Electric Market Risk:
Clearing Out The Cobwebs[1]

Robert McCullough
Managing Partner
McCullough Research
Portland, Oregon

America's greatest economist once wryly remarked "models are constraining on the underlying facts." To match his cynical tone, I would add that we are lucky that this is so since we would otherwise have to work much, much harder.

We are continuously assured that electric market risk started about three or four years ago. This is a corollary to the Veblen rule that goes "it isn't so if it isn't in the paper." This attitude has lead to an interesting form of schizophrenia in risk management in electricity with newly imported risk managers from the gas business dusting off their Black-Scholes models and leaping forward into the new world with little understanding of the structure or the data. As a number of nascent risk managers found out in Illinois last summer, the simple fact that Black-Scholes fundamental assumptions are completely inconsistent with the data and structure can making learning very expensive.

This paper summarizes four of the pitfalls where assumptions can overrule facts and then proceeds to evaluate possible solutions.

Identifying The Cobwebs

1. The Normal Distribution

A central part of the MBA's statistical education is the standard distribution. As we all remember, nature has kindly provided many examples of physical phenomena where this simple mathematical abstraction turns out to be a surprisingly good approximation of the real world. Some of us adopt an almost religious level of faith in the coincidence between real world data and this simple distribution.

Unfortunately, real market data reflects factors far more complex than those assumed for an intro level MBA statistics text. Actual market prices over the past twenty years aren't difficult to accumulate – in fact most utilities have such data. Before reliable market compilations started to appear, the source is from dispatch logs and similar sources. Actual data quality improves with age – recent indices are estimates and in the case of Dow's estimates, highly stylized estimates.

---

[1]Presented at tutorial on "Managing Risk in the New Power Business," IEEE PES Power Industry Computer Applications (PICA) Conference, Santa Clara, CA, 17 May 1999

Older data reflects actual market transactions. In a sense, since 1995 we have taken to using estimates of data which we used to observe. Yes, the world did exist prior to MegaWatt Daily.

The actual distribution has no resemblance to a normal distribution. This shouldn't surprise anyone – the actual market reflects a variety of operational and institutional factors that have no resemblance to a normal distribution.

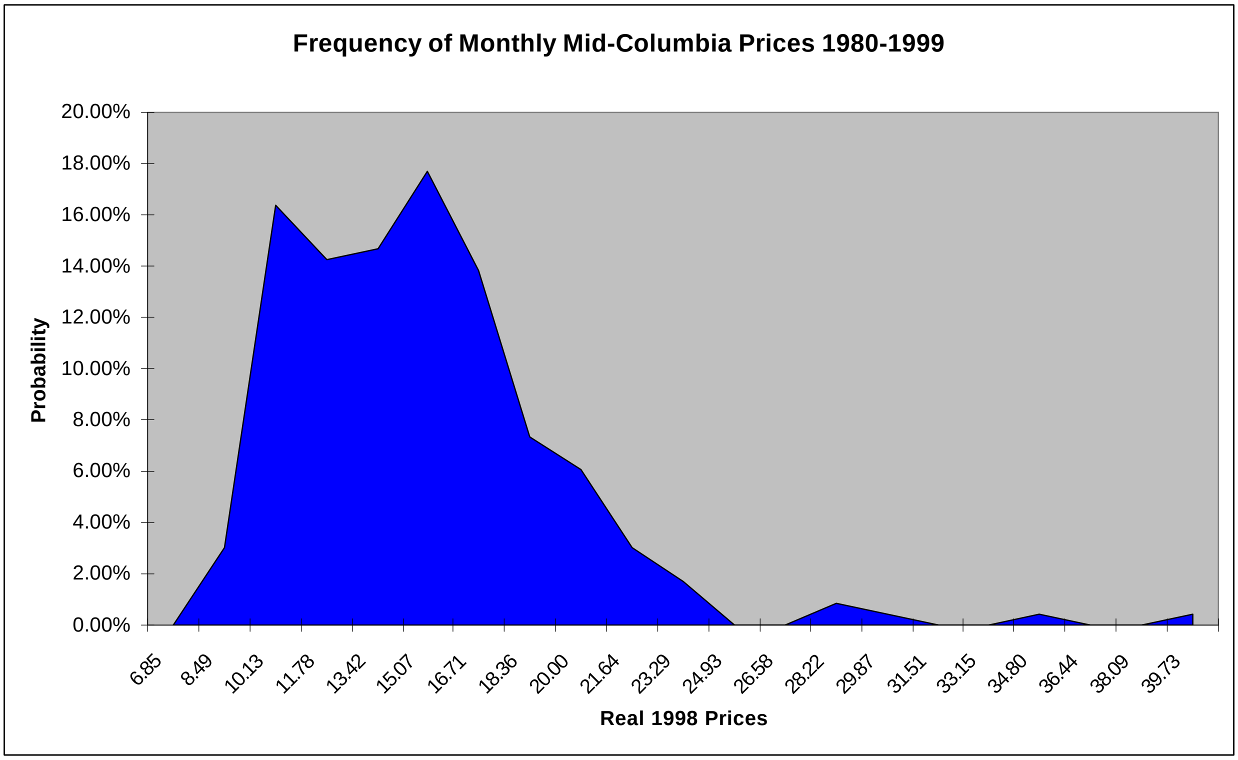

Lincoln Wolverton observed many years ago that using normal distributions for market data has little economic logic. He hypothesized that commodity prices with an inelastic supply curve would resemble a log-normal distribution – simply because the structure of pricing shifts dramatically at full capacity utilization. At levels below 100% of capacity, we can expect prices to reflect an economic balance between the participants. At 100% utilization, the terms of trade shift and prices are determined by scarcity value to buyers. As we saw in Illinois last summer, scarcity value can be very high indeed.

This chart is misleading, however. Institutional forces on the west coast shift one year ago. BPA's Larry Kitchen notes than on April 1st, 1998, California's large utilities stop focusing their energies on controlling the imported price of energy in favor of manipulating stranded cost. If we take the same chart and plot post April 1998 values in red we see a surprising result. The right hand tail of the distribution seems to reflect only prices in the past year. The easiest explanation simply doesn't work very well. The Pacific Northwest was near load resource balance during the early 1980s, so the balance of supply cannot be the only explanation.

