REALIZED BETA GARCH: A MULTIVARIATE GARCH MODEL WITH REALIZED MEASURES OF VOLATILITY

PETER REINHARD HANSEN,[a,b]* ASGER LUNDE[b,c] AND VALERI VOEV[c,d]

[a] *European University Institute, Florence, Italy*
[b] *CREATES, Aarhus, Denmark*
[c] *Department of Economics and Business, Aarhus University, Denmark*
[d] *Siemens Wind Power A/S, Brande, Denmark*

SUMMARY

We introduce a multivariate generalized autoregressive conditional heteroskedasticity (GARCH) model that incorporates realized measures of variances and covariances. Realized measures extract information about the current levels of volatilities and correlations from high-frequency data, which is particularly useful for modeling financial returns during periods of rapid changes in the underlying covariance structure. When applied to market returns in conjunction with returns on an individual asset, the model yields a dynamic model specification of the conditional regression coefficient that is known as the *beta*. We apply the model to a large set of assets and find the conditional betas to be far more variable than usually found with rolling-window regressions based exclusively on daily returns. In the empirical part of the paper, we examine the cross-sectional as well as the time variation of the conditional beta series during the financial crises. Copyright © 2014 John Wiley & Sons, Ltd.

Supporting information may be found in the online version of this article.

1. INTRODUCTION

Relatively accurate measurements of volatility and covariances can be computed from high-frequency data, and such statistics are commonly referred to as *realized measures*. Incorporating realized measures when modeling the dynamic properties of volatility is very beneficial, and leads to better empirical fit than conventional generalized autoregressive conditional heteroskedasticity (GARCH) models, which only use daily returns. The reason is that returns yield very weak signals about latent volatility, whereas realized measures provide accurate measurements. The latter is particularly useful during times of rapid changes in volatility and correlations.

In this paper we propose a multivariate GARCH-type model that utilizes and models realized measures of volatility and correlations. The model has a hierarchical structure where the 'market' return is modeled with a univariate realized GARCH model (see Hansen and Huang, 2012; Hansen *et al.*, 2012). A multivariate structure is constructed by modeling 'individual' returns conditional on the past and contemporary market variables (return and volatility). The resulting model has the structure of a dynamic capital asset pricing model (CAPM) that enables us to extract the 'betas' and study their dynamic properties. Moreover, the model is complete in the sense that all observables (returns, realized volatilities and realized correlations) are modeled. The latter enables us to infer the distribution of multi-period returns including the joint distribution of 'market' returns and 'individual' returns over longer horizons.

The main contributions of our paper are the following. We propose a flexible and tractable framework that enables the modeling of a potentially large set of assets. Unlike conventional multivariate GARCH models, which can suffer from the curse of dimensionality and estimation issues,

---

* Correspondence to: Peter Reinhard Hansen, European University Institute, Villa San Paolo, Via della Piazzuola 43, I-50133 Firenze, Italy. E-mail: reinhardhansen@gmail.com

we avoid such issues by incorporating realized measures and the use of measurement equations. The measurement equations tie realized measures to latent volatilities and correlations, and this leads to a useful regularization of the model. This particular structure was chosen for a number of reasons. First, the model provides a good empirical fit for the wide range of assets used in our empirical study; second, the structure of the model is amenable to a deeper analysis of secondary quantities such as betas; third, the model is simple to estimate, which is particularly important when a large set of assets are to be analyzed, as is the case in our empirical analysis.

The proposed model has a hierarchical structure in which the market return and a corresponding realized measure form the core of the model. The model can be extended to an arbitrarily large set of individual returns, by adding a conditional model for an individual return and two realized measures, one being a realized measure of return volatility, and the other a realized measure of the correlation between the individual return and the market. This yields a flexible model with a dynamic covariance structure that is constantly revised by using the information contained in the realized measures.

The concept of realized betas is not new. Bollerslev and Zhang (2003) carried out a large-scale estimation of the Fama–French three-factor model using high-frequency (5-minute) data on 6400 stocks over a period of 7 years. Their analysis showed that high-frequency data can improve the pricing accuracy of asset pricing models. Their approach differs from ours in important ways. For instance, they model raw realized factor loadings and use simple time series processes to forecast these. Thus there is no explicit link between realized and conditional moments of returns in their framework. Nor do they explicitly account for the measurement error (sampling error) in the realized quantities. Another related paper is Andersen *et al.* (2006), who study the time variation in realized variances, covariances and betas using daily returns to construct quarterly realized measures. They find evidence of long memory in the time series for variance and covariances, while the realized beta time series is less persistent and seemingly a short-memory process, which is indicative of fractional cointegration between realized volatility and realized covariance. Other related studies include Barndorff-Nielsen and Shephard (2004a), who established asymptotic results for realized beta, and Dovonon *et al.* (2013), who established the theory for bootstrap inference. MSE-optimal estimation of realized betas was analyzed in Bandi and Russell (2005), and Patton and Verardo (2012) studied the impact of news on betas. Morana (2009) uses realized betas to explain the variation in expected returns, and Corradi *et al.* (2011) use realized betas to extract the conditional 'alphas'. The importance of separating jump and continuous components of returns in relation to betas, as highlighted in Todorov and Bollerslev (2010) and Tsay and Yeh (2011), allow the dynamic beta to vary within the day.

The use of realized volatility measures in this context yields valuable insight about the degree of time variation in the betas, which has been up for debate in the literature. The studies by Ferson and Harvey (1991, 1993) and Shanken (1990) specify parametric relationships between betas and proxies for the state of the economy, and find support for time-varying betas. Gomes *et al.* (2003) provide a theoretical justification for a time-varying conditional beta specification in the context of a dynamic general equilibrium production economy. Conditional betas have been modeled by means of conventional GARCH models by Braun *et al.* (1995) and Bekaert and Wu (2000), among others. Lewellen and Nagel (2006) argue that variation in betas would have to be 'implausibly large' to explain important asset-pricing anomalies. In our empirical analysis we do find a substantial amount of time variation in the conditional betas. This is particularly the case during the global financial crisis period. We find the variation in betas to be substantial, even over short periods of time, such as a quarter.

The research devoted to high-frequency volatility measures was spurred by Andersen and Bollerslev (1998), who documented that the sum of squared intra-day returns, known as the realized variance, provides an accurate measurement of daily volatility. The theoretical foundation of realized variance was developed in Andersen *et al.* (2001a) and Barndorff-Nielsen and Shephard (2002). A large number of related estimators, such as realized bipower variation, realized kernels, multi-scale estimators, pre-averaging estimators and Markov chain estimators, have been proposed to deal with issues

such as jumps and market microstructure frictions (see Barndorff-Nielsen and Shephard (2004b); Barndorff-Nielsen *et al.* (2008); Zhang (2006), Jacod *et al.* (2009); Hansen and Horel (2009); and references therein. The multivariate extensions of the concept of realized volatility is theoretically developed in Barndorff-Nielsen and Shephard (2004a). Estimators that are robust to noise and/or asynchronous observations have been proposed by Hayashi and Yoshida (2005), Voev and Lunde (2007), Griffin and Oomen (2011), Christensen *et al.* (2010) and Barndorff-Nielsen *et al.* (2011). In this paper we will rely on the multivariate kernel estimator by Barndorff-Nielsen *et al.* (2011), which guarantees positive semi-definite estimates of the realized variance–covariance matrices.

While volatility is unobservable, the use of realized measures allows us to construct precise *ex post* volatility proxies. Currently, a growing body of research investigates the extent to which realized measures can be used to improve the accuracy of volatility forecasts. Hansen and Lunde (2011) categorize the existing approaches into two broad classes: reduced form and model based. Reduced-form volatility forecasts are based on a time series model for the series of realized measures, whereas a model-based forecast rests on a parametric model for the return distribution. Model-based approaches effectively build on GARCH models in which a realized measure is included as an exogenous variable in the GARCH equation (see, for example, Engle, 2002b). A complete framework that jointly specifies models for returns and realized measures of volatility was first proposed by Engle and Gallo (2006), who refer to their model as the multiplicative error model (MEM). A simplified MEM structure was proposed in Shephard and Sheppard (2010), who referred to their model as the HEAVY model. The realized GARCH model by Hansen *et al.* (2012) involves a different approach to the joint modeling of returns and realized volatility measures. A key component of the realized GARCH model is a measurement equation that relates the realized measure to the underlying conditional variance. This idea is generalized to a multivariate framework in this paper, where we introduce measurement equations for the realized measures of correlations.

The rest of the paper is structured as follows. The model and the underlying theory are presented in Section 2, and we discuss estimation of the model in Section 3. In Section 4 we show how multi-step predictions of volatilities and correlations as well as forecasts of return densities can be obtained from the model. Section 5 contains an empirical application of the model and Section 6 concludes. Detailed information about the data construction, and auxiliary results that are useful for the estimation are presented in two appendices.

2. A HIERARCHICAL REALIZED GARCH FRAMEWORK

Broadly speaking, our objective is the same as that of existing multivariate GARCH models, which is to model the conditional distribution of a vector of returns. Unlike conventional GARCH models, however, we also model the realized measures of volatility and correlation and make extensive use of these in the modeling of returns. The realized measures are highly informative about local (in time) levels of volatility and correlation. By tying all individual return series to the market return, we are implicitly imposing a factor structure on the volatility, where the variation in the correlation structure is driven by time variation in the correlations between the market return and the individual assets. This keeps the model relatively simple and parsimonious, facilitates estimation, and makes it easy to relate key variables in the model to dynamic betas.

Our model has a hierarchical structure. The core of our framework is a marginal model for the market return and its realized measure of volatility. Individual returns, their realized measures of volatility and correlation (with the market) are then modeled conditionally on market variables. The marginal model we use for the market-specific time series is a reparametrized version of the realized EGARCH model by Hansen and Huang (2012; see also Hansen *et al.*, 2012, section 6.3), which shares certain features with the EGARCH model by Nelson (1991). The conditional model we use in this paper is new.

Initially, we present the realized beta GARCH model in the simplest situation with a bivariate vector of returns (the market return and an individual asset return) and the corresponding 2×2 matrix of realized volatility measures. Subsequently, we discuss the straightforward extension to an arbitrary number of individual assets.

2.1. Notation and Modeling Strategy

Let $r_{0,t}$ and $x_{0,t}$ denote the market return and a corresponding realized measure of volatility, respectively. Similarly, we use the notation $r_{1,t}$ and $x_{1,t}$ for the same variables associated with an individual asset return, and use $y_{1,t}$ to denote a realized measure of correlation, where $y_{1,t} \in (-1, 1)$.

In this context with two returns, two realized measures of volatility and a realized measure of correlation, we have five observable variables to model. The natural filtration is given by

$$\mathcal{F}_t = \sigma\left(\mathcal{X}_t, \mathcal{X}_{t-1}, \ldots\right) \quad \text{with} \quad \mathcal{X}_t = \left(r_{0,t}, r_{1,t}, x_{0,t}, x_{1,t}, y_{1,t}\right)'$$

and we define the conditional variances, $h_{0,t} = \text{var}\left(r_{0,t}|\mathcal{F}_{t-1}\right)$ and $h_{1,t} = \text{var}\left(r_{1,t}|\mathcal{F}_{t-1}\right)$, much like in standard GARCH models, the key difference being that the information set, $\mathcal{F}_t$, is richer in the present framework. We also define the conditional correlation, $\rho_{1,t} = \text{corr}\left(r_{0,t}, r_{1,t}|\mathcal{F}_{t-1}\right)$, and it follows directly that the 'beta'

$$\beta_{1,t} = \text{cov}\left(r_{1,t}, r_{0,t}|\mathcal{F}_{t-1}\right) / \text{var}\left(r_{0,t}|\mathcal{F}_{t-1}\right) \tag{1}$$

is given by $\beta_{1,t} = \rho_{1,t}\sqrt{h_{1,t}/h_{0,t}}$. This establishes a connection to a dynamic CAPM and we are particularly interested in the dynamic properties of $\beta_{1,t}$.

