EPRI Powering Progress

WO3581

Technical Brief

POWER DELIVERY GROUP

Forward Curve Dynamics and Asset Valuation

Overview
Wholesale electricity and fuel markets are rapidly becoming highly competitive commodity markets where the techniques for establishing the value of assets and contracts differ from techniques used in regulated environments. This document discusses the use of simple models of forward curve dynamics to define valuation methods for assets whose value depends on the market prices of traded commodities.

Forward curve dynamics describes how the full curve of forward prices in a market changes over time. This contrasts with common practice in which models are defined in terms of how only the spot price changes over time. The benefit of moving from spot price dynamics to forward curve dynamics is that many details of the underlying physical supply system are contained within the initial forward curve and so do not need to be explicitly included in the price model. In addition, the ability to apply data on movements in the full forward curve in tuning the model parameters rather than relying solely on spot price data permits easier and better model calibration. Once price models are defined for relevant markets and their parameters set, the tools of finance theory can be brought to bear in order to compute the fair market value of assets or contracts whose cash flows depend on the prices in those markets.

This document will discuss in detail the use of the one and two factor versions of the N factor forward curve dynamics model,

$$\frac{dF_{tT}}{F_{tT}} = \sum_{i=1}^{N} \beta_i(t) e^{-\alpha_i(T-t)} dz_t^i$$

where F_{tT} denotes the forward price at time t for delivery at a later time T. The chief benefit of this family of models is that analytical solutions are available for many quantities of interest. The one and two factor versions also provide quite good fits to observed market data and are relatively tractable for a wide variety of numerical calculations. In particular, such models can be translated into equivalent spot price dynamics models so that existing tree methods can be used to value assets that require full dynamic information (e.g., American options).

Valuing Assets
Finance has provided a wealth of tools and theory for computing the fair market value of assets whose cash flows are dependent on the prices of various traded commodities. Finance's Risk Neutral Valuation rules show the critical role that forward prices and volatilities play in the valuation of such assets. Forward prices are expectations of future commodity prices adjusted for the risk associated with the uncertainty in those prices. The "forward curve" signifies a set of forward prices for all future delivery times of interest. These prices can be estimated from futures exchange prices and from the prices at which firm contract commitments are traded. Volatility is a measure of the uncertainty in a future commodity price at a particular time and can be estimated from the prices of traded options or from model fits to spot or forward price changes over time. The "volatility term structure" signifies values of the volatility for all future delivery times of interest. It should be noted that, for most market participants, forward curve data is much more easily obtained than volatility data.

Finance's Risk Neutral Valuation rules provide a general framework for asset valuation that simplifies into several special cases. The Risk Neutral Valuation rule for assets whose future deliveries and payments are known with certainty (e.g., forward contracts) is to consider all future deliveries of a commodity to be worth the forward price for that delivery date and to discount all future cash flows using the interest rate on risk-free investments of comparable term. For example, a forward contract that provides a payment of K for delivery of Q units of the commodity at time T will be worth

$$V_0^{Forward} = Q(F_{0T} - K)e^{-r_T T}$$

where r_T is the appropriate risk-free rate for zero-coupon bonds paying off at time T. Notice that the current forward curve F_{0t} is the only information about future commodity prices needed to value such assets. When such assets are freely traded, the prices at which they trade can be used to back out corresponding forward prices. Most market participants have a fairly good handle on estimating the forward curve, at least over short to medium time horizons.

The Risk Neutral Valuation rule for valuing assets whose cash flows at various points in time depend only on the commodity prices at that time (e.g., forward contracts, European option contracts) is to compute the expectation of future cash flows using a "risk neutral" probability distribution of future commodity prices and then discount those expected cash flows to the present using the risk-free rate. The risk neutral probability distribution is usually defined as the lognormal distribution[1] of future prices that at each point in time has its expected value equal to the forward price for that delivery date (instead of the true expected future price) and has the variance of the natural log of the price equal to the true variance of the natural log of the price. In particular, if an asset provides cash flow $CF(P_t,t)$

at each time t when the commodity price at that time is P_t then that asset will be worth

$$V_0 = \sum_t \langle CF(P_t,t) \rangle e^{-r_t t}$$

where <...> signifies the (risk neutral) expectation over the uncertain price P_t. P_t is assumed to be distributed at each time t according to a lognormal distribution with

$$\langle P_t \rangle = F_{0t}$$

$$\mathrm{var}(\ln(P_t)) \equiv \sigma(t)^2 t$$

(The relationship shown between the volatility $\sigma(t)$ and the variance of the natural logarithm of price defines the volatility.) Notice that the value of such assets depends solely on the forward curve F_{0t} and the volatility term structure $\sigma(t)$ since they fully define the risk neutral distribution at each point in time. For example, for a forward contract delivering Q units at time T for price K, the cash flow would be $Q\ (P_T - K)$ for which the (risk neutral) expectation is $Q\ (F_{0T} - K)$, leading to the result above for forward contracts.

