Application of Multi-Factor Models

Glen Swindle

June 24, 2013

Introduction

- The Issues
- Market Mechanics
- Spot Price Dynamics
- Application of a Two-Factor Model
- Extrapolation Taken Too Far?

The Issues

Why Do We Need Multi-Factor Models?

- Variety of Options Structures
  - Most franchise energy desks end up dealing in a broad spectrum of structures.
  - A unified vol model is usually better than separate one-factor models for each option type.
- Correlation Structures
  - By construction a single factor model has unit correlation of returns between contracts.
  - The presence of a nontrivial second factor (PCA) should clearly effect valuation of, for example, spread options.

The Issues

Variety of Options Structures

- Energy markets frequently have options with different frequency exercise for the same delivery period.
  - Do you treat each of these separately with its own one-factor model or do you couple them?

- Separate model approach: Independent models for each option type.
  - Benefits:
    - Simple and unambiguous to deploy.
    - Arguably adequate for prop or flow trading of standard options.
  - Liabilities:
    - If you are trading many options products (asians, quarterly swaptions,...) the number of vol surfaces can be very large.
    - All of the vegas of the different options types are unrelated; there is no unifying model and hedging a less liquid vega exposure with a more liquid is outside of the modeling paradigm.

The Issues

Variety of Options Structures

- Multi-Factor Approach: Single underlying modeling framework.
  - Benefits:
    - All options are related. Greeks can be systematically netted and managed.
    - The number of parameters changing daily (or real-time) is of manageable size.
    - Uses liquid market data to sensibly infer prices of less liquid instruments.
  - Liabilities:
    - Challenging to calibrate. Often the slower moving parameters require nonlinear optimizations.
    - Traders with a less analytical bent tend not to like it very much.
    - Overkill for vanilla flow books.

The Issues

Extrapolation

- The primary purpose of a "holistic" risk-neutral model is to extrapolate from liquid instruments to illiquid instruments.
- We have already seen this in several contexts.
  - Vol Backwardation
    - Our simple one-factor approach to this problem:

      $$dF(t,T) = \sigma(T-t)F(t,T)dB_t$$

      was used to extrapolate from the prices of vanilla options to valuation of options with non-standard expiry.
    - The risk was in the functional form used for $\sigma(T-t)$.
    - Only the magnitude of this function was calibrated to market data; the rate of backwardation was estimated statistically.
    - Given only say monthly options there is little reason to proceed to more complex models.
    - If, however, other options are liquidly traded (e.g. swaptions or daily options), a one-factor model does not have the flexibility to process such information and make useful inference on backwardation.

The Issues

Extrapolation

- Examples (cont)
  - Vol Skew (Single Commodity Vanilla)
    - A set of strikes is used to mark a vol surfaces.
    - Pricing a vanilla option at a different strikes involves interpolating (easy) or extrapolating (hard) the vol grid.
    - Extrapolation can be accomplished either by parametric fits of liquid surfaces or by using a stochastic process that generates "realistic" skews.
  - Vol Skew (Spread Options)
    - The use of vol-lookup by moneyness of opposing leg was purely heuristic approach.
    - There was no single underlying process that justified this approach.

The Issues

Extrapolation

- Examples (cont)
  - CSOs
    - Most activity is confined to the first few contract months.
    - From these we could calibration term correlations. But what do we do about inference to other non-traded pairs of contracts?
    - There is a non-trivial behavior of local correlation; see the plot below for daily WTI returns correlation (2007-2010).

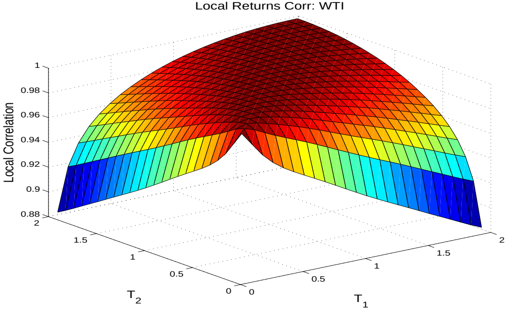

The Issues

Extrapolation

- Examples (cont)
  - Tolls
    - A similar question arises in cross-commodity spread options.
    - If there are broker quotes for tolls it is almost always within a tenor of 2 years.
    - There is a non-trivial behavior of local correlation between commodities.
    - The following shows the correlation of returns between July PJM and NYMEX NG as a function of tenor to expiry (2007-2010).

