Energy Trading: A Source of Profit for Investment Banks and Hedge Funds, a Source of Challenging New Problems for Applied Mathematicians

By Rene Carmona and Michael Ludkovski

The increase in energy prices that followed Hurricane Katrina restored energy trading to front-page news. Ever since, oil and natural gas prices have repeatedly set all-time record highs. This spring, oil was again trading above US $70 a barrel, having fallen as low as US $30 as late as January 2004. Increases for natural gas have been even more spectacular, from US $3.10 per MMBtu (million British thermal units) in August 2002 to an all-time high of US $14.75 in October 2005.

Accompanying the dramatic run-up in prices has been very strong growth in trading volume. The current environment has proved to be a bonanza for traders, with energy accounting for up to 50% of total trading profits on Wall Street. The figures are staggering: JP Morgan reported $1 billion of revenue from energy trading in the third quarter of 2005; Goldman Sachs and Morgan Stanley each had net earnings of about $1.4 billion from energy trading in 2004. Trading arms of oil majors are also power players; BP, for example, reported revenue of nearly $2 billion in 2004. Energy traders, as reported in The Financial Times (11/22/2005), are the highest paid traders on Wall Street.

Goldman Sachs and Morgan Stanley continue to lead the way. Responsible not only for the vast majority of trades, they are also redefining the rules of the game, by trading physical assets in addition to financial contracts. These banks site and build power plants, own pipelines, lease oil tankers, and purchase other assets---all for use as hedging tools. Goldman Sachs alone spent $4 billion to buy 30 power plants in 2004. Financial firms now own 5% of total electricity-generating capacity in North America, up from less than 1% five years ago.

Attracted by these successes, other banks embarked on a shopping spree, seeking to shore up their energy divisions. Merrill Lynch bought Entergy-Koch trading operations for $800 million in 2004; Bear Stearns formed an alliance with Calpine before deciding to go solo after Calpine filed for Chapter 11 in December 2005. Hedge funds, too, are fledgling participants. In 2005, Forbes (9/21) estimated that 420 hedge funds had energy-trading strategies. The increased participation is apparently responsible for some of the price run-up: Higher demand for oil futures has added an estimated $10-15 to prices.

The Enron Legacy

Despite its negative impact on public perception, Enron should be credited for being at the source of modern energy trading. The string of innovations fueling its spectacular rise left an indelible mark. Created in 1984 from the merger of two gas pipeline companies, Enron was based in Houston, Texas, far from New York City, where financial innovation is expected to occur. Enron's introduction of the Gas Bank in 1989 created a marketplace for pipeline operators, spearheading financial settlement of gas transmission contracts. The immediate and enormous success of this venture changed gas trading forever. In 1992 Enron introduced the energy industry to market-to-market accounting, a source of controversy once accounting irregularities were discovered. Building on early successes, Enron entered the wholesale electricity market in 1996, in a merger with Portland General Electric. Enron's clout as a market maker for gas and electricity transactions increased its visibility in the financial arena, and allowed the construction of many innovative instruments, swing options being one example.

By the late 1990s *Enron-On-Line* was the major exchange for energy trading; the associated financial wizardry caused Enron to be called an "*energy hedge fund masquerading as a utility.*" The company quickly became the darling of the market. At its peak, in the summer of 2000, Enron was the seventh-largest company by market capitalization in the U.S. While brash, Enron was a generous benefactor, funding multiple charities and even academic research.

Emboldened by the growth of the late 1990s, Enron tried to replicate its good fortune with more exotic ventures---in the water business, weather derivatives, and Internet broadband. Except for weather derivatives (which are now exchange traded), all failed miserably. To maintain the perception of steady growth, Enron engaged in increasingly aggressive and fraudulent accounting; this was based in part on fake projections of financial markets that allowed Enron to book future profits from its current trading positions. Pretty much the only market maker at the time, Enron could quote any price it wanted. The in-house risk analysts vehemently opposed making such projections without corresponding market liquidity but were eventually overruled by management. The ongoing trials of former chairman and CEO Ken Lay and former COO/CEO Jeff Skilling continue to shed new light on the practices and company culture which led this innovative giant to a disgraceful exit.

