Time-Varying Beta: The Heterogeneous Autoregressive Beta Model

Kunal Jain*

Professor George Tauchen, Faculty Advisor
Professor Tim Bollerslev, Faculty Advisor

Honors Thesis submitted in partial fulfillment of the requirements for Graduation with Distinction in Economics in Trinity College of Duke University.

**The author will be working for Moelis & Company in New York, New York and can be contacted at kunaljain44@gmail.com.*

Duke University
Durham, North Carolina
2011

Acknowledgements

I would like to thank Professor George Tauchen and Tim Bollerslev for their guidance, advice, help, and encouragement. Additionally, I would like to thank my peers in the 202FS Honors Workshop- Dylan Mingwei Lei and Angela Ryu- for their constructive feedback throughout the semester.

Abstract

Conventional models of volatility estimation do not capture the persistence in high-frequency market data and are not able to limit the impact of market microstructure noise present at very finely sampled intervals. In an attempt to incorporate these two elements, we use the beta-metric as a proxy for equity-specific volatility and use finely sampled time-varying conditional forecasts estimated using the Heterogeneous Autoregressive framework to form a predictive beta model. The findings suggest that this predictive beta is better able to capture persistence in financial data and limit the effect of microstructure noise in high-frequency data when compared to the existing benchmarks.

JEL classifications: C01, C13, C22, C29, C58,

Keywords: Beta, Heterogeneous Autoregressive, Persistence, Financial Markets

1. Introduction

An investor expects to be compensated in terms of the time value of money and the risk associated with a particular investment. The time value of money component is the concept that money available today is worth more than money available in the future due to multiple factors, including the potential earning capacity of the money via alternate investments, the possible inflation that could discount the present value, and the unrealized opportunity cost of receiving interest. Typically, investors associate the time value of money with the risk-free interest they could accrue by allocating the money today. For example, when governments issue bonds, they offer a specified return to the investor. Under the assumption that governments are free from default, this specified return signifies a guaranteed rate of return that is referred to as the risk-free rate. However, due to the minimal risk associated with this particular allocation of money, the associated return is relatively low. Alternatively, investors expect to be compensated for the incremental risk they take by investing in non risk-free, or risky assets. The rationale for investing in risky assets is the potential for excess returns they provide relative to the risk-free rate of return, referred to as the risk to reward ratio. The quantification of this excess risk can be done through a statistical measure called beta.

Beta is a commonly defined statistical measure that represents the volatility of a specific asset's returns relative to that of the market returns. It is defined as the covariance of an asset's returns with the market's returns, divided by the variation of the market returns. Specifically, beta measures the portion of an asset's statistical variance that cannot be mitigated by the diversification of a portfolio composed of many risky assets, or the market portfolio. By definition, the market portfolio has a beta of one because the covariance of the market with itself and its variation are the same. Alternatively, an asset with a beta of zero means that the returns of that asset change independently of changes in the market's returns over time. But, if an asset moves more than the market, then the asset's beta is higher than one and is said to be riskier with the potential of higher returns.

In securities analysis and company valuation, the accuracy of beta is crucial in determining investment strategies and the pricing of individual equities. An essential feature of the beta-metric is the explanative power it lends to assess portfolio risk and returns. Given a lack of explanative power in the beta-metric, portfolio managers cannot forecast returns and minimize risk to reward ratios (Klemkosky and Martin, 1975). For example, the Capital Asset Pricing Model (CAPM) uses beta as a key component in calculating the cost of equity for a specific equity. The cost of equity represents the discount rate used to calculate the present value of a company's future cash flows and eventually the company's share valuation. More specifically, the cost of equity measures the rate at which investors need to be compensated to invest in that equity. It does this by calculating the expected excess return on the market by multiplying beta with the market risk premium, or the difference between the expected return on the market and the risk-free rate. Holding all else constant, the higher an equity's beta, the higher the cost of equity discount rate and the lower the present value of the company's future cash flows. From this example it can be seen that slight variations in beta can impact the discount rate, which can lead to diverging company share valuation results.

Traditionally, the beta used within the CAPM is assumed to be constant over time and is typically estimated using moving estimation windows, typically of five to ten years. By constant, it is meant that betas are calculated on a set period-by-period basis, as oppose to a continuous evolution. But, numerous pieces of literature, including those from Bos and Newbold (1984), Collins, Ledolter and Rayburn (1987), Brooks, Faff and Lee (1992), and Choudhry (2002, 2004) have suggested that a constant-beta CAPM is unable to satisfactorily explain the cross-section of average returns on equities and the market to capture dynamics in volatility. Specifically, Adrian and Franzoni (2005) argued that models without time-evolving betas fail to capture investor characteristics and may lead to inaccurate estimates of the true underlying beta. Following this criticism, multiple time-varying beta models (Campbell and Vuolteenaho (2004), Fama and French (2005), Petkova and Zhang (2005), Lewellen & Nagel (2006), Ang and Chen (2007)) have been proposed to estimate a more precise method to measure volatility and calculate the true underlying beta of an asset. Many of these models utilize coarsely observed daily returns (typically squared returns) to extract information about the current level of volatility in returns and then utilize this to form expectations about the next period's return volatility. In doing this, these models treat beta as an endogenous variable that varies slowly and continuously over time. But again, this method does not fully capture the dynamics of volatility in situations where volatility changes rapidly from one level to the next.

One particular method utilized by Andersen, Bollerslev, Diebold and Wu (2006), highlights the usage of high-frequency data and realized measures to estimate the underlying beta. High-frequency data gives way to realized measures of volatility, including realized variance and realized covariance, which are computed on a finely sampled basis, such as 1 or 5 minutes. Unlike observed daily squared returns, realized measures are more informative about the current level of volatility because they instantaneously integrate new information. Anderson et al. explore the dynamic nature of beta by introducing the concept of a “realized betas”, computed from high-frequency intraday returns. The realized beta is an estimator of the underlying beta that is calculated by taking the ratio of the realized covariance of an equity’s returns with the market returns and the realized variance of the market returns. In line with realized measures, the usage of a realized beta allows for the instantaneous information adjustment to capture dynamics in volatility and an increased number of data points for the purposes of estimation.

In addition to capturing dynamics in volatility via high frequency data, it is imperative to take into account the long memory and persistence evident in financial data that are ideal for robust predictive measures. One way to incorporate long memory and persistence into models is through fractional integration or cointegration such as the autoregressive fractionally integrated moving average (ARFIMA), or the fractionally integrated generalized autoregressive conditional heteroskedasticity model (FIGARCH) developed by Baillie, Bollerslev, and Mikkelsen (1996). But, these models are not easy to estimate and they typically require nonlinear maximum likelihood estimation procedures. Literature by Andersen, Bollerslev, Diebold, and Labys (2003) and Andersen, Bollerslev, and Huang (2007) has empirically shown that linear models can oftentimes predict future volatility more accurately than nonlinear models such as the ARFIMA and FIGARCH models. The Heterogeneous Autoregressive (HAR) model developed by Corsi (2003) provides linear estimation via ordinary least squares and captures the effects of persistence in a parsimonious manner.

This paper provides the development and evaluation of an alternative time-varying conditional forecasted beta model that expands on previous literature to utilize high-frequency data to estimate beta. The time-varying model that is developed within this paper uses the Heterogeneous Autoregressive (HAR) framework that uses a linear combination of moving averages computed over varying time horizons, particularly days, weeks and months, to calculate realized averages. Specifically, the Heterogeneous Autoregressive Beta model (HAR-Beta) developed within this paper uses a linear combination of daily, weekly, and monthly-realized betas computed from high-frequency intraday returns to calculate a predictive beta. Then, the predictive beta is conditioned on the observed market return to give a conditional expected return for a particular asset. If beta were constant, analyzing the conditional expected return would not provide additional information other than that contained in the conditional mean. However, if beta is in fact time varying, then the conditional probability distribution of returns depends on the conditional probability distributions of the predicted betas. Hence, the use of conditional probability distributions revealed through the HAR-Beta model could better approximate the underlying beta and explain the excess returns on the market.

