Forecasting Beta Using Conditional Heteroskedastic Models

Andrew E. Bentley

Professor George Tauchen, Faculty Advisor

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

Duke University
Durham, North Carolina
2012

The Duke Community Standard was upheld in the completion of this report.

Acknowledgments

I would like to thank Professor George Tauchen for leading the Honors Finance Seminar and for his guidance throughout the development of this paper. Also, I would like to thank Professor Tim Bollerslev and my peers in the Honors Workshop for their advice and assistance in the research, development, and writing of this paper. Especially, I would like to thank my Honors Workshop classmate Jackson Pfeiffer for his constructive feedback over the course of the process. Finally, I would like to thank both my parents for all their love and support while writing this paper and throughout my education.

Abstract

Conventional measurements of equity return volatility rely on the asset's previous day closing price to infer the current level of volatility and fail to incorporate information concerning intraday fluctuations. Realized measures of volatility, such as the realized variance, are able to integrate intraday information by utilizing high-frequency data to form a very accurate measure of the asset's return volatility. These measures can be used in parallel with the traditional definition of the Capital Asset Pricing Model (CAPM) beta to better predict the time-varying systematic risk of an asset. In this analysis, realized measures were added to the General Autoregressive Conditional Heteroskedastic (GARCH) framework to form a predictive model of beta that can quickly respond to rapid changes in the level of volatility. The findings suggest that this predictive beta is better able to explain the stylized characteristics of beta and is a more accurate forecast of the realized beta than the GARCH model or the benchmark Autoregressive Moving-Average (ARMA) model used as a comparison.

JEL classifications: C03, C32, C53, C58
Keywords: Beta, GARCH, GARCHX, Realized Variance, High-Frequency Data

Contents

1 Introduction

In exchange for holding a security, investors expect to be compensated based upon the inherent systemic risk in the held asset as well as the time-varying opportunity cost of their invested money. This opportunity cost, referred to as the time value of money, implies that money available at the present time is worth more than the same amount in the future due to its potential earning capacity. This core principal of finance holds that, provided money can earn interest, any amount of money is worth more the earlier it is received. The opportunity cost of money also changes over time due to any combination of factors including the availability of potential earnings through alternative investments, decreases in relative value due to inflation, or the potential of unrealized earnings from interest.

When analyzing the potential return of a security, investors typically compare the investment's expected return relative to the guaranteed rate of return they could receive from an investment with zero risk. This albeit theoretical rate of return is known as the risk-free rate and represents the interest an investor would expect from an absolutely risk-free investment over a specified period of time. In practice, however, the risk-free rate does not exist because even the safest investments carry a small degree of risk. Thus, assuming that governments are free from the possibility of default, the interest rate of a short-term treasury bill is often used as a proxy for the risk-free rate of return. Given the minimal risk inherent with this type of investment, the associated return is similarly quite low. There exists a wealth of assets with greater inherent risk than the guaranteed rate of return that are attractive due to their potential to yield returns exceeding the risk-free rate. The ratio of relative risk between potential risky assets and the risk-free rate is captured in the statistical coefficient referred to as beta.

Beta is a frequently employed statistical measure that characterizes the volatility of an asset's returns in relation to the volatility of the market's returns. Specifically, beta represents the systemic risk present in an asset that can not be avoided through portfolio diversification. The beta of the market itself is always one, because, by definition, the covariance of the market when compared with itself is equal to one. Alternatively, a beta less than one indicates less movement by an asset compared to the market, and a beta larger than one designates a higher degree of movement. Finally, a beta of zero indicates that the returns of an asset are independent of the market's returns.

Within security analysis and valuation, as well as portfolio management, beta is central in determining the pricing of equities and managing the risk of investment portfolios. One of the fundamental principals of asset pricing theory, as noted by Andersen et al. [2006], is that only systematic risk should be priced. In the one-factor capital asset pricing model (CAPM), for example, systematic risk is determined by the measure beta [Sharpe, 1963; Lintner, 1965a,b], which is central in calculating the cost of equity for a particular security.

The cost of equity is used to calculate a firm's discount rate, which is in turn utilized to compute the present value of future cash flows and determine a valuation of the firm and the equity. Despite much criticism over the last 50 years, the CAPM remains at the forefront of academic research and industrial application.

This paper focuses on the use of beta within the CAPM and whether the systemic risk of an asset, measured by beta, is constant over time or evolves continuously. Beta, defined as the covariance between an asset's returns and the market's returns divided by the variance of the market's returns, has traditionally been assumed to be time invariant within the CAPM. Numerous pieces of literature including Bos and Newbold [1984], Jagannathan and Wang [1996], Groenewold and Fraser [1999], and Choudhry [2002] state that the static CAPM is unable to adequately explain the cross section of average returns on stocks and is unable to accurately capture the dynamics of volatility. However, several publications on economic volatility such as Bollerslev et al. [1988], Kim and Nelson [1989], Bollerslev and Engle [1993], Andersen et al. [2002], and Ewing and Malik [2005] have presented extensive evidence for time-varying fluctuations in conditional variances of stocks and their conditional covariances with the market. Thus, from a purely statistical standpoint, beta, comprised of these conditional variances and covariances, should be expected to inherit the fluctuations present in its constituent components.

Multiple time-varying beta models have been proposed to estimate the true underlying beta of an asset [Ebner and Neumann, 2005; Morana, 2009; Caporale, 2012]. Of these, most employ coarsely sampled daily returns (or more often squared daily returns) to extract information about the current degree of volatility and then utilize this information to forecast the next period's expected volatility. These betas are often slow to react and fail to capture the full dynamics of equity betas during periods of rapid movement in volatility. To address this issue, Andersen et al. [2006] presented a comprehensive framework which utilizes high-frequency data to estimate beta.

High-frequency data, measured at generally one- or five-minute intervals, fundamentally provides more information regarding the dynamics of the true underlying volatility. Andersen et al. [2006] introduced measures of the variance of an asset's returns as well as its covariance with market returns which utilize high-frequency data. These measures, which are known as realized measures, are more informative concerning the current level of volatility than observed daily square returns because these measures integrate new information almost instantly. Combining measures of variance and covariance, Andersen et al. [2006] developed a similar beta model computed from high-frequency intraday returns, which they call realized beta. Their results indicated that the use of high-frequency data and realized measures within the realized beta calculation allows for a more dynamic measurement of an asset's true underlying beta.

It has long been accepted that daily asset returns are essentially unpredictable, while on the other hand return volatility has been shown to be highly persistent and thus predictable. Therefore, because volatility is predictable and extremely important in financial market transactions and risk management, forecast models are both feasible and meaningful. There are two main approaches to modeling volatility. The first is parametric modeling such as stochastic volatility models. The second approach is to use volatility implied by derivative prices. Both methods have certain flaws. Traditional parametric models utilize the previous day's closing price of an asset and fail to incorporate intraday information. Implied volatility models have an advantage being derived directly from market prices; however, they rely on the validity of their underlying pricing models. The introduction of more informative realized measures of volatility has lead to a shift in volatility measurement, modeling, and forecasting and given increased robustness to parametric modeling. Because beta is defined as the ratio of the covariance of an asset's and the market's returns over the variance of the market's returns, these same methodologies can be adapted to form predictive models of beta utilizing high-frequency intraday data.

This paper provides a framework for a model to forecast time-varying conditional beta utilizing high-frequency data. The model developed in this paper builds on the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) framework of Bollerslev [1986] to calculate the time-varying conditional variances of equity and market returns and to combine these to form a framework for measuring and forecasting time-varying conditional beta. The structure of the paper is the following. In Section 2, the theory concerning the traditional time-invariant beta is reviewed. From there, the complementary theoretical continuous-time beta is developed and its lack of observability is illustrated. Finally, the realized measures used to approximate the theoretical continuous-time beta are presented and the concept of realized beta is introduced. In Section 3, the parametric models used to fit and forecast the constituent components of realized beta are developed and the final modified GARCH model with an exogenous variable (GARCHX) is presented, which incorporates realized measures as an exogenous explanatory variable in forecasting the conditional volatility. Section 4 provides the background on the Autoregressive Moving Average (ARMA) model used as a benchmark for the calculated betas from the GARCH and GARCHX models. Section 5 provides the criteria for the evaluation of the proposed GARCHX model. Finally, Section 6 summarizes the main results from the analysis and illustrates that the GARCHX model performs as well if not better than both the GARCH model and the baseline ARMA model. Some concluding remarks are offered in Section 7 as well as suggestions for future research.

2 Beta

2.1 Static Beta

The coefficient beta (β) is a key component of the CAPM and is a mathematical measure of the sensitivity of the rate of return of a portfolio or given security compared with the rate of return of the market portfolio[1]. According to asset pricing theory, beta represents the type of risk, know as systematic risk, that can not be minimized or removed by the diversification provided by the portfolio of several risky assets, due to the correlation of an asset's returns with the returns of all other assets within the portfolio. Beta is derived from the regression analysis of the returns of a specific asset against those of the market portfolio over a given time period. The results of this regression yield the Security Characteristic Line (SCL)

$$r_{i,t} - r_f = \alpha_i + \beta \times (r_{m,t} - r_f) + \epsilon_{i,t}, \tag{1}$$

where $r_{i,t}$ is the rate of return of the asset at time t, $r_{m,t}$ is the rate of return of the market at time t, and r_f is the risk-free rate typically taken as the rate of return on a short-term treasury note. Alpha (α_i) measures the excess return that a portfolio or security makes over and above what would be predicted by an equilibrium model such as the CAPM. Simply stated, alpha often is considered to represent the value that a portfolio manager adds to (or subtracts from) a fund's return. However, if it is assumed that markets are efficient and that investors are subsequently only compensated for the systematic risk of an asset, then the expected value of excess returns is zero for all assets ($E(\alpha_i) = 0$). Additionally, throughout this analysis high-frequency data sampled at discrete time intervals of one or five minutes is used. With the use of such finely sampled data, it can be assumed that the risk-free rate does not change significantly from one period to the next. Therefore, it is taken to be zero resulting in the following simplified equation for the SCL:

$$r_{i,t} = \beta \times r_{m,t} + \epsilon_{i,t}. \tag{2}$$

Equation 2 illustrates the relationship of the return of an asset i at time t with the return of market. Moreover, it explicitly portrays the characteristic of beta as a factor of the response of an asset's returns to those of the market. The linear regression analysis used to calculate the SCL and determine beta seeks to fit the best fit line to the data set such that the line minimizes the sum of the squared residuals of the linear regression model. The value corresponding to the best fit for the coefficient beta in Equation 2 can be shown mathematically to be the covariance of the asset and market returns scaled by the variance of the returns of

[1]The market portfolio should in theory include all types of assets that are held by anyone as an investment (including works of art, real estate, human capital, etc.) In practice, such a market portfolio is unobservable and people usually substitute a stock index such as the S&P500, as a proxy for the true market portfolio.

the market. Beta is therefore defined as

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

where $r_{i,t}$ is the return of an asset i at time t and $r_{m,t}$ is the return of the market at time t. By definition the market has a beta of one because the covariance of the market with itself is equal to the market variance. A positive beta indicates that an asset generally follows the market benchmark in that it tends to increase and decrease with the market. Conversely, a negative beta means that an asset moves opposite the market. Finally, an asset has a beta of zero if its price moves are completely uncorrelated with those of the market.

To compute the variance and covariance of market and asset returns, the change in security prices must first be measured. This analysis will utilize logarithmic prices and geometric returns in order to provide an accurate comparison between data points and to standardize the measurement scale. Suppose the price process is sampled equally S times per period, we then define the time interval S period geometric return as

$$r_{i,t,j} = p(t + \frac{j}{S}) - p(t + \frac{j-1}{S}), \quad j = 1, 2, ..., S \tag{4}$$

where p is the logarithm of the asset price and t is the specified time period. Using this definition of returns and the above equation for beta, the constant beta of a security over a specified time interval is able to be computed.

The beta discussed above and traditionally used within the CAPM is a static measure. Within financial economics, there has been significant amounts of research questioning whether the beta of a security is constant. Andersen et al. [2006] presents several pieces of evidence that indicate that betas are in fact time-varying. First, several publications assert this fact including Huang and Litzenberger [1988] and Andersen et al. [2006]. Second, in the absence of an explicitly defined time-varying framework, betas are typically estimated over moving time intervals of generally five or ten years, presumably to account for the variation of beta over the measurement period (see Fama [1976] and Campbell et al. [1997]). Finally from an empirical perspective, there has been extensive research within the field of financial econometric volatility on the presence of wide fluctuations and high persistence in individual equity conditional variances and covariances with the market [Andersen et al., 2002]. Betas, which are the ratios of time-varying variances and covariances, should be expected to display similar time-varying properties as discussed by Bollerslev et al. [1988]. In fact, unless there is a systematic cancellation between the variance and covariance terms, betas will inherit the features that are present in their constituent components [Andersen et al., 2006]. Therefore, it is necessary to shift from a static representation of beta to a model of beta based upon the time-varying nature of its underlying variance and covariance.

2.2 Continuous-Time Beta

In order to model beta as a continuous-time measurement, similar models for the prices of individual assets and the market are required. Throughout this paper the underlying theory and related empirical strategies developed in Andersen et al. [2001a,b, 2003] and by Barndorff-Neilson and Shephard [2004] will be followed. A stochastic model of returns is assumed where the efficient price process for an asset ($p_i(t)$) and the market ($p_m(t)$) are given by the following differential equations:

$$dp_i(t) = \mu_i(t)dt + \sigma_i(t)dw_i(t) \tag{5}$$

$$dp_m(t) = \mu_m(t)dt + \sigma_m(t)dw_m(t) \tag{6}$$

where $\mu_i(t)$ and $\sigma_i(t)$ give the instantaneous drift[2] and the volatility of an asset, respectively, and $dw_i(t)$ is a Weiner increment[3] for the standard diffusion process $w_i(t)$. Both $\sigma_i(t)$ and $\mu_i(t)$ are strictly stationary[4] and are jointly independent of $w_i(t)$. The previous also holds true for $\mu_m(t)$, $\sigma_m(t)$, $dw_m(t)$, and $w_m(t)$ concerning the market process. These equations imply that the efficient price of a security follows a general diffusion process. Therefore, the continuous-time beta can be defined as the instantaneous covariance of the change in price of both the market and the asset over the instantaneous variance of the change in market price,

$$\beta_t = \frac{Cov(dp_i(t), dp_m(t))}{Var(dp_m(t))}. \tag{7}$$

However, instantaneous price movements cannot be observed; and therefore, the continuous-time beta of an asset is unable to be calculated for a singular point in time. As an alternative, the continuous-time volatility can be modeled over a fixed period and approximated using discrete measurements.

