Analysis of High Frequency Financial Data: Models, Methods and Software. Part II: Modeling and Forecasting Realized Variance Measures.

Eric Zivot

July 4, 2005

1 Introduction

A key problem in financial econometrics is the modeling, estimation and forecasting of conditional return volatility and correlation. Having accurate forecasting models for conditional volatility and correlation is important for accurate derivatives pricing, risk management and asset allocation decisions. It is well known that conditional volatility and correlation are highly predictable. An inherent problem with modeling and forecasting conditional volatility is that it is unobservable, which implies that modeling must be indirect. Popular parametric models for latent volatility include the ARCH-GARCH family, the stochastic volatility family, and the Markov-switching family. In these models volatility is usually extracted from daily squared returns, which are unbiased but noisy estimates of daily conditional volatility. High frequency data is rarely utilized. The estimation of these models, however, often give unsatisfactory results. In particular, forecasts are imprecise. Moreover, standardized returns generally have fat-tails which has led to the search for appropriate error distributions that can adequately capture empirical return distributions. Furthermore, multivariate modeling of volatility and correlation can be extremely difficult and practical models are often only feasible for very low dimensions.

An exciting new area of research involves estimating, modeling and forecasting conditional volatility and correlation using high frequency intra-day data. The justification for using high frequency data follows from recent research that shows that daily conditional volatility and correlation can be accurately estimated using so-called realized volatility and correlation measures, which are based on summing high frequency squared returns and cross products of returns. Now, instead of using complicated models for unobserved volatility one can use more straightforward models for observed volatility. This use of high frequency data has the potential of revolutionizing the way volatility and correlation are modeled and forecasted.

This part of the lecture surveys the recent literature on modeling and forecasting realized variance and correlation using high frequency intra-day returns. Section 2 describes the construction of realized variance measures. Section 3 reviews the theoretical literature linking realized variance measures to quadratic variation processes derived from continuous-time arbitrage free price processes. The asymptotic distribution theory for realized variance measures is discussed in section 4. Sections 5, 6 and 7 survey some of the recent empirical literature on modeling and forecasting realized variance measures for foreign exchange and equity returns. Section 8 concludes with suggestions for future research.

2 Construction of Realized Variance

Let $p_{i,t}$ denote the log-price of asset i at time t, where each asset's log price has been aligned to a common regularly spaced time clock (e.g., every 5 minutes or every 30 minutes). In the multivariate context, let $\mathbf{p}_t = (p_{1,t}, \ldots, p_{n,t})'$ denote the $n \times 1$ vector of log prices at time t. Let Δ denote the fraction of a trading session associated with the implied sampling frequency, and let $m = 1/\Delta$ denote the number of sampled observations per trading session. For example, if prices are sampled every 30 minutes and trading takes place 24 hours per day then there are $m = 48$ 5-minute intervals per trading day and $\Delta = 1/288 \approx 0.0035$. If prices are sampled every 5 minutes and trading takes place 6.5 hours per day (e.g. the NYSE trades from 9:30 am EST until 16:00 p.m. EST) then there are $m = 78$ 5-minutes intervals per trading day and $\Delta = 1/78 \approx 0.0128$. Let T denote the number of days in the sample. Then there will be a total of mT observations on each asset $i = 1, \ldots, n$.

The intra-day continuously compounded (cc) return on asset i from time t to $t+\Delta$ is defined as

$$r_{i,t+\Delta} = p_{i,t+\Delta} - p_{i,t},\ i = 1, \ldots, n$$

The $n \times 1$ vector of cc returns from time t to $t + \Delta$ is defined as

$$\mathbf{r}_{t+\Delta} = \mathbf{p}_{t+\Delta} - \mathbf{p}_t$$

For notational simplicity, the daily returns are denoted using a single time subscript t, so that[1]

$$\begin{aligned} r_{i,t} &= r_{i,t-1+\Delta} + r_{i,t-1+2\Delta} + \cdots + r_{i,t-1+m\Delta},\ i = 1, \ldots, n \\ \mathbf{r}_t &= \mathbf{r}_{t-1+\Delta} + \mathbf{r}_{t-1+2\Delta} + \cdots + \mathbf{r}_{t-1+m\Delta} \end{aligned}$$

Realized variance (RV) for asset i $(i = 1, \ldots, n)$ on day t is defined as

$$RV_{i,t} = \sum_{j=1}^{m} r^2_{i,t-1+j\Delta},\ t = 1, \ldots, T$$

[1]This is the end-of-day return from the end of day $t-1$ until the end of day t.

Realized volatility (RVOL) for asset i on day t is defined as the square root of realized variance:

$$RVOL_{i,t} = \sqrt{RV_{i,t}}$$

Realized log-volatility (RLVOL) is the natural logarithm of RVOL:

$$RLVOL_{i,t} = \ln(RVOL_{i,t})$$

The $n \times n$ *realized covariance* (RCOV) matrix on day t is defined as

$$RCOV_t = \sum_{j=1}^{m} \mathbf{r}_{t-1+\Delta}\mathbf{r}'_{t-1+\Delta},\ t = 1, \ldots, T$$

It is clear that $RV_{i,t} = [RCOV_t]_{i,i}$ and $RCOV_{i,j,t} = [RCOV_t]_{i,j}$. The $n \times n$ matrix $RCOV_t$ will be positive definite provided $n < m$; that is, provided the number of assets is less than the number of intra-day observations. The realized correlation between asset i and asset j is computed using

$$RCOR_{i,j,t} = \frac{[RCOV_t]_{i,j}}{\sqrt{[RCOV_t]_{i,i} \times [RCOV_t]_{j,j}}} = \frac{[RCOV_t]_{i,j}}{RVOL_{i,t} \times RVOL_{j,t}}$$

Given daily measures of RV and RCOV, the corresponding non-overlapping measures over h days are computed as

$$\begin{aligned} RV_{i,t}^{h} &= \sum_{j=1}^{h} RV_{i,t},\ t = h, 2h, \ldots, T/h \\ RCOV_{i,t}^{h} &= \sum_{j=1}^{h} RCOV_{t},\ t = h, 2h, \ldots, T/h \end{aligned}$$

2.1 Practical Problems in the Construction of RV

There are a number of practical problems in the construction of RV measures. The foremost problem is the choice of sampling frequency Δ (or number of observations per day m). As will be shown below, the consistency of RV measures as estimators of underlying volatility depend on $\Delta \to 0$ ($m \to \infty$). However, it is not possible to sample continuously. As a result, RV measures contain measurement error. This point is emphasized in Bandi and Russell (2003), and they propose a data-based method for choosing Δ that minimizes the MSE of the measurement error. Additionally, as discussed in Bai, Russell and Tiao (2000), various market microstructure effects (bid/ask bounce, infrequent trading, calendar effects etc.) induce serial correlation in the intra-day returns $r_{i,t+\Delta}$ which may induce biases in RV measures. One way of correcting for these biases is to filter the intra-day returns using simple MA or AR models prior to constructing RV measures. These issues will be further discussed in the sections below.

3 Quadratic Return Variation and Realized Variance

Two fundamental questions about RV are:

Q1 What does RV estimate?

Q2 Are RV estimates economically important?

To answer these questions, in a series of papers, Andersen, Bollerslev, Diebold and Labys (2001, 2003), hereafter ABDL, and Barndorff-Nielsen and Shephard, (2002a,b, 2004a,b), hereafter BNS, have rigorously developed a theory connecting realized variance measures to quadratic return variation process derived from continuous time arbitrage-free log-price process. The general results apply to log-price processes belonging to the class of processes called *special semi-martingales*. Most of the continuous time processes utilized in financial models, including Itô diffusions, jump and mixed jump diffusion, are in this class.

To illustrate the main results, consider the univariate case and let $p(t)$ denote the univariate log-price process for a representative asset defined on a complete probability space (Ω, F, P), evolving in continuous time over the interval $[0, T]$[2]. Let F_t be the $\sigma-$field reflecting information at time t such that $F_s \subseteq F_t$ for $0 \leq s \leq t \leq T$. If $p(t)$ is in the class of special semi-martingales then it has the representation

$$p(t) = p(0) + A(t) + M(t),\ A(0) = M(0) = 0 \tag{1}$$

where $A(t)$ is a predictable drift component of finite variation, and $M(t)$ is a *local martingale*. Note that the predictability of the drift process, $A(t)$, allows for stochastic evolution. A detailed discussion of this type of decomposition is given in Protter (1990) and its economic significance is discussed in Back (1991). For notational convenience, let the unit interval denote one trading day. Then for mT a positive integer indicating the number of return observation obtained by sampling $m = 1/\Delta$ times per day for T days, the continuously compounded return on asset i over the period $[t - \Delta, t]$ is

$$r(t, t - \Delta) = p(t) - p(t - \Delta),\ t = \Delta, 2\Delta, \ldots, T$$

The daily cc return is

$$r(t, t - 1) = p(t) - p(t - 1)$$

and the cumulative return from 0 until t is

$$r(t) = p(t) - p(0)$$

[2]The following results generalize to the multivariate setting. See ABDL (2003) for details.

Using the notation from BNS (2002a), the *quadratic variation* (QV) of the return process at time t is defined as

$$[r](t) = p\lim \sum_{j=0}^{m-1} \{p(s_{j+1}) - p(s_j)\}^2 \tag{2}$$

where $0 = s_0 < s_1 < \cdots < s_M = t$ and the limit is for the mesh size

$$\max_{1 \leq j < m} |s_j - s_{j-1}| \to 0 \text{ as } m \to \infty$$

The QV process (2) measures the realized sample path variation of the squared return process. It is a unique and invariant ex-post realized volatility measure that is essentially model free.

The definition of QV implies the following convergence result:

$$RV_t \xrightarrow{p} [r](t) - [r](t-1) \equiv QV_t, \text{ as } m \to \infty \tag{3}$$

That is, *daily RV converges in probability to the daily increment in QV.* This answers the first question Q1.

As noted by ABDL (2001, 2003), QV_t defined in (3) is is related to, but distinct from, the daily conditional return variance. Specifically, they show that if

(i) the price process in (1) is square integrable;

(ii) the mean process $A(t)$ is continuous;

(iii) the daily mean process, $\{A(s) - A(t-1)\}_{s \in [t-1,t]}$, conditional on information at time t is independent of the return innovation process, $\{M(u)\}_{u \in [t-1,t]}$,

(iv) the daily mean process, $\{A(s) - A(t-1)\}_{s \in [t-1,t]}$, is a predetermined function over $[t-1, t]$,

then for $0 \leq t-1 \leq t \leq T$

$$var(r(t, t-1)|F_{t-1}) = E[QV_t|F_{t-1}] \tag{4}$$

That is, *the conditional return variance equals the conditional expectation of the daily QV process.* This result implies that QV_t is central to volatility measurement and forecasting. Furthermore, the ex post value of RV_t is an unbiased estimator for the conditional return variance $var(r(t, t-1)|F_{t-1})$:

$$E[RV_t|F_{t-1}] = E[QV_t|F_{t-1}] = var(r(t, t-1)|F_{t-1})$$

These results provide an answer the second question Q2.

ABDL (2003) argue that restrictions on the conditional mean process required for the result (4) allow for realistic price processes. In particular, the price process is allowed to exhibit deterministic intra-day seasonal variation. The mean process can be stochastic as long as it remains a function, over the interval $[t-1, t]$, of variables in F_{t-1}. Also, leverage effects caused by contemporaneous correlation between return innovations and innovations to the volatility process are allowed.