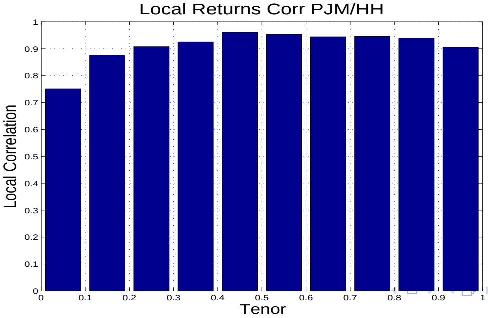

The Issues

Working Problem

- Price an Oct2010 forward starting daily straddle on PJM Western Hub 5x16 power as of pricing date 03Feb2010 .
- Details:
  - The quantity will be standard lot size of 50 MW.
  - Monthly Index: ICE PJM Western Hub 5x16 settlement on '-1b'.
  - Daily Index: PJM RT 5x16 daily average price as published by the PJM ISO.
- Comment:
  - This is not a common structure in power.
  - Forward starters are more common in NG markets than fixed strike daily options.

Market Mechanics

Linear Instruments

- Fixed price swaps settle on the monthly average realized spot price for a prescribed bucket.
- Eastern US:
  - Peak bucket is 5x16 (M-F 7AM to 11PM).
  - Offpeak (the "offpeak wrap") is the complement.
  - Often circumstances require trading (or at least viewing) the offpeak wrap as 2x16 and 7x8. These are highly illiquid swaps.
  - Standard contract size: 50MW.
- Western US:
  - Peak bucket is 6x16 (M-Sat 7AM to 11PM).
  - Offpeak (the "offpeak wrap") is the complement.
  - Standard contract size: 25MW.
- Texas (ERCOT):
  - Peak bucket is 5x16 (M-F 6AM to 10PM).
  - Offpeak (the "offpeak wrap") is the complement.
  - Standard contract size: 50MW.

Market Mechanics

Linear Instruments: Heatrates

- The term "heatrate" refers to a ratio of power price to a natural gas price.
  - The ratio can pertain to market prices, either spot or forward ("market heatrates").
  - It can also refer (as we'll see shortly) to engineering specs of a generator.
- Power swaps often trade as heat-rates:
  - A power buy/sell is associated with a sell/buy of a specified volume of natural gas, both over the same delivery period.
  - The trade will be quoted in heatrate units: $\frac{F_P(t,T)}{F_G(t,T)}$, where P and G denote power and gas respectively.

Market Mechanics

Linear Instruments: Heatrates

- The following figure shows PJM forward prices and PJM/Henry Hub forward heatrates.
  - The seasonality in both arises from (well-founded) expectations that loads will be higher in summer and to a lesser extent winter months.
  - This forces expected clearing prices higher up the stack into more expensive units.

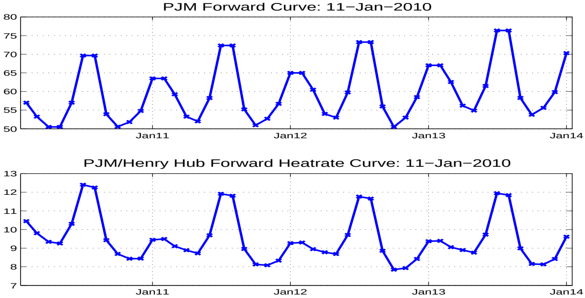

Market Mechanics

Options

- Annual Swaptions
  - Same structure as natural gas in mechanics/valuation.
- Monthly options:
  - Same structure as natural gas.
  - Standard expiry: '-2b'
- Daily Fixed Strike
  - Financially settling and usually manually exercised day-ahead ('-1b')
  - Example: Cash settled value of an auto-exercised call:

$$\sum_{d \in m} \max \left[F(d,d) - K, 0 \right]$$

  - Example: Cash settled value of standard exercise:

$$\sum_{d \in m} 1_{\{E_d\}} \left[F(d,d) - K \right]$$

    where E_d denotes exercise events.
  - Dominant vol exposure: daily (spot).

