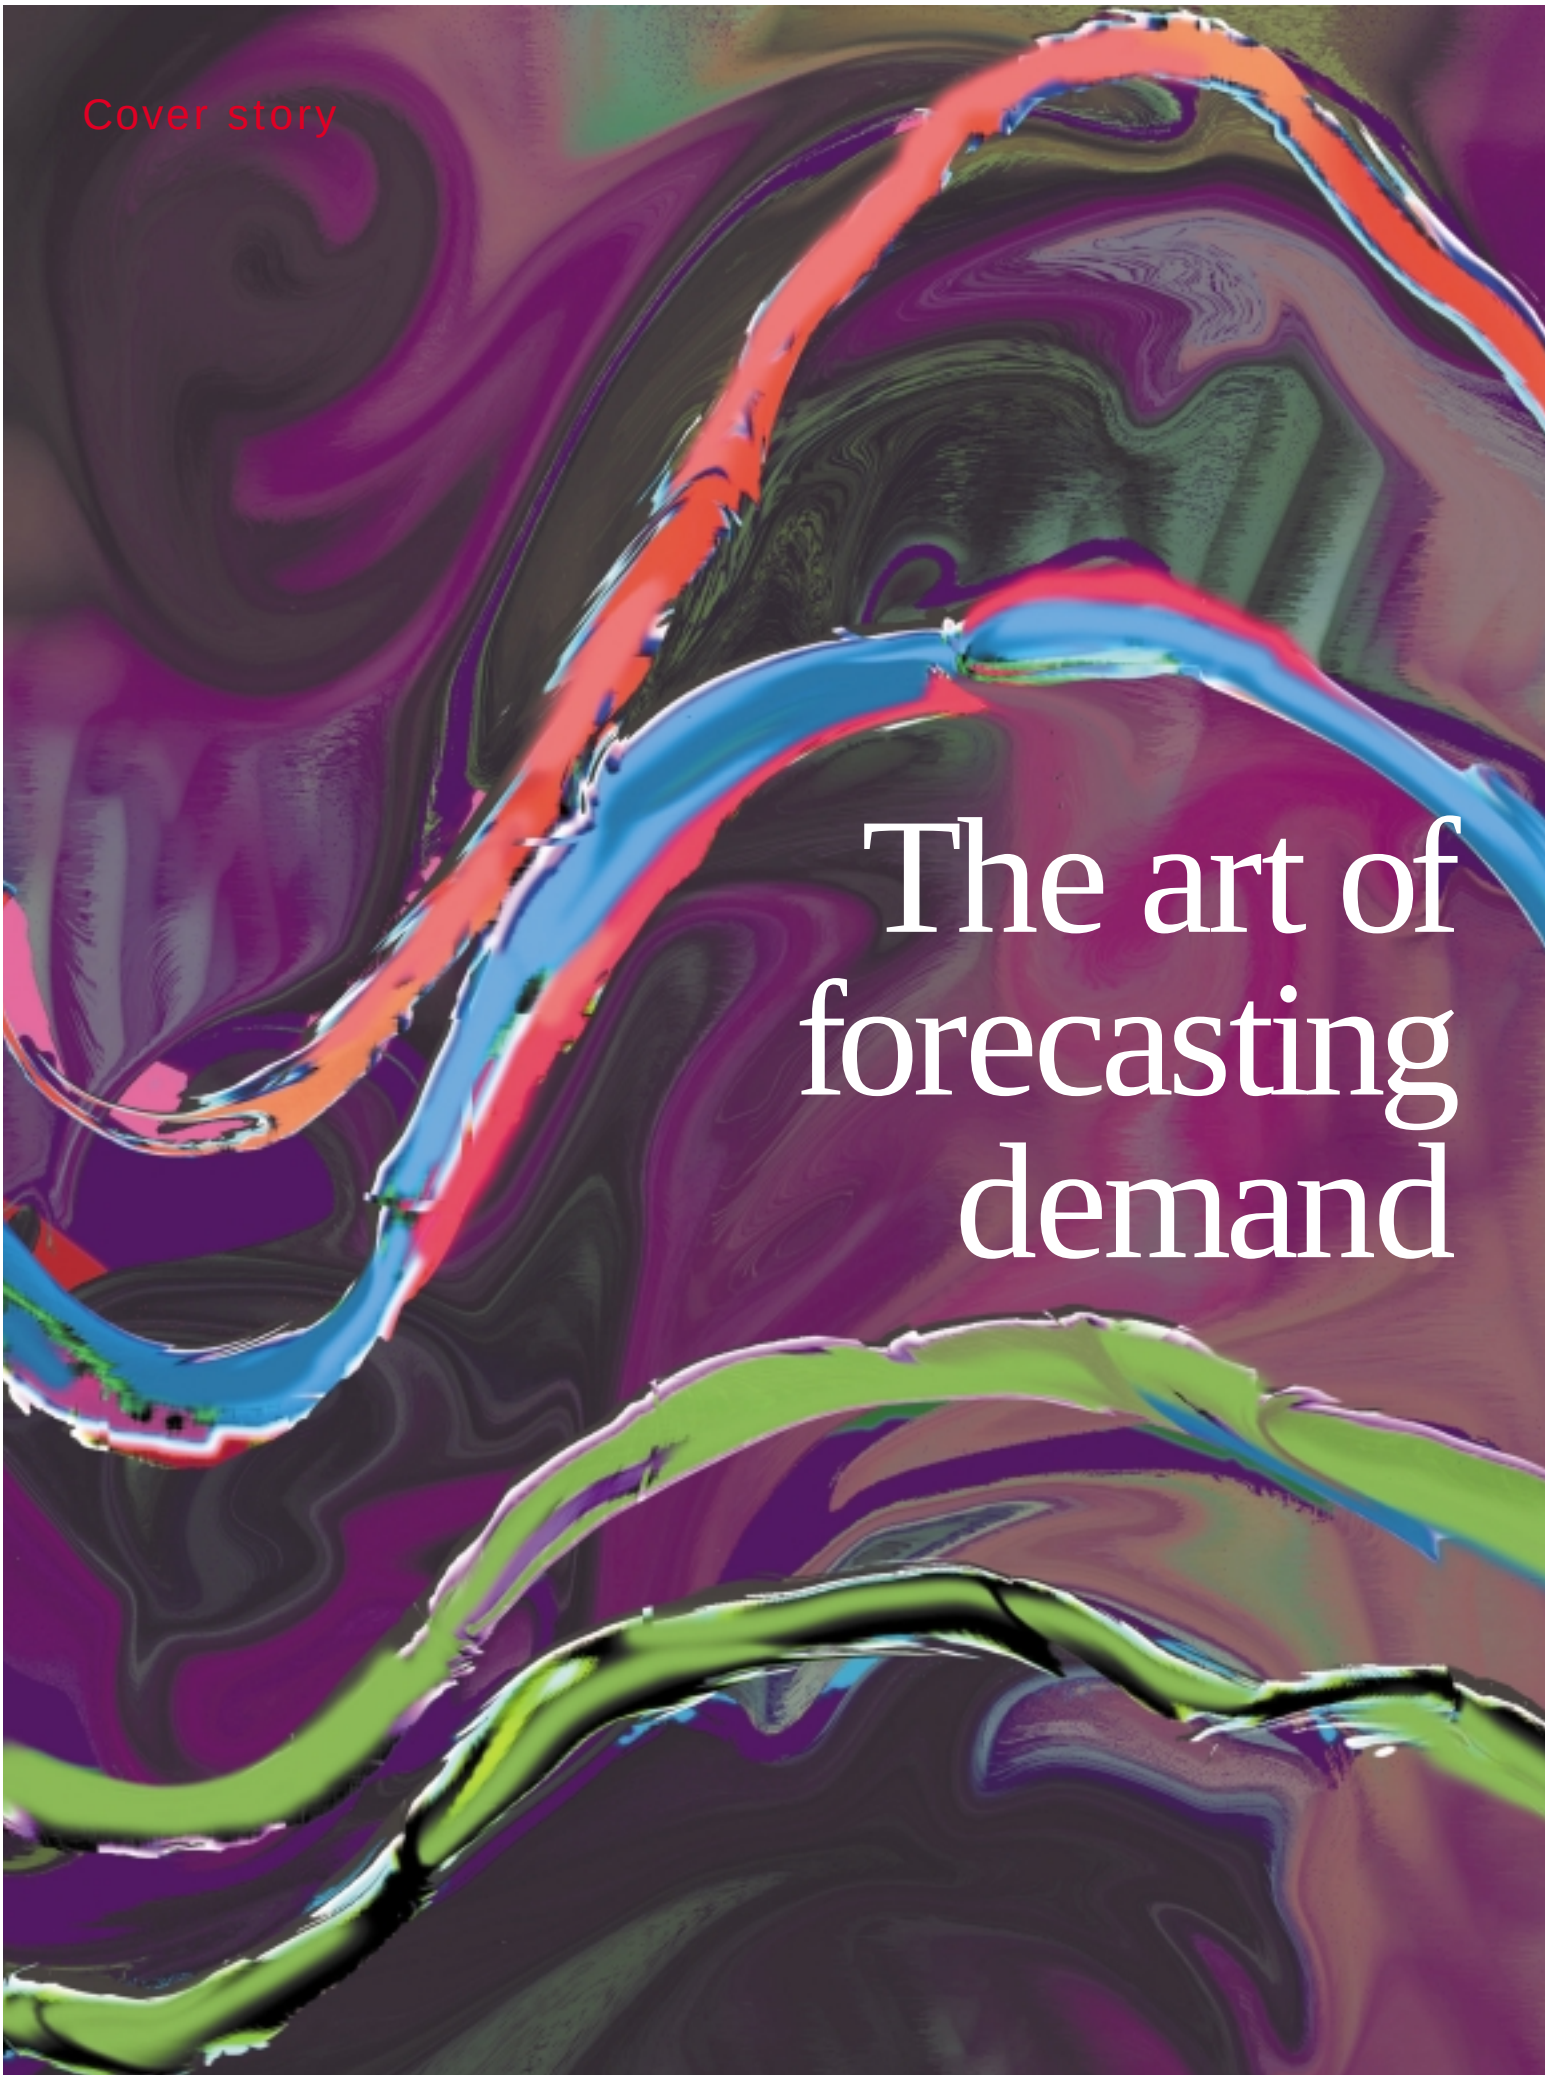

The art of forecasting demand

Winning in competitive electricity markets takes a fair amount of educated guesswork. Energy marketers cannot be certain that their future delivery of power at the price specified in a long-term contract will earn them a profit—because supply, demand, and the going rate may differ from expectations at the time the contract was signed. As a result, energy traders are only as good as the load forecasts they use.

With profit margins in many markets turning out to be as razor-thin as prices are volatile, some energy companies are taking steps to improve the accuracy of their demand forecasts. One is American Electric Power Co. (AEP), Columbus, Ohio. The predictions its traders rely on are continually refined and updated by mathematicians and statisticians in the company's new market analysis group. Significantly, the group was one of the first pieces of the wholesale trading operation that AEP began setting up years ago to prepare for competition.

What AEP and other energy marketers realize is that a forward price quote based on a poor demand forecast could translate into a loss rather than a profit. For a good example of the unpredictability of supply, demand, and price, consider what happened last year when Great Britain instituted its New Electricity Trading Arrangements (NETA; see box). For a time, market prices fluctuated wildly, between a low of minus $360/MWh (which marketers had to pay to sell to the system) to plus $70/MWh (which marketers had to pay to buy from the system). The energy and utilities team in Cap Gemini Ernst & Young's London office calculated that, in such a volatile environment, an improvement of 4% in the accuracy of a demand forecast could save the nation's electricity suppliers up to $29 million on balancing costs alone.

Because anomalies in demand tend to balance each other out over time, long-term demand forecasts tend to be more accurate than short-term ones. Although the accuracy of short-term forecasts has improved greatly in recent years, their sensitivity to unexpected events still defies even the most advanced modeling techniques. Take Sept. 11 for instance. The three minutes of silence observed in Great Britain last Sept. 14 caused the one of the largest drops in demand (2,700 MW, or 7% of system demand) in history.

