THE LONG AND THE SHORT OF IT – MEASURING AND FORECASTING VOLATILITY

Andrew Clark, Chief Index Strategist, Thomson Reuters Indices
September 23, 2009

THOMSON REUTERS

Topics

- In this webinar, we will look at volatility across a variety of time horizons
- This will allow us to build a volatility curve
- And forecast the points along that curve
- The basis of our work will be the Thomson Reuters (TR) U.S. Equity index
- And then be repeated using the Cyclical Consumer Good/Services sector in the same index
- Our reason for doing this will be to show the superiority of using TR indices and insight to build better volatility swaps and related derivatives

What is Volatility

- Many definitions of volatility exist such as:
  - Standard deviation
  - The difference in prices squared
  - The return squared
  - The absolute value of the difference in prices
  - The absolute value of the returns
- We will use the absolute value of returns as it is one of the most common academic definitions of volatility and gives the most stable results given varying sample sizes

What is a Volatility Swap

- Definition:
  - A volatility swap is a forward contract on the future realized volatility of the given underlying asset. Volatility swaps allow investors to trade the volatility of an asset directly, much as they would trade a price index
  - These instruments can be used to speculate on future volatility levels or to hedge the volatility exposure of other positions or businesses
- The difference between a volatility swap and a variance swap such as the VIX is:
  - The profit and loss from a variance swap depends directly on the difference between realized and implied volatility
  - Whereas a volatility swap depends only on the realized volatility

Current Volatility Models and the VIX

Standard Models

- Heston single-factor model
- Fonseca *et al* extension of same into a multi-factor framework
- The work of physicists to explain the "smile," skew and other stylized volatility facts. See works by J.L. McCauley and also J.W. Dash
- Multifractals
  - Calvert and Fisher demonstrate the efficacy of a multifractal model (one where the fractal value or Hausdorff dimension varies over time) to explain volatility
    - Multifractals capture the stylized facts of volatility – long memory, intermittency and the like – very well. Better than the prior 3 models

VIX: Interpretation

- The VIX is quoted in terms of percentage points and translates, roughly, to the expected movement in the S&P 500 index over the next 30-day period, on an annualized basis
- For example, if the VIX is at 15, this represents an expected annualized change of 15% over the next 30 days
- From this one can infer that the index option markets expect the S&P 500 to move up or down 4.33% over the next 30-day period ($15\%/\sqrt{12}$)
- That is, the S&P 500 index options are priced with the assumption of a 68% likelihood (one standard deviation) that the magnitude of the S&P 500's 30-day return will be less than 4.33% (up or down).

VIX: Criticism

- Often, when commentators discuss the option markets, the VIX is used to represent the overall sentiment for equity options
- However the relationship of the VIX to individual equity options can be easily overstated
- Different dynamics drive the volatility of index options compared to that of equity options and the two can often be uncorrelated
- In particular, the VIX is limited to a 30-day period, while for most non-index equity options, the most liquidity is usually found in the 2 to 6-month maturities
- In addition, volatility is often a function of market sector. For instance, volatility is usually assumed to be high in technology stocks and low in utility stocks
- Using a single number such as the VIX to represent the volatility for all equity options is usually overly simplistic.
- Finally, the VIX is a coincident indicator of volatility, not a forward one

Volatility Forecasting Models

GARCH Models

- Why GARCH Models
  - Because current and future volatility is recognized by most analysts as a linear or non-linear combination of past volatility (autoregressive heteroscedasticity)
    - And this is what GARCH is good at working with
  - IGARCH
    - Popularized by RiskMetrics
    - IGARCH is a actually a particular semigroup
      - An exponential moving average (EMA)
    - It's ease of implementation is its strongest appeal
    - It's main drawbacks are
      - The inherent bad trade off between the look-back period for composing the EMA and the desire for the EMA to respond quickly to market changes
      - Though a semigroup structure is used, it is used to compute the weights that are applied to actual returns, not to generate smoothed returns

GARCH Models

- GARCH Models (cont'd)
  - EMA GARCH w/better sampling
    - Compute volatility on a tick-by-tick basis
    - Simple returns replaced by smoothed returns (true semigroup process)
    - Good ability to generate forecasts intraday and also out a few days
- Other time series techniques are used as well to forecast volatility
- But for most analysts, GARCH predominates

Thomson Reuters Volatility Models

Decomposing Historical Volatility

- As volatility exists across a variety of trading horizons
- A volatility curve, like the yield curve, can be drawn
- Wavelets can decompose historical volatility into multiple time periods, ie, days, weeks and months
- All the methods mentioned so far, except for the EMA GARCH w/better sampling, deal with just one time period, maybe two
- So our wavelet decomposition method does things most others don’t

