Customer Risk from Real-Time Retail Electricity Pricing: Bill Volatility and Hedgability

Severin Borenstein*

One of the most critical concerns that customers have voiced in the debate over real-time retail electricity pricing is that they would be exposed to risk from fluctuations in their electricity cost. The concern seems to be that a customer could find itself consuming a large quantity of power on the day that prices skyrocket, resulting in a high monthly bill. I analyze the magnitude of this risk, using demand data from 1142 large industrial customers, and then ask how much of this risk can be eliminated through various straightforward financial instruments. I find that very simple hedging strategies—forward purchase contracts that are already used with many RTP programs—can eliminate more than 80% of the bill volatility that would otherwise occur. I then show that a slightly more sophisticated application of these forward power purchases can significantly enhance their effect on reducing bill volatility.

1. INTRODUCTION

Over the last few years, a great deal has been written about time-varying retail pricing of electricity. Many authors, myself included, have argued that real-time retail electricity pricing (RTP)—retail prices that change very frequently, *e.g.*, hourly, to reflect changes in the market's supply/demand balance—is a critical component of an efficient restructured electricity market. During the California electricity crisis in 2000-2001, RTP boosters pointed out the value of RTP in reducing the ability of sellers to exercise market power. While nearly all economists have

supported RTP, large-scale adoption has been slowed by at least three factors: cost of real-time metering and billing, large potential wealth transfers among customers due to RTP, and the potential volatility of bills that customers could face.

The metering and billing costs, while possibly a serious issue for smaller customers, are minor compared to the potential benefits of moving larger customers to RTP, as I have shown in Borenstein (2005). [1] In Borenstein (2006), I addressed the potential level of transfers among larger customers due to RTP. I found significant potential transfers and that more than half of all customers could be worse off, even compared with time-of-use pricing (a simple peak/off-peak pricing system). The transfers, however, can be greatly mitigated through a two-part RTP program that gives incumbent customers the right to buy their past load levels at regulated prices, but still charges/refunds at the real-time energy price for their deviations from that level.

The third area of concern one frequently hears about RTP is the possible volatility of costs to the end-use customer. This is often expressed as concern about the cost the customer would face for electricity consumed during an hour in which prices hit an extreme spike, such as $10,000/MWh. Two factors should reduce these concerns. First, customers almost certainly care about volatility of their payments, which are typically on a monthly or longer cycle, not the volatility of their hourly incurred liability. Second, in most RTP implementations, there is some opportunity for the customer to buy fixed-quantity, fixed-price contracts in advance to cover some of their demand, and then pay/receive the real-time price for deviations from the contracted quantity. In this paper, I address the question of how much bill volatility is caused by RTP and the extent to which hedging in advance, by purchasing forward contracts, might reduce that volatility.[2]

The concern about volatility is somewhat puzzling to hear from large corporations for whom electricity makes up a very small share of expenses. One might think that this sort of risk could be easily absorbed by such customers. There are a number of possible explanations, though none is entirely satisfying. One is that electricity may be only one or two percent of total expenses, but could still be equal to the entire profit margin of the firm, so the proportional effect on profits could be much more significant. Pressure from equity and debt holders could cause the firm to be squeamish about even short-term profit fluctuations.[3] A second explanation is organizational. For whomever within the firm is responsible for energy costs, this could be a large component of his or her responsibilities, so a sudden dramatic increase in costs could have significant career implications.

---

1. As technology improves over the next decade, the billing/metering cost arguments against RTP, even at the household level, will likely disappear.

2. Barbose, Goldman and Neenen (2006) give an overview of RTP programs currently in use and the hedging options available to customers.

3. This is part of the larger issue of why publicly traded firms worry about such idiosyncratic risk at all if their shareholders are holding diversified portfolios. I don't address this puzzle other than to note that companies often take costly actions to insure against risks that are much smaller than the fluctuations in electricity bills that are conceivable. Brown and Toft (2002) provide a number of citations to work on the reasons for corporate hedging.

A third explanation may be simply a cost/benefit analysis in which the customer sees tangible costs from increased complexity and risk, but doesn't understand the potential benefits.[4] In any case, bill risk is nearly always raised by opponents of RTP among even large customers, so it seems worth trying to estimate the size of that risk and consider how it might be mitigated.[5]

The issue of electricity bill volatility from RTP should also be considered in the context of the volatility that customers already face due to electricity consumption variation under historical pricing schemes. Even with a flat-rate tariff, a customer's bill will vary due to consumption variation. Under a time-of-use (TOU) tariff—a simple peak/off-peak pricing scheme—there will be more bill volatility for most customers due to higher prices in the peak hours of the peak months. From these comparison points, I examine how much RTP further increases bill volatility and how much that increase is mitigated by hedging.

The data I use for this analysis cover 1142 large industrial and commercial customers of Pacific Gas & Electric during 2000 through 2003.[6] I combine these customer consumption patterns with actual and simulated prices (based on data from the California Independent System Operator), as described in the next section, to determine the customers' bills under flat-rate, time-of-use, and real-time retail pricing programs, where retail rates are set in all cases to cover the full wholesale cost of the power. From these monthly bills, I calculate measures of monthly bill volatility for each customer.

Then I recalculate the customer bills and bill volatility under the assumption that they have engaged in actuarially-neutral hedging contracts.[7] By construction, the aggregate of all customer bills over the entire sample period is the same in all of these scenarios, but the monthly bill volatility of individual customers varies substantially. I show that while changing from flat-rate or TOU to RTP tar-

4. I have heard all three of these explanations from representatives of customers who are opposed to, or at least skeptical of, attempts to implement RTP.

5. A somewhat different concern about analyzing electricity bill volatility of large corporations is that the bill is determined in part by quantity, which is a choice variable of the firm. While that is correct, for many customers, variation in quantity is determined largely by weather factors that alter the amount of electricity necessary to maintain a given level of operations and comfort. Analogously, the quantity of autos one purchases is a choice, yet most people still insure against the need to purchase a new car because theirs has been stolen or destroyed in an accident. In the case of electricity, a customer's high weather-driven demand is also likely to be positively correlated with high electricity prices. Furthermore, one concern voiced by many industrial customers is that they will fail to monitor electricity prices hourly under RTP and will accidentally engage in high-consumption activities at a time when the price has spiked.

6. These are all customers for whom PG&E submitted data for all 1461 days of this time period to the California Energy Commission. The data are not publicly available, but were made available to me under a nondisclosure agreement.