In the case of European call and put options, the expectation of cash flows over the risk neutral distribution can be carried out in closed form, leading to the famous Black Equation.[2] For example, a European call option with strike price K and expiration at time T would have terminal cash flow $MAX(P_T - K, 0)$ which has expectation $(F_{0T}N(a) - K\,N(b))$ leading to

$$V_0^{Call} = e^{-rT}\left(F_{0T}N(a) - K\,N(b)\right)$$

where $N(x)$ is the cumulative standard normal distribution (tabulated in various math references and included as a special function in most major spreadsheets) and

$$a = \frac{\ln(F_{0T}/K) + \frac{1}{2}\sigma(T)^2 T}{\sigma(T)\sqrt{T}}$$

$$b = a - \sigma(T)\sqrt{T}$$

Similarly, the European put option with strike K and expiration T has terminal cash flow $MAX(K - P_T, 0)$ which leads to value

$$V_0^{Put} = e^{-rT}\left(K\,N(-b) - F_{0T}N(-a)\right)$$

If the forward curve and traded European option prices are available, the Black Equation can be inverted to determine the "implied" volatility term structure $\sigma(t)$. In other cases[3], the volatility must be estimated by fitting price dynamics models to historical spot or forward price changes. Accurate volatility estimation is a serious issue for most market participants since there is rarely enough option price data to create good implied volatility estimates.

The Risk Neutral Valuation rule in more general cases requires information about the full dynamics of spot price changes to carry out the calculations. Spot price dynamics are usually described using the tools of Ito's stochastic calculus.[4] Applying this mathematical representation, the change in spot price P_t over a small period of time dt is modled as

$$dP_t \equiv P_{t+dt} - P_t = \alpha(P_t,t)dt + \sum_{i=1}^{N} \beta_i(P_t,t)dz_t^i$$

where $\alpha(P_t,t)$ describes anticipated changes in the spot price over the time interval and the $\beta_i(P_t,t)$ describes possible uncertain changes in the price that could result from N types of new information that could arrive during the time interval from t to $t+dt$. The uncertainty enters the equation via the dz_t^i, which are uncertain variables with a mean of zero and a standard deviation of $\sqrt{dt}$. The dz_t^i at differing times t are uncorrelated, but the dz_t^i at the same time can have correlation $\rho_{ij}(t)$ between dz_t^i and dz_t^j. The variety of spot price dynamics models results from differing choices of the number of uncertain factors N and the functions $\alpha(P_t,t)$, $\beta_i(P_t,t)$, and $\rho_{ij}(t)$. A variety of tree methods have been developed to carry out the valuation of assets that depend on full spot price dynamics. These methods define an uncertainty tree of potential (risk neutral) spot price outcomes over time which can be used to identify potential cash flows for the asset. The asset value is then determined by discounting the expected cash flows back to the present using the risk-free rate. The key advantage of the tree approach is that it allows for intermediate decisions to change the cash flows over time. For example, an American option can have an early exercise decision that eliminates all future cash flows.

Why Forward Curve Dynamics?

The choice of a spot price dynamics model will affect the valuation of complex assets (i.e., assets that cannot be valued using just the forward curve and volatility term structure). For example, the decision to exercise an American option will depend not just on the price at a particular time but also on the expectations of how that price is likely to change over time. If European option prices are not available, the model choice will also affect how spot and forward price data is used to estimate the volatility term structure needed to value any asset with decision flexibility (e.g., all options). Models of spot price dynamics are typically adopted ad-hoc from financial markets and then fit to match the forward curve and volatility term structure of the commodity as closely as the model's free parameters allow. In most cases, there are numerous financial models of spot price dynamics to choose from and it is not entirely clear how to judge each model's assumptions in terms of the commodity market at hand.

The difficulty in judging spot price model assumptions in commodity markets is due to the fact that a commodity at two different points of time should really be considered two different commodities. Unlike financial markets where —in the absence of new information— the price of an asset at one point of time is simply related to its price a short time before, in commodity markets the price of the commodity at the later time can differ drastically from the earlier price. As an extreme example, consider electricity delivered at differing points in time. There is usually no economically feasible way to convert electricity at one time to electricity at another time, and so electricity must be generated on demand. Because of this non-storability, electricity spot prices hours apart can differ dramatically. Even in the case of storable commodities such as grain or petroleum products, details of production capacities, spoilage or degradation rates, storage facility costs and capacities, and other underlying market issues greatly complicate the direct modeling of the changes in spot price over time. Unfortunately, most financial models of spot price dynamics make strong assumptions about how the spot prices change over time. These assumptions are difficult to judge in the context of commodity markets.

The goal of the forward curve dynamics approach described in this document is to make simplifying assumptions about how the full forward curve changes over time rather than to make simplifying assumptions about how the spot price changes. Since the forward curve summarizes the relationship between prices of the commodity at differing points of time, it already contains detailed information about the underlying physical market complexities. This approach leverages both the full information content of the forward curve and the wealth of analytical tools (such as Ito calculus and Risk Neutral Valuation) developed in the financial markets.

Models of Forward Curve Dynamics

This section presents the general class of forward curve dynamics models that will be used in the rest of this document and derives their most important properties. This section is a bit heavy-going and the reader may wish to skip it on a first reading. The following section will focus on the specific members of this class appropriate for most situations and show how they can be used.

To describe forward curve dynamics, changes in the full set of forward prices F_{tT} over time must be characterized. This can be accomplished by using the tools of Ito calculus. The basic starting point for the forward curve dynamics approach is to treat each forward price as a separate market price. This assumption initially leads to consideration of a rather unwieldy set of price equations

$$dF_{tT} = \alpha(F_{tT},t,T)dt + \sum_{i=1}^{N} \beta_i(F_{tT},t,T)dz_t^{iT}$$

where everything can potentially depend on the delivery date T of the commodity. A number of simplifications can then be made due to the specific characteristics of forward price data.