Volatility Decomposed: TRX US

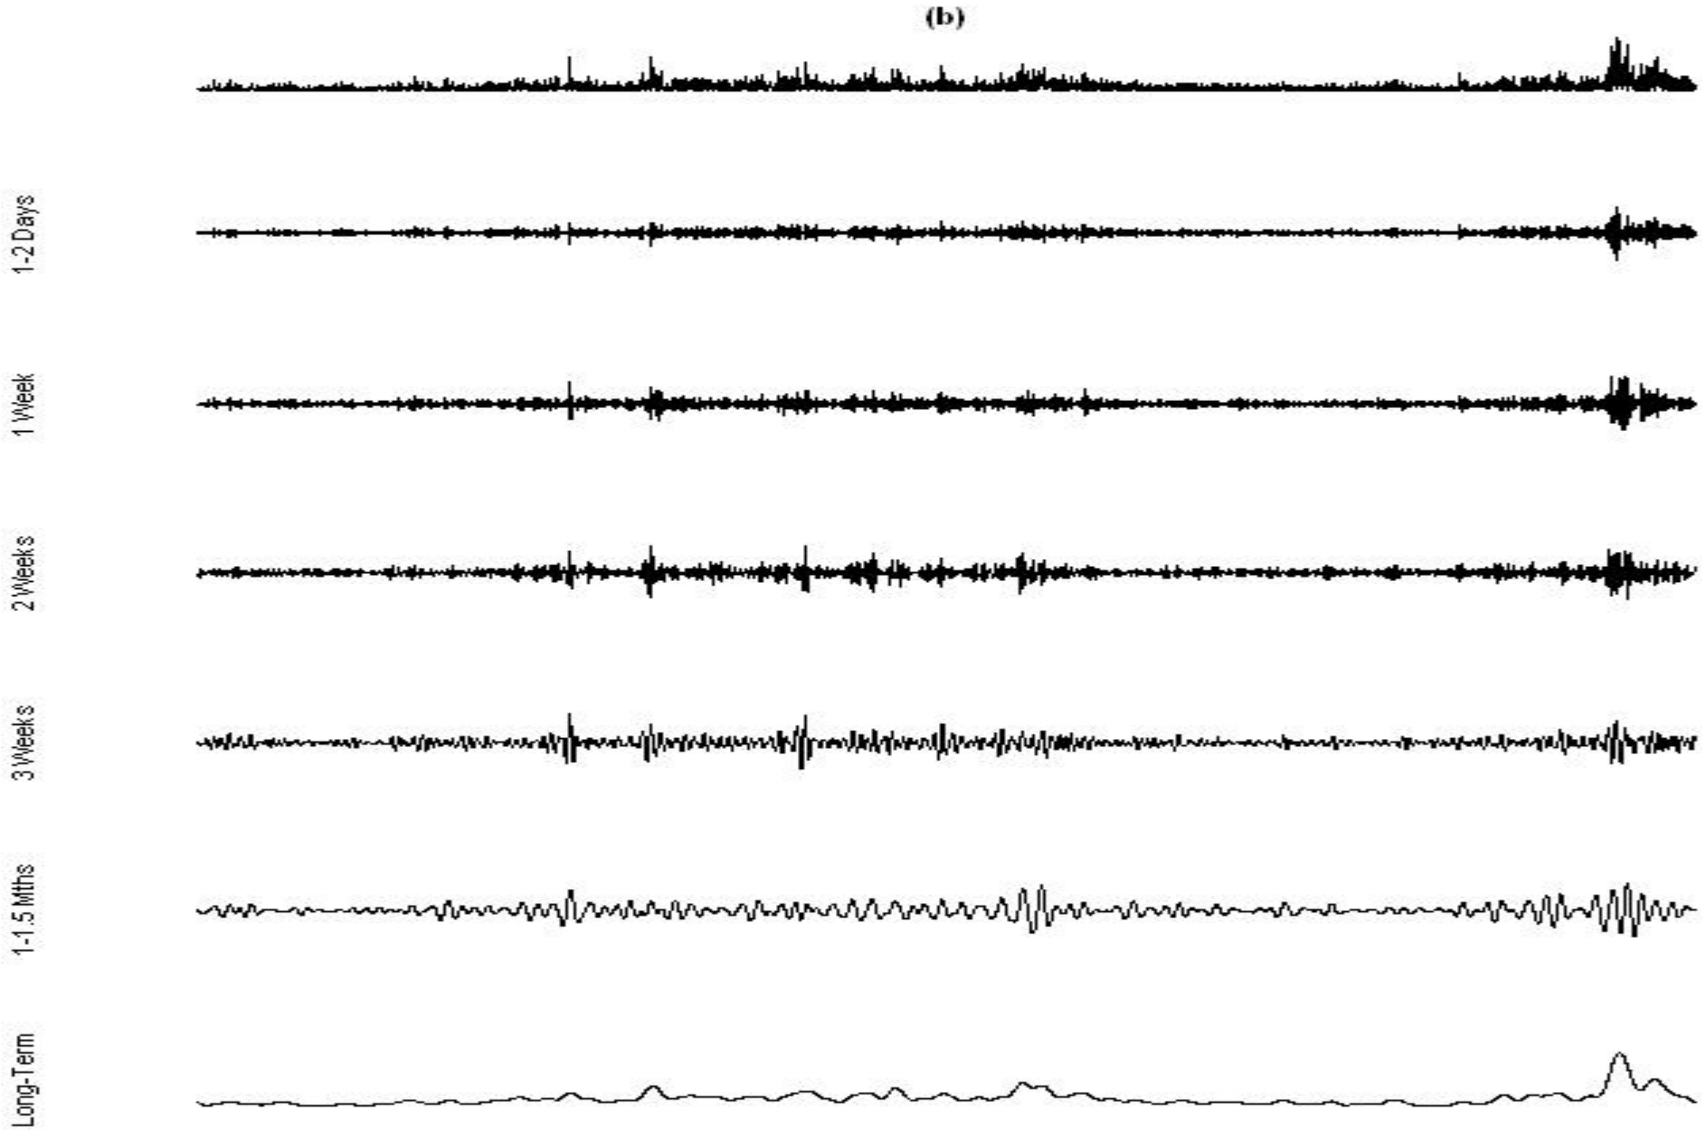

Volatility Curve: August 21, '09

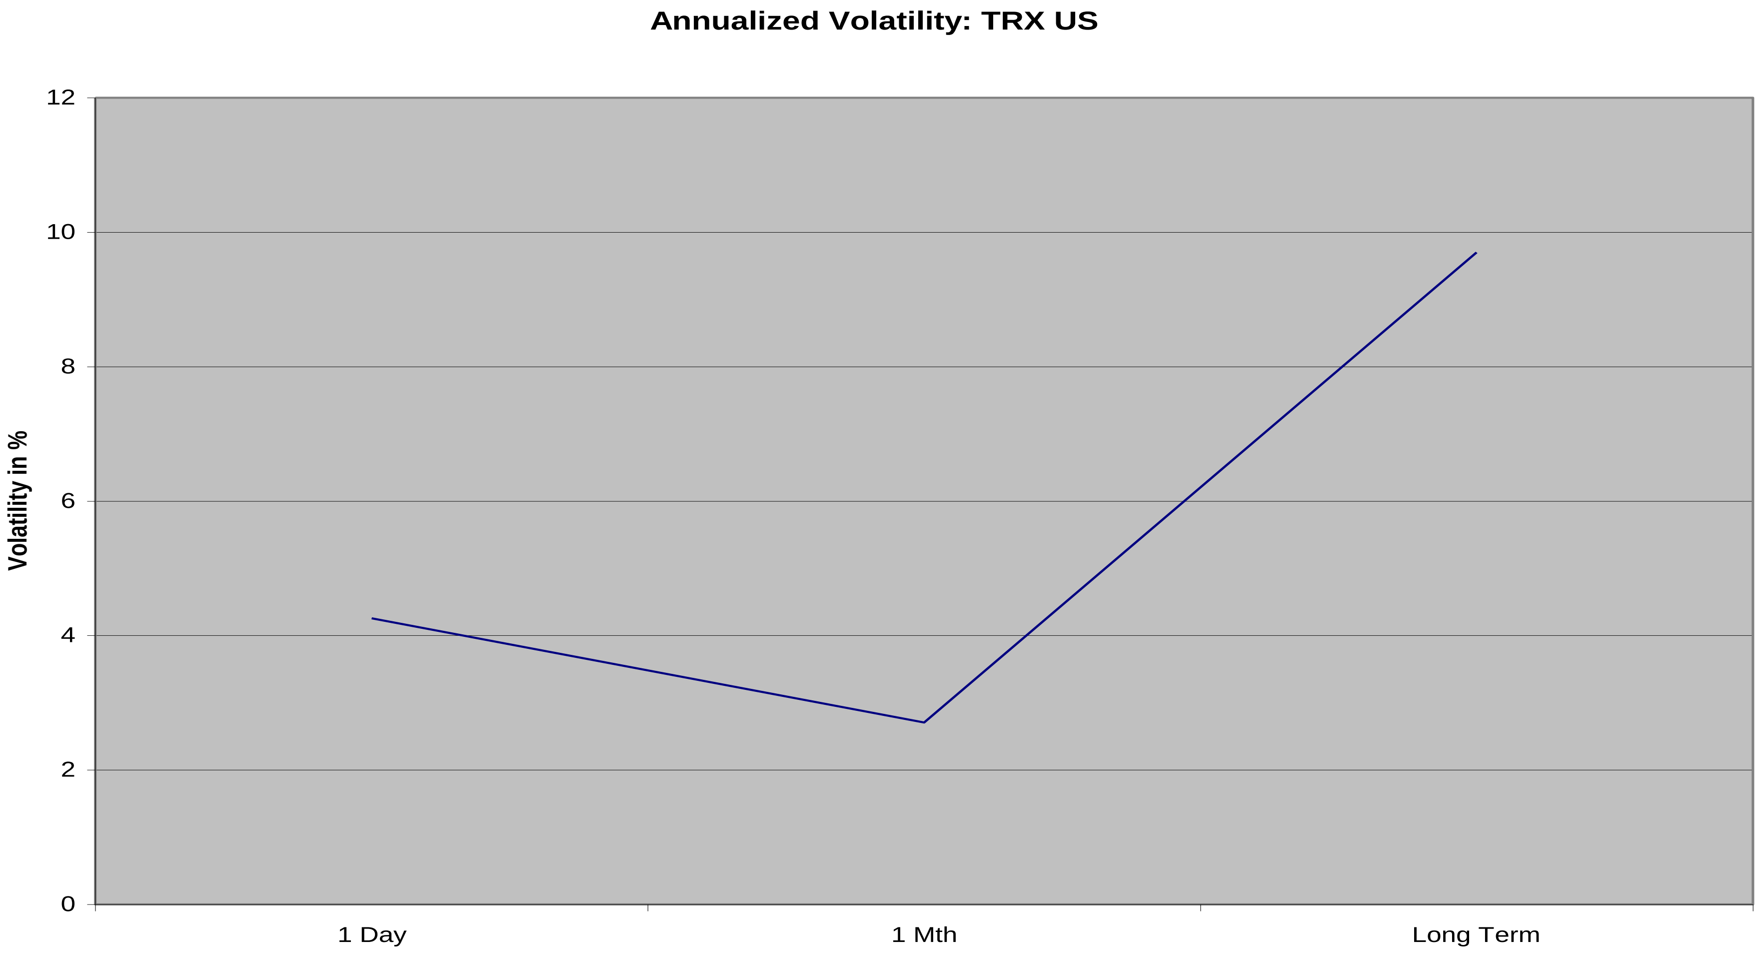

Forecasting vs Modeling

- Why the above are different
  - For modeling
    - The input data is the true signal plus (some or a lot of) noise
    - The analyst then identifies a *model equation(s)* that can duplicate the input data, ie, fit the signal *and the* noise
  - For forecasting, the same input is used
    - But here noise needs to be accounted for
    - Because predictions are expectations of the *signal*. Not the *noise*
    - So *denoising* the input data, if possible, is needed for good forecasts

Forecasting Volatility

- Using Dynamical Systems Theory
  - Denoise the data if possible
  - Determine topology of data by calc'ing embedding dimension and time lag
    - Specifically reconstruct the phase space or vector space of the data by using dimensional analysis
  - Calculate largest Lyapunov exponent
    - If negative, this implies a fixed stable point, so a wide variety of linear ad non-linear models are available
    - If zero (or near zero), a periodic or near-periodic function indictaed
    - If > 0, chaotic methods are available (both linear and non-linear)
  - Fit appropriate forecast model and back test