For a fixed interval of time, Andersen and Bollerslev [1998b] and Barndorff-Nielsen and Shephard [1998] developed the concept of the so-called integrated variance and integrated covariance as an approximation of the continuous-time processes. It follows that the interval-based beta of an individual security be defined as the ratio of the integrated covariance of the asset's and market's returns over the integrated variance of the market's returns. If the instantaneous volatility ($\sigma(t)$) were known, then the integrated variance and

[2] The drift (or drift rate in this case) is an adjustment that leads to the extension of the random walk model to include the tendency to move or "drift" in one direction or the other. In a random walk with drift model, the best forecast of the series tomorrow is the value of the series today plus the drift [Tsay, 2010].

[3] A Weiner increment is a special stochastic process with zero drift and variance proportional to the length of the time interval. As a result, the rate of change in expectation is the drift rate (μ) and the rate of change in variance is σ^2 [Tsay, 2010].

[4] A stationary process (or strictly stationary process) is a stochastic process whose joint probability distribution does not change when shifted in time or space. Consequently, parameters such as the mean and variance, if they exist, also do not change over time or position.

covariance could be calculated by computing the integrals

$$Integrated \quad Variance_t \equiv \int_{t-1}^{t} \sigma_m^2(\tau)d\tau \tag{8}$$

$$Integrated \quad Covariance_t \equiv \int_{t-1}^{t} \sigma_{m,i}(\tau)d\tau \tag{9}$$

$$\beta_{i,t} \equiv \frac{\int_{t-1}^{t} \sigma_{i,m}(\tau)d\tau}{\int_{t-1}^{t} \sigma_m^2(\tau)d\tau}. \tag{10}$$

From an empirical perspective, various trading frictions[5] limit the frequency at which returns can be sampled, making the extraction of instantaneous volatility estimates infeasible for a large number of points within an individual day [Andersen and Bollerslev, 1998b]. However, given that it is possible to observe prices at specific time periods, discrete measures of variation can be used to numerically approximate the integrated variance and covariance.

2.3 Realized Beta

Conceptually, assets are able to be traded at any point in time during which markets are open. Therefore, returns and corresponding volatilities should, in principle, be able to be obtained over arbitrarily short time intervals. The common method for approximating the ex-post instantaneous volatility is to use the squared return innovation as a proxy over the relevant time period. While the squared returns provide an unbiased estimate for the latent volatility factor, there is significant inherent noise in generating returns and they as a result are an extremely noisy measure of ex-post volatility. To address this concern, Andersen and Bollerslev [1998b] developed an alternative framework and demonstrated how high-frequency data allows for the construction of markedly improved ex-post volatility measurements via cumulative squared intraday returns.

Andersen and Bollerslev [1998b] proposed the ideas of quadratic variation and covariation as alternatives to the conditional variance and covariance measures of volatility. In contrast to conditional variance and covariance, the quadratic variation and covariation for asset prices depend solely upon the realization of squared returns. Moreover, regardless of the specific price process, the quadratic variation and covariation are obtained by summing the instantaneous squares and cross-products of returns. Therefore, unlike the conditional variance and covariance, the quadratic variation and covariation can theoretically be observed using high-frequency returns.

[5] Trading frictions, also referred to as microstructure noise, include frictions such as non-synchronous trading, liquidity effects, bid/ask bounce and misrecording spreads. One case detailing how such frictions distort prices is Roll [1984] which illustrates how the presence of bid-ask spread leads to a negative correlation in observed price changes.

If an adequately high sampling frequency is set and the intraday squared returns are summed over the period of interest, it is possible to approximate the theoretical quadratic variation. This observed measure is formally known as the realized variance. Given a set of geometric returns over a specified time frame, Andersen and Bollerslev [1998b] define the realized variance as follows:

$$RV_t(S) = \sum_{j=1}^{S} r^2_{t+\frac{j}{S}}. \tag{11}$$

From the definition of the realized variance, it follows that the integrated covariance between an asset's and market's returns can be similarly estimated using the realized covariance defined as the sum of the products of the intraday asset and market returns according to the following equation:

$$RCOV_{i,m,t}(S) = \sum_{j=1}^{S} r_{i,t+\frac{j}{S}} \cdot r_{m,t+\frac{j}{S}}. \tag{12}$$

Therefore, for any sampling frequency, the realized variance and covariance are directly observable in contrast to their underlying theoretical counterparts, the quadratic variation and covariation. Andersen and Bollerslev [1998b] also point out that for all t as the sampling frequency of the returns increases, the realized variance converges to the integrated variance. As a result, it is possible to construct ex-post realized measures of the underlying integrated variance that are asymptotically free of measurement error by summing sufficiently finely-sampled high-frequency returns [Andersen et al., 2006].

It is widely known that the efficient price processes for assets do not always follow a continuous path. Empirical evidence suggests that asset prices display infrequent large movements, too large to be normal shocks. Therefore, the time series of asset and market prices, given respectively by Equation 5 and Equation 6, will exhibit discontinuities which are referred to as jumps. To incorporate jumps into these equations, a new term $J(t)$ is simply introduced to arrive at

$$dp_i(t) = \mu_i(t)dt + \sigma_i(t)dw_i(t) + dJ_i(t) \tag{13}$$

$$dp_m(t) = \mu_m(t)dt + \sigma_m(t)dw_m(t) + dJ_m(t) \tag{14}$$

where $J(t) - J(s) = \sum_{s<\tau\leq t} \kappa(\tau)$ and $\kappa(\cdot)$ is the magnitude of a jump. Various nonparametric estimators have been developed to measure the integrated variance and variance of the jumps. The most common is the aforementioned realized variance. Andersen et al. [2002] first showed that

$$RV_t(S) = \sum_{j=1}^{S} r^2_{t+\frac{j}{S}} \rightarrow \int_{\tau=t-1}^{t} \sigma^2(\tau)d\tau + \sum_{t-1<\tau\leq t} \kappa^2(\tau), \tag{15}$$

where the first term in Equation 15 is the Integrated Variance given by Equation 8. Thus, the realized variance converges in probability to the integrated variance plus the sum of the squares of the jumps that occur within a day, termed the total variance TV_t[6]. In addition to validating the realized variance measurement, Huang and Tauchen [2005] confirmed that the realized variance is an accurate estimator of the integrated variance even in the presence of jumps. Besides the realized variance, there has been extensive research into jump robust estimators of realized volatility, which seek to better address the presence of jumps in the asset price process.

Barndorff-Nielsen and Shephard [2004] proposed a jump robust estimator, known as the bipower variation, which enables the impact of the presence of volatility to be separated from rare jumps. The bipower variation is expressed as

$$BV_t(S) = \mu_1^{-1} \frac{S}{S-1} \sum_{j=2}^{S} |r_{t,j}|\,|r_{t,j-1}| \tag{16}$$

where μ_1^{-1} is a scale factor usually taken to be $\pi/2$. It can be similarly shown that the bipower variation converges to the integrated variance and that in the absence of jumps the difference between the realized variance and the bipower variation converges to zero. The principle underlying the bipower variation is that if jumps are a rare occurrence, the probability of observing jumps in two consecutive returns approaches zero sufficiently fast as the sampling frequency increases. Consequently, the product of any two consecutive returns will be asymptotically driven by the diffusion component, thereby eliminating the contribution of jumps [Barndorff-Nielsen and Shephard, 2004].

Building upon the bipower variation of Barndorff-Nielsen and Shephard [2004], Mancini [2009] developed a technique for identifying instances where jumps larger than a suitably defined threshold occurred. This allowed for the development of the threshold estimator of the integrated variance known as the truncated variance. The truncated variance is denoted as

$$TV_t(S) = \sum_{j=1}^{S} |r_{t,j}|^2 \times I[|r_{t,j}| \leq cutoff_t], \tag{17}$$

where $I[\cdot]$ is a binary indicator equal to one if true and zero if false. The cutoff value is then defined based upon a latent measure of the previous day's volatility. In order to limit the impact of jumps, the series $cutoff_t$ utilizes the bipower variation, Equation 16, and is defined as

$$cutoff_t = \varphi \times \sqrt{\frac{1}{S} BV_{t-1}}, \tag{18}$$

[6] Since intraday returns are only used to calculate the realized variance, the overnight return is not taken into consideration. Since the majority of the jumps happen overnight as the closing price from previous day is very different from the opening price on the following day because of overnight announcements or macroeconomics policies, this specification excludes those jumps from realized variance calculations.

where φ is the minimum return magnitude relative to BV_t to be classified as a jump[7]. Therefore, three potential realized measures of volatility are available: the realized variance, the bipower variation, and the truncated variance. Throughout the remainder of this paper, the realized variance will be used. However, it is important to note that any of the three measures may be utilized and substituted for the realized variance in these equations.

Using the definition of time-invariant static beta given by Equation 3, the realized beta can be calculated by replacing the static variance and covariance with the realized variance and covariance. The equation for realized beta is subsequently defined as

$$R\beta_{i,t} = \frac{RCOV_{i,m,t}}{RV_{m,t}}, \tag{19}$$

where $RV_{m,t}$ and $RCOV_{i,m,t}$ are, respectively, the realized variance of the market returns and the realized covariance between the returns of asset i and the market over a specific time interval, at some sampling frequency S. In order to forecast both the market and asset components of beta in Section 3, beta is expressed in terms of the realized variance of the market and the realized variance of an asset, which can be accomplished by utilizing the definition of correlation.

Correlation is defined as the ratio of the covariance of an asset's and market's returns over the product of the square root of the variance of the asset's returns and the square root of the variance of the market's returns. Realized correlation is similarly defined using realized measures as

$$R\rho_{i,m,t} = \frac{RCOV_{i,m,t}}{\sqrt{RV_{i,t}} \cdot \sqrt{RV_{m,t}}}. \tag{20}$$

By rearranging the terms of Equation 20 to solve for the realized covariance and substituting this definition into Equation 19, a formula for the realized beta that is a factor of the realized variances of the market and individual equity returns as well as their realized correlation is derived.

$$\begin{aligned} R\beta_{i,t} &= \frac{R\rho_{i,m,t} \cdot \sqrt{RV_{i,t}} \cdot \sqrt{RV_{m,t}}}{RV_{m,t}} \\ &= \frac{R\rho_{i,m,t} \cdot \sqrt{RV_{i,t}}}{\sqrt{RV_{m,t}}} \end{aligned} \tag{21}$$

Finally, based on the underlying assumptions of Equation 5 and Equation 6, this realized beta measure is consistent for the underlying integrated beta

[7] The minimum magnitude in relation to the lagged bipower variation to be considered a jump is typically between 2 and 5.

$$R\beta_{i,t} \rightarrow \beta_{i,t} \equiv \frac{\int_{t-1}^{t} \sigma_{i,m}(\tau)d\tau}{\int_{t-1}^{t} \sigma_m^2(\tau)d\tau}, \tag{22}$$

for all t as S approaches infinity where σ_m^2 is the market variance and $\sigma_{i,m}$ is the individual stock covariance with the market. For more details on realized beta and realized correlation, see Andersen et al. [2006] and Vortelinos [2010], respectively.

It is now possible to calculate the ex-post beta of an individual asset based upon its realized variance, the realized variance of the market, and the realized correlation between the asset's and market's returns. However, it is desired to have information regarding future values of beta, which is accomplished by forecasting its three constituent components. For the purpose of simplification, the observable realized correlation of one period prior to the current is used as the forecasted input to Equation 21. Furthermore, the existence of volatility clustering between time periods is ubiquitous and has been a particularly active area of research. Many of these studies, including Ballie and Bollerslev [1989], Bollerslev [1987], and Engle and Bollerslev [1986] , found that the GARCH(1,1) model provides a good approximation for forecasting conditional volatility. Therefore, the GARCH(1,1) is a logical base measurement for forecasting the future market and asset realized variances.

3 Forecasting Beta

3.1 Characteristics of Volatility

As previously stated, a fundamental characteristic of stock volatility is that it is not directly observable. However, stock volatility has several stylized characteristics. First, there exists volatility clusters. For example: if an asset price changed significantly in the previous period, it is more likely to make a large move in the current period. Second, volatility evolves over time in a continuous manner implying that jumps in volatility are rare [Tsay, 2010]. These volatility characteristics are exemplified by the October 1987 stock market crash where the market value fell twenty percent in one day. Prior to the crash, the standard deviation of returns was approximately one percent per day. The crash resulted in a twenty standard deviation change and would not have been expected to occur in over 4.5 billion years if returns were normally distributed. In four of the five following days, the market moved over four percent. Therefore, volatility appeared to increase after the crash rather than remaining at one percent per day [Reider, 2009]. Volatility not only spikes up during a crisis, it eventually falls to approximately the same pre-crisis level. Additionally, volatility tends to fluctuate within a fixed range and is, therefore, said to be stationary. Finally, volatility is impacted differently by a large price increase as opposed to a large price decrease, referred to as the leverage effect. These properties play a vital role in the development of volatility models, which must seek to capture these stylized characteristics. The foundation for incorporating these factors into volatility models was first proposed by Engle [1982].

3.2 Autoregressive Conditional Heteroskedastic (ARCH) Model

The Autoregressive Conditional Heteroskedastic (ARCH) model is the basis for an entire class of volatility forecasting models and provides a systematic framework for volatility modeling. The name ARCH is derived from its main properties, the first of which is that it is an autoregressive model in squared returns. These models also utilize information from the previous period to model the current period's volatility. In other words, these models form an expectation of the volatility in the next period conditional upon the information currently available. Finally, heteroskedasticity implies that ARCH models allow for time-varying volatility over the estimation and forecast period.

The basic concept of the ARCH model is that the returns of an asset are serially uncorrelated[8], but dependent, and that such dependency can be described by a simple quadratic function of lagged values. Given a series of returns r_t, the shock of an asset is defined as the difference between the return of an asset at time t and the mean function μ_t of the return series. Specifically, an ARCH(m) model assumes that

$$a_t = r_t - \mu_t, \qquad a_t = \sqrt{h_t^{ARCH}} \cdot \epsilon_t, \tag{23}$$

$$h_t^{ARCH} = \alpha_0 + \alpha_1 a_{t-1}^2 + ... + \alpha_m a_{t-m}^2, \tag{24}$$

where ϵ_t is a sequence of independent and identically distributed (iid) random variables with a mean of zero and a variance of one and is often assumed to follow a normal or standardized Student-t distribution [Tsay, 2010]. From the structure of the model, it can be seen that large shocks in a previous period will imply a large conditional variance σ_t^2 for the current period. Therefore, under the ARCH framework, large shocks tend to be followed by large shocks. As a result, the ARCH model successfully captures a number of stylized facts of financial assets such as time-varying volatility and volatility clustering.

Although the ARCH model is able to capture certain characteristics of volatility, it also has some weaknesses. The model assumes that positive and negative shocks have the same effect on volatility because it depends on the square of the previous shocks. In practice, it is well known that the price of a financial asset responds differently to positive and negative shocks. Additionally, ARCH models are conditional exclusively on the previous period returns. Therefore, ARCH models respond slowly to large isolated shocks in the return series and likely overestimate volatility following such shocks [Tsay, 2010].

[8] Serial correlation is the relationship between a given variable and itself over various time intervals. Serial correlations are often found in repeating patterns when the level of a variable affects its future level.