First, in Section 2, this paper develops the elements of a time-varying beta model by explaining the concept of a beta and a realized beta. Then, Section 3 gives a brief description of the high frequency data used, including the estimation and out of sample datasets. To normalize the results, Section 3 also provides insight into market microstructure noise reduction and benchmark comparisons of a constant returns model and a constant beta model. Next, in Section 4, the theoretical basis behind the HAR-Beta model and the statistical methods for determining accuracy are developed. Within this section, key elements regarding the motivation for a time-varying beta, such as persistence and memory via statistical tools like standard deviation and first order autocorrelation are detailed. The HAR-Beta model is trained on the estimation dataset, and then the forecasts are produced over the out-of-sample period. The model is then evaluated based on the performance in the out-of-sample datasets, measured via the root mean squared error described in Section 4.3. Section 5 provides the estimation sample and out of sample results that are supplemented by final conclusion in Section 6.

2. The Elements of a Time-Varying Beta

The beta coefficient (β) is a key parameter in the one-factor Capital Asset Pricing Model (CAPM) and can be represented through the Security Characteristic Line (SCL)

$$r_{a,t} - r_f = \alpha_a + \beta \times (r_{m,t} - r_f) + \varepsilon_{a,t} \tag{1}$$

where $r_{a,t}$ is the rate of return on *asset a* at time *t*, $r_{m,t}$ represents the rate of return on the market at time *t* and r_f is the risk-free rate. For the ease of exposition, it will be assumed that markets are efficient and the expected value of the returns in excess of the compensation for the risk is zero for all portfolios, ($E(\alpha_{ia}) = 0$). Additionally, this analysis will use high frequency data taken over finely sampled discrete intervals of time, like 1 or 5 minutes. Given the usage of these finely sampled discrete intervals of time, it is also assumed that the effective risk-free rate, (r_f), does not change significantly and hence will be assumed to be zero. The resulting equation of the SCL is

$$r_{a,t} = (\beta \times r_{m,t}) + \varepsilon_{a,t}\,. \tag{2}$$

Now, the SCL represents the relationship between the return of a given asset *a* at time *t* with the return of the market ($r_{m,t}$) and a sensitivity measure of beta (β). Beta is a sensitivity measure that describes the relationship of an asset's return in reference to the return of a financial market or index. Beta is derived from linear regression analysis in which the returns of an individual asset (r_a) are regressed against the returns of the market (r_m) in a specific time interval to find the covariance of the asset's and the market's returns. Then, the covariance is scaled by the variance of the returns on the market (r_m) to measure the sensitivity of the asset's returns to the market's returns. Beta is defined as

$$\beta = \frac{Cov(r_{a,t}, r_{m,t})}{Var(r_{m,t})} \tag{3}$$

where $r_{a,t}$ is a measure of the rate of return on *asset a* at time *t* and $r_{m,t}$ is a measure of the rate of return on the market or index at time *t*. Specifically, beta measures the statistical variance or systematic risk of an asset that cannot be mitigated through diversification. By definition, the market portfolio has a beta of one because the covariance of the market with itself and its variance are the same. Alternatively, an asset with a beta of zero means that the returns of that asset change independently of changes in the market's returns over time.

To measure changes in stock prices, namely variance and covariance, this analysis will utilize the logarithmic of price and geometric returns as the scale of measurement. The rationale for logarithmic price and geometric returns is the comparability between two data points and the standardization of a scale that a logarithmic provides. Hence, the intraday geometric return is defined as

$$r_{i,j} = p(t-1+\frac{j}{k}) - p(t-1+\frac{j-1}{k}) \tag{4}$$

where p is the logarithmic of the stock price, t is the specific day, k is the frequency or sampling interval at which the logarithmic prices are sampled at, and $j = 1, 2, \ldots$, k. For intervals of time, the beta (β) of an asset is defined as the underlying ratio between the integrated stock and market return covariance and the integrated market variance. If the instantaneous volatility σ(t) were known, then the true variance and true covariance, called the *integrated variance* and *integrated covariance*, could be found by integrating the spot volatility over the time interval:

$$\text{Integrated Variance} = \int_{t-1}^{t} \sigma^2(u)du \tag{5}$$

$$\text{Integrated Covariance} = \int_{t-1}^{t} \sigma_{a.m}(u)du. \tag{6}$$

But, the underlying spot volatility is impossible to observe. Instead, given that it is possible to observe realized prices at discrete measures of time utilizing high frequency data, realized measures will be used to numerically approximate the *integrated variance* and *integrated covariance*.

Calculating realized measures, such as the realized variance and realized covariance is intuitive and parsimonious. Given a set of geometric returns at over a specified time interval, as defined above, the realized variance (RV) is computed as

$$RV = \sum_{j=1}^{K} r_{t,j}^2. \tag{7}$$

The realized variance represents the sum of the squared geometric returns sampled at a specified sampling frequency, *k.* Anderson and Bollerslev (1998) pointed out that as the frequency approaches infinity, $k \to \infty$, or the time between observations approximates zero, the realized variance converges to the integrated variance and can serves a measure of underlying volatility. Huang and Tauchen (2005) confirmed this and found that realized variance is a reasonable and appropriate estimator of the integrated variance, even in the presence of jumps. The realized covariance is calculated in the same manner. It is defined as

$$RCOV = \sum_{j=1}^{K} r_{a,t,j} r_{b,t,j} \tag{8}$$

where the sum of the covariance between the geometric return of *asset a* ($r_{a,t,j}$) and geometric return of *asset b* ($r_{b,t,j}$) is taken at a specific sampling frequency *k*. Overall, Anderson, Bollerslev, Diebold, and Labys (2003) reported that forecasting with realized measures outperforms a variety of more complicated stochastic volatility models in out-of-sample forecasting.

Hence, using the formula given above for beta, realized beta is calculated by

$$R\beta = \frac{RCOV_{a,mkt}}{RV_{mkt}} \tag{9}$$

where *RCOV* is the realized covariance of realized returns on an asset *a* and realized returns on the market over a sampling interval taken at a specific frequency, and RV is the realized variance of realized returns on the market over the same interval and frequency. The realized beta dismisses the null hypothesis of the period-by-period constant beta and allows for continuous evolution. More specifically, the realized beta is consistent for realizations of the underlying ratio between the integrated stock and market return covariance and the integrated stock market variance. Additionally, according to Anderson, Bollerslev, Diebold, Wu (2006), realized betas allows for a flexible econometric framework that avoids fractional integration and/or cointegration between the market variance and individual asset equity covariances with the market.

3. Data

The price data used in the analysis are based on minute-by-minute price quotes from a commercial vendor, price-data.com that includes every minute from 9:35 AM to 4:00 PM on trading days from 1997 to early 2009. The Standard & Poor's Repository Index 500 (SPY) was used as the market index for calculation. Additionally, there were a total of eight equities chosen for the analysis, including Coca Cola Company (KO), PepsiCo, Inc. (PEP), Microsoft Corporation (MSFT), JPMorgan Chase & Co. (JPM), Bank of America Corporation (BAC), Johnson & Johnson (JNJ), Wal-mart Stores Inc. (WMT), and Exxon Mobil

Corporation (XOM). These particular companies were chosen due to their liquidity, market capitalization and representation across industries including Consumer Goods, Technology, Financial, Healthcare, Services and Integrated Oil & Gas. The full time interval used for analysis is January 2, 2001 to January 3, 2009, which includes 1989 trading days, and was chosen due to the availability of data for the equities.

3.1 Market Microstructure Noise

One method used to calculate the price of an equity is the constant growth model which takes the sum of all the expected discounted future cash flows to get a net present value. The model assumes that a firm's profits per share, π, grow by a constant growth rate, g. Then, expected profits, π^e, are discounted by the difference between a cost of equity component, k, calculated through the Capital Asset Pricing Model (CAPM) and the growth rate to estimate the price of the equity

$$P = \frac{\pi^e}{(k-g)}. \tag{10}$$

By using a constant growth rate and discount rate to estimate prices, the constant growth model does not account for changes in growth-rate expectations and market frictions such as the bid-ask spread, instantaneous information asymmetry and other trading anomalies. These short-term market price movements of assets may reflect a value other than the true price of the asset.

