Optimizing Retail Contracts for Electricity Markets

Ross Guttromson

David Chassin

Pacific Northwest National Laboratory
PO Box 999 Richland, WA 99352

Ross.guttromson@pnl.gov

David.chassin@pnl.gov

Keywords: Retail Markets, Electricity, Efficient Frontier, Optimal

Abstract

The Olympic Peninsula Testbed Demonstration in Washington State allowed residential electricity customers the choice of three different types of retail electricity rates; a time-of-use price, a fixed price, and a 5-minute real-time price. Each of these rates structures has advantages and disadvantages for both the residential consumers as well as the utility offering them. This paper focuses on a methodology to select the mix of rate types a utility should offer to its residential consumers given the various objectives it seeks to achieve. The method used to determine an optimal mix was borrowed from stock market portfolio theory and results in what is referred to as the Efficient Frontier. This solution defines an optimal mix of contract types among many possible combinations.

Efficient Frontiers in Stock Portfolio Theory

The concept of efficient frontiers was introduced in 1957 by Nobel Prize winner Harry Markowitz as part of the Capital Asset Pricing Model (CAPM) for portfolio theory. The theory is based on the idea that combining several stocks into a portfolio will yield decreases in overall risk below that of any individual stock while retaining high returns.

Figure 1 depicts this idea. The dark shaded region shows all possible ways (weightings) to combine a group of stocks to make up a portfolio. Anywhere on the top leading edge of this region (called the efficient frontier) provides the optimal combinations (weightings) of these stocks for all possible portfolios that can be created. This edge provides the highest return for the lowest risk. Why would a person wish to invest in a portfolio in the central area of this curve? They wouldn't, since combining the same stocks in a different manner can always increase your return without increasing your risk or analogously, decrease your risk for the same return. In the case of stocks risk is defined as the volatility of a stock's price. We will see that the definition of risk differs for electricity rates.

In its truest form, Figure 1 simply shows how normal random variable distributions combine to form a unique random variable distribution. This concept can be used to estimate the best (or optimal) way to combine any set of normal random variables given a clear objective.

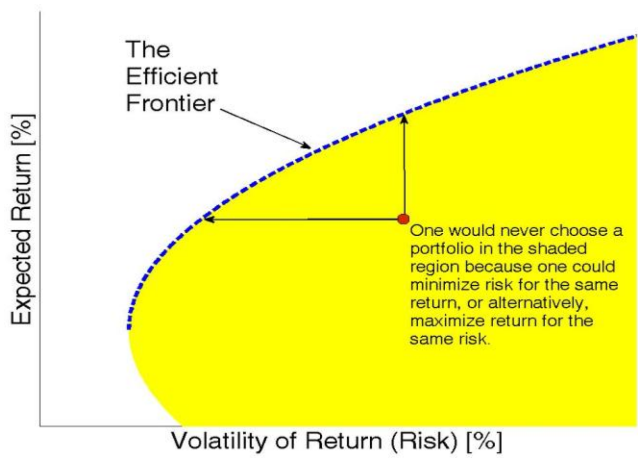

Figure 1. Efficient Frontier for a Stock Portfolio

We will use these principles to look at random variables generated by electricity markets, in the case retail-level markets. We pose the questions, "Given several types of rates that can be offered to customers, what is the optimal combination customer subscriptions to these rates given my objectives"? In the Olympic Peninsula Testbed Demonstration, we tested three retail electricity rates: a fixed price, a time-of-use price, and a 5-minute real-time price. Each of these rates offered electricity to customers who in different ways. Fixed-price customers paid the same price all year. Time-of-use customers paid a daily on-peak or off-peak price, which changed seasonally. Real-time price customers paid a price derived every 5 minutes based on the true cost of delivering the electricity during that period.

Each rate was characterized by data collected over a one year period that make up the random variables needed to perform the efficient frontier calculations. Note that the interpretation of the results for the calculations in utility markets that follow do not necessarily have the same implication as they do in stock portfolio analysis. For example, a point on the efficient frontier in portfolio analysis is by definition considered "good", whereas in evaluating utility market structures, the same part of the curves can only be judged as good or bad in the context of the utility's objectives. This analysis does not yield conclusive directives. Instead it provides a rich mechanism to evaluate the consequences of any given rate offering mix. Whether one rate offering mix is good or bad depends upon the objectives of the utility.