For the class of continuous-time *Itô processes* characterized by the stochastic differential equation

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) \tag{5}$$

where $W(t)$ is a Wiener process (standard Brownian motion), stronger results may be obtained. For this process, the daily return may be represented as

$$r(t, t-1) = \int_{t-1}^{t} \mu(s)ds + \int_{t-1}^{t} \sigma(s)dW(s)$$

In addition, the daily increment to QV has the form

$$QV_t = \int_{t-1}^{t} \sigma^2(s)ds = IV_t$$

where IV_t denotes daily *integrated variance* (IV). IV is central to certain models of option pricing with stochastic volatility (e.g., Hull and White (1987)). Since $RV_t \xrightarrow{p} QV_t$, it follows that for the Itô process

$$RV_t \xrightarrow{p} IV_t$$

ABDL (2003) further show that if mean process, $\mu(s)$, and volatility process, $\sigma(s)$, are independent of the Wiener process $W(s)$ over $[t-1, t]$ then

$$r(t, t-1)|\sigma\{\mu(s), \sigma(s)\}_{s\in[t-1,t]} \sim N\left(\int_{t-1}^{t} \mu(s)ds, IV_t\right) \tag{6}$$

where $\sigma\{\mu(s), \sigma(s)\}_{s\in[t-1,t]}$ denotes the $\sigma-$field generated by $(\mu(s), \sigma(s))_{s\in[t-1,t]}$. Given that $\int_{t-1}^{t} \mu(s)ds$ is generally very small for daily returns and that RV_t is a consistent estimator of IV_t, the result in (6) implies that daily returns should follow a *normal mixture distribution* with RV_t as the mixing variable. If there are jumps in (5), then it is still the case that $RV_t \xrightarrow{p} IV_t$ but returns are no longer conditionally normally distributed. As will be discussed in more detail below, ABDL demonstrate empirically that daily returns standardized by realized volatility

$$\frac{r(t, t-1)}{RVOL_t}$$

are approximately normally distributed which provides evidence that returns may be appropriately modeled by a jumpfree diffusion process.

The results presented above are for univariate returns. However, the results hold for multivariate returns as well. See ADBL (2001, 2003) for full details.

4 Asymptotic Distribution Theory for Realized Variance

Another fundamental question about RV is:

Q3 How precise is RV?

To help understand the answer to this question, consider the continuous diffusion process (5) where $\mu(t)$ is predictable and of finite variation and the $\sigma(t)$ process is independent of the Brownian motion $W(t)$. For this diffusion process, the consistency of RV_t for IV_t relies on the sampling frequency per day, Δ, going to zero. This theoretical convergence result, of course, is not attainable in practice as it is not possible to actually sample continuously. However, the theory suggests that one might want to sample as often as possible to get the most accurate estimate of IV_t. Unfortunately, market microstructure frictions will eventually dominate the behavior of RV as $\Delta \to 0$ which suggests that there is a practical lower bound on Δ for observed data. As a result, for $\Delta > 0$, RV_t will always be a noisy estimate of IV_t. The error for a given Δ may be represented as

$$u_{i,t}(\Delta) = RV_{i,t} - IV_t \tag{7}$$

BNS (2001) derive the asymptotic distribution of the error (7) as $\Delta \to 0$, or, equivalently, as $m = 1/\Delta \to \infty$. For the diffusion model (5), under the assumption that mean and volatility processes are jointly independent of $W(t)$ they show that

$$\sqrt{m}\frac{u_{i,t}(\Delta)}{\sqrt{2 \cdot IQ_{i,t}}} = \sqrt{m}\frac{(RV_{i,t} - IV_{i,t})}{\sqrt{2 \cdot IQ_{i,t}}} \xrightarrow{d} N(0,1)$$

where

$$IQ_{i,t} = \int_{t-1}^{t} \sigma^4(s)ds$$

is the *integrated quarticity* (IQ). This result shows that $RV_{i,t}$ converges to $IV_{i,t}$ at rate $\sqrt{m}$, and that asymptotic distribution of $RV_{i,t}$ is mixed-normal since $IV_{i,t}$ is random. Furthermore, BNS show that IQ_t may be consistently estimated using the following scaled version of *realized quarticity* (RQ)

$$\frac{m}{3}RQ_{i,t} = \frac{m}{3}\sum_{j=1}^{m} r^4_{i,t+\Delta}$$

Therefore, the feasible asymptotic distribution for $RV_{i,t}$ is

$$\frac{RV_{i,t} - IV_{i,t}}{\sqrt{\frac{2}{3} \cdot RQ_{i,t}}} \overset{A}{\sim} N(0,1) \tag{8}$$

This result suggests the following estimated asymptotic standard error for $RV_{i,t}$

$$\widehat{SE}(RV_{i,t}) = \sqrt{\frac{2}{3}\sum_{j=1}^{m} r_{i,t+\Delta}^4}$$

Using straightforward delta-method arguments, BNS also derive the asymptotic distribution of $RVOL_{i,t}$

$$\frac{RVOL_{i,t} - \sqrt{IV_{i,t}}}{\sqrt{\frac{2}{12} \cdot \frac{RQ_{i,t}}{RV_{i,t}}}} \overset{A}{\sim} N(0,1) \tag{9}$$

which suggests the feasible estimated standard error for $RVOL_{i,t}$

$$\widehat{SE}(RV_{i,t}) = \sqrt{\frac{2}{12} \cdot \frac{RQ_{i,t}}{RV_{i,t}}}$$

BNS find that the finite sample distribution of $RV_{i,t}$ and $RVOL_{i,t}$ can be quite far from their respective asymptotic distributions for moderately sized m. BNS (2003) derive the asymptotic distribution of $RLVOL_{i,t}^2$,

$$\frac{RLVOL_{i,t}^2 - \ln(IV_{i,t})}{\sqrt{\frac{2}{3} \cdot \frac{RQ_{i,t}}{RV_{i,t}^2}}} \overset{A}{\sim} N(0,1) \tag{10}$$

and show that the finite sample behavior of $RLVOL_{i,t}^2$ is closer to its asymptotic distribution than the finite sample behavior of $RV_{i,t}$ and $RVOL_{i,t}$. As a result, the log-based approximation (10) is likely to be preferred for constructing standard errors and confidence intervals in practice. This conjecture is consistent with the empirical evidence in ABDL (2001) who find that the unconditional distribution of RVOL is approximately log-normal.

BNS (2004) extend the above asymptotic results to cover the multivariate case, providing asymptotic distributions for $RCOV_t$ and $RCOR_{i,j,t}$, as well as realized regression estimates. These limiting distributions are much more complicated than the ones presented above, and the reader is referred to BNS (2004) for full details and examples.

5 Empirical Analysis of Realized Variance

Much of the published empirical analysis of RV has been based on high frequency data from two sources: The Trades and Quotation (TAQ) data for equity returns; and Olsen and Associates proprietary FX data sets for foreign exchange returns. Most studies of RV utilize the Olsen FX data. This is primarily due to the fact that the FX market for the major currencies is highly liquid and trades actively 24 hours per day. This guarantees many quotes per day per currency. For example, there are on average about 4,000 daily quotes for the DM/$ and about 2,000 daily quotes for the Yen/$ in the Olsen data sets. Also, FX quotes are revised constantly even in the absence of trading and this mitigates the negative autocorrelation induced by infrequent trading. These features make the FX data well suited for the analysis of RV. In contrast, the stocks covered by the TAQ data have varying amounts of liquidity. The U.S. equity markets are most active during the NYSE trading hours (9:30 a.m. EST to 4 p.m. EST). However, not all stocks are actively traded and so are not well suited for RV analysis. As a result, most studies of RV using equity tend to focus on a few actively traded stocks. The most comprehensive analysis of RV for equity to date only utilizes the 30 stocks in the Dow Jones Industrial average.

Before surveying the empirical analysis of RV, the following sections give more detail on the Olsen and TAQ data.

5.1 The Olsen FX Data

A number of authors have analyzed realized variance measures of foreign exchange returns computed from the either the Olsen HFDF-1993, HFDF-1996, Olsen HF-2000 data sets[3]. These data sets were made available for use in three conferences on the statistical analysis of high frequency data sponsored by Olsen and Associates. The Olsen HFDF-2000 data is the most commonly used data set, and it is briefly described here. This data set contains spot exchange rates sampled every 5 minutes for the U.S. dollar ($), the Deutschemark (DM), Swiss Franc (CHF), British Pound (BP), and the Japanese yen (Yen) over the period December 1, 1986 through June 30, 1999. The raw data consist of all interbank bid/ask indicative (non-binding) quotes for the exchange rates displayed on the Reuters FXFX screen during the sample period. The 5-minute DM/$ and Yen/$ returns over the sample period are constructed by Olsen Data following Dacorogna et al. (1993). Each quote consists of a bid and an ask price together with a time stamp to the nearest even second. After filtering the data for outliers and other anomalies using a proprietary filter technology, the log-price at each 5-minute tick is obtained by linearly interpolating from the average of the log-bid and the log-ask quotes for the two closest ticks, and the 5-minute cc return is computed as the difference in the log-price.

Prior to the computation of the realized variance quantities, the 5-minute return data is often further restricted to eliminate non-trading periods, weekends, holidays, and lapses of the Reuters data feed. The FX market is a 24 hour market but slows considerably during the weekend. As a result, the weekend period from Friday 21:05 GMT until Sunday 21:00 GMT is eliminated from the sample. Further, the following holidays are commonly removed: Christmas (December 24-26), New Year's (December 31- January 2), July 4th, Good Friday, Easter Monday, Memorial Day, Labor

[3]These data sets may be purchased directly from Olsen Data AG (`www.olsendata.com`).

Author	Series	Sample Period	T	m
AB (1998a)	DM/$, Yen/$	10/1/87-9/30/93	260	288
AB (1998b)	DM/$, Yen/$	10/1/87-9/30/93	260	288
ABDL (2000)	DM/$, Yen/$	12/1/86-12/1/96	2,445	48
ABDL (2001)	DM/$, Yen/$	12/1/86-11/30/96	2,449	288
ABDL (2003)	DM/$, Yen/$	12/1/86-6/30/99	3,045	48
ABDM (2005)	DM/$, Yen/$	12/1/86-6/30/99	3,045	48
BNS (2002a)	DM/$	12/1/86-11/30/96	2,449	various
BNS (2002b)	DM/$	12/1/86-11/30/96	2,449	288
Maheu et. al. (2002)	DM/$	12/1/86-12/1/96	2,465	288

Table 1: Summary of authors using Olsen data

Day, and Thanksgiving and the day after. In addition, days that contain long strings of zero or constant returns (caused by data feed problems) are also eliminated.

Table 1 below summarizes how a number of authors have analyzed the various Olsen data sets.

5.2 The TAQ Data

Most studies of RV measures to date have utilized the Olsen FX data. Only a few authors have studies RV measures computed from high frequency equity returns. Most of these authors use data from Trade and Quotation (TAQ) database. The TAQ data files contain intra-day trade and quotation information for all securities listed on the New York Stock Exchange (NYSE), American Stock Exchange (AMEX), and the National Association of Security Dealers Automated Quotation system (NASDAQ). The data start in January 1993 and is available monthly on DVD-ROM. The most active period for equity markets is during the trading hours of the NYSE between 9:30 a.m. EST until 4:00 p.m. EST. Most studies of RV measures restrict attention to these trading hours.

Andersen, Bollerslev, Diebols and Ebens (2001), hereafter ABDE, provide the most comprehensive analysis of RV measures based on the TAQ data. They compute and analyze RV measures for the 30 stocks in the Dow Jones Industrial Average (DJIA) over the period January 2, 1993 through May 29, 1998 ($T = 1,336$ days).

Equity returns are generally subject to more pronounced market microstructure effects (e.g., negative first order serial correlation caused by bid-ask bounce effects) than FX data. As a result, equity returns are often filtered to remove these microstructure effects prior to the construction of RV measures. A common filtering method involves estimating an MA(1) or AR(1) model to the returns, and then constructing the filtered returns as the residuals from the estimated model.

Table 2 below summarizes how a selection of authors have analyzed the TAQ data.

