Impacts of Responsive Load in PJM
Load Shifting and Real Time Pricing

Kathleen Spees and Lester Lave

Carnegie Mellon University
Carnegie Mellon Electricity Industry Center
January 31, 2007

kspees@cmu.edu
412.445.2694

lave@cmu.edu
412.268.8837

Abstract

Since bulk electricity cannot be stored economically, inefficient, low-capital cost generators satisfy peak demand. A fixed retail price encourages demand spikes whether the hourly marginal generation cost is $1 or $0 per kWh. For PJM, 15% of the capacity was used only 1.1% of the time in 2006. If retail price reflected hourly marginal cost, installed generation and transmission capacity could be reduced with a greater proportion of capacity in base-load plants that have higher capital costs, but lower average and marginal costs. We focus on current data from PJM to estimate the savings to both consumers and producers from two changes in pricing. The first is "real time pricing" (RTP) where retail prices reflect the hourly marginal cost of providing power and "time of use" (TOU) pricing that varies price by daily on-peak and off-peak periods to reflect the average annual price for each period. The surprising result is that a shift to RTP has a relatively small effect on price; consumer plus producer surplus rises 2.9%-3.3% with RTP and 0.6%-0.7% with TOU pricing. Peak capacity savings are seven times larger with RTP than with TOU. In the short-term, plausible consumer responses would drop peak load by 10.4-17.7% for RTP and 1.1-2.4% for TOU pricing. Half of all possible customer savings from load shifting can be obtained by shifting only 1.7% of all MWh to another time of day, indicating that small demand-side changes can make a large difference; only the largest customers need smart meters in order to get the majority of savings from RTP or TOU pricing.

1 Introduction

The electricity industry uses much of its generation and transmission capacity only a small fraction of the time. Over the calendar year 2006, 15% of the generation capacity in the Pennsylvania-New Jersey-Maryland (PJM) territory ran only 1.1% or fewer hours, and 20% of capacity ran only 2.3% or fewer hours [1] [1]. The result is tens of billions of dollars[2] invested in generation that has low capital cost, but high generation cost and life cycle social cost.

The expensive excess capacity has two causes. The first is technical: bulk electricity cannot be stored economically, so there must be enough generation and transmission capacity to satisfy demand in real time, or there will be a blackout. The second is regulatory: most customers pay a constant flat price that is averaged over a year rather than the changing hourly price of generation and transmission. For example, a customer whose retail rate is $0.10/kWh will pay that price no matter whether the wholesale price of electricity reaches its limit of $1/kWh or drops to $0/kWh. If customers faced the wholesale price volatility of electricity, as they do with gasoline, natural gas, fruits and vegetables, and many other products, they would buy less power at $1/kWh and more at

[1] This is based on the entire PJM hourly load profile in 2006 [1]. Even at peak load, the system had 17.5% excess available generation capacity. We do not include generation excess at coincident peak load in this calculation because some generation excess is necessary for reliability purposes.
[2] At $600/kW, a reasonable natural gas generator cost, this 15% of PJM's generation capacity is worth $13 billion. At $1800/kW, a reasonable price for a coal generator, 15% of PJM's capacity is worth.$39 billion.

$0/kWh. In doing so they would level demand over time, enabling society to diminish investments in peaking generators, instead focusing on base-load units that have higher capital costs, but low generation cost.

Some electricity customers face "time of use" (TOU) pricing that charges them a higher price during on-peak hours, with the fixed on-peak and off-peak rates calculated as the delivered cost averaged over months. Other customers face "real time pricing" (RTP) where the hourly wholesale generation price determines the retail price. The TOU price gives better information and incentives than a single fixed tariff, but does not account for the times when wholesale prices spike because of high demand or equipment problems. Some view a TOU rate as a good compromise that frees customers from having to be informed about constantly changing prices and adjusting their consumption accordingly.

Few end users have any ability to react to real-time market conditions or to the location-specific costs of generation and transmission. A PJM survey of load-serving entities (LSE) reported that only 5% of end users are on rates directly or indirectly related to the real-time or day-ahead locational marginal price (LMP) [2, 3][3]. Companies currently offering RTP rates usually have a variety of partial-hedging options as well [4]. Some additional customers are enrolled in direct load control, interruptible contracts, or other subsidy programs that offer curtailment incentives during the top few load hours per year. A Federal Energy Regulatory Committee (FERC) report estimates that 4% of peak load in Reliability*First* Corporation (RFC) territory[4] could potentially be curtailed via these programs, but the maximum response in 2005 was only 0.7% [5]. Actual reductions are usually much smaller than program enrollments, partly because reduction is often voluntary [6].

We view the current flat tariff as both inefficient and inequitable. It is inefficient in raising the cost of electricity and using much more capital equipment to deliver the same quantity of power to end users. It is inequitable since customers whose consumption pattern is flat over the day or greatest at off-peak hours subsidize customers whose consumption patterns raise coincident peak demand.

We present a short-run analysis of a change to a more responsive demand-side market. In Section 4.3, we use one year of PJM data to build a supply model that implicitly accounts for dispatch constraints and load variation conditions observed in a year. We use this model in three different simulations to determine the impacts of responsive load. The first in Section 5 is an assumed load-shifting scenario that finds the effects of small changes in load profile on overall price. The load-shifting simulation does not consider customer time preference, but does show how quickly savings could be achieved. The final two simulations in Section 6 are more realistic; they use hourly demand curves to predict short-run impacts from change toward TOU or RTP from flat-rate pricing

[3] Estimate is from 3653 MW on locational marginal price (LMP) based rates and 69,063 MW represented in survey responses. We do not include load listed as switched to third-party suppliers in the calculation.
[4] The RFC territory does not match up exactly with PJM territory.

2 Literature Review

Borenstein's long-run RTP analysis predicts more than double the peak load savings we predict, see Section 6.2 [7]. His conclusion results from using a long-term supply curve with a short-term demand curve to predict equilibrium conditions. Because Borenstein includes capital costs in his supply curves, he predicts hourly prices up to $90,772/MWh; this implies that customers would spend 22% of the yearly bill in one peak demand hour. Those high prices are impossible since prices are hard-capped at $1000/MWh[5] in all but one United States market [9].

Holland and Mansur predict less than half the short-term peak load savings that we predict from RTP, see Section 6.2 [10, 11]. The modest impact is due to their method of using one constant stacked marginal cost curve to represent supply[6]. The approach here is similar, although we use the actual observed prices to account for transmission and other constraints while they assume constraint-free economic dispatch of system generators at marginal cost. The system stacked marginal cost curve underestimates price in most hours and, more importantly, it also underestimates the *slope* of the real supply curve. The supply curve slope determines the impact that a small change in load has on price, see Appendix B.

Power engineers account for real-time transmission constraints by solving the security-constrained direct-current optimal power flow (DCOPF) problem in example cases. This approach is similar to how PJM sets market prices. Wang, Redondo, and Galiana used a DCOPF-based model to examine demand-side participation in wholesale energy and ancillary services markets [12]. Their results indicate that demand participation erodes generator market power. However, results from test systems with a few buses do not translate directly into implications for the PJM system with roughly 7800 pricing points. Fitting supply curves to daily market data incorporates these constraints.

Consumer responsiveness to price is not known with confidence. Most studies examine decreasing load after an increase in fixed price. With TOU or RTP, customers can also shift their use to hours with lower price. After 5 years of experience with default RTP for customers larger than 2 MW, Niagara-Mohawk Power Corporation has observed an average demand elasticity of substitution of -0.11 [13, 14]. Based on their experience and select other studies, the plausible reaction to TOU or RTP pricing is a short-term demand elasticity between 0 and -0.4, see Section 4.2 [15, 16]. As customers get more time to adjust and understand that the new tariffs will be permanent, the elasticity will rise, perhaps to -0.9.

---

[5] California ISO is the exception with a $400/MWh soft cap on energy and ancillary service bids [8]. Generators may bid above a soft price cap and will be paid as bid; other generators will receive payment only as high as the cap. The neighboring Western Electricity Coordinating Council (WECC) has the same price caps although WECC is not a market operator.

[6] Their stacked marginal cost curve is based on generator heat rates, fuel prices, emissions prices, and other publicly available data for the time frame in question.

3 Data

Our data are system-wide average PJM market clearing prices and loads in the day-ahead and real-time markets over 2006 [1]. Day-ahead demand bids L_{DA} from LSEs are charged at the day-ahead price P_{DA}, the real-time increment or decrement L_{RT}-L_{DA} is charged or credited at the real-time price P_{RT}. Overall revenue and price are calculated in Equations (1) and (2).

$$R = D_{DA} \cdot P_{DA} + (D_{RT} - D_{DA}) \cdot P_{RT} \qquad (1)$$

$$P_0 = \frac{\sum_{hours} R}{\sum_{hours} L_{RT}} \qquad (2)$$

Overall realized price and the real-time demand for each hour are the most accurate data for evaluating demand response. In the implementation of RTP rates, customers should have access to both day-ahead and real-time market prices. We assume that nearly all power continues to be purchased in the day-ahead market; both markets are counted as RTP.

4 Market Model

We construct a short-term equilibrium model accounting for producer, consumer, and local utility participation. Results from the full model for RTP and TOU pricing are in Section 6. The load-shifting scenario in Section 5 uses only the supply side developed in Section 4.3.

4.1 Short-Term Equilibrium Model

Almost all consumers currently pay a flat rate P_0 for all their power. The market operator assumes that demand is completely unresponsive to price, L_0. While each hour has wholesale price above or below retail price, the profits and losses sum to zero over the year.

This base case model represents the disconnect between wholesale and retail; the result is the same as if there were RTP but end users had an elasticity of zero. This is a good characterization of current behavior since few customers face RTP [2, 3, 5, 15, 16].

Under TOU the retail price takes on a value of p_{on} during on-peak hours and p_{off} during off-peak hours. In PJM off-peak hours are weeknights 11 PM to 7 AM and all day on weekends and the six NERC holidays [17]. On and off-peak prices are set so that local utility profit sums to zero over on-peak hours and off-peak hours separately.

When all customers are exposed to RTP, the retail price differs from the wholesale price only by the local delivery charge. If the generators bid their marginal cost, the RTP will produce the efficient, socially optimal price and consumption.

4.2 Demand Side

We assume that each hour has a unique demand curve with constant elasticity as shown in Equation (3) where the hourly offset parameters β are determined by base case price and hourly load [7, 10].

$$P_D(L) = \beta \cdot L^{1/E}$$
$$\beta = \frac{P_0}{L_0^{1/E}} \quad (3)$$

The left side of (3) is replaced with the retail price $P_D(L)$ that applies in the flat (4), TOU (5), or RTP (6) cases.

$$P_D(L) = p_0 \quad (4)$$

$$P_D(L) = P_{TOU} = \begin{cases} p_{on} \\ p_{off} \end{cases} \quad (5)$$

$$P_D(L) = P_S(L) \quad (6)$$

A high price will get a consumer to lower demand or to shift demand to a cheaper hour. Given more time to respond, the consumer can buy more energy efficient appliances or equipment that enables shifting load to a later hour. The first type of response is a price elasticity and the second an elasticity of substitution, either short- or long-run.