1.1. Random Variables and Normal Distributions

It is essential to understand what random variables are. Figure 2 shows a set of normally distributed random numbers. There is a portion of this data that appears random, such as the scattering effects of the points, and a portion of the data that does not appear to be random at all, such as clustering around the average, and the typical spread of the data around the average. These particular random variables have mean (or average) equal to 2, and their standard deviation (the average distance of each point from their collective average) of 0.2. The manner in which we describe these random variables (by a mean and a standard

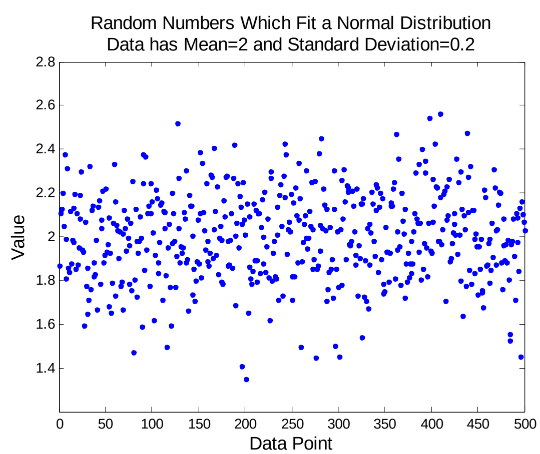

Figure 2 Normally Distributed Random Data

deviation) completely identifies this random variable set. For random variables, we don't care what the actual values are, but rather what the data as a whole looks like. When we consider the data in this manner, we allow for the fact that the next set of data will be completely different, yet will have the same mean and standard deviation. This allows us to evaluate results of events knowing that the particularly values we observe change, but their mean and the standard deviation remain constant.

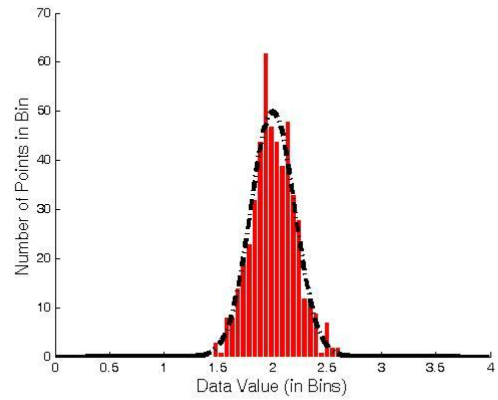

Figure 3 Histogram of Normally Distributed Random Data With Mean=2 and Stdev=0.2

Figure 3 shows a histogram of the same data set seen above. The solid bars constitute a histogram that characterizes the data. This format allows comparison of the data to the curve of a normal or Gaussian distribution. The dashed line is a mathematically defined probability density function based on the mean and standard deviation of the random variables above. If the data fits a normal distribution, then we can claim that the random data is "normal" supporting our conjecture that even though the values may change from one observation to the next, the mean and standard deviation remain constant. From this analysis, we conclude that the data is indeed normal (confirmed by the fact that the data set was created using normally distributed random variables).

The equation below defines a set of random variables by a probability density function. There are only two variables in this equation: the mean, μ; and the standard deviation, σ. These two parameters are sufficient to completely define a normal random variable distribution containing any number of data points.

$$N(\mu,\sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(x-\mu)^2/2\sigma^2}$$

Normal Probability Density Function

1.2. Portfolios: Combining Normal Distributions

Consider two different normal distribution curves, each completely defined by its respective mean and standard deviation, as shown in Figure 4. These two curves represent data from two independent sources, meaning that no observation in one is in any way related to an observation in the other. The first normal distribution curve might represent income from growing wheat, while the second normal distribution curve might represent income from growing barley. Suppose we want to know what income to expect if we grow both wheat *and* barley. We create a new distribution curve that is a weighted combination of the other two. We do this by combining the expected values, μ, and the variances, σ, of the two given normal distributions as follows:

$$\mu_{new} = \omega_1 \mu_1 + \omega_2 \mu_2$$

$$\sigma_{new}^2 = \omega_1^2 \sigma_1^2 + \omega_2^2 \sigma_2^2 + 2\omega_1 \omega_2 \rho_{12} \sigma_1 \sigma_2$$

Where ρ is the covariance between the two data sets and ω is a weighting factor to determine how much of a given distribution is to be added (note that $\omega_1 + \omega_2 = 1$)

The lightly marked normal distribution functions are obtained for various proportions of wheat and barley. There are many of these curves, each representing the sale of a different mix of wheat and barley. But together, all these curves represent all possible income levels by growing different combinations of wheat and barley.

