FORTNIGHTLY'S

SPARK

Letter #59
Nov. 2008

THE ON-LINE GATEWAY FOR READERS OF PUBLIC UTILITIES FORTNIGHTLY MAGAZINE.

Risk management is of great interest to utility companies in these precarious times. Based on past chaotic periods in electricity markets, Joel C. Gibbons, president of Logistic Research & Trading Co. draws helpful conclusions on the costs and benefits of risk-management programs. He looks at the role of power options and how they contribute to risk control. Also, he examines the complications introduced by the chaotic nature of prices and how to apply statistical models. Insight on how chaos hits markets is even more valuable in this troubled financial period.

L. A. Burkhart
Editor

RISK MANAGEMENT

The Wholesale Power Market

By Joel Clarke Gibbons

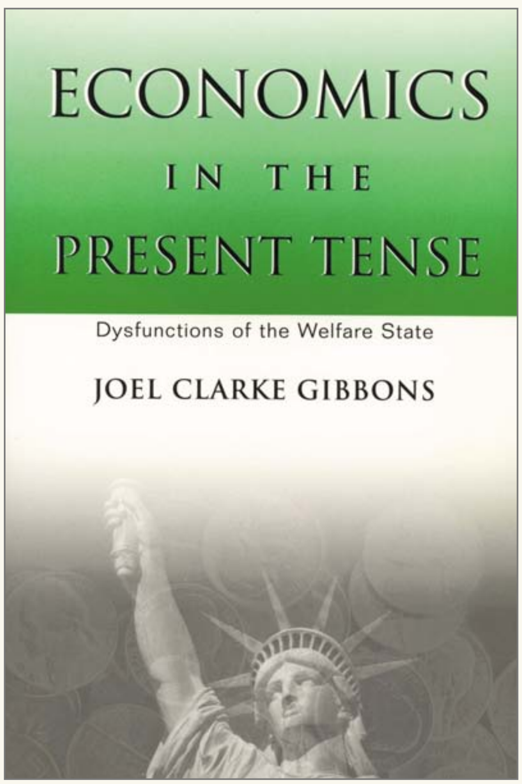

What can we learn about costs and benefits, from the last few years?

The industry has been marred by a record of crises and business failures. FERC on July 24, 2003 presented a bold blueprint for the future that promoted transparency and restrained risk, but was no panacea. Risk management is a permanent feature of management of power providers and power users. We can draw some conclusions about costs and benefits of a risk management program from available data on the power markets.

This essay explains the results of our historical study of wholesale power markets. Two questions in particular come to the fore: How would power options have performed in history, and what can we say about the distribution of spot prices? We will focus especially on those questions. (This essay has been extracted from a much »

PRINT

broader study of trading in electric power. Copies of the entire piece are available on request.)

Since deregulation of electric utilities and power transmission became a reality on a wide scale, the industry and consumers have had to cope with periodic crises in which the cost of power has risen explosively for brief periods of time. There are examples, probably familiar – not to say painfully familiar – to most who work in the industry, in which the spot price of power temporarily reached levels of ten or even a hundred times the usual cost of base-load power. These crises were not merely expensive in dollars. During the height of each one, essential pricing relationships between subregions failed, the physical availability of power became problematic, and consequently the risks that producers and consumers were confronting were too great even to quantify. In brief, crisis became panic. The resulting price histories exhibit in many ways the behavior of financial markets when panic sets in. The methods of modern finance that have been developed to get people through financial crises are also applicable to power, as we will endeavor to explain here. We need, however, to pay attention to the actual causes of crises, because they influence the flow of events. Unlike financial crises, power crises have their roots in the physical realities of the production and delivery of Watt-hours of power. The techniques of modern finance have a great deal to contribute to risk management, but to apply them successfully requires that they be applied in the unique context of electric power.