A 1984[7] review of 34 studies found short run and long run price elasticities to be approximately -0.20 and -0.90 respectively, implying that a 10% price increase would reduce consumption by 2% in the short-run and 9% in the long-run [15]. A Department of Energy study reviewed published price elasticities of substitution under TOU, critical peak pricing (CPP), and day-ahead RTP situations [16]. The range of elasticities of substitution was 0.02 to 0.27.

Even though responsiveness is uncertain, we judge the short-run response to be between 0 and -0.4 under RTP conditions. We examine the full range of elasticities to examine the possible effects of short-and long-term responses.

4.3 Wholesale Supply Side

At one extreme, we might hypothesize that the supply-side relationship between price and load is the same over an entire year. At the other extreme, we might hypothesize that the relationship is unique to each day. The bids for a specified load may differ from one day to another because some generating units or transmission lines are not available, fuel prices have changed, or weather is impeding supply. Fitting unique parameters for each day would give a better fit than insisting that one set of parameters must fit the entire

[7] The short run numbers were recently updated in another review of 36 estimates with a median of -0.28.

year. However, the former is not a parsimonious model and says nothing about what parameter values should be used in future days.

The delivered price of electricity for each hour in a day follows a predictable pattern of being low in the early morning and at night with one or two peaks during the day. We fit the price and load data for each day with a third-degree polynomial. To investigate the similarity of the polynomial parameters across days, we employ dummy variables, which function as on-off switches, taking on values of 0 or 1.

Equation (7) models price as a function of load represented by an intercept, load, load squared, and load cubed. The equation uses dummy variables δ_1 and δ_0 to allow for the possibility that the coefficient of load and the intercept might vary each day. We also examined the possibility that the coefficients of the squared and cubed terms take on unique values each day but determined that the additional dummy variables improved explanatory power very little. Appendix A presents the results from trying a range of models. We selected (7) as a model with good explanatory power, only half the number of parameters as employing the additional two dummy variables, and as a good fit to the plotted data.

$$P_S(L) = a \cdot L^3 + b \cdot L^2 + \sum_{t=1}^{n} \{\delta_1 \cdot c_t \cdot L + \delta_0 \cdot d_t\} \tag{7}$$

The adjusted R^2 is 0.949, the F-statistic of 223 is highly significant[8], and the estimated parameters *a* and *b* are highly significant[9] all with p-values $\ll 0.001$.

[8] Model significance test has $F(731,8028) = 223$ with p-value $\ll 0.001$.
[9] Studentized t-test have $t^a(8028) = 10.9$ and $t^b(8028) = 33.0$ with p-values $\ll 0.001$ in each case.

4.4 Economic Result Definitions

Changes in consumer surplus ΔCS and producer surplus ΔPS between flat rate and RTP conditions are calculated in Equations (8) and (9) and shown graphically in Figure 1. Producer surplus is easier to calculate by integrating over load than over price. Change in consumer surplus in Equation (8) can be calculated in the TOU case by replacing P^* with the retail TOU price p_{off} or p_{on}. Change in producer surplus calculated in Equation (9) is the same formula under a change toward TOU or RTP because the wholesale electric price determines the producer surplus.

$$\text{(8)} \qquad \Delta CS = \sum_{hours} \int_{P^*}^{P_0} L(P_D)\partial P = \sum_{hours} \int_{P^*}^{P_0} \left(\frac{P_D}{\beta}\right)^E \partial P = \sum_{hours} \left(\frac{1}{E+1}\right)\left(\frac{P_D}{\beta}\right)^{E+1} \Bigg|_{P^*}^{P_0}$$

$$\Delta PS = \sum_{hours} \int_{P_S(L_0)}^{P^*} L(P_S)\partial P = \sum_{hours} \left\{ P^* L^* - P_0 L_0 - \int_{L_0}^{L^*} P_S(L)\partial L \right\}$$

$$\text{(9)} \qquad \Delta PS = \sum_{hours} \left\{ P^* L^* - P_0 L_0 - \int_{L_0}^{L^*} \left(aL^3 + bL^2 + cL + d\right)\partial L \right\}$$

$$\Delta PS = \sum_{hours} \left\{ P^* L^* - P_0 L_0 - \left[\left(\frac{a}{4}L^4 + \frac{b}{3}L^3 + \frac{c}{2}L^2 + dL\right)\right]_{L_0}^{L^*} \right\}$$

With flat-rate or TOU pricing there is deadweight loss in both high-priced hours and low-priced hours. Because the RTP case has no deadweight loss, we calculate the deadweight loss in the flat rate and TOU cases based on the surplus changes in Equation (10). Both deadweight loss and LSE profit are shown in Figure 1 for a high-priced hour. Although the local utility may have a positive or negative profit in any one hour with TOU or flat rate, it has zero profit over the year under any of these pricing scenarios.

$$\text{(10)} \qquad \begin{aligned} DW_{flat} &= \Delta\Pi_{flat}^{RTP} + \Delta CS_{flat}^{RTP} + \Delta PS_{flat}^{RTP} = \Delta CS_{flat}^{RTP} + \Delta PS_{flat}^{RTP} \\ DW_{TOU} &= DW_{flat} - \Delta DW_{flat}^{TOU} = DW_{flat} - \left(\Delta CS_{flat}^{TOU} + \Delta PS_{flat}^{TOU}\right) \end{aligned}$$

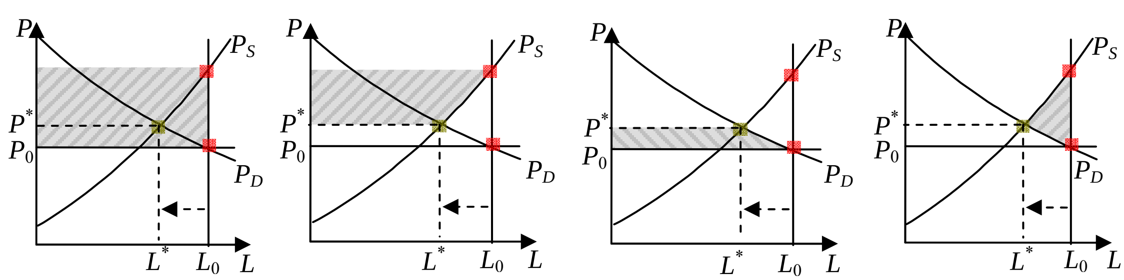

Figure 1. From left: local utility loss, producer surplus drop, consumer surplus drop, and deadweight loss.

5 Load Shifting

Assume, for whatever reason, that customers can be induced to shift their demand to be more level over the day. Although the resulting load profiles may not be realistic, the simulation shows how much shifting is necessary to flatten load and how quickly savings can be achieved.

5.1 Method

We scale possible consumer savings from demand response by incrementally shifting load to achieve a totally flat daily load profile without changing total consumption. Although this method does not consider real-world preference effects, it does set an upper bound on customer savings. The simulation allows load shifting to any other time of day but does not allow shifting from one day to another.

For a particular day, shift an increment of demand from the highest load hour to the lowest load hour. Continue shifting demand increments so that there is one price for the hours of greatest use and another (lower) price for the hours of least use. Stop shifting use when the quantity and price are the same for the high and low-priced hours. The maximum fraction *f* that is curtailed off the peak load hours is the same for all days.

Figure 2 illustrates the effects of shifting on load and price profiles of one week beginning Monday, June 19, 2006. This week originally exhibited moderately high load and price. Results are shown when 3% of all yearly MWh are shifted and after the maximum 5.3% of MWh are shifted. This method does not change total daily consumption in MWh, but the extremes of usage and price variation are reduced.

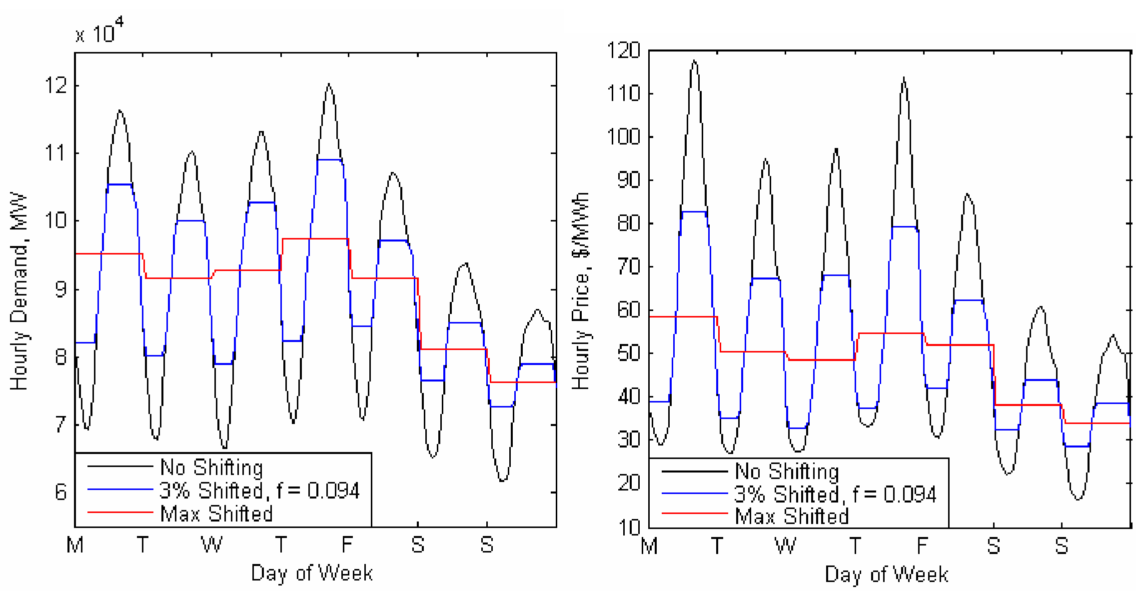

Figure 2. Load and price profiles for a July week; base case, 3% shifting ($f = 0.093$), and maximum shift.

5.2 Results

We remind the reader that the load shifts are imposed, rather than resulting from consumer preferences and so no conclusions can be drawn about consumers being better or worse off.

Customer expenditure savings from load shifting are shown in Figure 3. Savings are also split out by the amounts received by shifters and those received by free riders that do nothing. The left-hand plot in Figure 3 displays decreasing marginal savings with more shifting; when the daily load is leveled, there are no further savings. The right-hand plot of Figure 3 shows that shifters' percentage savings drop with increased shifting. This is because the price differential over a given day can be large under current conditions but approaches zero in the limit; small marginal savings steadily reduce average calculated savings. Total customer savings increase with the amount of shifting with an ultimate limit of 10.7% of the annual electric bill.