For example, suppose a company is expected to produce profits per share, π^e, of \$3.00 according to analyst via a 5% long term growth rate, g. Under the current market condition, investors feel confidently about the company's potential to meet these expectations and value of the cost of equity, k, at 8%. Then, according to the constant growth model, investors are valuing the price of the company's equity at

$$P = \frac{\$3.00}{(0.08-0.05)} = \$100.$$

Now, suppose that during this time, a public news release states that the company plans to launch a new product that could potentially increase expected profits. Now, the discount rate, or investor's cost of equity, k, decreases from 8% to 7.5% to reflect a change in the underlying price of the equity

$$P = \frac{\$3.00}{(0.075 - 0.05)} = \$120.$$

But, the market price does not instantly adjust to the new fundamental price due to market delays and frictions following the public news release. Hence, marginal changes in the input factors of the constant growth model, such as frictions and delays in the discount rate, can lead to diverging valuations in the true underlying price of an equity.

Typically, prices are modeled as a random process that is continuous in time and space. While this assumption of continuity may work for coarsely sampled datasets, it does not hold when using high frequency data. At finer scales of time, various market frictions (release of new information, bid-ask bounce, discreteness of prices, block trades, etc.) arise that impede the instantaneous adjustment to the fundamental price. Hansen and Lunde (2006) define these market frictions as market microstructure noise

$$p(t) = p^*(t) + u(t), \tag{11}$$

where $p(t)$ is the observed log price in the market at time t, $p^*(t)$ is the fundamental log price at time t, and $u(t)$ is the microstructure noise. Market microstructure noise is a pertinent consideration when using high-frequency data because as the frequency of data increases, the signal-to-noise ratio falls dramatically. This can be seen through changes in the underlying price defined as

$$p(t+\theta) - p(t) = [p^*(t+\theta) - p^*(t)] + [u(t+\theta) - u(t)] \tag{12}$$

where θ is a real number increment , $[p(t + \theta) - p(t)]$ represents the change in log price over a time interval and $u(t)$ is the i.i.d. microstructure noise applicable to the price change over the specified time interval. Now, as $\theta \rightarrow 0$ or the time interval is decreased, the magnitude of change in the fundamental log price, $[p^*(t + \theta)$ -

$p^*(t)]$, should decrease. Intuitively, a lack of new information available during small time periods or liquidity factors that force smaller price movement at decreased time intervals could be the explanation. Although, since $u(t)$ is i.i.d., the microstructure noise term, $[u(t + \theta) - u(t)]$, is not forced to decrease. Hence, at decreasingly small time intervals, as $\theta \to 0$, the change in fundamental log price is minimal and the observed microstructure noise is relatively large. The market microstructure noise can then distort the results and provide an inaccurate measure of the equity price.

Andersen, Bollerslev, Diebold and Labys (1999) recommend a graphical tool, called a *volatility signature plot*, as an approach to minimize the effect of market microstructure noise. Volatility signature plots display how average realized variance corresponds to sampling frequency. The x-axis represents the sampling frequency, or θ, and the y-axis maps the average realized variance calculated at each specific sampling frequency. If no microstructure noise were present, then the volatility signature plot would display a flat horizontal line that does not bias the estimates. The intuition to find the optimal sampling frequency, minimal microstructure noise and highest number of data points is as follows. As noted in Equation 12, as the time interval θ decreases, the realized variance will increase due to microstructure noise $[u(t + \theta) - u(t)]$. Though, as θ is increased by arbitrary amounts, given a liquid equity, the microstructure noises should diminish and subsequently lower the realized variance. Hence, at some point, the variance should stabilize, in which case the most number of data points can be used that are relatively robust to market microstructure noise. Figure 1A and Figure 1B show the volatility signature plots for all the equities used within the analysis. From these figures, it can be determined that below 10 minutes, there is a spike in average realized volatility as $\theta \to 0$. Utilizing the aforementioned technique and visually inspecting Figures 1A and 1B, a sampling frequency of 10 minutes was chosen as optimal for subsequent analysis.

3.2 Estimation Dataset

The estimation sample time interval chosen for the HAR-Beta model prediction coefficients was January 2, 2001 to January 2, 2006. The rationale for the particular estimation sample estimation interval was a balance between the overlap in the data available and the incorporation of a wide time interval containing numerous data points. The 5-year time period used to estimate the model is also useful in creating benchmark comparisons that will be explained in Section 4.1.1.

3.3 Out of Sample Dataset- A

The first out of sample period used was a two-year period mapped from January 3, 2006 to January 3, 2008. The two-year sampling period was used as a benchmark length for a corresponding estimation sample period that would necessitate a significant amount of data points. Additionally, this specific time interval is utilized as a time interval independent of the 2008 financial crisis, which could be characterized by a lack of persistence and predictability.

3.4 Out of Sample Dataset- B

The second out of sample period used was a three-year period mapped from January 3, 2006 to January 3, 2009. The reasoning for a secondary out of sample data period is to verify the robustness of the beta predictions in an alternate period characterized with the 2008 financial crisis. Hence, the reasoning for two out of sample periods is the comparison and depiction of the robustness of the beta predictions from the time-varying model given a time period of sporadic volatility.

4. Statistical Methods and Framework

4.1 Standard Deviations of Realized Beta

An elemental consideration when calculating beta is determining whether the asset's systematic risk, determined by its correlation with the market, is constant over time. But in order to assess this, it is necessary to characterize the persistence and predictability in the underlying components. Market betas, which are ratios of covariances to variances, are expected to display statistically persistent fluctuations over time (Andersen, Bollerslev, Diebold, and Wu, 2004). Tofalis (2008) provides a critique of the standard static beta by citing the associated volatility inherent within beta. By expressing beta from ordinary least squares (OLS) regression, one can find the resulting slope is

$$\beta = \rho \times (\frac{\sigma_a}{\sigma_m}) \tag{13}$$

where σ is the standard deviations of the rates of return on asset a and the market, and ρ is the coefficient of correlation between the rates of return. Equation 13 is equivalent to the ratio of the covariance between market and investment returns to the variance of the market returns as noted by Tofalis. Hence, a decomposition portrays the combination of volatility and correlation in the calculation of beta. This leads to an inequality (since $|r|$ is not greater than one) of

$$\sigma_a \geq |\beta_a| \times \sigma_m . \tag{14}$$

According to Equation 14**,** if realized betas, β_a, were constant over time, one would not expect the inequality to change statistically with time. But, Table 1 summarizes the standard deviations of the equity betas in the aforementioned interval that statistically deviate from the null hypothesis of no deviation. Additionally, Figure 2A and Figure 2B display a time-series plot of statistically significant fluctuations in realized betas over the whole sample interval used within this paper. Despite these fluctuations, Figures 2A and 2B show a mean reversion of the realized betas over the sampling interval that alludes to the persistence in betas. The movement in beta over time, evidenced through the statistically significant standard deviations, coupled with the mean reverting persistence, motivates a time-varying beta analysis.

4.2 Autoregressive AR(1) Model

The autoregressive process of order one, AR(1) model, is a linear time-series process used in statistics to capture dynamics. The model assumes that future values can depend on current and past values using linear approximations. Specifically, the AR(1) model is effective in its parsimonious usage of linear estimation and its ability to produce forecast results in line with more-complex forecasting models. It is defined as

$$X_t = \theta_o + \theta_1 X_{t-1} + \varepsilon_t . \tag{15}$$

The AR(1) model has the same form as a simple linear regression model where X_t is dependent and X_{t-1} is the explanatory variable, yet they have different properties. The mean and variance conditional on past returns are $E(X_t|X_{t-1}) = \theta_0 + \theta_1 X_{t-1}$ and $Var(X_t|X_{t-1}) = Var(\varepsilon_t) = \sigma_\varepsilon^2$. Hence, given the past return X_{t-1}, the current return is centered on $\theta_0 + \theta_1 X_{t-1}$ with variance σ_ε^2.

