Uncovering and pricing the hidden risks in power marketing

In competitive markets, all customers must be pursued, but all customers are not alike. Some bring more risk than reward to a marketer's overall portfolio. Profitability demands that marketers price these hidden risks appropriately

By Meredydd Rees and Richard Hooke

Many energy companies are just waking up to the intricacies of selling a commodity that may plunge a retailer into substantial losses if a given Monday turns out to be unexpectedly sunny or cold. That's risky business, indeed.

While most power marketers have learned how to use pricing to leverage brand and customer relationships and to diversify operations, many still do not fully appreciate the impact of price/volume risks on their revenues. Learning how to structure price quotes is the first step toward mitigating those risks. Size is no vaccine against price/volume risks; in fact, a large customer base often magnifies them.

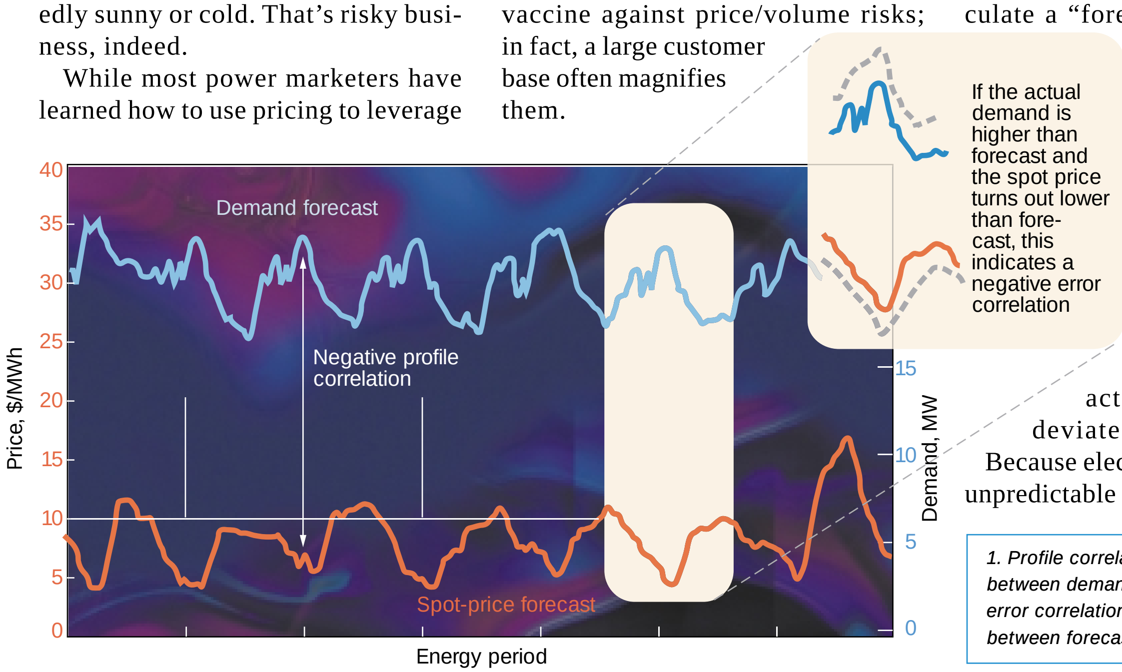

1. Profile correlation depicts the relationship between demand and price forecasts, while error correlation depicts the relationship between forecast and actual

Hidden risks of forecasting

At the heart of the problem lies the difficulty of accurately forecasting the spot price and demand for electricity for the medium term—typically between three months and two years ahead. Because the relationship between the forecast price and the forecast demand (see box)—called profile correlation—is unique for each customer, it represents the risk that customer brings to the portfolio. Since no forecast is perfect, there are bound to be differences between the forecast and the actual. Such "errors" relate to the second kind of risk, so-called error correlation.