3.3 General Autoregressive Conditional Heteroskedastic (GARCH) Model

Building upon the initial ARCH framework proposed by Engle [1982], Bollerslev [1986] introduced the GARCH model. The GARCH model extended the ARCH framework to allow for both a longer memory and more flexible lag structure. In ARCH processes, the conditional variance is specified as a linear function of past sample variances only, whereas the GARCH process allows lagged conditional variances to enter as well. Like ARCH, the GARCH model is also a weighted average of past squared residuals but also includes declining weights for the residuals that never reach zero. The most widely used GARCH specifications assert that the best predictors of the variance in the next period are the weighted average of the long-run average variance, the variance in the predicted period, and the new information captured in the current period by the most recent squared returns. In this way, the GARCH model exhibits adaptive or learning behavior as it progresses.

Standard GARCH models utilize lagged returns, typically in the form of squared returns, to model current volatility and to extrapolate this information to predict the expected volatility of the next period. Although the model lends itself to directly forecasting one-period in advance, a two-period forecast can be made by repeating the one-period process with the original forecast as the input for the previous day's input parameter. This process can be repeated indefinitely with the forecast moving toward the long-run average variance. This distant horizon forecast is the same for all periods and is merely the unconditional variance. Thus, the GARCH model is mean reverting with a constant unconditional variance.

For a log return series r_t, let $a_t = r_t - \mu_t$ be the innovation at time t. Then, the GARCH(m,s) model for a_t is specified as

$$a_t = \sqrt{h_t^{GARCH}} \cdot \epsilon_t, \qquad h_t^{GARCH} = \alpha_0 + \sum_{i=1}^{m} \alpha_i a_{t-i}^2 + \sum_{j=1}^{s} \theta_j h_{t-j}^{GARCH} \tag{25}$$

where again ϵ_t is a sequence of iid random variables with mean zero and variance of one. Moreover, the model implies that $\alpha_i = 0$ for all $i > m$ and $\theta_j = 0$ for all $j > s$. As with the ARCH model, ϵ_t is also assumed to follow a normal or standardized Student-t distribution [Tsay, 2010].

There has been an enormous amount of literature concerning the GARCH model and its ability to describe and forecast the volatility of asset returns. Despite encountering the same weaknesses as the ARCH model, such as a lagged response to isolated shocks and equal treatment of positive and negative shocks, Hansen and Lunde [2005] showed that the most basic GARCH model, GARCH(1,1), performed at least as well as 330 more complicated models. Therefore the GARCH(1,1) model given by

$$h_t^{GARCH} = \alpha_0 + \alpha_1 \cdot a_{t-1}^2 + \theta_1 \cdot h_{t-1}^{GARCH} \tag{26}$$

will form the basis for the proposed model in this paper. Equation 26 illustrates the strengths and weaknesses of the GARCH model. A large a_{t-1}^2 or h_{t-1}^{GARCH} will give rise to a large estimate of the conditional variance. This will again generate the well-known behavior of volatility clustering in financial time series. It also can be illustrated that, similar to the ARCH model, the tail distribution of the GARCH(1,1) process is heavier than a normal distribution. However, the model provides a simple parametric function that can be used to describe the evolution of volatility for financial time series [Tsay, 2010].

Several studies have evaluated GARCH models using intraday returns or their residuals (Andersen and Bollerslev [1997a], Andersen and Bollerslev [1997b], and Andersen and Bollerslev [1998a]) and support its use as an accurate estimator of volatility. Jorian [1995] and Figlewski [1997], however, questioned the usefulness of the GARCH model. They argued that such models were unable to explain much of the variability in squared returns when evaluated out-of-sample, despite the fact that the GARCH models had good in-sample fit, and concluded that GARCH models were of little value. By using the realized variance, which is a more accurate measure of volatility than squared returns, Andersen and Bollerslev [1998b] illustrated that the standard GARCH model performs very well. The apparent poor performance of the models can be attributed to the fact that squared returns are an extremely noisy proxy for the conditional variance.

While the GARCH model itself remains sound, its reliance on squared returns implies that standard GARCH models are not as well suited for periods characterized by rapid fluctuations in volatility. This is attributed to the fact that the GARCH model is slow to respond to sudden changes in the level of volatility and will take many periods for the conditional variance implied by the GARCH model to reach this new level as illustrated by Andersen et al. [2003]. The success in the ability of realized measures to better characterize the dynamic properties of volatility, this has prompted the development of GARCH type models that incorporate a realized measure as a parameter for forecasting conditional variance.

3.4 GARCH Model with Exogenous Variable (GARCHX)

While most ARCH and GARCH type models have been univariate, relating the volatility of a time series only to information contained in its own past history, additions of supplemental economic variables as covariates have been made to enhance the modeling of economic and financial time series volatility. Many of these works directly build upon the success of the GARCH(1,1) model and simply add a covariate as follows:

$$h_t^{GARCHX} = \alpha_0 + \alpha_1 \cdot y_{t-1}^2 + \theta_1 \cdot h_{t-1}^{GARCHX} + \gamma_1 \cdot x_{t-1}, \tag{27}$$

where y_t is a degraded time-series, h_t^{GX} is the variance conditional on the information available at time $t-1$, and x_t is the covariate exogenous to the model.

Given the relative ease required to incorporate the covariate into the model, a variety of different variables have been added in several studies in an attempt to better explain volatility. Glosten et al. [1993], Brenner et al. [1996], Gray [1996], Engle and Patton [2001], and Staikouras [2006] introduced interest rate levels as a covariate. Similarly, forward spot spreads and interest rate spreads were used respectively by Hodrick [1989] and Hagiwara and Herce [1999] as the covariates to the model. Moreover, the GARCHX model was also considered by Han and Park [2008] who employed the yield spread between bonds as the covariate. For current studies utilizing the GARCHX framework see Apergis and Rezitis [2011], Han and Kristensen [2012], Mulyadi and Anwar [2012], and Ben Sita and Abosedra [2012].

Most of the covariates used in GARCHX models are economic variables, but various realized volatility measures constructed from high-frequency data have recently been adopted with the rapid development seen in the field of realized volatility. As previously stated, realized measures of volatility such as the realized variance provide significantly stronger and more informative estimates concerning the current level of volatility than squared returns. Consequently, realized measures are very attractive for modeling and forecasting future volatility. The multiplicative error model (MEM) proposed by Engle [2002] was the first to use realized variance as the covariate to forecast conditional variance using the GARCHX model framework. Barndorff-Neilson and Shephard [2007] included both the realized variance and the bipower variation in their specification (See also Engle and Gallo [2006], Cipollini et al. [2009], Shephard and Sheppard [2010], and Hansen et al. [2012]). In particular, the high-frequency based volatility (HEAVY) model developed by Shephard and Sheppard [2010] and the Realized GARCH model of Hansen et al. [2012] included realized measures based on high-frequency data into their specification for the conditional variance.

In the end, the addition of the realized variance measure to Equation 26 is a simple method to enhance the conventional GARCH model. The generalized GARCHX(m,s,n) model is given below, where m, s, and n indicate the degree of lagged terms for the squared residuals, conditional variance, and realized measure, respectively. This is followed by the base GARCHX(1,1,1), where the realized variance, RV_t, is treated as an explanatory variable.

$$a_t = \sqrt{h_t^{GARCHX}} \cdot \epsilon_t \tag{28}$$

$$h_t^{GARCHX} = \alpha_0 + \sum_{i=1}^{m} \alpha_i a_{t-i}^2 + \sum_{j=1}^{s} \theta_j h_{t-j}^{GARCHX} + \sum_{k=1}^{n} \gamma_k \cdot RV_{t-k} \tag{29}$$

$$h_t^{GARCHX} = \alpha_0 + \alpha_1 \cdot a_{t-1}^2 + \theta_1 \cdot h_{t-1}^{GARCHX} + \gamma_1 \cdot RV_{t-1} \tag{30}$$

The relative contributions of the parameters within the model have been examined by many authors. Engle [2002] found that the coefficient for the squared return, α_1, was insignificant once the realized variance was included in the model. Barndorff-Neilson and Shephard [2007] arrived at the same conclusion if the bipower variation was used as the realized measure. Using ten years of U.S. and U.K. daily price data, Hwang and Satchell [2005] also illustrated that the return volatility of an individual stock can be better specified with GARCHX models than with GARCH models. There is strong evidence as a result that $\gamma_1 \neq 0$, implying that the realized measure is a useful predictor for the conditional variance. Furthermore, it is well known that the asymptotic properties of the GARCH(1,1) model are not fully appropriate in explaining the stylized facts in financial time series even though it is extensively used in practice. Both Han [2011] and Han and Kristensen [2012] illustrate that the stylized facts in financial time series can be significantly better, if not fully explained, in the framework of the GARCHX model.

3.5 GARCH and GARCHX Forecasted Beta

Recall from Equation 19 that the realized beta is given as the ratio of the realized covariance over the realized variance. Using GARCH and GARCHX models, the conditional variance for both the market's and asset's returns can be forecasted. Assuming that the conditional variances calculated with the model are accurate forecasts of the realized variance, we substitute these values for the realized variances in Equation 21 to arrive at the following forecasting estimate for beta:

$$\hat{\beta}_{i,t+1} = \frac{\hat{\rho}_{i,m,t+1} \cdot \sqrt{\hat{h}_{i,t+1}}}{\sqrt{\hat{h}_{m,t+1}}}, \text{ where} \tag{31}$$

$$\hat{h}_{m,t+1} = h^{GARCHX}_{m,t+1} = \alpha_{m,0} + \alpha_{m,1} \cdot a^2_{m,t} + \theta_{m,1} \cdot h^{GARCHX}_{m,t} + \gamma_{m,1} \cdot RV_{m,t}, \tag{32}$$

$$\hat{h}_{i,t+1} = h^{GARCHX}_{i,t+1} = \alpha_{i,0} + \alpha_{i,1} \cdot a^2_{i,t} + \theta_{i,1} \cdot h^{GARCHX}_{i,t} + \gamma_{i,1} \cdot RV_{i,t}, \tag{33}$$

$\hat{\rho}_{i,m,t}$ is the realized correlation between the market's and asset's returns at time t and is taken to be $\rho_{i,m,t}$ via the no-change model, $\hat{h}_{i,t}$ is the conditional variance of the asset, and $\hat{h}_{m,t}$ is the conditional variance of the market. The realized correlation between the market's and asset's returns is assumed to remain virtually unchanged from one period to the next. Given that the realized correlation is calculated directly from the realized covariance and the realized variances of an asset and the market, the realized correlation is expected to be highly sensitive to isolated shocks in the realized variance of the market or asset. As a consequence, a measurable degree of noise is anticipated to be introduced into the forecasted betas and is a noted limitation of the current model. Vortelinos [2010] explores in depth the properties of realized correlation and provides several alternative frameworks, including a heterogeneous autoregressive model, for examining the realized correlation that could be implemented in conjunction with the current model of beta in further research.

3.6 Beta Specification

As noted earlier, the realized beta is a measure of the factored response of an asset's returns to the returns of the market. However, this measure alone should not simply be used as the market beta because it is subject to extreme fluctuation. Due to the daily measurement period that is used within this study, there is reasonable variability when approximating the constituent components of the forecasted conditional beta in addition to the computed realized beta. For example, a large negative realized beta can be obtained for a trading day on which negative news for a specific equity is released despite positive market gains. Therefore, if the realized beta obtained on day t is simply used as the market beta, the results will be misleading. The useful information and significant noise both contained in the realized beta measurement must be reconciled. To accomplish this, a filter method similar to the GARCH model is employed. In other words, the realized beta is treated as an innovation on the latent underlying beta, which is specified as the market beta. The market beta is given by

$$\beta_{i,t+1} = \varphi_0 + \sum_{j=t-4}^{t-1} \varphi_j \beta_{i,j} + \varphi_t \hat{\beta}_{i,t} \tag{34}$$

where $\beta_{i,j}$ is the market beta of security i at time t, $\hat{\beta}_{i,t}$ is the calculated (or model forecasted) realized beta for asset i at time t, and φ_0 is a constant. For the sake of simplicity, an equal weighting is assigned to each of the coefficients φ_i and φ_t and the constant φ_0 is taken to be equal to zero. Therefore, the above equation reduces to the following and is a simple moving average of the previous market betas in addition to the realized beta of the current period.

$$\beta_{i,t+1} = \frac{1}{5}\left(\sum_{j=t-4}^{t-1} \beta_{i,j} + \hat{\beta}_{i,t}\right) \tag{35}$$

The choice for including five parameters is straightforward and follows from the desire to model the previous trading week. It is also recognized that each lagged term likely has a different weight in the ideal calculation, and that it is quite possible to calculate the optimal regression coefficients using a maximum likelihood estimator. However, the equal weighting approach was chosen for its simplicity while still allowing the realized beta to be an innovation on the market beta and limiting the exposure to large sudden changes in volatility.

4 Benchmark ARMA Model

In order to provide a benchmark for the betas calculated by the aforementioned framework, an ARMA model was used to forecast the realized betas and provide a comparison to the GARCH and GARCHX forecasts. As previously mentioned, there are several stylized characteristics of financial time series, namely the clustering of volatility and mean reversion, that must be considered when fitting such data to a model. The ARMA framework is a natural fit for such a time series and addresses precisely those two characteristics. Furthermore, given that the risk profile facing an individual asset is not expected to dramatically shift from one day to the next, an ARMA model is well suited to forecast the beta of an asset one period ahead. The ARMA model is a natural choice to compare the GARCH and GARCHX forecasts as the GARCH model can actually be written as an ARMA model itself. Descriptions of ARMA models can be found across financial volatility literature; however, this study will employ the definition given by Tsay [2010], which provides a description of a stochastic process in terms of two polynomials, one for the auto-regression (AR) and the second for the moving-averages (MA).

Given a time-series X_t with a statistically significant first-order autocorrelation, X_{t-1} may be useful in predicting X_t. A simple AR model that makes use of such predictive ability is as follows:

$$X_t = \varphi_0 + \varphi_1 \cdot X_{t-1} + \varepsilon_t, \tag{36}$$

where ε_t is assumed to be a serially independent series with mean zero and variance σ_ε^2. This model is similar in form to a simple linear regression model in which X_t is the dependent variable and X_{t-1} the explanatory variable. Equation 36 is referred to as an AR model of order one or simply an AR(1) model [Tsay, 2010]. Furthermore, the AR(1) model implies that

$$E(X_t|X_{t-1}) = \varphi_0 + \varphi_1 \cdot X_{t-1} \quad \text{and} \quad Var(X_t|X_{t-1}) = Var(\varepsilon_t) = \sigma_\varepsilon^2. \tag{37}$$

Given the past data point X_{t-1}, the current value is centered around $\varphi_0 + \varphi_1 \cdot X_{t-1}$ with a standard deviation of σ_ϵ. In other words, this indicates that, conditional on X_{t-1}, the value at X_t is not correlated with X_{t-i} for any i greater than one [Tsay, 2010].