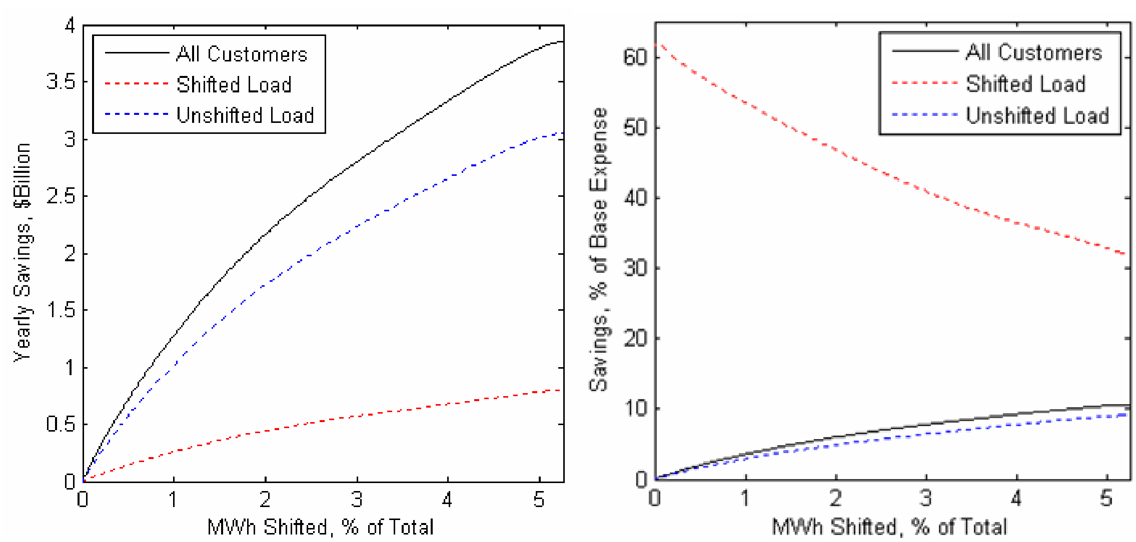

Figure 3. Savings to shifters, free riders, and total in dollars (left) and as a percentage of bill (right).

Load shifting reduces peak load dramatically as shown in Table 1, obviating the need for costly investment in generation and transmission.

Table 1. Peak load and overall cost savings with daily shifting.

Shifted Load, %	Peak Load, GW	Peak Load Saved	Total Expense, $Billion	Average Cost, $/MWh	Customer Bill Savings
0%	145	0.0%	$36.17	$51.96	0.0%
1%	138	4.8%	$34.90	$50.13	3.5%
2%	134	7.3%	$34.03	$48.88	5.9%
3%	131	9.3%	$33.37	$47.94	7.7%
4%	128	11.6%	$32.84	$47.17	9.2%
5%	122	15.8%	$32.38	$46.51	10.5%
5.3%	122	15.8%	$32.32	$46.43	10.7%

Table 2 shows how quickly customer savings are reached by load shifting. Half of all the possible savings from load shifting are achieved by shifting only 1.69% of all energy. This indicates that a small amount of demand response is all that is needed to get most of the benefits.

Table 2. Load shifting necessary to achieve a portion of limiting savings with daily shifting.

% of Savings in Limit	% Load Shifted	Maximum Hourly % Curtailed
25%	0.70%	3.9%
50%	1.69%	6.6%
75%	3.15%	9.6%
90%	4.26%	12.4%
95%	4.66%	14.0%
99%	5.06%	16.5%

6 Time of Use and Real Time Pricing

A simulation is used to determine the magnitude of effects from RTP and TOU in the wholesale market.

6.1 Sample Price and Load Profiles

The new price and load under RTP and TOU conditions are calculated as in Section 4.1. Figure 4 shows price and load profiles over a week in the base case, under TOU, and under RTP conditions with elasticity -0.2. For reference the retail rates p_0 and p_{TOU} are shown in dashed lines for the flat-rate and TOU cases respectively. The June week shown originally had moderately high load and wholesale price, so the RTP case shows steep drops in price and load during peak hours.

The left-hand graph in Figure 4 shows that RTP reduces peak loads much more than TOU pricing, which is only slightly better than flat rate pricing. The right-hand graph shows wholesale prices reflecting the marginal generation cost as solid lines; retail tariffs are in dashed lines. Under RTP the wholesale and retail prices are the same solid red line. Wholesale price peaks are moderated much more under RTP than under TOU pricing. A TOU rate actually exacerbates wholesale price peaks on weekends because end users see the off-peak price all day.

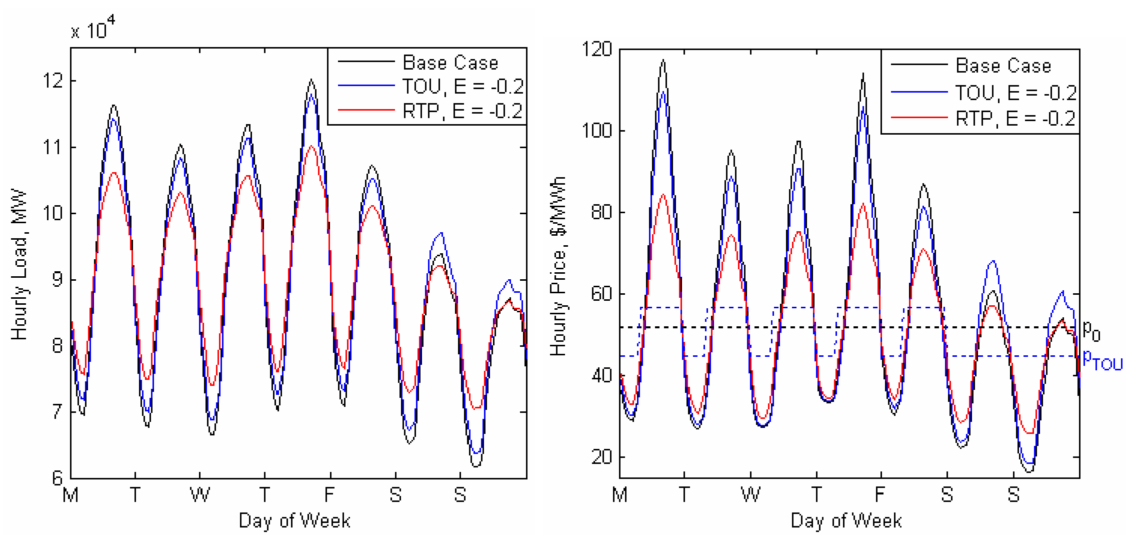

Figure 4. Load and price profiles with elasticity -0.2 for a July week with flat-rate, TOU, and RTP.

6.2 Economic Impacts

Market outcomes depend on the assumed demand elasticity[10]. On-peak, off-peak, and average wholesale prices are shown in the left-hand side of Figure 5 for TOU pricing and in the right-hand side for RTP. Prices drop more with RTP; they are about 4% lower. Both schemes moderate on-peak and off-peak prices on average, but these results say nothing about the most extreme prices. Table 3 shows the same prices as in Figure 5 at sample customer elasticities. A regulator looking only at prices might be deceived by the apparently small difference between RTP and TOU.

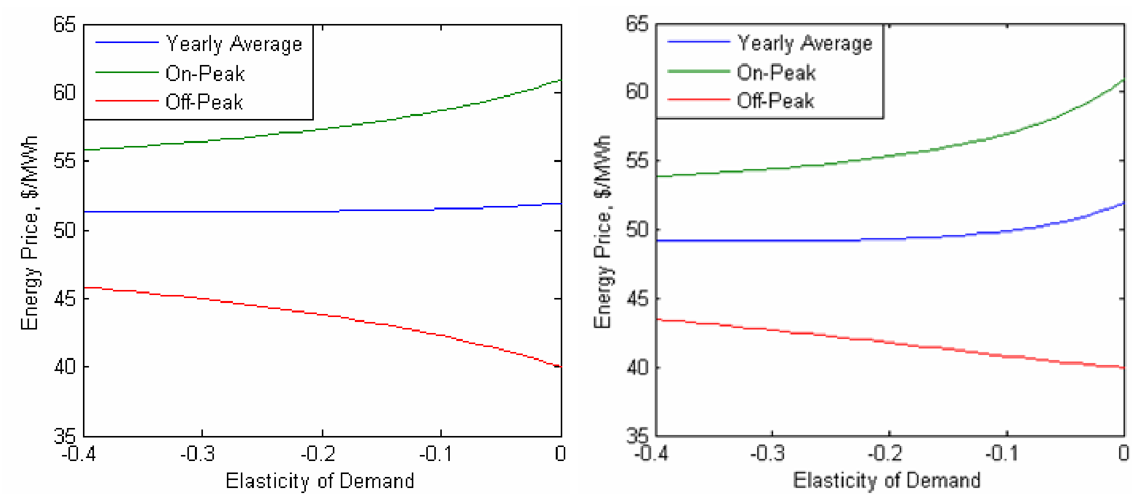

Figure 5. On-peak, off-peak, and average prices under the TOU scenario (left) and RTP scenario (right).

Table 3. Yearly prices with a change to TOU or RTP.

Elasticity of Demand	Average Price, $/MWh		On-Peak Price, $/MWh		Off-Peak Price, $/MWh	
	TOU	RTP	TOU	RTP	TOU	RTP
0	$51.96	$51.96	$60.92	$60.92	$40.01	$40.01
-0.05	$51.72	$50.82	$59.87	$58.86	$41.03	$40.28
-0.1	$51.54	$50.02	$58.86	$57.23	$42.08	$40.72
-0.15	$51.44	$49.59	$58.08	$56.17	$42.96	$41.20
-0.2	$51.38	$49.35	$57.46	$55.43	$43.70	$41.69
-0.25	$51.35	$49.23	$56.95	$54.89	$44.33	$42.16
-0.3	$51.34	$49.18	$56.53	$54.49	$44.87	$42.61
-0.35	$51.34	$49.18	$56.18	$54.17	$45.34	$43.04
-0.4	$51.34	$49.20	$55.87	$53.92	$45.75	$43.43

[10] Customers are more responsive when elasticity is more negative; responsiveness increases as one moves to the left in these plots.

Table 4 and Table 5 summarize impacts on consumption, expense, and peak load with TOU and RTP rates respectively. The impacts from TOU pricing are a fraction of those from RTP. Impacts from TOU in peak load shaved, consumption increase, and consumer expense saved are never more than 14.4%, 22.3%, and 21.9% respectively of the impacts from changing to RTP at any elasticity.

Impacts on consumer expense and consumption increase are small under either rate structure change. The most striking result in these tables is that with RTP, peak load reductions are large even with highly (but not completely) inelastic demand. We estimate a 10.4% reduction in peak demand at elasticity E = -0.1, a huge reduction at a modest assumed responsiveness. Holland and Mansur's prediction is less than half ours at 3.91%, while Borenstein's estimate is more than twice the size at 24.5% [11] [7, 10].

Table 4. Load increase, peak shaving, and price savings with TOU pricing.

Elasticity of Demand	Peak Load, GW	Peak Load Saved	Total Energy, TWh	Consumption Increase	Total Expense, $Billion	Consumer Expense Saved	Average Price, $/MWh
0	145	0.0%	696	0.0%	$36.17	0.0%	$51.96
-0.05	144	0.6%	697	0.1%	$36.04	0.4%	$51.72
-0.1	143	1.1%	697	0.2%	$35.95	0.6%	$51.54
-0.15	143	1.5%	698	0.3%	$35.91	0.7%	$51.44
-0.2	142	1.9%	699	0.4%	$35.90	0.8%	$51.38
-0.25	142	2.2%	699	0.4%	$35.90	0.7%	$51.35
-0.3	141	2.4%	700	0.5%	$35.91	0.7%	$51.34
-0.35	141	2.6%	700	0.5%	$35.93	0.7%	$51.34
-0.4	141	2.8%	700	0.6%	$35.95	0.6%	$51.34

Table 5. Load increase, peak shaving, and price savings with RTP.