Author	**Series**	**Sample Period**	T	m
AB (2001)	Dow Jones 30 stocks	10/1/87-9/30/93	1,366	79
Bandi et. al. (2003)	IBM	2/1/2002-2/28/2002	1/12	various
Hansen et. al. (2005)	Dow Jones 30 stocks	1/29/2001-12/31/2004	986	various

Table 2: Summary of authors using TAQ data

	RV_D	RV_Y	$RVOL_D$	$RVOL_Y$	$RLVOL_D$	$RLVOL_Y$	$RCOV$	$RCOR$
Mean	.529	.538	.679	.684	-.449	-.443	.243	.435
Variance	.234	.272	.067	.070	.120	.123	.073	.028
Skewness	3.71	5.57	1.68	1.87	.345	.264	3.78	-.203
Kurtosis	24.1	66.5	7.78	10.4	3.26	3.53	25.3	2.72

Table 3: Summary statistics for daily RV measures. Source ABDL (2001).

6 Empirical Analysis of FX Returns

The properties of RV variance measures for FX returns from the Olsen data are studied in a number of papers by Andersen, Bollerslev, Diebold and Labys. The main results from these papers are summarized in this section.

6.1 Unconditional Distribution of RV measures

ABDL (2001, 2003) study the properties of RV measures ($RV_{i,t}$, $RVOL_{i,t}$, $RLVOL_{i,t}$, $RCOV_{ij,t}$, $RCOR_{ij,t}$) for the DM/$ and Yen/$ returns over the ten year period from December 1986 through December 1996. In ABDL (2001) they compute RV measures using 5-minute returns ($m = 228$), and in ABDL (2003) they compute RV measures using 30-minute returns ($m = 48$). ABDL (2001) focus on the distributional properties of RV measures, whereas ABDL (2003) focus on modeling and forecasting RV measures.

Table 3 below gives summary statistics for the RV measures, and Figure 1 shows kernel density estimates of the distributions. The distributions of RV_t, $RVOL_t$ and $RCOV_t$ are non-normal and skewed right, whereas the distributions of $RLVOL_t$ and $RCOR_t$ appear to be approximately normal. The apparent non-normality of RV_t and $RVOL_t$ cast some doubt on the accuracy of the asymptotic distribution theory developed for these measures by BNS (2002). However, the approximate normality of $RLVOL_t$ is in line with the asymptotic theory developed by BNS (2003).

Table 4 shows the sample correlation matrix between the RV measures. The measures of volatility between the two currencies are highly positively correlated. That is, when the volatility of DM/$ is high the volatility of Yen/$ also tends to be high. As ABDL (2001) point out, this suggests a common factor driving volatility for the two currencies. Interestingly, the volatility measures are also positively correlated with the correlation measures. ABDL call this the “correlation-in-volatility” effect.

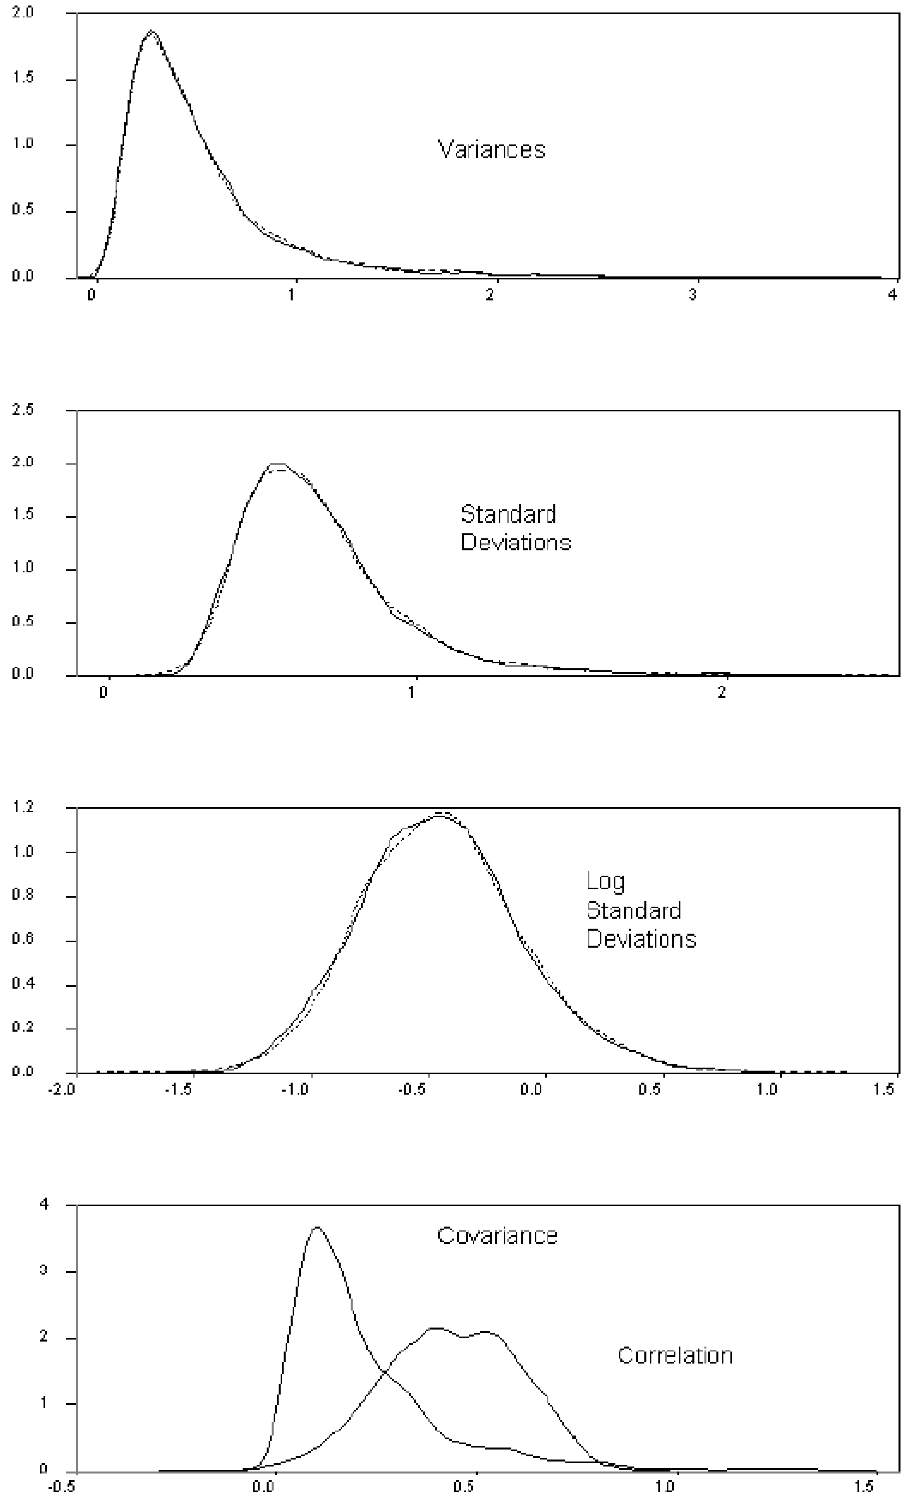

Figure 1: Distributions of daily realized exchange rate volatilities and correlations. Source: ABDL (2001).

	RV_Y	$RVOL_D$	$RVOL_Y$	$RLVOL_D$	$RLVOL_Y$	$RCOV$	$RCOR$
RV_D	.539	.061	.552	.860	.512	.806	.341
RV_Y	1.00	.546	.945	.514	.825	.757	.234
$RVOL_D$		1.00	.592	.965	.578	.793	.383
$RVOL_Y$			1.00	.589	.959	.760	.281
$RLVOL_D$				1.00	.604	.720	.389
$RLVOL_Y$					1.00	.684	.294
$RCOV$						1.00	.590

Table 4: Correlation matrix for daily RV measures. Source: ABDL (2001)

In particular, high RV seems to increase the RCOV and RCOR measures. This effect is illustrated in Figure 2, which shows kernel density estimates of $RCOR_t$ conditioned on high and low volatility days.

6.2 Accuracy of RV Measures

BNS (2002), using the same Olsen FX data as ADBL (2001), investigate the accuracy of RV measures. In particular, using intra-day DM/$ returns they compute RV_t for values of m ranging from 1 to 288 as well as approximate 95% confidence intervals based on the log-approximation (). Figure reproduces these results for the first 9 days of the dataset. Two features stand out: (1) the confidence intervals narrow as m increases as predicted by theory; (2) when RV_t is low it is estimated precisely, and when RV_t is large it is not estimated very precisely.

6.3 Conditional Distribution of RV Measures

Since RV measures are close connected to unobserved conditional volatility and correlation, the properties of the conditional distribution of RV measures gives information about the distribution of conditional volatility and correlation. Figure 4 shows time series plots of $RVOL_{i,t}$ $(i = D, Y)$ and $RCOR_{DY,t}$, and Figure 5 shows the sample autocorrelations (SACFs) of these measures. Both $RVOL_{i,t}$ and $RCOR_{DY,t}$ vary considerably over time. The slow decay of the SACFs reveal very strong persistence in these measures suggestive of long-memory or even unit root behavior.

ABDL (2001) reject the presence of unit roots in $RVOL_{i,t}$ and $RCOR_{DY,t}$. However, they find strong evidence for long-memory behavior. Recall, a stationary process y_t has *long memory*, or *long range dependence*, if its autocorrelation function behaves like

$$\rho(k) \to C_\rho k^{-\alpha} \text{ as } k \to \infty$$

where C_ρ is a positive constant, and α is a real number between 0 and 1. Thus the autocorrelation function of a long memory process decays slowly at a *hyperbolic* rate. Granger and Joyeux (1980) and Hosking (1981) independently showed that a long