The first simplification is to recognize that since the forward price is the (risk neutral) expectation of future price, the expected change in forward price F_{tT} must be zero. This eliminates all of the $\alpha(F_{tT},t,T)$ terms from the price dynamics equations. The next simplification is to notice that the forward curve shifts in a fairly smooth manner. This implies that the full set of potential dz_t^{iT} can be productively replaced by a much smaller set of uncertain factors dz_t^i, where each one affects the forward curve over a range of delivery times. The smoothness of the curve shifts can be understood by recognizing that storage capability links the prices in differing times, so that an increase in the price at one time will raise the prices of nearby times. More generally, the curve shifts can be understood by realizing that much of the news altering the supply-demand balance will persist over a broad period of time. Another simplification is to model each uncertainty term as $F_{tT}\beta_i(t,T)dz^i$, so that the resulting forward price distribution is lognormal and therefore the forecasted forward prices are always positive. The final simplification is to note that short dated forward prices are observed to be much more volatile than long dated ones, suggesting that the uncertainty terms be modeled with

$$\beta_i(t,T) = \beta_i(t)e^{-a_i(T-t)}.$$

Collectively, these simplifications motivate the study of forward curve dynamics models of the form

$$\frac{dF_{tT}}{F_{tT}} = \sum_{i=1}^{N} \beta_i(t)e^{-a_i(T-t)}dz_t^i$$

where different models depend on the number N of uncertain factors and the forms of the functions $\beta_i(t)$. These forward curve dynamics models are extremely analytically tractable, can exactly fit observed forward curves (as can all forward curve dynamics models), can fit observed volatility data and historical data on forward price changes over time quite well, and can be translated into equivalent spot price dynamics models to leverage existing tree solution algorithms for valuing complex assets. The one and two factor versions will be explored in some detail in the following section, but first some facts about this full set of models will be established.

First, the risk neutral probability distribution of the spot price P_t corresponding to these forward curve dynamics models will be derived. This spot price distribution will show how the volatility term structure $\sigma(t)$ is related to the model parameters. To carry out this derivation the forward curve dynamics model must first be integrated to relate the (uncertain) forward curve F_{tT} at all future times t to the current forward curve F_{0T}, the model parameters, and outcomes of the uncertain factors dz_t^i. Namely,

$$F_{tT} = F_{0T}\exp\left(-\tfrac{1}{2}\sum_{i,j=1}^{N}\int_0^t ds\,\rho_{ij}(s)\beta_i(s)\beta_j(s)e^{-(a_i+a_j)(T-s)} + \sum_{i=1}^{N}\int_0^t \beta_i(s)e^{-a_i(T-s)}dz_s^i\right)$$

The relationship between the spot price P_t and the forward curve that $P_t = F_{tt}$ can then be used to immediately find that

$$P_t = F_{0t}\exp\left(-\tfrac{1}{2}\sum_{i,j=1}^{N}\int_0^t ds\,\rho_{ij}(s)\beta_i(s)\beta_j(s)e^{-(a_i+a_j)(t-s)} + \sum_{i=1}^{N}\int_0^t \beta_i(s)e^{-a_i(t-s)}dz_s^i\right)$$

In this expression, all of the uncertainty arises through the terms

$$w_t^i \equiv \int_0^t \beta_i(s)e^{-a_i(t-s)}dz_s^i$$

which contain the uncertain factors dz_t^i. Since the factors dz_t^i are joint normally distributed with a mean of zero at each time, are uncorrelated at differing times, and have correlations $\rho_{ij}(t)$ between dz_t^i and dz_t^j at the same time, the w_t^i are joint normally distributed at each time t with

$$\langle w_t^i \rangle = 0$$

$$\mathrm{var}(w_t^i) = \int_0^t \beta_i(s)^2 e^{-2a_i(t-s)}ds$$

$$\mathrm{covar}(w_t^i, w_t^j) = \int_0^t \rho_{ij}(s)\beta_i(s)\beta_j(s)e^{-(a_i+a_j)(t-s)}ds$$

The spot price distribution depends on the sum of the w_t^i, which is itself normally distributed with a mean of zero and a variance equal to the sum of $\mathrm{covar}(w_t^i, w_t^j)$ over all i,j. This allows the risk neutral probability distribution of P_t to be written as the lognormal distribution

$$P_t = F_{0t}\exp\left(-\tfrac{1}{2}\sigma(t)^2 t + \sigma(t)\sqrt{t}\,W\right)$$

where the volatility is related to the variance of the natural logarithm of price (which is equal to the variance of the sum of the w_t^i) through

$$\sigma(t)^2 \equiv \tfrac{1}{t}\mathrm{var}(\ln(P_t)) = \tfrac{1}{t}\sum_{i,j=1}^{N}\int_0^t ds\,\rho_{ij}(s)\beta_i(s)\beta_j(s)e^{-(a_i+a_j)(t-s)}$$

and where W is distributed according to the standard normal distribution. This equation for the volatility $\sigma(t)$ is the basis for relating the volatility term structure to the model parameters or, as will be shown below, for fitting the model parameters to observed implied volatility data.