Elasticity of Demand	Peak Load, GW	Peak Load Saved	Total Energy, TWh	Consumption Increase	Total Expense, $Billion	Consumer Expense Saved	Average Price, $/MWh
0	145	0.0%	696	0.0%	$36.17	0.0%	$51.96
-0.05	137	5.7%	699	0.4%	$35.52	1.8%	$50.82
-0.1	130	10.4%	702	0.8%	$35.11	2.9%	$50.02
-0.15	126	13.3%	705	1.2%	$34.94	3.4%	$49.59
-0.2	123	15.1%	707	1.6%	$34.90	3.5%	$49.35
-0.25	121	16.6%	709	1.9%	$34.93	3.4%	$49.23
-0.3	119	17.7%	711	2.2%	$34.99	3.3%	$49.18
-0.35	118	18.7%	713	2.4%	$35.07	3.0%	$49.18
-0.4	117	19.5%	715	2.7%	$35.16	2.8%	$49.20

[11] Holland and Mansur also predict a 5.88% peak load reduction at E = -0.2, where we predict a 15.1% savings. Borenstein also predicts 35.2% peak load reduction at E = -0.3 where we predict a 17.7% savings. The modest impacts predicted by Holland and Mansur are largely dictated by their method of using a stacked bid curve, see Appendix B. Borenstein's large projected peak reduction has to be understood knowing that his supply curve comprised of three generator types results in a load duration curve that is completely chopped off on the high end; he does not argue that this is a realistic resulting load duration curve.

Costumers buy more energy under RTP or TOU conditions as shown in Figure 6; marginal impacts diminish with more responsive load. Customer expenditure on electricity decreases steeply if elasticity is low in magnitude as shown in Figure 7. With inelastic demand most of the changes in consumption patterns are small reductions at peak prices. With greater elasticity, savings drop as the effect of greater consumption dominates the overall expense.

These RTP results are explained by the large positive skew in electricity prices and the increasing steepness of supply curves at high load. Large price reductions from small amounts of curtailment at high prices dominate results at elasticities near zero. With increasing responsiveness, the load profile becomes flatter and flatter but overall consumption increases. Under these conditions, the effect of the consumption increase dominates other results. Results with TOU pricing have similar characteristics but only a fraction of the magnitude.

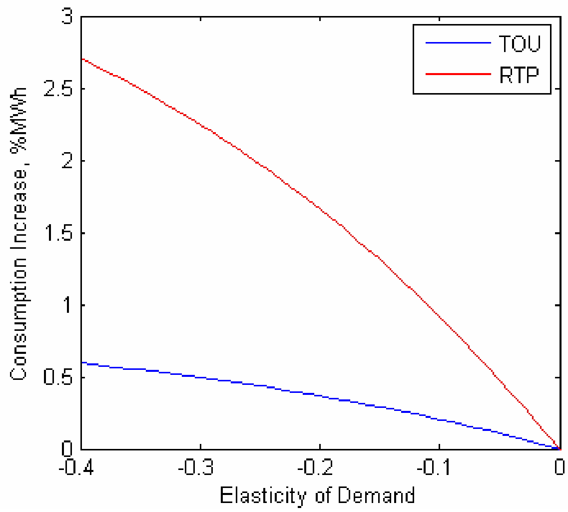

Figure 6. Consumption increase, TOU and RTP.

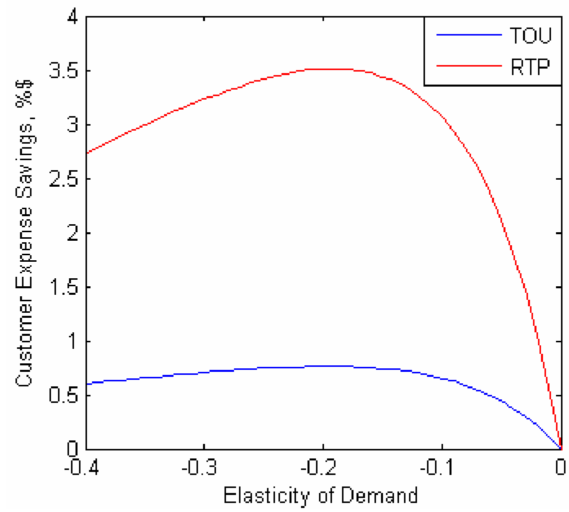

Figure 7. Customer bill savings, TOU and RTP.

Because consumers are buying more energy with less total expenditure, the overall impact on consumers is more easily understood by looking at a customer who refuses to change behavior as others do under TOU or RTP. In Figure 8, savings are shown for a single customer who has elasticity zero, while the aggregate system has an elasticity shown on the x-axis. We show savings for three types of customers:

Flat – Customer uses a constant level of power during all hours of the year.
Typical – Customer load profile is proportional to the original system load profile.
50% More Extreme – During each hour, the customer demands the typical customer's load *plus* an additional 50% of the difference between the typical customer's load for that hour and the minimum load for the day.

An unchanging typical customer saves less per unit than a responsive customer, but slightly more overall because she does not increase consumption[12]. More interesting is that a flat customer would save 7.0% of her annual electric bill even if no one responded to price. She would save the amount that currently goes to subsidize the excesses of more peaky customers. The more extreme customer loses money under RTP if no one responds, but will have net savings if the aggregate elasticity is even slightly responsive, $E \leq -0.04$.

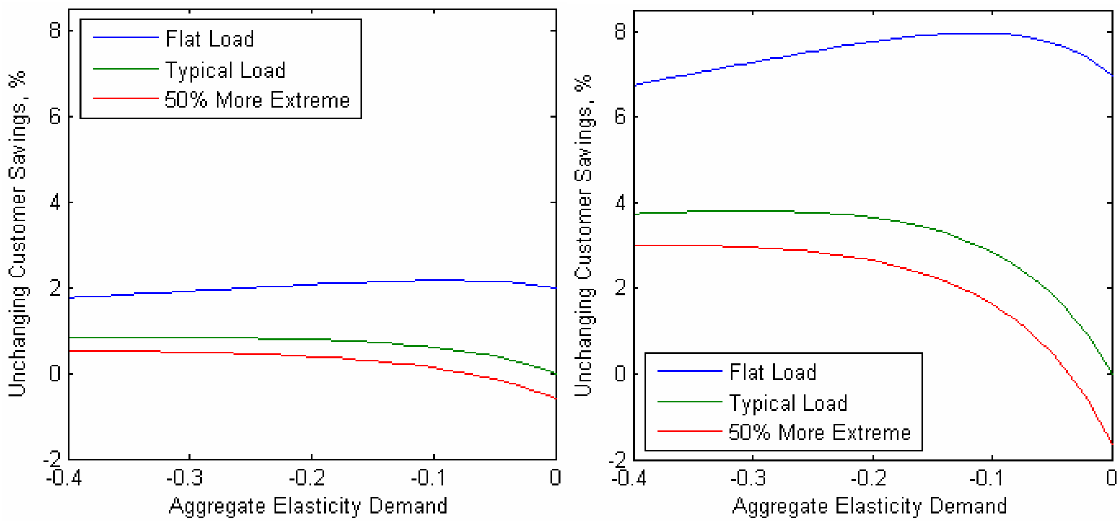

Figure 8. Expense savings to an unresponsive customer when others respond, TOU (left) and RTP (right).

Peak load reductions are extreme with a small amount of responsiveness but marginal savings taper with greater responsiveness as shown in Figure 9. Discontinuities in Figure 9 are caused by a change in the day upon which peak load is observed.

The large peak load savings under RTP have huge implications for the total system cost. Peak load determines the total capacity investment necessary for the system to operate reliably. Although no savings will be made on peak capacity that has already been built, there will be savings via unneeded capacity investment as generators have to be replaced or load increases over time. At elasticity -0.2, peak load drops by 15.1% with RTP. At that level, an overnight capacity value of $600/kW or $1800/kW, corresponding roughly with the capital costs of gas and coal generation, translates into a dollar savings of $13-$39 billion from a change to RTP. A change to TOU pricing would reduce $1.7 to $5.0 billion in capacity investments under the same conditions.

At $13 billion, capacity savings would be $257 for each of the 51 million people in PJM territory [18]. Compared to the hardware and installation costs of $123-$215 per unit for the advanced metering infrastructure required to implement RTP, these capacity savings justify RTP rates starting with the largest and most responsive customers [5]. We conjecture that only large customers, those responsible for perhaps 50% of total load, need to face RTP to achieve these savings. If only 10% of customers need smart meters,

[12] At E = -0.2, the typical responsive customer saves 5.0% per unit and 3.5% overall; the typical unresponsive customer saves 3.6% although her quantity consumed is constant.

RTP, and automatic energy managers to respond to price, the cost of implementing RTP would be much smaller than the social benefit, with large benefits to all customers.

In previous work using data covering the summer of 2005, we found smaller, but still sizeable, peak load reductions. Because 2005 had a mild summer and lower peak load, RTP would not have had as big an impact on peak reductions. Since the highest load over many years determines the capacity investment necessary for system stability, the peakier 2006 data are more useful for estimating RTP impacts on necessary peak generation. Other results do not change appreciably using the updated 2006 data.

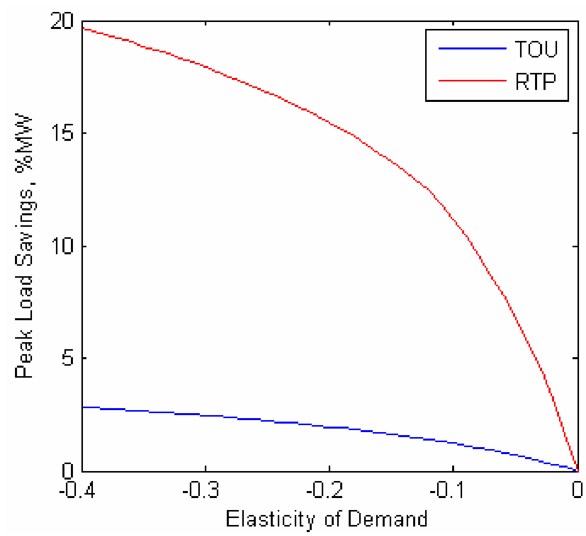

Figure 9. Peak load reductions, TOU and RTP.

Figure 10 and Table 6 show surplus increases with a time-varying rate. Neither consumer nor producer surplus changes monotonically with elasticity. Producer surplus drops slightly with peak price reductions but then increases with overall consumption. Producer surplus is equal to revenue minus operating costs and so indicates profitability if capital costs are not considered. Because we see almost no change in producer surplus, these results indicate that producers will not see the large reduction in profits that they might have feared from RTP. There is no change in consumer surplus for an elasticity of zero, but for an elasticity of -0.2, consumer surplus increases 0.7% for TOU pricing and 3.2% for RTP. We find that TOU pricing has only 20.3%-21.8% the impact in increasing total surplus that RTP would have[13]. No matter what the assumed elasticity, consumer surplus increases with RTP or TOU[14].

[13] Although the magnitude of our surplus estimates are much smaller than Borenstein's and much larger than Holland and Mansur's, the ratio of surplus increases between TOU and RTP are remarkably close given the different definitions of TOU used in each case. Borenstein predicted that TOU would have 8-25% the effect of RTP on surplus; Holland and Mansur predicted 15% [7, 10].