Market Mechanics

Options

- Monthly Options:
  - Exercises at time (T_e) before the beginning of the delivery month into either:
    - A physical forward with delivery during the contract month at a price that is the strike K;
    - Cash settlement based on the value $F_m(T_e) - K$.
  - For example, the value of a call (in either case) is: $d(0, T)\tilde{E}\left(\max\left[F_m(T_e) - K, 0\right]\right)$.

- Daily Options:
  - A set of daily options usually exercising one business day before delivery.
    - These are usually financial settling on the spread between spot price and the strike: $F(t, t) - K$ where t indexes the delivery day.
  - For example, the value of a call is: $d(0, T)\tilde{E}\left(\sum_{t \in m} \max\left[F(t, t) - K, 0\right]\right)$.
  - The notation $t \in m$ denotes the active delivery days in the option during month m. For peak (5x16) power options the sum would span the (roughly 20-22) business days in the particular month.

Market Mechanics

Options

- Forward Starters:
  - As with daily options these constitute a sequence of daily options.
  - The difference is that the strike is set based on the prevailing price for monthly delivery just before the month begins; hence the term "forward starter" as the strike is floating up until the delivery month.
    - We will assume that this is the same as the standard monthly option expiry date $F_m(T_e)$.
    - Settlement is on the spread between spot price and the strike: $F(t,t) - F_m(T_e)$ where t once again indexes the delivery day.
  - The value of a call is: $d(0,T)\tilde{E}\left(\sum_{t\in m} \max\left[F(t,t) - F_m(T_e), 0\right]\right)$.

Spot Price Dynamics

Vernacular and Convention

- The terms "monthly vol" ($\bar{\sigma}_M$) and "daily vol" ($\bar{\sigma}_D$) will always refer to implied vols pertaining to the fixed strike options (the first two categories above).
- Market dichotomy:
  - Power: Monthly and daily fixed strike options are the vanilla options.
  - Natural Gas: Monthly and the forward starter are the vanilla options.

Spot Price Dynamics

Volatility Structure in Power

- Power and natural gas exhibit seasonality in volatility similar to seasonality in price.
  - The following figure shows the ATM monthly and daily implied vols for PJM Western Hub on03Feb2010 .

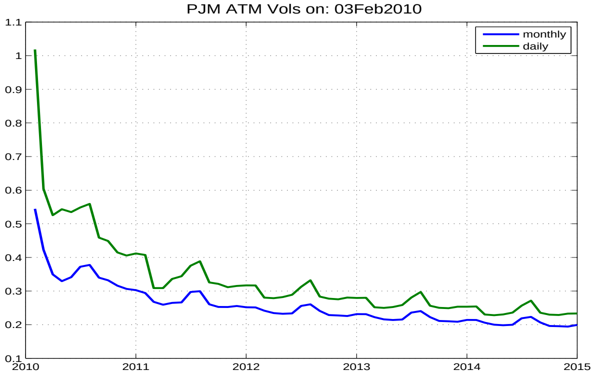

Spot Price Dynamics

Volatility Structure in Power

- The following figure shows the skew for monthly and daily options in PJM for Jul10 .

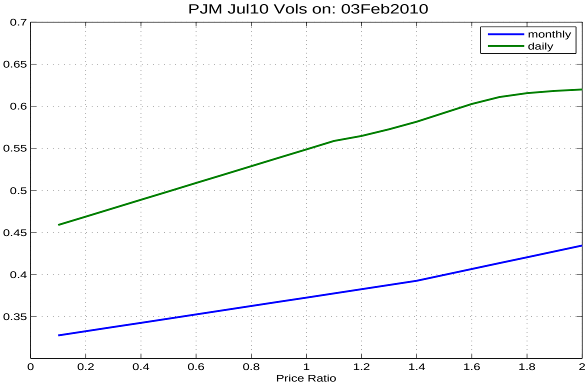

Spot Price Dynamics

Spot Volatility

- A simple caricature of spot (daily) prices processes is for the daily returns for the daily spot prices to be i.i.d. normal:

$$F(t,t) = F_m(T_m)e^{\zeta Z_t - \frac{1}{2}\zeta^2}$$

  where

  - F_m denotes the contract month containing day t
  - Z_t is standard normal.
  - ζ is the spot volatility.