To portray how this model captures dynamics, it is necessary to understand the *coefficient of first-order autocorrelation*. The coefficient of first-order autocorrelation, or the autocorrelation parameter, represents a scaled measure of linear dependence of the present value, X_t on the past value, X_{t-1} and is defined as

$$Corr(X_t) = \frac{Cov(X_t, X_{t-1})}{Var(X_t)}. \tag{16}$$

To put this into context, now consider time indexed by t and let t $\in\{-\infty,\ldots 1,2\ldots\ \infty\}$. Let $\varepsilon_t \sim N(0,1)$ and *i.i.d.* which means that no matter how t varies, ε_t retains a standard normal distribution. The random variable, X_t, can now be represented by the AR(1) model as

$$X_t = (1-\lambda)\theta \ + \lambda\ X_{t-1} + \sigma\varepsilon_t \tag{17}$$

where θ, λ, and σ are fixed scalars, or parameters of the process. If we were to now substitute this expression for X_t into the numerator of the autocorrelation parameter

$$\begin{aligned} Cov(X_t, X_{t-1}) &= Cov((1-\lambda)\theta + \lambda X_{t-1} + \sigma\varepsilon_t,\ X_{t-1}) \\ &= \lambda Cov(X_{t-1}\ ,\ X_{t-1}) \\ &= \lambda Var(X_{t-1}).^{1} \end{aligned}$$

Now, since X_t and X_{t-1} are unconditional moments of a stationary time series, *Var(X_{t-1}) = Var(X_t)*. Therefore, Equation 16 can be expressed as

[1] The formal presentation and proof for this illustration can be found in the textbook *Time Series Analysis* by James Hamilton, Princeton University Press, 1994.

$$Corr(X_t) = \lambda \tag{18}$$

Hence, we know that the dynamics in the AR(1) model, are summarized by the autocorrelation parameter, λ. If $\lambda = 0$, then there is said to by no dynamics in the time series and X_t is said to be *i.i.d.* and normal. Alternatively, if $\lambda > 0$, then there is said to be persistent dynamics in the time series which can lend itself to predictability. The first order autocorrelation for the ε_t and ε_{t+1} of equity betas used within the analysis are significantly positive and are summarized in Table 2**.** These positive autocorrelations suggest the persistence and predictability of beta and further the reasoning behind a time-varying beta analysis.

4.3 HAR-Beta Regression Model

Empirical tests evidenced in Andersen, Bollerslev, Diebold, and Labys (2003) and Andersen, Bollerslev, and Meddahi (2004) utilize the Heterogeneous Autoregressive (HAR) class of volatility models to show that simple linear forecasts can outperform the more complicated stochastic volatility models. Numerous alternative articles have suggested additional autoregressive moving averages, such as the Autoregressive Fractionally Integrated Moving Average (ARFIMA) or the Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity (FIGARCH), to supply better persistence in data. These models typically require nonlinear maximum likelihood estimation procedures other than the conventional Ordinary Least Squares (OLS). Andersen, Bollerslev and Huang (2007) have empirically tested models such as FIGARCH and ARFIMA, and deduced that simple linear models can oftentimes predict future volatility more accurately. To summarize, the HAR class of models have two main advantages over other models. Firstly, they are parsimonious in estimation. Secondly, they capture the extreme persistence of variation in a method that is intuitive and simple to calculate. Given this theoretical backing, the Heterogeneous Autoregressive (HAR) framework was adapted for the subsequent time-varying beta analysis.

This paper formulates a Heterogeneous Autoregressive Beta model (HAR-Beta) that uses a linear combination of historical betas calculated over different time horizons, to capture the persistence of financial data. Specifically, the HAR-Beta model uses latent integrated variances and integrated covariances to compute realized betas over different time horizons. Then, given these realized betas, the model estimates normalized sums of designated time intervals to calculate a simple average of the designated quantity

$$R\beta_{t,t+k} = \frac{1}{n}\sum_{i=t+1}^{n+k} R\beta_i . \tag{19}$$

In this paper, we estimate 1-day forward beta forecasts, which corresponds directly to one trading day within a calendar month. These 1-day forward beta predictions are estimated using the basic HAR model of

$$R\beta_{t,t+1} = \beta_0 + \alpha_1 R\beta_{t-1,t} + \alpha_5 R\beta_{t-5,t} + \alpha_{22} R\beta_{t-22,t} + \varepsilon_{t+1} \tag{20}$$

where the dependent variables correspond to lagged daily (t =1), weekly (t = 5) and monthly (t = 22) regressors. The specific time horizons for the lagged explanatory variables are chosen for their intuitive appeal and natural economic interpretation of trading days in a day, week and month. The convention of lagged regressors is taken directly from the formulization of the HAR-RV model in Corsi (2003) and Muller et. al (2007).

4.3.1 Calculating Beta Predictions

First, realized betas, as described in Section 2, were computed for each of the chosen equities over the whole sample time interval, January 2, 2001 to January 3, 2009 at the optimal sampling frequency of 10 minutes. These realized betas, specifically those from January 2, 2001 to January 2, 2006, were the first necessary component in the estimation required for the subsequent regression analysis. Then, the HAR-framework from Equation 19 was used to estimate normalized average daily (β_{t-1}), weekly (β_{t-5}), and monthly (β_{t-22}) realized betas over the whole sample time interval of January 2, 2001 to January 3, 2009 at the optimal sampling frequency.

Once the realized betas and the HAR realized betas were computed over the whole sample, regression analysis was used to estimate the beta coefficients on the estimation sample period. Specifically, the realized betas from the estimation sample period of January 2, 2001 to January 2, 2006 were regressed on the HAR daily, weekly and monthly-realized betas from the estimation sample period. The results of this regression analysis were daily, weekly and monthly beta coefficients from the HAR-Beta in-sample regression. The computed coefficients are summarized in Table 3.

Subsequently, the HAR-Beta framework from Equation 20 was employed to get the 1-day forward beta predictions. To do this, the computed beta coefficients from the estimation sample time interval, Table 3, were used to approximate the weight of the corresponding out of sample realized beta. Specifically, the daily, monthly and weekly beta coefficients obtained from the regression in the estimation sample were multiplied by the corresponding HAR realized beta in the out of sample period. Then, the sum of these products was taken to formulate the corresponding 1-day forward beta prediction for the out of sample period.

4.4 Evaluation Criteria

4.4.1 Benchmarks of Comparison

One of the benchmarks of comparison used within this time-series analysis will be the constant mean logarithmic return model. The model assumes that the value at time t-1 is independently drawn from a common probability distribution and values at time t will be drawn from the same distribution. The model is independent of a beta coefficient and a conditional beta distribution that tracks the movement of an asset with the market. Instead, this comparison assumes that returns follow a model such that

$$R\beta = \frac{1}{n}\sum_{i=1}^{n} R\beta_t \tag{21}$$

where logarithmic returns at time *t*, *R(t)*, are observed as a sum of all latent logarithmic returns leading up to time *t-1, R(t - 1),* divided by the number of observations *n*. Particularly, in this assumption, logarithmic returns at time *t* are thought to be estimated by the mean of all latent logarithmic realized returns noted at a corresponding sampling frequency. Hence, the natural forecast for future values is therefore the sample mean of past intraday return data. While the constant mean model may seem to be simplistic, it is in fact that building block for a number of sophisticated models including the random walk and Autoregressive Moving Average Models (ARIMA).

An alternate benchmark comparison to the time-varying beta used within this analysis is the conventional one-factor constant beta calculated for the Capital Asset Pricing Model (CAPM). Typically, beta is estimated using moving estimation windows of 5-10 years, presumably to guard against beta variation (Fama 1976, Campbell, Lo, & MacKinlay 1997). The constant beta is calculated on a set period-by-period basis, as oppose to the allowance of continuous evolution in a time-varying method. Specifically, Banz (1981) cites the usage of a beta computed from monthly returns over a 5-year time period. Hence, given the estimation dataset of the 5-year period between January 2, 2001 and January 2, 2006, a constant beta was computed for all equities on the basis of monthly returns over the corresponding time interval. Analysis corresponding to the time-varying beta was done on the computed constant betas to gauge results.

4.4.2 Root Mean Squared Error

Accuracy measures on the predictability of beta involve the beta predictions computed using the simplified SCL from CAPM, referenced in Equation 2,

$$r_{a,t} = (\hat{\beta} \times r_{m,t}) + \varepsilon_{a,t} \tag{22}$$

where $r_{m,t}$ is the observed market return at time t, $\hat{\beta}$ is defined as the 1-day forward corresponding beta prediction at time t, $\varepsilon_{a,t}$ is a random error term and $r_{a,t}$ represents the predicted return on *Asset a* at time *t*. The Mean Square Error (MSE) measures the test of return accuracy from the predicted beta that is conditioned on the observed market return. MSE is defined as

$$MSE = R(t) = \frac{1}{n}\sum_{i=1}^{n}(R_{a,t} - \hat{R}_{a,t})^2 \tag{23}$$

where n is the number of predictions contained, $R_{a,t}$ is the predicted out of sample realized return on *asset a* at time t and $\acute{R}_{a,t}$ is the observed realized return on *asset a* at the corresponding time t. For the benchmark comparisons, the temporally corresponding out of sample constant mean return is used as the $R_{a,t}$.. Alternatively, the predictive return calculated as the sum of the temporally corresponding constant beta conditioned on the observed market return takes the form of $R_{a,t}$ for MSE comparison.