Forecast of Long Term Volatility: Feb 4 '08 – Mar 23 '08

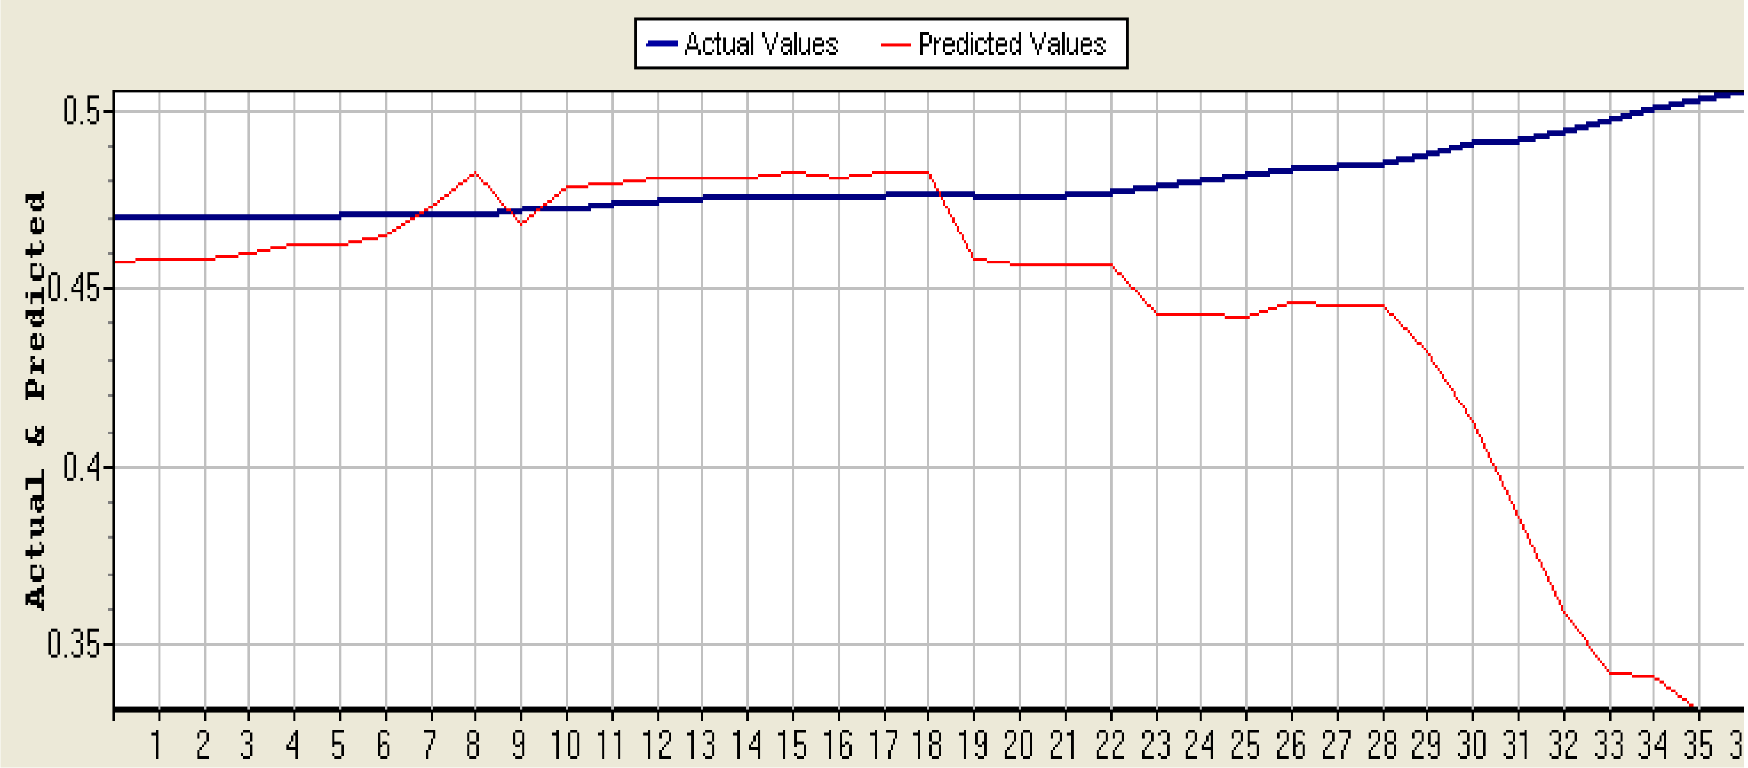

THOMSON REUTERS

Forecast of Long Term Volatility: Sept 18 ’08 – Nov 5 ‘08

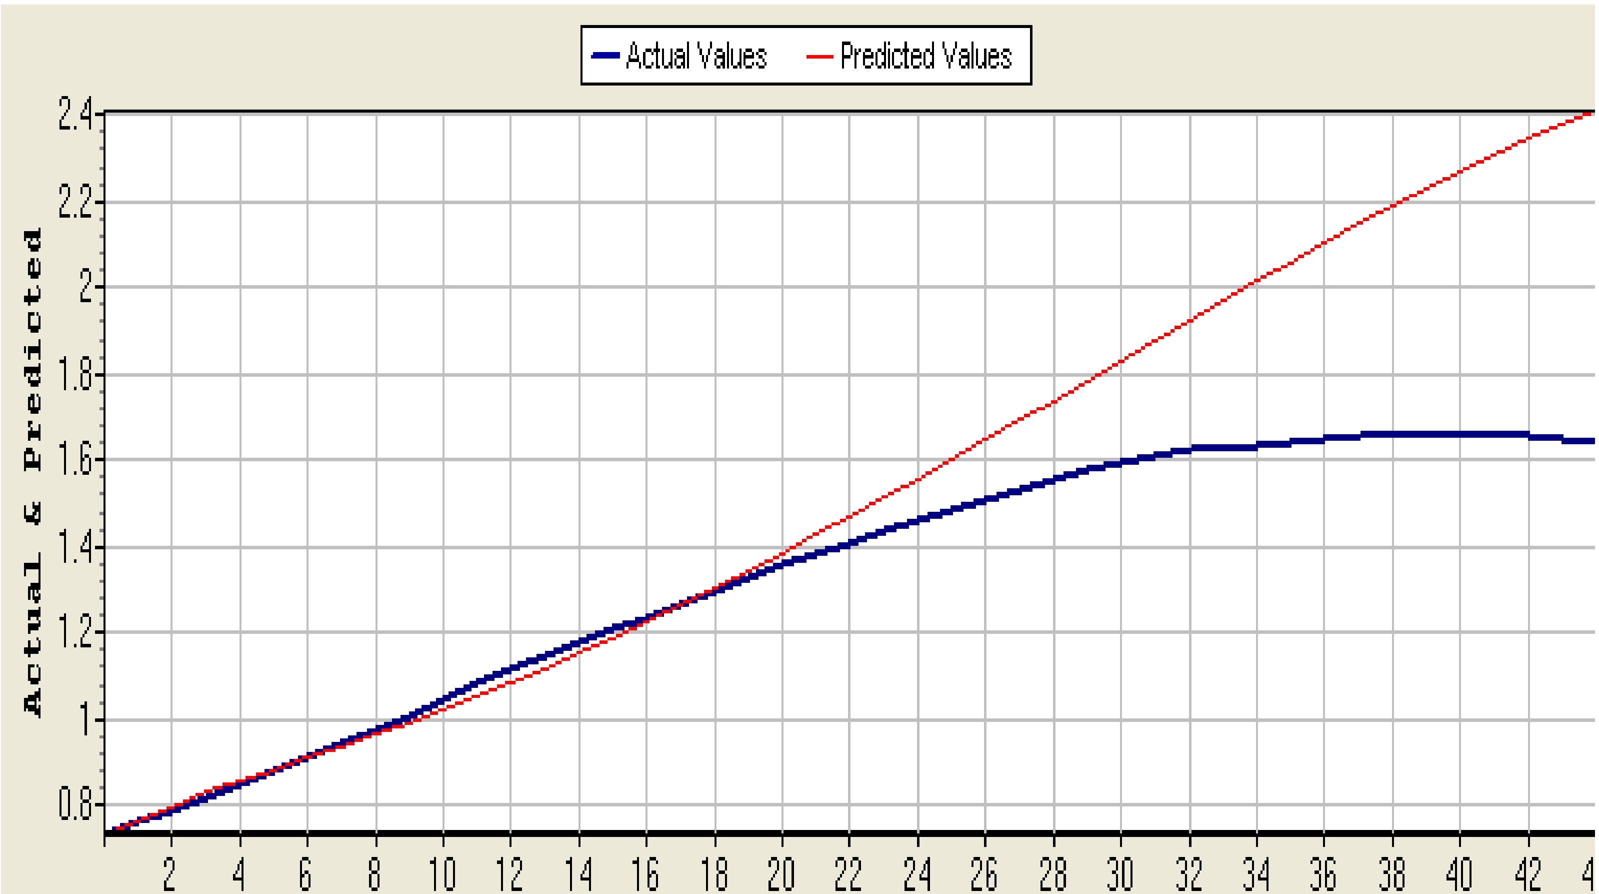

Long Term Volatility Forecast Model

- The features of the model are
  - Time Lag of 44
  - Embedding Dimension of 4
  - 44 x 4 = 176: # of days of the average cycle
    - For example 176 days from high to low to next high
    - So Long Term volatility is volatility over an ~ 8 mth period
    - Formally 176 is the time span of the embedding vector
  - Largest Lyapunov ~ 0.15
    - Process is chaotic so it has some level of determinism and limited forecast ability
    - In this case about 18 – 20 days
  - Best fit using RMSE - Locally Linear forecast model
    - Euclidean distance metric
  - Neighborhood size: 6
    - IE, 6 nearest neighbors found. The # of vectors needed to produce good forecasts