Second, methods for fitting the model parameters to market price data will be discussed. Two approaches can be taken to fit the model parameters to market data. In the first approach, historical data on simultaneous changes in forward prices at various delivery times T are used to fit the parameters using the forward curve dynamics equation. This approach is best when there is access to historical forward curve data. In the second approach, the implied volatility term structure is fit using the above equation relating the volatility to the model parameters. This approach is best when the user has access to the prices of options which trade in liquid markets and can therefore be used to obtain good implied volatility estimates. In either approach, a specific parametrized functional form is assumed for the unknown functions $\beta_i(t)$ and $\rho_{ij}(t)$. (In most cases they are assumed either constant or repeating on an annual basis.) In the more common case of assuming constant $\beta_i(t)$ and $\rho_{ij}(t)$, the first approach attempts to fit parameters to match data on forward curve changes at differing delivery dates T,S through

$$\tfrac{1}{dt}\mathrm{var}\left(d\ln(F_{tT})\right)=\sum_{i,j=1}^{N}\rho_{ij}\beta_i\beta_j e^{-(a_i+a_j)(T-t)}$$

$$\tfrac{1}{dt}\mathrm{covar}\left(d\ln(F_{tT}),d\ln(F_{tS})\right)=\sum_{i,j=1}^{N}\rho_{ij}\beta_i\beta_j e^{-a_i(T-t)-a_j(S-t)}$$

$$d\ln(F_{tT})\equiv\ln(F_{t+dt\,T})-\ln(F_{tT})$$

while the second approach fits the parameters to match the implied volatility through

$$\sigma(t)^2 t=\sum_{i,j=1}^{N}\rho_{ij}\beta_i\beta_j\left(\frac{1-e^{-(a_i+a_j)t}}{a_i+a_j}\right)$$

With either approach, the model parameters should be varied so as to minimize the difference between the right sides of the equations (which depend on the model parameters) and the left sides of the equations (which depend on the market data). Most major spreadsheets contain tools that can be applied to minimize a measure of the fit error (e.g., the sum over all data points of the squared differences between the left and right sides of the equations) by varying the model parameters.

Finally, the spot price dynamics models corresponding to these forward curve dynamics models will be derived. These spot price dynamics models are used to construct decision trees for valuing more complex assets than can be valued using the forward curve and volatility alone (e.g., American options). The easiest way to accomplish this is to return to the first expression above for the risk neutral distribution of P_t and substitute in the definitions of w^i_t and $\sigma(t)$ to find that

$$P_t=F_{0t}\exp\left(-\tfrac{1}{2}\sigma(t)^2 t+\sum_{i=1}^{N}w_t^i\right)$$

where

$$w_t^i\equiv\int_0^t\beta_i(s)e^{-a_i(t-s)}dz_s^i$$

$$\sigma(t)^2=\tfrac{1}{t}\sum_{i,j=1}^{N}\int_0^t ds\,\rho_{ij}(s)\beta_i(s)\beta_j(s)e^{-(a_i+a_j)(t-s)}$$

The expression for w^i_t can be differentiated with respect to time to show that it follows the dynamics equation

$$dw_t^i=-a_i w_t^i dt+\beta_i(t)dz_t^i$$

where $w^i_0=0$. Tree methods can be used to characterize the joint outcomes of w^i_t at each time t, which can then be turned into the corresponding spot price outcome using

$$P_t=F_{0t}\exp\left(-\tfrac{1}{2}\sigma(t)^2 t+\sum_{i=1}^{N}w_t^i\right)$$

to calculate the resulting cash flows at that time and the uncertainty outcome. The value of the asset is then computed as the expected present value of the cash flows where all discounting is done at the appropriate risk-free interest rate.

Forward Curve Dynamics Models for Everyday Use

Most practitioners limit themselves to the use of one or two factor models of price dynamics. This is due both to the sparseness of data available for fitting the model parameters and to limits of computing power and memory. However, even in the realm of one and two factor models, there are more than enough models available to meet most analysis needs. In this section, one factor and two factor forward curve dynamics models are presented. The models will differ in terms of their assumptions as to what types of new information typically arrive concerning the commodity market.

Long-Term Supply-Demand Shifts

The simplest forward curve dynamics model assumes that all new information affects the whole forward curve equally. Intuitively, the assumption is that the new information describes changes in the supply-demand balance that will persist over the long term. Mathematically, the model is

$$dF_{tT}/F_{tT}=\beta(t)dz_t$$

where the initial forward curve F_{0T} is assumed to be known and the volatility term structure is related to the model parameters through $\sigma(t)^2=\tfrac{1}{t}\int_0^t\beta(s)^2ds$

Notice that the volatility is constant and equal to β when $\beta(t)$ is constant. Many commodities exhibit decreasing volatility over time, which requires that $\beta(t)$ decrease steadily over time in order to fit the volatility term structure.

The model parameters can be matched to market data by fitting $\beta(t)$ to the above volatility equation or by fitting it to changes in the forward prices over time through

$$\tfrac{1}{dt}\mathrm{var}\left(d\ln(F_{tT})\right)=\tfrac{1}{dt}\mathrm{covar}\left(d\ln(F_{tT}),d\ln(F_{tS})\right)=\beta(t)^2$$

$$d\ln(F_{tT})\equiv\ln(F_{t+dt\,T})-\ln(F_{tT})$$

which is usually carried out assuming $\beta(t)$ is constant.

The spot price dynamics model for this forward curve model is

$$P_t\equiv F_{0t}\exp(-\tfrac{1}{2}\sigma(t)^2t+w_t)$$

$$dw_t=\beta(t)dz_t$$

$$w_0=0$$

which in the financial literature is known as the "Random Walk Model". A more familiar (and equivalent) form for this model is

$$\frac{dP_t}{P_t}=\frac{d\ln(F_{0t})}{dt}dt+\beta(t)dz_t$$

which shows that the forward curve characterizes the expected changes in price over time. For valuing complex assets, either the Binomial Tree Method or the Finite Difference (Trinomial Tree) Method can be used for valuation.[6]

This model is most appropriate when there is very little data available (i.e., as a first cut estimate), when the volatility term structure is known and appears to be fairly flat, or when changes in short dated forward prices are strongly correlated and of similar magnitude to simultaneous changes in longer dated forward prices.