Spot Price Dynamics

Spot Volatility

- Spot volatility is implied by the monthly and daily vols

$$\zeta^2 = \bar{\sigma}_D^2 T_d - \bar{\sigma}_M^2 T_m$$

where $T_d \approx T_m + \frac{1}{24}$.

  - The following figure shows the ATM spot vol by contract month for PJM

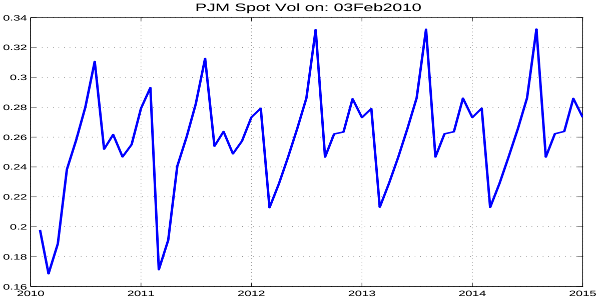

Spot Price Dynamics

Spot Volatility

- For comparison the following figure shows historical spot volatility for PJM and TETM3.
  - The reference price is the BOM contract fixing so that spot returns are defined as $\log\left[\frac{p_d}{F_m(T_e)}\right]$.

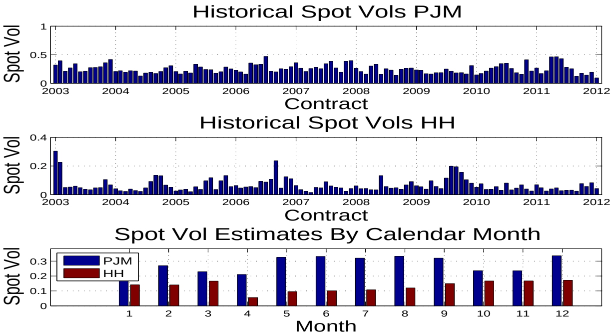

Application of a Two-Factor Model

Two-Factor Gaussian Exponential

- Recall this form:

$$\frac{dF(t,T)}{F(t,T)} = \sum_{j=1}^{2} \left[\sigma_j(T) e^{-\beta_j (T-t)} dB_s^{(j)} \right]$$

  - We have written this in "Big-T" form for now.
  - We will often assume that the BMs are independent for simplicity.

- Intuition:

  - If $\sigma_2 \equiv 0$, this is a one-factor model identical to that described in the first section:

$$\sigma(T-t) = \alpha e^{-\beta(T-t)}$$

  - The second factor will typically have $\beta_2 \ggg \beta_1$ and is intended to represent shorter time-scale forward returns.

Application of a Two-Factor Model

Movation

- The motivation behind this model is that:
  - The first factor will be the "long" slowly decaying perturbation to forwards.
  - The second "short" factor will have a larger β and will accommodate the relatively high values of spot volatility.
  - The previous simple model with i.i.d. spot perturbations is formally this model with $\beta_1 = 0$ and $\beta_2 = \infty$

Application of a Two-Factor Model

Calibration: General Comments

- Modeling framework:
  - Number of factors
  - Choice of β's

  depends upon the market and what you trade.

- Moving to multi-factor models rarely eliminates the need for estimation and/or judgement regarding parameter selection.
  - There are rarely enough liquid tradeables to reliably calibrate all parameters.
  - Your goal is a model with enough flexibility to calibrate to common tradables and get reasonably close to infrequently observed prices of less standard structures.