The structure of our model will take advantage of the simple decomposition of the conditional density:

$$f\left(r_{0,t}, x_{0,t}, r_{1,t}, x_{1,t}, y_{1,t}|\mathcal{F}_{t-1}\right) = f\left(r_{0,t}, x_{0,t}|\mathcal{F}_{t-1}\right) f\left(r_{1,t}, x_{1,t}, y_{1,t}|r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) \tag{2}$$

which serves to illustrate the hierarchical structure of our model. We will adopt the realized EGARCH model as our specification of the first term, $f\left(r_{0,t}, x_{0,t}|\mathcal{F}_{t-1}\right)$. The specification for the second conditional density, $f\left(r_{1,t}, x_{1,t}, y_{1,t}|r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right)$, defines how the time series associated with the individual asset evolves, conditional on contemporary market variables. Our specification of this conditional density has a structure that is similar to that of the univariate realized GARCH model, but has some important adaptations for the modeling of the correlation structure. Realized correlation measures between the individual assets are not needed, because it is implicitly assumed that these correlations are characterized through the correlations between the individual returns and the market return. In our empirical analysis, we investigate the validity of this assumption.

2.2. Realized EGARCH Model for Market Returns

The realized EGARCH model for market returns and realized measures of volatility is given by the following three equations:

$$r_{0,t} = \mu_0 + \sqrt{h_{0,t}}z_{0,t}, \tag{3}$$

$$\log h_{0,t} = a_0 + b_0 \log h_{0,t-1} + c_0 \log x_{0,t-1} + \tau_0\left(z_{0,t-1}\right) \tag{4}$$

$$\log x_{0,t} = \xi_0 + \varphi_0 \log h_{0,t} + \delta_0\left(z_{0,t}\right) + u_{0,t} \tag{5}$$

Here, we do not follow the conventional GARCH notation, because we want to reserve the notation 'β' for the market-beta variable we defined in equation (1).

In our likelihood analysis we specify $z_{0,t} \sim$ i.i.d. $N(0, 1)$ and $u_{0,t} \sim$ i.i.d. $N\left(0, \sigma_{u_0}^2\right)$. The properties of the estimators do not critically hinge on this Gaussian specification being correct. Still, the Gaussian specification for $u_{0,t}$ can be motivated by findings in Andersen *et al.* (2001a, 2001b, 2003), who document that realized volatility is approximately log-normal. Furthermore, Andersen *et al.* (2001b) find that returns standardized by realized volatility are approximately normally distributed.

The functions $\tau(z)$ and $\delta(z)$ are called leverage functions because they model aspects related to the *leverage effect*, which refers to the dependence between returns and volatility. Hansen *et al.* (2012) found that a simple second-order polynomial provides a good empirical fit. We will adopt this structure in our framework, and set $\tau(z) = \tau_1 z + \tau_2 \left(z^2 - 1\right)$ and $\delta(z) = \delta_1 z + \delta_2 \left(z^2 - 1\right)$. This leads to a GARCH equation that is somewhat similar to that of an EGARCH model. Unlike the EGARCH model, however, we also utilize the realized measure, $x_{0,t-1}$, to model the dynamics of volatility.

We refer to the first two equations, (3) and (4), as the *return equation* and the *GARCH equation*, respectively. These two equations define a GARCH-X model, similar to those estimated by Engle (2002b), Barndorff-Nielsen and Shephard (2007) and Visser (2011). See also Chen *et al.* (2011) for additional variants of the GARCH-X model and some related models.

The third equation, (5), called the *measurement equation*, completes the specification of the density, $f\left(r_{0,t}, x_{0,t} | \mathcal{F}_{t-1}\right)$. Tying the realized measure, x_t, to the conditional variance, h_t, is motivated by the fact that the GARCH equation trivially implies that

$$\log\left(r_t - \mu\right)^2 = \log h_t + \log z_t^2$$

Since the realized measure, x_t, is similar to r_t^2 in the sense of being a measurement of volatility (just far more accurate), it is natural to expect that $\log x_t \approx \log h_t + g\left(z_t\right) + \text{error}_t$, for some function g. Because we may compute realized measures of volatility over a shorter period of time than the one spanned by the return (e.g. if we use only data from the trading session, which often excludes the overnight period), some flexibility in the specification may be required motivating the 'intercept' ξ_0 and the 'slope' φ_0. So long as $x_{0,t}$ is roughly proportional to $h_{0,t}$, we should expect $\varphi_0 \simeq 1$, and $\xi_0 < 0$, which is always the case empirically.

2.3. Conditional Model for an Individual Asset Return and its Realized Measures

To extend the framework to a joint model for the market returns/volatility and another asset's return/volatility and their correlation, we shall formulate a model for the time series associated with the individual asset, conditional on contemporaneous 'market' variables, i.e. a specification for $f\left(r_{1,t}, x_{1,t}, y_{1,t} | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right)$. We utilize a further decomposition of this conditional density, specifically

$$f\left(r_{1,t}, x_{1,t}, y_{1,t} | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) = f\left(r_{1,t} | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) f\left(x_{1,t}, y_{1,t} | r_{1,t}, r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right)$$

The first part, $f\left(r_{1,t} | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right)$, is modeled with

$$r_{1,t} = \mu_1 + \sqrt{h_{1,t}} z_{1,t}, \tag{6}$$

where the dependence on $\left(r_{0,t}, x_{0,t}\right)$ operates through, $\rho_{1,t} = \text{cov}\left(z_{0,t}, z_{1,t} | \mathcal{F}_{t-1}\right)$, the conditional correlation. This reveals the 'factor' structure, since we have

$$z_{1,t} = \rho_{1,t} z_{0,t} + \sqrt{1 - \rho_{1,t}^2} w_{1,t}$$

where $w_{1,t} = (z_{1,t} - \rho_{1,t} z_{0,t}) / \sqrt{1 - \rho_{1,t}^2}$ has mean zero, unit variance and is uncorrelated with $z_{0,t}$. Hence the studentized returns for the individual asset is a linear combination of the studentized market return and the idiosyncratic component $w_{1,t}$, where the relative weighting (defined by $\rho_{1,t}$) is time varying.

To complete this part of the model, we need to specify the dynamics for $h_{1,t}$ and $\rho_{1,t}$. For $h_{1,t}$ we use the GARCH equation:

$$\log h_{1,t} = a_1 + b_1 \log h_{1,t-1} + c_1 \log x_{1,t-1} + d_1 \log h_{0,t} + \tau_1 (z_{1,t-1}) \tag{7}$$

which only differs from equation (4) by the presence of the term, $d_1 \log h_{0,t}$. Recall that $h_{0,t}$ is $\mathcal{F}_{t-1}$-measurable, so that the presence of $h_{0,t}$ on the right-hand side does not contradict our definition of $h_{1,t}$. The parameter d_1 can be interpreted as a spillover effect that measures the extent to which the market's volatility affects the volatility of the individual asset while accounting for the asset-specific volatility dynamics.

For the dynamic modeling of $\rho_{1,t}$ we shall use the Fisher transformation (also known as the inverse hyperbolic tangent, arctanh), $\rho \mapsto F(\rho) \equiv \frac{1}{2} \log \frac{1+\rho}{1-\rho}$, which is a one-to-one mapping from $(-1, 1)$ into $\mathbb{R}$. The GARCH equation for the transformed correlations is given by

$$F(\rho_{1,t}) = a_{10} + b_{10} F(\rho_{1,t-1}) + c_{10} F(y_{1,t-1})$$

Finally, to specify the conditional density, $f(x_{1,t}, y_{1,t} | r_{1,t}, r_{0,t}, x_{0,t}, \mathcal{F}_{t-1})$, for the last two realized measures we use the measurement equations

$$\log x_{1,t} = \xi_1 + \varphi_1 \log h_{1,t} + \delta_1(z_{1,t}) + u_{1,t} \tag{8}$$

$$F(y_{1,t}) = \xi_{10} + \varphi_{10} F(\rho_{1,t}) + v_{1,t} \tag{9}$$

These measurement errors will be assumed to be independent of $(z_{0,t}, z_{1,t})$, implying that the conditioning on $(r_{1,t}, r_{0,t})$ is captured by $\delta_1(z_{1,t})$. The three measurement errors are allowed to be correlated and we define their covariance matrix:

$$\Sigma = \text{var} \begin{pmatrix} u_{0,t} \\ u_{1,t} \\ v_{1,t} \end{pmatrix} = \begin{bmatrix} \sigma_{u_0}^2 & \sigma_{u_0,u_1} & \sigma_{u_0,v_1} \\ \cdot & \sigma_{u_1}^2 & \sigma_{u_1,v_1} \\ \cdot & \cdot & \sigma_{v_1}^2 \end{bmatrix}$$

We find significant correlation across all measurement errors in our empirical application.

2.4. The Extensions to Multiple Individual Assets

We have specified the model structure for a market return and a single individual asset (along with their corresponding realized volatility variables). Next, we discuss the extension to multiple assets. Fortunately, the existing structure is amendable to this extension, albeit some additional assumptions are needed before certain interpretations carry over to the general context. First, we need to redefine the natural filtration, $\mathcal{F}_t = \sigma(\mathcal{X}_t, \mathcal{X}_{t-1}, \ldots)$, to be defined by the full set of variables:

$$\mathcal{X}_t = (r_{0,t}, r_{1,t}, \ldots, r_{N,t}, x_{0,t}, x_{1,t}, \ldots, x_{N,t} y_{1,t} \ldots, y_{N,t})'$$

where N is the number of individual assets in our analysis. The conditional model for the individual asset is assumed to be invariant to this enhancement of the information set. This implicitly assumes that the dynamic variation in correlations between individual assets is fully explained by the individual

asset's correlation with the market return. Put differently: The variation in the $(N + 1) \times (N + 1)$ conditional covariance matrix is fully described by the $N + 1$ conditional variances and the N conditional correlations. This structure has testable implications that we return to in our empirical section.

In practice, the estimation proceeds by first estimating the model for the market data $(r_{0,t}, x_{0,t})$ and then estimating each conditional model for $(r_{i,t}, x_{i,t}, y_{i,t})$ separately for $i = 1, 2, \ldots, N$. This can be done for a very large number of assets. For instance, in the empirical analysis we estimated the realized beta GARCH model for about 600 assets.

For model diagnostics, in particular the validity of the single-factor structure, we define conditional studentized residuals:

$$\hat{w}_{i,t} = \frac{\hat{z}_{i,t} - \hat{\rho}_{i,t}\hat{z}_{0,t}}{\sqrt{1 - \hat{\rho}_{i,t}^2}}, \ i = 1, \ldots, N$$

So far the model structure has been silent about the dependence structure across the population equivalents of these residuals, and the same is true for the conditional error terms, $(u_{i,t}, v_{i,t}|u_{0,t})$, across individual assets. We cast light on this dependence structure in our empirical section.