Short-Term Supply-Demand Shifts

An alternative to assuming that new information affects the whole forward curve is to assume that new information only affects the near-term forward prices. This assumption is based on the observation that near-term forward prices are typically much more volatile than the longer term ones for many commodities. Intuitively, this assumes that most of the information coming in only affects the short-term supply-demand balance and does not affect longer term forecasts. Mathematically, this model is

$$\frac{dF_{tT}}{F_{tT}} = \beta(t)e^{-a(T-t)}dz_t$$

where the initial forward curve F_{0T} is assumed to be known and the volatility term structure is related to the model parameters through

$$\sigma(t)^2 = \frac{1}{t}\int_0^t e^{-2a(t-s)}\beta(s)^2 ds$$

When $\beta(t)$ is constant, the volatility is equal to

$$\sigma(t) = \beta\sqrt{\frac{1-e^{-2at}}{2at}}$$

which decreases smoothly over time and which fits the volatility term structure of many commodities quite well over short to medium time scales.

The model parameters can be matched to market data by fitting a and $\beta(t)$ to the above volatility equation or by fitting them to changes in the forward prices over time through

$$\frac{1}{dt}\mathrm{var}\left(d\ln(F_{tT})\right) = \beta(t)^2 e^{-2a(T-t)}$$

$$\frac{1}{dt}\mathrm{covar}\left(d\ln(F_{tT}), d\ln(F_{tS})\right) = \beta(t)^2 e^{-a(T-t)}e^{-a(S-t)}$$

$$d\ln(F_{tT}) \equiv \ln(F_{t+dt\,T}) - \ln(F_{tT})$$

Both of these approaches are usually carried out assuming $\beta(t)$ is constant.

The spot price dynamics model for this case is

$$P_t \equiv F_{0t}\exp(-\tfrac{1}{2}\sigma(t)^2 t + w_t)$$

$$dw_t = -aw_t dt + \beta(t)dz_t \quad w_0 = 0$$

which differs from the previous case due to the new term which makes w_t revert back to zero with a time scale of a^{-1} rather than allowing it to randomly walk to any value. It is this term which causes the volatility term structure to decrease over time and which leads to this model being called "mean reverting." Hull and White have developed a tree algorithm for very efficiently modeling the outcomes of w_t over time, which can be used to value assets using this model.[6]

This model is most appropriate when either the volatility term structure is known and seen to be decreasing over time or changes in short dated forward prices are reasonably correlated with, but of much larger magnitude than, simultaneous changes in longer dated forward prices.

Short- and Long-Term Supply-Demand Shifts

A better but more complicated assumption for how information arrives is to assume that it typically affects both short term and long term supply-demand shifts. This model combines the two previous models into the following two factor model

$$\frac{dF_{tT}}{F_{tT}} = \beta_S(t)e^{-a(T-t)}dz_t^S + \beta_L(t)dz_t^L$$

which includes uncertain factors dz^S_t and dz^L_t describing the arrival of short-term and long-term supply-demand shift information, respectively. This model has a volatility term structure related to its parameters through

$$\sigma(t)^2 = \frac{1}{t}\int_0^t \left(e^{-2a(t-s)}\beta_S(s)^2 + \beta_L(s)^2 + 2e^{-a(t-s)}\rho_{LS}(s)\beta_L(s)\beta_S(s)\right)ds$$

In the case where the unknown functions $\beta_L(t)$, $\beta_S(t)$, and $\rho_{LS}(t)$ are constants, the user finds that

$$\sigma(t)^2 = \beta_L{}^2 + \beta_S{}^2\left(\frac{1-e^{-2at}}{2at}\right) + 2\rho_{LS}\beta_L\beta_S\left(\frac{1-e^{-at}}{at}\right)$$

which fits the volatility term structure of many commodities quite well over fairly wide time scales. Notice that this model allows the short and long dated forward prices to move largely independently of each other.

The model parameters can be matched to market data by fitting a, $\beta_L(t)$, $\beta_S(t)$, and $\rho_{LS}(t)$ to the above volatility equation or by fitting them to changes in the forward prices over time

$$\frac{1}{dt}\mathrm{var}(d\ln(F_{tT})) = \beta_L(t)^2 + \beta_S(t)^2 e^{-2a(T-t)} + 2\rho_{LS}(t)\beta_L(t)\beta_S(t)e^{-a(T-t)}$$

$$\frac{1}{dt}\mathrm{covar}(d\ln(F_{tT}), d\ln(F_{tS})) = \beta_L(t)^2 + \beta_S(t)^2 e^{-a(T-t)}e^{-a(S-t)}$$

$$+\rho_{LS}(t)\beta_L(t)\beta_S(t)(e^{-a(T-t)} + e^{-a(S-t)})$$

$$d\ln(F_{tT}) \equiv \ln(F_{t+dt\,T}) - \ln(F_{tT})$$

Both approaches are usually carried out assuming $\beta_L(t)$, $\beta_S(t)$, and $\rho_{LS}(t)$ constants.

The spot price dynamics model for this case is

$$P_t \equiv F_{0t}\exp(-\tfrac{1}{2}\sigma(t)^2 t + w_t^S + w_t^L)$$

$$dw_t^S = -aw_t^S dt + \beta_S(t)dz_t^S$$

$$dw_t^L = \beta_L(t)dz_t^L$$

$$w_0^S = w_0^L = 0$$

which differs from the previous case in that there are now two uncertain factors for which simultaneous outcomes must be tracked. Notice that the short term factor is mean reverting, just as in the previous case while the long term factor follows a random walk, as in the first case. Hull and White have also developed a tree algorithm for very efficiently modeling the joint outcomes of w^S_t and w^L_t over time, which can be used to value assets using this model.[7] This model is about at the limit of what can be reasonably solved and possibly beyond it for the case of assets whose value depends on two or more commodity prices.