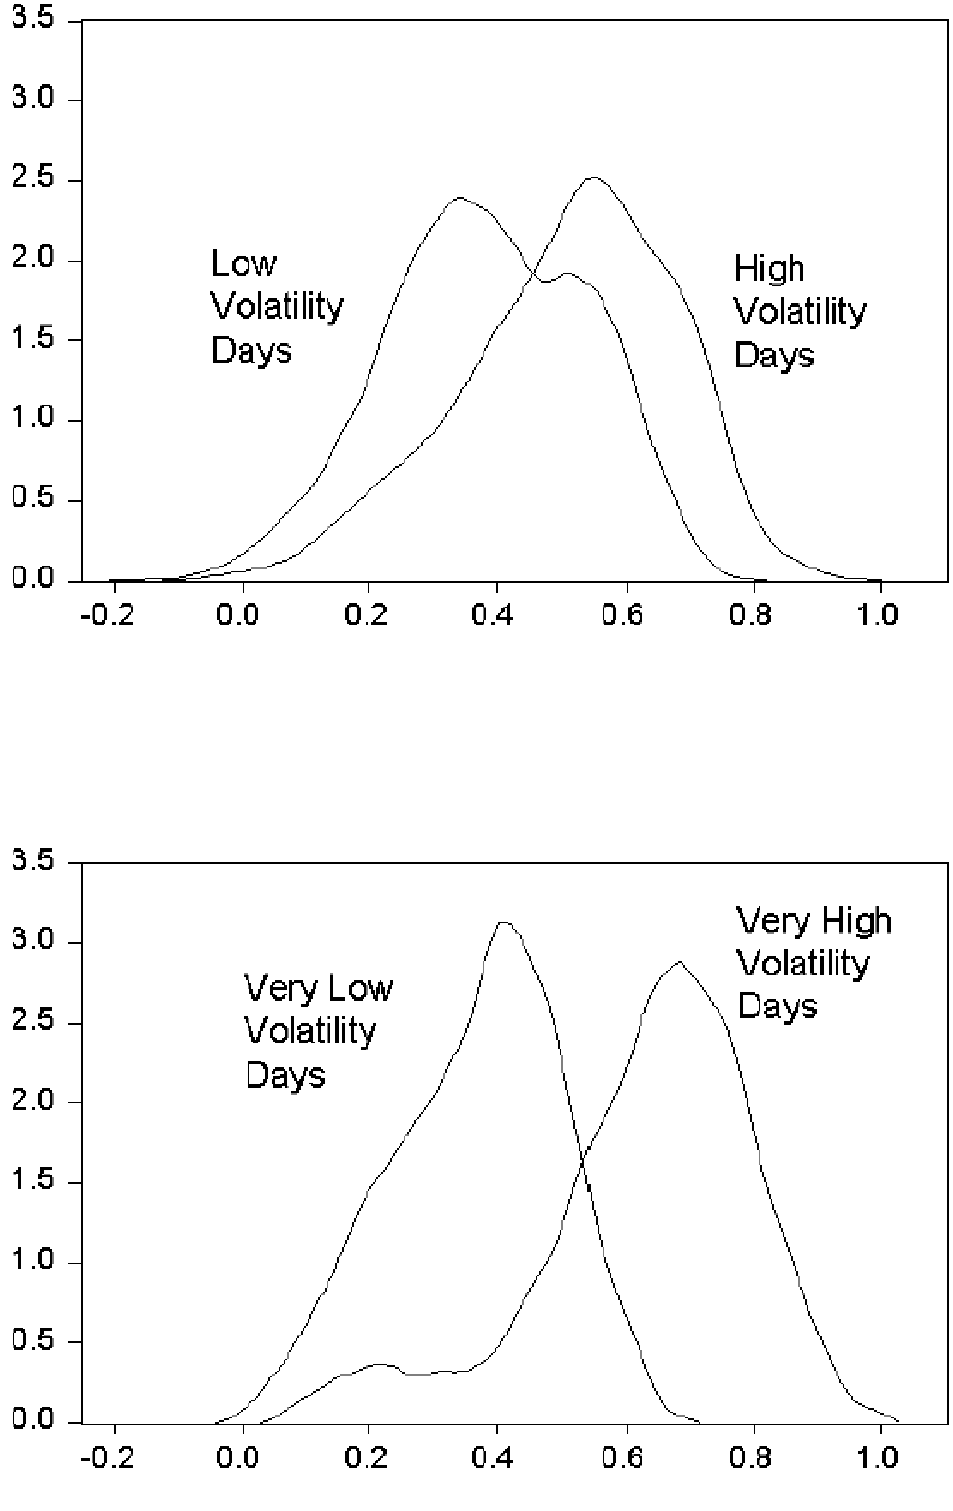

Figure 2: Distributions of Realized Correlations: Low Volatility vs. High Volatility Days. Source: ABDL (2001).

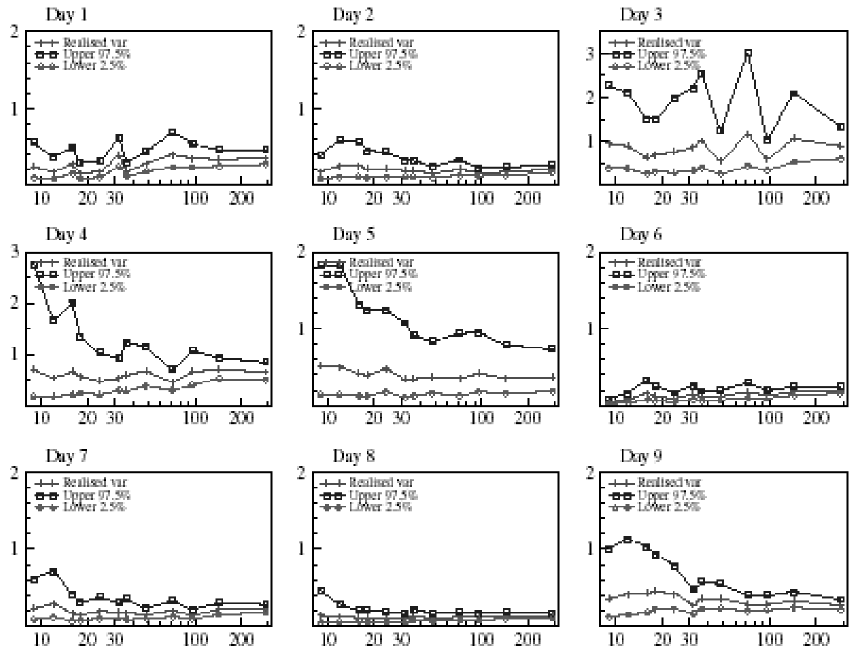

Figure 3: Daily RV, plotted against m, and 95% confidence intervals computed from asymptotic distribution. Source: BNS (2002).

memory process y_t can also be modeled parametrically by extending an integrated process to a *fractionally integrated process*:

$$(1-L)^d(y_t-\mu)=u_t$$

where L denotes the lag operator, d is the fractional integration or fractional difference parameter, μ is the expectation of y_t, and u_t is a stationary short-memory disturbance with zero mean. It can be shown that when $|d|>1/2$, y_t is non-stationary; when $0<d<1/2$, y_t is stationary and has long memory; when $-1/2<d<0$, y_t is stationary and has short memory, and is sometimes referred to as *anti-persistent.* The fractional integration parameter d may be estimated non-parametrically using the log-periodogram regression of Geweke and Porter-Hudak (1983), or it may be estimated parametrically from a fully specified fractional ARIMA model.

Table 5 shows estimates of the fractional differencing parameter, d, for the RV measures obtained from the GPH log-periodogram regression. The typical estimate of d is around 0.4 which indicates stationary long-memory behavior in all of the RV measures. ABDL (2003) use this evidence for long memory to build simple forecasting models for $RLVOL_{i,t}$.

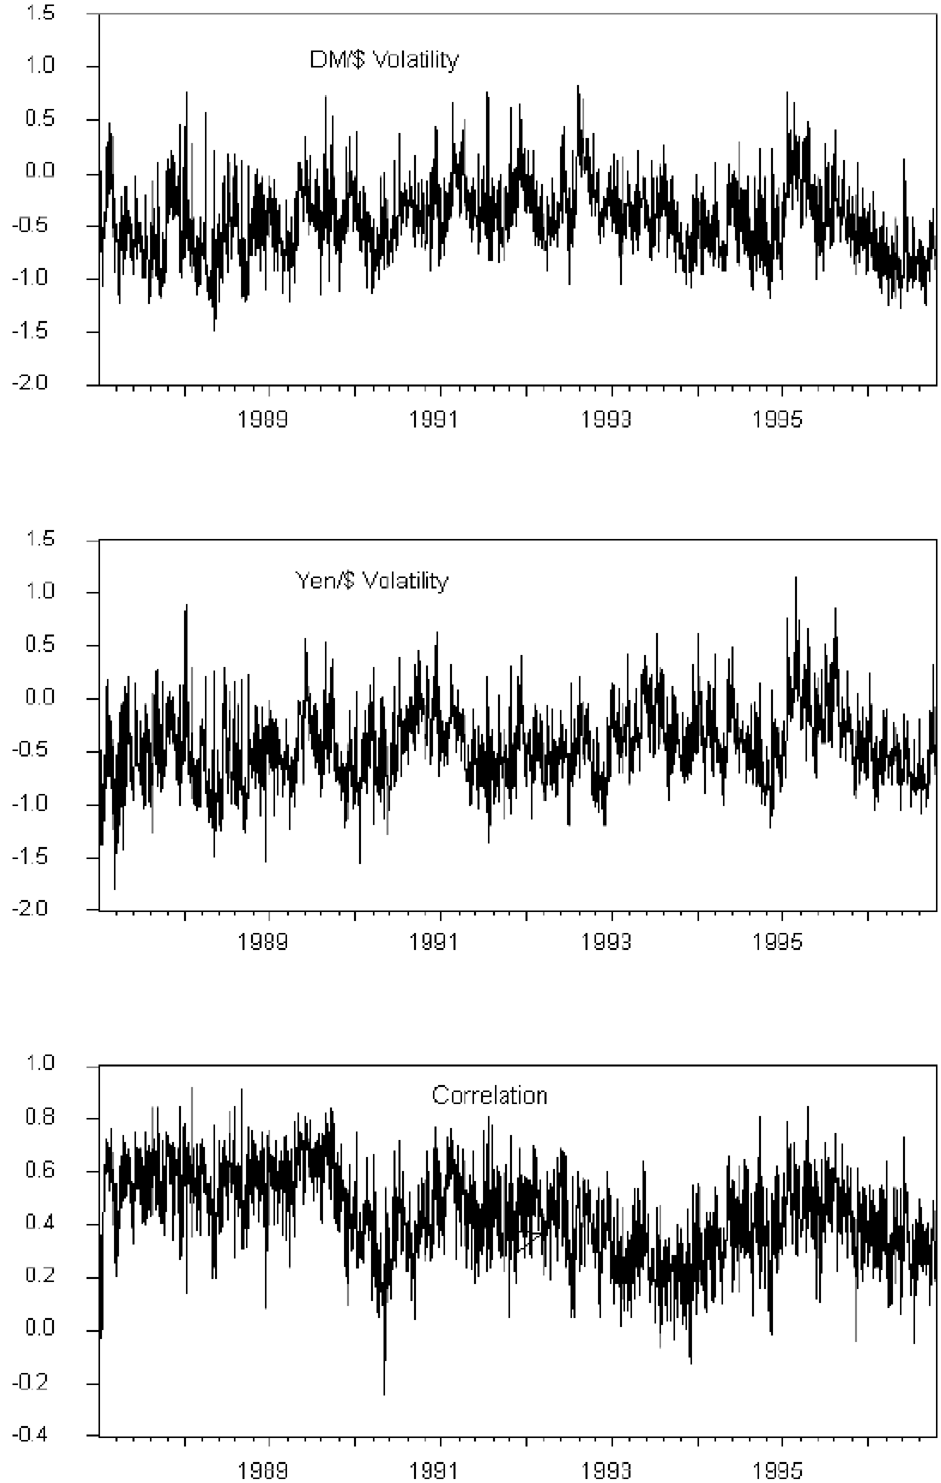

Figure 4: Time Series of Daily Realized Volatilities and Correlations. Source: ABDL (2001).

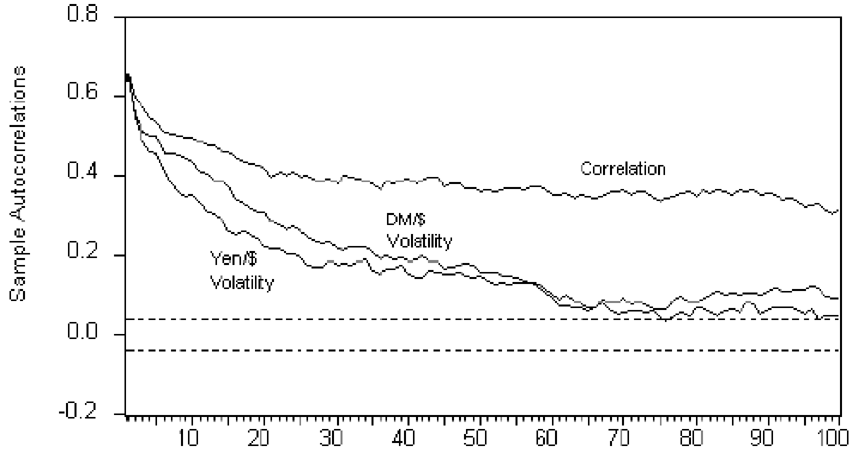

Figure 5: Sample autocorrelations of realized volatilities and correlations. Source: ABDL (2001).

	RV_D	RV_Y	$RVOL_D$	$RVOL_Y$	$RLVOL_D$	$RLVOL_Y$	$RCOV$	$RCOR$
$\widehat{d}$	.356	.339	.381	.428	.420	.455	.334	.413

Table 5: Long memory parameter estimates for daily RV measures. Source ABDL (2001).