3. ESTIMATION

In this section, we define the quasi log-likelihood function and exploit its structure to simplify the estimation problem. We have five observed variables $(r_{0,t}, x_{0,t}, r_{1,t}, x_{1,t}, y_{1,t})$ and we consider their joint density conditional on past information, $\mathcal{F}_{t-1}$. Without loss of generality we can decompose this 'joint' density as stated in equation (2), and, for the purpose of estimation, we adopt Gaussian specifications for the 'marginal' and 'conditional' densities, $f(r_{0,t}, x_{0,t}|\mathcal{F}_{t-1})$ and $f(r_{i,t}, x_{i,t}, y_{i,t}|r_{0,t}, x_{0,t}, \mathcal{F}_{t-1})$, $i = 1, \ldots, N$. Moreover, we assume that the studentized returns, $(z_{0,t}, z_{1,t})$, are independent of the error terms in the three measurement equations, $(u_{0,t}, u_{1,t}, v_{1,t})$. This enables us to decompose the quasi log-likelihood function into four terms as we discuss below.

3.1. The Marginal Model for Market Variables

The marginal model is essentially that of Hansen *et al.* (2012), which implicitly entails a further decomposition of the conditional density:

$$f(r_{0,t}, x_{0,t}|\mathcal{F}_{t-1}) = f_{r_0}(r_{0,t}|\mathcal{F}_{t-1}) f_{x_0|r_0}(x_{0,t}|r_{0,t}, \mathcal{F}_{t-1})$$

The two conditional densities are given from $r_{0,t} \sim N(\mu_0, h_{0,t})$ and $\log x_{0,t} \sim N(\xi_0 + \varphi_0 \log h_{0,t} + \tau_0(z_{0,t}), \sigma_{u_0}^2)$, which leads to the following two contributions to (minus two times) the log-likelihood function:

$$\ell_{z_0} = \sum_{t=1}^{T} \log h_{0,t} + \frac{(r_{0,t} - \mu_0)^2}{h_{0,t}} = \sum_{t=1}^{T} \log h_{0,t} + z_{0,t}^2,$$

$$\ell_{u_0} = \sum_{t=1}^{T} \log \sigma_{u_0}^2 + \frac{(\log x_{0,t} - \xi_0 - \varphi_0 \log h_{0,t} - \tau_0(z_{0,t}))^2}{\sigma_{u_0}^2} = \sum_{t=1}^{T} \log \sigma_{u_0}^2 + \frac{u_{0,t}^2}{\sigma_{u_0}^2}$$

3.2. The Conditional Model for Individual Assets

Next, we consider the likelihood contributions from the conditional model. The conditional model also permits a further decomposition of the conditional density:

$$
\begin{aligned}
f\left(r_{1,t}, x_{1,t}, y_{1,t} | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) = {} & f_{r_1 | r_0, x_0}\left(r_{1,t} | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) \\
& \times f_{x_1, y_1 | r_1, r_0, x_0}\left(x_{1,t}, y_{1,t} | r_{1,t}, r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right)
\end{aligned}
$$

The first term is the density of the individual asset's return conditional on the contemporaneous market variables (and the past). Due to the Gaussian specification we only need to derive the conditional mean and variance of $r_{1,t}$ in order to compute the appropriate likelihood term. The assumed independence between $\left(z_{0,t}, z_{1,t}\right)$ and $u_{0,t}$ and the i.i.d. assumptions imply that

$$
\mathrm{E}\left(g\left(r_{1,t}\right) | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) = \mathrm{E}\left(g\left(r_{1,t}\right) | z_{0,t}, u_{0,t}, \mathcal{F}_{t-1}\right) = \mathrm{E}\left(g\left(r_{1,t}\right) | r_{0,t}, \mathcal{F}_{t-1}\right)
$$

for any function g for which the conditional mean is well defined. Hence

$$
\mathrm{var}\left(r_{1,t} | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) = \mathrm{var}\left(r_{1,t} | r_{0,t}, \mathcal{F}_{t-1}\right) = h_{1,t} - \left(\rho_{1,t}\sqrt{h_{0,t}h_{1,t}}\right)^2 / h_{0,t} = \left(1 - \rho_{1,t}^2\right)h_{1,t}
$$

since $\mathrm{cov}\left(r_{1,t}, r_{0,t} | \mathcal{F}_{t-1}\right) = \rho_{1,t}\sqrt{h_{0,t}h_{1,t}}$. Next, the conditional mean of $r_{1,t}$ is

$$
\begin{aligned}
\mathrm{E}\left(r_{1,t} | r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) &= \mu_1 + \beta_{1,t}\left(r_{0,t} - \mu_0\right) = \mu_1 + \frac{\rho_{1,t}\sqrt{h_{0,t}h_{1,t}}}{h_{0,t}}\left(r_{0,t} - \mu_0\right) \\
&= \mu_1 + \rho_{1,t}\sqrt{h_{1,t}}z_{0,t}
\end{aligned}
$$

The contribution to (minus two times) the log-likelihood function from this conditional density is

$$
\ell_{z_1 | z_0} = \sum_{t=1}^{T} \log\left[\left(1 - \rho_{1,t}^2\right)h_{1,t}\right] + \frac{\left(r_{1,t} - \mu_1 - \rho_{1,t}\sqrt{h_{1,t}}z_{0,t}\right)^2}{\left(1 - \rho_{1,t}^2\right)h_{1,t}}
$$

The last likelihood term, $\ell_{u_1, v_1 | u_0}$, which relates to the two measurement equations is associated with the conditional density, $f_{x_1, y_1 | r_1, r_0, x_0}\left(x_{1,t}, y_{1,t} | r_{1,t}, r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right)$. First, we note that the conditional distribution of $\left(u_{1,t}, v_{1,t}\right)$ given $\left(u_{0,t}, z_{0,t}, z_{1,t}\right)$ is Gaussian with mean

$$
\begin{pmatrix} \sigma_{u_1, u_0} / \sigma_{u_0}^2 \\ \sigma_{v_1, u_0} / \sigma_{u_0}^2 \end{pmatrix} u_{0,t}
$$

and variance

$$
\Omega = \begin{bmatrix} \sigma_{u_1}^2 & \sigma_{u_1, v_1} \\ \cdot & \sigma_{v_1}^2 \end{bmatrix} - \begin{bmatrix} \sigma_{u_1, u_0} \\ \sigma_{v_1, u_0} \end{bmatrix} \frac{1}{\sigma_{u_0}^2} \begin{bmatrix} \sigma_{u_0, u_1} & \sigma_{u_0, v_1} \end{bmatrix}
$$

Thus it does not depend on $\left(z_{0,t}, z_{1,t}\right)$ owing to the assumed independence. The implication is that

$$
f_{x_1, y_1 | r_1, r_0, x_0}\left(x_{1,t}, y_{1,t} | r_{1,t}, r_{0,t}, x_{0,t}, \mathcal{F}_{t-1}\right) = f_{x_1, y_1 | r_1, r_0, x_0}\left(x_{1,t}, y_{1,t} | u_{0,t}, \mathcal{F}_{t-1}\right)
$$

and that the last term in (minus two times) the log-likelihood is given by

$$\ell_{u_1,v_1|u_0} = \sum_{t=1}^{T} \log \det \Omega + U_{1,t}' \Omega^{-1} U_{1,t}$$

where we have defined

$$U_{1,t} = \begin{pmatrix} u_{1,t} \\ v_{1,t} \end{pmatrix} - \begin{pmatrix} \sigma_{u_1,u_0}/\sigma_{u_0}^2 \\ \sigma_{v_1,u_0}/\sigma_{u_0}^2 \end{pmatrix} u_{0,t}$$

The parameters are now estimated by maximizing $\ell = -\frac{1}{2} \left(\ell_{z_0} + \ell_{u_0} + \ell_{z_1|z_0} + \ell_{u_1,v_1|u_0} \right)$. Fortunately, the structure of this likelihood permits a number of simplifications, which greatly simplify the computational burden in the estimation. These simplifications are detailed in Appendix B (supporting information).

4. FORECASTING

In this section we discuss how multi-step predictions of volatilities and correlations, as well as return density forecasts, can be obtained with our model. Denote $\tilde{h}_{0,t} \equiv \log h_{0,t}$, $\tilde{h}_{i,t} \equiv \log h_{i,t}$ and $\tilde{\rho}_{i,t} \equiv F(\rho_{i,t})$. Point forecasts turn out to be very easy to obtain owing to the fact that the vector $\left(\tilde{h}_{0,t}, \tilde{h}_{i,t}, \tilde{\rho}_{i,t} \right)$ can be represented as a VARMA(1,1) system. Substituting each of the measurement equations (5), (8) and (9) in the equations for the corresponding conditional moments and imposing, $\varphi_0 = 1$, one obtains

$$
\begin{aligned}
\tilde{h}_{0,t+1} &= a_0 + c_0 \xi_0 + (b_0 + c_0 \varphi_0) \tilde{h}_{0,t} + c_0 \delta_0 (z_{0,t}) + \tau_0 (z_{0,t}) + c_0 u_{0,t} \\
\tilde{h}_{i,t+1} &= a_i + c_i \xi_i + (b_i + c_i \varphi_i) \tilde{h}_{i,t} + d_i \tilde{h}_{0,t+1} + c_i \delta_i (z_{i,t}) + \tau_i (z_{i,t}) + c_i u_{i,t} \quad (10)\\
\tilde{\rho}_{i,t+1} &= a_{i0} + c_{i0} \xi_{i0} + (b_{i0} + c_{i0} \varphi_{i0}) \tilde{\rho}_{i,t} + c_{i0} v_{i,t}
\end{aligned}
$$

Let $V_t = \left(\tilde{h}_{0,t}, \tilde{h}_{i,t}, \tilde{\rho}_{i,t} \right)'$, then by substituting the equation for $\tilde{h}_{0,t+1}$ into that for $\tilde{h}_{i,t+1}$, one can show that

$$V_{t+1} = C + A V_t + B \varepsilon_t$$

where $\varepsilon_t = \left(\delta_0 (z_{0,t}), \tau_0 (z_{0,t}), \delta_i(z_{i,t}), \tau_i (z_{i,t}), u_{0,t}, u_{i,t}, v_{i,t} \right)'$ and

$$
C = \begin{bmatrix} a_0 + c_0 \xi_0 \\ a_1 + c_i \xi_i + d_i (a_0 + c_0 \xi_0) \\ a_{i0} + c_{i0} \xi_{i0} \end{bmatrix}, \qquad A = \begin{bmatrix} b_0 + c_0 \varphi_0 & 0 & 0 \\ d_i (b_0 + c_0 \varphi_0) & b_i + c_i \varphi_i & 0 \\ 0 & 0 & b_{i0} + c_{i0} \varphi_{i0} \end{bmatrix},
$$

$$
B = \begin{bmatrix} c_0 & 1 & 0 & 0 & c_0 & 0 & 0 \\ d_i c_0 & d_i & c_i & 1 & d_i c_0 & c_1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & c_{i0} \end{bmatrix}
$$

The innovation process, ε_t, is a martingale difference sequence but is slightly heterogeneous, because time-variation in $\rho_{i,t}$ induces time variation in the distribution of ε_t. It follows that $E(V_{t+k}|V_t) = A^k V_t + \sum_{j=0}^{k-1} A^j C$, which can be used to produce a k-step-ahead forecast of V_{t+k}.