1 Week Volatility Forecast: Aug 18 ’08 – Sept 5 ‘08

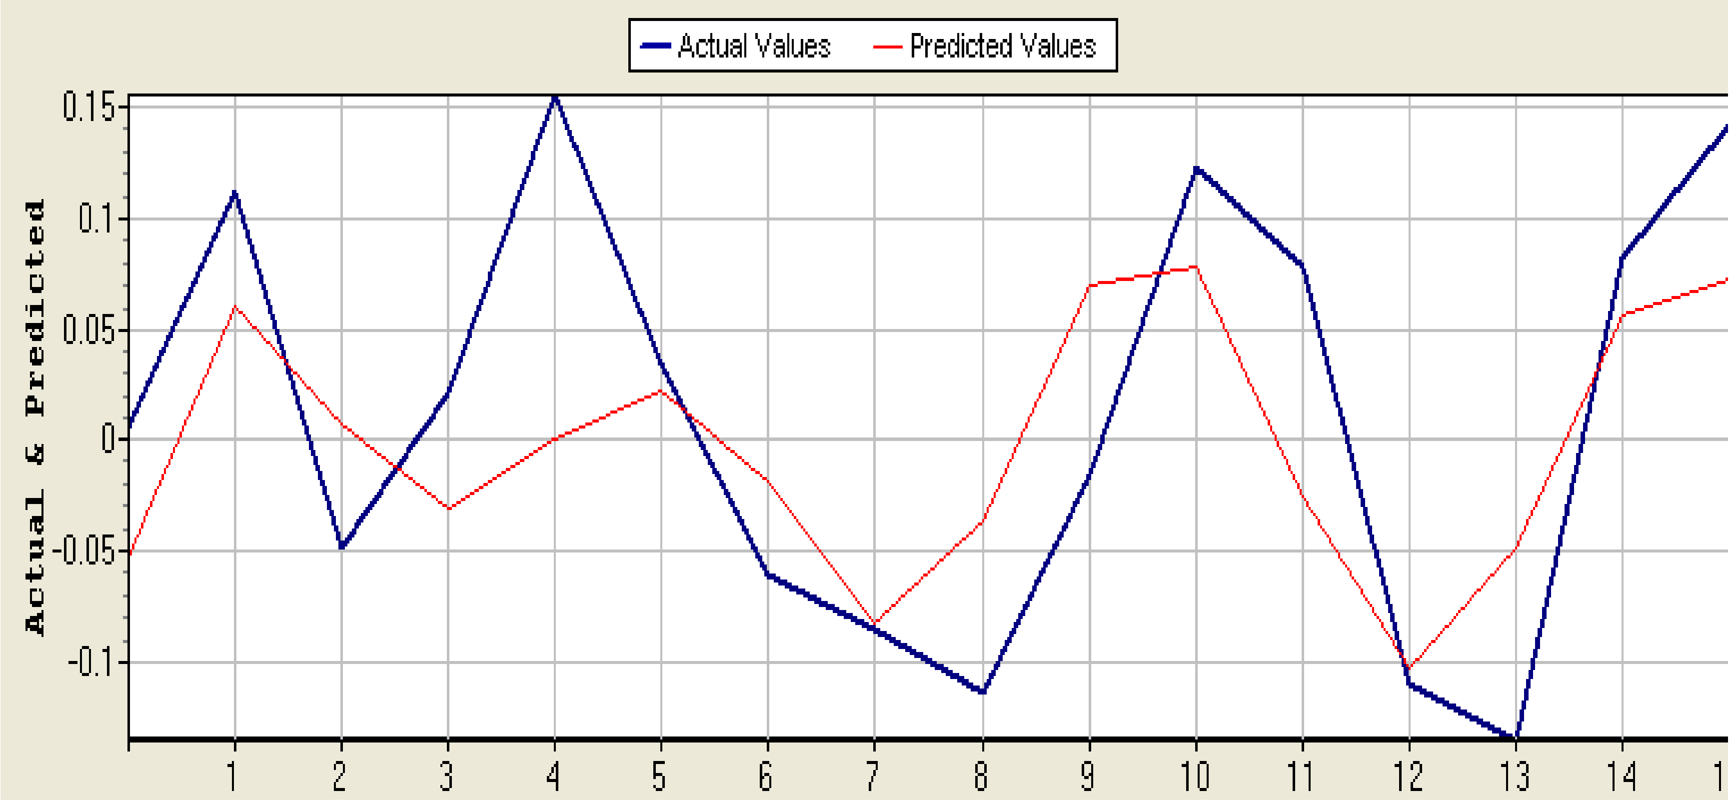

1 Week Volatility Forecast Model

- The features of the model are
  - Time Lag of 1
  - Embedding Dimension of 5
  - 1 x 5 = 5: # of days of the average cycle
  - Largest Lyapunov ~ 0.46
    - Process is chaotic so it has some level of determinism and limited forecast ability
    - In this case about 12 – 15 days
  - Best fit: Locally Constant forecast model
    - Cosine distance metric
  - Neighborhood size: 16
    - IE, 16 nearest neighbors. The # of vectors needed to produce good forecasts