Post-deregulation, load and demand forecasting are no longer black arts practiced by a chosen few. More power companies are recognizing that the ability to correctly predict demand and load—allowing them to buy and sell electricity at the best prices—means the difference between staying in business and going bankrupt. So now, rather than continuing to rely on off-the-shelf software packages to churn out forecasts, some companies are developing their own forecasting expertise in-house

Long-term demand vs. short-term load

Complicating the task for energy traders is the fact that long-term demand forecasting and short-term load forecasting require totally different skills and inputs, says Nigel B. Lewis, a managing consultant in Cap Gemini Ernst & Young's London office. Long-term demand forecasting requires a company to do scenario planning: what their portfolio will look like at some time in the future; what competitors might do; whether next winter will be warmer than usual; and so on. Lewis explains that, by comparison, short-term load forecasts are easier to get right because a company can base them on more precise information about customer behavior and shorter-term weather reports. However, the accuracy of short-term forecasts is bedeviled by the impact of major events—such as a Super Bowl or World Cup final telecast—that can be a dominating influence on consumption over five- to 10-minute intervals.

Deregulation has done more than make forecasting necessary; it has also changed who does it. When electricity sectors were regulated, utility monopolies used short-term load forecasts to ensure the reliability of supply, and long-term demand forecasts as the basis for planning and investing in new capacity from one to five years out.

That is no longer the case where competition has been or is being introduced. In Brazil, for example, the national utility Eléctrobras is now only interested in what demand will be well into the future, 10 to 20 years ahead. Shorter-term load forecasting is now the province of competitive power marketers.

Similarly, in Great Britain, prior to NETA the National Grid was responsible for producing forecasts and everyone else paid the same price for each half-hour in the UK Pool. Although load forecasting was used in structuring commercial supply contracts, it wasn't until NETA that suppliers had a financial incentive to forecast it on a regular basis.

The timeliness of forecasts can itself be problematic. U.S. power marketers, for example, say that the scarcity of timely aggregate load forecast data outside the five regions where an independent system operator (ISO) compiles them makes it nearly impossible for them to accu-