Forecast of the conditional distribution of $V_{t+k}|\mathcal{F}_t$, which can be used to deduce unbiased forecasts of the non-transformed variables, e.g. $h_{0,t} = \exp\left(\tilde{h}_{0,t}\right)$, can be obtained by simulation or bootstrap methods. In the simulation approach, we first generate

$$\eta_t := \begin{pmatrix} z_{0,t} \\ \tilde{z}_{i,t} \\ u_{0,t} \\ u_{i,t} \\ v_{i,t} \end{pmatrix} \sim N_5\left(0, \begin{bmatrix} I_2 & 0 \\ 0 & \Sigma \end{bmatrix}\right), \qquad t = 1, \dots, N$$

Given an initialization for $\rho_{i,0}$, one can produce the entire time series $\{\tilde{\rho}_{i,t}\}$ from $\{v_{i,t}\}$ using equation (10). Next one can define $z_{i,t} = \rho_{i,t}z_{0,t} + \sqrt{1 - \rho_{i,t}^2}\,w_{i,t}$, which has the proper correlation with $z_{0,t}$, and thus finally ε_t can be computed. The alternative approach is to bootstrap the residuals, $(\hat{\eta}_1, \dots, \hat{\eta}_N)$, rather than simulating them from the Gaussian distribution. Both approaches were explored in Lunde and Olesen (2013), who produced 1-, 5- and 20-day-ahead predictions using the realized exponential GARCH model. In their application, which used returns on energy forwards, the out-of-sample losses for the two approaches (simulations or bootstrap) were indistinguishable, with either method being clearly superior to the forecasts produced by the conventional EGARCH models. In the present context, where the system may be large, the bootstrap approach does have a clear advantage, in that it does not require one to formulate explicit assumptions about the correlation structure of the $w_{i,t}$-variables, nor specify the correlation structure of measurement errors across individual assets.

5. EMPIRICAL ANALYSIS

5.1. Data Description

The model is estimated for a large cross-section of assets. We included any asset that was a constituent of the S&P 500 index at some point between 19 January 2006 and 25 June 2010, albeit excluding assets for which we had fewer than 1000 daily observations during our sample period from 3 January 2002 to the end of 2009. The data were obtained by merging information from the TAQ dataset and the CRSP daily stock files (see Appendix A (supporting information) for further details), and resulted in a total of 594 time series with distinct CRSP permanent company numbers (PERMNOs) and a sample size that ranged from 1000 to 2008 observations for each of the individual stocks.

The high-frequency transaction data are cleaned according to the filtering algorithm described in Barndorff-Nielsen *et al.* (2009), and the multivariate realized kernel by Barndorff-Nielsen *et al.* (2011) is used as our realized measures of volatility and co-volatility. We use the exchange traded fund, SPY, as a proxy for the market index in our empirical analysis, making the total number of assets in our analysis 595.[1]

5.2. Empirical Results

A summary of the estimation results is presented in Table I and Figure 1.[2] The first row in Table I contains the estimates for the marginal model for the market return, as defined by equations (3-5), and

---

[1] The model takes about 3 seconds to estimate for each asset using a workstation (XEON 3.07 GHz) when the simplifications detailed in Appendix B (supporting information) are implemented. The structure is such that the conditional models (one for each individual asset) can be estimated separately, so that one can take advantage of parallel computing.

[2] The results reported are the estimates when imposing the restrictions $\varphi_i = 1, i = 0, 1, \dots, N$, which did not result in a significant reduction of the log-likelihood function; see Hansen and Huang (2012) for a discussion on this. The initial values for the latent variables, $h_{0,t}$, $\tilde{h}_{i,t}$ and $\rho_{i,t}$, are treated as unknown parameters.

Table I. Parameter estimates of the realized beta GARCH model

	Volatility parameters										Correlation parameters				
SPY	μ_0 μ_i	a_0 a_i	b_0 b_i	c_0 c_i	d_i	τ_{01} τ_{i1}	τ_{02} τ_{i2}	ξ_0 ξ_i	δ_{01} δ_{i1}	δ_{02} δ_{i2}	a_{i0}	b_{i0}	c_{i0}	ξ_{i0}	φ_{i0}
	0.001	0.175	0.601	0.377		-0.092	0.020	-0.469	-0.104	0.018					
Mean	0.023	0.208	0.585	0.335	0.045	-0.041	0.010	-0.346	-0.030	0.064	0.036	0.708	0.285	-0.099	0.982
Median	0.018	0.193	0.592	0.340	0.034	-0.041	0.009	-0.352	-0.032	0.063	0.036	0.701	0.278	-0.066	0.915
Min	-0.157	-0.024	0.350	0.093	-0.023	-0.084	-0.019	-0.963	-0.100	0.011	-0.103	0.536	0.027	-1.661	0.420
Min(1)	-0.130	-0.011	0.354	0.146	-0.019	-0.080	-0.018	-0.895	-0.091	0.011	-0.097	0.544	0.032	-1.522	0.445
Min(2)	-0.121	-0.000	0.381	0.166	-0.016	-0.080	-0.014	-0.880	-0.088	0.012	-0.096	0.549	0.043	-1.516	0.488
Min(3)	-0.082	0.006	0.382	0.176	-0.014	-0.078	-0.012	-0.860	-0.085	0.018	-0.095	0.556	0.048	-1.379	0.505
Min(4)	-0.079	0.017	0.391	0.187	-0.013	-0.078	-0.010	-0.779	-0.080	0.018	-0.089	0.558	0.051	-0.955	0.514
Min(5)	-0.077	0.023	0.397	0.204	-0.012	-0.076	-0.010	-0.760	-0.080	0.020	-0.086	0.559	0.051	-0.952	0.518
1%	-0.073	0.023	0.402	0.204	-0.012	-0.076	-0.009	-0.759	-0.079	0.021	-0.082	0.562	0.060	-0.942	0.529
5%	-0.042	0.063	0.455	0.256	-0.002	-0.065	-0.005	-0.577	-0.060	0.033	-0.035	0.601	0.121	-0.407	0.603
95%	0.102	0.402	0.701	0.407	0.125	-0.016	0.028	-0.094	0.002	0.095	0.107	0.834	0.484	0.148	1.463
99%	0.140	0.522	0.756	0.442	0.163	-0.006	0.035	-0.016	0.016	0.112	0.139	0.887	0.600	0.227	2.539
Max(-5)	0.149	0.523	0.762	0.447	0.164	-0.005	0.036	-0.012	0.019	0.113	0.140	0.890	0.600	0.229	2.560
Max(-4)	0.153	0.525	0.763	0.447	0.166	-0.005	0.037	-0.009	0.019	0.113	0.143	0.905	0.602	0.238	2.599
Max(-3)	0.164	0.537	0.769	0.449	0.173	-0.005	0.037	-0.005	0.023	0.116	0.157	0.925	0.615	0.239	2.897
Max(-2)	0.168	0.565	0.789	0.459	0.175	0.000	0.038	-0.002	0.026	0.120	0.185	0.939	0.631	0.242	2.988
Max(-1)	0.220	0.649	0.828	0.473	0.190	0.001	0.039	0.056	0.026	0.128	0.193	0.958	0.640	0.247	3.331
Max	0.229	0.800	0.887	0.483	0.203	0.020	0.048	0.082	0.037	0.154	0.205	0.963	0.647	0.252	3.956

Note: A summary of the estimated parameters. First row has the estimates from the market return model (3)–(5), and information about the cross-section of estimates for the 594 conditional models (6)–(9) are presented in the remaining rows.
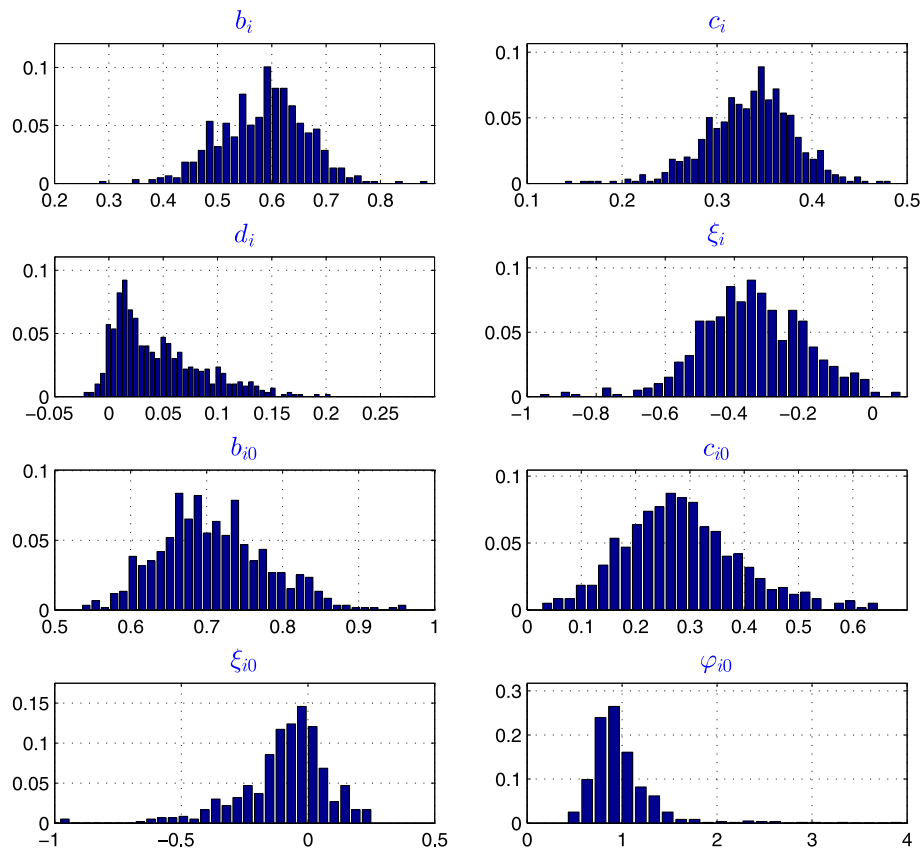
Figure 1. Histograms of the 594 parameter estimates for some selected parameters

the rest of the table presents a summary of the estimation results for the 594 conditional models, each defined by equations: (6)–(9).[3] To conserve space the cross-sectional statistics for the estimates of Σ are omitted, but some selected estimates of Σ will be presented in Table II.

The parameter c_i, which captures the effect of the lagged realized measure on the conditional variance, is large and significant, whereas the GARCH parameter, b_i, is much smaller than is usually the case for conventional GARCH models. The reason is that the realized measure is far more informative about volatility than the squared return, which makes the model far more adaptive to abrupt changes in volatility, which in turn leads to a better empirical fit and more accurate forecast. The negative estimates of τ_{i1} and positive estimates of τ_{i2} indicate the presence of a leverage effect; see Hansen *et al.* (2012) for the relation of these leverage functions to the news impact curve. Examining the parameters of the measurement equation, we find that ξ_i is negative. This is to be expected because the realized measures are computed over the open-to-close period, which only captures a fraction of daily (close-to-close) volatility. The conditional model for the individual stocks has the additional parameter, d_i, in the GARCH equation. This parameter measures the spillover effect from market volatility to individual stock volatility. The mean and the median of this coefficient is positive, and so are the vast majority of the individual estimates. Altogether this shows that market volatility tends to have a positive contemporaneous effect on individual asset volatility.

---

[3] Note that a row in Table I does not present the estimates for a particular stock. For example, the 1% quantile of b_0 and c_0 may not be estimates for the same asset.