7. One can make plausible arguments that buyers would have to pay a premium or would be able to obtain a discount on average when they lock in purchase prices and quantities on a forward basis. While a positive or negative forward premium would affect the customer's cost of hedging, it would not affect the benefit that the customer would receive in decreased bill volatility, which is the subject I examine here. Siegel and Siegel (1988) present a general discussion of the effect of futures purchases in reducing cost volatility. O'Sheasy (2002) presents an example specific to electricity and RTP.

iffs increases customer average bill volatility by two to four times, simple hedging strategies eliminate the vast majority of that volatility. I conclude that bill volatility should not be a significant impediment to implementation of a well-designed RTP program, but that hedging instruments are likely to be important to building customer acceptance.

I then explore the levels of hedging in which a customer could engage. A simple example illustrates the intuition that if price is positively correlated with a customer's consumption quantity, hedging more than expected quantity can reduce bill volatility by more than hedging just the expected quantity. In fact, I show that such "over-hedging"' can reduce volatility to below the level faced by a customer subject to a flat-rate tariff. I evaluate the empirical importance of this result by examining the effect of various levels of over-hedging on the bill volatility of the observed customers.

2. CONSUMPTION AND PRICING DATA

The analysis is based on a sample of 1142 large industrial and commercial customers of Pacific Gas & Electric company, which has a service territory that covers most of northern California, during the period 2000-2003. All of the customers are in the sample for the entire period. With these data, it is straightforward to construct a customer's monthly bills by combining the quantity data with any given price series for the same time period.

By constructing the customer bill this way, I am assuming zero price-elasticity of the customer's demand. Customer price elasticity would reduce the bill volatility induced by RTP, because customers would respond to the higher-priced hours by consuming less. Thus, assuming no such elasticity biases the calculation in the direction of finding a larger impact of RTP on bill volatility.

I carried out the calculations using three different wholesale price scenarios. The first is the actual wholesale spot prices that were observed in the control area of the California Independent System Operator (CAISO) for northern California during the observed time period. The other two are based on a long-run breakeven wholesale price simulation that uses the CAISO's actual load during the sample period. In all three scenarios, I assume that the retail real-time prices customers face are equal to the wholesale prices plus $40/MWh for transmission and distribution (T&D).[8]

The simulation model, which is described in detail in Borenstein (2005), establishes a long-run perfectly competitive equilibrium in capacity and wholesale prices for a given demand profile (load duration curve), assumed demand elasticity, and costs of different types of production capacity. The data used for

---

8. This is a fairly typical assumption for T&D costs. In a more sophisticated implementation of RTP, the T&D charge would also be subject to real-time variation that reflects congestion, but that is not part of most current or planned RTP programs. Most utilities impose a charge related to distribution capacity, called a demand charge, that is based on the customer's peak usage during a billing period regardless of when that peak usage occurs.

generating the wholesale price series for this paper are not exactly the same as in Borenstein (2005). First, I use different cost data than those in the earlier paper, reflecting changes in capital and fuel costs since that paper was written.[9] Second, I use only demand data from the 4-year period 2000-2003. By limiting the time period of simulation to just that period, I can impose that the resulting prices are sufficient in aggregate to cover the amortized capital and variable costs of all generators during the sample time period.

Absent large elasticity of aggregate electricity demand, much of the capital costs are recovered in peak hours, though exactly how many hours and how peaky the prices are depends on the exact elasticity of aggregate demand. I create two wholesale price series with differing elasticities of aggregate demand and different resulting peakiness of prices.

By some measures the simulated wholesale prices are spikier than actually occurred during most of 2000-2003 period. Due in part to price caps that were in place, the highest actual prices are lower than the highest prices in either simulated scenario, but the simulated prices exhibit fewer total hours with price above $200/MWh. The actual prices average slightly higher than the simulated prices over the entire period, but this masks a significant change in mid-2001. Prior to June 2001, actual prices are considerably higher than simulated, but after June 2001 that reverses. The earlier period was the California electricity crisis. The latter was a period in which the state was widely viewed as having significant excess capacity, so prices remained low. These wholesale prices generally were viewed as too low to support generation investment.[10]

The simulation model assures that generators cover their variable plus amortized fixed costs during the sample period. The simulation is of an energy-only revenue model; there are no separate capacity payments. Thus, fixed cost recovery occurs during the highest-demand hours when price exceeds the variable cost of even the most costly generation units.

The two simulated scenarios differ in the degree of demand elasticity that within-market producers are assumed to face. Demand elasticity may come about from actual end-user adjustments, but it can also come from import supply elasticity or the system operator utilizing out-of-market resources to provide supply if the market prices rises high enough. With extremely inelastic demand, the simulated market equilibrium includes a very small number of hours in which prices are extremely high. These hours produce the net revenues (scarcity rents) necessary for peaker generation units to cover their amortized fixed costs. With

---

9. The assumptions I use here for annual production cost are: Baseload (coal) Cost =$208247/MW + $25/MWh; Mid-merit (CCGT) Cost=$93549/MW +$50/MWh; and Peaker (Combustion Turbine) Cost =$72207/MW + $75/MWh. These figures are taken from the PJM (2005), pages 82-83. California does not have coal plants, but (a) there are coal plants in the western grid and (b) the results are not affected substantially by fixing the level of baseload capacity in advance to reflect nuclear and other must-take capacity.

10. The actual prices also include hours in which the real-time wholesale price was negative, which can occur due to non-convexities in the production process and the fact that electricity generated cannot always be disposed of costlessly.

somewhat greater demand elasticity, the long-run equilibrium involves the peaker generators collecting scarcity rents over more hours, but a lower level of scarcity rents and a lower wholesale price in any one of those hours. In scenario I, I assume that the demand elasticity faced by within-market producers is -0.025. In scenario II, I assume an elasticity of -0.1.[11] Summary statistics for the three wholesale price scenarios are presented in table 1.[12]

I focus primarily on the analysis of results from scenario I simulated prices, because this is the sort of scenario under which bill volatility from RTP would be of the greatest concern. Bills are less volatile under scenario II due to the higher assumed demand elasticity. Bills are more volatile under scenario III, the actual prices in Northern California over 2000-2003. That volatility, however, was due to a combination of extreme market power, sparse forward contracting at the wholesale level, and almost no demand response at the retail level that led to the 2000-01 California electricity crisis, a combination that is unlikely to ever be repeated. I include results for scenarios II and III, for which the qualitative conclusions are the same, in the appendix.