[14] The reason for the lack of monotonicity in consumer surplus can be understood by seeing what happens to the area representing ΔCS in Figure 1 with extremely steep, moderate, and extremely flat demand curves. A similar figure should be drawn and examined for the case in which load and price increase with RTP.

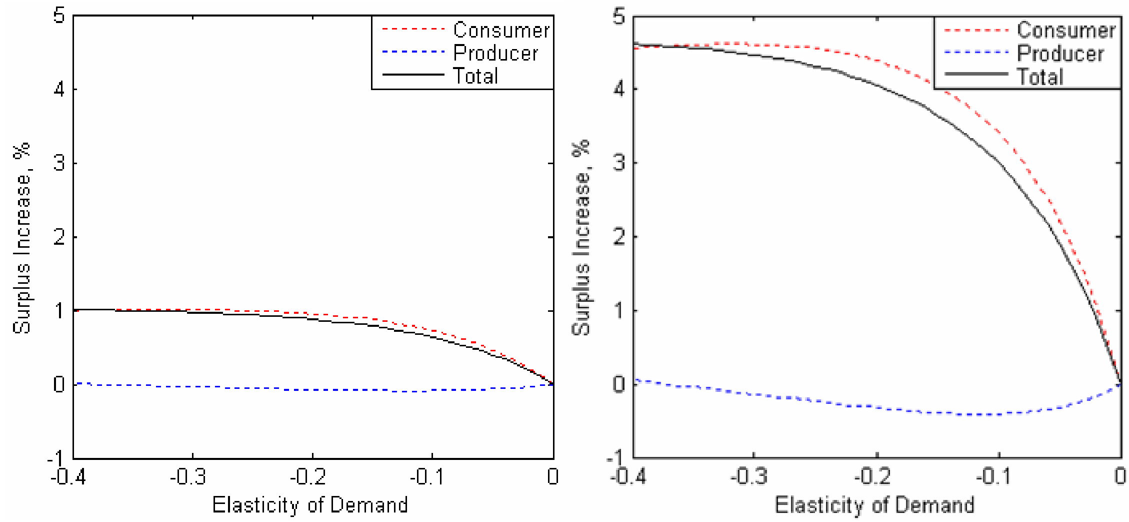

Figure 10. Surplus increases with TOU (left) and RTP (right) as a percent of baseline expense.

Table 6. Economic outcomes with RTP as a percentage of baseline expenditure.

Elasticity of Demand	Flat-Rate Deadweight Loss	TOU Rate Deadweight Loss	Surplus Increase with TOU			Surplus Increase with RTP		
			Consumer	Producer	Total	Consumer	Producer	Total
0	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
-0.05	1.6%	1.3%	0.4%	-0.1%	0.3%	1.9%	-0.3%	1.6%
-0.1	2.8%	2.2%	0.7%	-0.1%	0.6%	3.2%	-0.4%	2.8%
-0.15	3.5%	2.8%	0.8%	-0.1%	0.8%	3.9%	-0.4%	3.5%
-0.2	4.0%	3.1%	0.9%	-0.1%	0.9%	4.3%	-0.3%	4.0%
-0.25	4.3%	3.3%	1.0%	-0.1%	0.9%	4.5%	-0.3%	4.3%
-0.3	4.4%	3.5%	1.0%	0.0%	1.0%	4.6%	-0.2%	4.4%
-0.35	4.5%	3.6%	1.0%	0.0%	1.0%	4.6%	-0.1%	4.5%
-0.4	4.6%	3.6%	1.0%	0.0%	1.0%	4.6%	0.0%	4.6%

Before looking at these results, a regulator might be concerned about charging RTP for customers who have no ability to respond. It would seem unfair to charge customers high RTPs if they could not react. These results indicate that even if customers had no means of knowing or responding to the RTP, the adverse effect of extremely high prices would not cause any problems on average over the year. Flat and countercyclical customers would benefit by not having to subsidize the excesses of others. Even customers with problematic load profiles would not have a large change in overall cost and would actually save money from other customers' responses. These results indicate that regulators do not need to exercise great caution in implementing RTP at the retail level.

7 Conclusions and Recommendations

Pretending that consumers demand the same amount of electricity no matter what the price or that consumers cannot vary their demand as prices change has cost consumers dearly and led to large, unnecessary investments in peaking plants. In 2006, 15% of the generation capacity in PJM territory ran only 1.1% or fewer hours, and 20% of capacity ran only 2.3% or fewer hours [1] [15]. These under-utilized peak generation investments are a luxury that neither providers nor customers should have to pay for.

The good news is that the peak load problem can be mitigated by moving flat rate customers onto RTP tariffs. Even with little price responsiveness, surprisingly large peak load reductions can be achieved; at elasticities -0.1 and -0.2, 10.4% and 15.1% respectively can be shaved off of coincident peak consumption. Most other quantities of interest such as generator profitability, overall consumption, and average end user expense will not be affected greatly by a change toward RTP. However, policy makers will be disappointed with the short-term reduction in overall bills. A move toward RTP should be driven by concerns about peak load and fairness among end users, with larger savings from the lower investment in peak generators and transmission lines.

Under current conditions counter-cyclical end users subsidize the problematic load profiles of others. Even if some customers do not want to alter consumption habits, the rest of the system should not have to pay for their excesses. When problematic customers are confronted with higher bills, they will want to make small but important changes. Just as consumers have learned to respond to the volatile prices of gasoline, fruits and vegetables, and other commodities, so they can learn to respond to electricity prices. The largest difference is that customers purchase electricity every hour of the year and therefore need automated devices to react to changing prices.

Because only modest aggregate price elasticities are necessary for large peak capacity savings, most of the benefits can be achieved by shifting only large, responsive customers to RTP. These are the customers who would benefit the most by installing the relatively inexpensive equipment necessary for automated response to RTP. With RTP, each customer is free to react in the ways that best serve her interest.

8 Acknowledgements

Kathleen Spees acknowledges the National Science Foundation Graduate Research Fellowship Program and the Achievement Rewards for College Scientists Foundation of Pittsburgh for funding her graduate studies. Additional funding was provided by the Carnegie Mellon Electricity Industry Center.

[15] This is based on the entire PJM hourly load profile in 2006 [1]. Even at peak load, the system had 17.5% excess available generation capacity. I do not include generation excess at coincident peak load in this calculation because some generation excess is necessary for reliability purposes.

9 References

[1] PJM Market Data. Available: http://www.pjm.com/markets/market-monitor/data.html

[2] *2005 Price Responsive Load Survey Results.* Available: http://www.pjm.com/committees/working-groups/dsrwg/downloads/20060615-05-price-responsive-load-survey.pdf

[3] *Assessment of PJM Load Response Programs.* PJM Market Monitoring Unit. Report to the Federal Energy Regulatory Commission, Docket No. ER02-1326-006. August 29,2006. Available: http://www.pjm.com/markets/market-monitor/downloads/mmu-reports/dsr-report-2005-august-29-%202006.pdf

[4] Barbose, Galen, Charles Goldman, and Bernie Neenan. *A Survey of Utility Experience with Real Time Pricing.* Lawrence Berkeley National Lab. December 2004. Available: http://www.osti.gov/bridge/servlets/purl/836966-SZe2FO/native/836966.pdf

[5] *Assessment of Demand Response and Advanced Metering*. Federal Energy Regulatory Commission, Staff Report. Docket Number: AD-06-2-00. August 2006. Available: http://www.ferc.gov/legal/staff-reports/demand-response.pdf

[6] Heffner, Grayson C. "Configuring Load as a Resource for Competitive Electricity Markets - Review of Demand Response Programs in the US and Around the World." Lawrence Berkeley National Laboratory. November 2002. Available: http://repositories.cdlib.org/cgi/viewcontent.cgi?article=1747&context=lbnl

[7] Borenstein, Severin. "The Long-Run Efficiency of Real-Time Electricity Pricing." *The Energy Journal.* 2005; 26; 3. pp 93. Available: http://www.ucei.berkeley.edu/PDF/csemwp133r.pdf

[8] *FERC Order on Section 206 Investigation into WECC-Wide Price Cap and the CAISO Ancillary Service Capacity Bid Cap*. February 13,2006. Available: http://www.caiso.com/179e/179ead8256e10.pdf

[9] *RTO-ISO Handbook, Revised - February 28, 2006.* Federal Energy Regulatory Commission. Available: http://www.ferc.gov/industries/electric/indus-act/rto/handbook.asp#skipnavsub

[10] Holland, Stephen P, and Erin T Mansur. "The Short-Run Effects of Time-Varying Prices in Competitive Electricity Markets." *The Energy Journal*, Vol. 27, No. 4.

[11] Holland, Stephen P. and Erin T. Mansur. *Is Real-Time Pricing Green? The Environmental Impacts of Electricity Demand Variance.* Available: http://repositories.cdlib.org/cgi/viewcontent.cgi?article=1039&context=ucei/csem

[12] Wang, Jing, Nuria Encinas Redondo, and Francisco D Galiana. "Demand-Side Reserve Offers in Joint Energy/Reserve Electricity Markets." *IEEE Transactions on Power Systems*, Vol. 18, No. 4, November 2003. Available: http://ieeexplore.ieee.org/iel5/59/27910/01245550.pdf?arnumber=1245550

[13] Boisvert, Richard N, Peter Cappers, Charles Goldman, Bernie Neenan, and Nicole Hopper. "Customer Response to RTP in Competitive Markets: A Study of Niagara Mohawk's Standard Offer Tariff," *The Energy Journal,* Vol. 28, No. 1. 2007.

[14] Neenan, Bernie, Donna Pratt, Peter Cappers, and Richard Boisvert. *Does Real-Time Pricing Deliver Demand Response? A Case Study of Niagara Mohawk's Large Customer RTP Tariff.* Lawrence Berkeley National Laboratory. August 2004. Available: http://eetd.lbl.gov/ea/EMS/reports/54974.pdf

[15] King, Chris S, and Sanjoy Chatterjee. Predicting California Demand Response: How do Customers React to Hourly Prices? Public Utilities Fortnightly, July 1, 2003. Available: http://www.americanenergyinstitutes.org/research/CaDemandResponse.pdf

[16] *Benefits of Demand Response in Electricity Markets and Recommendations for Achieving Them: A Report to the United States Congress Pursuant to Section 1252 of the Energy Policy Act of 2005*. US Department of Energy, February 2006. Available: http://www.oe.energy.gov/DocumentsandMedia/congress_1252d.pdf

[17] *PJM Manual 06: Financial Transmission Rights, Revision 8.* Effective Date: March 08, 2006. Available: http://www.pjm.com/contributions/pjm-manuals/pdf/m06v08.pdf

[18] *PJM Overview*. Available: http://www.pjm.com/about/overview.html

[19] Navidi, William. *Statistics for Engineers and Scientists.* McGraw-Hill. 2006.