This model is most appropriate when either the volatility term structure is known and seen to level off after decreasing over time or when changes in short dated forward prices are not highly correlated with simultaneous changes in longer dated forward prices.

Short- and Medium-Term Supply-Demand Shifts

A final, related model of note is the two factor model corresponding to both short-term and medium-term information arriving. This model is identical to the above two factor model except for the addition of a slow decay at rate b for the long-term uncertainty (making it really medium term). The forward curve dynamics equation is

$$\frac{dF_{tT}}{F_{tT}} = \beta_S(t)e^{-a(T-t)}dz_t^S + \beta_L(t)e^{-b(T-t)}dz_t^L$$

and the volatility is related to the model parameters through

$$\sigma(t)^2 = \frac{1}{t}\int_0^t \Big(e^{-2a(t-s)}\beta_S(s)^2 + e^{-2b(t-s)}\beta_L(s)^2 + 2e^{-(a+b)(t-s)}\rho_{LS}(s)\beta_L(s)\beta_S(s)\Big)ds$$

In the case where the unknown functions $\beta_L(t)$, $\beta_S(t)$, and $\rho_{LS}(t)$ are constants, the user finds that

$$\sigma(t)^2 = \beta_L^2\left(\frac{1-e^{-2bt}}{2bt}\right) + \beta_S^2\left(\frac{1-e^{-2at}}{2at}\right) + 2\rho_{LS}\beta_L\beta_S\left(\frac{1-e^{-(a+b)t}}{(a+b)t}\right)$$

The main difference in the behavior of this volatility term structure compared with the previous model's behavior is that the volatility never levels off in the long term but continues to decrease slowly.

The model parameters can be matched to market data by fitting a, b, $\beta_L(t)$, $\beta_S(t)$, and $\rho_{LS}(t)$ to the above volatility equation or by fitting them to changes in the forward prices over time through

$$\frac{1}{dt}\operatorname{var}(d\ln(F_{tT})) = \beta_L(t)^2e^{-2b(T-t)} + \beta_S(t)^2e^{-2a(T-t)} + 2\rho_{LS}(t)\beta_L(t)\beta_S(t)e^{-(a+b)(T-t)}$$

$$\frac{1}{dt}\operatorname{covar}(d\ln(F_{tT}), d\ln(F_{tS})) = \beta_L(t)^2e^{-b(T-t)}e^{-b(S-t)} + \beta_S(t)^2e^{-a(T-t)}e^{-a(S-t)} + \rho_{LS}(t)\beta_L(t)\beta_S(t)\left(e^{-a(T-t)}e^{-b(S-t)} + e^{-b(T-t)}e^{-a(S-t)}\right)$$

$$d\ln(F_{tT}) \equiv \ln(F_{t+dt,T}) - \ln(F_{tT})$$

Both approaches are usually carried out assuming $\beta_L(t)$, $\beta_S(t)$, and $\rho_{LS}(t)$ are constants.

The spot price dynamics model for this case is

$$P_t \equiv F_{0t}\exp(-\tfrac{1}{2}\sigma(t)^2t + w_t^S + w_t^L)$$

$$dw_t^S = -aw_t^S dt + \beta_S(t)dz_t^S$$

$$dw_t^L = -bw_t^L dt + \beta_L(t)dz_t^L$$

$$w_0^S = w_0^L = 0$$

which differs from the previous case only in that now both uncertain factors are mean reverting. The same tree methods as above can be used with the same caveats.

This model would be appropriate when the volatility term structure decreases quickly and then switches to a slower decrease. The model would also be appropriate when changes in short dated forward prices are not highly correlated with simultaneous changes in longer dated forward prices and the magnitude of the changes for long dated forwards decreases with delivery date at a slower rate than for short dated forwards. However, it is doubtful that good forward price estimates will extend to sufficiently distant delivery dates to allow the user to estimate the parameters for this model. In general, it is recommended that one of the previous models be used unless the data suggesting the application of this model is truly compelling.

Example

Natural gas is probably the most complex commodity for which forward prices are readily available.[8] Natural gas prices are highly seasonal and deliverability constraints lead to drastically differing prices in differing locales. Daily forward price data for natural gas can be found in the financial pages of many newspapers.[9] For example, the forward price curve for natural gas deliveries to Henry Hub on August 21, 1997, can be seen in *Figure 1*. Notice the significant peaks in the two winters, illustrating the seasonality of the commodity.

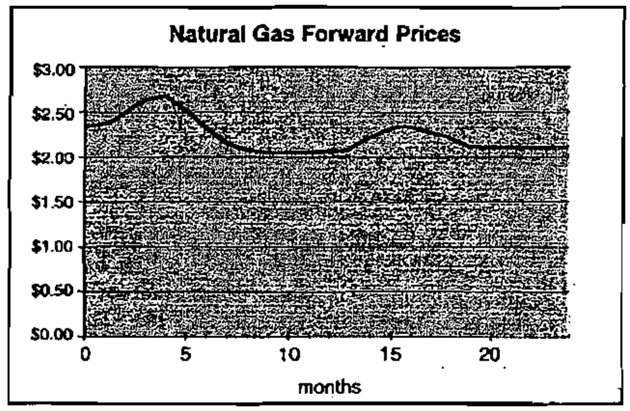

Figure 1: Natural gas forward prices on August 21, 1997.