Table II. Selected point estimates of the measurement error variance

CVX			UTX			WMB		
0.113	*0.493*	*0.309*	0.114	*0.434*	*0.391*	0.114	*0.409*	*0.436*
0.066	0.159	*0.016*	0.064	0.194	*0.117*	0.065	0.221	*0.118*
0.011	0.001	0.012	0.019	0.007	0.020	0.023	0.009	0.025

EK			SNV			MSFT		
0.114	*0.524*	*0.376*	0.114	*0.428*	*0.365*	0.113	*0.356*	*0.372*
0.068	0.149	*0.157*	0.060	0.170	*0.046*	0.058	0.236	*0.097*
0.019	0.009	0.022	0.018	0.003	0.021	0.018	0.007	0.021

Note: The estimated measurement error variance matrix, $\Sigma = \mathrm{var}(u_{0,t}, u_{i,t}, v_{i,t})$, for six selected assets: CVX, Chevron; UTX, United Technologies; WMB, Williams Companies; EK, Eastman Kodak; SNV, Synovus Financial; MSFT, Microsoft. The numbers in italics (upper right triangular) are correlations.
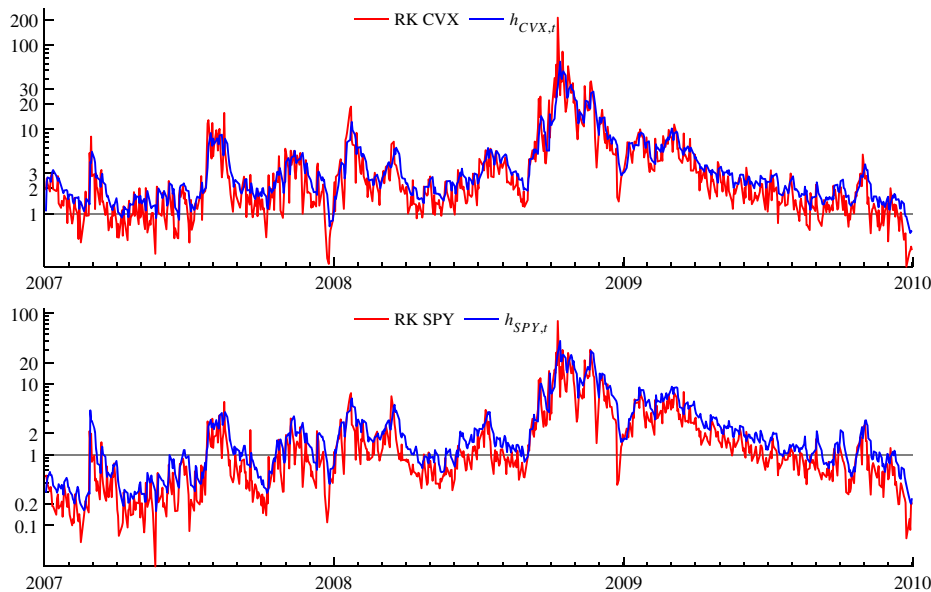

Figure 2. Realized kernel (RK) variance and conditional variance of CVX (upper panel) and SPY (lower panel) over the period 2007–2009

The cross-sectional variation of parameter estimates are presented in the histogram plots in Figure 2. Both Table I and Figure 1 show that the parameter estimates are quite stable in our cross-section of stocks. Only $\hat{\varphi}_{i0}$ is estimated to have an extreme value in some cases, but even in these cases we have verified that the estimated conditional variances and correlations are in agreement with their corresponding realized measures.

Table II presents the estimates of Σ for six selected assets. The upper left element in these matrices is the estimated variance of $u_{0,t}$ in the measurement equation for the realized measures associated with the market return. This point estimate varies slightly across the six matrices due to variation in the sample period for the six assets. Note that the measurement error variance for the individual assets tends to be larger than that of the market. This is to be expected because the realized measure for the market return is based on a larger number of high-frequency data. Also note that there is substantial correlation across the measurements error, in particular for the realized measures of volatility.
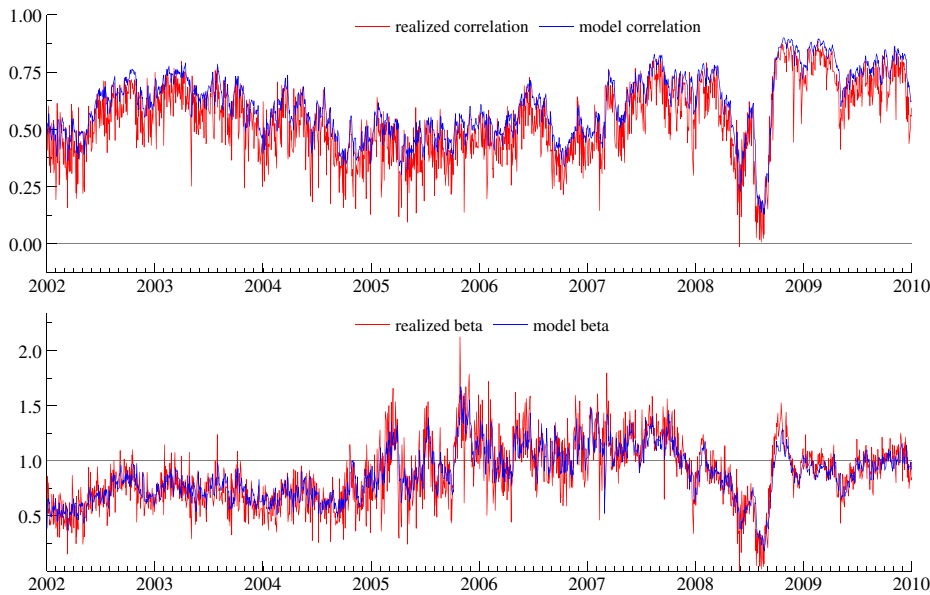
Figure 3. Realized and conditional correlation (upper panel) and beta (lower panel) of CVX (permno number 14,541) for the period 2002–2009

In Figure 2 we present the realized variance of CVX and SPY against the model-implied conditional variance. Clearly, the conditional variance tracks the realized series closely but has less high-frequency variation. Naturally, this relation is largely imposed by the model's structure, because the measurement equation implies a (noisy) relationship between the conditional variance and realized measure. The apparent downward bias of the realized measure is due to the fact that it is computed over a fraction of the day (roughly 6.5 hours where assets are actively traded). This aspect of the realized measures explains why the coefficients ξ_0 and ξ_1 are negative.

We turn next to the model-implied betas given by

$$\hat{\beta}_t = \hat{\rho}_t \sqrt{\hat{h}_{1,t}/\hat{h}_{0,t}} \tag{11}$$

where $\hat{\rho}_t = \rho_t\left(\hat{\theta}\right)$ and $\hat{h}_{i,t} = h_{i,t}\left(\hat{\theta}\right)$, $i = 0, 1$, denote the estimated quantities. The time series can be contrasted to the realized betas:

$$\tilde{\beta}_t = y_{1,t} \sqrt{x_{1,t}/x_{0,t}} \tag{12}$$

that are computed exclusively from high-frequency data on day t.

The model-implied betas take into account the presence of measurement error in the realized quantities as well as the dynamic linkages between realized measures and conditional moments. To get an idea of the time variation of $\hat{\beta}_t$ in our model compared to its raw realized counterpart, we continue with our previous example, and present graphic results for the realized and the conditional beta and correlation of CVX in Figure 3. The correlation changes rapidly during the sample period, which carries over to the systematic risk of CVX, as defined by its beta. In fact, the beta for CVX ranges from about 0 to more than 1.5 over this period.

In Figure 4 we present quantile time series plots of the cross-sectional variation in the conditional correlation and beta during the financial crisis, where the time of the collapse of Lehman Brothers is
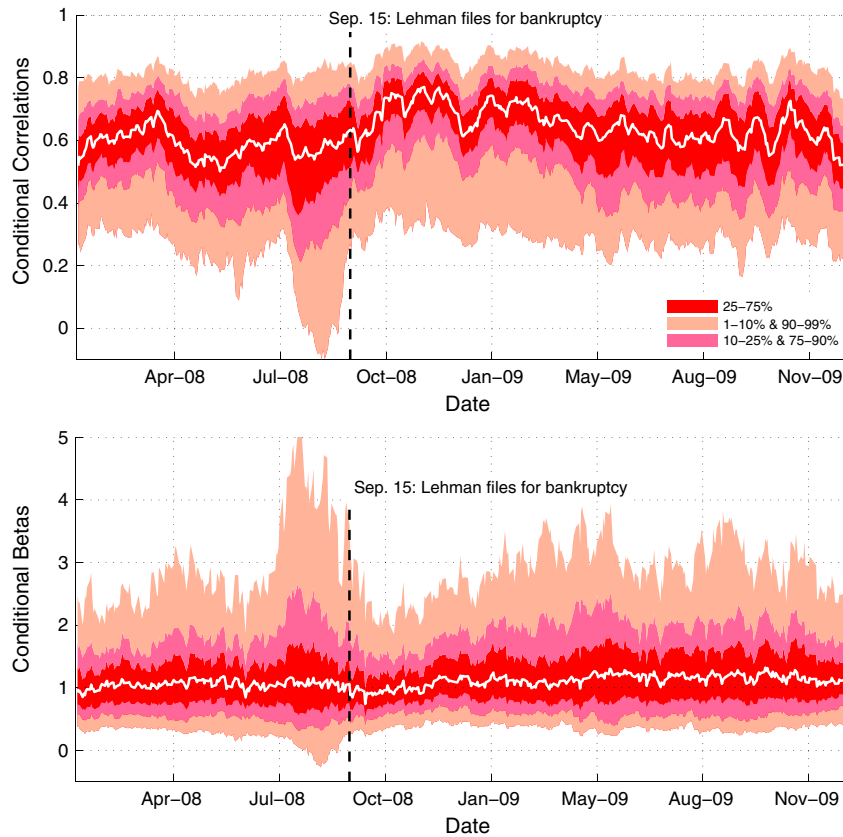
Figure 4. Quantile time series plot of conditional realized GARCH correlations for the period June 2008 to December 2009

clearly identified. It is, perhaps, the period leading up to the collapse of Lehman Brothers that stands out most. On 15 July the SEC temporarily prohibited naked short selling in the securities of Fannie Mae and Freddie Mac. The time of this announcement coincides with some major changes in the cross-sectional distribution of correlations and betas, although we do not claim any causal relation in this matter. In the subsequent period correlations decreased (on average) and the cross-sectional distribution became increasingly left skewed. This might suggest that assets became somewhat more susceptible to idiosyncratic shocks and less so to market-wide shocks. It is therefore perhaps surprising that the distribution of conditional betas became more right skewed. The explanation is that individual asset volatility increased relatively more than market volatility, which more than offset the reduction in average correlation. The mechanics of this are easily understood from the definition of the conditional beta, $\beta_{i,t} = \rho_{i,t}\sqrt{h_{i,t}/h_{0,t}}$. After this initial chaotic period correlations started to increase and the variation in betas decreased. Eventually, correlations peaked around mid November with a median value of over 70%—well above the 55% value at the beginning of June.