At almost any place in the chain that starts with fuel acquisition and ends at the wall outlet, there is potential for disruption. In practice, however, the chief villain has been failure of the interregional transmission grid. Everyone in the utility industry is familiar with Transmission Line Load Relief (TLR) orders. The aggregate statistics on the frequency are, or should be, disturbing. The various subregions differ widely in the incidence of these TLRs. Taking the years 2000 and 2001, the subregion containing Cinergy, EMSC, reported 652 load reliefs. The Mid-America Connected Network (MAIN, an electric regional reliability council of NERC that ceased operation), which included the city of Chicago, reported 773 relief orders. A simple aggregate count like this obscures important differences in severity, but nonetheless it is hard to ignore a record that amounts to more than one each day! (In fact, the extremely poor technical performance of MAIN Inter as measured by frequency of Load Reliefs, may explain the demise of that gateway.) In many subregions the physical facilities needed to maintain the integrity of the power grid appear to be less than satisfactory. FERC wisely addressed that topic in its final rulemaking. [*See Standardization of Generator Interconnection Agreements and Procedures, Docket No. RM02-1-000, Order No. 2003, 104 FERC ¶ 61,103, 18 CFR Part 35, July 24, 2003 (FERC)*]

The root of the problem is that »

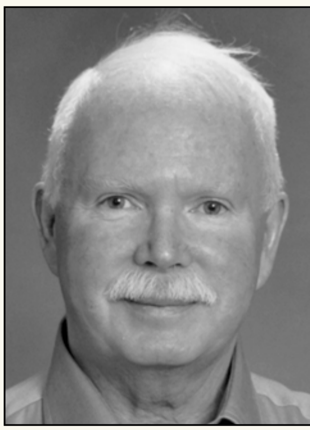

There is one clear culprit behind a breakdown of the price correlation between subregions, and that is congestion or outright failure of the interconnect grid.

—Joel Clarke Gibbons

FORTNIGHTLY'S
SPARK

Public Utilities Reports, Inc.
8229 Boone Blvd., Suite 400
Vienna, VA 22182
Phone: 703–847–7720
800–368–5001
Fax: 703–847–0683
http://www.pur.com

Bruce Radford, Publisher
radford@pur.com
Lori A. Burkhart, Editor
lab@pur.com
Horia Stefanescu, Designer
ahstephen@pur.com

E-mailed to all *Fortnightly* subscribers.
Call: 800–368–5001

For e–mail address changes or other information, contact jcole@pur.com or 800–368–5001.

the major utilities that are responsible for providing interconnection simply have underinvested in this vital resource. The strength of the FERC rules is to create a free market in transmission and interconnection services so that the law of competitive supply can elicit adequate investment. As the rulemaking took effect and reshaped the industry, the crises were expected to become less common and less severe. We can of course only look ahead in a limited way because we are largely confined to rethinking the past and its data. With this proviso, and its implication that any analysis needs constant updating to maintain its relevance and accuracy, we can begin.

Crisis and Panic

When we look a little closely at what actually transpires in a crisis we gain some important insight into both the cause and the dynamics of wholesale power prices. We will not delve into a broad discussion of power crises because there is only one key point that we need to observe. In normal times, prices of power between neighboring subregions are highly correlated. Consider the pair consisting of SP15 and Palo Verde: Southern California and Arizona, respectively. This example is used because unlike the interconnection between northern California and Mid Columbia, this particular interface was not considered to be a cause of problems or crises. In normal times the price basis between them is almost negligible, and it is clear that they inhabit a single market for wholesale power. In a crisis, not only are prices high, but they only loosely correlated between subregions (*see Figures 1 and 2*). Figure 1 shows daily price pairs in SP15 and Palo Verde over the period from March 15, 1999 to April 30, 2002.[1]

Most of the time, the price of wholesale power was less than $35 in both subregions and nearly was equal between them. (*See Figure 1.*) The right side within this chart shows, however, the relationship between the two price series attenuates. The loss of correlation may not show as clearly as it should in this chart, because the dominant regime of high correlation trains the naked eye to see correlation. Just by extracting those days of extreme prices – the upper right hand portion of this graph – the loss of correlation becomes more obvious.

The correlation of wholesale price for these crisis days is not zero of course. It is about 50%, as contrasted with a correlation of essentially 100% for the full sample. Nonetheless, the pure randomness of the individual prices – their capacity to change independently of each other – contributes not only to the Crisis, but to a deeper sense of Panic because familiar relationships that decision makers rely on daily no longer seem to be reliable at