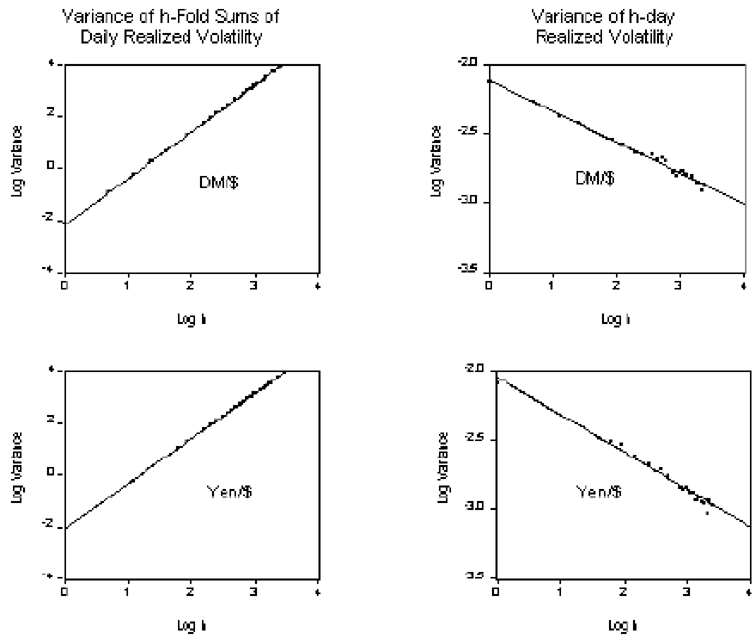

Figure 6: Scaling laws under temporal aggregation. Source: ABDL (2001).

6.4 Temporal Aggregation and Scaling Laws

ABDL (2001) investigate the conditional distribution of RV measures at different levels of aggregation ($h = 5, 10, 15$ and 20 days) and establish some simple scaling laws that further reinforce the evidence for long-memory behavior. They note that for the class of fractionally integrated models, the fractional differencing parameter d is invariant under aggregation. They compute log-periodogram estimates of d for the RV measures for different values of h and find little difference from the estimates based on $h = 1$. In addition, they compute $h-$fold partial sums of the form

$$[x_t]_h = \sum_{j=1}^{h} x_{h(t-1)+j},\ t = 1, 2, \ldots, h/2$$

and make use of the fact that if x_t is fractionally integrated with parameter d then

$$var([x_t]_h) = c \cdot h^{2d+1} \tag{11}$$

for some constant c. This result implies that plots of the logarithm of the sample variances of the partial sums of RV_t versus the logarithm of the aggregation level h should be linear. Figure 6 reproduces this plot taken from ABDL (2001), and indicates strong evidence for the long-memory scaling law (11).

6.5 Returns Standardized by RV

ABDL (2000) study the properties of returns standardized by RV measures computed from 30-minute returns. They motivate the analysis by assuming daily returns r_t may be decomposed following a standard conditional volatility model

$$r_t = \sigma_t \varepsilon_t \tag{12}$$

where σ_t represents the unobservable standard deviation of returns conditional on time t information, and $\varepsilon_t \sim iid\ (0, 1)$. They study the properties of raw returns r_t, as well as estimates of the standardized returns

$$\hat{\varepsilon}_t = \frac{r_t}{\hat{\sigma}_t}$$

where $\hat{\sigma}_t$ represents either a RV or normal-GARCH(1,1) estimate of σ_t. The normal-GARCH(1,1) model has the form

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

where $\varepsilon_t \sim iid\ N(0, 1)$.

Figure 7 illustrates the differences between RV and normal-GARCH(1,1) estimates of σ_t^2. The top panel shows squared returns, the middle panel shows the normal-GARCH(1,1) estimates, and the bottom panel shows the RV estimates. From (12), the squared returns are $r_t^2 = \sigma_t^2\varepsilon_t^2$ and so $E[r_t^2] = \sigma_t^2$ since $E[\varepsilon_t^2] = 1$ by assumption. Although r_t^2 is an unbiased estimate for σ_t, it is clearly a very noisy estimate. Comparing the GARCH and RV estimates of σ_t^2, it can be seen that the GARCH estimates are quite a bit smoother than the RV estimates. The GARCH estimate of σ_t^2 is essentially an exponentially weighted average of squared returns starting at $t-1$, and does not make use of information between $t-1$ and t. The RV estimate, in contrast, focuses exclusively on high frequency squared returns between $t-1$ and t. As a result, it can more accurately estimate volatility at time t than the GARCH model.

Table 6 and Figure 8 summarize the distributions of the unstandardized and standardized returns. Unstandardized daily returns tend to be roughly symmetric but leptokurtic. The returns standardized by the normal-GARCH(1,1) model are also roughly symmetric and slightly less leptokurtic. This is a typical finding with normal-GARCH models, which has motivated the use of GARCH models with fat-tailed innovations. The returns standardized by $RVOL_t$, in contrast, are approximately normally distributed. This result supports the theoretical prediction from a jumpless continuous-time diffusion model that returns standardized by RV should be normally distributed. It also supports the *mixture-of-distributions-hypothesis* for returns originally proposed by Clark (1973) and further developed by Tauchen and Pitts (1983) and Taylor (1986).

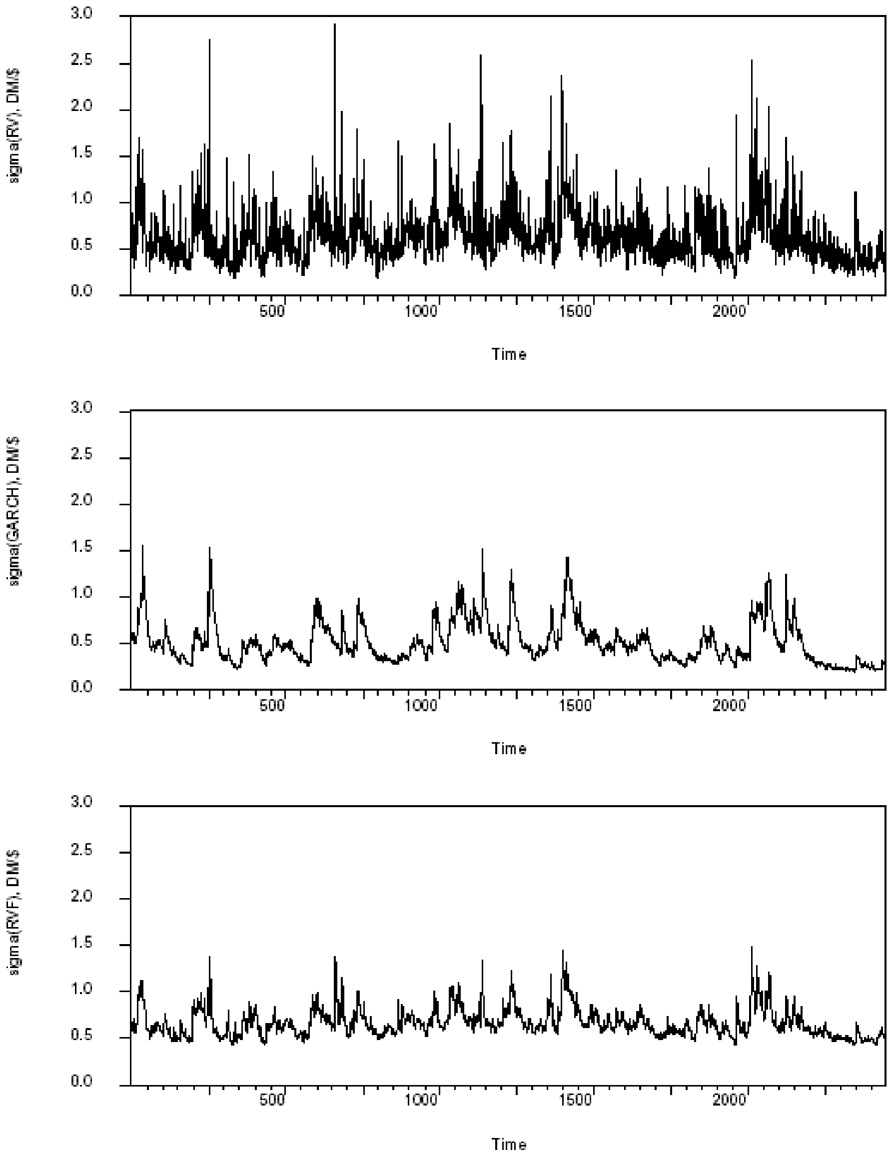

Figure 7: Time Series of Alternative Volatility Measures. Source: ABDL 2000.

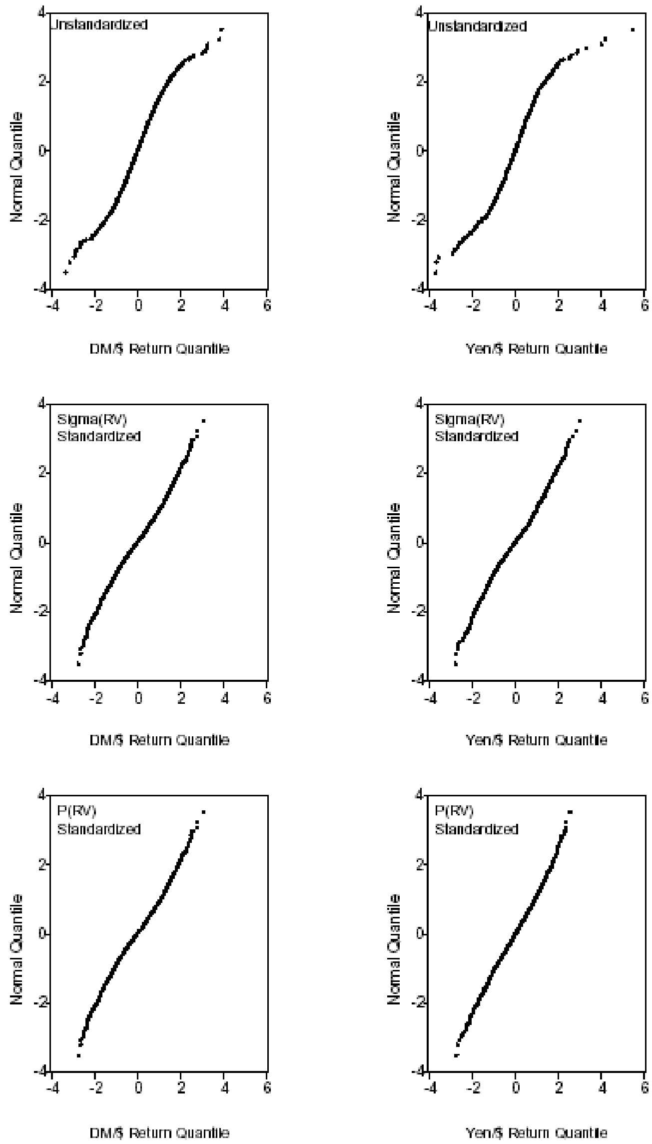

Figure 8: Normal QQ-plots for daily returns and returns standardized by RV measures. Source: ABDL (2000).

	r_t		$\frac{r_t}{\hat{\sigma}_t^{GARCH}}$		$\frac{r_t}{RVOL_t}$	
	DM/$	Y/$	DM/$	Y/$	DM/$	Y/$
Mean	-.007	-.009	-.002	-.011	-.007	.007
Std. Dev.	.710	.705	1.00	1.00	1.01	.984
Skewness	.033	.052	-.027	-.139	.015	.002
Kurtosis	5.40	7.36	4.75	5.41	2.41	2.41
Correlation	.659		.661		.661	

Table 6: Descriptive statistics for returns. Source: ABDL (2000)

Figure 9 shows scatterplots of the daily DM/$ and Yen/$ returns, as well as scatterplots of returns standardized by $RVOL_t$ and returns standardized using

$$\begin{pmatrix} \hat{\varepsilon}_{D,t} \\ \hat{\varepsilon}_{Y,t} \end{pmatrix} = RCOV_t^{-1/2} \begin{pmatrix} r_{D,t} \\ r_{Y,t} \end{pmatrix}$$

where $RCOV_t^{1/2}$ is the Choleski factorization of the 2×2 realized covariance matrix. The unstandardized returns are positively correlated with a correlation coefficient of 0.66. The bivariate distribution is clearly non-normal. One approach to modeling a non-bivariate bivariate distribution is through the use of *copulas*[4]. However, copula methods may not be necessary. The bivariate distribution of returns standardized by $RVOL_t$ appears to be approximately bivariate normal with a correlation of about 0.66, and the distribution of returns standardized by $RCOV_t^{1/2}$ appears to be approximately bivariate normal with no correlation.