In order to standardize the results, the MSE is multiplied by the number of trading days in a year, 252, to yield annualized units. Additionally, to further standardize the results and yield an ease of interpretation, the MSE is converted into standard deviation units via conversion to the Root Mean Square Error (RMSE) :

$$RMSE = \sqrt{MSE \times 252} \tag{24}$$

The resulting RMSE can be interpreted in annualized standard deviation units and produces a reliable estimator of the variance, known as the standard error. RMSE will be used as the measure of accuracy and comparison for the predictive beta calculated from the model, the mean return, and the constant beta.

5. Findings

Table 1 summarizes the standard deviations of the estimation sample betas across all equities using the 10-minute sampling frequency. As described in section 4.1, if beta were constant over time, then the inequality between the variation of an asset with respect to market variation would not statistically deviate over time. But, in account for the microstructure noise at the 10 minute sampling frequency, the average standard deviation of realized beta over the whole sample interval is still 0.3737, which is statistically significant under the null hypothesis of no deviations. Additionally, Table 2 summarizes the results of the first order autocorrelations from the whole sample interval. As described in section 4.2, positive first order autocorrelations suggest the persistence in the realize beta. The average first order autocorrelation is approximately 0.3484, which substantiates the predictability in the time-varying element. It is notable that

Microsoft Corporation (MSFT) and Exxon Mobil Corporation (XOM) display a relatively low first order autocorrelation of 0.0866 and 0.1521, respectively.

The literature suggests that it is reasonable to expect that there is some instability in most econometric relationships across time or space. Typically in cross sections with market data, there is likely some degree of heterogeneity amongst assets (Elliott and Müller, 2006). There are numerous factors that contribute to the heterogeneity of a time series including regulation, economic and monetary policy, and exchange rates. Giacomini and White (2006) contest that as long as this heterogeneity is not "too strong", standard regression methods still have reasonable properties. The resulting regression coefficients can then approximate the "true" values of the coefficients with averages of the individual and/or inter-temporal true values. Table 3 illustrates the HAR-Beta coefficients computed from the regression of the estimation sample realized betas on the estimation sample HAR-Beta estimates and contains the corresponding R^2 values. The sum of the beta coefficients across all equities approximates one. Again, this alludes to the persistency of the regression coefficients and the accuracy in the resulting beta predictions. One exception is Microsoft Corporation (MSFT) that has a sum of coefficients equal to -0.2387. Notably, Exxon Mobil Corporation (XOM) has a sum of regression coefficients equal to 0.7623 that is relatively below one, especially when compared to the other results. These results make intuitive sense given the results from Table 2 which suggested that MSFT and XOM betas were relatively less predictable over the estimation sample period.

5.1 Findings for Dataset A

The out of sample dataset A refers to the time interval between January 3, 2006 and January 3, 2008. Table 4 summarizes the root mean squared errors (RMSE) calculated from the difference between observed realized returns and predicted returns conditioned on the observed market return and the predictive beta from the HAR-Beta model. The results from the constant beta benchmark comparison and the constant return benchmark comparison are also included for assessment. The direct comparison of RMSE between the hypothesized HAR-Beta model and the benchmark comparisons display a significant reduction. When compared to the constant returns benchmark, the HAR-Beta model gives an approximate 21.94% reduction in RMSE. Additionally, when compared to the constant beta benchmark comparison, the HAR-Beta model gives an approximate reduction of 6.62%.

Notably, Wal-mart Stores (WMT), has a 35.63% increase in RMSE when the HAR-Beta model is compared to constant returns. The result seems surprising when referring back to the results from Table 1, Table 2 and Table 3 in which WMT highest first order autocorrelation (0.5701) and the highest R^2 value (0.5391). Perhaps, the result could be the outcome of outliers that complicate the extrapolation of OLS regression models. Poon and Granger (2003) note the common problems of possible sample outliers in volatility estimation and suggest alternate methods of robust regressions.

The results from Table 2 and Table 3, which suggested that MSFT and XOM were relatively less predictable given the lack of persistence, are concurred by the subsequent analysis. The R^2 values for MSFT and XOM, as displayed in Table 3 are also the lowest at 0.0113 and 0.1106, respectively. As noted in Table 4, when constant returns of MSFT and XOM are compared to the HAR-Beta model, there is a relatively low 0.07% and 5.98% reduction, respectively. But, when compared to the constant beta benchmark comparison, MSFT and XOM display the largest RMSE reductions of 13.45% and 7.96%, respectively. This point is interestingly quite salient. Despite the low predictability of beta measured by the results from Table 2 and Table 3, which are then concurred with the low R^2 values in Table 3, the RMSE of the HAR-Beta model are still reduced in terms of both benchmark comparisons. Alternatively, in both the cases of MSFT and XOM, when the constant beta is compared to constant returns, there is a significant increase in RMSE. The overall reduction in the RMSE highlights the information revealed through the conditional probability distributions of the HAR-Beta model that gives better estimates of the true underlying beta. The results from these equities seemingly confirm the results displayed in Table 1 that beta deviates over time, despite the indicators of lacking persistence and predictability.

5.2 Findings for Dataset B

The out of sample dataset B refers to the time interval between January 3, 2006 and January 3, 2009. Table 5 summarizes the root mean squared errors (RMSE) calculated from the difference between observed realized returns and predicted returns conditioned on the observed market return and the predictive beta from the HAR-Beta model. The results from the constant beta benchmark comparison and the constant return benchmark comparison are also included for assessment. The driver for this out of sample interval, as defined earlier, is to check the robustness of the HAR-Beta during a period characterized with financial uncertainty and sporadic volatility. It is important to note, the increase in RMSE present when comparing constant returns to the HAR-Beta model of Microsoft Corporation (MSFT). The point once again can be surmised to the presence of an outlier that leads to extrapolation errors. Aside from this result, the RMSE are reduced by approximately 19.67% when compared to the constant Beta benchmark, and approximately 39.28% (including the MSFT increase) when compared to the constant return benchmark.

6 Conclusion

In line with recent literature that suggests alternatives to the constant beta conventionally used within the Capital Asset Pricing model, this paper generates evidence for an alternative model that explains variation in beta over time. At a determined optimal sampling frequency of 10 minutes chosen to limit the effect of market microstructure noise, time-series plots and first order autocorrelations of beta reveal persistence in realized beta over time. The adaptation of the Heterogeneous Autoregressive framework in regression analysis yields predictive betas that produce returns with one-half to two-thirds the fraction of root mean squared error when compared to the constant mean or constant beta models. The conditional probability distributions of the predictive betas generated from the HAR-Beta model better represent the distribution of the true underlying beta and hence better explain changes in market and asset returns.

It is important to note the inherent weakness of ordinary least squares regression analysis when dealing with outliers. Given one or two outliers, the R^2 of a sample can inaccurately represent the actual fit of the regression and give imprecise approximates of the coefficient. Specifically, OLS may be bias to the outliers and artificially inflate the variance of the estimates. The misspecification can then lead to an incorrect probability distribution that may inaccurately represent the true data. But, this issue can be partially remedied with the usage of a logarithmic scale that normalizes the data points. Specifically within this analysis, the logarithmic scale is coupled with the optimal sampling frequency to further dismiss the effects of distortions and outliers present within the dataset. A similar but more effective approach is to use robust regressions that are designed to circumvent limitations in parametric and non-parametric methods, such as outliers. To the extent that outliers exist within the dataset, this approach may better produce estimates of the true underlying beta than the ordinary least squares approach.

Overall, the results from this analysis emphasize the importance of looking beyond the constant beta to accurately asses the risk within a particular investment. By using temporally varying realized measures to integrate instantaneous information in the calculation of conditional beta probability distributions, the true underlying beta can be better approximated. The replication of the true underlying beta can be instrumental to risk-advisors and portfolio managers in that they can better explain their rationale for taking incremental risk to investors.