The demise of Enron in late 2001 was a huge setback for the industry. Energy trading evaporated overnight, and many of Enron's smaller competitors, such as Dynegy and Mirant, and eventually Calpine, were dragged into bankruptcy. The energy market began to recover only with the arrival of a new oil shock in 2004, and some Enron trading platforms have never come back.

The New Energy Marketplace

The health of energy markets depends on their continued growth, which depends in turn on continued liberalization and deregulation. The California electricity deregulation fiasco in 2001, which ended with repeated rolling blackouts and cost the state more than $20 billion, had long-lasting consequences. The failed experiment caused panic among policy makers and voters worldwide (Switzerland rejected electricity deregulation in a vote last year). Enron, by the way, played a significant role in the California crisis, first through strong lobbying for deregulation and then as one of the most aggressive traders, exposing loopholes in the proposed market design and earning billions in the process. After the collapse of Enron, its top electricity traders were indicted and pled guilty to market manipulation.

Deregulation continues, however, albeit at a significantly slower pace. A case in point is the ongoing expansion of the PJM (originally Pennsylvania-New Jersey-Maryland, but now covering 13 states) electricity market in the eastern U.S., which currently serves more than 50 million people. PJM illustrates the potential for maintenance of a stable trading environment.

Recent increases in natural gas prices are also driving development of the Liquefied Natural Gas (LNG) infrastructure. LNG facilities already exist on the northeastern U.S. and Texas Gulf coasts. However, these no longer have enough capacity and several projects on the California coast (in Eureka,Ventura county, and all the way to Ensenada, Mexico) had to be abandoned after heated battles conducted by local communities and green activists. Nevertheless, it appears that Chevron-Texaco and Sempra's projects in Baja, Mexico, will be approved soon and become operational by 2008. The importance of LNG for energy markets lies in the potential for continued unbundling. A buzzword of early deregulation rhetoric, "unbundling" is meant to evoke the separation of services---generation versus transmission versus distribution versus retail---with the goal of creating more competition and greater liquidity. With separate transportation, storage, and land transmission facilities, all of which can be managed (and traded) piecemeal, LNG is a perfect candidate for unbundling.

On this note, it is important to recognize the fractured nature of gas and electricity markets, in which a vast variety of intricately interdependent contracts are traded. Most trades are done in *forward markets*, in which the up to 72 different maturities include contracts of more than 20 years. Multiple short-term markets, such as weekly, day-ahead (and in the case of power, real-time) prices, operate as well. Each is a separate price universe, with no clear price convergence, so that the day-ahead price can be very different from, say, the real-time price. One reason for this is storage. Electricity cannot be stored, and continuous balancing is required to ensure that demand matches supply. Accordingly, alongside the real-time market is a separate electricity reserve market for power plants kept on stand-by, ready to come online in the case of a sudden imbalance. For gas, storage is possible, but costly and generally inflexible. Facilities with flexible gas storage form yet another market that can be used to speculate on month-to-month fluctuations of gas prices.

Geographic location also matters---electricity in Los Angeles is not equivalent to electricity in Dallas (as Californians learned in 2001, when transmission capacity was not sufficient to bring power to the state from outside). Each town has its own electricity price: The PJM system, for example, quotes more than 3000 locational prices every day. The same applies to a lesser degree to gas, with each pipeline hub having a different price.

To make matters more complicated, gas and electricity are closely intertwined. Nearly all new power plants in North America run on gas; with an increase in gas prices comes an increase in the marginal cost of producing electricity and, thus, in power prices. In any given season (summer, say), electricity and gas demands are themselves dependent on temperature through a nearly linear function. This well-documented empirical fact makes temperature options a prime candidate for hedging volume risk, and leads to an interdependence of weather and energy derivatives. Because power plants also generate a lot of emissions, commodity prices are related to emissions trading, which is also picking up steam, especially in Europe.