With these prices and customer consumption quantities, I can then calculate each customer's monthly bill under RTP for each month it is in the sample. I then divide each bill by the number of hours in the month in order to eliminate variation due to varying lengths of months. Henceforth, I refer to this average hourly bill during each month as the customer's "monthly bill."[13]

In order to compare bill volatility under RTP with the alternatives, I also need to create a flat-rate tariff and a time-of-use (TOU) tariff to use in calculating the monthly bills for the observed customers. I do this by continuing to assume no demand elasticity on the part of the observed customers and calculating the prices that would cover the costs of the observed group of customers as if they were a distinct tariff class.

For the flat-rate tariff, this is completely straightforward: I calculate the total revenue that would be required to purchase the consumption of the observed customers in the wholesale market over the 48-month period (using each of the wholesale price series discussed above) and then divided that revenue requirement by the aggregate consumption of these customers over the 48-months. This yielded a flat-rate tariff that generated sufficient revenues to purchase the power

11. In both simulations, I assume zero cross-hour price elasticity. Incorporating such cross-elasticity would likely dampen real-time price volatility and the increase in bill volatility due to RTP. See Borenstein (2005) for a discussion of this assumption.

12. The average price under scenario II is somewhat lower than under scenario I, because greater price elasticity leads to lower capacity investment and higher capacity utilization in equilibrium, as discussed in Borenstein and Holland (2005).

13. I drop nearly all hours in which the consumption is reported as exactly zero. For those in which the adjacent hours on each side are present (a single missing hour) and are of typical level for the customer, I interpolate the missing data. For many hours or days in a row of zero readings, I set these observations to missing and calculation the "monthly bill" based only on the remaining hours. These are either meter failures or complete power shutoffs, which could be due to on-site generation, plant shut-down, or other causes. Including these zeros—which are not common in the sample—makes very little quantitative difference to the results and does not change any of the qualitative findings.

Table 1. Wholesale Prices in Alternative Scenarios (all prices in $/MWh)

Time Period: 2000-2003 (35064 total hours)

		Scenario I Very Volatile Simulated Prices	Scenario II Less Volatile Simulated Prices	Scenario III Actual No. Cal. Spot Prices
Flat-Rate Tariff		93.50	93.41	103.54
Fixed-Ratio TOU Tariff -- Maintaining Actual Price Ratios Among TOU Periods				
Winter	Off-Peak	79.45	79.38	87.98
Winter	Peak	98.42	98.32	108.98
Summer	Off-Peak	79.65	79.56	88.19
Summer	Shoulder	92.96	92.87	102.94
Summer	Peak	151.71	151.56	168.00
Cost-Based TOU Tariff -- Breakeven within Each TOU Period				
Winter	Off-Peak	68.80	68.80	101.16
Winter	Peak	88.57	88.56	110.21
Summer	Off-Peak	74.41	74.90	91.82
Summer	Shoulder	97.07	104.38	107.58
Summer	Peak	203.52	192.50	120.31
Real-time Pricing Tariff				
Minimum Price		65.00	65.00	-285.61
Median Price		90.00	89.13	76.59
Mean Price		88.77	88.77	100.95
Maximum Price		6,321.66	1,051.08	790.00
No. of Hours Price>200		287	918	2,725
No. of Hours Price is Above Highest Simulated Generation Marginal Cost		383	1,713	N/A

demanded by this group.

To create a TOU tariff requires first that one determine the different rate periods. I use nearly the same rate periods that PG&E used in its standard TOU tariffs during the observed time period. There are two periods during the "winter" months, in effect November through April: the peak rate is in effect from 8am to 10pm on non-holiday weekdays and the off-peak rate is in effect at all other times. Summer rates, which cover May through October, have three components: Peak period is noon-6pm on non-holiday weekdays; Shoulder period is 8am-noon and 6pm-10pm on non-holiday weekdays, and off-peak is all other times.[14]

Creation of a TOU tariff is less straightforward than the flat-rate because for a given demand history there are an infinite number of tariffs that would cover the total wholesale power costs. A natural TOU tariff would be for the rate within each of the five TOU periods to be set to exactly cover the wholesale power procurement costs for that period, *i.e.,* no cross-subsidy across TOU periods. I use that

14. In the four-year period I study, the number of hours each rate is in effect are: winter off-peak, 10,512 hours; winter peak, 6,888 hours; summer off-peak, 10,440 hours; summer shoulder, 4128 hours; summer peak, 3096 hours. The actual tariff rate changes at 8:30am, but I have altered that because the data are aggregated to the hour level. Similarly, the evening change occurs at 9:30pm, but I've altered that to 10pm.

as one basis of analysis, but the TOU tariff that results from this calculation exhibits much larger peak to off-peak rate differences than have historically been utilized by PG&E or most other regulated utilities. After adding a $40/MWh charge for transmission and distribution, the ratio of the highest TOU-period prices to the lowest would be 3.11 (using the scenario I real-time prices), while the same figure for PG&E's standard TOU rate is 1.91. Thus, this constructed TOU tariff, which I refer to as "cost-based TOU," is economically appealing, but it would yield more volatile bills than are likely to occur under actual TOU tariffs.

To address this, I also create a TOU tariff that mimics the inter-period rate ratios that existed under the standard TOU tariff that PG&E offered during the sample period. To construct this tariff, I calculated the ratios between the rates in the PG&E tariff for the five periods. I then adjusted all of the TOU rates together, maintaining these ratios, until the resulting revenue covered the wholesale procurement costs of all power consumed by the observed customers in aggregate. I refer to this as "fixed-ratio TOU." The resulting flat-rate and TOU tariffs, along with summary statistics of the RTP prices, are shown in table 1.

By assuming throughout these rate calculations that all of the observed customers exhibit zero price elasticity, I am implicitly assuming that moving these customers to RTP does not further dampen wholesale price volatility. That price dampening effect must instead be recognized in evaluating the results of the exercise. It suggests that the extremely volatile price scenario, scenario I, is less likely to occur once some customers are on RTP.

Finally, comparing a customer's bill volatility under different pricing and hedging regimes requires a measure of volatility. Because the customer's concern is with risk, I assume that predictable monthly variation is not at issue. Thus, I measure the "unexpected deviation" component of the bill as the difference between the actual monthly bill and the customer's average bill for that month of the year. I use two measures of bill risk. The first I refer to as a "coefficient of deviation", which is the standard deviation of the customer's monthly unexpected bill deviation divided by the mean of the customer's monthly bill.[15] The second is an attempt to capture some customers' concerns about extreme events, particularly bill spikes when the wholesale price increases drastically for a short period. The second measure, which I refer to as the "coefficient of maximum deviation" is the customer's maximum monthly unexpected bill deviation over the sample period divided by its mean monthly bill.[16]

---

15. This is nearly the coefficient of variation of monthly bills, but not quite, because I am using the deviation from the monthly *seasonal-adjusted* average bill.