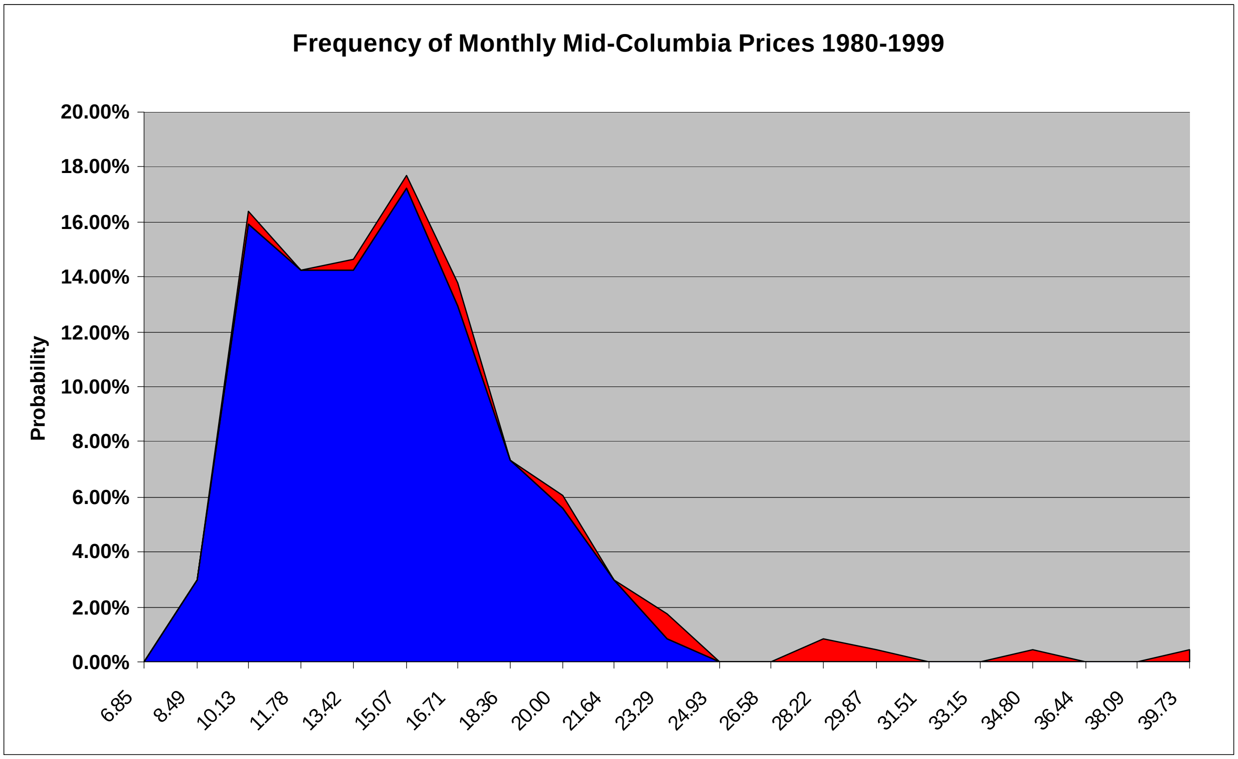

My recent article, "Electric Competition, One Year Later: Winners and Losers in California," argues that much of the increase in price is due to a shift in the terms of trade between California and the Pacific Northwest.[2]

We are currently wrestling with the implications of a new application of Thorsten Veblen's dicta that the model determines the data. The results for a number of months under California's new Independent System Operator/California Power Exchange are clearly anomalous. The ISO, in particular, has imposed its own institutional beliefs on power prices in California. This was amply demonstrated in the first week of September 1998. The ISO decided that the system was facing emergency conditions and launched into 250 mill purchasers throughout the week. Actually, no such emergency existed. The actual situation simply reflected market participants gaming the ponderous ISO/PX market mechanism.

As an issue for risk management, the ISO/PX institutional beliefs are just as relevant as real market forces. If the ISO and the PX believe that prices should be high, the market will rush to sell energy to Californians at inflated prices and the results – for market participants – are identical to those that would occur if the system really faced emergency conditions. All in all, this is a 90s version of agricultural price controls.

The bottom line is straightforward. This is not a normal distribution. Worse yet, it is fairly clear that it isn't a simple random variable at all.

---

[2]Electric Competition, One Year Later: Winners and Losers in California, Public Utilities Fortnightly, March 1[st], 1999.

This discussion is only necessary since the simplistic models that are relevant to interest rate risk were ported to natural gas and are now being applied to a market where they are wholly inappropriate. A fascinating monograph representing the most dangerous of intuitive leaps in statistical usage was published last year by Dragana Pilipoviƒ.[3] Entitled "Energy Risk", it subtitle goes on to say that it "Includes Electricity, Natural Gas, and Other Energy Markets." Application of the tools recommended in this book is the moral equivalent of issuing full loaded M-16s to a playground of small children.

2. Terms of Trade

All of the market models in our beginning economics text showed smoothly sloping demand and supply curves. Indeed, given enough time, it is arguable that the existence of open entry and exit as well as possible substitution will always lead to elastic demand and supply curves.[4]