With so many markets, owning physical assets like power plants and refineries is crucial. As mentioned earlier, some of the most successful Wall Street investment banks have bought hard assets to back financial contracts. With ownership of a physical asset, however, comes the need for a scheduling desk to optimize operations. Game-theoretic issues arise in the pricing of power, in the functioning of oligopolies, and in the design of appropriate markets. Plant outages are an example; directly affecting available supply, they are critical for determining prices. Because they have been used to manipulate prices, it is important to understand their impact.

Failure to hedge energy risk properly can be disastrous. In the aftermath of Katrina, losses for some hedge funds are rumored to have exceeded $100 million. Airlines have collectively lost billions of dollars because of their refusal to hedge jet fuel costs. Not so long ago, Southwest, the only airline to hedge its exposure to fuel risk, became far more profitable than most of its competitors. Similarly, the inability of Pacific Edison to hedge its electricity purchases led to bankruptcy in 2002.

Mathematical Modeling Issues

As hinted at earlier, energy markets harbor a wealth of new analytical problems that are ripe for mathematical modeling. Quantitative research on energy markets began with ideas imported from traditional financial engineering. A prime example is the cross-product spread option used in equity and foreign-exchange markets. The spread option has found many natural applications in energy derivatives, including the crack-spread (price difference between crude oil and gasoline) and the spark-spread (between natural gas and electricity). Some of the modeling and computational aspects of spread options are discussed in a 2003 SIAM Review paper, "Pricing and Hedging Spread Options" [1]. More exotic derivatives have also gained popularity, including swing options and multiple American exercise contracts.

Second-generation mathematical energy models are sorely needed. The obvious lack of market completeness, low liquidity, and poor transparency in price discovery are some of the stumbling blocks. On a more fundamental level, traditional financial mathematics revolves around the idea of no-arbitrage. But in the case of electricity, which is non-fungible and non-storable, what does arbitrage mean? The extreme price spikes for electricity seem to contradict the standard Markovian assumption. Pronounced seasonality (natural gas is always more expensive in winter) and mean-reversion must also be considered.

Unlike traditional financial assets, there is no consensus as to appropriate models for energy price dynamics. All of the more than a dozen proposed stochastic models for the evolution of electricity prices, for example, have limited credibility. Similarly, there is no consensus on the way to model storable commodities. Historically, the notion of *convenience yield* was introduced to measure the net benefit of holding a physical asset compared with owning a forward contract (with the benefit arising from the increased flexibility---the owner of a physical asset does not face the possibility of commodity shortage). This convenience yield opens a can of worms, however: Is it a real concept or just a modeling artifact? Is it observed? Is it stochastic? How does it relate to theory of storage and other economic paradigms? Can we do without it?

An understanding of *market-wide* equilibrium must include engineering concepts. For instance, operational constraints, such as time required to bring a power plant online, are often crucial for understanding price evolution. Similarly, the *geometry* of the transmission grid plays a large role in determining congestion and locational prices.

Because of all the aforementioned complications, computational approaches, especially those based on Monte Carlo simulation, are the tool of choice among practitioners. Increasingly, energy trading is one of the main sources of new *computational* challenges in financial mathematics. Statistical challenges arise as well. First, the amount of data is enormous (recall, for example, the 3000 daily electricity prices in the PJM market). Second, data are often available for too short a time to permit full analysis, and because of rapidly changing conditions, time series are highly non-stationary: Should we really count the prices from California 2001?