An examination of the daily forward price changes reveals substantial information about the market price dynamics. The volatility of a particular forward price can be estimated as the square root of the variance of daily changes in the natural logarithm of prices ("log prices") multiplied by $\sqrt{250}$ to convert it to an annual volatility.[10] *Figure 2* shows the volatility of the forward prices in *Figure 1* computed over 20 trading days[11] of prior data. Notice that the volatility of the near-term prices is much higher than the volatility of the long-term ones. Notice also that the volatility seems to level off after about a year or so. The user can also examine how much the daily changes in log prices are correlated between differing contracts. *Figure 3* shows the correlation of log price changes of various forward prices with the front month price from the same data set. Notice that the later prices are not highly correlated with the front ones.

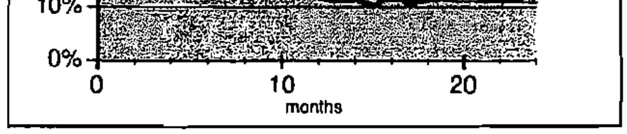

Figure 2: Natural gas forward price volatilities.

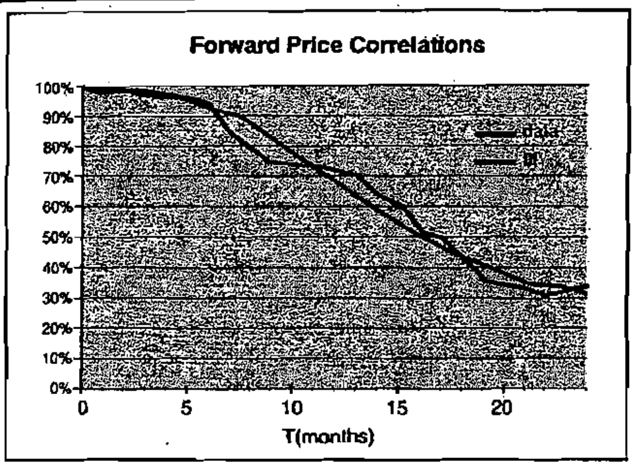

Figure 3: Correlations of natural gas forward prices with spot.

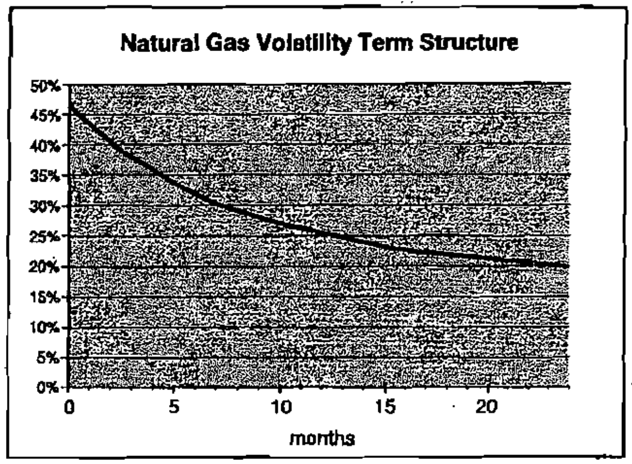

Figure 4: Natural gas volatility term structure.

Both the leveling off of the forward price volatilities and the low correlation between the long dated and near dated forward prices, suggest that the Short- and Long-Term Supply-Demand Shift Model is the most appropriate over the two-year time scale shown. The four parameters needed to describe this model[12] are the long-term volatility β_L, the short-term volatility β_S , the mean reversion parameter a, and the correlation between long- and short-term curve shifts ρ_{LS}. These parameters can be fit[13] to the forward price volatility and correlation[14] data of *Figures 2 and 3* using

$$\frac{1}{dt}\operatorname{var}(d\ln(F_{tT}))=\beta_L^{\,2}+\beta_S^{\,2}e^{-2a(T-t)}+2\rho_{LS}\beta_L\beta_S e^{-a(T-t)}$$

$$\frac{1}{dt}\operatorname{covar}(d\ln(F_{tT}),d\ln(F_{tS}))$$

$$=\beta_L^{\,2}+\beta_S^{\,2}e^{-a(T-t)}e^{-a(S-t)}+\rho_{LS}\beta_L\beta_S(e^{-a(T-t)}+e^{-a(S-t)})$$

$$d\ln(F_{tT})\equiv\ln(F_{t+dt,T})-\ln(F_{tT})$$

The values of the parameters can be found in *Table 1* and the model fit can be seen compared to the data in *Figures 2 and 3*. Notice the visibly excellent fit to the data despite the simplicity of the model.

Table 1
Parameters of the model fit to the natural gas data in Figures 2 and 3.

Model Parameters	
β_L	11%/yr
β_S	45 %/yr
a	0.5 yr
ρ_{LS}	0%

Once the model parameters have been fit, the volatility term structure can be computed using

$$\sigma(t)^2=\beta_L^{\,2}+\beta_S^{\,2}\left(\frac{1-e^{-2at}}{2at}\right)+2\rho_{LS}\beta_L\beta_S\left(\frac{1-e^{-at}}{at}\right)$$

This, in turn, allows one to compute the prices of various European options to be computed. The volatility term structure corresponding to the model parameters in *Table 1* can be seen in *Figure 4*.

The model parameters can also be used to create tree-based evaluation models to determine the value of American options and other complex assets. In this case, the dynamics equations for w^S_t and w^L_t

$$dw_t^S=-aw_t^S dt+\beta_S dz_t^S$$
$$dw_t^L=\beta_L dz_t^L$$
$$w_0^S=w_0^L=0$$

would be used to construct a two factor Hull-White tree to obtain their joint outcomes at each point in time and then

$$P_t\equiv F_{0t}\exp(-\tfrac{1}{2}\sigma(t)^2t+w_t^S+w_t^L)$$

would be used to obtain the corresponding spot price outcomes and hence the cash flows at various times and uncertainty outcomes.