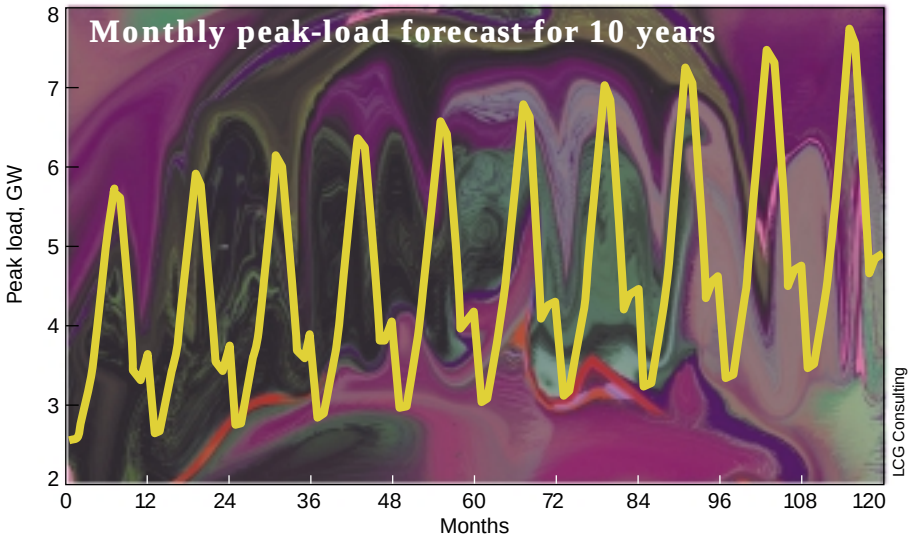

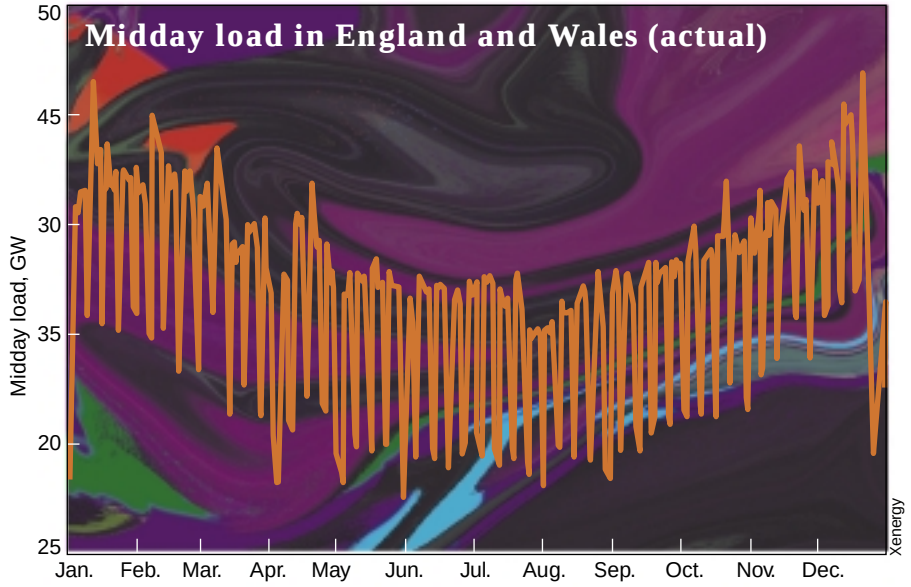

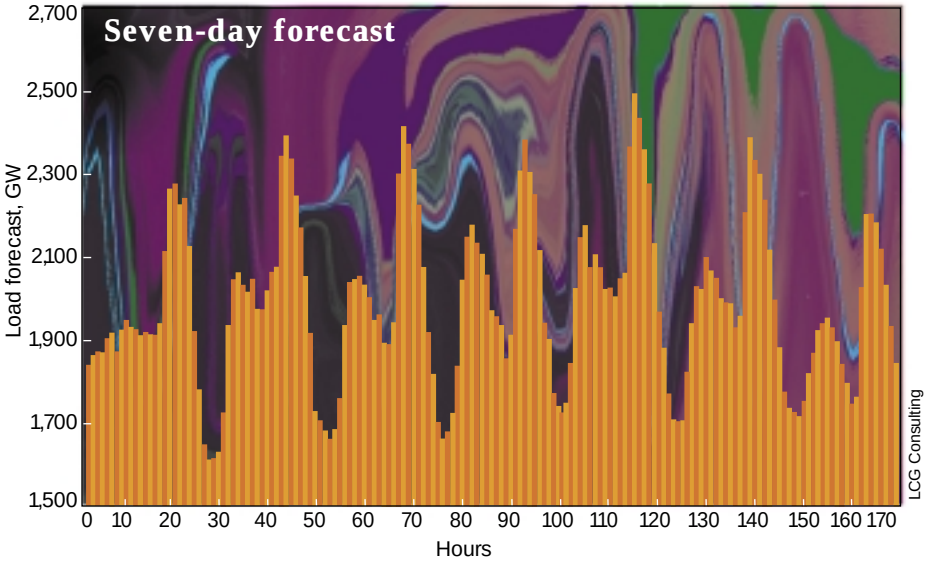

rately predict market swings in near real-time.

What about accuracy?

In some sense, the above discussion of the relative difficulty of making long- and short-term forecasts is disingenuous. It ignores the first question that anyone inside or outside the energy industry would ask about any forecast: How accurate is it?

"The one thing you can guarantee about forecasts is that they will be wrong," says Lewis. The best that he has seen in Great Britain is a 2% error within any half-hour period. In other words, if the actual load is 100 MW, the best forecast is either 98 or 102 MW, but not closer.

If the inputs to your forecasting model are poor, it's difficult or impossible to come up with a good forecast no matter how good your model is. In Brazil, for example, the consumption data recorded on a minute-by-minute by Eléctrobras basis are irregular and full of missing points and outliers. A research team led by Professor Reinaldo Castro Souza of Pontificia Universidade Católica-Rio is developing corrective filters to produce continuous observations to serve as input to the models.

The accuracy of forecasts also depends on the type of customer doing the consuming. Lewis observes that residential loads are easier to forecast because of the sheer number of residential customers. If one customer does something strange, it has less impact. On the other hand, a large industrial customer may behave unpredictably enough to belie predictions,—for example, by generating its own electricity or adding an extra work shift.

Micro and macro factors also affect load and demand forecasts, explains Paul Corby, senior VP, Planalytics, Wayne, Pa. By way of example, he reminds that commercial electricity consumption falls in depressed real estate markets because there's no one in empty offices using power. During recessions, industrial usage also falls because if no one is buying hard goods, factories shut down. The reality is that there is more elasticity in the indus-

Load forecasting under NETA

The New Electricity Trading Arrangements (NETA) is a bilaterally traded or OTC brokered market with a balancing mechanism; it replaced the U.K. Pool in March 2001. One difference between the U.K. Pool and NETA is buyers' submission of load requirements. These buyers include the former regional electricity companies (ex-RECs in their new identities, after the wave of mergers and acquisitions), Centrica, small niche suppliers, and Internet-based suppliers. Load forecasting is no longer just a concern of the system operator NGC, which balances the system.