[20] Krapels, Edward N., and Paul Flemming. *Impacts of the PJM RTO Expansion. November 2005.* Available: http://www.pjm.com/documents/downloads/reports/20051101-impact-pjm-expansion.pdf

[21] Kirschen, Daniel S. "Demand-Side View of Electricity Markets." *IEEE Transactions on Power Systems,* Vol. 18, No. 2, May 2003. Available: http://ieeexplore.ieee.org/iel5/59/26968/01198281.pdf?arnumber=1198281

[22] "Natural Gas Summary." *Natural Gas Navigator.* Energy Information Administration. Available: http://tonto.eia.doe.gov/dnav/ng/ng_sum_lsum_dcu_nus_a.htm

Appendix A Detail on Building a General Supply-Side Model

A.1 Possible Model Structures

We have examined several possible models for predicting price from load using variations on the third-degree polynomial in Equation (11).

$$(11) \qquad P_S(L) = \sum_{t=1}^{n} \left\{ \delta_3 \cdot a_t \cdot L^3 + \delta_2 \cdot b_t \cdot L^2 + \delta_1 \cdot c_t \cdot L + \delta_0 \cdot d_t \right\}$$

Each term in (11) is multiplied by a dummy variable δ_i that has possible values one and zero. These dummy variables act as on-off switches for the term parameter based on the time period t. For example, if we want to assume that each day has a unique third degree supply curve, then the number of time periods is $n = 365$ and each term will have $n = 365$ different parameters a_t, b_t, c_t, and d_t, one for each day. The dummy variables δ_0, δ_1, δ_2, and δ_3 ensure that only the parameters appropriate for the time period in question are considered; all others are zeroed out. The resulting values of $P_S(L)$ are then the same as they would have been had we fit 365 different third degree polynomials to the data; overall goodness of fit statistics for that model would have $4 \cdot n = 1460$ parameters.

Another advantage of the dummy variable approach is that I am able to selectively drop dummy variables from the model to simplify it. For example, my conclusion from these regressions is that the daily fits are the same shape in the second and third degree terms as long as a linear offset is applied to each day as in (7). The simplified model includes only the zero and first degree dummy variables δ_0 and δ_1, and has $2 \cdot n+2 = 732$ p.

We have examined models with eight different definitions of time period *t* and consequent values for *n*. Table 7 is a summary description of each of the eight time period definitions and the number of model parameters resulting from including a given number of dummy variables from 1 to 4. In each case except for the yearly model there are 15 ways to combine dummy variables.

Table 7. Descriptions of the eight time period definitions examined.

Time Period	***n***	**Number of Parameters by Included Dummy Variables**				**Description**
		1 δ_0	**2** δ_0, δ_1	**3** $\delta_0, \delta_1, \delta_2$	**4** $\delta_0, \delta_1, \delta_2, \delta_3$	
Year	1	NA	2	3	4	One curve. Dropping a dummy means dropping the entire term.
Month of Year	12	15	26	37	48	One curve for each month.
Week of Year	53	56	108	160	212	Week is Mon-Sun. Data begin and end with Wed.
Day of Year	365	368	732	1096	1460	One curve for each day.
Week or Weekend of Year	105	108	212	316	420	One curve for each week Mon-Fri; one curve for each weekend Sat-Sun.
Week or Weekend, Holidays as Weekend	105	108	212	316	420	Append 6 NERC holidays[16] to closest weekend, all happen to fall on Mon or Fri.
Day of Week	7	10	16	22	28	One curve for each day of week.
Hour of Day	24	27	50	73	96	One curve for each hour of day.

[16] North American Electric Reliability Council (NERC) holidays are considered off-peak hours in PJM [17]

A.2 Statistical Significance and Goodness of Fit

For each of the 109 models we have evaluated the goodness of fit statistics. By examining adjusted R^2 values indicating explanatory power, we have concluded that the best way to drop dummy variables is starting with the highest order term and working downward. That is to drop δ_3, then δ_3 and δ_2, then δ_3, δ_2, and δ_1. This ordering is consistent for almost all model types[17]. Table 8 displays these adjusted R^2 values for the 31 models consistent with this drop ordering. Models are listed in order of decreasing explanatory power; the ordering of models by explanatory power is identical no matter how many dummy variables are included. Even after dropping two dummy variables, the daily model has more explanatory power than any other model with all four dummy variables.

Table 8. Model adjusted R^2 values.

Model Sorted in Order of Descending Adjusted R^2	Dummy Variables Included			
	1 δ_0	2 δ_0, δ_1	3 $\delta_0, \delta_1, \delta_2$	4 $\delta_0, \delta_1, \delta_2, \delta_3$
Day of Year	0.9096	0.9488	0.9630	0.9661
Week/WeekendorHoliday	0.8866	0.9124	0.9223	0.9241
Week/Weekend	0.8859	0.9118	0.9221	0.9240
Week of Year	0.8725	0.8961	0.9061	0.9079
Month of Year	0.8521	0.8774	0.8853	0.8887
Hour of Day	0.7990	0.8151	0.8208	0.8225
Day of Week	0.7942	0.8001	0.8085	0.8088
Year	--	0.6925	0.7453	0.7805

The same ordering for dropping dummy variables is dictated by the F-statistic for overall model significance. Because of the large DOF, the p-values associated with these F-statistics are vanishingly small and therefore uninformative.

Table 9. Overall model F-statistics.

Model Sorted in Order of Descending Adjusted R^2	Number of Dummy Variables Included			
	1 δ_0	2 δ_0, δ_1	3 $\delta_0, \delta_1, \delta_2$	4 $\delta_0, \delta_1, \delta_2, \delta_3$
Day of Year	241	223	210	172
Week/WeekendorHoliday	641	433	331	256
Week/Weekend	637	430	330	255
Week of Year	1091	707	533	410
Month of Year	3607	2509	1879	1490
Hour of Day	1340	789	558	428
Day of Week	3758	2338	1762	1374
Year	--	19732	12815	10387

[17] In the day of week model, keeping higher order terms is preferred. Hour of day and month of year models also prefer a higher order term when only two dummy variables are included. These models are poor representations based on the adjusted R^2 values, and so I do not consider these issues further.

Model ordering is largely dictated by the number of parameters, the only exception being the month of year and versus hour of day models. Because the theoretical import of the model as decreases as the number of parameters increases, it may be a good idea to accept a model with less explanatory power to obtain a more elegant model. Table 10 shows the explanatory power as calculated by the adjusted R^2 value lost by dropping to the next best model.

Temporal resolution always improves the explanatory power of the model, but the meaning of this observation is clouded by the fact that the higher temporal resolution models use more parameters. The largest drop in explanatory power occurs when moving from sequential time-series to non-sequential bunches of data. That means that Mondays have no interesting common characteristics, but that hours within one day or one week *do* have common characteristics. We conclude from this observation that system conditions change slowly over time and that grouping consecutive hours is a good way to capture these effects.

Table 10. Adjusted R^2 lost by dropping to next best model.

Model Sorted in Order of Descending Adjusted R^2	**Number of Dummy Variables Included**			
	1 δ_0	**2** δ_0, δ_1	**3** $\delta_0, \delta_1, \delta_2$	**4** $\delta_0, \delta_1, \delta_2, \delta_3$
Day of Year	0.0230	0.0364	0.0407	0.0420
Week/WeekendorHoliday	0.0007	0.0006	0.0002	0.0001
Week/Weekend	0.0134	0.0157	0.0160	0.0161
Week of Year	0.0204	0.0187	0.0208	0.0192
Month of Year	0.0531	0.0623	0.0645	0.0662
Hour of Day	0.0048	0.0150	0.0123	0.0137
Day of Week	--	0.1076	0.0632	0.0283
Year	--	--	--	--

In deciding how many dummy variables to drop, it is useful to examine the explanatory power lost in dropping the least important dummy variable as shown in Table 11. We look primarily at the models with time-sequential data groupings. Dropping the δ_3 variable drops explanatory power a miniscule amount. Dropping δ_2 is only slightly worse. Based on this assessment, we conclude that including only linear offsets is a powerful way to represent price and load data. By dropping the number of dummy variables to two, the number of parameters in the model is roughly halved.

Table 11. Adjusted R^2 lost in dropping one dummy variable.

Model Sorted in Order of Descending Adjusted R^2	Adjusted R^2 Loss from Dropping One Dummy			
	δ_0 to Year	δ_0,δ_1 to δ_0	$\delta_0,\delta_1,\delta_2$ to δ_0,δ_1	$\delta_0,\delta_1,\delta_2,\delta_3$ to $\delta_0,\delta_1,\delta_2$
Day of Year	0.1291	0.0392	0.0142	0.0031
Week/WeekendorHoliday	0.1061	0.0258	0.0099	0.0018
Week/Weekend	0.1054	0.0259	0.0103	0.0019
Week of Year	0.0920	0.0236	0.0100	0.0018
Month of Year	0.0716	0.0253	0.0079	0.0034
Hour of Day	0.0185	0.0161	0.0057	0.0017
Day of Week	0.0137	0.0059	0.0084	0.0003
Year	--	--	0.0528	0.0352

For further insight in determining how many dummy variables to drop, we have calculated an F-statistic for model improvement with and without each dummy variable according to (12) from [19]. The variable k represents the number of parameters a_t through d_t; the variable *SSE* represents the sum of squared error between the real data and model prediction; *N* is the number of data. Subscripts *full* and *reduced* refer to the models with and without the dummy variable respectively.

$$(12) \qquad \frac{\left(SSE_{reduced} - SSE_{full}\right) / \left(k_{full} - k_{reduced}\right)}{SSE_{full} / \left(N - k_{full}\right)}$$

Calculated F-statistics are in Table 12; associated p-values are again vanishingly small. This indicates that keeping additional dummy variables would be justified, although the higher order dummy variables are less important.

Table 12. F-Statistic for testing the hypothesis that a model is no better than the next best model.

Model Sorted in Order of Descending Adjusted R^2	F-Statistic for Dropping to the Next Best Model			
	δ_0 to Year	δ_0,δ_1 to δ_0	$\delta_0,\delta_1,\delta_2$ to δ_0,δ_1	$\delta_0,\delta_1,\delta_2,\delta_3$ to $\delta_0,\delta_1,\delta_2$
Day of Year	35	19	10	3
Week/WeekendorHoliday[18]	80	25	11	3
Week/Weekend	79	25	12	3
Week of Year	123	39	19	4
Month of Year[19]	386	120	97	26
Hour of Day	36	23	24	4
Day of Week	13	20	65	113
Year	--	--	1814	1409

[18] The next best model for the F-statistic is “Week of Year”. The “Week/WeekendorHoliday” model has the same number of parameters and a smaller *SSE* than the “Week/Weekend” model, rendering the F-statistic meaningless for that pair.

[19] The next best model for the F-statistic is “Year”. The “Month of Year” model has fewer parameters and a smaller *SSE* than the “Hour of Day” model, rendering the F-statistic meaningless for that pair.

Each parameter in each model has a t-statistic and an associated p-value measuring its significance in improving the model. Some of the models we examined have upwards of one thousand parameters, so we have grouped the parameters d_t through a_t corresponding to term order zero through three respectively. Table 13 and Table 14 show mean and median t-test p-values in the daily and week or weekend-holiday model respectively.

When the dummy variable associated with each parameter is included, the number of parameters is large and the mean and median p-values are displayed without shading. Shaded p-values indicate that the associated dummy variable has been dropped and there is just one parameter of that order that applies to the entire model. In those cases mean and median are the same by definition and so only one is displayed. Bolded results indicate significance at the $p<0.05$ level.

Table 13. Daily model t-test p-values by parameter order and dummy variables included.