Fig. 1 **Scatter Plot: Same Day Spot Price of Wholesale Power**

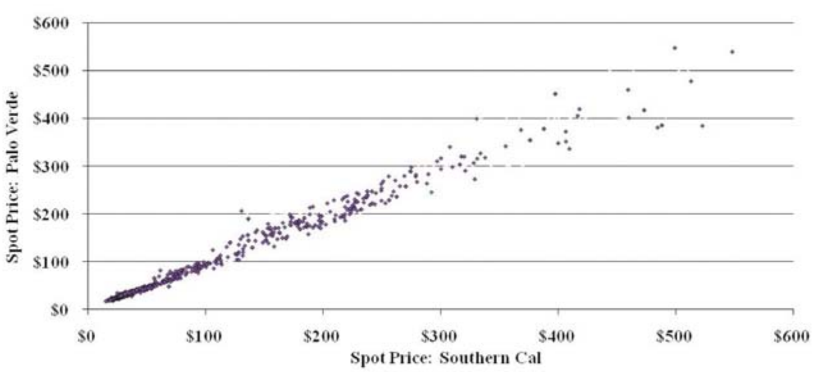

all. Crisis is the event that demands prompt and decisive action, but panic is the condition that makes decisiveness very difficult to achieve. (*See Figure 2.*)

There is one clear culprit behind a breakdown of the price correlation between subregions, and that is congestion or outright failure of the interconnect grid. Here again we are brought to the point that the FERC rulemaking addressed: The need to invest in the grid.

This history has important implications for risk management. First of all, it points clearly to an area of pressing need of risk control. FERC and utility managements address the full range of risk-control measures, but that is much too broad an agenda for our purposes. We need only to draw one inference, which is that in a crisis, price become chaotic. It is difficult, if not impossible, at that time to make rational decisions regarding prices and trading decisions. The sorry example of the State of California, which seems to have panicked and locked in very high rates well into the future, provides a graphic illustration of this. It does not pay to try to make trading decisions in a crisis. What are needed are sound risk-control and risk-measurement regimes that have been put in place beforehand. The chaotic nature of price introduces a further complication, which is that familiar and accepted statistical models » cannot be applied blindly. Better statistical methods are needed. We will come back to this point.

Power Options and Forward Contracts

It is possible to purchase wholesale power for the coming month at a preset price negotiated between the buyer and the seller. The terms of the contract typically specify the amount of power to be delivered each day of the month and the price. In the case of a few subregions there exists the alternative of market-traded futures contracts. The advantage of a futures contract is that it removes counterparty risk: The purchaser and the seller each deal with the exchange – the Nymex Exchange in New York – rather than with each other. The exchange is liable to both. Otherwise there is no economically significant difference between these two types of arrangements. Forward contracting is useful for many purposes but it does not really serve the needs of risk management for any entity that expects ordinarily to provide for its own load. Sometimes the purchased power will be a welcome addition, but most of the time it will simply replace power that could have been produced more profitably in-house.

It is possible but expensive to purchase power one day ahead in the spot market. That freedom does directly tackle the needs of risk management for coping with some kinds of emergencies.[2] The management of a utility can for instance offset temporary curtailments of its own native generation in this way. For crises with wide impact on the spot price however, the daily spot market provides no control of price risk. It would be better and cheaper to purchase in advance, and for a fixed term, an option to take power any day at a predetermined price. Not only is the cost of purchased power known in advance, but the need for daily trading and negotiating are eliminated. Power options fill this need by giving the purchaser the right to buy up to a fixed amount of power on any day at a fixed price. Typically, the unit of time for power options is one month.

Risk management has two complementary parts. The role of power options is to contribute to risk control. Risk control – which is of course only partial control under the best of circumstances – consists of achieving a degree of operational control in the face of conditions that face the firm. Power options actually provide a high degree of management control over price spikes in the wholesale market. The one important qualification that must be made is that they are useful only when the interconnect grid is itself operational. Most financial options are settled, or can in effect be settled, in cash. Because an electric utility is required, to the degree possible, to deliver power on demand, electric-power options must settle up in actual Watts of power. The terms of the option cannot insure that this will be possible in a crisis. As a consequence, the FERC rulemaking will have the effect of making power options more useful to the degree that it makes the performance of the grid more reliable. Actually, not only is it true that power options have an important part to play in risk control, a stronger statement could be made. Control of price risk always amounts to replicating the behavior of options, either by actually purchasing options or by adopting trading discipline that replicates the behavior of option payoffs.

The other part of risk management is risk measurement, which consists of quantifying the costs of the hazards that confront the management of the business. The essence of risk is mystery; risk is what we do not know. Still it is possible to quantify the statistical distribution of potential losses by using actual past experience in combination with statistical methods and models. The nature of the price history of recent years and unavoidable features of power options, require that we rely less on conventional models and rely more on analysis of the data directly.