To Summarize

This document reviews financial rules for valuing assets whose value depends on traded commodities. The rules show that only information about the forward curve and volatility term structure is needed to value a fairly wide range of assets. Although forward curve data is relatively available, volatility information is more difficult to attain and often requires a price dynamics model to interpret available data on changes in spot or forward prices over time. Furthermore, other common assets (e.g., American options) require additional information from a price dynamics model to value them.

Commonly used price dynamics models, adopted from the financial markets, describe in simplified terms how spot prices change over time. Commodity markets typically have quite complicated spot price behaviors, much of which is characterized in the shape of the forward price curve. This insight leads to the forward curve dynamics approach which begins with models of how the forward curve shifts over time and uses them as the basis for valuation.

In the forward curve dynamics approach, the forward curve is an input to the valuation as is either the volatility term structure (when available) or historic movements of forward prices. If the volatility term structure is available, then only it and the forward curve are needed to value many assets. If the volatility term structure is not available, then the forward curve dynamics approach provides a modeling framework for using historical shifts in the forward curve to estimate the volatility term structure. In either case, other common assets require models of the full spot price dynamics to value them, and forward curve dynamics provides such models based on the detailed shape and behavior of the forward curve.

The document discusses several forward curve dynamics models and provides guidance was as to which model is appropriate based on particular aspects of the available data inputs. The Long-Term Supply-Demand Shift Model is the basic starting point for valuations, with extremely limited data available to calibrate it. This model is just the Random Walk Model from finance in another guise. In cases where there is a reasonable amount of calibration data, the Long-Term Supply-Demand Shift Model or the Short-Term Supply-Demand Shift Model will be the best choice, depending on the relative magnitudes of variation in short dated and long dated forward prices. In cases where there is sufficiant data to ensure that the variations in short dated and long dated forwards are not highly correlated or to see the volatility term structure level off, the user may consider switching to the Short- and Long-Term Supply-Demand Shift Model if it will have significant impact on the results. Finally, the document provides an example of applying these models to natural gas data. For that data set, the Short- and Long-Term Supply-Demand Shift Model was the most appropriate and provided excellent fit to the market data. ■

For Further Information

EPRI's Power Markets & Resource Management (PM&RM) Target is working to develop methods and tools for calculating the market value of assets, customers, and contracts, and to manage their associated risks. For further information about this or other PM&RM work, please feel free to contact

Richard E. Goldberg
Manager, Pricing & Risk Management
650-855-2397
rgoldber@epri.com

Victor Niemeyer
Manager, Power Markets
650-855-2744
niemeyer@epri.com

Art Altman
Manager, Risk Management
650-855-8740
aaltman@epri.com

Footnotes

[1] The lognormal distribution is a probability distribution in which the natural logarithm of the variable is distributed according to the normal distribution. Accordingly, defining the expected value of the variable and the variance of its natural logarithm fully define the distribution. One of the benefits of this distribution is the fact that only positive values of the variable are allowed.

[2] The Black Equation is the version of the Black-Scholes Equation for commodity markets.

[3] Most exchange-traded options for energy commodities are American "options on futures" meaning that striking the option leads to exchange of a futures contract rather than the commodity itself. The prices of these American options are also commonly inverted using the Black Equation to define implied volatilities which can lead to an overestimate of the volatility but the error is not expected to be severe.

[4] Ito calculus provides rules for integrating uncertain functions over time. These rules differ from those of standard calculus due to the effects of uncertainty accumulating over time. In the context of price modeling, Ito calculus defines the procedure by which a model describing the instantaneous, uncertain changes in a price over time ("price dynamics") can be used to calculate the probability distributions of future prices. See Hull, "Options, Futures, and other Derivative Securities", Prentice-Hall for an introduction to these methods.

[5] See the boxed set EPRI AP-107748 "Valuing Generation Assets in Uncertain Markets" for an introduction to and examples of both the Binomial Tree and Finite Difference methods for valuing assets. (Please note: to apply the Binomial Tree method with time-valuing volatility, one needs to use time-varying time steps so that over each time interval $\beta(t)^2 \Delta t$ is constant.)

[6] See J. Hull & A. White, *The Journal of Derivatives*, Fall 1994, pp. 7-16. Caveat lector: almost all of the literature in this area is geared towards interest rate or cross-currency derivatives.

[7] See J. Hull & A. White, *The Journal of Derivatives*, Winter 1994, pp. 37-48.

[8] Electricity is expected to be an even more complex commodity than natural gas but forward price data is currently too limited to make a good example.

[9] Futures exchange prices are good estimates of the forward price (when interest rates are much less uncertain than prices of the commodity in question). In particular, the New York Mercantile Exchange (NYMEX) Henry Hub futures prices for natural gas are widely available and will be used for this example.

[10] 250 represents the approximate number of trading days in a year.

[11] When computing historical volatilities one is forced to decide on how long a series to examine: too short and the result is overly biased by near term events, too long and irrelevant past events will distort the results. 20 days is an arbitrary choice meant to incorporate about a month's worth of data.

[12] Assuming that the volatilities and correlation do not change seasonally.

[13] The fit was done by eye but a least-squares fit didn't appreciably change the results.

[14] Recall that the correlation between two variables is defined as their covariance divided by the square root of the product of their variances.

EPRI
Electric Power
Research Institute
Post Office Box 10412, Palo Alto, California 94303 (650) 855-2000

October 1997 TB-108991