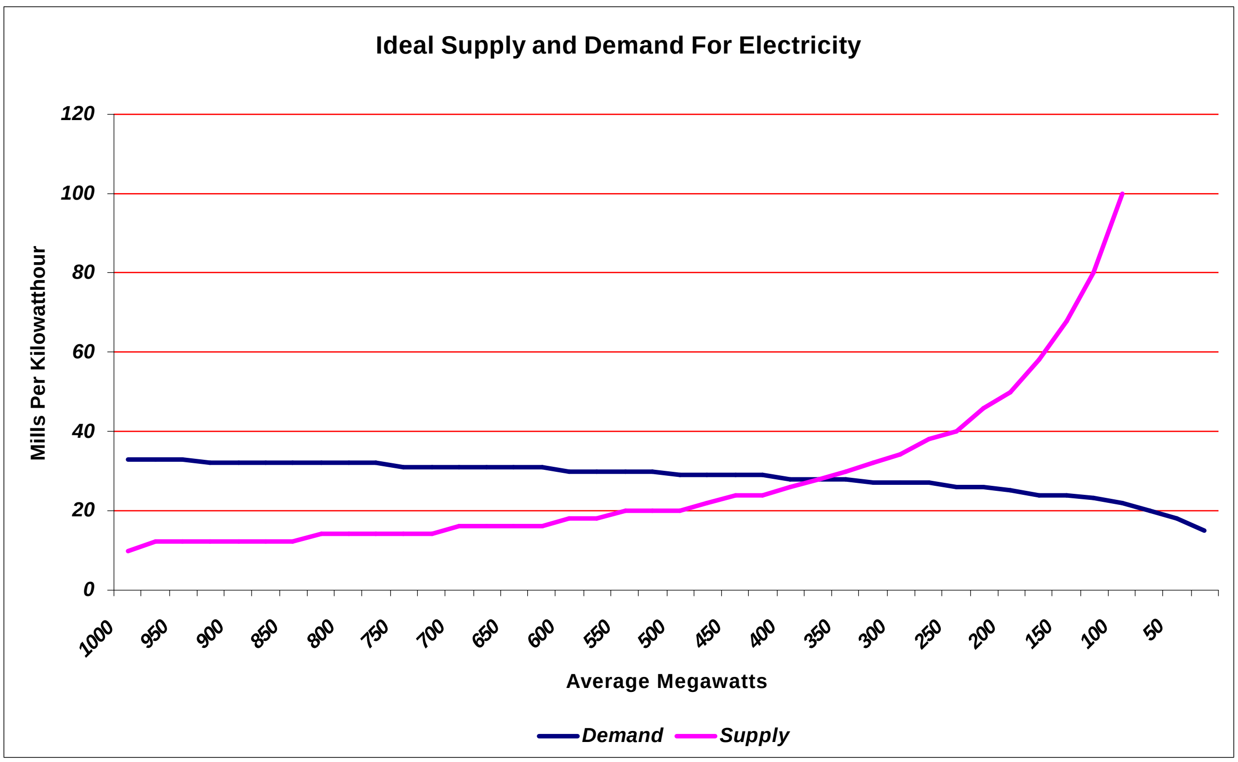

Recent deregulation efforts have shown a touching belief in the a priori slope of the electric

---

[3]Energy Risk, Dragan Pilipoviƒ, McGraw-Hill, 1998. See especially the discussion on page 67 for the representation of seasonal effects as a sine curve.

[4]Economists use the term "elasticity", (change in quantity/quantity)/(change in price/price) as shorthand to describe a world where quantities respond smoothly to changes in price. Everyone but monopolists hope that supply and demand are elastic. Monopolists, on the other hand, streak towards markets where quantities aren't elastic. This is one of the original arguments for natural monopoly in the electric business.

supply and demand curves. The Harvard Energy Project, for example, has generally mandated that all of the curves will be elastic without the benefit of experience or research. The reality is very, very different.

We know that the supply curve for electricity has a kink at the upper end – we simply can't produce more electricity after a certain point. When Sam Insull, the founder of the electric industry, invented the terms "energy" and "capacity", he did so because the issue of the finite availability of capacity was terribly important to customers and suppliers. To an extent, this problem is unique to the electric business since our ability to store electricity is terribly limited. Other industries with limited storage at least can fall back on the costs and speed of delivery, but this isn't terribly helpful for a business whose product moves at the speed of light.[5]

The very reassuring graph shown above actually looks like a gentle slope until it meets two near vertical lines:

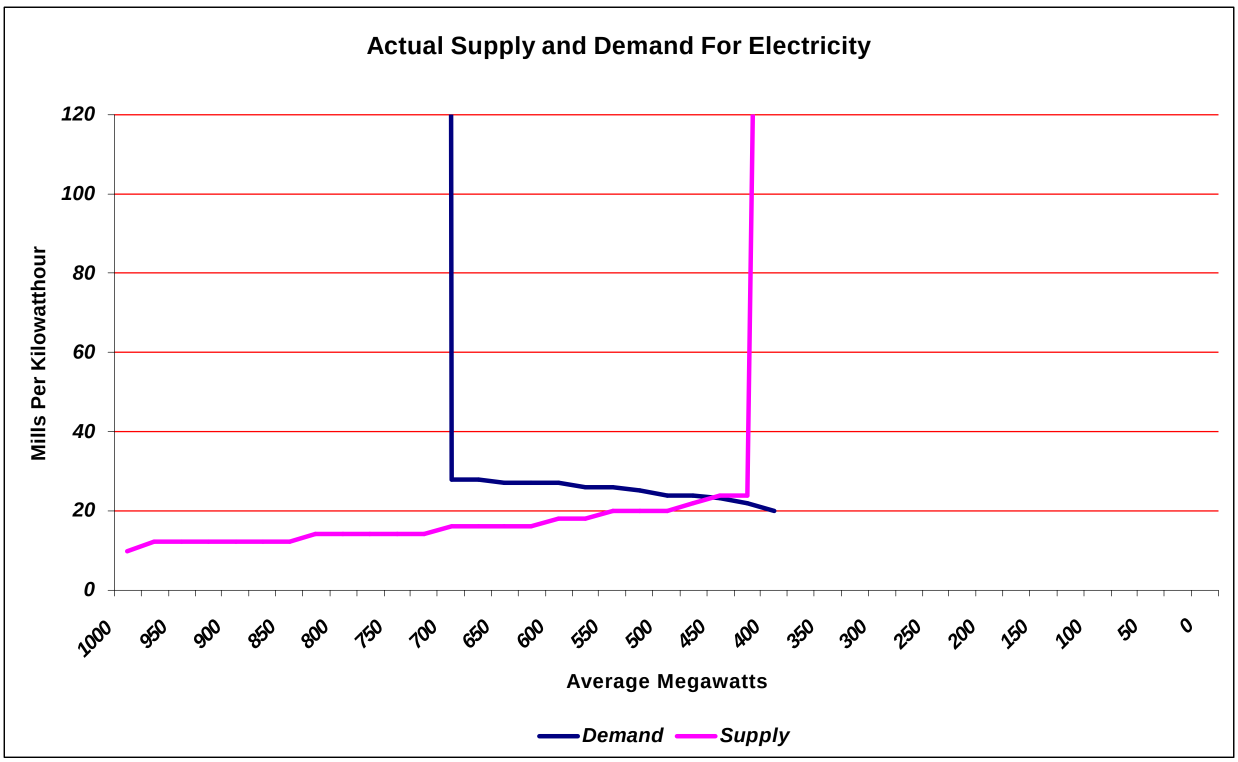

This represents demand and supply curves more representative of the real world. The elastic portion of the demand curve crosses the elastic portion of the supply curve. Conditions in Illinois

---

[5]Storage of all commodities is limited. When the grain elevators fill up and the boats are full, we simply can't store more grain. We can pay for the existing grain to be delivered more quickly or to be delivered from farther and farther away. Electricity simply doesn't fit into this model. Our market is the same scale as the available grid and the delivery is immediate. Electricity is simply less forgiving than corn or crude oil.

last summer looked like this:

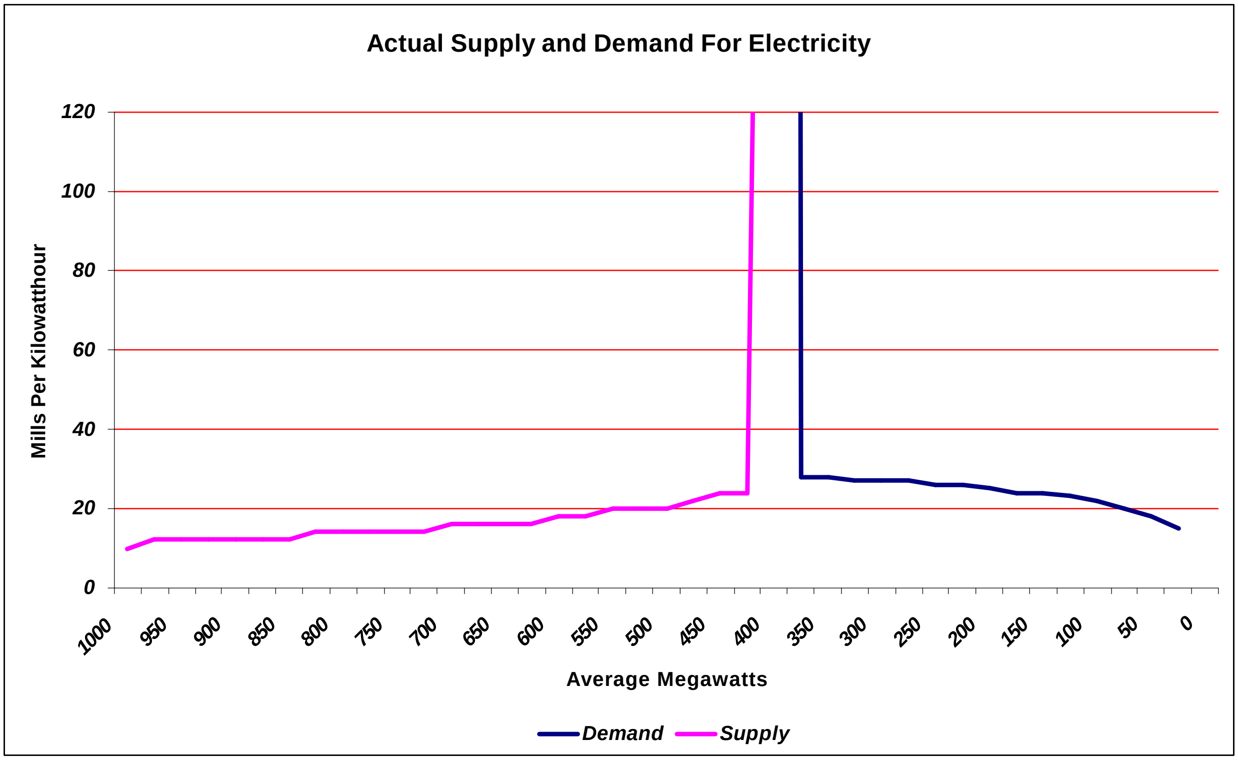

This is the case where traditional economic theory may not provide a solution. Such boundary conditions are common in the real world. They represent conditions where traditional market economics can fail. We are all familiar with these conditions in the real world – this is why taxicabs aren't allowed to set their rates according to the desperation of the customer and police protection is furnished by society instead of a local entrepreneur. In either of these cases, circumstances exist where the supplier would simply set the price to the net worth of the buyer. Imagine bargaining with the taxi to carry you out of the bad part of town as night approaches.

Although the vertical portion of the supply curve is obvious – when the available generating capacity is fully committed, new generation is several years away. That the demand curve approaches the vertical is less common knowledge. Most uses of electricity are simply not elastic in the short run. Part of this is due to the difficulty of communicating prices to consumers. We don't know how to tell the consumer that his hot shower will cost $100 today, instead of $.25. Part of its also reflects economic rationality. Industry's cost for interrupting operations is the value of the lost output divided by the energy saved.

For energy intensive operations like aluminum or chlorine, the cost might only be three of four times the going electric rate. Simple arithmetic can derive the value for aluminum – 100 mills – or steel – 300 mills. Less obvious is the value for wheel turning loads in heavy manufacturing -3,000 mills – or commercial – 5,000 to 10,000 mills. Simply stated, voluntary curtailment is difficult to achieve.

Another way to look at this is that the terms of trade move so strongly against the consumer that society changes the rules. This is why we set taxi fares, have the government provide police service, and regulate utility rates.

The shift in the terms of trade is a qualitative change – not differentiable or continuous at the boundary – with very different characteristics. My suspicion is that we won't observe many prices at the boundary because political solutions will dominate. In terms of risk, this political dimension supports continuations of capacity markets as opposed to an "energy only" solution.

3. Capacity Markets

The third "cobweb" to be removed is the very strong hidden assumption about the structure of energy markets in the future. An increasing number of players in the electric industry have come to believe that the one-hundred year tradition of pricing both energy and capacity will soon be replaced by energy-only prices. The phrase "monomic" means that only one price will characterize electric trades.

While this idea has a seductive simplicity, our research indicates that it is unlikely to come to pass. The same simplicity that makes it attractive obscures serious operating and economic issues that would make monomic pricing a very expensive set of prices indeed. This debate wouldn't be significant – after all, buyers and sellers can choose what ever definitions of price they choose in an open market – except that some theorists actually propose implementing central energy markets where the only permissible prices are based on energy.

A good first step is to set out the terminology. Traditional electric markets contain prices for both energy and capacity. Energy prices pertain to the actual kilowatt-hours. Capacity prices pertain to the right to take energy. Purchases from a thermal unit usually include prices that will cover fixed costs (capacity) and payments to cover fuel (energy.)

The gist of the energy-only pricing school is that the volatility of spot prices is an adequate replacement for capacity. Their argument is that capacity will be implicitly guaranteed by spot energy markets in two ways:

> In the short run, energy prices will climb to the point where users will curtail their needs and release capacity for those customers without the ability to curtail their demand.
>
> In the long run, the probability of high price periods in the spot markets will create an opportunity for entrepreneurs to build new capacity.

Modeling of a monomic future critically depends on curtailment charges. These charges usually reflect the price that causes a customer to curtail their load <u>or</u> liquidate the damages to the customer of an outage.

Although the models used to predict the future in this way can be complex, the relationship between frequency of curtailments, curtailment cost, and alternative resources is very simple.