Application of a Two-Factor Model

Calibration: A Market Dichotomy

- Long Time-Scale Markets
  - Most structures involve time-scales at or beyond monthly delivery.
  - Example: Crude oil in which swaps and options refer to current nearby forward contracts.
  - The primary modeling goal is valuation/hedging of cross-contract correlation structures.
  - Parameters: Mean-reversion rates can be estimated by MLE using historical forwards/futures or set with by fitting Asian options and swaptions.
  - Typically $\beta_1 \in [.1, .5]$ and $\beta_2 \in [1, 10]$

Application of a Two-Factor Model

Calibration: A Market Dichotomy

- Long Time-Scale Markets
  - Allowing for the BMs to be correlated the two factor model implies a correlation surface:

$$\rho(T,S) = \frac{\sigma_1^2 e^{-\beta_1(T+S)} + \sigma_1\sigma_2\rho\left[e^{-(\beta_1 T+\beta_2 S)} + e^{-(\beta_2 T+\beta_1 S)}\right] + \sigma_2^2 e^{-\beta_2(T+S)}}{\sigma(T)\sigma(S)}$$

  - Since the returns of the T contract can be written as:

$$\sigma_1 e^{-\beta_1 T}\left[dB_t^{(1)} + \frac{\sigma_2}{\sigma_1} e^{(\beta_1-\beta_2)T} dB_t^{(2)}\right]$$

    the free parameters in $\rho(T,S)$ are:
    - The difference in the decay rates $\beta_2 - \beta_1$.
    - The volatility ratio $\lambda = \frac{\sigma_2}{\sigma_1}$
    - The correlation ρ

Application of a Two-Factor Model

Calibration: A Market Dichotomy

- Long Time-Scale Markets
  - Minimizing the Frobenius norm of the difference of the empirical and model covariance matrices using WTI returns from Jan2007 to Dec2010 yields:
    - The optimal decay rates are: $\bar{\beta} = [0.106, 1.528]$
    - The estimated correlation between the factors is $\rho = 0.119$

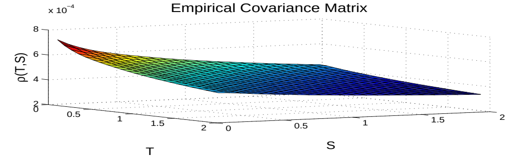

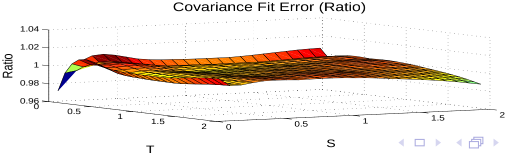

Application of a Two-Factor Model

Calibration: A Market Dichotomy

- Multiple Time-Scale Markets:
  - Common tradeables involve daily time-scales usually referencing spot prices.
  - Example: Power and, often by association through tolls, natural gas.
  - The primary modeling goal is valuation/hedging of cross-commodity correlation structures usually with spot (daily) price exposure.
  - Parameters: Mean-reversion rates are harder to estimate due to relevance of both forward and spot prices.
  - Typically $\beta_1 \in [.1, .5]$ and $\beta_2 \ggg 10$.

Application of a Two-Factor Model

Working Problem: Calibration

- We have two factors with two volatilities per contract month, namely monthly and daily volatility.
- Calibration involves setting σ_1 and σ_2 to be consistent with the observed monthly and daily option prices.
- Applying a few reasonable approximations this ammounts to solving:

$$\begin{pmatrix} D_1 & D_2 \\ M_1 & M_2 \end{pmatrix} \begin{pmatrix} \sigma_1^2 \\ \sigma_2^2 \end{pmatrix} = \begin{pmatrix} \bar{\sigma}_D^2 \\ \bar{\sigma}_M^2 \end{pmatrix}$$

- Procedure:
  - Obtain $\bar{\sigma}_M$ and $\bar{\sigma}_D$ from market data.
  - Compute D_1, D_2, M_1, M_2.
  - Solve the linear system above.
  - Use the resulting σ_1 and σ_2 for structure valuation.