16. Using a monthly seasonal adjustment with only four years of data suggests that a degrees of freedom correction should be applied to all estimates of standard deviations, multiplying the sample statistic by 4/3 to correct for the fact that the mean is calculated from the same data. Such a correction would, of course, make no difference to the comparisons of volatility across pricing regimes. A similar correction to the coefficients of maximum deviation, however, is less clear since these are essentially order statistics. To maintain some comparability between the two measures, I have not performed a degrees of freedom correction on either.

3. BILL VOLATILITY WITH NO HEDGING

Table 2 presents the measures of bill volatility under the four tariff re-gimes—flat-rate tariff, cost-based TOU, fixed-ratio TOU and RTP—for whole-sale price scenario I, the most volatile wholesale prices. Results for the two other wholesale price scenarios are presented in the Appendix. As noted, one would expect to observe bill volatility even absent time-varying prices due to consump-tion variation. This is reflected in the variation under a flat-rate tariff. Indeed, the left-hand column indicates that under the flat-rate tariff the average customer coefficient of deviation is 0.155. The volatility is almost exactly the same under either TOU pricing plan. While prices are more volatile under these plans, the dif-ferences are captured in the predictable seasonal variation. This is not surprising in that the prices are, by construction, identical across years for a given month. Because the coefficient of deviation is bounded below at zero, the distribution is highly skewed. The median variation is about half the mean.

Bill volatility under RTP is, however, many times greater than under the other billing arrangements. The mean volatility is more than double and the me-dian is about four times greater than under the alternatives.

The table presents the statistics on the full distribution of customer bill volatilities, indicating that there are some customers with extremely high volatil-ity. The fact that this results even with a flat-rate tariff indicates that this bill-volatility is a result of large fluctuations in the amount of electricity the customer uses. In fact, it appears that the more volatile bills are due more to consumption volatility than price volatility. Moving up the distribution from lower to higher bill volatility, it is clear that the ratio of volatility under RTP to volatility under the conventional billing arrangements declines. In other words, RTP exacerbates volatility for the bills that have low and medium volatility under conventional billing, but has a relatively smaller effect on volatility for the bills that are most volatile under conventional billing.

Table 2. Monthly Bill Volatility Under Alternative Billing Arrangements

(1142 Customers in Sample)

Coefficient of Deviation in Monthly Bill

	Mean	1st	10th	50th	90th	99th
			PERCENTILES			
Tariff						
Flat-Rate	0.155	0.021	0.035	0.087	0.364	1.048
TOU - Fixed-ratio	0.156	0.021	0.037	0.086	0.365	1.057
TOU - Cost-based	0.159	0.022	0.039	0.088	0.371	1.083
RTP	0.419	0.187	0.264	0.367	0.602	1.272

Same-customer Comparison of Coefficient of Deviation Under RTP and TOU-F

	Mean	1st	10th	50th	90th	99th
			PERCENTILES			
Tariff						
cd (RTP) / cd (TOU-F)	4.737	0.963	1.393	4.300	8.854	12.843

The bottom panel of table 2 presents a by-customer analysis, showing the ratio of bill volatility (as measured by the coefficient of deviation) under RTP to volatility under fixed-ratio TOU (TOU-F). Consistent with the first panel, the median customer sees its bill volatility increase by more than four times. For more than ten percent of the sample, bill volatility is at least eight times higher under RTP than under TOU-F.

While the coefficient of deviation is a good measure of overall volatility, what many customers seem to be concerned about is the outlier event, when they are hit with a huge unanticipated bill spike. I attempt to capture than in table 3. Table 3 has the same structure as table 2 except it presents the coefficient of maximum deviation, the ratio of a customer's maximum bill deviation (from it's seasonally-adjusted expectation) in the sample to its average bill.

Table 3 confirms that there is very little variation in unanticipated bill costs among flat rates, fixed-ratio TOU and cost-based TOU, but RTP opens up much more possibility for bill shock from a single month. In this sample of 1142 customers, the coefficient of maximum deviation averages between 30% and 35% under flat rate or TOU. But the coefficient of maximum deviation under RTP is on average more than 150%, suggesting that a typical customer would at some time during this sample receive a bill that is more than two and a half times the expected level. For more than 10% of the sample the largest unexpected bill differential would cause more than a tripling of the bill.

These analyses confirm that real-time pricing *without hedging* could significantly increase customer bill volatility, at least if wholesale prices are quite volatile. As mentioned earlier, this wholesale price volatility is probably greater than would obtain if a significant number of customers were on RTP. Still, the volatility increase could very well be sufficient to cause substantial unease among some customers.

Table 3. Coefficient of Maximum Deviation Under Alternative Billing Arrangements

(1142 Customers in Sample)

Coefficient of Deviation in Monthly Bill

	Mean	1st	10th	50th	90th	99th
			PERCENTILES			
Tariff						
Flat-Rate	0.318	0.042	0.072	0.182	0.658	2.359
TOU - Fixed-ratio	0.324	0.043	0.078	0.186	0.668	2.392
TOU - Cost-based	0.346	0.048	0.088	0.199	0.690	2.455
RTP	1.595	0.592	0.994	1.417	2.294	4.301

Same-customer Comparison of Coefficient of Deviation Under RTP and TOU-F

	Mean	1st	10th	50th	90th	99th
			PERCENTILES			
Tariff						
cx (RTP) / cx (TOU-F)	8.869	0.909	2.551	8.050	16.603	25.545

4. BILL VOLATILITY WITH SIMPLE HEDGING

Because of the potential for large unexpected bills under RTP, those utilities that offer or require an RTP plan for some customers usually also offer a hedging option. The hedging program allows the customer to purchase some of its power months in advance at a less-volatile expected price. These programs are often used as inducements by setting the hedge price below the expected spot price, thereby combining a hedging program with a transfer, which is often done to compensate the customer for the loss of a cross-subsidy it was receiving under conventional billing.[17] The analysis here does not encompass this aspect of long-term power purchases; instead I set hedging prices to be actuarially fair. I implement this not on a forward-looking probabilistic basis, but simply by setting the price of the hedge equal to the unweighted average price during the hours that it covers in the sample.