Daily returns are approximately uncorrelated over time but squared and absolute returns exhibit substantial autocorrelation. The high persistence in squared returns, for example, indicates time varying conditional volatility in support of (12) where σ_t is modeled with a GARCH process. Figure 10 shows the sample autocorrelations of unstandardized squared returns, returns standardized by $RVOL_t$ and returns standardized by $RCOV_t^{1/2}$. The squared returns standardized by $RVOL_t$ are essentially uncorrelated, but the cross products $\hat{\varepsilon}_{D,t}\hat{\varepsilon}_{Y,t}$ exhibit slight autocorrelation. This autocorrelation in the cross products is eliminated for the returns standardized by $RCOV_t^{1/2}$.

6.6 Modeling and Forecasting Realized Variance

Traditional statistical approaches to modeling and forecasting daily conditional volatility treat conditional volatility as unobservable. Commonly used models for describing

[4]See chapter 19 in Zivot and Wang (2005) for an introduction to modeling bivariate distributions with copulas.

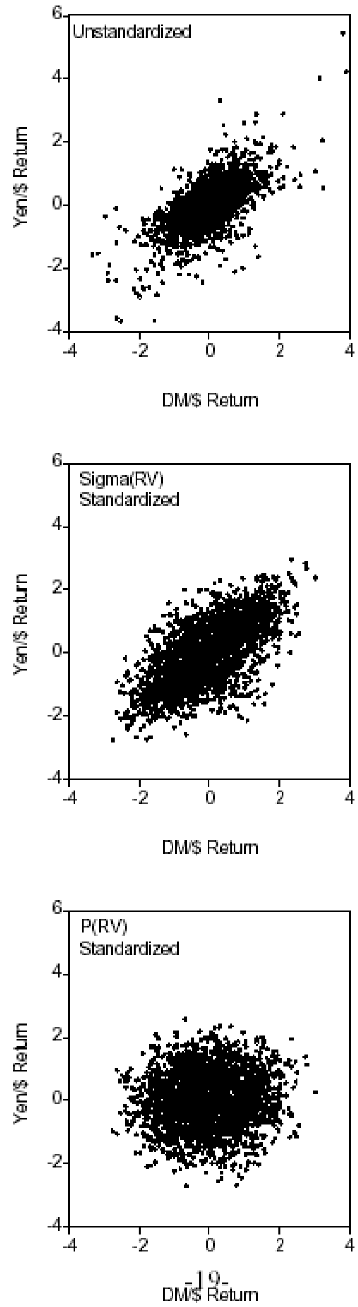

Figure 9: Scatterplots of returns and returns standardized by RV measures. Source: ABDL (2000).

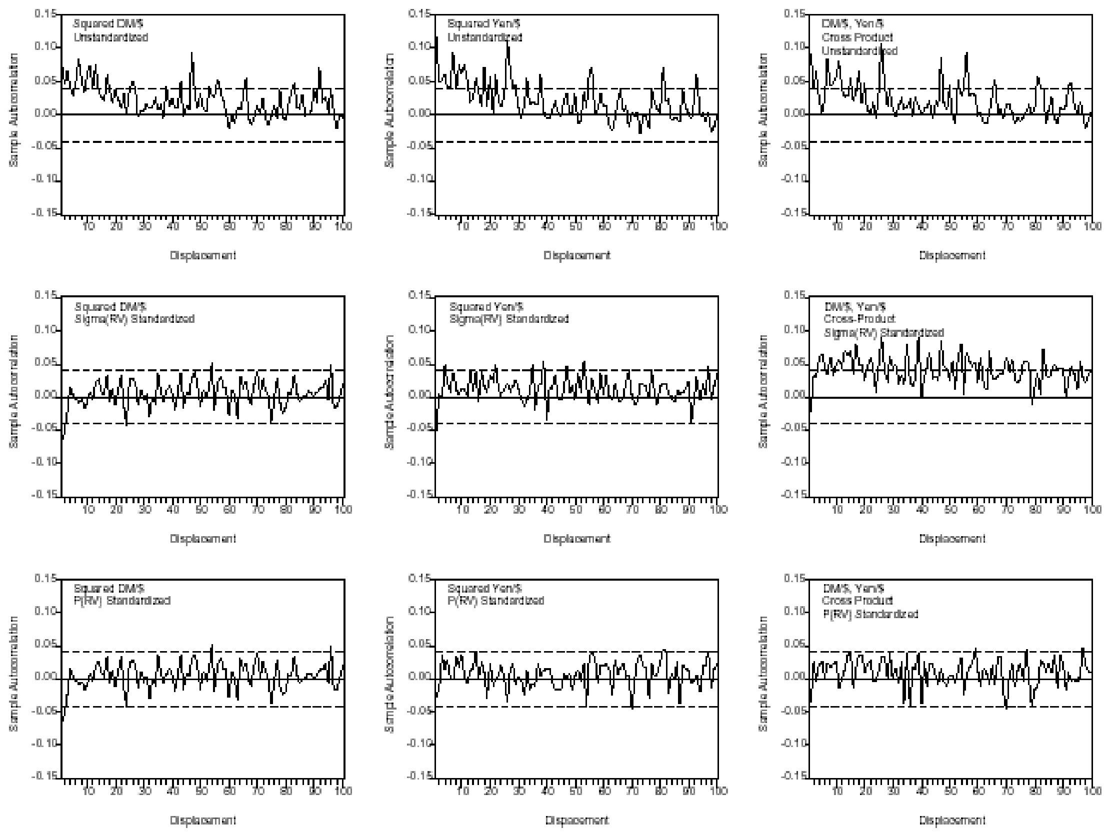

Figure 10: Sample autocorrelations of squared returns and squared returns standardized by RV meausres. Source: ABDL (2000).

daily returns are the normal-GARCH(1,1) model

$$\begin{aligned} r_t &= \sigma_t \varepsilon_t,\ \varepsilon_t \sim N(0,1) \\ \sigma_t^2 &= \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \end{aligned}$$

and the log-normal stochastic volatility (SV) model

$$\begin{gathered} r_t = \sigma_t \varepsilon_t,\ \varepsilon_t \sim N(0,1) \\ \ln \sigma_t^2 = \omega + \beta \ln \sigma_{t-1}^2 + \sigma_u u_t,\ u_t \sim N(0,1) \end{gathered} \tag{13}$$

In both models, the daily conditional variance σ_t^2 is unobserved and is allowed to evolve stochastically over time. The unobservability of σ_t^2 complicates the estimation of the models, particularly the SV model whose likelihood function conditional on observed returns must be formed by integrating out the unobserved conditional volatility. The problem become much worse with multivariate models, and practical multivariate models must be of very low dimension. Another drawback of the daily GARCH and SV models is that forecasts of σ_{t+1} based on information at time t cannot accommodate the information in intra-day data. GARCH and SV models may be specified directly using intra-day data, but doing so requires accounting for intra-day seasonalities and other market microstructure effects. Furthermore, these models often do not forecast as well as models specified for daily data and they cannot forecast very well beyond a day.

ABDL (2000) and ABDL (2003) argue that modeling and forecasting conditional volatility based on RV measures has many advantages over traditional approaches. The main advantage is that RV measures may be treated as observable estimates of conditional volatility. This allows the use of simple time series models (e.g. ARMA models) for describing the behavior of observed RV measures. In the multivariate context, the observability of RV measures allows for the possibility of modeling and forecasting very high dimensional covariance matrices.

ABDL (2003) illustrate the modeling and forecasting of RV measures using a system of three exchange rates (DM/\$, Yen/\$, Yen/DM) taken from the Olsen data. Making use of the empirical result that the logarithm of realized volatility is approximately normally distributed, they consider modeling and forecasting

$$\mathbf{y}_t = \begin{pmatrix} RLVOL_{D/\$,t} \\ RLVOL_{Y/\$,t} \\ RLVOL_{Y/D,t} \end{pmatrix} \tag{14}$$

where $RLVOL_{i,t}$ $(i = D/\$, Y/\$, Y/D)$ is computed from equally spaced 30-minute returns. The system (14) may be used as a model for the elements of the 2×2 realized covariance matrix

$$RCOV_t = \begin{pmatrix} RV_{D/\$,t} & RCOV_{D/\$,Y/\$,t} \\ - & RV_{Y/\$,t} \end{pmatrix}$$

since, by triangular arbitrage,

$$RCOV_{D/\$,Y/\$,t} = \frac{1}{2}\left(RV_{D/\$,t} + RV_{Y/\$,t} - RV_{Y/D,t}\right)$$

ABDL (2003) fit various models for y_t using the in-sample period 12/1/86 - 12/1//96, and construct forecasts for the out-of-sample period 12/2/96 - 6/30/99.

6.6.1 Long-Memory VAR Model

In ABDL (2001), it was shown that $RLVOL_{i,t}$ exhibits long-memory behavior. Using the GPH estimator, ABDL (2003) report estimates of the fractional integration parameter d to be close to 0.4 for the different elements of $\mathbf{y}_t$. Figure 11 shows the sample autocorrelations for the elements of $\mathbf{y}_t$ as well as the sample autocorrelations of the fractionally differenced series $(1-L)^{0.4}RLVOL_{i,t}$. The autocorrelations for $RLVOL_{i,t}$ decay very slowly whereas the autocorrelations of $(1-L)^{0.4}RLVOL_{i,t}$ die out quite quickly. In addition the elements of $\mathbf{y}_t$ are all moderately positively correlated.

Based on the above results, ABDL (2003) propose the simple fractionally differenced VAR(5) model to model and forecast $\mathbf{y}_t$:

$$\begin{gathered}\boldsymbol{\Phi}(L)(1-L)^{0.4}(\mathbf{y}_t - \boldsymbol{\mu}) = \boldsymbol{\varepsilon}_t \\ \boldsymbol{\varepsilon}_t \sim iid\ N(\mathbf{0}, \boldsymbol{\Omega}) \\ \boldsymbol{\Phi}(L) = 1 - \boldsymbol{\Phi}_1 L - \cdots - \boldsymbol{\Phi}_5 L^5\end{gathered}$$

ABDL denote this model VAR-RV. They fit the model using daily data for $\mathbf{y}_t$ over the ten year period 12/1/86 - 12/1/96. They do not report the estimates of the VAR(5) model parameters. However, they mention that the lag length of the VAR was chosen to capture dynamic effects that may be present up a week. Also, they mention that the VAR(5) model has an approximate diagonal structure that is not much different than a system of stacked univariate AR(5) models for each element of $(1-L)^{0.4}(\mathbf{y}_t - \boldsymbol{\mu})$.