Profile correlation for shape

As history reveals, customer demand patterns are sensitive to consumption, weather, and other events. Plotting a customer's future demand against power price forecasts on a half-hourly or hourly basis shows how well these patterns match up. If a customer consumes most of his or her power at off-peak times, a negative relationship between the spot price and the demand shape pertains, because demand is highest when prices are lowest. This "profile correlation" (Fig. 1) is well understood by power marketers and can be incorporated into their pricing strategies. For example, they can calculate a "forecast volume weighted spot-price" by multiplying out each element of half-hourly spot-price and load forecasts

Error correlation for sensitivity

A less understood relationship is the way that actual price and demand deviate from forecast levels. Because electricity cannot be stored, unpredictable weather and other unex-

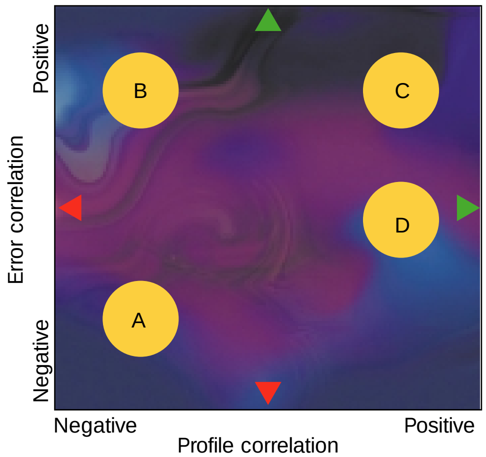

A: Billingham Salt Inc.: An integrated salt mine and chemical plant is a price-sensitive industrial consumer; it can modulate electricity consumption in response to prices and generally consumes more during off-peak hours

B: Linton Towers Hotel: A large hotel in a hot climate with a significant air conditioning load that peaks when guests are in their rooms from around 7:00 p.m. to 8:00 a.m.

C: Rocoll Insurance: A typical office building that is occupied between 9:00 a.m. and 5:00 p.m. and has significant air conditioning load

D: Blaw Diffusion Devices: A manufacturer of extruded aluminum products; most of its consumption is between 9:00 a.m. and 5:00 p.m. when electricity demand is 10 times higher than average

2. Graphic at left shows examples of hypothetical customers with negative and positive correlation

pected events can lead to sudden price spikes. A classic example is a heat wave that causes a price spike. In such a market, a power marketer caught by surprise could be bankrupted or incur substantial losses. Adding to the problem is the uncertainty of individual customers' demand; their power purchases reflect their own circumstances and external influences. As a result, three-month spot power price forecasts typically miss the mark by anywhere from 25 to 100%, while customer demand forecasts over a similar horizon are typically off by 25%.

Better forecasts only go so far

What steps can a power marketer take to mitigate the effects of all this uncertainty? One solution is to try to improve the quality of spot price and demand forecasts. But it's impossible to predict whether, when, and for how long an industrial customer might start up or shut down its production line—for example, in response to a strike or severe weather. Price spikes are equally difficult to predict. That's not to say that forecasting efforts are futile. A better characterization is that they produce diminishing returns: After some initial improvements, a brick wall is soon hit, beyond which little progress is possible.

Determining customer type and risk

For power marketers, one solution to uncertainty problems is to finesse their marketing strategy using a mathematical concept called standard error correlation. It can be explained by the following example:

Consider the impact on electricity demand of unexpectedly hot weather. All of the air conditioners in a region are going full-blast at the very moment that spot prices are peaking. Price and demand forecasts often underestimate the impact of such event. This type of customer (hotels, office buildings, and residential customers) exhibits a positive correlation between the inaccuracies in the price forecast and those in the demand forecast (Fig. 2).

Now compare retail consumer demand to that from a price-insensitive industrial company (designated Blaw Diffusion Devices) under the same external conditions. The hypothetical company is mostly automated and immune to the effects of the weather. So it continues to churn out extruded widgets as the mercury rises.