	Number of Dummy Variables Included			
Median p-Values	1 δ_0	2 δ_0, δ_1	3 δ_0, δ_1,δ_2	4 δ_0, δ_1,δ_2,δ_3
a_t	**0.000**	**0.000**	**0.000**	0.257
b_t	**0.000**	**0.000**	**0.000**	0.131
c_t	**0.000**	**0.008**	**0.018**	0.058
d_t	**0.001**	**0.000**	**0.000**	**0.002**
Mean p-Values	1	2	3	4
a_t				0.318
b_t			**0.046**	0.254
c_t		**0.000**	**0.000**	**0.000**
d_t	0.139	0.111	0.106	0.134

Table 14. Week or weekend-holiday t-test p-values by parameter and dummy variables included.

	Number of Dummy Variables Included			
Median p-Values	1 δ_0	2 δ_0, δ_1	3 δ_0, δ_1,δ_2	4 δ_0, δ_1,δ_2,δ_3
a_t	**0.000**	**0.000**	**0.000**	0.277
b_t	**0.000**	**0.000**	**0.001**	0.144
c_t	**0.000**	**0.000**	**0.001**	**0.033**
d_t	**0.000**	**0.000**	**0.001**	**0.003**
Mean p-Values	1	2	3	4
a_t				0.311
b_t			0.067	0.234
c_t		**0.000**	**0.000**	**0.000**
d_t	0.147	0.101	0.087	0.136

From Table 13 and Table 14 it is clear that if a dummy variable is dropped, then including a single parameter for that order term is always a statistically significant improvement to the overall model. The y-intercept, first order, and second order dummy variable parameters are statistically significant in the median but not always in the mean.

The median number is a more useful measure because these distributions have strong positive skews. The third order dummy variable parameters do not show statistical significance. Examination of these t-test results justifies dropping one dummy variable and keeping the remaining three.

A.3 Visual Examination of Model Characteristics

Adjusted R^2 results suggest moving ahead with the daily model and only two dummy variables. The t-test results indicate that including three dummy variables is justified. F-statistics suggest reducing the time resolution.

Figure 11 through Figure 14 display predictions from models including 1 through 4 dummy variables respectively. Original data are plotted in black in the background; curves representing each time-period *t* are plotted in red in the foreground. Prices for each time period *t* are plotted over the range of loads observed in that time period. Left-hand plots represent the daily models, right-hand plots represent the week/weekend-holiday models.

From Figure 11 it is clear that more than one dummy variable must be included in order to get a decent representation of the overall data characteristics. The weekly/weekend-holiday models in Figure 12 through Figure 14 do appear to represent general characteristics of the data but do poorly in the extremes. Especially obvious is the inability of the weekly/weekend-holiday models to capture the excessively high prices that are an important part of a demand-response analysis. The daily models are able to capture these high-price characteristics by including two to four dummy variables.

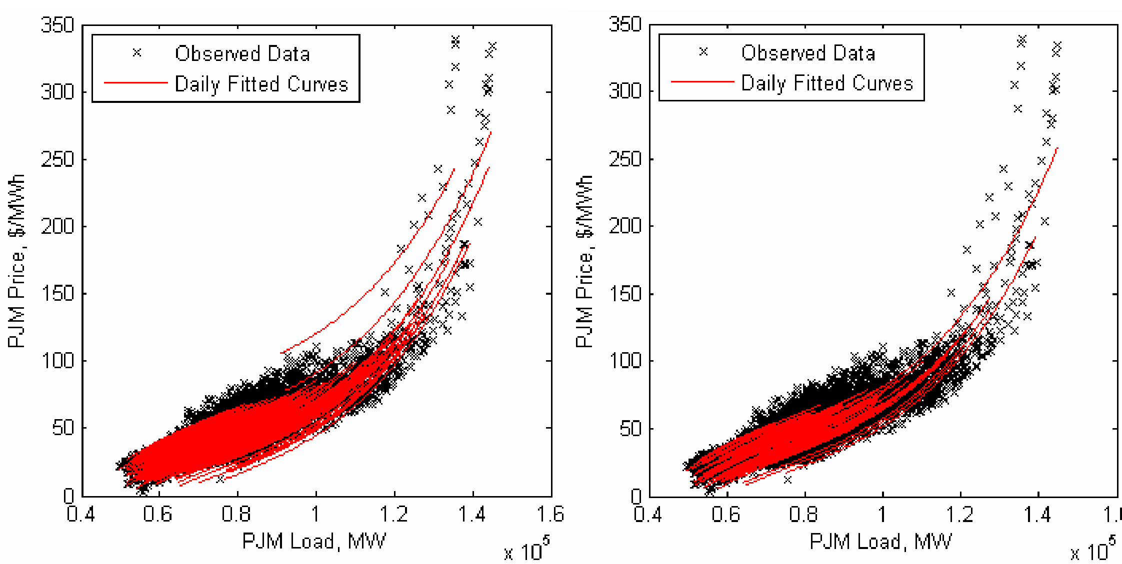

Figure 11. Data plotted with model curves using 1 dummy variable, daily (left) and week/weekend (right).

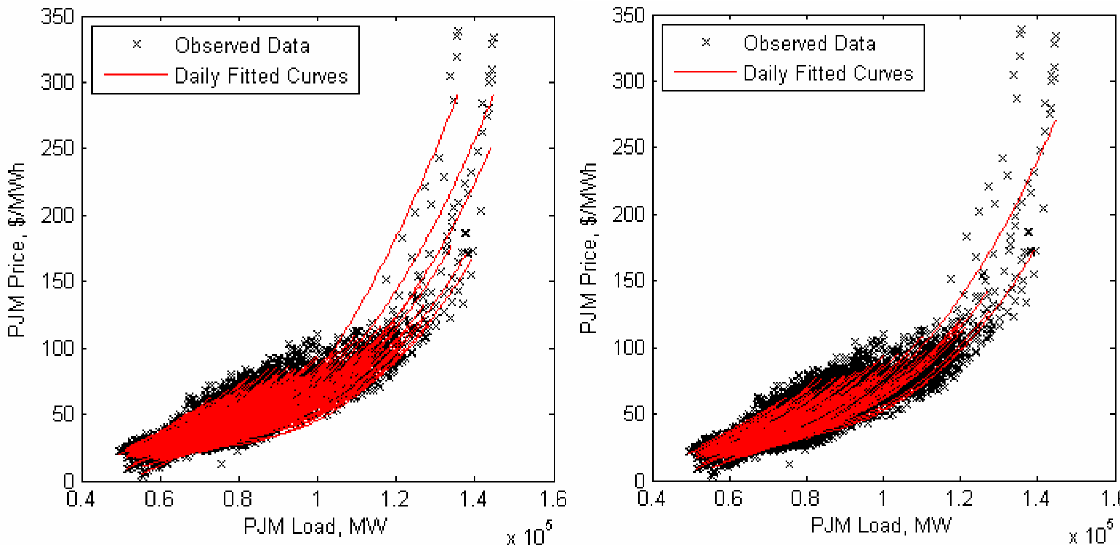

Figure 12. Data plotted with model curves using 2 dummy variables, daily (left) and week/weekend (right).

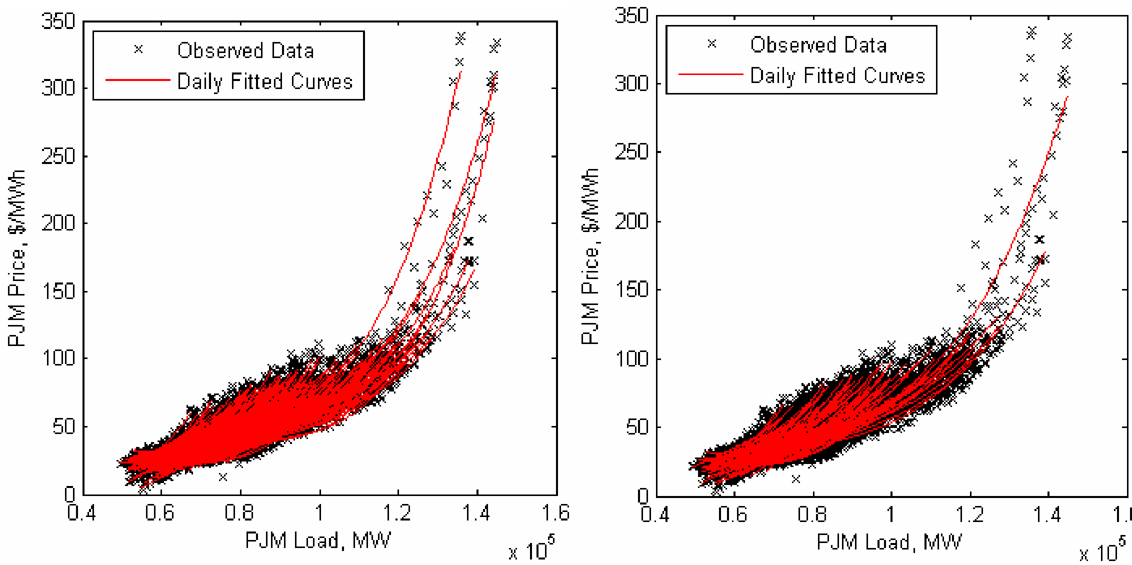

Figure 13. Data plotted with model curves using 3 dummy variables, daily (left) and week/weekend (right).

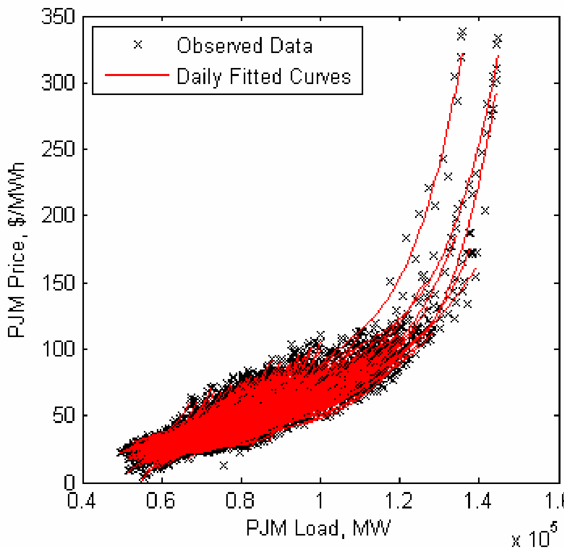

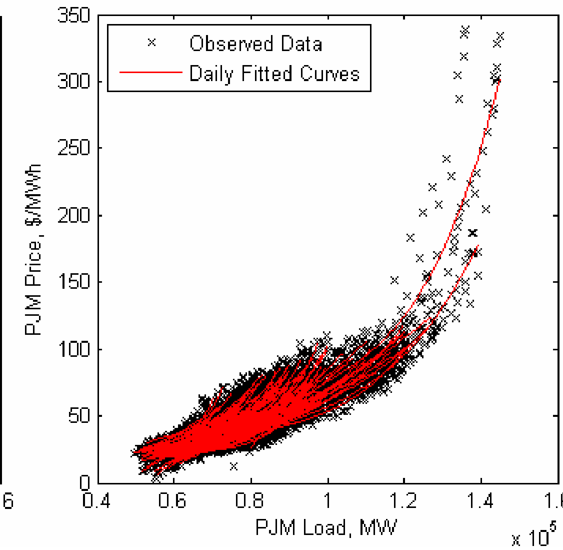

Figure 14. Data plotted with model curves using 4 dummy variables, daily (left) and week/weekend (right).