6.6.2 Alternative Forecasting Models

ABDL consider the following alternative forecasting models for $\mathbf{y}_t$:

1. VAR-ABS: VAR(5) fit to $|\mathbf{r}_t|$
2. AR-RV: univariate AR(5) fit to $(1-L)^{0.4}RLVOL_{i,t}$
3. Daily GARCH(1,1): normal-GARCH(1,1) fit to daily returns $r_{i,t}$
4. Daily RiskMetrics: exponentially weighted moving average model for $r_{i,t}^2$ with $\lambda = 0.94$

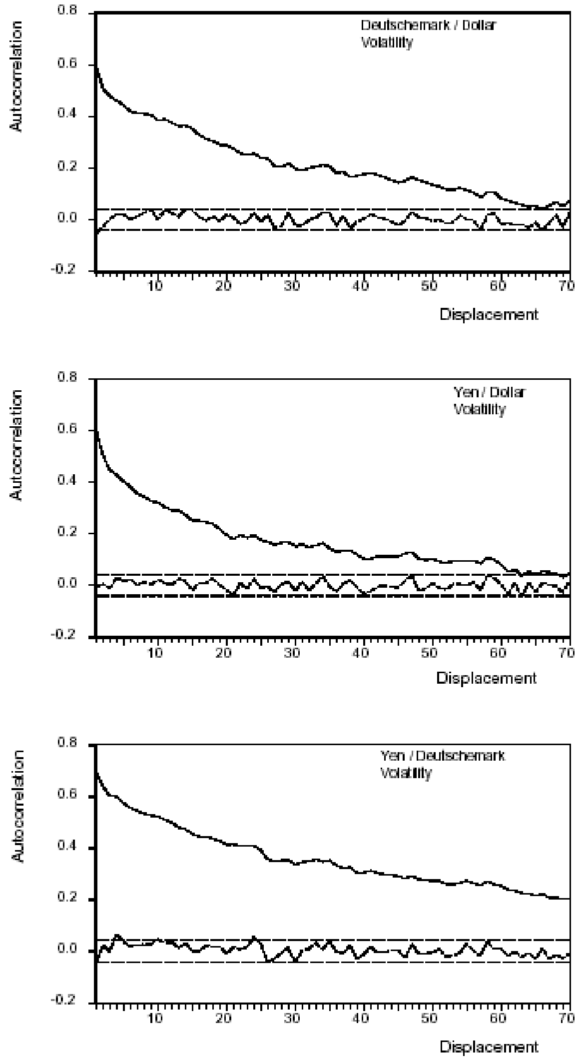

Figure 11: Sample autocorrelations of $RLVOL_{i,t}$ and $(1 - L)^{0.4}RLVOL_{i,t}$. Source: ABDL (2003).

5. Daily FIEGARCH(1,1): univariate fractionally integrated exponential GARCH(1,1) fit to $r_{i,t}$

6. Intra-day FIEGARCH deseason/filter: univariate fractionally integrated exponential GARCH(1,1) fit to 30-minute filtered and deseasonalized returns $\tilde{r}_{i,t+\Delta}$.

6.6.3 Forecast Evaluation and Comparison with Alternative Models

ABDL compute out-of-sample forecasts from models described above over the period 12/2/96 - 6/30/99. Figure 12 illustrates the forecasting accuracy of the preferred VAR-RV model for $\mathbf{y}_t$, and Figure 13 shows illustrates the forecasting accuracy of the daily GARCH(1,1) models. The VAR-RV forecasts track actual RV remarkably well, whereas the daily GARCH forecasts are much smoother. Both forecasts are based on lagged estimates of conditional volatility. The VAR-RV model uses lagged realized volatilities which are based on intra-day data and are accurate estimates of conditional volatility, whereas the GARCH model uses lagged squared returns which are very noisy estimates of conditional volatility.

ABDL evaluate the RV forecasts using the so-called *Mincer-Zarnowitz regression*

$$RVOL_{i,t} = b_0 + b_1\widehat{RVOL}_{i,t}^{VAR-RV} + b_2\widehat{RVOL}_{i,t}^{\text{model}} + \text{ error}_t \tag{15}$$

where $\widehat{RVOL}_{i,t}^{VAR-RV}$ denotes the 1-day-ahead out-of-sample forecast of $RVOL_{i,t}$ based on the VAR-RV model, and $\widehat{RVOL}_{i,t}^{\text{model}}$ denotes the 1-day-ahead out-of-sample forecast of $RVOL_{i,t}$ based on an alternative model. If VAR-RV is the best forecasting model, then one should find that the R^2 from (15) using just $\widehat{RVOL}_{i,t}^{VAR-RV}$ is higher than the R^2 from (15) when using any other model. In addition, if VAR-RV is an unbiased forecasting model, then one should find that $b_0 = 0, b_1 = 1$ and $b_2 = 0$. Using (), ABDL find that, indeed, the VAR-RV model is the best forecasting model. For in-sample regressions, they find that () estimated with just RV-VAR has the highest R^2. Moreover, they rarely reject the null hypothesis that $b_0 = 0, b_1 = 1$ and $b_2 = 0$. They find similar results for out-of-sample regressions.

ABDL provide evidence that their VAR-RV model also produces accurate $h-$step ahead forecasts. Figure 14, reproduced from ABDL (2003), shows RV_t for the DM/$ as well as forecasts from the VAR-RV model and the GARCH(1,1) model for four 35 day episodes. The first 25 days shows one-day-ahead in-sample forecasts, and the remaining 10 days shows $h-$day ahead forecasts. Notice how the VAR-RV model tracks RV_t both in-sample and out-of-sample, whereas the GARCH models performs quite poorly.

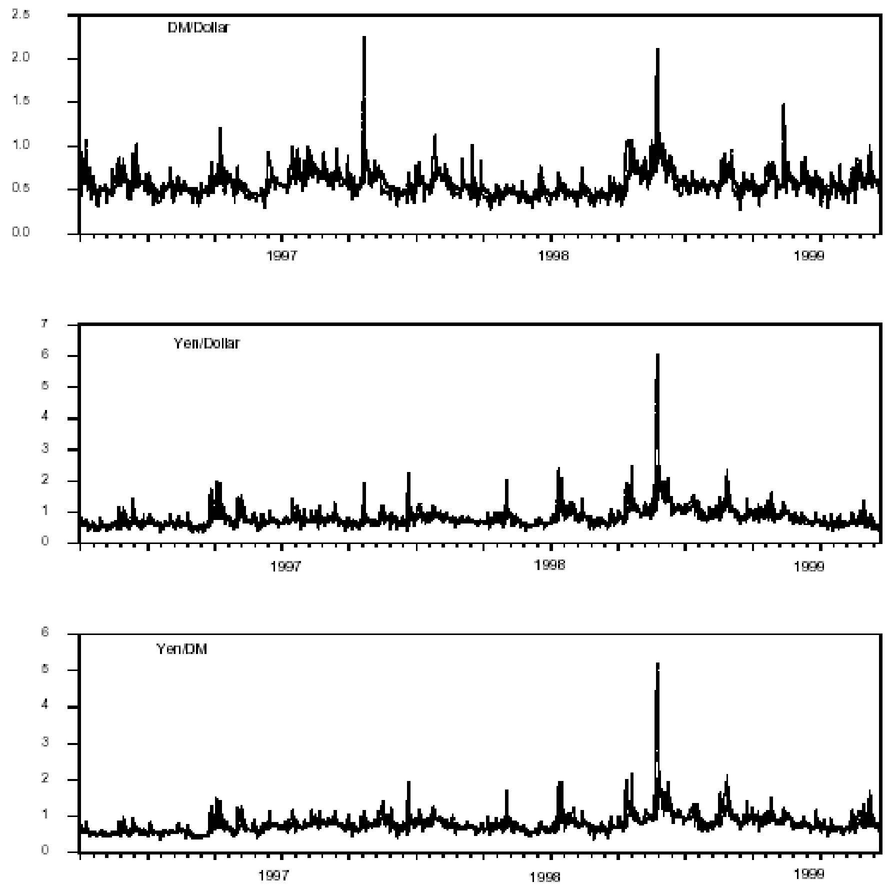

Figure 12: Realized volatility and 1-day-ahead out-of-sample VAR-RV forecasts. Source: ABDL (2003).

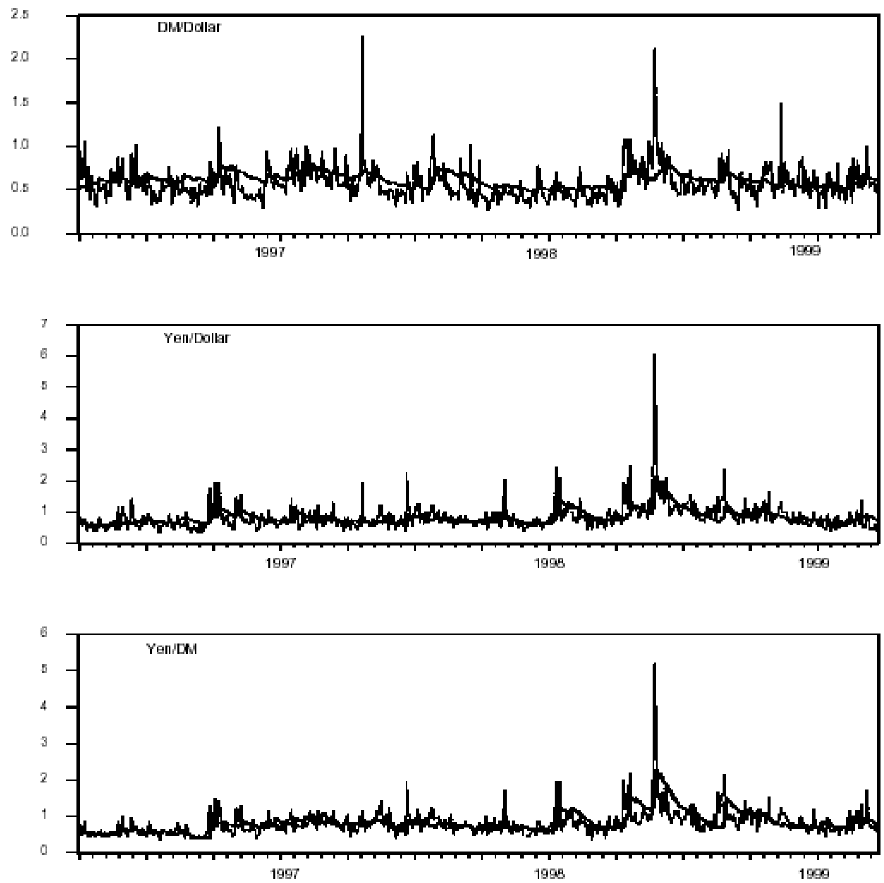

Figure 13: Realized volatility and 1-day-ahead out-of-sample forecasts from daily GARCH model. Source: ABDL (2003).

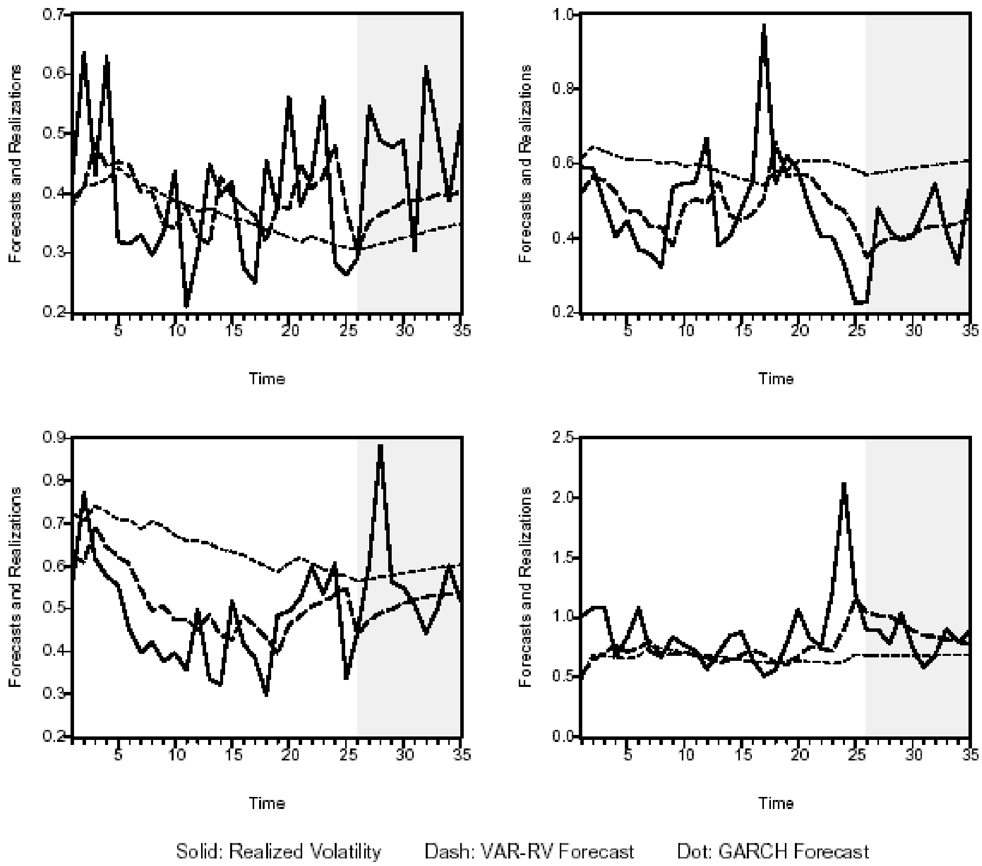

Figure 14: Realized volatility and h-step-ahead out-of-sample forecasts from VAR-RV and daily GARCH models. Source: ABDL (2003).