The hedging products that I examine here are modeled after structures in some of the RTP hedge programs: different products for peak and off-peak periods. For simplicity, I model five hedge products that correspond to the five TOU periods over the year: a peak and an off-peak product during the winter, and a peak, shoulder and off-peak product during the summer. The product itself is defined as purchase of one unit of power during each of the hours in the sample that falls under the specified TOU period. The price of the hedge, per MWh, is the unweighted average price of power during the hours covered. To prevent cherry-picking of the hours with the highest expected prices within a TOU period, a customer cannot buy different quantities for different hours within the same TOU period. The prices for the hedges under price scenario I are: Summer-Peak (3,096 hours): $198.69, Summer-Shoulder (4,128 hours): $96.70, Summer-Off-peak (10,440 hours): $73.80, Winter-Peak (6,888 hours): $88.36, Winter-Offpeak (10,512 hours): $68.43. A comparison of these prices with table 1 reveals that they are closer to cost-based TOU than to fixed-ratio TOU, but they are a bit more compressed than TOU-C. The reason is that the hedge prices are unweighted averages of the wholesale prices while the TOU-C prices are weighted averages in order to recover the full wholesale cost. These are still breakeven hedge prices, because the hedge product is for the same quantity (one MWh) in all hours within the period.

These are obviously not very sophisticated hedge products. Hedges could potentially be offered and priced for even individual hours at some expected cost for that specific hour. Without resorting to a forecasting model, I can't mimic that process. Instead, I'm using a broad hedging product based on the average prices

---

17. Such advance purchases of a set quantity are usually referred to as customer baseline (CBL) programs and the quantity that the customer purchases in advance is its CBL quantity. Unlike a voluntary hedging program, the CBL quantity is not a choice of the customer, but is usually based on the customer's past usage. Braithwait and O'Sheasy (2002) and O'Sheasy (2002) discuss conceptually the effect of CBL programs on bill volatility. See Borenstein (2006) for further discussion of the use of CBL quantity purchases to compensate for loss of a cross-subsidy.

over a long time span, the four-year sample period. Because the wholesale price simulations generate different prices only due to demand variation, this approach approximates the distribution of prices a customer might anticipate in the spot market going into a period, before information about abnormal demand (or supply) is revealed.

To examine the magnitude of bill volatility under RTP with hedging, I also need to determine the quantity of power purchased as part of the hedge. Here again I take a very simple approach. The quantity that a customer purchases is assumed to be the customer's naive expected consumption in all hours to which the hedge product applies. For example, for all hours that belong to the Summer-Peak TOU period, the customer purchases the same hedge quantity and that quantity is the customer's average consumption during Summer-Peak TOU hours during the four-year sample period. In an actual implementation, the customer could almost certainly do better than this in predicting it's consumption for, for instance, a given summer period. The customer might know if it is ramping up its own production or closing most of the operations for some period of time. Thus, this approach is likely to underestimate the risk reduction from hedging.

A customer on RTP who hedges is assumed to have a power cost for hour t that is in TOU period h equal to $Cost_t = \hat{q}_h P_h + (q_t - \hat{q}_h) P_{RTPt}$, where $\hat{q}_h$ is the customer's average consumption during hours in that TOU period and P_h is the hedge price of a MWh of power during that TOU period.

From these hourly power costs, the customer's monthly bill (actually hourly average bill during the month, as explained earlier) is created for each month and, just as in the previous section, the deviation from expected cost, conditional on month of the year, is then calculated for each month.

The results are shown in table 4, which repeats table 2, but with a line for RTP-Hedged. The effect of hedging is quite strong: On average, this simple hedging plan removes about 90% of the unanticipated bill volatility that would otherwise occur as a result of changing from TOU-F to RTP. Table 5 repeats table 3, but with a line for RTP-Hedged. Measuring now the coefficient of maximum deviation, hedging again removes most of the volatility due to RTP. On average RTP-H eliminates more than 80% of the additional volatility that would otherwise occur under RTP. These effects are also present with pricing scenarios II and III, the results for which are shown in the appendix. Under all three pricing scenarios, hedging eliminates the vast majority of volatility that would otherwise result from implementation of RTP.

5. BILL VOLATILITY WITH VARIABLE HEDGING

In the previous section, I assumed a naive hedging strategy and demonstrated that even such a simple approach could eliminate a great deal of the bill volatility that would be associated with real-time pricing. More granular hedging products and precise demand forecasts would further reduce the unanticipated

Table 4. Monthly Bill Volatility Under Alternative Billing Arrangements

(1142 Customers in Sample)

Coefficient of Deviation in Monthly Bill

			PERCENTILES			
	Mean	1st	10th	50th	90th	99th
Tariff						
Flat-Rate	0.155	0.021	0.035	0.087	0.364	.048
TOU - Fixed-ratio	0.156	0.021	0.037	0.086	0.365	1.057
TOU - Cost-based	0.159	0.022	0.039	0.088	0.371	1.083
RTP - Hedged	0.187	0.028	0.048	0.114	0.416	1.120
RTP	0.419	0.187	0.264	0.367	0.602	1.272

Same-customer Comparison of Coefficient of Deviation Under RTP and TOU-F

			PERCENTILES			
	Mean	1st	10th	50th	90th	99th
Tariff						
cd (RTP-H) / cd (TOU-F)	1.308	0.900	0.982	1.230	1.732	2.420
cd (RTP) / cd (TOU-F)	4.737	0.963	1.393	4.300	8.854	12.843

Table 5. Coefficient of Maximum Deviation Under Alternative Billing Arrangements

(1142 Customers in Sample)

Coefficient of Deviation in Monthly Bill

			PERCENTILES			
	Mean	1st	10th	50th	90th	99th
Tariff						
Flat-Rate	0.318	0.042	0.072	0.182	0.658	2.359
TOU - Fixed-ratio	0.324	0.043	0.078	0.186	0.668	2.392
TOU - Cost-based	0.349	0.049	0.089	0.202	0.694	2.483
RTP - Hedged	0.559	0.063	0.134	0.351	1.121	3.444
RTP	1.595	0.592	0.994	1.417	2.294	4.301

Same-customer Comparison of Coefficient of Deviation Under RTP and TOU-F

			PERCENTILES			
	Mean	1st	10th	50th	90th	99th
Tariff						
cx (RTP-H) / cx (TOU-F)	1.961	0.817	0.975	1.870	2.977	4.563
cx (RTP) / cx (TOU-F)	8.869	0.909	2.551	8.050	16.603	25.545

volatility.[18] If the customer's goal is to minimize its bill volatility, however, there is additional value in considering the correlation between its own unanticipated high demand and high electricity prices.