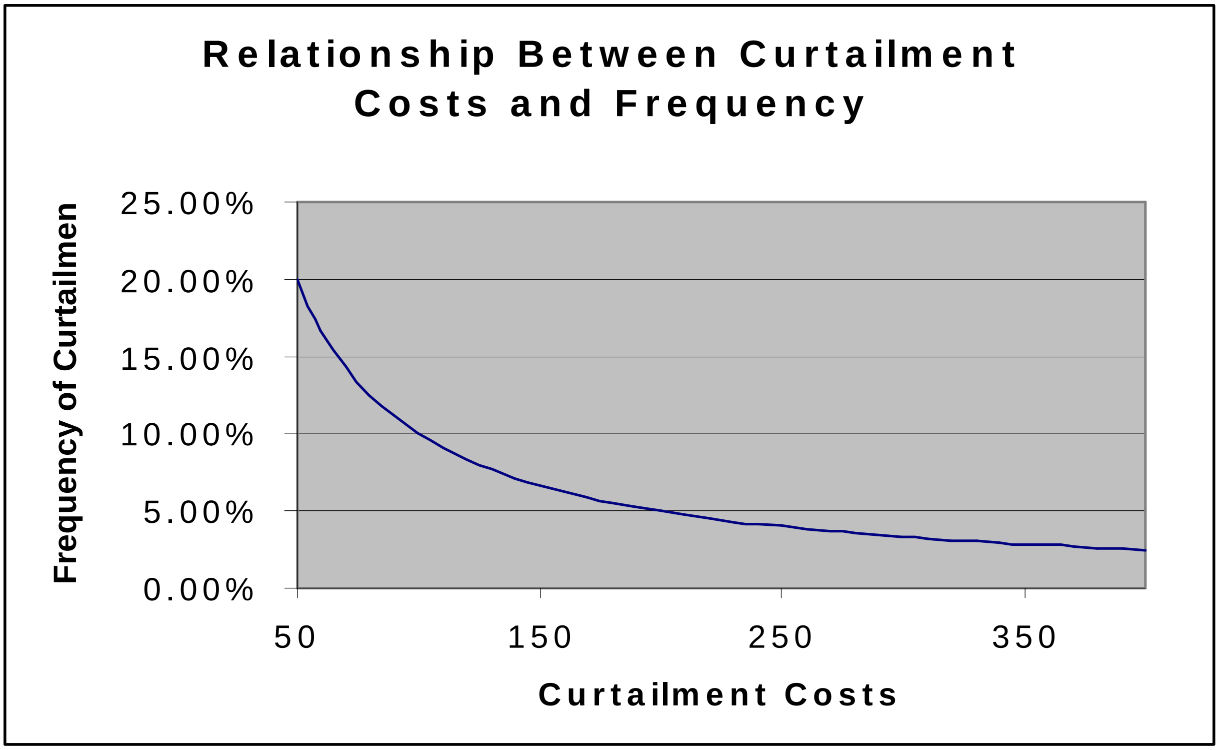

What do the curtailment values represent? Our extensive experience with end users indicate that curtailment is seldom a very good option. Although many industries currently have accepted interruptible rates, the fact is that most of these industries never actually expect to be interrupted. Recent growth of contracts including interruptibility has been coincident with increasing levels of surplus in most of the U.S. and Canada. The only major use of interruptibility in the last five years lead to lawsuits and the elimination of the interruptible clauses in subsequent contracts.

In the short run, curtailment costs strongly reflect the share of electricity in total costs.

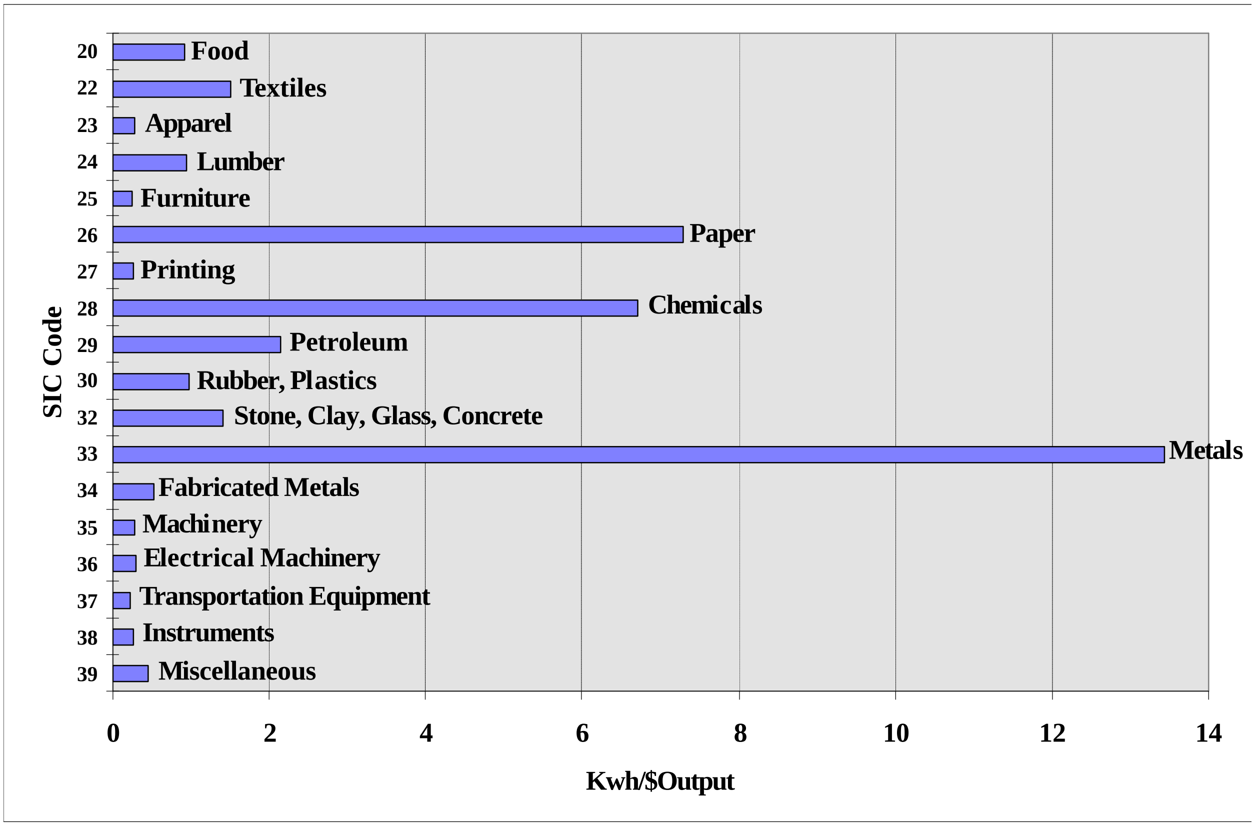

The chart above shows the energy intensity of different major industries. Actual curtailment costs for metals can be relatively low – 90 mills. Curtailment costs for food and lumber climb to 1000 mills and curtailment costs for the vast majority of industry range from 3000 to 5000 mills per kilowatt-hour.

The credibility of monomic pricing depends on the depth of the curtailment as well as its frequency. The following charts shows the scale (percentage of the market curtailed) for the West Coast of the U.S. and Canada for 1996, assuming that the regions started in load/resource balance.