At this time, Blaw's electricity consumption remains unchanged, but if it were to shut down for unscheduled maintenance, its consumption would vary from the forecast. In general, for Blaw, there appears to be no coincidence of inaccuracies for our demand and spot price forecasts.

This example illustrates the importance of understanding whether deviation from forecast demand will occur at the same time as a deviation from the forecast price. Failure to take into account the error correlation could result in a gross margin around 5% lower than expected—a significant amount in a low-margin business like power retailing.

Pricing for the right mix of customers

Most customers, however, exhibit positive error correlation: actual demand and spot prices are either both higher or both lower than the forecast values. When the error correlation is positive, having many customers alone doesn't provide any protection against price/volume risks. In fact, if all of a retailer's customers are of the same type and in the same geographical region, the effects could be reinforced by an unexpected event that affects both overall system demand and spot price. Having some customers with negative profile and error correlation would help diversify the overall risk for the power marketer. For example, the discount in a competitive quote to a customer like Billingham Salt would be more than offset by the diversification that it brings to the portfolio.

It is to the energy company's advantage, then, to understand the likely behavior of its customer base and to maintain a good balance of customers with a low error correlation—that is, whose need for electricity is steady and doesn't vary with weather shifts and other unexpected events. Because cor-

Correlation explained

Correlation is a mathematical measure of the relationship between two data series. Its range is from –1 to +1. Zero correlation means no relationship. Negative correlation means as one goes up, the other goes down. Positive correlation means as one goes up, the other goes up. The lower the correlation, the closer to zero the number is. A high correlation is close to 1 or –1. Correlation does not indicate causality.

relation reflects the customer's sensitivity to changes, it is important to price this sensitivity accordingly. Otherwise, the customer will get a free option—as economists call it—to vary its consumption patterns substantially from forecasted. Since customers do not typically pay a premium for retail price quotes (effectively the right to buy any quantity at fixed prices), sellers have to pay for this option either by increasing their prices or giving price quotes with shorter half-lives. The latter allows sellers to reprice dynamically against the latest forward curves.

To minimize overall portfolio risk, power marketers should first characterize their customers by their profile and error correlation to quantify the risks they represent. They can then price the premium or discount offered to those customers in terms of the risk they add to the marketer's portfolio. ■

Meredydd Rees is a strategy consultant to London-based KWI (www.kwi.com). Richard Hooke is KWI's chief technology officer.

Load forecasting
in practical terms

By Phil Inje Chang

Competitive markets put a premium on information, and that information can become scarce when it's owned by market participants. A case in point is load forecasts, which utilities have perfected over many years of providing power to their service areas. Suddenly, market participants are active in areas outside of their service areas, or new entities are buying and selling power between regions, and no one has a clue what load will look like