Long Term Volatility: Cyclical Consumer Goods/Services

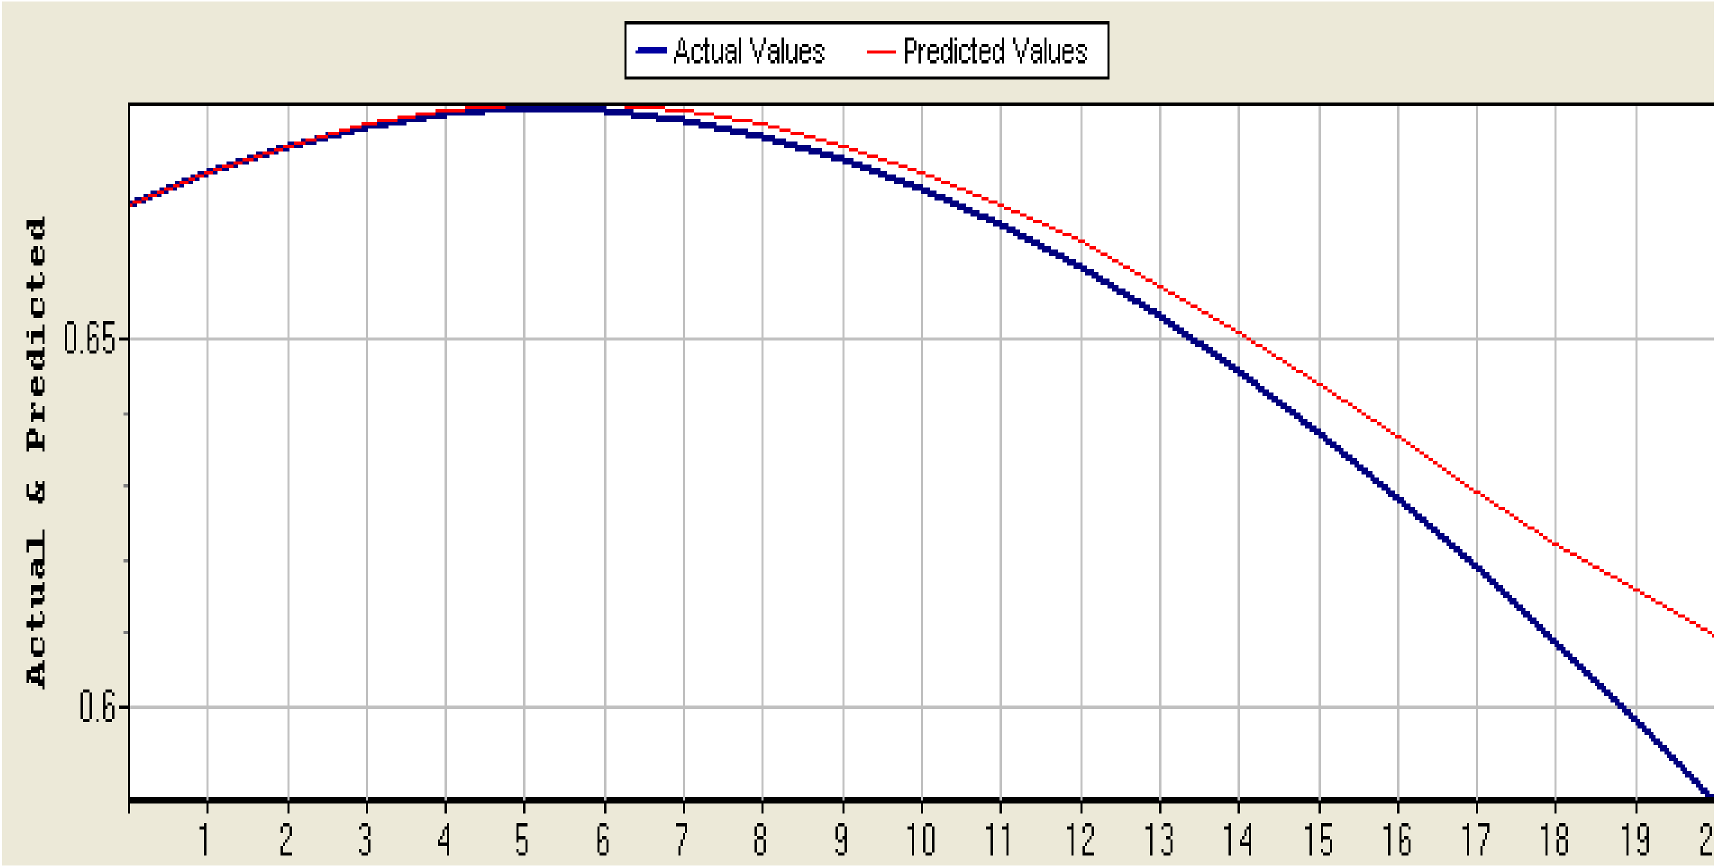

Long Term Volatility: Cyclical Consumer Goods/Services

- The features of the model are
  - Time Lag of 28
  - Embedding Dimension of 4
  - 28 x 4 = 112: # of days of the average cycle
  - Largest Lyapunov ~ 0.02
    - Process is probably more periodic than chaotic
    - Forecast horizon: 15 – 20 days
  - Best fit: Local Weighted Linear
    - Kernel: Tricube
    - Euclidean distance metric
  - Neighborhood size: 5
    - IE, 5 nearest neighbors. The # of vectors needed to produce good forecasts

CONCLUSIONS

Conclusions

- The US TRX Equity index volatility can be easily decomposed into multiple horizons: day(s), weeks and months
- As can the sector indices that make up the larger index
- For the larger index and for the sectors, a good long term volatility indicator is produced
- And that indicator reflects the 6 mth – 8 mth volatility horizon
- The availability of a long term measure of future realized volatility is something not in the market
- Forecast accuracy of the long term volatility and other volatility time horizons is good, often extending out 1 month
- This decomposition and the forecast accuracy across a variety of horizons means better volatility swaps and options based on these swaps is available

Conclusions

- An underlying is needed to construct the swap. In this case it would be the TRX US Equity Index or its sector (TRBC) component
- As the future realized volatility is more than likely well known 1 month away (20 business days), the use of these swaps as a hedge vehicle would clearly outstrip the usefulness of existing swaps out there
- Finally, the wavelet and dynamical systems method can be used on individual stocks to generate their volatility horizons and possibly models that could forecast as well

QUESTIONS

References

- S. L. Heston, “A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options,” *The Review of Financial Studies*, Vol 6, Issue 2
- J. Fonseca, M. Grasselli and C. Tebaldi, “A Multifactor Volatility Heston Model,” *Quantitative Finance*, Vol 8, No 6
- J. L. McCauley, *Dynamics of Markets*, Cambridge University Press, 2004
- L.E. Calvert & A.J. Fisher, *Multifractal Volatility*, Academic Press, 2008
- J. W. Dash, *Quantitative Finance and Risk Management*, World Scientific, 2005