It is important to understand that the high degree of variation that we find in the betas should be attributed to variation in the true conditional betas, rather than an artifact of the variation in the realized quantities. The reason is that the variation in the beta produced by the model is driven by features of daily returns—features that the model attempts to explain to the extent it is possible. From the maximization of the likelihood function it is evident that the realized measures are useful predictors of the covariance structure of daily returns. Conversely, had there been little variation in the covariance

structure of daily returns, then the likelihood estimation would not produce estimates that caused the realized measures to induce a high degree of variation in the conditional covariance structure. The finding that correlations and betas exhibit a high degree of variation is an important observation. The conventional approach to estimating dynamic betas typically involves rolling window regression methods using daily returns. This approach induces some degree of smoothing, so it is not surprising that the regression-based estimates result in a variation of a much smaller magnitude. Based on rolling window estimated betas of this kind, Lewellen and Nagel (2006) concluded that time variation in beta is insufficient to explain certain asset pricing anomalies. Given the large variation we observed in the systematic risk of individual companies, it could be interesting to revisit this question.

To illustrate the degree of variation in beta, Figure 5 presents the time series of conditional betas for four selected stocks during the second half of 2008. An interesting example is Williams Companies (WMB), which moved from the lower 10% to the upper 10% in the fall of 2008. The example of SNV shows how some financial companies became relatively more risky as the financial crisis approached in the early fall of 2008. Finally, we included UTX and EK to show that the betas of some companies are relatively stable.

5.3. Beta Comparison

An advantage of the realized beta model is that it readily produces the time series of the betas, as we illustrated in Figure 5. Similar quantities can be obtained by multivariate GARCH-type models that do not utilize realized measures. A prime example of such a model is the dynamic conditional correlation (DCC) model by Engle (2002a).
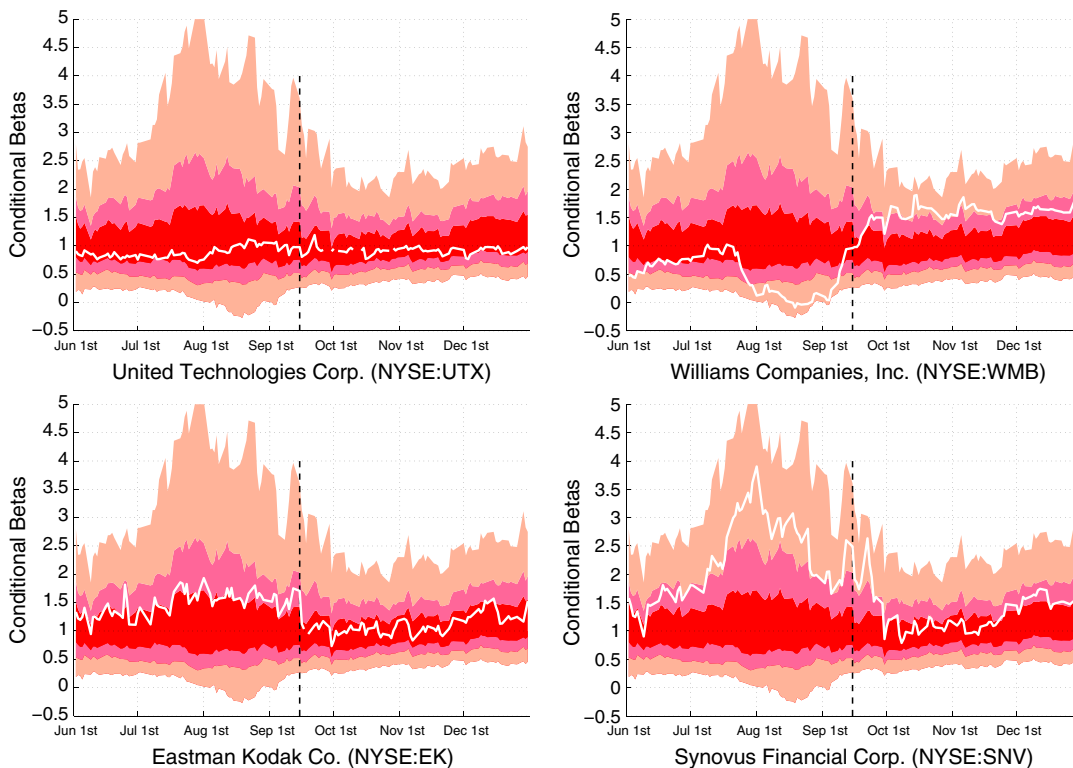

Figure 5. Quantile time series plot of conditional realized GARCH betas for the period June 2008 to December 2008

In this section we compare the betas produced by the realized beta GARCH model with those produced by the DCC model, as well as a static model that has β constant. The comparison is done using two methods: first, we use a regression-based comparison that was proposed by Engle (2012). Second, we compare the betas using a beta hedging-tracking exercise.

The regression-based comparison is based on the auxiliary variables

$$Z_{i,t}^{\mathrm{RBG}} = \hat{\beta}_{i,t}^{\mathrm{RBG}} r_{0,t}, \quad Z_{i,t}^{\mathrm{DCC}} = \hat{\beta}_{i,t}^{\mathrm{DCC}} r_{0,t}, \quad Z_{i,t}^{\mathrm{CAPM}} = \hat{\beta}_{i} r_{0,t}$$

and the estimated regression model

$$r_{it} = \alpha + \delta_i^{\mathrm{RBG}} Z_{i,t}^{\mathrm{RBG}} + \delta_i^{\mathrm{DCC}} Z_{i,t}^{\mathrm{DCC}} + \delta_i^{\mathrm{CAPM}} Z_{i,t}^{\mathrm{CAPM}} + \varepsilon_{i,t} \tag{13}$$

for each of the assets. A specification that yields an ideal time series for β_{it} would correspond to its δ-coefficient being one, while the two other δ-coefficients would be zero. Engle (2012) compared the beta of a DCC model to a constant beta using in-sample regressions, and found strong support for the DCC specification. In our comparison, we add the third time series of betas (that of the realized beta GARCH model) to the comparison, and make the regression-based comparison using out-of-sample data. This is done as follows: First, we estimated each of the three models using data up until the end of 2007. The estimated models are then used to compute betas out-of-sample, from which we construct the auxiliary variables. For the static model, $\hat{\beta}_i$ is simply obtained by regressing individual returns on the market return and a constant over the in-sample period.

We estimate equation (13) by least squares and test the following three hypotheses:

$$
\begin{aligned}
H_{\mathrm{RBG}} &: \delta_i^{\mathrm{RBG}} = 1, \ \delta_i^{\mathrm{DCC}} = \delta_i^{\mathrm{CAPM}} = 0, \\
H_{\mathrm{DCC}} &: \delta_i^{\mathrm{DCC}} = 1, \ \delta_i^{\mathrm{RBG}} = \delta_i^{\mathrm{CAPM}} = 0, \\
H_{\mathrm{CAPM}} &: \delta_i^{\mathrm{CAPM}} = 1, \ \delta_i^{\mathrm{RBG}} = \delta_i^{\mathrm{DCC}} = 0
\end{aligned}
$$

using robust inference in the sense of White (1980). In addition, we make a direct comparison between the realized beta GARCH model and the DCC model by excluding $Z_{i,t}^{\mathrm{CAPM}}$ from the regression in equation (13) and test the hypotheses

$$H_{\mathrm{RBG}}' : \delta_i^{\mathrm{DCC}} = 0 \quad \text{and} \quad H_{\mathrm{DCC}}' : \delta_i^{\mathrm{RBG}} = 0$$

Our second comparison of the betas is based on a beta hedging-tracking exercise. Here we compute the time series, $\left(r_{it} - Z_{i,t}^{\mathrm{RBG}}\right)$, $\left(r_{it} - Z_{i,t}^{\mathrm{DCC}}\right)$ and $\left(r_{it} - Z_{i,t}^{\mathrm{CAPM}}\right)$ for each asset, and compute their sample variances. In this exercise we seek the one with the smallest variance. For this multiple-comparisons problem we apply the model confidence set (MCS) methodology of Hansen *et al.* (2011), to identify a specification whose variance is significantly smaller than that of the competitors.

The out-of-sample comparison is based on an extended sample that spans 4 years from the beginning of 2008 to the end of 2011. Our results are based on the 450 assets for which we had data for the full sample period (in-sample as well as out-of-sample).

The results are summarized in Table III. Panels A and B show results from the regression-based comparisons, and panel C shows MCS results from the beta hedging-tracking exercise. Panel A presents the average estimated δ-coefficients, from the regression (13), and four subset regressions, also with statistics that summarize the cross-sectional variation in these estimates, across the 450 assets. The beta deduced from the realized beta model has, on average, the largest weight. In the

Table III. Beta comparison

	Full model			Excluding $Z_{i,t}^{CAPM}$		Single		
	δ_i^{RBG}	δ_i^{DCC}	δ_i^{CAPM}	δ_i^{RBG}	δ_i^{DCC}	δ_i^{RBG}	δ_i^{DCC}	δ_i^{CAPM}
Panel A: Summary of parameter estimates								
Mean	0.640	0.067	0.318	0.884	0.183	1.050	1.128	1.126
SD	0.517	0.335	0.498	0.363	0.336	0.220	0.134	0.408
1%	−0.502	−0.781	−0.950	0.104	−0.574	0.640	0.785	0.453
5%	−0.093	−0.514	−0.518	0.285	−0.367	0.731	0.924	0.541
Median	0.609	0.087	0.320	0.845	0.177	1.014	1.112	1.070
95%	1.655	0.603	1.111	1.531	0.779	1.477	1.366	1.873
99%	1.980	0.828	1.555	1.882	1.037	1.729	1.492	2.214
Panel B: Rejection frequencies								

Panel B: H_{RBG}: 0.72 H_{DCC}: 0.94 H_{CAPM}: 0.92 H_{RBG}': 0.40 H_{DCC}': 0.95

Panel C: Frequency in model confidence set
RBG: 0.88   DCC: 0.53   CAPM: 0.45

Note: Comparisons of the betas of the RBG, DCC and CAPM models. Panel A presents summary statistics of the estimated regression coefficients in equation (13) and four submodels (one that excludes $Z_{i,t}^{CAPM}$ and the three models that only include one of three regressors. The rejection frequencies for five hypotheses are given in panel B, and panel C reports the frequency (across assets) that each of the models are included in the MCS in the beta hedging exercise.

regressions using the three auxiliary variables, the DCC model gets little weight. In fact, the average coefficient is much smaller than that of the CAPM model, even though the DCC is better than the CAPM in a direct comparison (see Engle, 2012). This is explained by a crowding-out effect that occurs when the RBG regressor is included in the analysis. The middle part of panel A has the estimates from the direct comparison of the RBG and DCC betas, which also points to the RBG betas being superior to those of the DCC. This is confirmed in the results for the hypotheses tests given in panel B. The hypotheses that the DCC or the CAPM produce ideal betas, H_{DCC} and H_{CAPM}, are almost always rejected. The hypothesis that the RBG model produces ideal betas, H_{RBG}, is also rejected in the majority of cases, but cannot be rejected for 28% of the assets. In the direct comparison between RBG and DCC we reject the H_{DCC}' for the vast majority of assets; however, H_{RBG}' cannot be rejected for 60% of the assets. Thus for 60% of the assets there is no significant evidence that $\hat{\beta}_{i,t}^{DCC} r_{0,t}$ can explain variation in $r_{i,t}$, that is not already explained by $\hat{\beta}_{i,t}^{RBG} r_{0,t}$.

The general conclusion is further corroborated by the results in panel C. Here we report the frequency (across assets) by which the tracking error associated with each of the betas is in 90% MCS. The MCS is constructed such that it contains the competitor with the smallest tracking error with a probability no less than 90%. The DCC is found to be significantly inferior for 47% of the assets, and the CAPM is significantly inferior for 55% of the assets. Again, the RBG does better, and there are only 12% of the assets for which the RBG is found to be significantly inferior, which is only slightly more than can be attributed to Type 1 error.