Other implications of monomic pricing are equally surprising. By definition, if the only incentive to build base load prices comes from spot electric markets, the price of spot must quickly increase to the fully allocated cost of a combustion turbine and stay at that level (on an annual average basis) in perpetuity.

By comparison, historical spot market prices for electricity have never approximated the level of fully allocated costs of new resources. The Bonneville Power Administration, one the continent's largest utilities, maintains data on both spot sales revenues and the fully allocated costs of new resources.

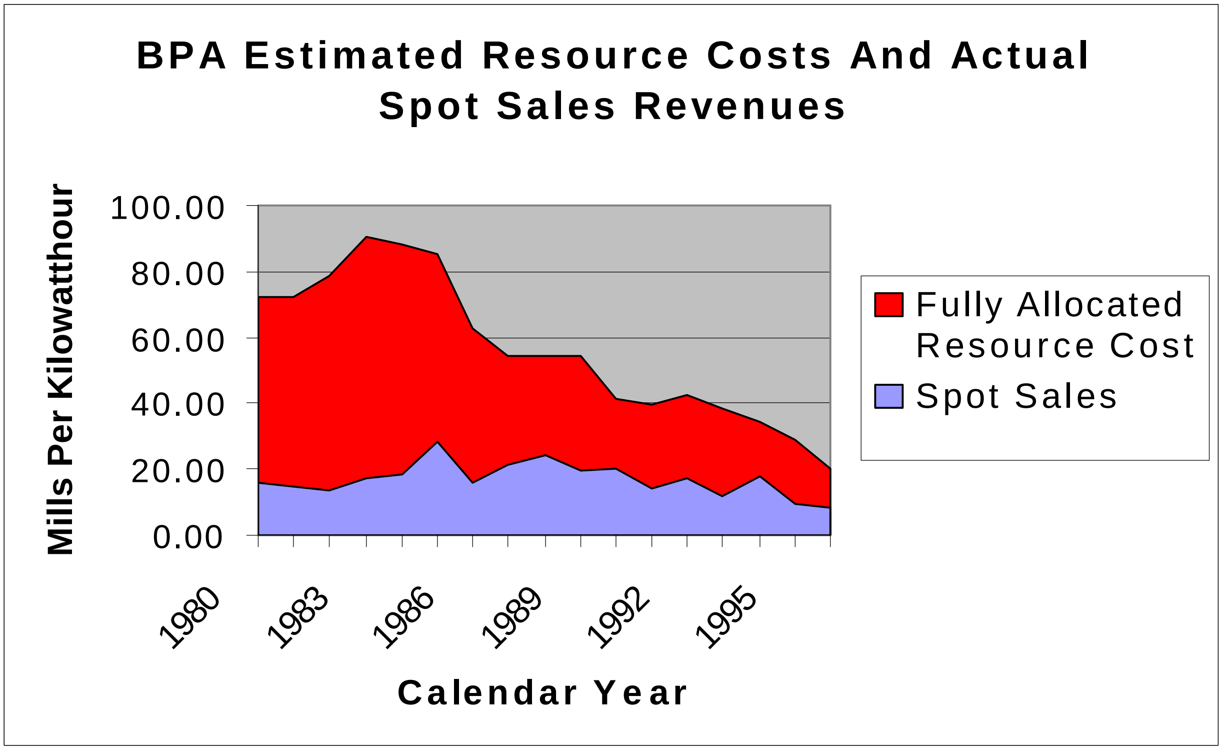

The falling level of fully allocated costs reflects the shift from nuclear to natural gas based units. BPA's spot sales reflect market forces since it is allowed to sell into the wholesale market at prices up to its average costs.

Devotees of monomic prices believe that the two lines shown above will converge in the twenty-five to forty mill range. The value for the spot market is completely determined by the requirement that average spot rates climb steeply enough to cover the fully allocated cost of new generation.

How realistic is this forecasted dramatic shift in spot prices? After a little thought it is obvious that the primary concept isn't very accurate at all. By assumptions, entrepreneurs will forecast the high prices caused by curtailment costs and construct baseload resources in response. Wiser entrepreneurs would certainly preempt the activities of the baseload entrepreneurs by constructing simple cycle turbines. Since simple cycle turbines are a vastly more efficient choice for serving short duration high cost periods (curtailments), the simple cycle turbines would become the resource option of choice.

Logically, all load growth would be served by simple cycle turbines (and other inexpensive peakers) until the percentage of the time where simple cycle units were operating on the margin was sufficient to pay for combined cycle units.

Using standard assumptions, simple cycle units are less expensive to operate than combined cycle units for loads that occur as much as forty percent of the time.[6] One absurd result would be for the variable costs of simple cycle units to define marginal production costs for much of the year. If this were the case, the price might never increase high enough to build combined cycle units.

Several recent "energy only" analyses assume that generators are the active parties in the electric markets and consumers are simply passive recipients of market prices. This tends to suggest that the engineering orientation of the industry still has enormous influence. The truth is that consumer's preferences have an enormous impact of market prices and the cost of generation. In a fully deregulated market, consumers always have the choice to find their own solutions regardless of the preconceptions of the suppliers.

The following represents a simple example of what utility planners call a resource scoping curve.