7 Figures

Figure 1A: Volatility Signature Plots (January 2, 2001- January 3, 2009)

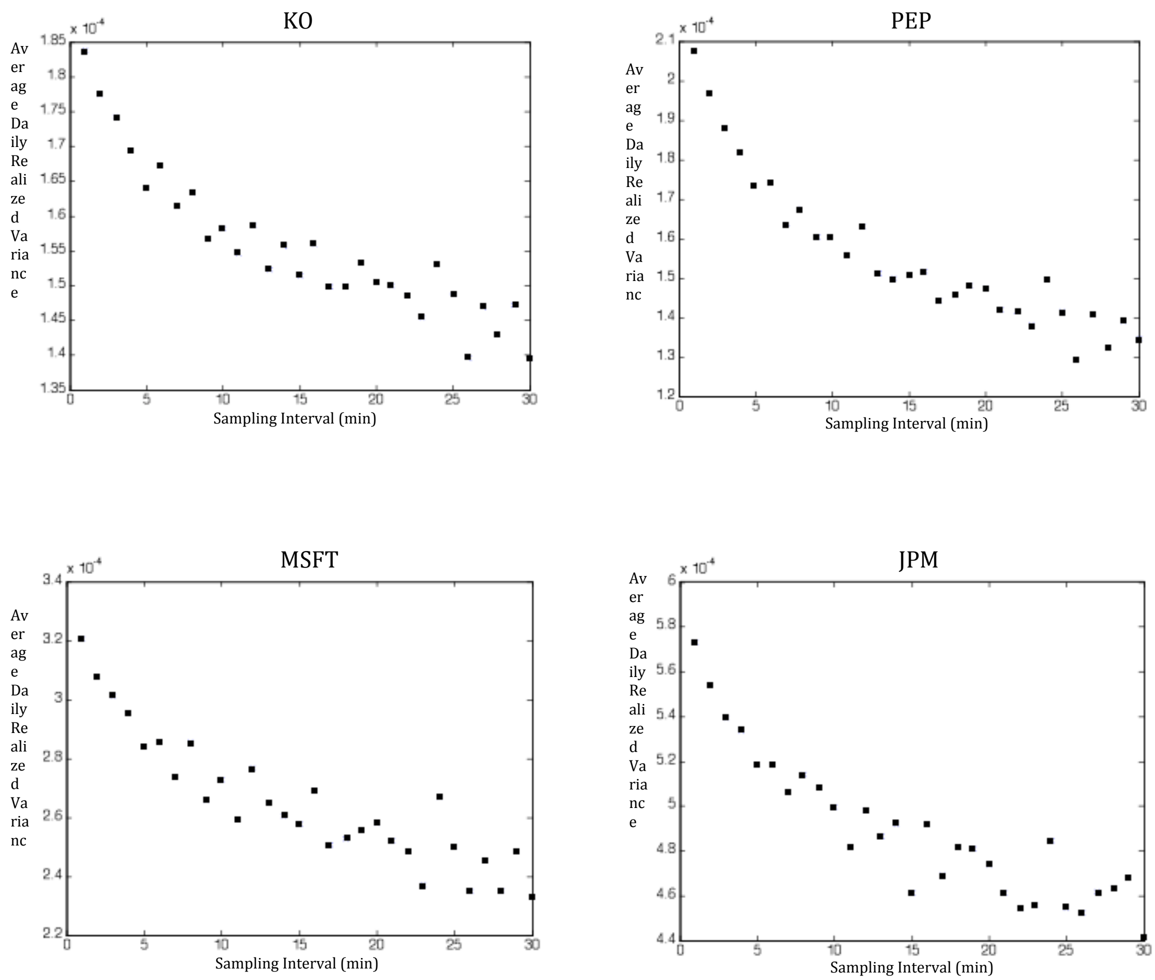

Figure 1B: Volatility Signature Plots (January 2, 2001- January 3, 2009)

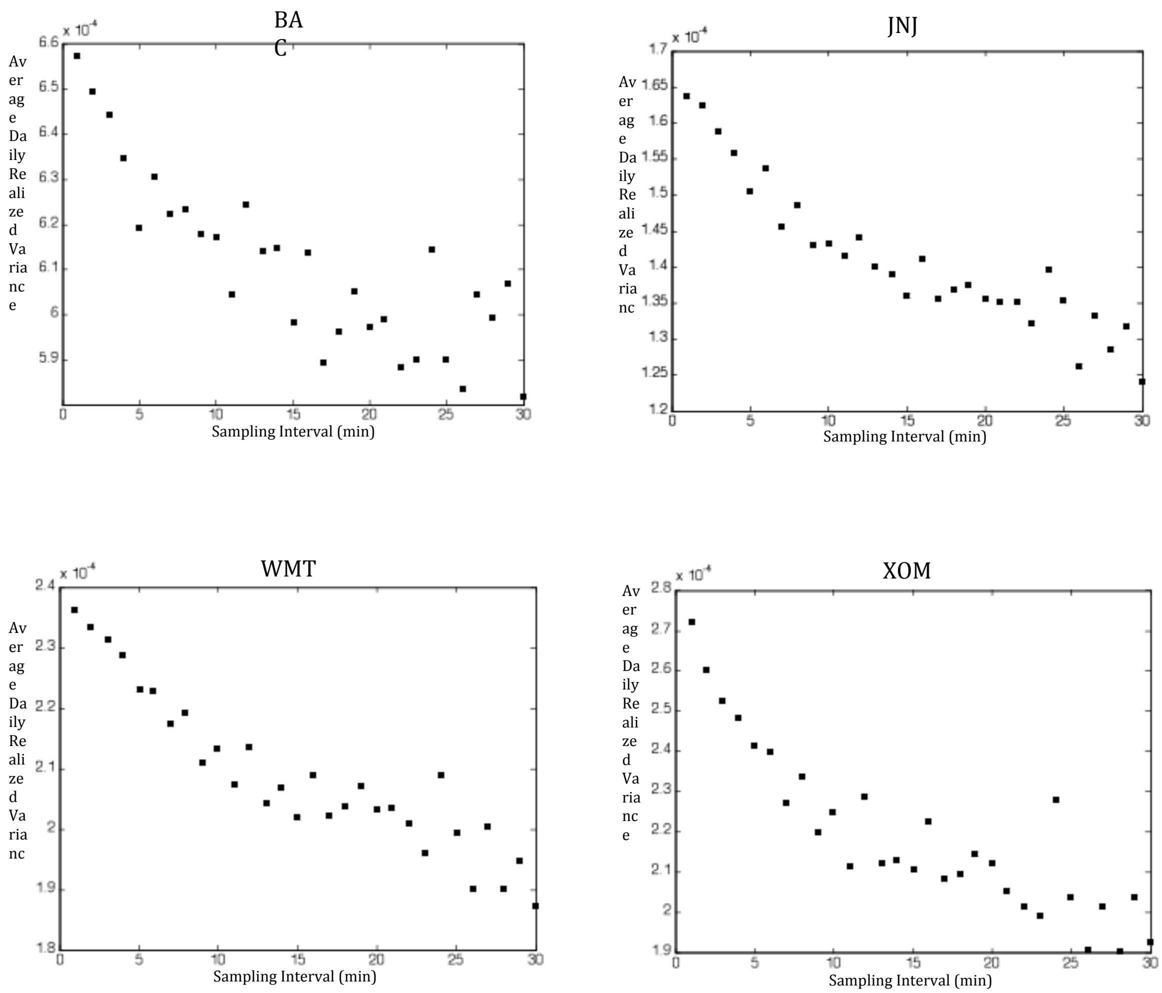

Figure 2A: Time Series- Realized Betas (January 2, 2001 – January 3, 2009)