Application of a Two-Factor Model

Mean-Reversion Rates

- Selection of β's requires estimation.
  - Liquid asians or non-standard expiry options would potentially mitigate the need for an empirical basis for β selection.
  - In the absence of such we can either examine term structure of volatility or consider "fundamental" reasons for selection.
  - The following plot shows the implied β versus tenor for PJM WH from historical returns using data from 01Jan2007.
  - The asymptotic implied β is approximately .3.

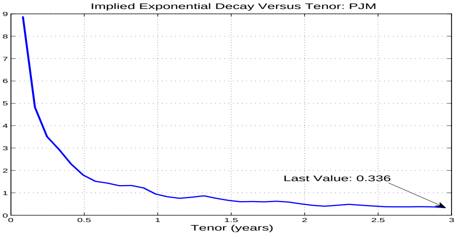

Application of a Two-Factor Model

Mean-Reversion Rates

- The following figure shows σ_1 and σ_2 for Oct2010 for varying β_2 with $\beta_1 = .3$.
- Note the increase in both σ's as the higher β_2 requires more factor vol to hit the required implied vols.

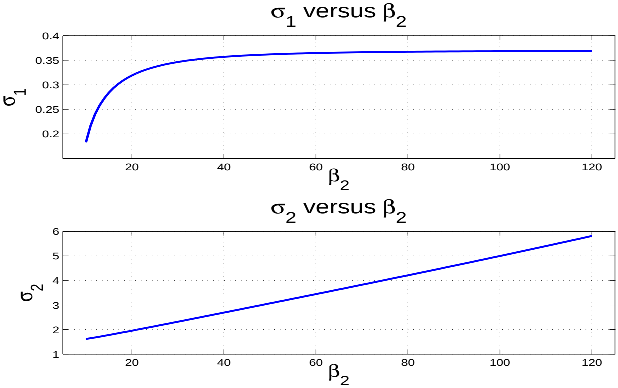

Application of a Two-Factor Model

ACF of Spot Returns

- Spot returns statistics give us some guidance to β_2:
- What are spot returns for a non-storable commodity?
  - A useful definition for spot returns: $\log\left[\frac{p_d}{F_m(T_e)}\right]$
  - The plot below shows both the ACF as well as the implied β by lag.
  - The results are of comparable magnitude to the weather analysis.

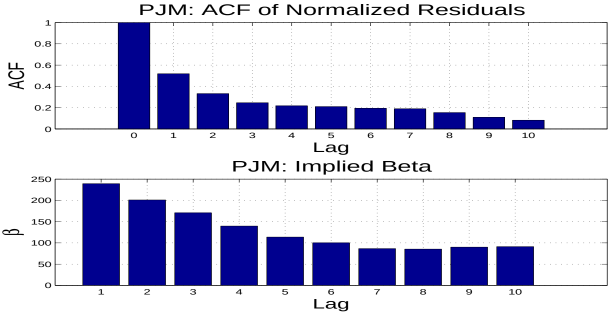

Application of a Two-Factor Model

Working Problem

- Recall that our goal is to price a Jul08 forward starting daily option on PJM Western Hub 5x16 power.
  - We need to evaluate:

$$\tilde{E}\left[\sum_d (F(T_d, T_d) - F_m(T_f))^+\right]$$

    where T_f refers to the date of fixing ('-1b') F_m denotes the monthly contract price.
  - Focusing on one particular day d and conditioning on $\mathcal{F}_{T_f}$ we first evaluate:

$$\tilde{E}\left[(F(T_d, T_d) - F_m(T_f))^+ \,|\mathcal{F}_{T_f}\right]$$

Application of a Two-Factor Model

Working Problem

- Conditional Value:
  - Since our processes are log-normal, this is rote application of Black for an ATM option—for example, the call value is:

$$C(T_f, F) = e^{-r(T_d - T_f)} \left(F_m(T_f)\Phi(d_1) - F_m(T_f)\Phi(d_2) \right)$$

    where $d_{1,2} = \pm\frac{1}{2}\sigma_{T_f,T_d}\sqrt{T_d - T_f}$.