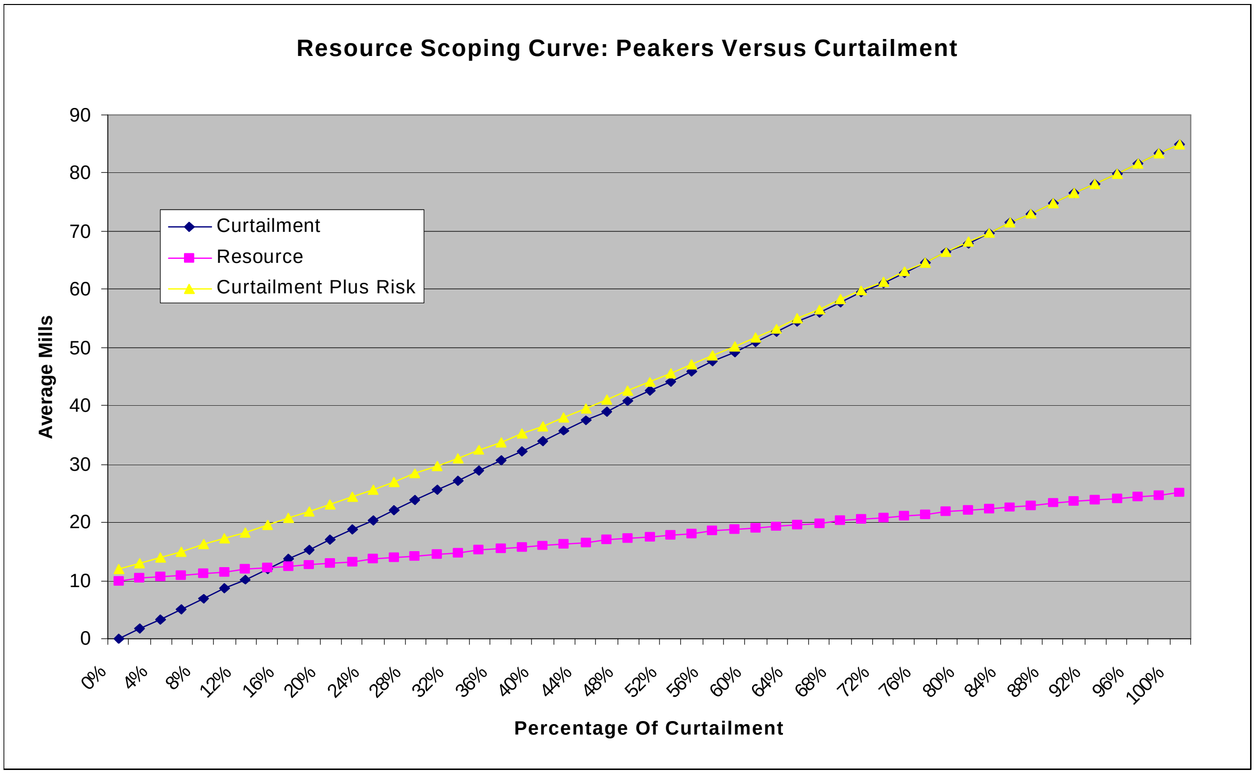

These curves are a standard industry tool that identifies the optimal operating range for each type of resource. Resources can be plotted with expected operations from 0% to 100% of the year. The resource with the lowest cost will be the best choice for loads whose durations match the percentages along the bottom of the chart.

[6]For example, if combined cycle units cost $500/kilowatt with have a heat rate of 7000 btu/kWh and simple cycle units cost $300/kilowatt with a heat rate of 10,000 btu/kWh, simple cycle units will be a more efficient choice to meet loads with durations of 45.6% of the year and less. Combined cycle units would be more efficiently constructed to serve loads of 54.4% duration and more.

For example, in this chart comparing the cost of curtailment at low percentages with the cost of a peaker, it is clear that curtailment – whose cost line goes through the origin – is always the best choice for loads with little annual duration. In fact, since the cost line for a peaker always starts out at the capital cost of the peaker, it is a mathematical certainty that curtailment is always the best option – regardless of the curtailment cost.

Given this effect, why is it that utilities (and their customers) haven't planned for substantial degrees of curtailment? The simplest answer is also a bit facile: utilities don't plan for curtailment because rigid operating rules promulgated by regulators don't allow them to. This answer is facile because rules denying curtailment as a resource choice aren't an international law, or even a Federal regulation. Operating standards promulgated in Canadian provinces, U.S. cities, rural cooperatives, Federal Marketing Agencies, and the state regulatory bodies of fifty states have all set operating rules at the near zero curtailment level. The simple argument requires that you believe that all of these parties have made the same mistake – not once but annually over the past hundred years.

A more insightful answer is that governments, utilities, cooperatives, and regulators have all perceived that the customer is not indifferent to risk. Electric customers are risk averse. Their risk aversion is reflected though their utilities, coops, elected representatives, and regulatory bodies.

Estimating risk aversion is never easy. The chart above contains a second resource cost line for curtailment. The addition to cost for risk is high for small levels of curtailment and diminishes rapidly as curtailment becomes more frequent. This is not as counter-intuitive as it seems. A completely passive consumer would not respond to additional interruptions. A rational customer would quickly determine ways to respond. A utility planning 20% curtailments would find that the customers had provided themselves with emergency backup equipment and alternative fuels.

The level of risk aversion is clearly very high. One tool of modern economists, "revealed preference", uses customers actual choices to reveal customers' degree of risk preference. We know from universal experience over the past one hundred years that the degree of risk aversion is greater than the cost of a simple cycle turbine simply because all of the choices made in all of the jurisdictions over all of the years have never revealed a preference for curtailment.

4. Fuel Dependency

The final "cobweb" is both the simplest and the most difficult one to correct. At a deep philosophic level many market participants believe that since some electric plants burn natural gas, the price of natural gas determines the price of electricity. After all, the price of eggs determines the price of omelettes, right?

This assumption makes the mathematics so easy that the fact that it is wrong is often overlooked. The easier part of the fallacy is that the link between fuels is only present when gas is the marginal fuel:

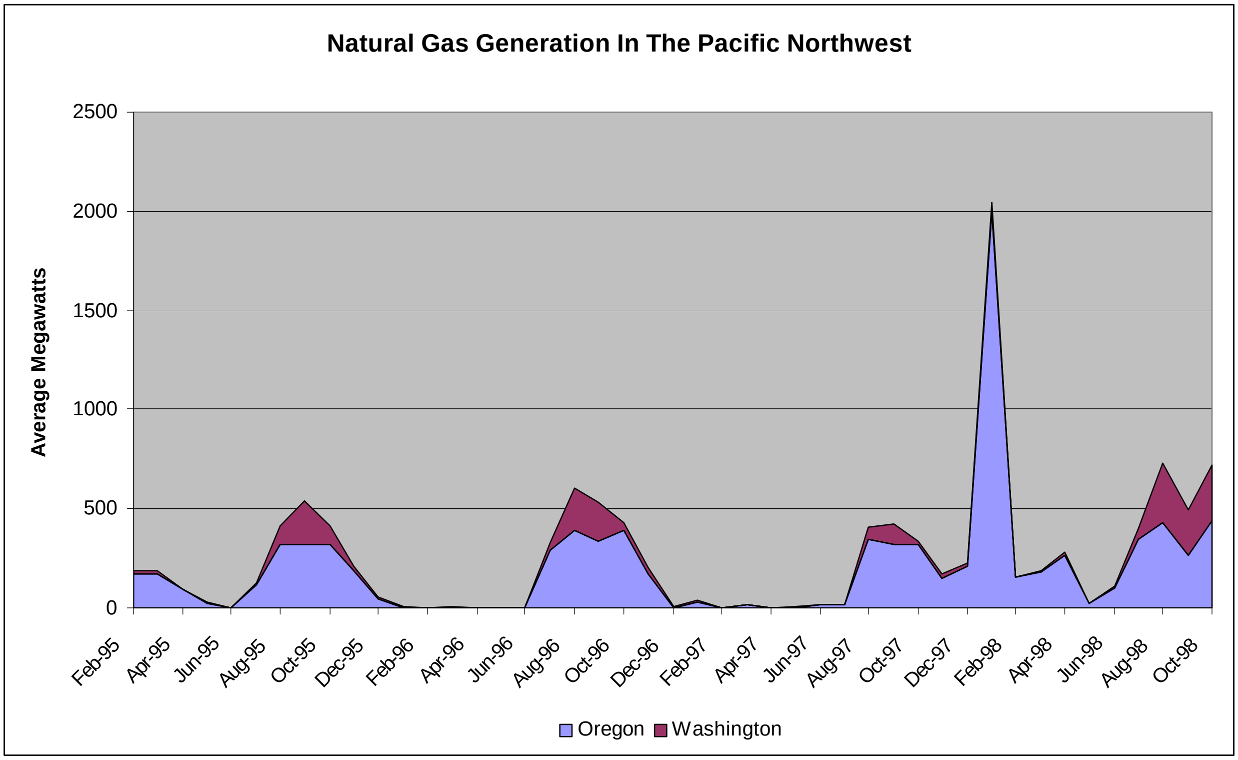

The correlation between Pacific Northwest prices and gas has always been poor. The reason tends to reflect the uneven use of gas north of the California border. For example, we would expect a very poor correlation in months where no gas was actually used.

The shift from one marginal fuel choice to the next means that natural gas will work well as an explanation in some periods and be completely irrelevant in others.

The problem goes deeper than that, however. As any farmer will tell you, cheap omelettes make cheap eggs as well. Nationally the generation of electricity is approximately 10% of the generation mix. It is also significant to note that gas purchased for electric generation is 10% of the total gas market. Since electricity is a larger market than gas, we can expect the shares to diverge given current planning – it is logical to forecast that electricity will become a larger part of the gas market than gas is of the electric market.

One of our utility clients is the largest user of gas in their area. When they make an operating decision the price of gas responds immediately. The real question is: Which commodity is driving which?

The two markets have very different fundamentals. Gas, once developed, is inexpensive to produce and can be easily stored. Electricity is expensive to produce and cannot be stored. Electric demand is inelastic. Natural gas demand is quite elastic with several competing fuels for both industry and electric generation.

Historical data isn't much use at this point. Assuming natural gas prices were fixed historically and then looking at electric prices simply assumes that we know the direction of causality. The correct approach has to reflect the causes of natural gas price fluctuations simultaneously with electric price fluctuations. As a matter of statistical practice, use of natural gas prices in the statistics is inferior to the use of a simultaneous forecast of natural gas prices – although the mathematics of this type of forecasting becomes involved. Colloquially put, natural gas prices are not a fundamental determinant of electric prices. The fundamentals of natural gas prices are fundamentals of electric prices.

Brushing The Cobwebs Aside

Estimates of electric market risk is primarily a forecasting problem. We know that we can explain approximately 80% to 90% of Mid-Columbia prices.[7] Residual risk – the price changes not explainable by fundamentals – does begin to look more like a normal distribution once the clearly unique distributions of the underlying data have been accounted for.

The first step is simply to introduce fundamentals into the forecasting. The use of a sine curve to represent seasonality is a very bad idea. Use of actual seasonal data is a good idea.

The second problem, Terms of Trade, is quite a bit more difficult. To most market participants, the actions of the PX and the ISO seem like noise. Without a detailed understanding of the beliefs of the market makers in California, their actions appear completely random. Frankly, we knew more about electric markets before April 1st, 1998. Before that date forecasts were better and intuitions about market forces tend to match results. Sam Van Vactor's market survey, Energy Market Report, contained an interesting aside in its narrative a month or so ago.

> Place Your Bets
>
> Western power prices for Tuesday shot sharply higher against a backdrop of limited Northwest supplies, unit outages, hot Southwest weather, and strong CalPX averages. Sure there are a lot of bullish factors out there right now, but not enough to warrant prices at these levels, said one marketer. When asked why prices reached the levels they did, most players were at a loss to explain, except for one trader who sheepishly said, More buyers than sellers?! Several market pundits were anticipating a market correction on Tuesday, but most indicated that they were not willing to put money on it.[8]

Our work in this area has lead to simply using a dummy variable for the PX/ISO price changes.

---

[7]Anyone can explain 100% of the variance of anything by simply exploiting the mathematical property that sufficient explanatory variables can match every data point. It isn't unusual to come across forecasters attempting to prove this homily. As a practical matter, our 80% to 90% model reflects good operating practice – theoretical argument backed by significant estimators.

[8]Energy Market Report -- April 20, 1999.

Values of the dummy evolve continuously, by have tended to show up in the 7 mill on-peak, 3 mill off-peak range.

Suffice it to say that we don't like this approach much. Its naivete is its only defense – since we are not ascribing structure where none exists. On the other hand, it is like avoiding a student driver's bad driving decisions by simply avoiding streets where he (or she) might appear. At the moment, the California market authorities simply are market risk.

The third problem is simple to solve, but requires an act of faith. This is not a situation where the right answer falls from the sky. If the Harvard Energy Project gets its way, we should outlaw capacity markets and pursue risky and less efficient solutions. In their defense, I suspect that they simply don't understand the problem well, and there proposals haven't been well fleshed out. If we eliminate capacity markets, market risk becomes enormous. In effect, eliminating capacity markets is synonymous with outlawing the insurance value that capacity provides.

We have the belief that preferences shown over many years and many locations will prevail and that capacity markets will continue to function efficiently. For those of you with experience in the Midwest, capacity is now at a premium.

Market risk estimates under the two different regimes differ enormously. Simply put, monomic markets are expensive. Volatility estimates in monomic markets can be twice those in traditional energy markets.

The fourth problem is easy. Do not accept market risk estimates from anyone who tells you that he knows that gas drives electric prices. For that matter, flee from charlatans who tell you that electric prices drive gas. Both groups should be institutionalized for their own good.

As a matter of statistical practice, if you see gas prices in the mix, remove it. It simply doesn't belong in the answer. Its fundamentals do, of course, so the answer is to proceed with caution.

The Bottom Line

Theories are constraining on the underlying facts. We all knew that. Aggressive and well funded state deregulatory initiatives do disrupt smoothly functioning markets. They do add market risk. This is one of the prices we pay for living in a free and dynamic society.

One of my college professors once said that the history of the U.S. economy was an unrelenting war on the economy by the government. He noted that every major decision ever made in Washington D.C. – from declaring war on our only trading partner in 1812 to subsidizing the purchase of OPEC oil during the oil embargo – was clearly wrong.

With deregulatory theories popping out from every nook and cranny, market risk is going to increase as institutional factors add noise to the markets. As Betty Davis put it "hold on to your seats, it is going to a bumpy ride."