Of course, there are circumstances in which X_{t-1} alone is not satisfactory to determine the conditional expectation of X_t. Therefore, it is necessary to provide a more flexible model. A straightforward generalization of the AR(1) model is the AR(p) model:

$$X_t = \varphi_0 + \sum_{i=1}^{p} \varphi_i \cdot X_{t-i} + \varepsilon_t, \tag{38}$$

where $\{\varphi_0 ... \varphi_p\}$ are parameters to the model and p is a nonnegative integer. The AR(p) is of the same form as a multiple liner regression model with lagged values serving as its explanatory variables [Tsay, 2010].

There is no particular reason, besides simplicity, to assume that the order of an AR model is finite. In fact, such a model is easily illustrated by allowing the integer p in Equation 38 to be replaced by infinity. However, such a model is unrealistic in practice because it is not possible to estimate infinitely many parameters. The model can be made more practical by assuming that the coefficients $\{\varphi_0, \varphi_1, ..., \varphi_p\}$ are determined by a finite number of parameters [Tsay, 2010].

Another class of simple models that are useful in financial time series analysis are moving-average (MA) models. The simplest example of these models says that, except for the constant term, X_t is a weighted average of shocks a_t and a_{t-1}. Therefore, the model is called a MA model of order one or MA(1) model. The general form of a MA(1) model is

$$X_t = \varphi_0 - \theta_1 a_{t-1} + a_t, \tag{39}$$

where φ_0 is a constant and a_t is also a serially independent series [Tsay, 2010]. Taking into account the expectation and variance of the model, it follows from Equation 39 that

$$E(X_t) = \varphi_0 \quad \text{and} \quad Var(X_t) = \sigma_a^2 + \theta_1^2 \sigma_a^2, \tag{40}$$

given that a_t and a_{t-1} are uncorrelated and both $E(X_t)$ and $Var(X_t)$ are time invariant [Tsay, 2010].

As was the case for the AR model, the MA(1) may not be adequate to fully describe the time series, but it can be similarly expanded to include q lagged periods as follows:

$$X_t = \varphi_0 + \sum_{i=1}^{q} \theta_i \cdot a_{t-i} + a_t, \tag{41}$$

where $\{\theta_1 ... \theta_q\}$ are the model parameters, φ_0 is the expectation of X_t, and $\{a_t, a_{t-1}, ..., a_{t-q}\}$ are noise terms [Tsay, 2010].

In some cases, the AR and MA models discussed above become cumbersome and often require high orders of lagged parameters to accurately describe a dynamic time series. Therefore, the two models can be combined in a compact form, known as an ARMA model, so that the number of parameters used remains small [Tsay, 2010]. The simplest ARMA(1,1) model is trivial to derive as it is a simple juxtaposition of the AR(1) and MA(1) models above. The expanded $ARMA(p, q)$ equation is denoted as follows:

$$X_t = \varphi_0 + \varepsilon_t + \sum_{i=1}^{p} \varphi_i \cdot X_{t-i} + \sum_{i=1}^{q} \theta_i \cdot \varepsilon_{t-i}, \tag{42}$$

where ε_t is a white noise series[9] and p and q are nonnegative integers. For more information on the

[9] A white noise series is a random signal or process characterized as a sequence of serially uncorrelated random variables with

characteristics of ARMA models, the determination of their orders, and illustrative examples see Tsay [2010].

5 Methodology and Evaluation Framework

5.1 Data

In practice, the discreteness of actual securities prices can cause continuous time models to be poor approximations when sampled at very high frequencies. For example, discrete prices are generally only available at unevenly spaced points in time. High-frequency returns must then be calculated based upon a form of interpolated prices. The theory of quadratic variation implies the necessity of sampling at very high frequencies in order to model the ideal of continuously-observed prices. On the other hand, the reality of trading frictions suggests not sampling too frequently. Within this study, a sampling frequency of 288 times per day is used. This data is then converted into five-minute returns. This frequency is seemingly high enough such that the realized measures are free from measurement error, yet low enough such that biases from trading frictions are not of great concern.

This paper utilizes five-minute geometric returns data for eight stocks in the S&P 500. These stocks were chosen in order to represent the various major industry sectors—technology, finance, and food/agricultural, automotive, and retail. Each trading day contains returns from 9:40 AM to 4:00 PM, leading to 76 intra-day observations. Most stock data runs from 1997 to 2010, although data for certain securities are available for only a subset of that time period.

Ticker	Industry	Start Date	End Date	Number of Days
AAPL	Tech	4/16/97	12/30/10	3408
EXC	Energy	10/23/2000	12/30/10	2524
F	Automotive	4/9/97	12/30/10	3405
IBM	Tech	4/9/97	12/30/10	3413
JPM	Finance	4/9/97	12/30/10	3411
KFT	Food/Agricultural	6/13/01	12/30/10	1949
PG	Manufacturing	4/9/97	12/30/10	3412
WMT	Retail	4/9/97	12/30/10	3409

Table 1: Stocks used in the analysis, with start and end dates and the number of days included in the data.

5.2 Calculating Beta Predictions

In order for the GARCH and GARCHX forecasts to be calculated, realized betas were computed by calculating the realized variances and correlations of Equation 21. These realized betas were computed for the eight chosen equities over their respective time periods at the selected five-minute optimal sampling frequency.

mean zero and finite variance.

These realized measures were requisite for the subsequent steps in the analysis. With these calculations, the GARCH and GARCHX frameworks, given by Equation 26 and Equation 31, respectively, were used to estimate the GARCH and GARCHX forecast models for the conditional variance of the asset and market. These forecasts were then combined with the previously computed realized correlation to arrive at a calculation of the daily realized beta time series.

Once the realized betas were computed, the market beta was calculated using the framework in Equation 35. First, the market beta was calculated using the realized beta of the asset in question as the innovation on the current period's market beta. Then similar series were computed utilizing the conditional betas forecasted by the GARCH and GARCHX models to arrive at the one-day forward beta predictions.

In order to provide a benchmark to compare the proposed GARCHX model, an ARMA model was calculated using the daily market betas as its parameter for the model. As noted earlier, the realized beta time series had high, positive, first-order autocorrelations and partial autocorrelations, promoting the use of an ARMA(1,1) model. The calculated ARMA process is represented by the following:

$$\hat{\beta}_t^{ARMA} = \alpha + \theta \cdot \beta_{t-1} + \varphi \cdot \varepsilon_{t-1} + \varepsilon_t \tag{43}$$

$$\varepsilon_t = \beta_t - \hat{\beta}_t^{ARMA} \tag{44}$$

where β_t is the market beta at time t and ε_t is the error of the forecast at time t. Table 4 denotes the coefficients that were calculated for each of the eight stocks in the data set (see Appendix).

5.3 Root Mean Squared Error

There is a wide spectrum of measures designed to evaluate ex-post forecasts. These include mean squared error (MSE), root mean squared error (RMSE), mean error (ME), mean absolute error (MAE), mean squared percent error (MSPE), and root mean squared percent error (RMSPE). The most common of these, MSE and MSPE, are defined as follows:

$$MSE = \frac{1}{n}\sum_{t=1}^{n} e_t^2 \tag{45}$$

$$MSPE = \frac{1}{n}\sum_{t=1}^{n} p_t^2 \tag{46}$$

where e is the forecast error term defined as the difference between the actual value and the forecasted value and p is the percentage form of the forecast error.

Examination of two measures of the forecast accuracy are of interest. The first is the degree to which the models are able to directly forecast the next day's realized beta. The second is the accuracy in forecasting the market beta as specified by Equation 35. The MSE measure of beta forecast accuracy from the predicted beta is defined as

$$MSE_{\beta} = \frac{1}{n}\sum_{t=1}^{n}(R\beta_{i,t} - \hat{\beta}_{i,t})^2 \tag{47}$$

where n is the number of days in the estimation interval, $R\beta_{i,t}$ is the underlying realized beta of asset i at time t, and $\hat{\beta}_{i,t}$ is the forecasted beta of asset i at time t. The MSPE is similarly defined as

$$MSPE_{\beta} = \frac{1}{n}\sum_{t=1}^{n}\left(\frac{R\beta_{i,t} - \hat{\beta}_{i,t}}{R\beta_{i,t}}\right)^2. \tag{48}$$

The MSE and MSPE can be converted into the RMSE and RMSPE, respectively, to standardize the results and to convert the MSE and MSPE into more easily interpreted standard deviation units, known as the Root Mean Squared Error (RMSE) and Root Mean Percent Squared Error (RMPSE). MSE, MSPE, RMSE, and RMSPE were used as measures of accuracy and comparison for the predictive betas of the GARCH, GARCHX, and ARMA models against the true underlying beta. However, for simplicity, only the MSE's of the forecasts are presented.

6 Results

Table 2 summarizes the variances of the estimated realized betas (shown graphically in Figure 5) for the eight equities using the five-minute sampling frequency. If beta were constant with respect to time, the covariance between an asset's returns and the market's returns would not be expected to statistically deviate over time; however, the average variance of realized beta across the equities is 0.1396. Furthermore each of

Equity	Variance	95% Confidence Interval
AAPL	0.3441	0.3283 - 0.3610
EXC	0.0912	0.0864 - 0.0965
F	0.2084	0.1988 - 0.2186
IBM	0.0710	0.0678 - 0.0745
JPM	0.1371	0.1308 - 0.1438
KFT	0.0643	0.0604 - 0.0685
PG	0.0907	0.0865 - 0.0951
WMT	0.1102	0.1051 - 0.1156
Average	0.1396	—

Table 2: Analysis of the Variance of Realized Beta

the equities has a non-zero variance which is statistically significant at the 95% confidence level. Therefore, the hypothesis that realized beta has a variance of zero and is time-invariant can be rejected. Analysis of the values in Table 2 also reveals interesting characteristics concerning the variance of betas for stocks in different industries. The betas with the higher variances belong to firms in historically more volatile industries. For instance, the beta with the highest variance belongs to Apple Inc. (AAPL) within the highly volatile technology sector. However, equities such as Exxon Corp. (EXC) and Kraft Foods Inc. (KFT) have considerably lower variance in their betas. This is likely due to the consistency of demand in the energy and food industries. Therefore, not only does beta fluctuate over time but the degree to which it does so appears to vary across stocks and industry sectors[10].

The proposed model relies on the forecasts of the realized variance given in Figure 1 in order to construct a predictive model for the conditional beta. Table 3 presents the GARCH and GARCHX model coefficients[11] and is informative concerning the statistical significance of the parameters and the performance of the two models. Additionally, The forecasted conditional variances pertaining to these GARCH and GARCHX coefficients are presented in Figure 3.

	GARCH Coefficients				GARCHX-RV Coefficients				
Equity	α_0	α_1	θ_1	R^2	α_0	α_1	θ_1	γ_1	R^2
AAPL	0.0306	0.0577	0.9392	0.2994	0.1967	0.0277	0.7015	0.2008	0.4153
	(0.4485)	(0.1430)	(0.0000)		(0.1811)	(0.2875)	(0.0000)	(0.0000)	
EXC	0.0285	0.1061	0.8833	0.3952	0.0585	0.0927	0.6533	0.1910	0.4694
	(0.3937)	(0.0297)	(0.0000)		(0.2780)	(0.0396)	(0.0000)	(0.0000)	
F	0.0571	0.0730	0.9155	0.2363	0.1566	0.0215	0.6862	0.2101	0.3212
	(0.4464)	(0.3121)	(0.0000)		(0.3481)	(0.4393)	(0.0000)	(0.0000)	
IBM	0.0075	0.0698	0.9285	0.3138	-0.0026	0.0210	0.7838	0.1739	0.4365
	(0.4634)	(0.0327)	(0.0000)		(0.5140)	(0.2700)	(0.0000)	(0.0000)	
JPM	0.0152	0.0976	0.9048	0.3416	-0.0059	0.0289	0.7412	0.2252	0.4639
	(0.4670)	(0.0820)	(0.0000)		(0.5141)	(0.3241)	(0.0000)	(0.0000)	
KFT	0.0247	0.0522	0.9295	0.2566	0.0470	0.0078	0.6843	0.2325	0.3755
	(0.3727)	(0.0704)	(0.0000)		(0.2585)	(0.4049)	(0.0000)	(0.0000)	
PG	0.0053	0.0514	0.9455	0.2191	0.0160	0.0260	0.5094	0.3481	0.2946
	(0.4728)	(0.1143)	(0.0000)		(0.4141)	(0.2607)	(0.0000)	(0.0000)	
WMT	0.0066	0.0620	0.9364	0.2716	0.0010	0.0232	0.8257	0.1176	0.3369
	(0.4712)	(0.0713)	(0.0000)		(0.4954)	(0.2833)	(0.0000)	(0.0000)	
Average	0.0219	0.0712	0.9228	0.2917	0.0584	0.0311	0.6982	0.2124	0.3892

The GARCH model and GARCHX model equations for an asset are defined respectively as follows:
$h^{GARCH}_{i,t+1} = \alpha_{i,0} + \alpha_{i,1} \cdot a^2_{i,t} + \theta_{i,1} \cdot h^{GARCH}_{i,t}$ and $h^{GARCHX}_{i,t+1} = \alpha_{i,0} + \alpha_{i,1} \cdot a^2_{i,t} + \theta_{i,1} \cdot h^{GARCHX}_{i,t} + \gamma_{i,1} \cdot RV_{i,t}$.

Table 3: Coefficients of the GARCH and GARCHX Equations

It has been mentioned several times throughout this paper that daily squared returns, corresponding to α_1, are poor predictors for the conditional variance. The results presented in Table 3 reinforce this statement. The average value for the GARCH and GARCHX estimations of α_1 across the eight equities examined are 0.0712 and 0.0311, respectively. The regression thus indicates that the squared returns are able to explain

[10]Given that only one stock from each sector is utilized in this analysis, the findings from this small subset of equities may not apply to other equities. Therefore, stocks in similar sectors as those employed in this analysis may exhibit different characteristics. Further research regarding the variance of equity realized betas is needed to form a substantive conclusion.

[11]The coefficients were estimated using maximum likelihood. The likelihood equation used to calculate both the GARCH and GARCHX parameters is presented in the Appendix.

little of the conditional variance. The p-values of each coefficient, given in parentheses under each respective parameter in Table 3, detail the statistical significance of each of the regressors. For the GARCH estimates, only two of the equities have α_1 coefficients that are significant at the five percent level and none significant at the one percent level. Even more so, under the GARCHX model, only one of the eight equities has a coefficient α_1 that is statistically significant at the five percent level. Therefore, it is evident that the shocks from daily squared returns offer little utility in forecasting volatility.

Table 3 also illustrates the statistical significance of the coefficient for the lagged conditional variance θ_1 (also called the GARCH coefficient) as well as the coefficient for the realized measure γ_1 included in the GARCHX model. For the GARCH model, the lagged conditional variance explains nearly all of the information concerning the future conditional variance, as expected. With an average θ_1 value of 0.9228, this illustrates the well known volatility clustering characteristic in financial time series. Additionally, all eight of the θ_1 coefficients are highly significant, as measured by their p-values of effectively zero. Moreover, the coefficient results agree with those given by Hansen et al. [2010] who stated that parameter values of $\alpha_0 = 0$, $\alpha_1 = 0.05$, and $\theta_1 = 0.95$ are in line with typical estimates for the GARCH(1,1) model, when estimated with daily returns.