All too often, quantification means in practice scaling the parameters of some pre-shrunk model. We do not »

FIG. 2 SCATTER PLOT: OUTLIERS

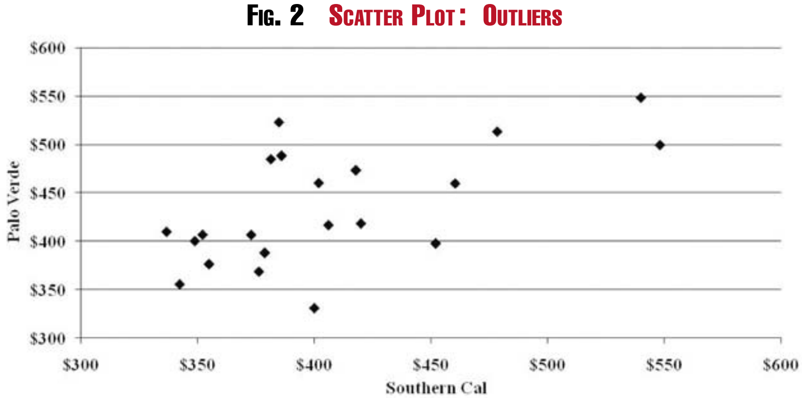

by any means reject models where we think that they are likely to succeed and to provide valuable insight. They are not, however, the heart of our approach to risk measurement. Rather, our approach is to pose important questions that arise from the actual decisions that management needs to be prepared to make, and to apply the historical data directly to them. What this means in practice will become clearer upon addressing two central issues, which both bear on the costs and benefits of power options.

The first question is: What can we say about the process that describes the Maximum price that will be observed in any given month? To what extent is the maximum predictable at the end of the preceding month – in time to purchase power options if that seems warranted – and what factors are leading indicators of the maximum? Given a best forecast of the maximum, moreover, what can we say about the statistical distribution of the actual maximum achieved in relation to the forecast?

The second question has to do directly with options. The value of options depends critically on the statistical properties of the process that generates prices over time. Because the statistical process that describes the price of wholesale power is so very chaotic, conventional models of option pricing cannot be applied. We can, however, compute the actual historical payoffs that would have been achieved by options. Each day the option payoff is either zero – in case the option is out of the money – or it is the difference between the actual price of wholesale power and the strike price of the option. Once we specify a rule that dictates the choice each month of an option strike, it is a simple chore to recreate the historical payoffs. The costs and benefits of power options can then be judged on the basis of how they would actually have performed, free of any dubious assumptions about stochastic processes.

Study of the Maximum Price

The business risk associated with price spikes is directly related to the highest price. Accordingly, no single bit of information is as useful as the distribution of the maximum price that can occur within the horizon of the decisions that must be made. We will simply take that horizon to be one month. What have we discovered about the statistical behavior of the maximum price?

Before discussing the predictive model, it is instructive to summarize the behavior of the residuals – the unsystematic component of the maximum price – from a regression model of the maximum. The chaotic natures of price shocks can be illustrated (*see Figure 3*). It is a histogram of the residuals from a simple regression in logarithms of the maximum price within a month. The right-hand side variables are also logarithms; the log of the average price within the same month, the average price the preceding month, and the maximum price the preceding month. These variables were chosen for this test to help understand the relationship that exists between the maximum that occurs in a month and the average price for the same month. It is not a predictive model of course, since it includes the average spot price within the month. Since the contemporaneous average price has, understandably, great explanatory power for the maximum, the fitted residuals from this regression model are as "tame" as they could be. The data that underlies this study consists of all months for which full data exists, covering all power subregions. There were 568 monthly observations in all, of which, for example, fifty-eight come from TVA and seventy-five come from the Mid-Columbia subregion (in the Northwest electric market).

Even when using information about the average price in the same month, and even when taking logarithms of all the factors, the unexplained residuals are still highly positively skewed »