5.4. Residual Correlations and Test for Constant Correlations

The realized beta GARCH model implies that the correlation between the individual studentized returns, z_{it} and z_{jt}, is time varying. Recall the decomposition

$$z_{i,t} = \rho_{i,t} z_{0,t} + z_{i,t} - \rho_{i,t} z_{0,t} = \rho_{i,t} z_{0,t} + \sqrt{1 - \rho_{i,t}^2} w_{i,t}$$

where $w_{i,t}$ and $z_{0,t}$ are uncorrelated, both have mean zero and unit variance, and in the likelihood analysis we modeled both as standard Gaussian random variables. It follows that

$$\mathrm{corr}(z_{i,t}, z_{j,t}) = \rho_{i,t}\rho_{j,t} + \sqrt{\left(1 - \rho_{i,t}^2\right)\left(1 - \rho_{j,t}^2\right)}\mathrm{E}\left(w_{i,t}w_{j,t}\right)$$

which is time varying unless $\mathrm{E}\left(w_{i,t}w_{j,t}\right)$ behaves in a rather unlikely way that offsets the variation in $\rho_{i,t}$ and $\rho_{j,t}$. We have not stated explicit assumptions about the correlation, $\mathrm{E}\left(w_{i,t}w_{j,t}\right)$, which induces additional dependence between $z_{i,t}$ and $z_{j,t}$, beyond that inherited from their correlations with the market return. This additional channel for dependence is ignored in our estimation (in order to make the estimation of large systems feasible). A non-zero correlation between $w_{i,t}$ and $w_{j,t}$ is evidence that the realized beta GARCH model does not fully characterize the complete system, so that the estimated model will need to be enhanced to capture such effects. It would also suggest that the estimation is inefficient to some extent, albeit this is to be expected with a relatively simple estimation procedure in a highly complex model.

In this section we study the magnitude of $\mathrm{E}\left(w_{i,t}w_{j,t}\right)$ and the potential evidence of time variation in this correlation. Since our model implies time variation in the correlation between $z_{i,t}$ and $z_{j,t}$ we shall evaluate the empirical evidence of this.

First, we consider a test for constant correlation that is based on the general theory by Nyblom (1989). This is the underlying framework of several tests for parameter constancy, including that of Hansen (1992) (linear regression models) and that of Hansen and Johansen (1999) (cointegration VAR).

Consider a bivariate process (x_t, y_t) of studentized variables, $\mathrm{E}(x_t) = \mathrm{E}(y_t) = 0$ and $\mathrm{E}\left(x_t^2\right) = \mathrm{E}\left(y_t^2\right) = 1$; so that the correlation is given by

$$\rho_t = \mathrm{E}\left(x_t y_t\right)$$

We are to construct tests for constant correlation and zero correlation. The maintained hypothesis is that the partial sum

$$W_T(u) \equiv T^{-\frac{1}{2}} \sum_{s=1}^{\lfloor uT \rfloor} (x_s y_s - \rho_s), \qquad u \in [0, 1]$$

satisfies a functional central limit theorem, so that $W_T(u) \Rightarrow \sigma_W B(u)$, where $B(u)$ is a standard Brownian motion and σ_W^2 is the long-run variance of $x_t y_t - \rho_t$.

Under the null hypothesis, $H_0 : \rho_t = \rho$ (constant correlation) it follows that

$$NB_c = \frac{T^{-1} \sum_{t=1}^{T} \left(T^{-1/2} \sum_{s=1}^{t} (x_s y_s - \bar{\rho})\right)^2}{\hat{\sigma}_W^2} \xrightarrow{d} \int_0^1 B_b(u)^2 \mathrm{d}u$$

where $B_b(u) = B(u) - uB(1)$ is a standard Brownian bridge, $\bar{\rho} = T^{-1} \sum_{t=1}^{T} x_t y_t$ and $\hat{\sigma}_W^2$ is some consistent estimator of σ_W^2. Under the null hypothesis $H_0 : \rho_t = 0$ (zero correlation) we have

$$NB_0 = \frac{T^{-1} \sum_{t=1}^{T} \left(T^{-1/2} \sum_{s=1}^{t} x_s y_s\right)^2}{\hat{\sigma}_W^2} \xrightarrow{d} \int_0^1 B(u)^2 \mathrm{d}u$$

where $\hat{\sigma}_W^2 \xrightarrow{p} \sigma_W^2$. In the absence of serial dependence we can use the estimator $\hat{\sigma}_W^2 = T^{-1} \sum_{t=1}^{T} (x_t y_t - \bar{\rho})^2$, which is consistent for σ_W^2 under both null hypotheses. The 5% critical values of these limit distributions are 0.462 and 1.656, respectively (see Nyblom, 1989).

Table IV. Sector statistics

Sector	Company counts	Min. beta	Median	Max beta
Energy	46	0.299 (0.174)	1.009 (0.145)	2.396 (0.510)
Materials	37	0.382 (0.104)	0.873 (0.117)	2.347 (0.836)
Industrials	63	0.264 (0.133)	0.931 (0.144)	2.293 (0.427)
Consumer discretionary	103	0.163 (0.156)	0.912 (0.136)	2.395 (0.664)
Consumer staples	46	0.336 (0.127)	0.970 (0.135)	2.188 (0.422)
Healthcare	64	0.358 (0.147)	1.072 (0.158)	2.586 (0.632)
Financials	101	0.309 (0.127)	1.037 (0.141)	2.633 (0.665)
Information technology	86	0.335 (0.156)	0.944 (0.145)	2.496 (0.734)
Telecommunication services	9	0.727 (0.166)	1.082 (0.186)	1.620 (0.307)
Utilities	40	0.284 (0.140)	0.988 (0.175)	2.332 (0.579)

Note: The table gives summary statistics of the sectoral aggregation. The third to fifth columns give the time series average of the minimum, median and maximum beta for each sector. Standard deviations are given in parentheses.

In our application we shall apply the test for constant correlation to $z_{i,t} z_{j,t}$ and $w_{i,t} w_{j,t}$, and the test for zero correlation to $w_{i,t} w_{j,t}$.

With 594 stocks in our cross-section there are 176,121 distinct correlation series to consider. In order to make the presentation manageable, we aggregate the results by industrial segmentation, based on the Global Industry Classification Standard (GICS). The 10 sectors are listed in Table IV along with the number of companies and summary statistics for their betas within each of the sectors.

Next we turn to the constancy of correlations within and across sectors. These results are presented in Tables V–VII. Table V gives the sample average of the unconditional correlations for residuals sorted by GICS. The upper panel is for the studentized returns, $\hat{z}_{i,t}$ and $\hat{z}_{j,t}$, and serves as a benchmark measure. It is interesting to compare these correlations with those in the lower panel, which are based on $\hat{w}_{i,t}$ and $\hat{w}_{j,t}$. The differences reveal how much of the correlations between individual stocks can be attributed to the market factor. The average correlation across $\hat{w}$-variables is small for stocks in different sectors, which suggests that the market may account for much of these correlations. However, for stocks within the same sector there is a substantial amount of unexplained correlation, as is evident from the diagonal entries in the lower panel. Thus an additional factor (beyond the market factor) is needed to explain the correlation structure of stocks within the same sector. For now we will simply investigate some statistical properties of the residual co-variation and leave the modeling for future work.

In Table VI we apply the NB_c test for constant correlation to our residual series. We report the rejection frequencies for a 5% significance level. In the upper panel we present the frequencies for the product of the studentized returns, $\hat{z}_{i,t} \hat{z}_{j,t}$. For example, in the case of the 46 energy companies there are 1035 tests and the null hypothesis of constant correlation is rejected for almost 65% of these test. Once we account for the market factor, and test the hypothesis of constant correlation for the idiosyncratic studentized returns, w_{it}, the rejection frequencies are much smaller, as can be seen from the lower panel of Table VI. This is especially true for assets in different sectors. For stocks within the same sector there continues to be substantial evidence of non-constant correlations, and the same is the case when one (or both) of the assets belongs to the energy sector.

One key message to take away from Table VI is that the evidence of time-varying correlations across sectors is greatly reduced by accounting for their associations with market returns. Within sectors there is substantial residual time variation and, evidently, for some sectors there is a need for additional factors if a larger fraction of the time variation is to be accounted for.

Table VII presents the tests for zero correlation. The table reports the rejection frequencies for the NB_0 test applied to the $\hat{w}_{i,t} \hat{w}_{j,t}$ series. Given the results in Table VI, it is not surprising that the test is frequently rejected for assets belonging in the same sector. Across sectors the zero correlation is also

Table V. Unconditional correlations (sorted by GICS)

	Energy	Materials	Industrials	Consumer discretionary	Consumer staples	Healthcare	Financials	Information technology	Telecommun. services	Utilities
Energy	0.598	0.324	0.276	0.203	0.158	0.180	0.231	0.217	0.196	0.298
Materials		0.417	0.376	0.308	0.240	0.242	0.341	0.297	0.273	0.285
Industrials			0.400	0.330	0.258	0.262	0.356	0.320	0.287	0.286
Consumer discretionary				0.322	0.234	0.232	0.326	0.281	0.251	0.238
Consumer staples					0.260	0.207	0.260	0.203	0.213	0.232
Healthcare						0.260	0.258	0.227	0.215	0.221
Financials							0.429	0.301	0.289	0.300
Information technology								0.356	0.267	0.229
Telecommun. services									0.369	0.251
Utilities										0.487
Energy	0.483	0.091	0.015	-0.035	-0.044	-0.018	-0.056	-0.022	-0.026	0.099
Materials		0.160	0.082	0.042	0.008	0.004	0.023	0.018	0.006	0.021
Industrials			0.104	0.058	0.021	0.020	0.026	0.036	0.012	0.013
Consumer discretionary				0.098	0.032	0.020	0.049	0.033	0.010	-0.005
Consumer staples					0.105	0.036	0.021	-0.015	0.016	0.040
Healthcare						0.096	0.012	0.009	0.010	0.022
Financials							0.145	0.002	0.011	0.024
Information technology								0.130	0.023	-0.023
Telecommun. services									0.187	0.029
Utilities										0.341

Note: Average sample correlations for residuals sorted by industry classification (GICS). Upper panel is for $\hat{z}_{i,t}$ and $\hat{z}_{j,t}$ and the numbers in the lower panel are based on $\hat{w}_{i,t}$ and $\hat{w}_{j,t}$.