As Cap Gemini Ernst & Young's Nigel B. Lewis observes, when NETA went live, people were mainly interested in getting something in place. To get load numbers, some people went out and bought software packages. A year later, they've come to realize that they need to know what's inside the black box.

Prior to NETA, the ex-RECs had forecasting capability but no financial incentive to improve. Their interest was more in long-range planning than in short-term forecasting—but that's what's needed to balance the system. The main driver then was to feed into business planning

NETA gave marketers a financial incentive to produce more accurate load forecasts. The penalty for an incorrect forecast can translate to as much as a 30% premium on electricity prices. Buyers also face a double-whammy problem: Power is most expensive when you need it most. There's a penalty whether you over- or underforecast. As a result, some traders purposely go long or short, but this doesn't avoid the penalty.

Shanti Majithia, head of operational forecasting at National Grid Co. (NGC), Coventry, England, observes that the market is usually long: The system buy price is higher than the system sell price. He adds that people tend to overforecast. NGC also does its own load forecasting for balancing purposes. If an imbalance occurs, NGC needs to contract for expensive generation, and this translates to heavy penalties for those who used more electricity than they bought.

trial and commercial sectors than there is in the residential sector.

Follow the weather

Amidst all this darkness, there is one beacon that gives power marketers hope that they can predict load and demand more accurately. According to Corby, there is virtually a one-to-one correlation between weather and residential power consumption.

It makes sense for marketers to seek out good weather forecasts; otherwise they'll have garbage in, garbage out. However, human nature makes most people take weather forecasts for granted. Cap Gemini's Lewis notes that nearly all of Britain's electricity marketers use Britain's Meteorological Office's weather forecasts verbatim in producing their own load forecasts. By doing so, they ignore at their peril the statistical uncertainties inherent in any forecast, much less one about something as fickle as weather. For instance, temperature and other numerical data get fed into the models without translating the probabilistic qualifiers that often accompany them. A forecast which states, "There's a 30% chance of rain, with an average temperature of 5 degrees Celsius," usually gets condensed and translated into the single-digit input "5 degrees."

One way around this problem is through the use of so-called weather ensemble predictions. Rather than using point forecasts, it makes use of multiple scenarios for the future value of a weather variable. Dr. James W. Taylor of the Business School at Oxford University, and Dr. Roberto Buizza, principal scientist at the European Centre for Medium-Range Weather Forecasts (ECMWF) in Reading, England, have found that weather ensembles greatly improve the accuracy of load forecasting by neural networks. The ensemble conveys the degree of uncertainty in the weather variable.

A more intuitive way of forecasting loads is to use a "similar day" analysis, suggests Pushkar Wagle, senior economist at LCG Consulting, Los Altos, Calif. This involves identifying a day in the past with similar weather conditions and projecting that particular day's loads onto the day of interest, without using any econometric technique. For daily (short-term) load forecasts for any given month, some of his clients are ready to pay up to $15,000.

Living with errors

If no forecast is perfect, what can be done about errors? Corby advises marketers to adopt a disciplined hedging policy, regardless of which model they use. One way to hedge the risk of errors is through portfolio diversification. Another is to "pass through" the errors. Lewis reports one British supplier trying to "incentivize" its clients to report their own forecasts as accurately as possible.

More and more, market participants are looking for a better solution to this problem. Instead of spending more time trying to produce better forecasts, they are doing transactions that recognize that forecasts can be far from accurate. Some companies have started to look for energy products that are tied to actual power pool load. Dr. Anil K. Suri, CEO of E-lectrade, Tarrytown, N.Y., explains that E-lectrade's "percentage of pool" structured energy product is designed to help manage the business risk associated with hard-to-predict loads.

Load and demand forecasting doesn't have to be treated either as a software "black box" or rocket science. Lewis thinks it's better to start with a simple model and understand and improve upon it, rather than start with a complicated one. If a complicated model breaks, you may not be able to fix it. He adds that forecasting isn't something you want to outsource because doing it helps you understand your customers. For that reason, load and demand forecasting are taking their place among the industry's best-kept secrets—secrets which companies are keeping in-house. ■

—Anne Ku