Fig. 3 Histogram of Fitted Residuals

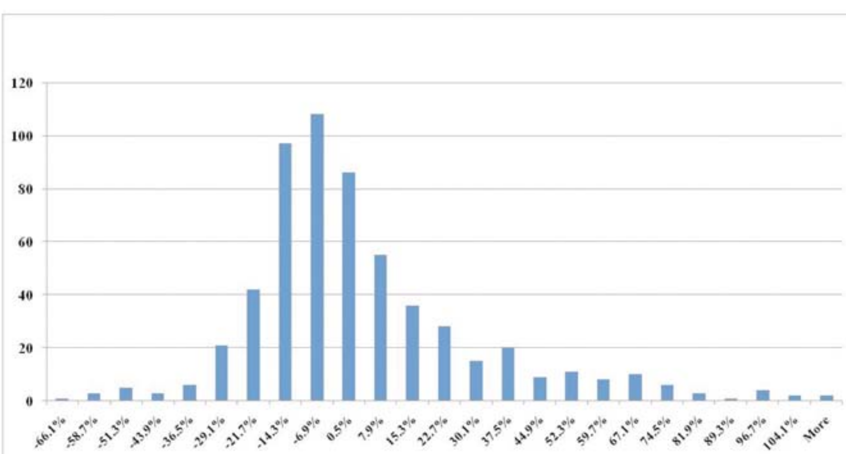

(*see Figure 3*). What this means in practice can be made clear in a simple contrast. In twenty-eight of the months, out of a total of 568 months, the maximum price was two times the fitted value from the model. On the other hand, there are no months in which the fitted value was two or more times the actual maximum price. These twenty-eight monthly observations are of course the twenty-eight most interesting – or most frightening – months in the whole sample, and yet even a model that controls for the average price each month would under predict the maximum by 100%.

We do not propose to try to "correct" the data or the models for the skewness of the model residuals. Quite the contrary, this skewness – what previously is termed the chaotic nature of the price shocks – is real. In any given month, the markets sample from a positively skewed distribution like this one. It is entirely sufficient for purposes of risk management to have quantified the residual distribution, and especially to have an estimate of the degree of skewness. The significance of this distribution of residuals, for our purposes, is that it permits us to quantify the distribution, which is precisely the distribution of price shocks that occur in times of crisis. Residual distributions like this one, therefore, are an integral part of risk measurement.

For crises with wide impact on the spot price however, the daily spot market provides no control of price risk.

The only way to deal with random innovations is through quantification of their stochastic properties. The maximum price is not entirely random, however. A component of it can be anticipated before the month has begun, and this part – the systematic component – is both a useful management tool and a source of insight into the process that causes price shocks. Two further questions can be addressed to the available data. How significant is weather in causing price spikes, and are available price forecasts useful in predicting them? The power-trading market provides the industry with forward contracting prices that are in the nature of forecasts of the average price to prevail of a given month in the future. It is an interesting question how good these forecasts are, and what is the source of their insight if any. We have looked at this issue.

Temperature

It is obvious that price shocks are to a large extent provoked by extremes of weather that generate an exceptional drain on the power grid. This is borne out in our data, when we add the average monthly temperature as an explanatory factor.[3] The average temperature makes a significant contribution to explaining the maximum price observed during the month.

Forward Prices

For the purposes of this study, the forward price was used quoted on the last day of the preceding month, which is of course the last day in which it is still purely a forecast. Since these prices should have the attributes of price forecasts, it is of some interest to see what light they shed on the maximum price. Several conclusions can be drawn. The available data on forward prices is unfortunately somewhat limited, and as a result we could investigate the relationship for only a subsample consisting of about 200 monthly observations.

The forward price contains useful information about the maximum price even when other predictive factors are held constant, such as the average price and the maximum price in the preceding month. Evidently, energy traders possess some information that is not simply an extrapolation of the recent past. When average price in the month also is held constant – *i.e.*, the average that the forward price is a forecast of – the forward price becomes insignificant. This suggests, reasonably enough, that whatever information is embedded in the forward price is information about the average price over a month, and that the forward has no additional insight specifically about the maximum price. It further strengthens the conviction that the magnitude of price spikes is in fact entirely random and unpredictable.

While the forward price adds useful information about future prices, it is not by any means an efficient forecast. That is to say, any forecast that uses the forward price can be improved upon by adding some other factors as well. The insight that they provide is not reflected in the forward price. One factor that the forward price does, however, anticipate is the mean temperature. In any » model that contains the forward price of electricity, adding the mean temperature to the model does not add anything to it. As noted previously, temperature is itself very much related to the maximum price. The point made here is that temperature adds nothing further to any model that controls for the forward price, and that is because the trading market is able to forecast the mean temperature pretty accurately.

These statistical models serve to identify that part of the maximum price that can be anticipated, and that serves as a best point forecast. They go part of the way to reducing or explaining the rather chaotic component of the maximum. Typically, a model of this sort explains between 50% and 60% of the variance in the logarithm of the maximum price. The skewed residuals remain, and as previously explained, already they can parameterize their distribution as well.