Each dummy variable adds some small predictive ability to the daily model, especially in the high and low price extremes. Another important issue is whether the model can make small extrapolations outside observed daily loads. In addressing that issue we have looked at data and prediction plots for all 365 days for each set of dummy variables. We plotted along with those curves the most extreme daily demand curves[20] to determine the largest amount of extrapolation required. When we include all 4 dummy variables the price predictions can go off-course with extrapolation, but with fewer dummy variables the extrapolative ability improves. Every single day appears to have acceptable extrapolative ability when using only two dummy variables.

Based on these observations, we conclude that the best overall supply-side model for analyzing RTP effects is the third-degree polynomial model with linear daily offsets as in (7).

[20] From (3) with elasticity -0.4 and minimum or maximum daily load.

Appendix B Split between Supply Curve and Stacked Bid Curves

We have claimed that using generator marginal cost curves to approximate supply curves underestimates both price and slope. In order to support that claim, we have used PJM data on generator bids into the market to construct day-ahead hourly bid curves [1]. Most generators supply one bid curve into the market that will apply for the entire 24 hours, but others self-schedule their generation amounting to an hourly zero-price offset. A small number of hourly increments or decrements are bid at a non-zero price. We have constructed bid curves for every hour of the year from June 1, 2005 to May 31, 2006 by accounting for each of these bid types.

In this Appendix we examine an earlier time period than in the rest of this paper because the generator bid data are released only after a six-month delay and are unavailable at this time[21].

Aggregate bid curves vary little over the course of one day; the maximum total available load offset in the time frame we observed was 5.6% between daily maximum and daily minimum. On the left-hand side of Figure 15 we have plotted the bid curves for noon of every day on the left. The bid curves have the hockey-stick shape typical of system marginal cost curves. As noted earlier it is common in literature to find that the bid curve is assumed to be the true supply curve [10, 16, 21]. When this assumption is made the conclusion is that small changes in load have almost no effect on price except at high loads above the “elbow”.

On the right-hand side of Figure 15 real market clearing results are shown along with the same noon bid curves. In this second plot, we have shown only the section of bid curve corresponding with the actual daily load range. If the bid curve were a good approximation of the system supply curve, then real market results would be close to the bid curve range. This graph makes it clear that bid curves are a poor approximation of overall supply. This is because the real-time constraints on generator dispatch including unit commitment, transmission constraints, and operating reserves are ignored. Real market prices are much higher than would be predicted by these bid curves.

[21] Dominion merged with PJM on May 1, 2005 and increased system peak load increased by 18.6% [20]. The data start date allows us to study one contiguous years’ worth of data without a territory expansion; the end date allows access to the generator bid data on a six-month delay.

Figure 16 shows the same data along with daily fitted supply curves from (7); this model has overall adjusted $R^2 = 0.942$ and an overall model F-statistic of 194. By comparing Figure 16 with Figure 15, it is clear that after accounting for real-time system constraints, the supply curves have a much steeper slope than the bid curves even at moderate and low load. This implies that small changes in load can have large impacts on price that would not be predicted by examining bid curves. Supply curve slope is the most influential factor in determining the impact of a small change in load on price.

These data covering the summer of 2005 have some qualitative differences from the 2006 data examined elsewhere. High price extremes were greater in 2006 because load extremes were also higher. We also see that high prices were observed even on days when load was moderate or low. This is because electric generators faced high natural gas prices in the fall of 2005 [22]. Natural gas generators are more versatile in load-following and are scheduled during a few hours of day even when overall demand is not high.

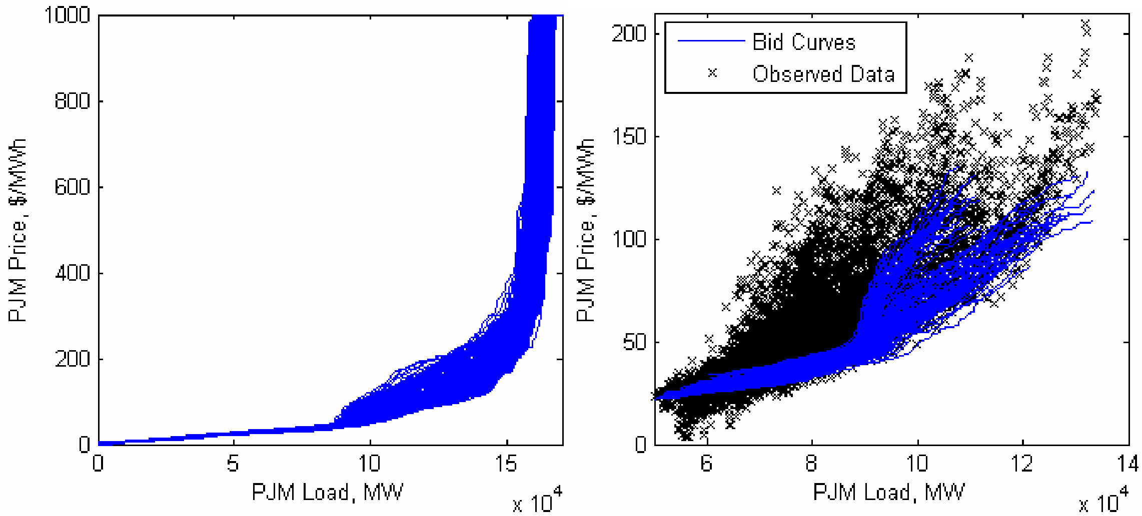

Figure 15. Daily bid curves at noon (left); noon bid curves with observed data (right).

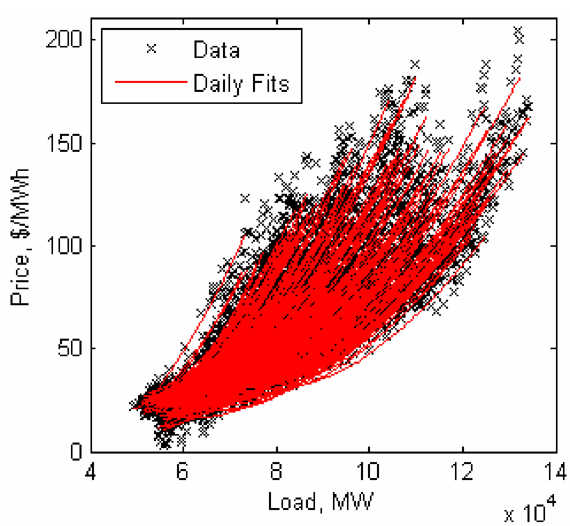

Figure 16. Data plotted with daily fits from (7).

Many analyses of demand response have assumed away the effects of system constraints outright and taken the implications of a constraint-free stacked bid curve to their logical conclusions [7, 10, 16, 21]. By examining aggregate system results shown here, the magnitude of the discrepancy between the bid curves and actual system results becomes clear.

With a fitted supply model and the observed load, we can predict what price would have been in any hour. Table 15 compares the average, minimum, and maximum prices predicted by a bid-curve model, a supply-curve model, and the actual observed prices. Based on this comparison it is clear that the bid-curve model predicts prices that are much too low, although not as low as they ought to be in low-load hours. The comparison implies a $15.88/MWh average premium for system constraints.

Table 15. Prices observed versus predicted by bid curve and supply curve models.

	Bid Curve	Supply Curve	Observed
Minimum Price	$21.87	$10.66	$3.34
Maximum Price	$138.73	$181.64	$204.46
Median Price	$36.97	$54.17	$53.10
Average Price	$48.57	$69.44	$69.44

The comparison of prediction and observed price duration curves in Figure 17 again show that bid curve price predictions are almost always too low. The duration curve predicted by the supply curves is indistinguishable from the curve actually observed.

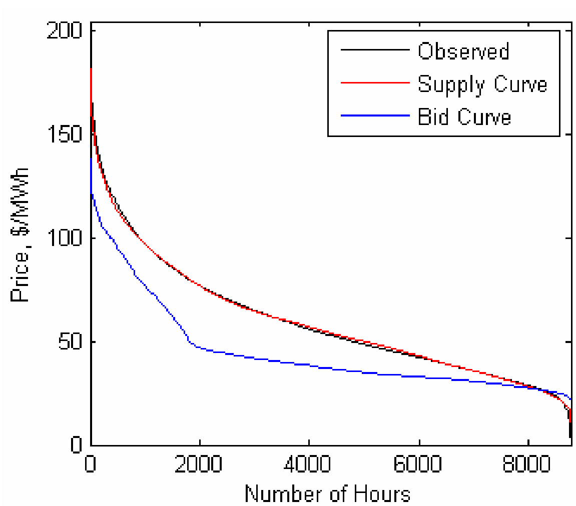

Figure 17. Predicted and observed price duration curves.

By plotting observed prices against predicted prices in Figure 18, a richer comparison of model quality can be made. If either model were perfect, then the scatterplot of predicted price and observed price would fall along the identity line. In order to show how close each model comes to the identity line I have plotted that along with the line outputted from a linear least-squares regression. Bid curve predicted prices are systematically lower than real prices, and are never observed in the high price region. The odd-looking heteroscedasticity in the left graph can be understood by comparing it with Figure 18, but the general conclusion that the bid curve does not accurately represent the characteristics of observed prices and has poor predictability with adjusted $R^2 = 0.673$. The supply curve model best fit line is indistinguishable from the identity line and error appears to be evenly distributed up and down except at the most extreme prices.

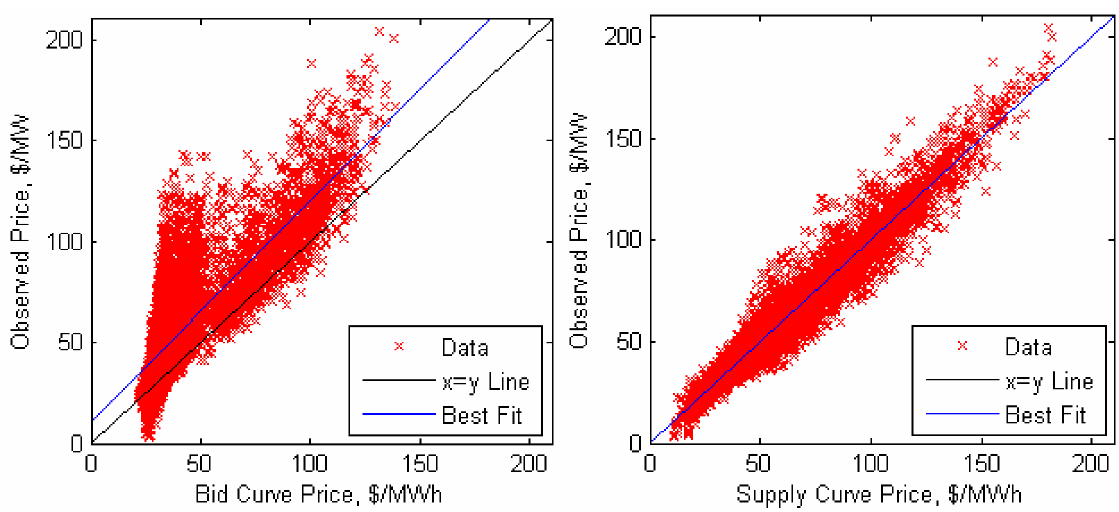

Figure 18. Observed price versus prices predicted by bid curves (left) and supply curves (right).