The inclusion of the realized variance as a covariate has a considerable impact on the weights of the lagged conditional variance coefficients. Hansen et al. [2010] also indicated that typical estimates for the GARCHX coefficients are $\alpha_0 = 0$, $\alpha_1 = 0$, $\theta_1 = 0.5$, and $\gamma_1 = 0.5$. The coefficients calculated for the GARCHX model in this analysis are noticeably different from these specified parameters. Across the examined equities, the average θ_1 coefficient for the GARCHX model drops to 0.6982, in comparison to the GARCH model average of 0.9228 , with each individual stock's θ_1 coefficient remaining highly significant. For all stocks, the null hypothesis that the coefficient γ_1 is equal to zero is rejected at the one percent significance level, indicating that the realized measure provides a measurable degree of information concerning the future values of the conditional variance. The average γ_1 coefficient for the realized variance in the GARCHX model is 0.2124, which is significantly lower than the 0.5 suggested by Hansen et al. [2010]. One possible explanation may have to do with the stocks used in this analysis. All of the stocks, with the exception of AAPL, have average betas less than one over the time period (see Table 6 in the Appendix). A beta below one indicates that the equity is less volatile than the market and exhibits a diminished response to market fluctuations. Consequently, the security is less prone to sudden changes in volatility making the current value of the conditional variance rely more heavily on the previous day's variance and less on the realized measure of volatility. Table 3 also presents the R^2 of the GARCH and GARCHX forecasts indicating the overall performance of the two models. With an increase in R^2 from an average of 0.291 to 0.389, the GARCHX model clearly represents a better fit for the true conditional variance. Finally, the sum of the coefficients across each of the regressions is approximately equal to one, indicating the persistence of the regression coefficients and supporting the accuracy in the resulting beta predictions.

Table 4 summarizes the results of the first-order ACF and PACF for the entire sample interval. Positive

Equity	First-Order Autocorrelation	ACF 95% CI	First-Order Partial Autocorrelation	PACF 95% CI
AAPL	0.4232	±0.0664	0.4233	±0.0343
EXC	0.3572	±0.0662	0.3575	±0.0399
F	0.3867	±0.0615	0.3867	±0.0343
IBM	0.3187	±0.0536	0.3187	±0.0343
JPM	0.4603	±0.0672	0.4606	±0.0343
KFT	0.2054	±0.0598	0.2056	±0.0454
P	0.5483	±0.0846	0.5484	±0.0343
WMT	0.4716	±0.0722	0.4720	±0.0343
Average	0.3964	—	0.3966	—

Table 4: First-Order Autocorrelations and Partial Autocorrelations of Realized Beta

first-order autocorrelations suggest the persistence in the realized beta. The average first-order autocorrelation is 0.396, which substantiates the predictability in the time-varying beta. Similarly, the average partial first-order autocorrelation is 0.397, which indicates that realized beta is linearly related to its first lagged value. Plots of the ACF and PACF (Figure 7) indicate that the realized beta time-series has significant autocorrelation for up to five or even ten lagged terms. Therefore, an ARMA model with at least order five is likely optimal for fitting the realized beta series. However, an ARMA(1,1) model is used for simplicity. Due to the variance of the realized beta from one period to the next, the application of an ARMA(1,1) model directly on this time series yields nothing more than a series of noise with no meaningful information. The ARMA(1,1) is therefore applied to the market beta given by Equation 35. The autoregressive nature of the market beta is similar to the higher order ARMA process fit to the lagged realized beta and, therefore, more accurately models the optimal ARMA process. When determining the joint order of an ARMA model, Tsay [2010] state that the ACF and PACF are not well suited to calculate the optimal orders. They suggest an alternative approach using the extended autocorrelation function (EACF), which is significantly more intensive. Therefore, despite their short-comings, the more basic ACF and PACF were employed to specify the order of an ARMA process.

Table 5 presents the MSE of the three forecast models examined in this paper each of which are evaluated for their ability to forecast two different values. The first of these is the daily realized beta given by Equation 21. Because both the GARCH and GARCHX models are proven to be accurate forecasters of volatility, these models are expected to outperform the ARMA model when forecasting the next period's realized beta. Moreover, because the GARCHX model incorporates the previous period's realized volatility, it is similarly expected to outperform the GARCH model, which is indeed what the results illustrate.
Utilizing the GARCH and GARCHX model, the MSE is reduced by an average of 11.03% and 11.39%, respectively, when compared to the ARMA model. Similarly, the GARCHX results in a limited average

	Realized Beta			Market Beta		
Equity	GARCH	GARCHXRV	ARMA	GARCH	GARCHXRV	ARMA
AAPL	0.2898	0.2866	0.3159	0.0374	0.0338	0.0186
EXC	0.0835	0.0832	0.0927	0.0103	0.0096	0.0056
F	0.1793	0.1782	0.1961	0.0217	0.0199	0.0118
IBM	0.0740	0.0738	0.0858	0.0095	0.0087	0.0047
JPM	0.1260	0.1263	0.1428	0.0162	0.0144	0.0080
KFT	0.0625	0.0628	0.0672	0.0068	0.0064	0.0042
PG	0.0597	0.0593	0.0700	0.0082	0.0075	0.0040
WMT	0.0879	0.0870	0.1025	0.0118	0.0113	0.0058
Average	0.1203	0.1197	0.1341	0.0152	0.0140	0.0078

Table 5: Mean Squared Error for GARCH, GARCHX, and ARMA Forecasts of Daily Realized Beta and Market Beta

reduction in MSE of 0.42% when compared with the GARCH model. While a reduction in MSE is observed for six of the eight individual equities when moving from the GARCH model to the GARCHX model, the MSE of two of the stocks, KFT and IBM, increased. However, both the GARCH and GARCHX model still resulted in a decrease in MSE for these two stocks when compared to the ARMA model.

The second value against which the forecasted betas are compared is the market beta defined by Equation 35, which utilizes the calculated realized beta (or forecasted realized beta) as an innovation on the true underlying market beta. Overall, the GARCHX model forecasts perform systematically better than the GARCH model when predicting future values of the specified market beta. Table 5, however, suggests that the ARMA model in fact outperforms both the GARCH and GARCHX models. This is, however, not surprising, given that the ARMA process uses the market beta as its time series to forecast the next period's beta. The GARCH and GARCHX models both aim to model the daily realized variances of asset and market returns (given in Figure 1), which are known to vary extensively from one period to the next. Furthermore, the simplification of the model to use the previous day's realized correlation as a proxy for a forecasted conditional correlation introduces a measurable degree of variability to the GARCH and GARCHX predicted betas. However, the ARMA process is based upon the market beta, which by its nature is a smoother time series. Because the ARMA model is fit to the market beta, its coefficients are optimized for that series versus the GARCH and GARCHX models where the coefficients are derived from the actual realized beta. Therefore, less variability is expected in the forecasted ARMA beta, which over the large data set leads to lower MSE. Equally weighting the terms in the specification of the market beta is likely to be a significant factor in the relatively better performance of the ARMA model for this specific scenario. If optimized, rather than equally weighted, coefficients for the beta specification were used for each equity, the GARCH and GARCHX models are likely to produce results comparable to the ARMA model. Moreover, the specification for the market beta is a conceptual measure used within this analysis and is not a robust measure to the same degree as realized beta, which was shown to be most accurately forecasted using the proposed GARCHX framework.

7 Conclusion

Extensive research has centered on exploring alternatives to the constant beta traditionally used in the CAPM. There is significant data supporting the fact that beta can indeed be better represented as a continuous-time process. This paper generates further evidence that beta is indeed time-varying and provides a model for explaining its underlying nature. The development of realized measures of volatility enables the use of high-frequency data to extract more information concerning intraday equity returns and more appropriately model an asset's volatility. The GARCH model provides a simple parametric equation for modeling and forecasting the conditional variance of an asset. Realized measures of volatility can then be incorporated into the GARCH model as a covariate, forming a GRCHX model, to more accurately measure and forecast volatility. The GARCHX framework yields predictive betas with 6.55% to 15.29% less mean squared error than the benchmark ARMA model used in this analysis.

It is important to note a few key limitations to the proposed GARCHX model. The first, as previously mentioned, is the assumption of equal weighting on the lagged terms in the specification for the market beta. It is likely that, rather than being equal, the coefficients for the regressors vary considerably and, moreover, have different values for each stock. By assuming equal weighting, the information from the GARCHX model, namely from the inclusion of the realized measure of variance, is not appropriately incorporated into the value for the market beta. This limits the ability of the market beta to adjust to abrupt changes in the level of volatility. Optimizing the values for the market beta coefficients is likely to significantly improve the prediction accuracy of the GARCHX model in comparison to both the GARCH and ARMA processes.

The use of maximum likelihood and ordinary least squares analyses in calculating and determining the coefficients for the GARCH and GARCHX frameworks are also potential weaknesses. The maximum likelihood function used in this analysis assumes that the GARCH process follows a standard normal distribution[12]. However, extensive research has shown that this is not the case and that the tail distribution of a GARCH(1,1) process is heavier than a normal distribution. In addition, ordinary least squares analysis is inherently weak when dealing with outliers, which may lead to inaccurate estimates of R^2 and give imprecise representation of the actual fit of the regression. Due to the capacity of volatility to exhibit sudden movements, this may be of particular concern. This issue may be partially addressed in future research by using a logarithmic scale to normalize the data points.

Finally, while the proposed model forecasts the realized volatility of both asset and market returns, it assumes a no-change model for the realized correlation. The expected realized correlation should theoretically not change significantly from one day to the next. However, because the realized correlation is calculated

[12]For the statistical properties of the GARCH and GARCHX models see Han [2011]

each day from realized variances and covariances, sudden changes in these measures will result in a similar change in the realized correlation. There is, therefore, a certain level of variability expected to be present in the realized correlation, as illustrated in Figure 4. This is likely to introduce a measurable degree of noise into the forecasts for the conditional beta (Figure 6). This limitation may be addressed in future research by modeling the realized correlation with a GARCH or GARCHX forecast model itself. Alternatively, rather than represent the realized beta in terms of the asset's and market's respective realized variances and the realized correlation between the two, the definition of realized beta as the ratio of the realized covariance between the asset's and market's returns over the realized variance of the market's returns may be employed. The proposed GARCHX model could then be adapted to forecast the realized covariance.

Overall the results of this paper illustrate the time-varying dynamics of beta and the need to move away from its traditional static definition within the CAPM. Realized measures computed with high-frequency data enable intraday information to be instantaneously integrated into the calculation of the conditional beta. The GARCHX framework, utilizing these realized measures, has been illustrated to be a better predictor of the daily realized beta when compared to either the base GARCH model or benchmark ARMA model. These more accurate forecasts of the realized beta can be instrumental in company valuation and portfolio management and allow investors to better analyze the dynamic, underlying volatility of risky assets.

8 Appendix

8.1 Tables

Equity	Variance	95% Confidence Interval
AAPL	0.3441	0.3283 - 0.3610
EXC	0.0912	0.0864 - 0.0965
F	0.2084	0.1988 - 0.2186
IBM	0.0710	0.0678 - 0.0745
JPM	0.1371	0.1308 - 0.1438
KFT	0.0643	0.0604 - 0.0685
PG	0.0907	0.0865 - 0.0951
WMT	0.1102	0.1051 - 0.1156
Average	0.1396	—

Table 1: Analysis of the Variance of Realized Beta

	GARCH Coefficients				GARCHX-RV Coefficients				
Equity	α_0	α_1	θ_1	R^2	α_0	α_1	θ_1	γ_1	R^2
AAPL	0.0306 (0.4485)	0.0577 (0.1430)	0.9392 (0.0000)	0.2994	0.1967 (0.1811)	0.0277 (0.2875)	0.7015 (0.0000)	0.2008 (0.0000)	0.4153
EXC	0.0285 (0.3937)	0.1061 (0.0297)	0.8833 (0.0000)	0.3952	0.0585 (0.2780)	0.0927 (0.0396)	0.6533 (0.0000)	0.1910 (0.0000)	0.4694
F	0.0571 (0.4464)	0.0730 (0.3121)	0.9155 (0.0000)	0.2363	0.1566 (0.3481)	0.0215 (0.4393)	0.6862 (0.0000)	0.2101 (0.0000)	0.3212
IBM	0.0075 (0.4634)	0.0698 (0.0327)	0.9285 (0.0000)	0.3138	-0.0026 (0.5140)	0.0210 (0.2700)	0.7838 (0.0000)	0.1739 (0.0000)	0.4365
JPM	0.0152 (0.4670)	0.0976 (0.0820)	0.9048 (0.0000)	0.3416	-0.0059 (0.5141)	0.0289 (0.3241)	0.7412 (0.0000)	0.2252 (0.0000)	0.4639
KFT	0.0247 (0.3727)	0.0522 (0.0704)	0.9295 (0.0000)	0.2566	0.0470 (0.2585)	0.0078 (0.4049)	0.6843 (0.0000)	0.2325 (0.0000)	0.3755
PG	0.0053 (0.4728)	0.0514 (0.1143)	0.9455 (0.0000)	0.2191	0.0160 (0.4141)	0.0260 (0.2607)	0.5094 (0.0000)	0.3481 (0.0000)	0.2946
WMT	0.0066 (0.4712)	0.0620 (0.0713)	0.9364 (0.0000)	0.2716	0.0010 (0.4954)	0.0232 (0.2833)	0.8257 (0.0000)	0.1176 (0.0000)	0.3369
AVERAGE	0.0219	0.0712	0.9228	0.2917	0.0584	0.0311	0.6982	0.2124	0.3892

Table 2: Coefficients of the GARCH and GARCHX Equations

Equity	First-Order Autocorrelation	ACF 95% CI	First-Order Partial Autocorrelation	PACF 95% CI
AAPL	0.4232	±0.0664	0.4233	±0.0343
EXC	0.3572	±0.0662	0.3575	±0.0399
F	0.3867	±0.0615	0.3867	±0.0343
IBM	0.3187	±0.0536	0.3187	±0.0343
JPM	0.4603	±0.0672	0.4606	±0.0343
KFT	0.2054	±0.0598	0.2056	±0.0454
P	0.5483	±0.0846	0.5484	±0.0343
WMT	0.4716	±0.0722	0.4720	±0.0343
Average	0.3964	—	0.3966	—

Table 3: First-Order Autocorrelations and Partial Autocorrelations of Realized Beta

Equity	θ	φ
AAPL	0.9946	-0.0963
EXC	0.9937	-0.1120
F	0.9929	-0.1032
IBM	0.9969	-0.0770
JPM	0.9967	-0.1325
KFT	0.9881	-0.0116
PG	0.9959	-0.0917
WMT	0.9962	-0.0842
Average	0.9944	-0.0886