18. I have explored the effect of using more granular hedging instruments by assuming that a customer can buy a different hedging product for each month-hour-weekday/weekend. More precise hedging of this sort does further reduce the measures of volatility, but the effect is relatively small.

There is an intuition, often expressed by large electricity customers, that a customer is at increased risk of high bills if its demand is positively correlated with price. This is more than a theoretical concern for a customer that, for instance, runs its air conditioners hardest on the hottest days of the year when prices are also likely to be highest.

A simple example demonstrates both that this intuition is correct and that this correlation can make hedging especially valuable. Consider the examples in figure 1. A customer consumes a quantity of either 4 or 6 with equal probability. The market price is either 8 or 12, also with equal probability. In example A, the customer's consumption is uncorrelated with the system price. A bit of arithmetic reveals that the customer's expected bill is 50 with a standard deviation of 14.28. An actuarially fair hedge contract in this market would sell for 10, the expected price. Assume that the customer purchases a $q^=5$ unit hedge, equal to its expected demand, for a price of $p=10$ before either its demand or the market price is revealed. If its demand turns out to be 6, it will buy one additional unit on the spot market, and if its demand turns out to be 4, it will sell the extra unit on the spot market. A bit more arithmetic shows that the expected bill is still 50, but now with a standard deviation of 10.20. In fact, it is straightforward to show that a hedge of $q^=5$ minimizes the standard deviation of the distribution of possible bills.

Now consider example B in which the prices and customer's quantity each have the same distribution as in A, but are now perfectly correlated with one another. Either the customer consumes 4 when the price is 8 or it uses 6 when the price is 12. The customer's expected bill is now 52 with a standard deviation of 20. If the customer buys a hedge of $q^=5$, still at the actuarially fair price of 10, it reduces it's standard deviation to 10, less than the standard deviation resulting from the same hedge position when price and quantity are uncorrelated. Now, however, consider if the customer took an even larger hedge position, $q^=10$. This hedge would cost the customer 100. On low price/quantity days, it would sell back 6 units at a price of 8. On high price days, it would sell back 4 units at a price of 12. With probability 0.5, its cost of power would be $100-6x8 = 52$ and with probability 0.5, its cost of power would be $100-4x12 = 52$. That is, over-hedging would not only compensate for the positive price/quantity correlation, but would take advantage of it to reduce bill volatility to zero.

The result is surprising (to most people, at least), but in retrospect it is just the mirror image of a familiar result. Consider the case of perfect negative correlation, shown in example C. In this case, most people are not surprised that complete *under-hedging*, $q^=0$, purchasing all quantity on the spot market, results in zero bill volatility. Increasing the hedge position above zero results in greater bill volatility. McKinnon (1967) develops optimal hedging for a risk-averse farmer trying to stabilize income when price and his crop yield are negatively correlated. He shows that "the more highly negatively correlated are price and output the smaller will be the optimal forward sale." More generally, he derives the optimal hedge position as

Figure 1. Example of the Effect of Over-Hedging

Market Price = $\begin{cases} prob = 0.5 & P=12 \\ prob = 0.5 & P=8 \end{cases}$ Customer Quantity = $\begin{cases} prob = 0.5 & q=4 \\ prob = 0.5 & q=6 \end{cases}$

Example A: Zero Correlation Between Market Price and Customer Quantity

No Hedging : $\begin{cases} prob = 0.25 & q=4, P=8, \quad \text{Bill}=32 \\ prob = 0.25 & q=4, P=12, \quad \text{Bill}=48 \\ prob = 0.25 & q=6, P=8, \quad \text{Bill}=48 \\ prob = 0.25 & q=6, P=12, \quad \text{Bill}=72 \end{cases}$

Average Bill = 50   Standard Deviation = 14.28

Hedging $\hat{q} = 5$ at $\hat{p} = 10$: $\begin{cases} prob = 0.25 & q=4, P=8, \quad \text{Bill} = 5 \cdot 10 - 1 \cdot 8 = 42 \\ prob = 0.25 & q=4, P=12, \quad \text{Bill} = 5 \cdot 10 - 1 \cdot 12 = 38 \\ prob = 0.25 & q=6, P=8, \quad \text{Bill} = 5 \cdot 10 + 1 \cdot 8 = 58 \\ prob = 0.25 & q=6, P=12, \quad \text{Bill} = 5 \cdot 10 + 1 \cdot 12 = 62 \end{cases}$

Average Bill = 50   Standard Deviation = 10.20

Example B: Perfect Positive Correlation Between Market Price and Customer Quantity

No Hedging : $\begin{cases} prob = 0.5 & q=4, P=8, \quad \text{Bill}=32 \\ prob = 0.5 & q=6, P=12, \quad \text{Bill}=72 \end{cases}$

Average Bill = 52   Standard Deviation = 20

Hedging $\hat{q} = 5$ at $\hat{p} = 10$: $\begin{cases} prob = 0.5 & q=4, P=8, \quad \text{Bill} = 5 \cdot 10 - 1 \cdot 8 = 42 \\ prob = 0.5 & q=6, P=12, \quad \text{Bill} = 5 \cdot 10 + 1 \cdot 12 = 62 \end{cases}$

Average Bill = 52   Standard Deviation = 10

Hedging $\hat{q} = 10$ at $\hat{p} = 10$: $\begin{cases} prob = 0.5 & q=4, P=8, \quad \text{Bill} = 10 \cdot 10 - 6 \cdot 8 = 52 \\ prob = 0.5 & q=6, P=12, \quad \text{Bill} = 10 \cdot 10 - 4 \cdot 12 = 52 \end{cases}$

Average Bill = 52   Standard Deviation = 0

Example C: Perfect Negative Correlation Between Market Price and Customer Quantity

No Hedging : $\begin{cases} prob = 0.5 & q=4, P=12, \quad \text{Bill}=48 \\ prob = 0.5 & q=6, P=8, \quad \text{Bill}=48 \end{cases}$

Average Bill = 48   Standard Deviation = 0

$$\hat{q} / E[q] = 1 + \rho \,(cv(q)/cv(p)) \tag{1}$$

where *cv(q)* and *cv(p)* are the coefficients of variation of quantity and price, respectively, and ρ is the correlation between the price and the firm's quantity. Equation (1) implies that 100% hedging is optimal when price and quantity are uncorrelated, $\rho = 0$. When ρ is positive, optimal hedging will be greater than 100% and

the degree to which it increases is an increasing function of the variability of the firm's q and a decreasing function of the variability of the market p. Likewise, if $\rho<0$, the optimal hedge is less than 100% and by a greater amount as the variability of q increases relative to the variability of p. The result makes intuitive sense. For instance, if a firm's quantity is highly predictable, $cv(q)$ near zero, it is best off to just hedge that expected quantity, regardless of the correlation between price and quantity. McKinnon's result has been applied broadly in an agricultural economics literature on optimal hedging. I have not found an application to a buyer facing positive price/quantity correlation, a case that McKinnon does not discuss, but that follows immediately from his analysis.