Table VI. Testing for constant correlations (sorted by GICS)

	Energy	Materials	Industrials	Consumer discretionary	Consumer staples	Healthcare	Financials	Information technology	Telecommun. services	Utilities
Energy	0.649	0.781	0.854	0.636	0.589	0.423	0.718	0.794	0.679	0.570
Materials		0.619	0.676	0.619	0.470	0.341	0.612	0.609	0.616	0.588
Industrials			0.767	0.665	0.616	0.415	0.623	0.628	0.626	0.683
Consumer discretionary				0.645	0.470	0.387	0.647	0.575	0.595	0.536
Consumer staples					0.507	0.453	0.516	0.552	0.577	0.674
Healthcare						0.367	0.409	0.369	0.410	0.475
Financials							0.738	0.575	0.570	0.439
Information technology								0.682	0.562	0.513
Telecommun. services									0.778	0.597
Utilities										0.474
Energy	0.722	0.435	0.200	0.373	0.248	0.286	0.608	0.090	0.244	0.354
Materials		0.386	0.221	0.207	0.150	0.125	0.260	0.135	0.072	0.125
Industrials			0.147	0.159	0.110	0.094	0.168	0.132	0.063	0.142
Consumer discretionary				0.182	0.107	0.079	0.130	0.128	0.137	0.142
Consumer staples					0.173	0.153	0.141	0.204	0.143	0.183
Healthcare						0.166	0.107	0.093	0.113	0.130
Financials							0.382	0.131	0.086	0.411
Information technology								0.260	0.087	0.094
Telecommun. services									0.444	0.097
Utilities										0.268

Note: Rejection frequencies for the NB_c test for constant correlation using a 5% significance level. The upper panel shows the results for $\hat{z}_{i,t}, \hat{z}_{j,t}$ and the lower panel are those for $\hat{w}_{i,t}, \hat{w}_{j,t}$.

P. R. HANSEN, A. LUNDE AND V. VOEV

Table VII. Testing for zero correlations (sorted by GICS)

	Energy	Materials	Industrials	Consumer discretionary	Consumer staples	Healthcare	Financials	Information technology	Telecommun. services	Utilities
Energy	1.000	0.605	0.304	0.343	0.389	0.153	0.484	0.277	0.326	0.868
Materials		0.943	0.746	0.454	0.223	0.183	0.344	0.176	0.186	0.226
Industrials			0.744	0.545	0.236	0.210	0.298	0.335	0.145	0.142
Consumer discretionary				0.692	0.299	0.189	0.427	0.259	0.149	0.170
Consumer staples					0.736	0.283	0.259	0.298	0.130	0.435
Healthcare						0.706	0.157	0.270	0.123	0.162
Financials							0.816	0.238	0.111	0.443
Information technology								0.853	0.265	0.281
Telecommun. services									0.889	0.267
Utilities										1.000

Note: Rejection frequencies for the NB_0 test for zero correlation applied to $\hat{w}_{i,t}\,\hat{w}_{j,t}$.

frequently rejected. By introducing sector-specific factors it may be possible to explain the correlation structure of stocks within the same sector. Since additional factors would change the definition of the residual studentized returns, it is also plausible that sector specific factors could mitigate the residual correlation we find for assets in different sectors. We shall pursue this issue in future research.

6. CONCLUSION

In this paper we propose a multivariate GARCH model that utilizes realized measures of volatility and correlation, and entails a complete modeling of their dynamic properties. The model builds on a self-contained system of equations that link realized measures to the appropriate population quantities of volatility. The structure implies a dynamic model of the conditional betas that are popular measures of risk in finance. The proposed framework allows for leverage effects and spillover effects between the assets and the market volatility. In this respect the model combines the flexibility of the GARCH modeling framework with the statistical precision in volatility measurement resulting from the use of high-frequency data. Importantly, the realized beta GARCH model has a hierarchical structure that makes it easy to apply to a vast number of assets.

Our empirical study revealed some interesting features of the cross-sectional variation of the conditional betas, as well as their time series variation. In particular, we find that the betas exhibit substantial variation at a daily frequency—variation that is largely concealed in the rolling-window estimates of β that one can obtain with regression methods using daily returns. In our empirical analysis we also found that the realized beta GARCH model explains a great deal of the time variation in the correlation structure, and in a comparison of the dynamic betas of the model, with betas of alternative models, we found the betas of the realized beta GARCH to have be best empirical properties, on average.

Despite the advances brought by the realized beta GARCH model, it does not capture all of the dynamics of the large system. This was evident from some Nyblom tests, which revealed variation in the correlation structure that could be be explained by the underlying market factor. This was particularly the case for correlations between assets that belong to the same sector. For this reason, it will be interesting to consider a generalized structure where additional sector-specific correlation factors are used. We shall pursue this generalization in future research.

REFERENCES

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

Andersen TG, Bollerslev T, Diebold FX, Labys P. 2001a. The distribution of exchange rate volatility. *Journal of the American Statistical Association* **96**(453): 42–55. Correction published 2003; **98**: 501.

Andersen TG, Bollerslev T, Diebold FX, Ebens H. 2001b. The distribution of stock return volatility. *Journal of Financial Economics* **61**: 43–76.

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

Andersen T, Bollerslev T, Diebold F, Wu G. 2006. Realized beta: persistence and predictability. In *Advances in Econometrics: Econometric Analysis of Economic and Financial Time Series*, Fomby T, Terrell D (eds), Vol. 20. Elsevier Science: Amsterdam; 1–37.

Bandi FM, Russell JR. 2005. *Realized Covariation, Realized Beta, and Microstructure Noise*, Working paper. University of Chicago: Graduate School of Business.

Barndorff-Nielsen OE, Shephard N. 2002. Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society B* **64**: 253–280.

Barndorff-Nielsen OE, Shephard N. 2004a. Econometric analysis of realised covariation: high frequency based covariance, regression and correlation in financial economics. *Econometrica* **72**: 885–925.

Barndorff-Nielsen OE, Shephard N. 2004b. Power and bipower variation with stochastic volatility and jumps (with discussion). *Journal of Financial Econometrics* **2**: 1–48.

Barndorff-Nielsen OE, Shephard N. 2007. Variation, jumps and high frequency data in financial econometrics. In *Advances in Economics and Econometrics: Theory and Applications—Ninth World Congress*, Blundell R, Persson T, Newey WK (eds), Econometric Society Monographs. Cambridge University Press: Cambridge, UK; 328–372.

Barndorff-Nielsen OE, Hansen PR, Lunde A, Shephard N. 2008. Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica* **76**: 1481–536.

Barndorff-Nielsen OE, Hansen P, Lunde A, Shephard N. 2009. Realized kernels in practice: trades and quotes. *Econometrics Journal* **12**(3): C1–C32.

Barndorff-Nielsen OE, Hansen P, Lunde A, Shephard N. 2011. Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *Journal of Econometrics* **162**: 149–169.

Bekaert G, Wu G. 2000. Asymmetric volatility and risk in equity markets. *Review of Financial Studies* **13**(1): 1–42.

Bollerslev T, Zhang BYB. 2003. Measuring and modeling systematic risk in factor pricing models using high-frequency data. *Journal of Empirical Finance* **10**(5): 533–558.

Braun PA, Nelson DB, Sunier AM. 1995. Good news, bad news, volatility, and betas. *Journal of Finance* **50**: 1575–1603.

Chen X, Ghysels E, Wang F. 2011. HYBRID-GARCH a generic class of models for volatility predictions using mixed frequency data. Available: http://ssrn.com/abstract=1815762 [23 February 2014].

Christensen K, Kinnebrock S, Podolskij M. 2010. Pre-averaging estimators of the ex-post covariance matrix in noisy diffusion models with non-synchronous data. *Journal of Econometrics* **159**(1): 116–133.

Corradi V, Distaso W, Fernandes M. 2011. Conditional alphas and realized betas. Working paper. Available: http://www.birmingham.ac.uk/Documents/college-social-sciences/business/economics/events/conditional-alphas.pdf [23 February 2014].

Dovonon P, Goncalves S, Meddahi N. 2013. Bootstrapping realized multivariate volatility measures. *Journal of Econometrics* **173**: 49–65.

Engle RF. 2002a. Dynamic conditional correlation: a simple class of multivariate GARCH models. *Journal of Business and Economic Statistics* **20**(3): 339–350.

Engle RF. 2002b. New frontiers of ARCH models. *Journal of Applied Econometrics* **17**: 425–446.

Engle RF. 2012. Dynamic conditional beta. Available: http://ssrn.com/abstract=2084872 [23 February 2014].

Engle RF, Gallo G. 2006. A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics* **131**: 3–27.

Ferson WE, Harvey CR. 1991. The variation of economic risk premiums. *Journal of Political Economy* **99**(2): 385–415.

Ferson WE, Harvey CR. 1993. The risk and predictability of international equity returns. *Review of Financial Studies* **6**(3): 527–566.

Gomes J, Kogan L, Zhang L. 2003. Equilibrium cross section of returns. *Journal of Political Economy* **111**(4): 693–732.

Griffin JE, Oomen RCA. 2011. Covariance measurement in the presence of non-synchronous trading and market microstructure noise. *Journal of Econometrics* **160**(1): 58–68.

Hansen BE. 1992. Testing for parameter instability in linear models. *Journal of Policy Modeling* **14**: 517–533.

Hansen H, Johansen S. 1999. Some tests for parameter constancy in cointegrated VAR- models. *Econometrics Journal* **2**: 306–333.

Hansen PR. 2009. Quadratic variation by Markov chains, Research Paper 2009-13, CREATES.

Hansen PR, Huang Z. 2012. Exponential garch modeling with realized measures of volatility, Research Paper 2012-44, CREATES.

Hansen PR, Lunde A. 2011. Forecasting volatility using high-frequency data. In *The Oxford Handbook of Economic Forecasting*, Clements M, Hendry D (eds). Blackwell: Oxford; 525–556.

Hansen PR, Lunde A, Nason JM. 2011. The model confidence set. *Econometrica* **79**: 453–497.

Hansen PR, Huang Z, Shek H. 2012. Realized GARCH: a joint model of returns and realized measures of volatility. *Journal of Applied Econometrics* **27**: 877–906.

Hayashi T, Yoshida N. 2005. On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli* **11**: 359–379.

Jacod J, Li Y, Mykland PA, Podolskij M, Vetter M. 2009. Microstructure noise in the continuous case: the pre-averaging approach. *Stochastic Processes and their Applications* **119**(7): 2249–2276.

Lewellen J, Nagel S. 2006. The conditional CAPM does not explain asset-pricing anomalies. *Journal of Financial Economics* **82**(2): 28900314.

Lunde A, Olesen KV. 2013. Modeling and forecasting the volatility of energy forward returns, Research Paper 2013-15, CREATES.

Morana C. 2009. Realized betas and the cross-section of expected returns. *Applied Financial Economics* **19**: 1371–1381.

Nelson DB. 1991. Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* **59**(2): 347–370.

Nyblom J. 1989. Testing for the constancy of parameters over time. *Journal of the American Statistical Association* **84**: 223–230.

Patton A, Verardo M. 2012. Does beta move with news? Firm-specific information flows and learning about profitability. *Review of Financial Studies* **25**: 2789–2839.

Shanken J. 1990. Intertemporal asset pricing : an empirical investigation. *Journal of Econometrics* **45**(1–2): 99–120.

Shephard N, Sheppard K. 2010. Realising the future: forecasting with high frequency based volatility (HEAVY) models. *Journal of Applied Econometrics* **25**: 197–231.

Todorov V, Bollerslev T. 2010. Jumps and betas: a new framework for disentangling and estimating systematic risks. *Journal of Econometrics* **157**(2): 220–235.

Tsay RS, Yeh J-H. 2011. Random aggregation with applications in high-frequency finance. *Journal of Forecasting* **30**: 72–103.

Visser MP. 2011. GARCH parameter estimation using high-frequency data. *Journal of Financial Econometrics* **9**: 162–197.

Voev V, Lunde A. 2007. Integrated covariance estimation using high-frequency data in the presence of noise. *Journal of Financial Econometrics* **5**: 68–104.

White H. 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* **48**: 817–838.

Zhang L. 2006. Efficient estimation of stochastic volatility using noisy observations: a multi- scale approach. *Bernoulli* **12**: 1019–1043.