Table 4: ARMA Model Coefficients

	Realized Beta			Market Beta		
Equity	GARCH	GARCHXRV	ARMA	GARCH	GARCHXRV	ARMA
AAPL	0.2898	0.2866	0.3159	0.0374	0.0338	0.0186
EXC	0.0835	0.0832	0.0927	0.0103	0.0096	0.0056
F	0.1793	0.1782	0.1961	0.0217	0.0199	0.0118
IBM	0.0740	0.0738	0.0858	0.0095	0.0087	0.0047
JPM	0.1260	0.1263	0.1428	0.0162	0.0144	0.0080
KFT	0.0625	0.0628	0.0672	0.0068	0.0064	0.0042
PG	0.0597	0.0593	0.0700	0.0082	0.0075	0.0040
WMT	0.0879	0.0870	0.1025	0.0118	0.0113	0.0058
Average	0.1203	0.1197	0.1341	0.0152	0.0140	0.0078

Table 5: MSE for GARCH, GARCHX, and ARMA Forecasts of Daily Realized Beta and Market Beta

	GARCH Model-Implied Beta					*GARCHX-BV Model-Implied Beta*				
Equity	Mean	Median	Min	Max	Std. Dev.	Mean	Median	Min	Max	Std. Dev.
AAPL	1.1044	1.1061	0.0688	2.0850	0.3778	1.1071	1.1239	0.1765	2.0776	0.3440
EXC	0.5674	0.5447	0.2043	1.1217	0.1666	0.5686	0.5421	0.1850	1.0607	0.1544
F	0.7399	0.6916	0.1605	1.6989	0.2616	0.7524	0.7123	0.1475	1.6500	0.2670
IBM	0.7420	0.7177	0.4432	1.3027	0.1965	0.7453	0.7248	0.4878	1.1537	0.1517
JPM	0.9901	0.9713	0.5388	2.0896	0.2794	0.9965	0.9510	0.5550	2.3172	0.2813
KFT	0.3051	0.3051	0.0226	0.6015	0.0991	0.3159	0.3076	0.0558	0.5995	0.1122
PG	0.5381	0.4755	0.2029	1.1945	0.2148	0.5371	0.4844	0.2069	1.1225	0.1929
WMT	0.6947	0.6828	0.1631	1.3329	0.2571	0.6949	0.7018	0.2622	1.2411	0.2143
	GARCHX-RVModel-Implied Beta					*Realized Beta*				
Equity	Mean	Median	Min	Max	Std. Dev.	Mean	Median	Min	Max	Std. Dev
AAPL	1.1008	1.1156	0.1823	2.0261	0.3349	1.1303	1.1087	0.1284	2.1516	0.3724
EXC	0.5662	0.5449	0.1944	1.0394	0.1507	0.5817	0.5634	0.1146	1.1530	0.1699
F	0.7478	0.7066	0.1493	1.6585	0.2568	0.7712	0.7247	0.1380	1.8923	0.2747
IBM	0.7440	0.7208	0.4828	1.1811	0.1523	0.7678	0.7510	0.4899	1.2409	0.1419
JPM	0.9939	0.9555	0.5393	2.3279	0.2762	0.9753	0.9404	2.1437	0.5809	0.2334
KFT	0.3147	0.3073	0.0573	0.5942	0.1075	0.3256	0.3159	.0226	0.6980	0.1206
PG	0.5377	0.4777	0.2031	1.1537	0.1997	0.5941	0.5399	0.2007	1.3209	0.2224
WMT	0.6951	0.6995	0.2594	1.2590	0.2158	0.7567	0.7816	0.2400	1.2698	0.2218

Table 6: Descriptive Statistics of Realized Beta and GARCH and GARCHX Conditional Betas

8.2 Figures

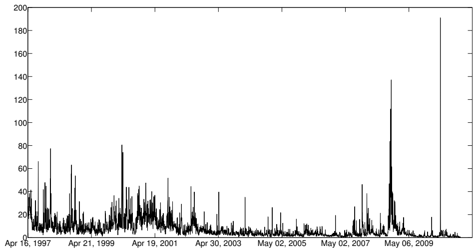

(a) AAPL

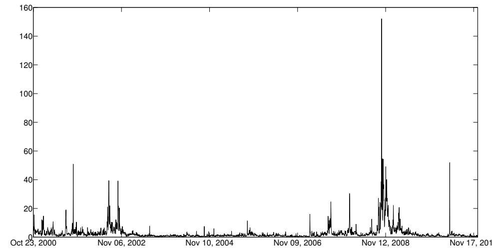

(b) EXC

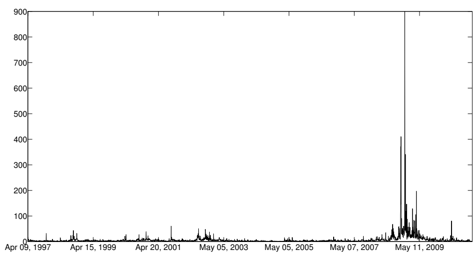

(c) F

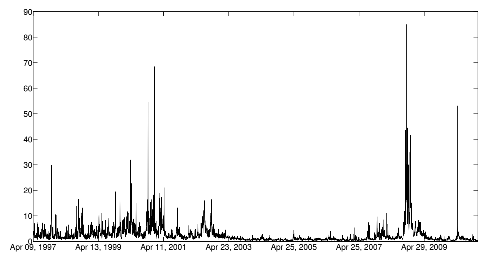

(d) IBM

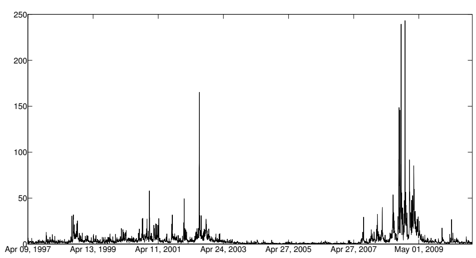

(e) JPM

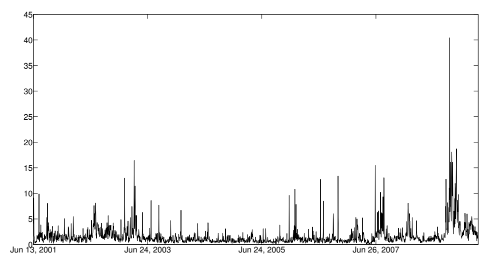

(f) KFT

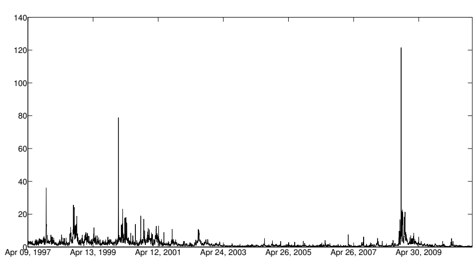

(g) PG

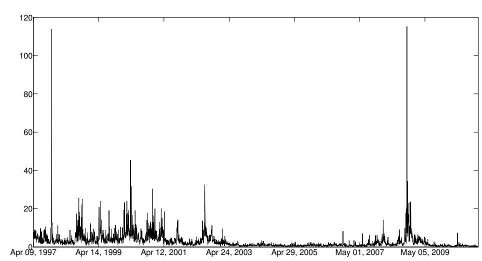

(h) WMT

Figure 1: Equity Realized Variances

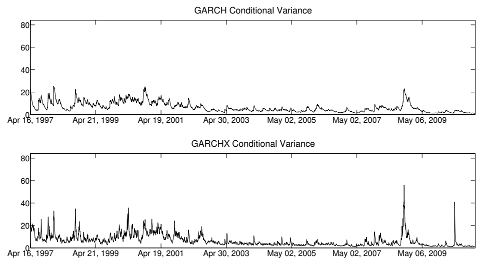

(a) AAPL

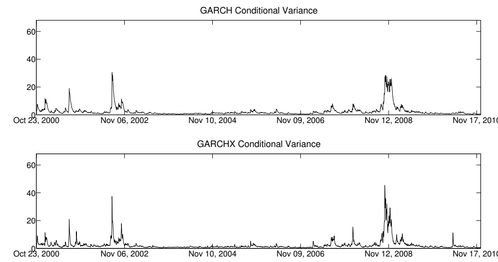

(b) EXC

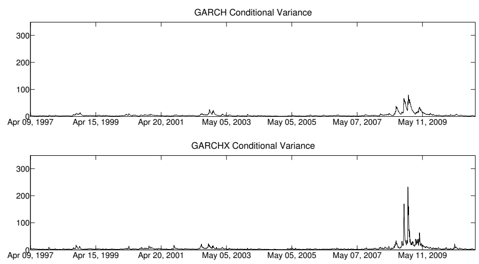

(c) F

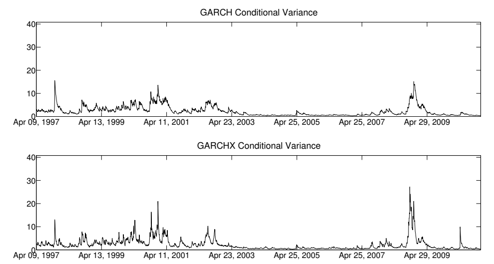

(d) IBM

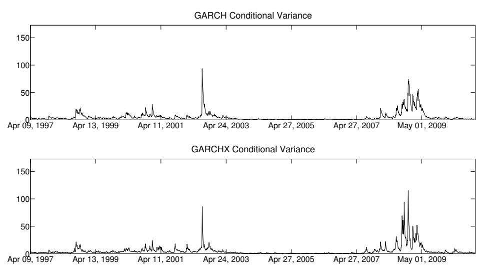

(e) JPM

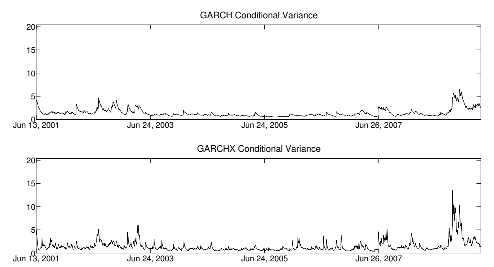

(f) KFT

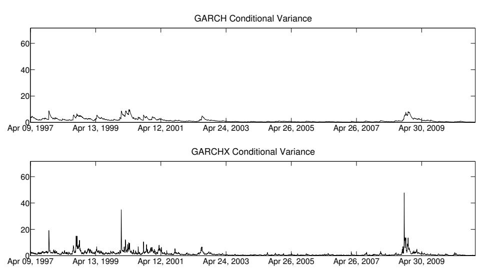

(g) PG

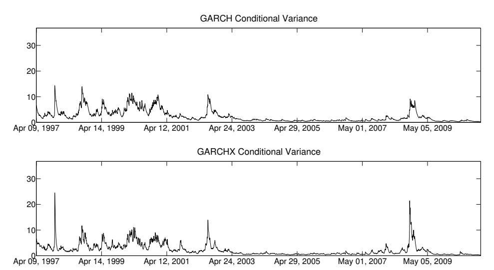

(h) WMT

Figure 2: GARCH and GARCHX Forecasted Conditional Variances

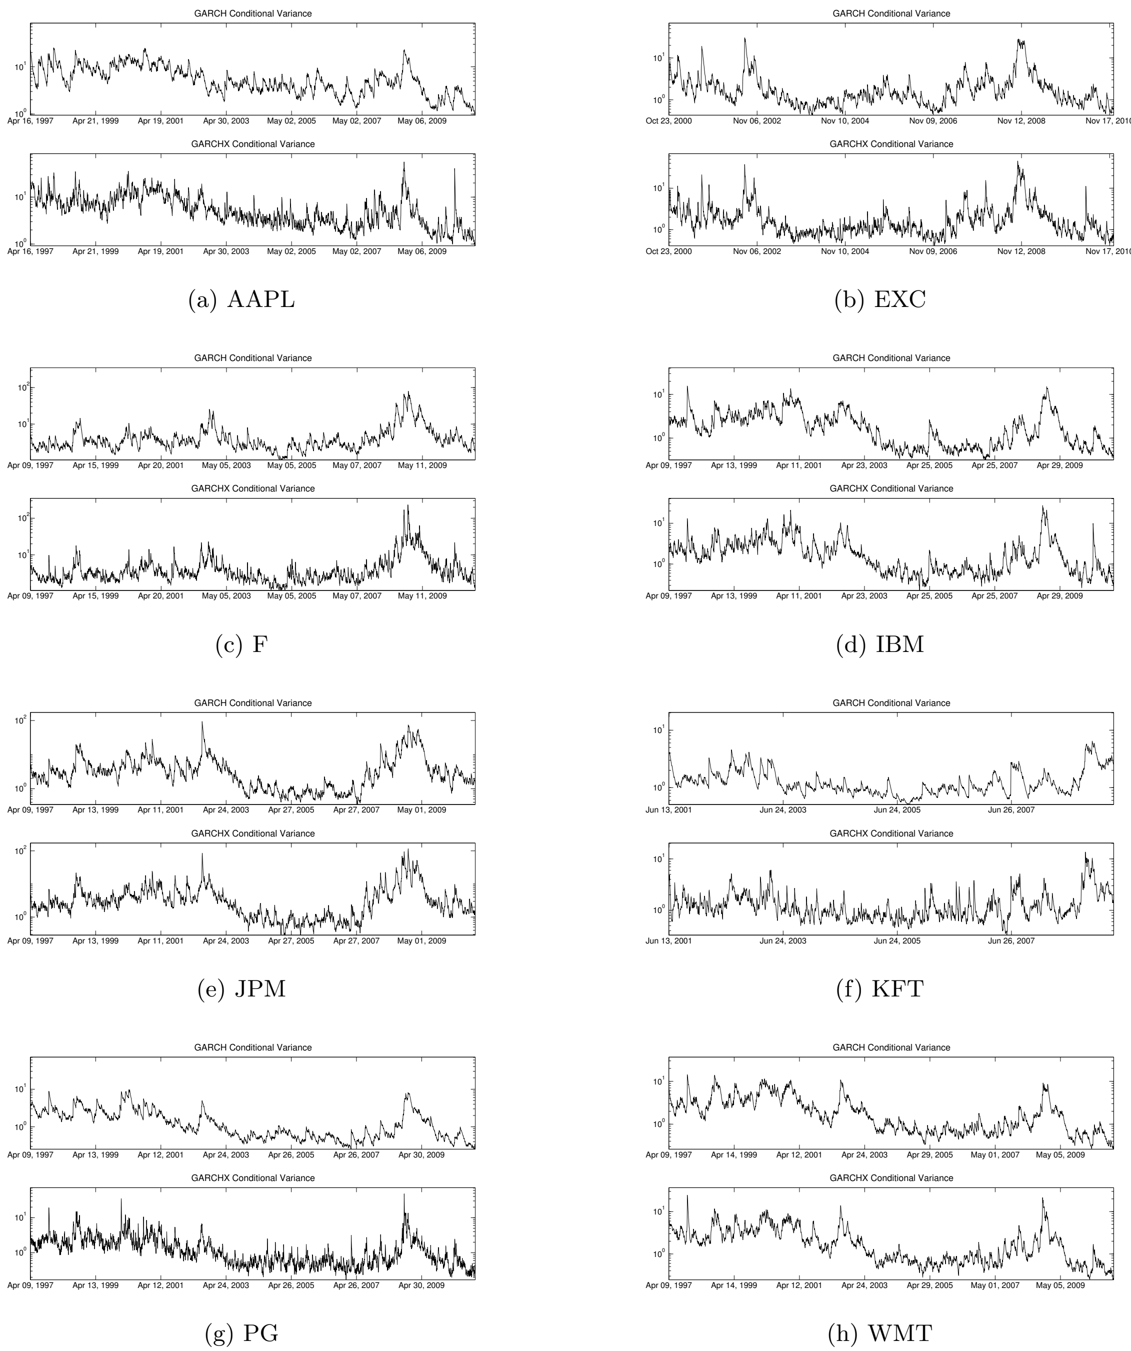

(a) AAPL (b) EXC

(c) F (d) IBM

(e) JPM (f) KFT

(g) PG (h) WMT

Figure 3: GARCH and GARCHX Forecasted Conditional Variances (Logarithmic Scale)