One might assume that the mean value of each curve would simply follow a relatively straight line between the two curves, but this is not at all what happens. This result confirms that something very important is going on.

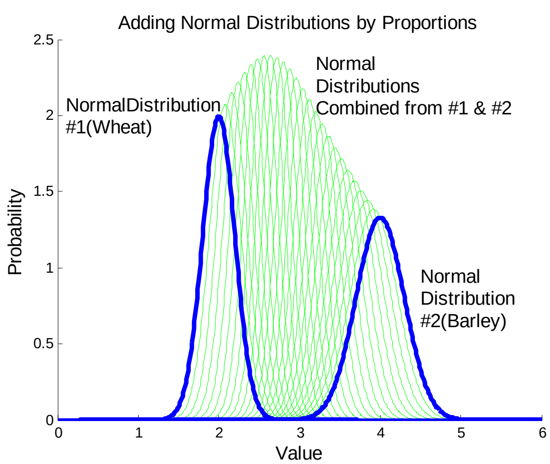

Figure 4 Combinations of Two Normal Random Variable Distributions

Looking at this a different way, we recall that each normal distribution data set can be defined by just two numbers–a mean and a standard deviation. Using these only, we develop what is called the efficient frontier, as shown in Figure 5. Mathematically, this process is simply combining the probability density functions together by the proportions listed.

In the wheat/barley example, we consider what mixture of wheat and barley to grow knowing the expected income (mean) and variability in income (standard deviation) of all possible combinations of wheat and barley. Clearly over a period of many years, the mean income is maximized when only barley is grown. But what about the variability of income? In this context, the standard deviation refers to how far from the mean income each year's income is. If the uncertainty of income is not important, then it is clear that barley is the preferred growing strategy.

However, if the variability of income is a consideration, such as when a steady income is desired, then some income must be sacrificed in exchange. This is not much different than paying interest on a short term loan to cover seasonal operating expenses. The variability in income is minimized at the optimal mix of wheat and barley with a mean of 2.6 and a standard deviation near 0.165. The portion of the graph made as small circles constitutes the efficient frontier. Where along that continuum you decide to operate is a matter of preference. You would never want to drop below this optimal point, however, as you would be increasing your variability of income while decreasing your income. What if you sold wheat exclusively? Given what we see above, by selling a little barley along with your wheat, you would increase your income and make it more stable too.

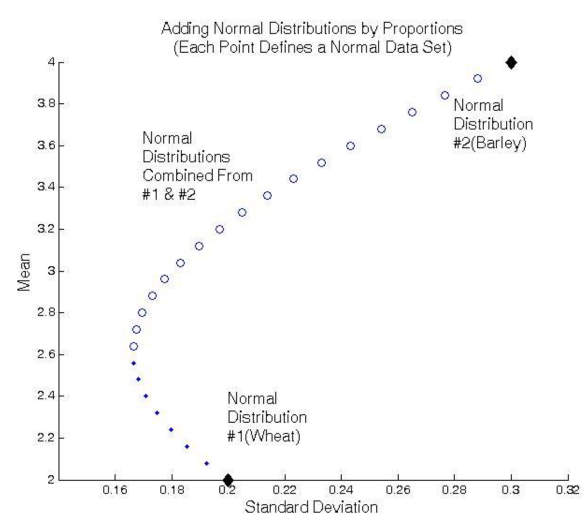

Figure 5 The Efficient Frontier

1.3. Retail Electricity Markets

Now we can consider the electric power utility industry. In the Olympic Peninsula Testbed Demonstration, there were three types of residential rate contracts offered to consumers of electric power: fixed price, which charged a fixed rated for electricity usage in dollars per amount of energy used, time-of-use price, which charged two different seasonal rates for electricity that were consistently applied for specific hours of each day, and the real-time price, which charged higher rates for electricity usage when the power system capacity was at or near its capacity limit and lower prices when the system had excess capacity.