This model documents the fact that the historical data – even data as unconventional as this – yields valuable information when one takes off the blinders of predetermined models and puts on the decision maker's spectacles, through which we pose question of the history. Following in this vein, there is another way to look at this data which, if anything, bears even more directly on hedging and risk management.

Any forecast that uses the forward price can be improved upon by adding some other factors as well.

Valuing Call Options.

As observed above, the tools for managing price risk are largely an exercise in valuing options to cover the risk from outliers. In general, the contribution that quantitative research makes is two-fold: to calculate values of options and to formulate trading rules that replicate the outcomes of options. Conventional theory requires one to specify at the outset the kind of distribution that governs prices. The well-known Black-Scholes Model, for instance, starts with the assumption that the logarithm of price is normally distributed. As we know all too well, the price of wholesale power is far too erratic to succumb to any convenient model. Indeed, the times where management is most in need of reliable valuations are precisely those times when conventional models are the most inaccurate! As usual, we want to start not with assumptions, but with the history.

The behavior of prices discussed above shows clearly that conventional models do not apply to options on wholesale power. There is hardly any need to stress how much the distribution of price deviates from a normal distribution because of the random incidence of large spikes. There is another, equally important way in which familiar option models misrepresent this data. In finance theory, options are valued indirectly, by computing the payoff of a dynamic hedging strategy that replicates option outcomes. We assume that price evolves by some sort of random walk, in which the price tomorrow is related to the price today by the addition of some sort of increment – whether from a normal distribution or not – but wholesale power is anything but a random walk. It would be more accurate to say that each day is drawn from the same distribution as the previous day. The random walk assumption must be made in order to justify any hope of dynamic hedging. Random-walk price paths have the property – essential for conventional option models, that all price changes are permanent, but the price spikes that occur in electric power, by contrast, are essentially transient phenomena. If dynamic hedging is not feasible, the Black – Scholes – Merton approach to modeling options is not valid. Under these circumstances, it is meaningless to propose any dynamic hedging method, and equally meaningless to attempt to value options by that method.[4]

Using the price data, we instead constructed the history of option payoffs, by sub-region and by month. The terms of the standard call option on wholesale power are very familiar: Each single option entitles the bearer to draw one megawatt-hour of power each day of the month at a predetermined strike price. Thus, it is like a strip of daily options, one for each day of the month. The option payoffs are simply the realized value each month of such a compound option. There are many questions to be posed, but perhaps the most central one is this: Historically, have power options been more valuable in months where we would predict a high maximum price than in months » where we would predict a normal maximum?

Does our forecast of the maximum price for a given month tell us anything in advance about how valuable the option contract will turn out to be?

The answer is a very resounding affirmative. Before arriving at any numbers, one must specify the terms of the contract. Specifically, there must be some rule that chooses a strike price each month. In order to be sure that the results are not an accident of the way the choice is made, we tested various rules, but the simplest one is indicative of all. We tested a rule that fixed the option strike at a percentage spread over the expected average spot price for the coming month. Thus, for instance, if at the start of a month we forecast that the daily spot price would average \$24, and if we fix the spread at 10%, then the option outcome consists that month of the realized value of an option to buy power at \$26.40. As the expected spot price changes from month to month and from subregion to subregion, the dollar strike adjusts accordingly, but the ten percent spread rule remains the same.

The answer to the question posed above lies in the correlation between the realized value of call options and the expected maximum, as of the start of the month. It is not appropriate, however, to simply correlate them, because both are extremely skewed by those rare months of radical price spikes. To obtain a more reliable measure of the correlation, we made two adjustments: First, we related not the raw factors, but their logarithms. This reduces the importance of the few extreme months. Second, we added a constant to the option value each month, before calculating the logarithm. The exact value of the constant intercept is another statistic that we estimate. The model therefore takes the following form:

$$\ln(\text{Option Outcome} + \text{Constant}) = A + B * \ln(\text{expected Maximum}) + u.$$