  - It remains to calculate σ_{T_f,T_d} in our multi-factor setting.
    - Note that in the "low-budget" model life is simple: $\sigma_{T_f,T_d} = \zeta$.
  - The evolution of F from T_f to T_d is:

$$F(T_f, T_d) = F_m(T_f) e^{\sum_{j=1}^{2} \left[\int_{T_f}^{T_d} \sigma_j(m) e^{-\beta_j(T_d - s)} dB_s^{(j)} - \text{"ITO-BLOB"} \right]}$$

    .

  - A little integration yields:

$$\sigma^2_{T_f,T_d}(T_d - T_f) = \sum_j \sigma_j^2(m) \frac{1 - e^{-2\beta_j(T_d - T_f)}}{2\beta_j}$$

Application of a Two-Factor Model

Working Problem

- Unconditioned Value:
  - We now have all components for the conditional value:

$$
\begin{aligned}
& \tilde{E}\left[\max\left(F(T_d, T_d) - F_m(T_f), 0\right) \middle| \mathcal{F}_{T_f}\right] \\
= \quad & e^{-r(T_d - T_f)} F_m(T_f)\left(\Phi(d_1) - \Phi(d_2)\right)
\end{aligned}
$$

  - The only random part for the unconditioned expectation is $F_m(T_f)$ and we know that

$$\tilde{E}\left[F_m(T_f)\right] = F_m(0)$$

  - The hard (mechanical) part is putting this all together on precisely the right calendar days.
  - The following figure shows the pricing results for various β_2 given market data as of 03Feb2010.
  - The decrease in value at high values of β_2 is an artifact of weekends:
    - At high mean-reversion rates averaging vols of "daily optionlets" to hit broker vols becomes problematic.

Application of a Two-Factor Model

Working Problem

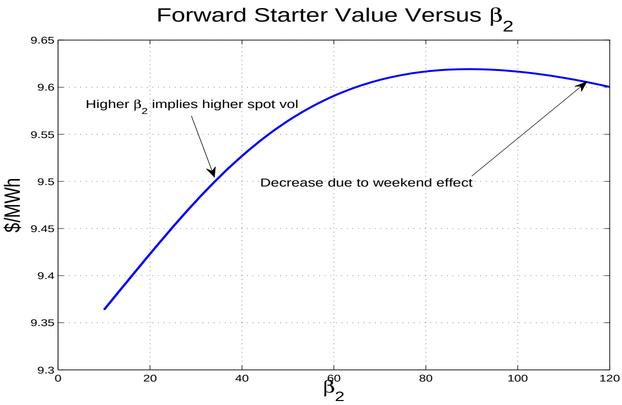

Application of a Two-Factor Model

- The modeling framework above (arguably) successfully accommodates two important features of the energy markets commonly encountered:
  - Backwardation of volatility.
  - Separation of time-scales—slow mean-reversion of general price level (at all tenors); fast mean-reversion of spot prices.
- There remain open issues:
  - Returns are log-normally distributed and skew is handled via "vol-lookup-by-price ratio."
  - The factors were uncorrelated.
    - This is trivial to fix (the formulas are simply optically less pleasant) and harder to calibrate (especially in the spot price setting.
- One other useful feature of this "HJM"-like approach:

  The model is a-priori consistent with the current forward curve.

Extrapolation Taken Too Far?

A Variation of Peaking Options

- Peaking options
  - Generic expression for daily options in which the holder can select a subset of the number of delivery days for exercise.
  - These can be:
    - Auto-exercise (look-back) or manual
    - Fixed strike or forward-starting
- Working Problem:
  - The current date is 01Apr2011 .
  - Value a Henry Hub Gas Daily look-back "peaking" option paying the best J daily forward-starting call payoffs out of the 31 days in Jan12 .

Extrapolation Taken Too Far?

A Variation of Peaking Options

- Working Problem:
  - Let $\mathcal{I}_J$ denotes the subset of the 31 days in Jan12 with the largest J values of

$$\max\left[p_d - F_m(T_e), 0\right]$$

    - T_e denotes the FOM price
    - t_d denotes the spot delivery time for day d
    - p_d the Gas Daily spot index for day d
  - The structure will pay:

$$\frac{1}{J}\sum_{d\in\mathcal{I}_J}\max\left[p_d - F_m(T_e), 0\right]$$

    - We are normalizing the payoff by $\frac{1}{J}$ to afford comparison between structures with different J

Extrapolation Taken Too Far?