1.3.1. Optimizing Contract Selection for Peak Power

Figure 6 shows the peak energy usage. Only data from the times of the year and day when energy consumption was high were used for this analysis– specifically the time of year from November 1st to December 8th and the hours of the day from 6am to 9am and from 6pm to 9pm. This data represents the times when the electric power system was at its highest load relative to the available capacity, and therefore represents the best time to evaluate at how the different contract types influenced (both suppliers and consumers) the systems response to capacity constraints.

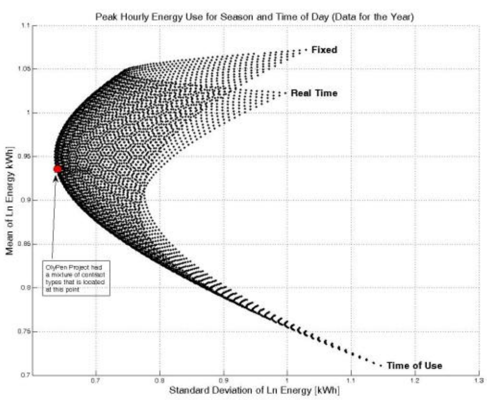

Figure 6 Contract Type Impacts on Peak Power

Figure 6 shows the efficient frontier analysis based on this peak condition data. The shaded area represents all possible proportions of combining the three markets types- and the sharp points at the ends of the shaded regions represent a market that is not mixed at all, but consists only of that type of contract. For example, the word "Fixed" appears next to a corner point near (1.05kWh, 1.07kWh). That point represents a mixture of contract types where Fixed type = 100%, Time of Use type = 0% and Real Time type = 0%. Moving away from the extreme points in the shaded region, other contact types start getting mixed together. The details of how to determine what the mix is will be explained below. The Olympic Peninsula project had a mixture of roughly 1/3 of each contract type.

To interpret Figure 6, the utility's objective must be known. Presumable, the utility would like to reduce peak energy during its times of high load or limited capacity. Doing so allows the utility to defer very expensive system capacity upgrades to accommodate the increase in electricity use during these periods. This objective tells us that the peak energy use (y-axis) should be as low as possible.

So what about the variability (x-axis) of the peak energy use? At first thought we might say that we want the variability to be low. However, if we take it as given that the peak energy is low, we would want the participants to be responsive- that is- to change their energy use as a result of price changes. This implies that we actually desire a high variability. Together, the evaluation above points to the desired market structure as the Time of Use rate by itself- not mixed in combination with any of the two other rates structures. But this evaluation is incomplete, and does not consider other objectives of the utility, such as Gross Margin. Let's consider these affects next.

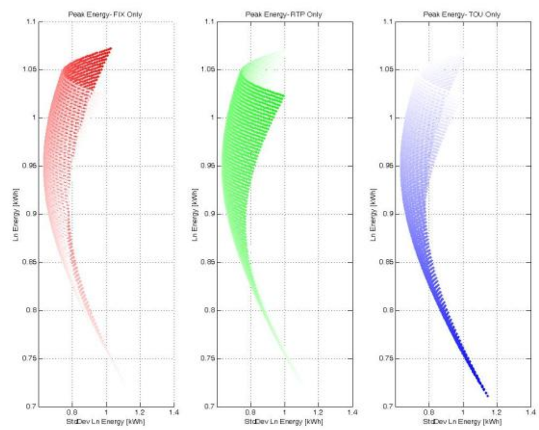

Figure 7 Contract Type Impacts on Peak Power (Contract Types Separated)

Figure 7 is simply Figure 6 'taken apart', allowing one to see what the constituent contract mixes are that make up the entire curve. For example, the Olympic peninsula project is shown in Figure 6 at the point Stdev=0.64 and Peak Energy=0.93. Looking at all three graphs in Figure 7, one can see that this point has shading in all three curves. This implies that all three curves participate in the contract mix at that point. Analogously, one can use Figure 7 to help determine the contract options for a desired Stdev and Peak Energy. Figures 6 and 7 are better represented by using a single colored version of Figure 6, allowing the color to represent the contract type- as is represented below in Figure 10 for the Gross Margin analysis. The x-axis on these three graphs has been squeezed in order to allow sufficient space in the document, however, the range of the axes in Figure 7 are the same as those in Figure 6.