APPENDIX

Important Steps in the Last 10 - 15 Years

- Heston single factor model
  - Among single factor stochastic volatility models, the most popular and easy to implement is certainly the Heston (1993) model
    - In which the volatility follows an Ornstein-Uhlenbeck process – a common economic assumption for both stock prices and stock options
    - The associated pricing and hedging problem can be efficiently solved by performing a fast Fourier transform on the characteristic function – the defining function of any probability distribution
- Fonseca et al
  - Extends the Heston model to a multifactor specification for the volatility process where both long and short term volatility are realized, along with the implied and skew volatilities

Important Steps in the Last 10 - 15 Years

- McCauley
  - McCauley presaged some of Fonseca *et al's* work with his derivation of option pricing via exponentials vs Gaussians
  - The connection between the two lies buried in the math of each and has strong ties back to Heston. All 3 methods have to do with inverting an important function that helps solve for volatility

Important Steps in the Last 10 - 15Years

- Multifractal Volatility
  - Calvert and Fisher demonstrate the efficacy of a multifractal model (one where the fractal value or Hausdorff dimension varies over time)
    - Multifractals capture the stylized facts of volatility – long memory, intermittency and the like – very well. Better than the prior 3 models
  - Calvert and Fisher assume that shocks to volatility will be scale invariant – daily shocks affect longer term volatility in the same way as it effects short term volatility
  - This webinar will show that this is not the case
  - Which means there is a strong need to modify a basic underlying assumption common to both the multifractal interpretation and most standard econ and finance models

The Impact of News on Prices and Volatility

- Standard theory says agents integrate new price information *instantaneously*
- So no attention is paid to the time needed for a piece of information to be available to the agents in a market
- And the diversity of interpretations different agents have to the same piece of news
- This webinar and other work show that markets need a finite time to adjust
- And that movements in volatility are not the same at all time scales