_S (GW)	1.863	0.16
S_b (GW)	43.68	