1.3.2. Optimizing Contract Selection for Gross Margin

Gross Margin is defined as the revenue generated by the sale of electricity minus the cost of that electricity. It does not include costs of infrastructure, labor, taxes, overheads, or other fixed costs. It simply gives an early preview of what profits might look like. Omitting these other charges helps keep this financial metric relevant to a broader range of companies- all of which can add back in their own specific fixed charges. Unlike the previous analysis which looked only at peak periods of electricity use, this analysis uses data for the residential homes for the entire year at 24 hours per day and 7 days per week. Note that this is the exact same customer set, but we are now considering data from a different period.

One might expect the same curve as before, but that is not the case. Earlier we considered peak power usage, and now we are looking at gross margin–both important to a utility, but each the basis for completely different objectives.

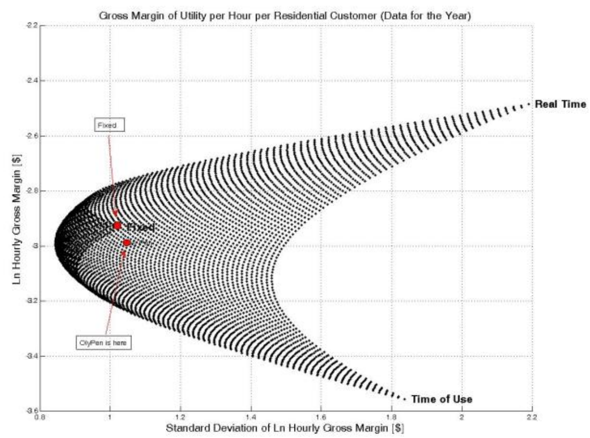

Figure 8 Contract Type Impacts on Gross Margin

Regarding the utility's objectives, we assume that they would like a high gross margin. Regarding the variability of this gross margin, we might consider that all else being equal, the utility would like it minimized. However it is probably not very important in this analysis since seasonal affects will likely have more impact on gross margin variability than would contract type. Because low variability implies a lower gross margin, each utility must establish for itself where on this upper leading edge it would prefer to operate.

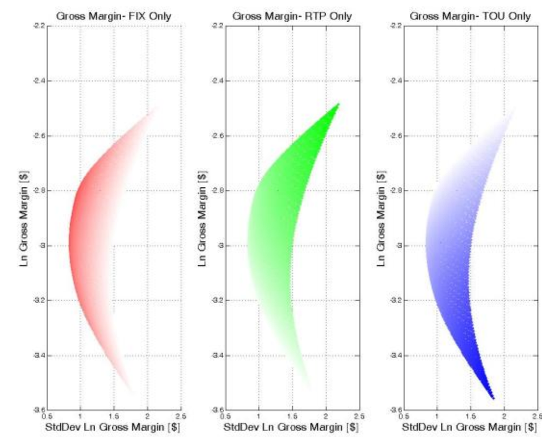

Figure 9 Contract Type Impacts on Gross Margin (Contract Types Separated)

As was the case with Figure 7, Figure 9 'takes apart' the contract types embedded in Figure 8, allowing one to observe which contract types contribute to a given contract mix. Again, the x-axis on these curves have been squeezed to allow sufficient space to insert them into this document.

1.3.3. Comparing Gross Margin analysis to Peak Power Analysis

By now, it seems obvious that picking a contract mix which minimizes peak power does not necessarily result in an optimal contract mix that maximizes gross margin. Earlier, we concluded that the data supported time-of-use contracts as best for reducing peak power and thus deferring capacity upgrades. Follow up analysis shows that gross margin is maximized by emphasizing real-time contacts whereas the time-of-use contract type minimizes gross margin.

It should now be apparent that this problem requires optimization, but not all the information needed is available. It would be very helpful to know how a point on one graph is translated to the other graph. The graphs so far have not shown this information,

Figure 10 shows a 3-D surface situated above the region taken from Figure 8. The 3-D surface reveals information about the mix of contracts for all points on the efficient frontier plot. The mixture has been represented by color where all red indicates fixed-price contracts, all blue indicates time-of-use contracts, and all green indicates real-time contracts. Mixtures of contracts are represented by mixing the colors in similar proportions. Therefore, areas on the map with a some red, a some blue, and a some green indicated a contract mix in that proportion. As we can see, the Olympic Peninsula Testbed Demonstation point is one such representation, with about 1/3 each. The numeric value along the z-axis represents the combined colors and can be ignored–it is simply a convenient mathematical method of separating multiple solutions along this axis.