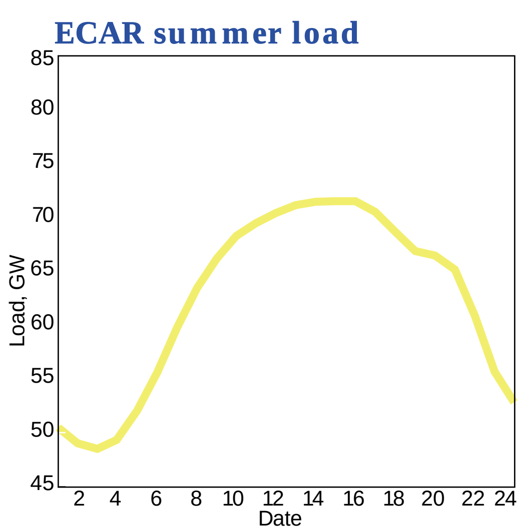

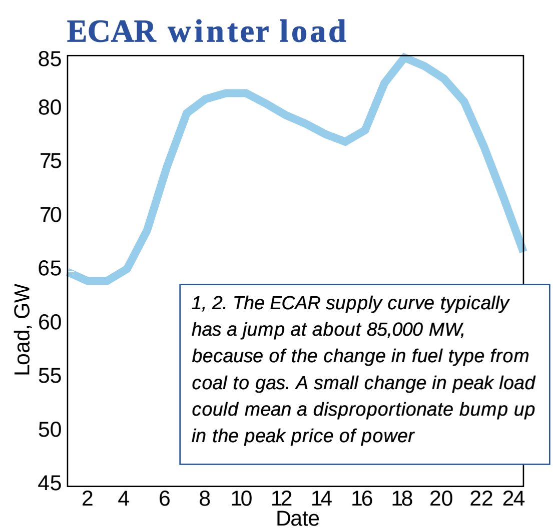

1, 2. The ECAR supply curve typically has a jump at about 85,000 MW, because of the change in fuel type from coal to gas. A small change in peak load could mean a disproportionate bump up in the peak price of power

Why is load forecasting important? Price is dictated by the market, not by load, right? The U.S. power markets have seen some cunning examples of this logic. But competitive markets are defined by supply and demand, so demand is not only important, it is also particularly tricky when it comes to power. Since power can't be stored, plants have to be on when its needed. And there is only so much power in a given region, an amount that can be reduced by forced outages and plant maintenance.

It turns out that certain rules of thumb become second nature to most successful power traders. Industrial load is relatively constant during the week. What causes large variations in load are sudden shifts in temperature, especially during the summer (Fig. 1), when the de facto solution is to turn on air conditioners. So traders are apt to hover over weather forecasts, reasoning that there is