A brief discussion of a specific problem---asset valuation---serves as an illustration. The theory of *real options* was developed to take into account the optionality of abandonment in the valuation of projects. When applied to the valuation of a power plant, this theory implies that the plant should run when the spark-spread is positive (i.e., when the cost of fuel is less than the revenue from the power produced) and be shut down whenever the spread turns negative. Moreover, if ramp-up and shutdown costs (and the safety issues associated with frequent switches) are ignored, this strategy is optimal. In computations of the expected present value of this operating scheme, the value of the plant appears as the sum of a string of daily spark-spread options over the period of ownership of the plant. This approach was the main motivation for the analysis presented in [1].

A more sophisticated model would explicitly model the operating constraints and allow continuous output levels. Moreover, the cost of a specific operating strategy might be path-dependent, making the control infinite-dimensional both in state space and in time. To simplify, we can assume that the Markov property holds, and reduce the modes of operation of the plant to a finite number of states, the operator of the plant being able to switch from one state to another at well-understood economic costs. In this way, the difficult stochastic control problem is reduced to an optimal switching problem that can be studied theoretically and solved numerically; see [2]. Many challenges remain: development of models that can take price impact into account, or of tractable methods for valuation of a fleet of power plants in very distant locations. These are not abstract concerns: A recently reported deal was a complex US $130 million spark-spread hedge between Morgan Stanley and Calpine involving 13 power plants.

Academic Activities and Opportunities for Research

As the business side picks up, there is opportunity for increased collaboration with industry. Enron was an early pioneer in supporting academic research on energy-related problems, generously funding in-house mathematicians and academic consultants. Nowadays, the investment banks have returned to university campuses, hiring many mathematics and statistics graduates (PhD or master's-level) for their energy trading operations.

The energy industry has traditionally been run by engineers. The gradual shift to probabilistic and stochastically oriented approaches now under way makes this an opportune time to contribute the academic perspective. For probabilists, numerical analysts, computer scientists, game theorists, optimization experts, and industrial engineers, a goldmine of tailor-made research topics lies in the modeling of energy markets, managing resources, and quantifying and mitigating monetary risk.

The relevant mathematical literature has been growing steadily, and even in the absence of dedicated academic journals, many technical articles are published in such peer-reviewed journals as SIAM Review, SIAM Journal on Control and Optimization, Mathematics of Operations Research, and Journal of Economic Dynamics & Control. From the practitioner's perspective, there is Energy Risk magazine, which also publishes Energy Risk Management handbooks. Textbooks have also started to appear; see [3,5,6].

Conferences have been devoted to energy risk, and two minisymposia on the topic are planned for the upcoming SIAM financial mathematics conference, which will be held in Boston, July 9-12, in conjunction with the 2006 SIAM Annual Meeting. Many other workshops have been held, including two in Princeton (2001 and 2003). This exciting research area is inherently interdisciplinary. It will thrive through synergistic interactions between mathematicians, economists, and engineers bringing together the stochastic, financial, and operational aspects. As such, it is a perfect new testbed for applied mathematics.

[1] R. Carmona and V. Durrleman (2003): Pricing and Hedging Spread Options. SIAM Review 45(4): 627-685.

[2] R. Carmona and M. Ludkovski (2005): Optimal Switching and Application to Tolling Contracts (tech. rep. FROM WHERE?).

[3] L. Clewlow and C. Strickland (2000), *Energy Derivatives: Pricing and Risk Management*, Energy and Power Risk Management.IS THAT A JOURNAL? OR A CHAPTER IN A BOOK? IF JOURNAL, NEED VOL, PAGE NOS.

[4] A. Dixit and R. Pyndick (1994), *Investment under Uncertainty*, Princeton University Press.

[5] A. Eydeland and K. Wolyniec (2003), *Energy and Power Risk Management:* New *New Developments in Modeling, Pricing and Hedging*, Wiley.

[6] H. Geman (2005). *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*, Wiley.

Rene Carmona, chair of the SIAM Activity Group on Financial Mathematics and Engineering and co-chair of the organizing committee for the group's July meeting, is Paul M. Wythes ’55 Professor of Engineering and Finance at Princeton University. Michael Ludkovski is Assistant Professor at the University of Michigan.