Folding over of the 3-D surface implies that multiple solutions can be found for some regions. One can clearly see the fixed-price contract type is folded up and back over the minimum volatility region. This superimposed area implies that more than one contract mix exists for a number of points near the efficient frontier. This means that is more than one choice of contract mix to achieve a given objective on the efficient frontier curve (in the areas where these solutions overlap each other).

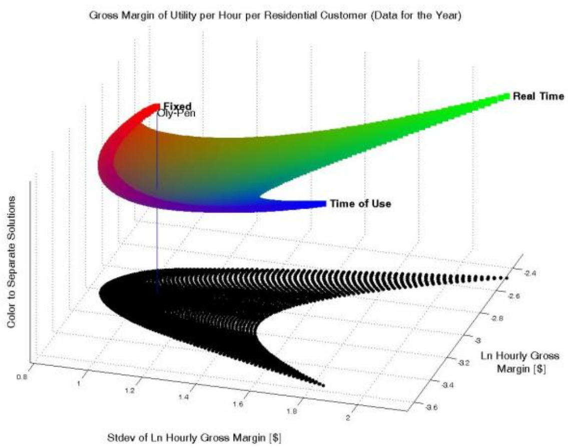

Figure 10 Contract Type Impacts on Gross Margin Shown in 3D With Color Denoting Contract Type

Qualitative Optimization

We observed above that time-of-use contracts were more effective for reducing peak power and the real-time contracts for maximizing gross margin. We now have a method to determine what mix of contracts best suits any the objectives of the utility- and are able to understand the resulting tradeoffs in peak shaving and gross margin for any given portfolio of contracts. Of course, there are more than two criteria that may be considered in selecting a contract mix. The technique below shows how to create other efficient frontiers, and it is up to the reader to evaluate these based on specific objectives.

References

[1] Siegel, A., "Practical Business Statistics", 5th Edition, 2002, McGraw-Hill/Irwin

[2] Papoulis, A., S. Pillai, "Probability, Random Variables and Stochastic Processes", 4th Edition, 2002, McGraw Hill

[2] Markowitz, Harry M. "Portfolio Selection", 1952, Journal of Finance 7 (1): 77-91

[3] Higgins, R., "Analysis for Financial Management", 8th Edition, 2007, Irwin/McGraw-Hill

[4] Brealey, R., "Principles of Corporate Finance" 2005, McGraw-Hill/Irwin,

Biographies

Ross T. Guttromson (M'01, SM'04) received a BSEE and MSEE degrees from Washington State University, Pullman. He is currently pursuing an EMBA degree from the University of Washington. Ross manages the Electricity Infrastructure Operations Center and is a senior research engineer with the Energy & Environmental Directorate at the Pacific Northwest National Laboratory, Richland, WA. He was with R. W. Beck Engineering and Consulting, Seattle, WA, from 1999 to 2001 and with the Generator Engineering Design Group, Siemens-Westinghouse Power Corporation, Orlando, FL, from 1995 to 1999. Ross has two US patents, and is author and co-author of several papers on power systems analysis. He is a senior member of the IEEE and a member of the WECC Planning Coordination Committee. Ross is a U.S. Navy Submarine Veteran, having served on the USS Tautog (SSN 639) from 1987 to 1991.

David P. Chassin (M'02, SM'05) received his B.S. of Building Science from Rensselaer Polytechnic Institute in Troy, New York. He is a staff scientist with the Energy Science & Technology Division at Pacific Northwest National Laboratory, where he has worked since 1992. He was Vice-President of Development for Image Systems Technology from 1987 to 1992, where he pioneered a hybrid raster/vector computer aided design (CAD) technology called CAD OverlayTM. He has experience in the development of building energy simulation and diagnostic systems, leading the development of Softdesk Energy and DOE's Whole Building Diagnostician. He has served on the International Alliance for Interoperability's Technical Advisory Group and chaired the Codes and Standards Group. His recent research focuses on emerging theories of complexity as they relate to high-performance simulation and modeling.