A Variation of Peaking Options

- Working Problem:
  - Market Data
    - NYMEX natural gas for Jan12 delivery is trading at $5.131/ MMBtu.
    - The ATM implied vols are: $\bar{\sigma}_M = 0.298$ and $\bar{\sigma}_D = 0.324$
  - Calibration
    - We will use the Gaussian two-factor model with $\bar{\beta} = [0.30, 40]$
    - Calibration to the implied vols yields: $\sigma_{1,2} = [0.34, 1.17]$.
  - Simulation
    - Using Big-T calibration $F(t, T)$ is driven entirely by two OU processes of the form:

$$dY_j = -\beta_j Y_j dt + dB_t^{(j)}$$

    - Given $F_m(0)$ and a realization of the OU processes:

$$\left[\bar{Y}_{T_F}, \bar{Y}_{t_1}, \ldots, Y_{t_{31}}\right]$$

      we compute $F_m(T_e)$ by integrating $F(T_e, T)$ over $T \in [T_1, T_2]$ where T_1 and T_2 are the start and end of the delivery month.
    - Recalling that $F_m(T_e)$ is a bid-week price we have taken T_e to be four business days before the start of the delivery month.

Extrapolation Taken Too Far?

A Variation of Peaking Options

- Working Problem:
  - Simulation
    - The daily spot prices p_d for each delivery day d is given by $F(t_d^*, t_d)$ where t_d^* is one-business day before each delivery time t_d since the fixing for the GD index occurs for trading on this day.
  - The results are shown below using 100,000 realizations (MC error $\sim$ .3% of estimated values.
  - The top figure shows the value in \$ / MMBtu as a function of swing number.
    - Note that as we have normalized by J the value decreases with J as lower realized outcomes are included in the average. structures. structures.
  - The second and third plots are related to just how effectively one can expect to hedge these structures by constructing min-variance hedges using the standard forward starter $J = 31$.

Extrapolation Taken Too Far?

A Variation of Peaking Options

- Working Problem:

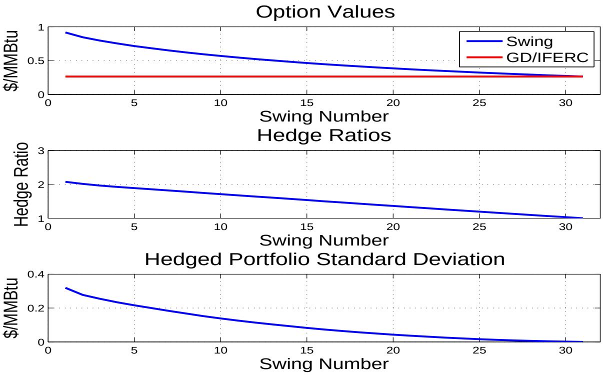

Extrapolation Taken Too Far?

A Variation of Peaking Options

- The moral of the story: It is very difficult to hedge the low J structures.
  - The standard forward-starter is ineffective.
  - Δ-hedging within the month is a joke:
    - You have next day prices and (perhaps) balance of of month swaps
    - There is no balance-of-month options markets
  - All that you have really done is used the risk-neutral measure as a reference measure for hedge construction in an incomplete market setting.

Key Points

- If possible model the forward curve explicitly.
  - Spot price models can be compelling intuitively but are burdened with the need to invert inhomegenous drifts for consistency with forward curves.
- It is in general challenging to calibrate all parameters from traded instruments.
  - Long time-scale markets are more amenable to likelihood based estimation.
  - Mixed time-scale markets often rely on more fundamental analysis.
- Multi-factor models provide (arguably) plausible methods for extrapolation of correlations and forward vols.
- Beware of incompleteness in extrapolation
  - Ask the question: How would I actually hedge this payoff given the instruments available?

On Deck

- Lessons from Tolling Deals
- Variable Quantity Swaps and Econometric Models