The three parameters – A, B, and Constant – are estimated simultaneously. The estimated correlation is 50%, which has a t-ratio of 13.5. This correlation coefficient is somewhat inflated by price correlation across neighboring sub-regions means, because the data used are not entirely independent observations. It is appropriate to adjust – to deflate – the t-ratio. Even by taking the conservative tack to dividing the t-ratio in half, it is still extremely significant. The actual estimate of the Constant is \$3.85, and the estimate of B is 68%. A one-percent increase in the expected maximum price adds about two-thirds of a percent to the expected option outcome plus \$3.85. An example may help to clarify this calculation. In October, 1999 the expected maximum price – expected as of the end of September – was \$46.35. The following month, the expected maximum had grown to \$70.65. The expected value of an option contract of October therefore was:

Expected Option Value October =

$$\exp(-.73 + .68 * \ln(\$46.35)) - \$3.85 = \$2.70.$$

The Following month,
Expected Option Value for November =

$$\exp(-.73 + .68 * \ln(\$70.65)) - \$3.85 = \$4.61$$

The expected option value would be about \$1.89 higher each day of November, because the expected maximum price was also higher in November.

Conclusion

Quantitative data can yield a rich harvest of information that adds value to business decisions, but in order to extract the information it is necessary to view the data as evidence and the study as an exercise in uncovering statistical evidence. This is the approach that is termed the Stochastic approach. All too often studies begin with a commitment to force the data into one or another standard sort of model whose only virtue is that one would know how to interpret the results. When starting off with a particular sort of model it is nearly certain that the actual data will violate the assumptions that the model requires, and the results will actually be useless.

The place to start is with an understanding of the business problem at hand: What do we want or need to know, and what sort of evidence and estimates would help us to make better decisions? No data and no model will make all the decisions. On the contrary, our only expectation is a more modest one, which is that the evidence we uncover will be of some value. The next step is to let the data speak for itself. Does it provide any basis for useful evidence or valuable estimates? What can it tell us that will make for better decisions?

This is the way we have approached the history of the price of electric power. Our study of this topic is of course by no means finished. It is a continuing exercise in rethinking the questions we wish to pose and the methods that are most appropriate to address them. Even at this relatively early stage however we feel that what we have uncovered should be interesting to the utility industry and to power traders, and perhaps to regulators as well. ■

Joel C. Gibbons is the president of Logistic Research & Trading Co, which is a » *registered Commodity Trading Advisor that engages in trading across a wide range of commodities, and engages in state-of-the-art statistical research on a consulting basis. Mr. Gibbons holds a doctorate in business from the University of Chicago and a doctorate in mathematics from Northwestern University, and he has lectured extensively on finance and economics. He is a member of the International Association of Financial Engineers. Contact him at: 269-408-1511 or E-mail: jgibbons@logisticresearch.com.*

The article is excerpted from *Economics in the Present Tense: Dysfunctions of the Welfare State*, Vantage Press: New York, N.Y. 2008.

ENDNOTES

1. Our price data comes from the *Power Markets Week* database.
2. The daily spot market does not solve all critical problems, because understandably it is necessary to purchase power that is needed for a given day some time on the preceding day. This limitation is intrinsic in the physical realities of power generation and transmission.
3. Temperature, for this purpose, is the average daily temperature in the largest metropolitan area within each subregion.
4. This statement is a bit oversimplified. Certainly, periods of high prices are grouped together because the causes – weather and the like – are somewhat persistent phenomena. Spikes however generally persist for only a few days, after which price drops precipitously.

Next Month's FORTNIGHTLY

The December issue of *Fortnightly* magazine foregoes the usual columns for an exclusive look at generation. This issue is a must read. Editor-in-Chief Michael T. Burr gathered utility CEOs to debate the merits of a retail surcharge to fund clean-tech research and development.

Here is more of what you will find:

- **The Capture-Committed Power Plant**
  'Capture readiness' hasn't helped coal projects move forward, but a firm commitment might make the difference.
- **Memo to the President-Elect (Part 2)**
  Addressing climate change will require extending the life of today's nuclear fleet and laying the foundation for new plants.
- **Water Worries**
  Cooling water shortages might force nuclear project developers to get creative.
- **Ontario's Standard Offer**
  The province's renewable program was vastly oversubscribed. But was it successful?
- **Squeezing Energy From a Rock**
  Low-temperature closed-loop generators promise vast growth in geothermal power.
- **Riding on the Wind**
  Plug-in hybrid vehicles (PHEVs) open a new intersection between wind power and transportation.
- **Back to Gas**
  Utilities are turning to natural gas as a bridge fuel, and to support non-dispatchable renewables.
- **The New Gas Wisdom**
  Unconventional gas and LNG are changing the outlook for future gas prices.