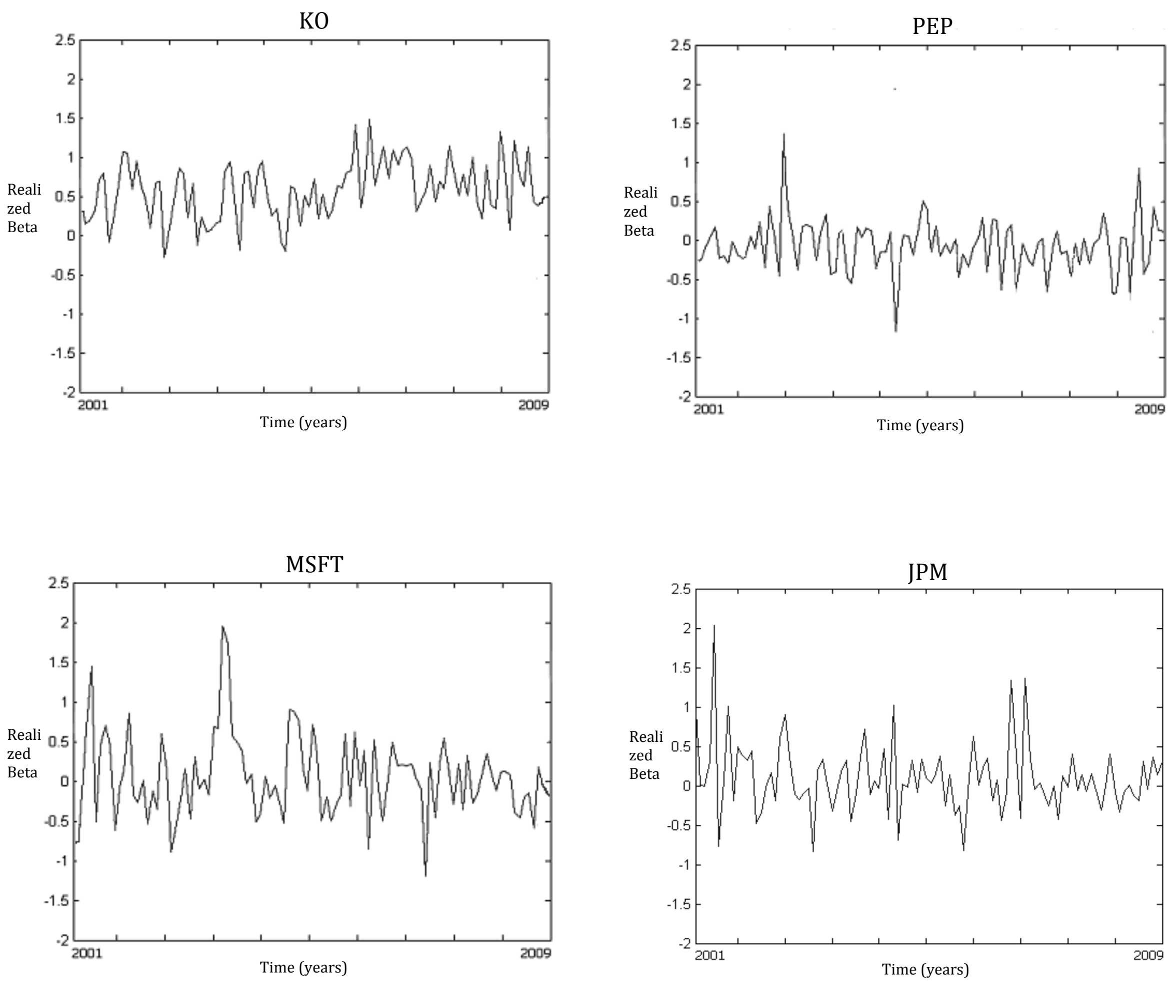

Figure 2B: Time Series- Realized Betas (January 2, 2001 – January 3, 2009)

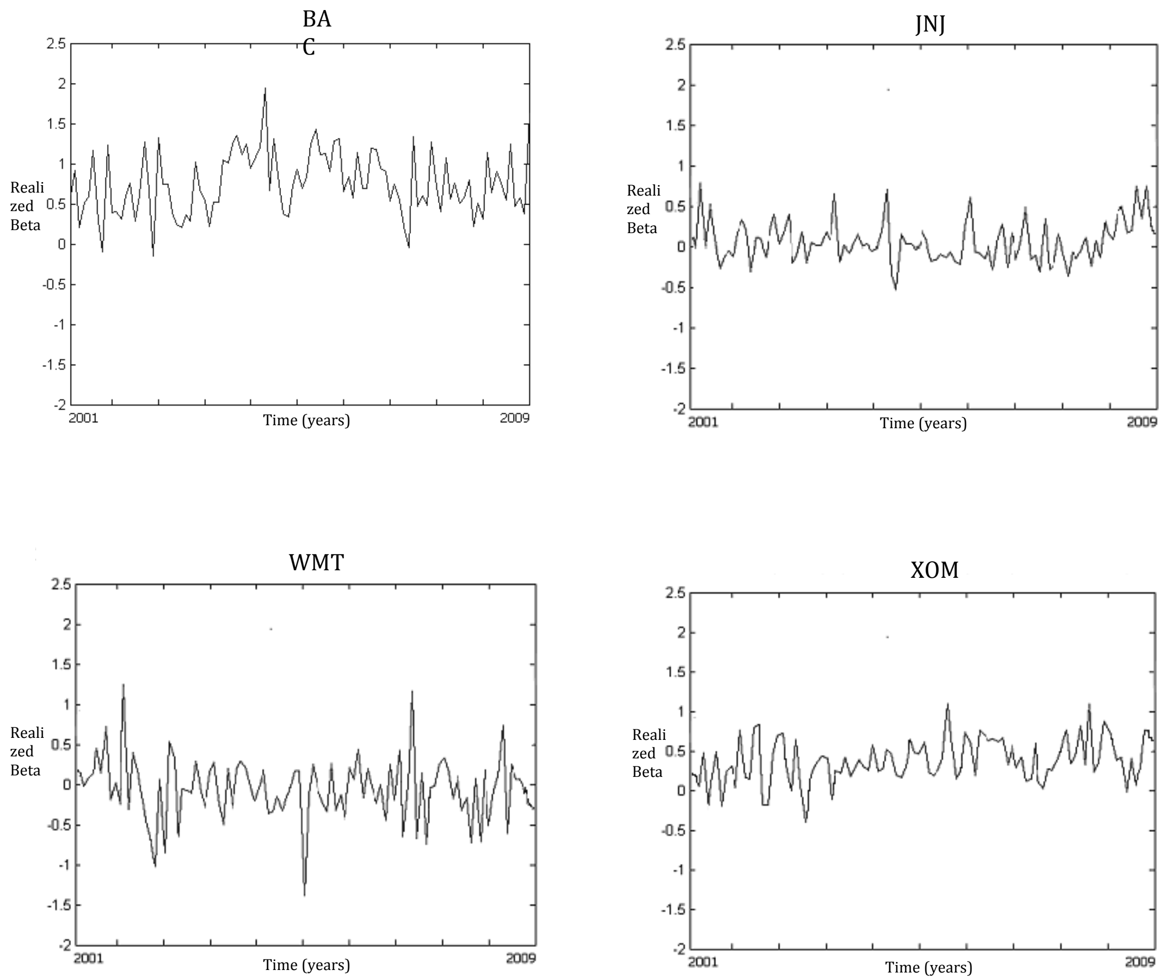

8 Tables

All tables display statistics that use 10-minute sampling intervals.

Table 1: Standard Deviation of Beta

Company	**Standard Deviation**
Coca Cola Company (KO)	0.3925
PepsiCo, Inc. (PEP)	0.2745
Microsoft Corporation (MSFT)	0.3417
JPMorgan Chase & Co. (JPM)	0.4485
Bank of America Corporation (BAC)	0.3721
Johnson & Johnson (JNJ)	0.3822
Wal-mart Stores Inc. (WMT)	0.4783
Exxon Mobil Corporation (XOM)	0.2997

All tables display statistics that use 10-minute sampling intervals.

Table 2: AR(1)- First Order Autocorrelation of Beta

Equity	**First Order Autocorrelation**
Coca Cola Company (KO)	0.5345
PepsiCo, Inc. (PEP)	0.2910
Microsoft Corporation (MSFT)	0.0866
JPMorgan Chase & Co. (JPM)	0.3296
Bank of America Corporation (BAC)	0.3550
Johnson & Johnson (JNJ)	0.4673
Wal-mart Stores Inc. (WMT)	0.5701
Exxon Mobil Corporation (XOM)	0.1531

All tables display statistics that use 10-minute sampling intervals.