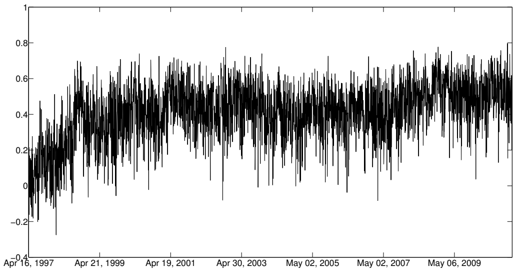

(a) AAPL

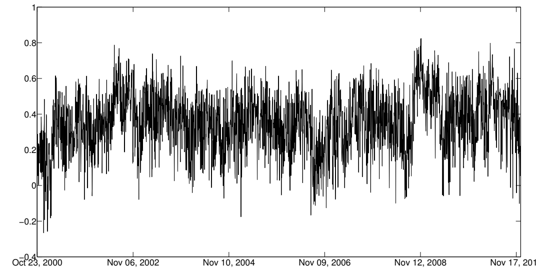

(b) EXC

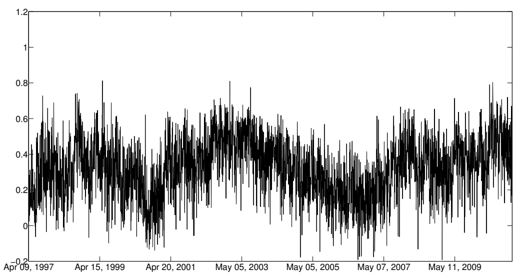

(c) F

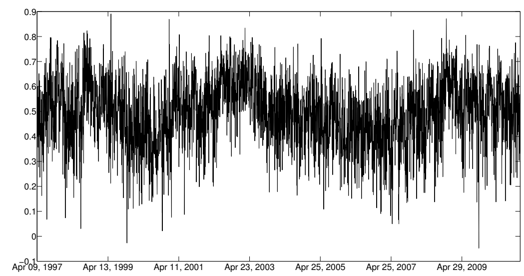

(d) IBM

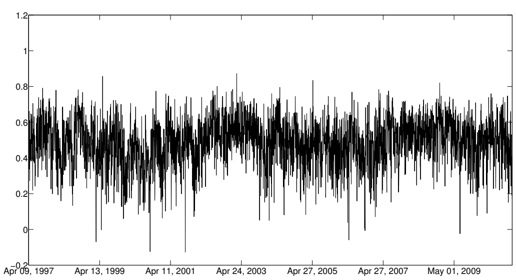

(e) JPM

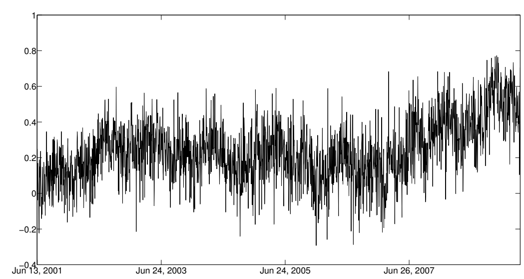

(f) KFT

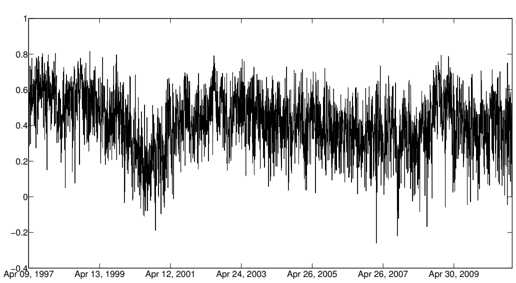

(g) PG

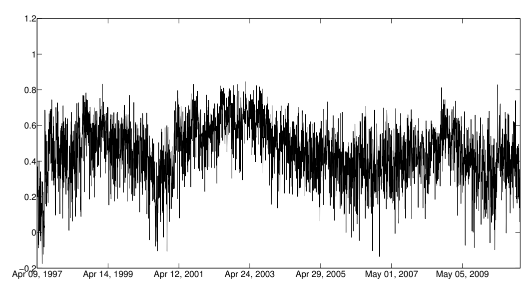

(h) WMT

Figure 4: Equity Realized Correlations with S&P500

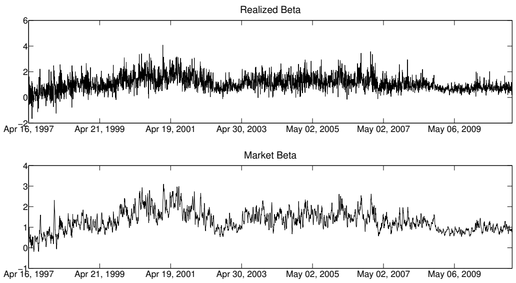

(a) AAPL

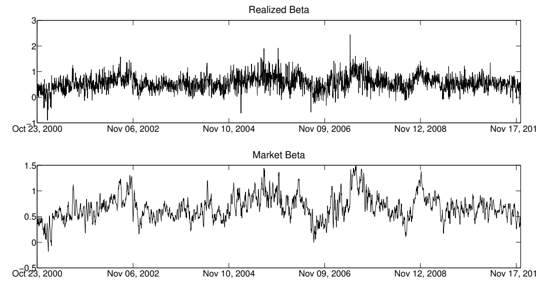

(b) EXC

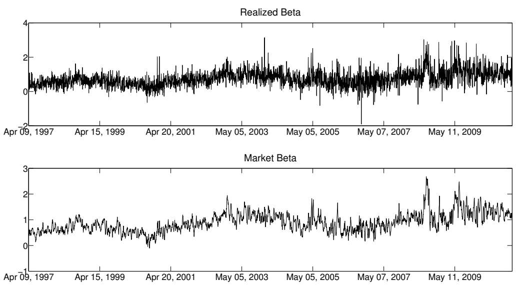

(c) F

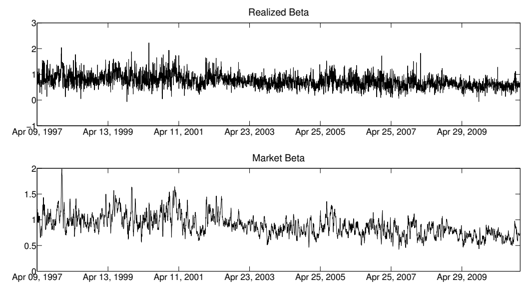

(d) IBM

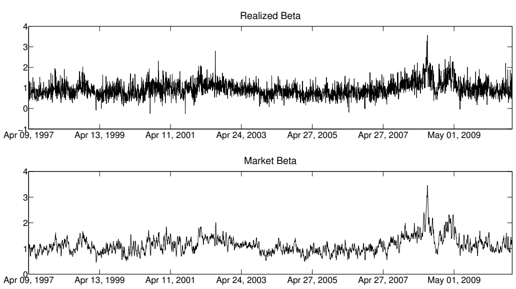

(e) JPM

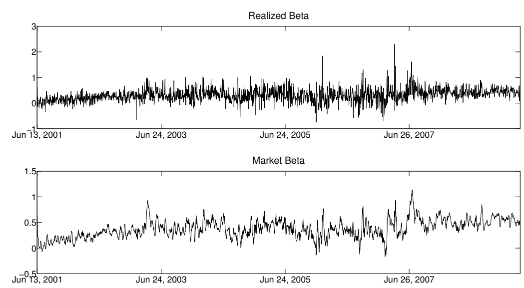

(f) KFT

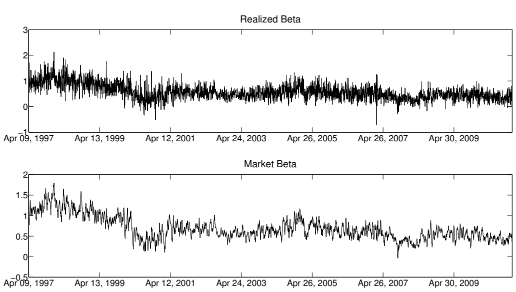

(g) PG

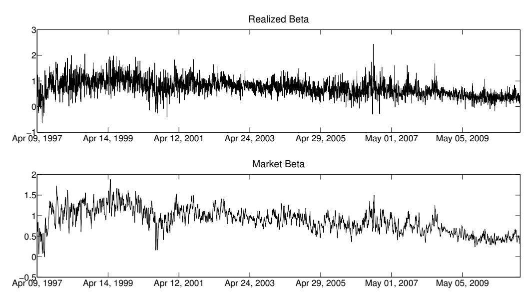

(h) WMT

Figure 5: Realized Beta and Model Specified Market Beta

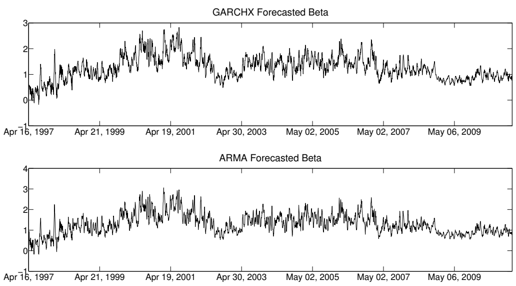

(a) AAPL

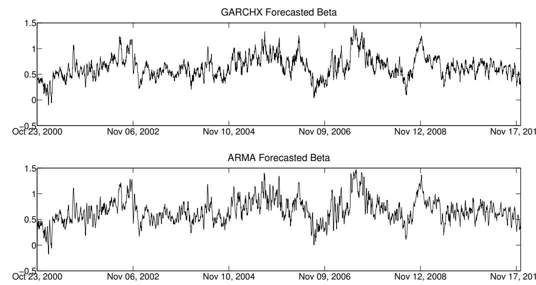

(b) EXC

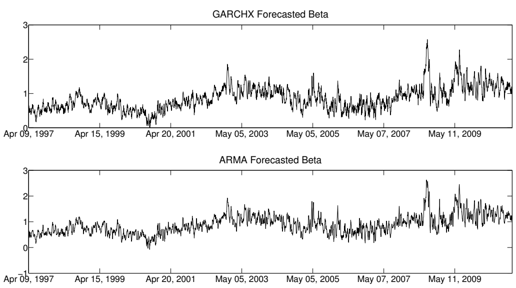

(c) F

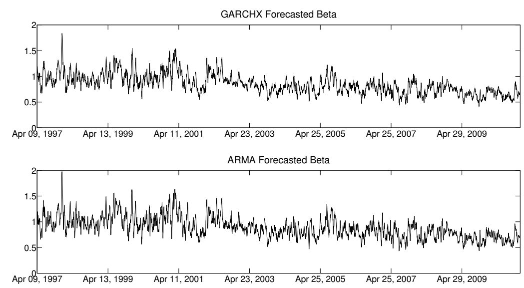

(d) IBM

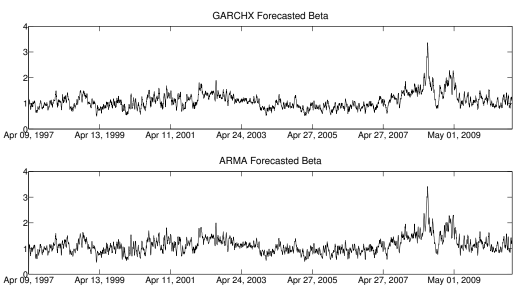

(e) JPM

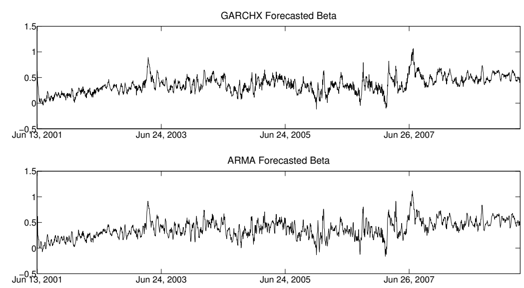

(f) KFT

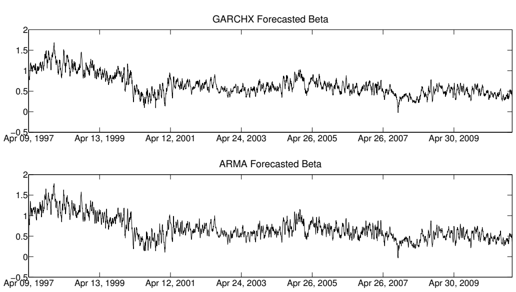

(g) PG

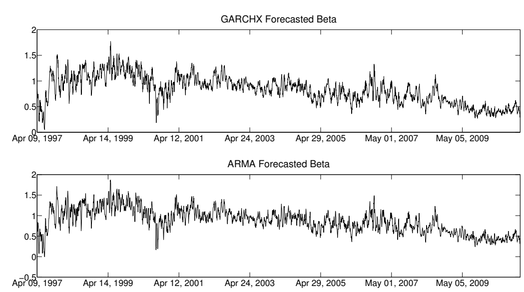

(h) WMT

Figure 6: GARCHX and ARMA Forecasted Beta

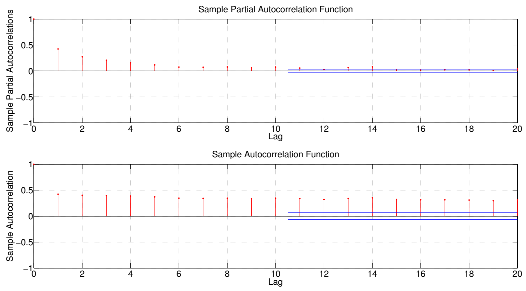

(a) AAPL

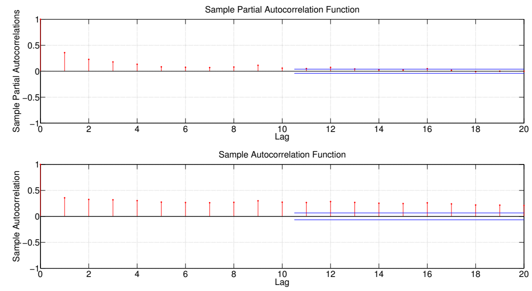

(b) EXC

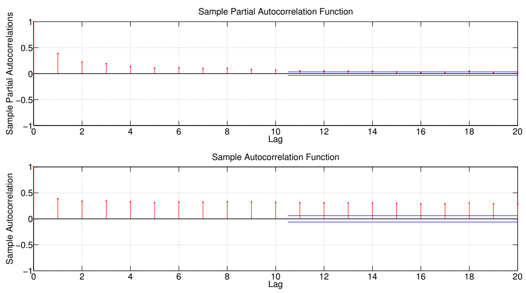

(c) F

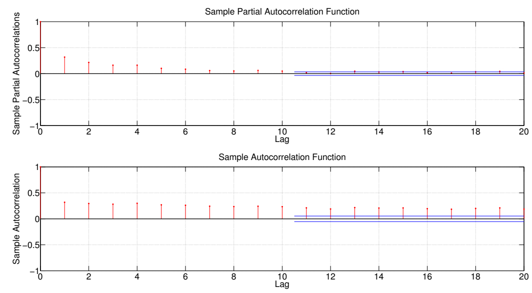

(d) IBM

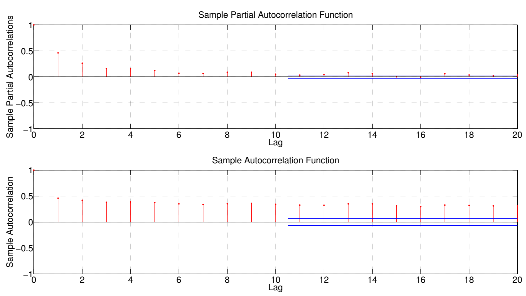

(e) JPM

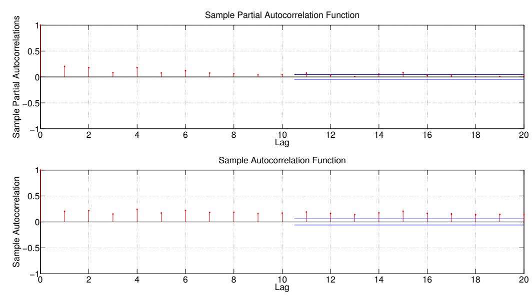

(f) KFT

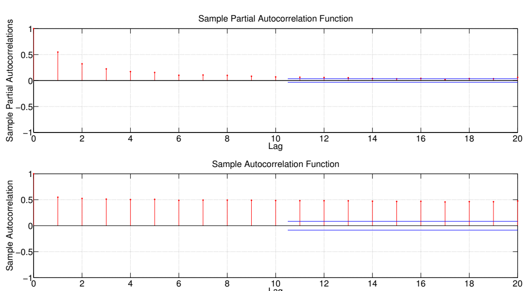

(g) PG

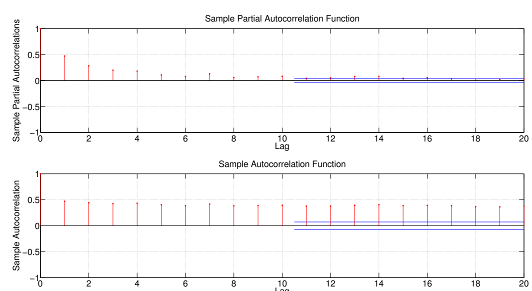

(h) WMT

Figure 7: Autocorrelations and Partial Autocorrelations of Realized Beta

8.3 Maximum Likelihood Estimation of GARCH and GARCHX Parameters

Index the observations as

$t = 0, 1, ..., T-1$

Let β_t, h_t be the market beta and conditional variance for time t so their values are determined at time $t-1$. Other variables are the following:

r_t : open-to-close return at day t of the asset under study

$\hat{\beta}_t$: the realized beta at day t of the asset under study

RV_t : the realized variance at day t of the asset under study

The dynamics of h_t and β_t are described in the following equations:[13]

$a_t = r_t - \mu$

$h_t^{GARCHX} = \alpha_0 + \sum_{i=1}^{m} \alpha_i a_{t-i}^2 + \sum_{j=1}^{s} \theta_j h_{t-j}^{GARCHX} + \sum_{k=1}^{n} \gamma_k \cdot RV_{t-k}$

Then define

$\omega = (\alpha_0, \alpha_1, ...\alpha_m, \theta_1, ..., \theta_s, \gamma_1, ..., \gamma_n)$

$z_t = (1, \mu_t^2, ..., \mu_{t-m+1}^2, h_t, ..., h_{t-m+1, RV_t, ..., RV_{t-k+1}})$

Then the dynamics of h_t can be written as

$h_t = \omega' z_{t-1}$

The log likelihood function is

$$L(\omega) = \sum_{t=0}^{T-1} l_t(\omega) = \sum_{t=0}^{T-1} \left[-\frac{1}{2} ln(2\pi) - \frac{1}{2} log h_t - \frac{1}{2} \frac{a_t^2}{h_t} \right] \tag{49}$$

Since the first term $ln(2\pi)$ does not involve any parameters, the log likelihood function becomes

$$L(\omega) = \sum_{t=0}^{T-1} l_t(\omega) = \sum_{t=0}^{T-1} \left[-\frac{1}{2} log h_t - \frac{1}{2} \frac{a_t^2}{h_t} \right] \tag{50}$$

The estimates obtained by maximizing the prior log likelihood function are referred to as the conditional maximum likelihood estimates under normality.

[13] This specification is for the GARCHX framework. Removing the RV_t summation above will result in the standard GARCH equation for the conditional variance. This can be similarly optimized using the following likelihood estimation.

References

T. Andersen, T. Bollerslev, F. Diebold, and H. Ebens. The distribution of realized stock return volatility. *Journal of Financial Economics*, 61(1):43–76, 2001a.

T. Andersen, T. Bollerslev, and F. Diebold. Parametric and nonparametric volatility measurement, 2002.

T. Andersen, T. Bollerslev, F. Diebold, and P. Labys. Modeling and forecasting realized volatility. *Econometrica*, 71 (2):579–625, 2003.

T. Andersen, T. Bollerslev, F. Diebold, and G. Wu. Realized Beta: Persistence and Predictability. *Advances in Econometrics: Econometric Analysis of Economic and Financial Time Series*, 20:1–39, 2006.

T. G. Andersen and T. Bollerslev. Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns. *Journal of Finance*, 53(3):975–1005, 1997a.

T. G. Andersen and T. Bollerslev. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4:115–158, 1997b.

T. G. Andersen and T. Bollerslev. Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance*, 53(1):219–265, 1998a.

T. G. Andersen and T. Bollerslev. Answering the skeptics: Yes standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4):885–905, 1998b.

T. G. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, 96:42–55, 2001b.

N. Apergis and A. Rezitis. Food price volatility and macroeconomic factors: Evidence from garch and garch-x estimates. *Journal of Agricultural and Applied Economics*, 43(1):95–110, 2011.

R. Ballie and T. Bollerslev. The message in daily exchange rates: A conditional variance tale. *Journal of Business and Economic Statistics*, 7:297–305, 1989.

O. E. Barndorff-Neilson and N. Shephard. Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics. *Econometrica*, 72(3):885–925, 2004.

O. E. Barndorff-Neilson and N. Shephard. Variation, jumps, and high frequency data in financial econometrics. *Advances in Economics and Econometrics. Theory and Applications*, 2007.

O. Barndorff-Nielsen and N. Shephard. Aggregation and model construction for volatility models, 1998.

O. E. Barndorff-Nielsen and N. Shephard. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2(1):1–37, 2004.

B. Ben Sita and S. Abosedra. Volatility spillovers: Evidence on us oil product markets. *Journal of Applied Business Research (JABR)*, 28(6):1237–1242, 2012.

T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327, 1986.

T. Bollerslev. A conditional heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics*, 69:542–547, 1987.

T. Bollerslev and R. Engle. Common persistence in conditional variances. *Econometrica: Journal of the Econometric Society*, 61(1):167–186, 1993.

T. Bollerslev, R. F. Engle, and J. M. Wooldridge. A capital asset pricing model with time-varying covariances. *The Journal of Political Economy*, 96(1):116–131, 1988.