To investigate how much empirical bite this insight has in the context of RTP, I evaluated the bill volatility of the same 1142 PG&E customers under various levels of under- and over-hedging. For this exploration, I simply fixed the level of hedging for a customer at different percentages of the customer's average consumption during the hedge/TOU period. So, for example, for a single customer, I calculated its bill volatility if it bought a hedge for each hedge/TOU period equal to 130% of its average consumption (over the four-year sample) during that hedge/TOU period. For each customer, I did this calculation for every 10% interval from 0% to 200% hedging.[19] I've carried out this analysis using price scenario I.

Defining "optimal hedging" for a customer as the level of hedging that minimizes its bill coefficient of deviation, the bars in figure 2 show the distribution of optimal hedge levels for the 1142 customers. About 77% of all customers are best off with some amount of over-hedging, *i.e.*, a hedge level above 100% of their average consumption for each hedge period. With every customer hedging optimally, the average coefficient of deviation drops to 0.147, down from 0.187 that resulted when all customers hedged 100% of their average consumption. In fact, unexpected bill volatility is lower on average under RTP with optimal hedging than it is under even a flat-rate tariff. The three lines in figure 2 show the average coefficient of deviation for customers in each of the hedging bins under three different billing arrangements. The difference between the "100% hedged" line and the "optimal hedge" line shows that the reduction in volatility for RTP customers can be quite substantial for those customers with optimal hedge positions that are far from 100%. Comparing the "optimal hedge" to the "TOU-F" line illustrates that optimally hedged RTP customers of all types have bill volatility at least as low as under the fixed-ratio TOU tariff they currently face.

Figure 3 presents the same analysis for the coefficient of maximum deviation. The value of over-hedging appears to be even greater in reducing the bill shock from an extreme outlier. About 84% of customers reduce their maximum unexpected bill deviation by over-hedging. The bill volatility lines show that optimal hedging is a substantial improvement over simple 100% hedging for most customers in reducing the highest unexpected bills. This somewhat more sophis-

19. The 0% hedge results correspond to those shown in tables 4 and 5 for RTP and the 100% hedge results correspond to those for RTP-H.

Figure 2. Distribution of Optimal Hedging and Effect on Coefficient of Deviation
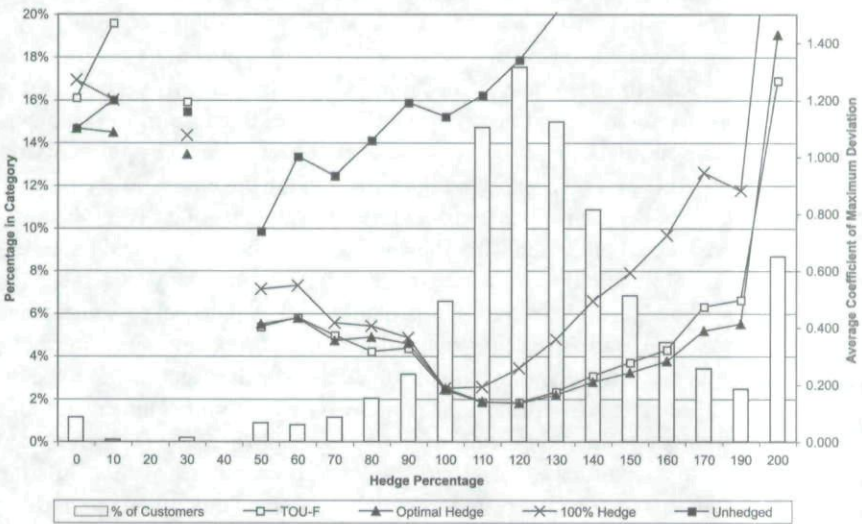

Figure 3. Distribution of Optimal Hedging and Effect on Coefficient of Maximum Deviation

ticated hedging strategy reduces the highest bills for most customers to about the same level they now face under TOU-F.

While this analysis suggests that variable hedging may be important in reducing customer bill risk, it likely overstates the potential benefits, because it considers the effect of the optimal *ex post* hedging strategy for each customer. In reality, customers will not know *ex ante* exactly how much hedging will be optimal for their future demand pattern. Their *ex ante* choices will necessarily be suboptimal on average. Still, variable hedging with even imperfect information would almost certainly be an improvement over naive hedging of 100% of a customer's expected demand. Further research into methods for *ex ante* determination of (constrained) optimal hedging levels is likely to be valuable.

6. CONCLUSION

Despite the fact that electricity bills are a small part of costs for most commercial and industrial customers, there is a great deal of concern with the volatility of this cost component. Such concerns have prompted some large customers to oppose real-time pricing on the grounds that it could increase the risk they face from volatility in their monthly bills.

I have shown that RTP without hedging could indeed substantially increase customer bill volatility, potentially leading to bills that in some months are double or triple the level that the customer would normally expect. That increased bill risk, however, can be almost entirely eliminated through use of simple forward purchase contracts that hedge price risk for a fixed quantity of power. I've demonstrated that the simple strategy of a customer buying its expected demand quantity through a forward purchase contract eliminates more than 80% of the additional bill volatility.

I then consider a somewhat more sophisticated approach to hedging, recognizing that the optimal amount of hedging for a customer will depend on the correlation between its demand and the real-time price. Because most customers exhibit a positive correlation, it is optimal for most to over-hedge, *i.e.*, purchase forward more than 100% of their expected demand. Among customers in the sample, the optimal level of hedging varies considerably; for some, less that 100% hedging would have minimized bill volatility. I show that optimal hedging would have substantially reduced bill volatility compared to 100% hedging for many customers. On average, optimal hedging even reduced bill volatility below the level customers face on conventional TOU tariffs. *Ex post* optimal hedging is probably not achievable, but this suggests that it would be worthwhile to explore further approaches to determining *ex ante* a customer's best hedging strategy based on limited information.