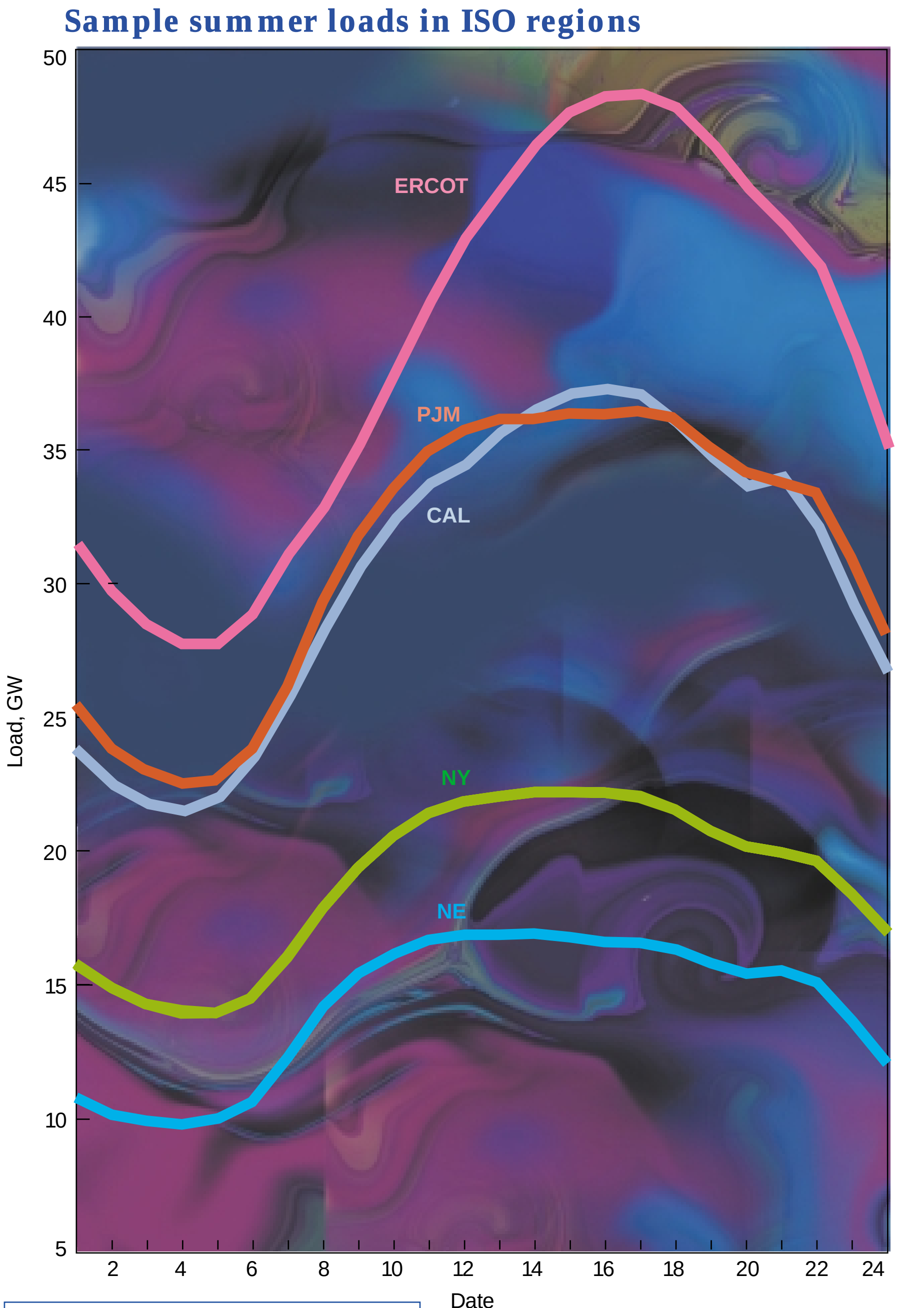

3. *In the U.S., only the five ISO region loads are made available on a timely basis*

some strong correlation between shifts in temperature and shifts in price.

This kind of logic only goes so far. There are problems. The first is the accuracy of the weather forecast. The second problem concerns some reliable estimation of load based on the temperature and other variables. The third is how load intersects supply, as the supply curve is discontinuous and linked to fuel types and prices (Fig. 2). The fourth problem is how power will flow over transmission lines. Getting a complete picture of power fundamentals is just plain difficult.

While there are solutions to getting at least parts of this picture, the point here is to address the second problem. The problem with predicting load is the scarcity of data, because the only way to get a truly accurate load forecast is to know the near-real-time load in a region. In the U.S., there are only five regions where (aggregate) loads are made available on a timely basis: the independent system operators ISOs (Fig. 3). So in those regions, it is quite possible to generate a reasonable load forecast with a decent temperature forecast, standard statistical techniques, and reasonably frequent runs. It won't be as good as what the local utility can do, but it can get traders in the ballpark.

Outside of the ISOs there is a more difficult problem: historical loads are made available on an annual basis. So in those regions, more approximate methods can get traders into a bigger ballpark, with variation caused by current temperature forecasts. Lots of considerations go into this, including seasonal patterns, similar days in the past, population changes, and so on.

Two other issues ought to be considered. Exotic techniques for predicting load often appear in technical literature, citing neural nets or other methods. A neural net is a computer forecasting technique based on learning through patterns of behavior. Neural nets are known for being insensitive to rapid changes, but besides that, most of these techniques have a high sensitivity to available data. Again, the data are not available, so what good is a fancy model that requires minute-by-minute actual loads and highly accurate wind speed forecasts—when such data are not available?

The second issue is a bit more philosophical. If the goal is to predict price, wouldn't it make sense to figure out the relationship between temperature and price? Remember, this is what traders are trying to do in the first place. Yes, this leaves load out of the picture. It turns out that the relationship is complex, a lot more complex than traders would realize intuitively. So it is still quite common for traders to take positions on a forecasted change in temperature and end up with losses. Given this, the fundamentals are decent guidelines for market behavior. The credo of successful traders ought to be: know what you can know. ■

Phil Inje Chang is vice president of business development at San Francisco-based e-Acumen, Inc. (www.e-acumen.com).