T. Bos and P. Newbold. An empirical investigation of the possibility of stochastic systematic risk in the market model. *Journal of Business*, 57(1):35 – 41, 1984. ISSN 00219398.

R. J. Brenner, R. H. Harjes, and K. F. Kroner. Another look at models of the short-term interest rate. *The Journal of Financial and Quantitative Analysis*, 31(1):pp. 85–107, 1996. ISSN 00221090.

J. Campbell, A. Lo, and A. Mckinlay. *The Econometrics of Financial Markets*. Princeton Univers. Press, 1997. ISBN 9780691043012.

T. Caporale. Time varying capm betas and banking sector risk. *Economics Letters*, 115(2):293 – 295, 2012. ISSN 0165-1765.

T. Choudhry. The stochastic structure of the timevarying beta: Evidence from uk companies. *The Manchester School*, 70(6):768–791, 2002. ISSN 1467-9957.

F. Cipollini, R. Engle, and G. Gallo. A model for multivariate non-negative valued processes in financial econometrics. *Available at SSRN 1333869*, 2009.

M. Ebner and T. Neumann. Time-varying betas of german stock returns. *Financial Markets and Portfolio Management*, 19:29–46, 2005. ISSN 1555-4961.

R. Engle. Autoregressive conditional heteroscedasity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.

R. Engle. New Frontiers for ARCH Models. *Journal of Applied Econometrics*, 17(4):425–446, 2002.

R. Engle and A. Patton. What good is a volatility model? *Quantitative Finance*, 1(2):237–245, 2001.

R. F. Engle and T. Bollerslev. Modelling te persistence of conditional variances. *Econometric Reviews*, 5:1–50, 1986.

R. F. Engle and G. M. Gallo. A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics*, 131(1 - 2):3 – 27, 2006. ISSN 0304-4076.

B. T. Ewing and F. Malik. Re-examining the asymmetric predictability of conditional variances: The role of sudden changes in variance. *Journal of Banking and Finance*, 29(10):2655 – 2673, 2005. ISSN 0378-4266.

E. Fama. *Foundations of Finance: Portfolio Decisions and Securities Prices*. Basic Books, 1976. ISBN 9780465024995.

S. Figlewski. Forecasting volatility. *Financial Markets, Institutions & Instruments*, 6(1):1–88, 1997.

L. R. Glosten, R. Jagannathan, and D. E. Runkle. On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5):pp. 1779–1801, 1993. ISSN 00221082.

S. F. Gray. Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics*, 42(1):27 – 62, 1996. ISSN 0304-405X.

N. Groenewold and P. Fraser. Time-varying estimates of capm betas. *Mathematics and Computers in Simulation*, 48 (4-6):531 – 539, 1999. ISSN 0378-4754.

M. Hagiwara and M. A. Herce. Endogenous exchange rate volatility, trading volume and interest rate differentials in a model of portfolio selection. *Review of International Economics*, 7(2):202–218, 1999. ISSN 1467-9396.

H. Han. Asymptotic properties of garchx processes. *Manuscript, National University of Singapore*, 2011.

H. Han and D. Kristensen. Asymptotic theory for the qmle in garch-x models with stationary and non-stationary covariates. *Creates research paper*, 25, 2012.

H. Han and J. Y. Park. Time series properties of arch processes with persistent covariates. *Journal of Econometrics*, 146(2):275 – 292, 2008. ISSN 0304-4076.

P. R. Hansen and A. Lunde. A forecast comparison of volatility models: Does anything beat a garch(1,1)? *Journal of Applied Econometrics*, 20(7):pp. 873–889, 2005. ISSN 08837252.

P. R. Hansen, A. Lunde, and V. Voev. Realized beta garch: A multivariate garch model with realized measures of volatility and covolatility. CREATES Research Papers 2010-74, School of Economics and Management, University of Aarhus, Nov 2010.

P. R. Hansen, Z. Huang, and H. H. Shek. Realized garch: a joint model for returns and realized measures of volatility. *Journal of Applied Econometrics*, 27(6):877–906, 2012. ISSN 1099-1255.

R. J. Hodrick. Risk, uncertainty, and exchange rates. *Journal of Monetary Economics*, 23(3):433 – 459, 1989. ISSN 0304-3932.

C. Huang and R. Litzenberger. *Foundations for financial economics*. Prentice Hall, 1988.

X. Huang and G. Tauchen. The relative contribution of jumps to total price variance. *Journal of Financial Econometrics*, 3(4):456–499, 2005.

S. Hwang and S. E. Satchell. GARCH model with cross-sectional volatility: GARCHX models. *Applied Financial Economics*, 15:203–216, 2005.

R. Jagannathan and Z. Wang. The conditional capm and the cross-section of expected returns. *The Journal of Finance*, 51(1):pp. 3–53, 1996.

K. Jain. Time-varying beta: The heterogeneous autoregressive beta model. 2011.

P. Jorian. Predicting volatility in the foreign exchange market. *The Journal of Finance*, 50:507–528, 1995.

C.-J. Kim and C. R. Nelson. The time-varying-parameter model for modeling changing conditional variance: The case of the lucas hypothesis. *Journal of Business and Economic Statistics*, 7(4):pp. 433–440, 1989.

J. Lintner. Security prices, risk, and maximal gains from diversification. *The Journal of Finance*, 20(4):pp. 587–615, 1965a. ISSN 00221082.

J. Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1):pp. 13–37, 1965b. ISSN 00346535.

C. Mancini. Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics*, 36(2):270–296, 2009. ISSN 1467-9469.

C. Morana. Realized betas and the cross-section of expected returns. *Applied Financial Economics*, 19(17):1371 – 1381, 2009. ISSN 09603107.

M. Mulyadi and Y. Anwar. Return and volatility spillover across usa and europe (study of american and eu crisis period). *African Journal of Business Management*, 6(19):5916–5926, 2012.

R. Reider. Volatility forecasting i: Garch models. 2009.

R. Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance*, 39 (4):pp. 1127–1139, 1984. ISSN 00221082.

W. F. Sharpe. A simplified model for portfolio analysis. *Management Science (pre-1986)*, 9(2), 1963.

N. Shephard and K. Sheppard. Realising the future: forecasting with high-frequency-based volatility (heavy) models. *Journal of Applied Econometrics*, 25(2):197–231, 2010. ISSN 1099-1255.

S. K. Staikouras. Testing the stabilization hypothesis in the uk short-term interest rates: Evidence from a garch-x model. *The Quarterly Review of Economics and Finance*, 46(2):169 – 189, 2006. ISSN 1062-9769.

R. Tsay. *Analysis of Financial Time Series*. Wiley Series in Probability and Statistics. Wiley, 2010. ISBN 9781118017098.

D. I. Vortelinos. The properties of realized correlation: Evidence from the french, german and greek equity markets. *The Quarterly Review of Economics and Finance*, 50(3):273 – 290, 2010. ISSN 1062-9769.