Table 3: HAR-Beta Regression Coefficients

	β_0	$R\beta_{t,\,t-1}$	$R\beta_{t,\,t-5}$	$R\beta_{t,\,t-22}$	Sum of Coefficients	R^2
KO	0.0096 (0.0199)	*0.0289 (0.0148)	*0.2695 (0.1396)	**0.6367 (0.2457)	0.9447	0.4784
PEP	**0.0051 (0.0122)	*0.0201 (0.0101)	*0.0880 (0.0447)	**0.6754 (0.2330)	0.7886	0.0019
MSFT	0.0107 (0.0149)	**0.1023 (0.0499)	- 0.0226 (0.1142)	- 0.3291 (0.2669)	- 0.2387	0.0113
JPM	*0.0078 (0.0039)	*0.0210 (0.0104)	**0.4598 (0.1121)	**0.3476 (0.1144)	0.8362	0.2547
BAC	*0.0011 (0.0002)	*0.0416 (0.2133)	*0.0864 (0.0422)	**0.7470 (0.1241)	0.8761	0.2797
JNJ	*0.0066 (0.0034)	*- 0.0375 (0.0193)	**0.3764 (0.1202)	**0.5924 (0.1155)	0.9379	0.4361
WMT	*0.0094 (0.0048)	-0.0151 (0.0525)	**0.4014 (0.1196)	**0.5613 (0.1115)	0.9570	0.5391
XOM	*0.0034 (0.0017)	0.0094 (0.0504)	*- 0.0934 (0.0457)	*0.8429 (0.4323)	0.7623	0.1106

*The significance levels of the coefficients are denotes by the asterisk: * → $p < 0.05$, ** → $p < 0.01$*

All tables display statistics that use 10-minute sampling intervals.

Table 4: Root Mean Squared Error (RMSE)-HAR-Beta, Constant Beta, Constant Return, and Standard Deviation of Beta (*Out of Sample Dataset-A*)

	RMSE HAR-Beta	**RMSE Constant Beta**	**Constant Returns**	**Standard Deviation of Beta (Out Sample)**
KO	0.1732	0.1867	0.4935	0.2628
PEP	0.1963	0.2095	0.2511	0.2703
MSFT	0.2831	0.3271	0.2833	0.3800
JPM	0.4347	0.4763	1.1433	0.4410
BAC	0.3340	0.3410	0.5289	0.3479
JNJ	0.1650	0.1674	0.2716	0.2425
WMT	0.3441	0.3652	0.2215	0.3614
XOM	0.2892	0.3142	0.3076	0.3669

All units are expressed in Annualized Standard Deviation Units.

All tables display statistics that use 10-minute sampling intervals.

Table 5: Root Mean Squared Error (RMSE)-HAR-Beta, Constant Beta, Constant Return, and Standard Deviation of Beta (*Out of Sample Dataset-B*)

	RMSE HAR-Beta	**RMSE Constant Beta**	**Constant Returns**	**Standard Deviation of Beta (Out Sample)**
KO	0.1953	0.2496	0.5994	0.2442
PEP	0.1986	0.2195	0.3570	0.2575
MSFT	0.2291	0.3796	0.1774	0.3613
JPM	0.4262	0.6329	1.2493	0.4884
BAC	0.4706	0.4924	0.6348	0.4208
JNJ	0.2149	0.2230	0.5767	0.2207
WMT	0.2293	0.3106	0.3274	0.3339
XOM	0.2188	0.2719	0.4135	0.3476

All units are expressed in Annualized Standard Deviation Units.

References

[1] Adrian, T., Franzoni, T. (2005). Learning about beta: time-varying factor loadings, expected returns, and the conditional CAPM. *Federal Reserve Bank of New York.*

[2] Andersen T.G., Bollerslev, T. (1998). Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts. *International Economic Review*, *39*(4).

[3] Anderson, T.G., Bollerslev, T., Diebold, F.X., Labys, P. (1999). Realized Volatility and Correlation. *Risk*, 104-108.

[4] Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P. (2003). Modeling and Forecasting Realized Volatility. *Econometrica*, *71*(2), 579-625.

[5] Andersen, T.G., Bollerslev, T., Diebold, F.X., Wu, G. (2004). Realized Beta: Persistence and Predictability. *Advances in Econometrics*, *20*(2), 1-39.

[6] Anderson, T.G., Bollerslev, T., Diebold, F.X, Wu, J. (2006). Realized beta: Persistence and Predictability. *Advances in Econometrics*, *20*, 1-39.

[7] Andersen, T.G., Bollerslev, T., Huang, X. (2007). A Reduced Form Framework for Modeling Volatility of Speculative Prices based on Realized Variation Measures. *The Review of Economics and Statistics*, *89*(4), 701-720.

[8] Anderson, T.G., Bollerslev, T., Meddahi, N. (2011). Realized volatility forecasting and market microstructure noise. *Journal of Econometrics*, *160*(1), 220-234.

[9] Ang, A., Chen, J. (2007). CAPM over the long run: 1926-2001. *Journal of Empirical Finance*, *14*(1), 1-40.

[10] Amatyakul, P. (2009). The Relationship Between Trading Volume and Jump Processes in Financial Markets. Duke University Senior Honors Thesis.

[11] Baillie, R.T., Bollerslev, T., Mikkelsen, H.O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedastciity. *Journal of Econometrics*, *74*(1), 3-30.

[12] Banz, R.W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics*, *9*(1), 3-18.

[13] Bos,T., Newbold, P. (1984). An Empirical Investigation of the Possibility of Stochastic Systematic Risk in the Market Model. *Journal of Business*, *57*(1), 35-41.

[14] Brooks, R.D., Faff, R.W., Lee, J.H.H. (1992). The form of time variation of systematic risk: Some Australian Evidence. *Applied Financial Economics*, *2*, 191-198.

[15] Campbell, J.Y., Lo, A. W., MacKinlay, A.C. (1997). The Econometrics of Financial Markets. Princeton New Jersey.

[16] Campbell, J.Y., Vuolteenaho, T. (2004) Inflation Illusion and Stock Prices. *The American Economic Review*, *94*(2), 19-23.

[17] Choudhry, T. (2002). The Stochastic Structure of the Time Varying Beta: Evidence from UK Companies. *The Manchester School*, *70*(6), 768-791.

[18] Choudhry, T. (2004). Time-varying beta and the Asian financial crisis: Evidence from Malaysian and Taiwanese firms. *Pacific-Basin Finance Journal*, *13*(1), 93-118.

[19] Collins, D.W., Ledolter, J., Rayburn, J. (1987). Some Further Evidence on the Stochastic Properties of Systematic Risk. *The Journal of Business*, *60*(3), 425-448.

[20] Corsi, F. (2003) A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*.

[21] Fama, E.F. (1976). Foundations of Finance: Portfolio decisions and securities prices. New York.

[22] Fama, E.F., French, K.R. (2005). Financing decisions: who issues stock?. *Journal of Financial Economics*, *76*(3), 549-582.

[23] Giacomini R., White H. (2006) Tests of Conditional Predictive Ability. *Econometrica*. *74*(6), 1545-1578.

[24] Hamilton, J. (1994). Time Series Analysis. Princeton University Press.

[25] Hansen, P.R., Lunde, A. (2006). Realized Variance and Market Microstructure Noise. *Journal of Business and Economic Statistics*, *24*(2), 127-161.

[26] Huang, X., Tauchen, G. (2005). The Relative Contribution of Jumps to Total Price Variance. *Journal of Financial Econometrics*, *3*(4), 456-499.

[27] Klemkosky, R.C., Martin, J.D. (1975). The Adjustment of Beta Forecasts. *Journal of Finance*. *30*(4), 1123-1128.

[28] Lewellen, J., Nagel, S. (2006). The conditional CAPM does not explain asset-pricing anomalies. *Journal of Financial Economics*, *82*(2), 289-314.

[29] Müller, U.K., Elliot, G. (2007). Efficient Tests for General Persistent Time Variation in Regression Coefficients. *The Review of Economic Studies*, *73*(4), 907-940.

[30] Petkova, R., Zhang, L. (2005). Is value riskier than growth?. *Journal of Financial Economics*, *78*(1), 187-202.

[31] Poon, S., Granger C.W.J. (2003). Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature*, 41(2), 478–539.

[32] Song, D. (2009). Volatility Forecasting in the Presence of Market Microstructure Noise. Duke University Senior Honors Thesis.

[33] Tofallis, C. (2008). Investment Volatility: A Critique of Standard Beta Estimation and a Simple Way Forward. *European Journal of Operational Research*, *187*(3), 1358–1367.