I have not considered the implications of variable hedging for the forward market. While it seems likely that sellers would also be interested in stabilizing their revenues—so would be interested in selling significant quantities forward—it is not clear that the demand for over-hedging that this analysis suggests would be met by an equally willing supply of over-hedging without a significant premium. In fact, to the extent that producers are subject to random outages and that they occur more frequently when prices are high, McKinnon's analysis suggests that

they might want to hedge less than 100% of their expected output if there were no risk premium or discount. Bessembinder & Lemon (2002) suggest that if only risk-averse industry participants trade in a forward market for electricity, an equilibrium risk premium will result that will be negative during low-demand/low-volatility times and positive during high-demand/high-volatility times. This would, of course, leave speculative opportunities that have positive net present value. Borenstein, Bushnell, Knittel, and Wolfram (forthcoming) find significant forward premia/discounts in the California electricity market, but argue that they could not plausibly be attributable to risk aversion. Whether entry from other financial participants, which has occurred to some extent in electricity, would be sufficient to drive forward prices to expected spot price levels is an open question.

I also have not yet analyzed the role that option contracts could have in addressing bill volatility. It might at first appear that a call option contract would eliminate all bill volatility for the buyer that is due to price fluctuations, because it gives the option to buy (or not) at a fixed price. This intuition, however, is not accurate. A financial option yields positive returns whenever the market price is above the strike price, regardless of the demand of the customer holding the option at that particular time. So, profit-maximizing exercise of options would not eliminate bill volatility (net of profits from holding the option). The degree to which the contingent payoff structure of option contracts would help to mitigate bill volatility merits further research.[20]

Appendix. Bill Volatility Using Alternative Wholesale Price Scenarios

Price Scenario II: Less Volatile Simulated Prices

Coefficient of Deviation in Monthly Bill

			PERCENTILES			
	Mean	1st	10th	50th	90th	99th
Tariff						
Flat-Rate	0.155	0.021	0.035	0.087	0.364	1.048
TOU - Fixed-ratio	0.156	0.021	0.037	0.086	0.365	1.057
TOU - Cost-based	0.159	0.022	0.039	0.088	0.365	1.083
RTP - Hedged	0.171	0.025	0.044	0.097	0.385	1.083
RTP	0.284	0.132	0.163	0.228	0.449	1.176

Coefficient of Maximum Deviation in Monthly Bill

			PERCENTILES			
	Mean	1st	10th	50th	90th	99th
Tariff						
Flat-Rate	0.318	0.042	0.072	0.182	0.658	2.359
TOU - Fixed-ratio	0.324	0.043	0.078	0.186	0.668	2.392
TOU - Cost-based	0.346	0.048	0.088	0.199	0.691	2.456
RTP - Hedged	0.447	0.056	0.110	0.266	0.908	2.897
RTP	0.920	0.372	0.536	0.769	1.368	3.366

20. Brown & Toft (2002) and Oum, Oren & Deng (2005) consider optimal hedging with options when price and quantity are stochastic.

Appendix. Bill Volatility Using Alternative Wholesale Price Scenarios, cont.

Price Scenario III: Actual Northern California Real-Time Wholesale Prices

Coefficient of Deviation in Monthly Bill

Tariff	Mean	1st	PERCENTILES 10th	50th	90th	99th
Flat-Rate	0.155	0.021	0.035	0.087	0.364	1.048
TOU - Fixed-ratio	0.156	0.021	0.037	0.086	0.365	1.057
TOU - Cost-based	0.155	0.021	0.036	0.086	0.363	1.045
RTP - Hedged	0.186	0.025	0.044	0.107	0.427	1.093
RTP	0.637	0.367	0.519	0.586	0.811	1.425

Coefficient of Maximum Deviation in Monthly Bill

Tariff	Mean	1st	PERCENTILES 10th	50th	90th	99th
Flat-Rate	0.318	0.042	0.072	0.182	0.658	2.359
TOU - Fixed-ratio	0.324	0.043	0.078	0.186	0.668	2.392
TOU - Cost-based	0.317	0.043	0.073	0.181	0.672	2.350
RTP - Hedged	0.508	0.051	0.100	0.280	1.127	3.894
RTP	1.783	0.809	1.380	1.585	2.261	5.036

REFERENCES

Barbose, Galen, Charles Goldman and Bernie Neenan (2006). "The Role of Demand Response in Default Service Pricing," *The Electricity Journal*, April: 64-74.

Bessembinder, Hendrik and Michael L. Lemmon (2002). "Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets," *Journal of Finance*, 57(3): 1347-1382.

Borenstein, Severin (2005). "The Long-Run Efficiency of Real-Time Electricity Pricing," *Energy Journal*, 26(3): 93-116.

Borenstein, Severin (2007). "Wealth Transfers Among Large Customers from Implementing Real-Time Retail Electricity Pricing," *Energy Journal*, 28(2): 131-148.

Borenstein, Severin, James Bushnell, Christopher R. Knittel and Catherine Wolfram (forthcoming). "Inefficiencies and Market Power in Financial Arbitrage: A Study of California's Electricity Markets," *Journal of Industrial Economics*.

Borenstein, Severin and Stephen P. Holland (2005). "On the Efficiency of Competitive Electricity Markets With Time-Invariant Retail Prices," *RAND Journal of Economics*, 36(3):469-493.

Braithwait, Stephen and Michael O'Sheasy (2002). "RTP Customer Demand Response: Empirical Evidence on How Much Can You Expect," in Faruqui and Eakin, eds., *Electricity Pricing in Transition*, Boston: Kluwer Academic Publishers.

Brown, Gregory W. and Klaus Bjerre Toft (2002). "How Firms Should Hedge," *Review of Financial Studies*, 15(4): 1283-1324.

McKinnon, Ronald I (1967). "Futures Markets, Buffer Stocks and Income Stability for Primary Producers," *Journal of Political Economy*, 75(6): 844-861.

O'Sheasy, Michael (2002). "Is Real-Time Pricing a Panacea?" *The Electricity Journal*, December: 24-34.

Oum, Yumi, Shmuel Oren and Shijie Deng (2005). "Hedging Quantity Risk with Standard Power Options in a Competitive Wholesale Electricity Market," mimeo, U.C. Berkeley, October.

Siegel, Daniel R. and Diane F. Siegel (1988). *Futures Markets*, Orlando, FL: Dryden Press.

PJM Interconnection Market Monitoring Unit (2005). "2004 State of the Market Report." Available at http://www.pjm.com/markets/market-monitor/downloads/mmu-reports/pjm-som-2004.pdf.