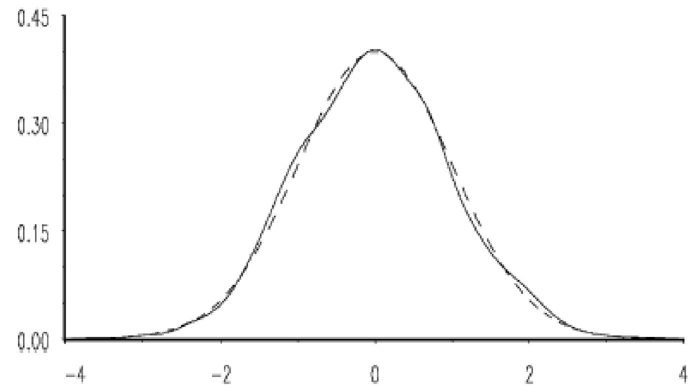

Figure 15: Unconditional distribution of $RLVOL_t$ from Alcoa returns. Source: ABDE (2001).

7 Empirical Analysis of Equity Returns

This section briefly summarizes the results of ABDE (2001) who study RV measures computed for the 30 DJIA stocks. They study the distribution of RV measures, but do not consider modeling and forecasting. The results are quite similiar to those found by ABDL (2000, 2001, 2003) for the Olsen FX returns.

7.1 Unconditional Distribution of RV Measures

ABDE compute and analyze univariate and multivariate RV measures ($RV_{i,t}$, $RVOL_{i,t}$, $RLVOL_{i,t}$, $RCOV_{ij,t}$, $RCOR_{ij,t}$) for the 30 stocks in the DJIA over the period January 2, 1993 through May 29, 1998 ($T = 1,336$ days). They align all intra-day returns to a common 5-minute clock starting at 9:30 a.m. EST until 4:30 p.m. EST giving $m = 79$ 5-minute returns each day. The total data set contains $3,237,420$ observations. Returns are de-meaned and MA(1)-filtered prior to the construction of the RV measures.

The unconditional distributions of the RV measures are similar to those found for the Olsen FX data. The distributions of $RV_{i,t}$, $RVOL_{i,t}$ and $RCOV_{ij,t}$ are non-normal and right skewed, whereas the distributions of $RLVOL_{i,t}$ and $RCOR_{ij,t}$ are approximately normal. See Figures 15 and 16.

7.2 Conditional Distribution of RV Measures

Figures 17 and 18 show a representative time series of $RLVOL_{i,t}$ and its sample autocorrelation function. There is clearly high persistence and evidence of long memory behavior.

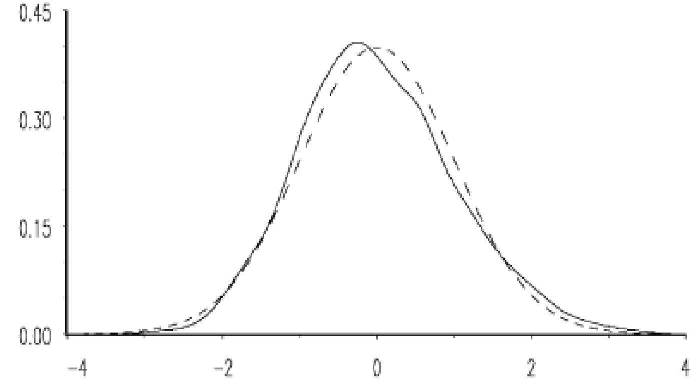

Figure 16: Unconditional distribution of $RCOR_{ij,t}$ from Alcoa and Exxon returns. Source: ABDE (2001).

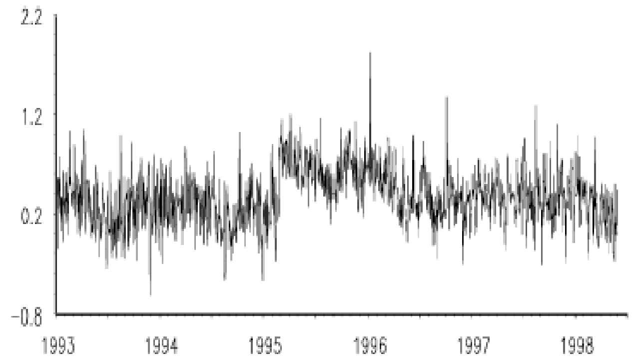

Figure 17: Time series of $RLVOL_{i,t}$ from Alcoa returns. Source: ABDE (2001).

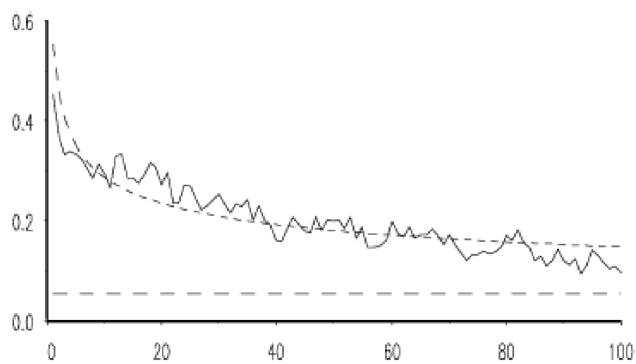

Figure 18: Sample autocorrelations of $RLVOL_t$. Source: ABDE (2001).

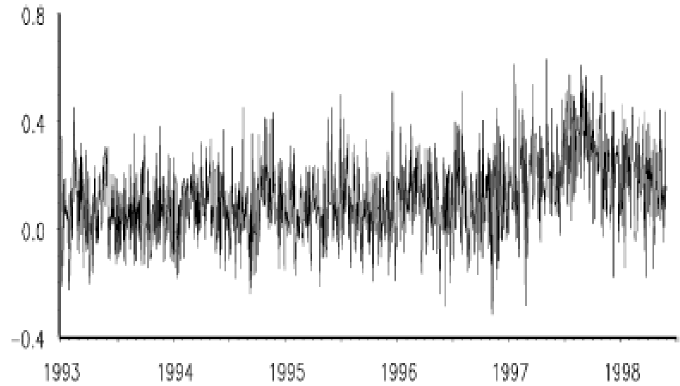

Figure 19: $RCOR_{ij,t}$ computed from Alcoa and Exxon returns. Source: ABDE (2001).

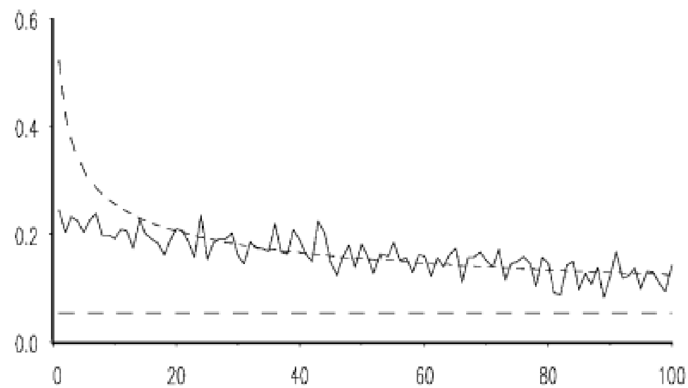

Figure 20: Sample autocorrelations of $RCOR_{ij,t}$ computed from Alcoa and Exxon returns. Source: ABDE (2001).

Similarly, Figures and show a representative time series of $RCOR_{ij,t}$ and its sample autocorrelations. The correlations are extremely variable and persistent, and also show evidence of long memory behavior.

As with the FX RV measures, the equity RV measures also follow the scaling law (11) implied by a long memory process.

7.3 Returns Standardized by RV

Daily equity returns are highly non-normal and right skewed. However, daily returns standardized by $RVOL_{i,t}$ are approximately normally distributed. This is illustrated for a representative stock in Figure 21.

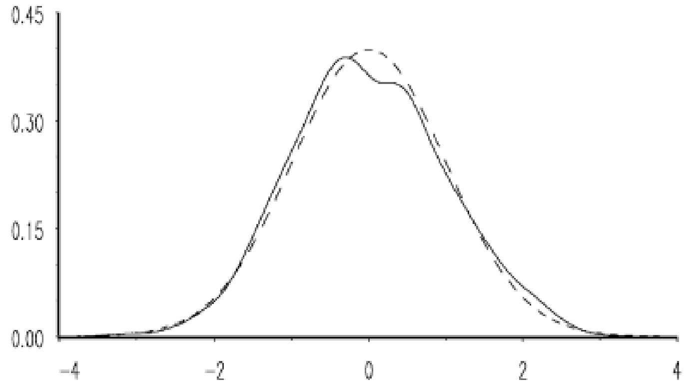

Figure 21: Density estimate for daily returns standardized by $RVOL_t$. Source: ABDE (2001).

8 Directions for Future Research

In their forthcoming *Handbook of Financial Econometrics* chapter "Parametric and Nonparametric Volatility Measurement," Andersen, Bollerslev and Diebold (2005) conclude with the following:

> In the last ten years, there has been a movement toward the use of newly-available high frequency asset return data, and away from restrictive and hard-to-estimate parametric models toward flexible and computationally simple nonparametric approaches. Those trends will continue. Two related, directions for future research are apparent: (1) continued development of methods for exploiting the volatility information in high-frequency data, and (2) volatility modeling and forecasting in the high-dimensional multivariate environments of practical financial economic relevance. The realized volatility concept readily tackles both: it incorporates the powerful information in high-frequency data while dispensing with the need to actually model the high-frequency data, and it requires only the most trivial of computations, thereby bringing within reach the elusive goal of accurate and high-dimenstional volatility measurement, modeling and forecasting.

9 References

1. Ait-Sahalia, Y., and P. Mykland (2003). "How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise," NBER working paper.

2. Andersen, T., and T. Bollerslev (1997a). "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns, *Journal of Finance* 52, 975-1005.

3. Andersen, T., and T. Bollerslev (1997b). "Intraday Seasonality and Volatility Persistence in Foreign Exchange and Equity Markets," *Journal of Empirical Finance*, 52, 115-158.

4. Andersen, T., and T. Bollerslev (1998a). "DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer-Run Dependencies," *Journal of Finance*, 53, 219-265.

5. Andersen, T., and T. Bollerslev (1998b). "Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts," *International Economic Review*," 39, 885-905.

6. Andersen, T., and T. Bollerslev, P.F. Christoffersen, and F.X. Diebold (2005). "Practical Volatility and Correlation Modeling for Financial Market Risk Management," NBER Working Paper No. 11069.

7. Andersen, T., and T. Bollerslev, P.F. Christoffersen, and F.X. Diebold (2005). "Volatility Forecasting," NBER Working Paper No. 11188,

8. Andersen, T., T. Bollerslev, and F.X. Diebold (2004). "Parametric and Nonparametric Volatility Measurement," in *Handbook of Financial Econometrics*, ed. by L.P. Hansen and Y. A-Sahalia. Amsterdam: North Holland, forthcoming.

9. Andersen, T., T. Bollerslev, F.X. Diebold, H. Ebens (2001). "The Distribution of Realized Stock Return Volatility," *Journal of Financial Economics*, 61, 43-76.

10. Andersen, T., T. Bollerslev, F.X. Diebold, P. Labys (2000a). "Exchange Rate Returns Standardized by Realized Volatility Are (Nearly) Gaussian," *Multinational Finance Journal*, 4, 159-179.

11. Andersen, T., T. Bollerslev, F.X. Diebold, P. Labys (2000a). "Great Realizations," *Risk*, 13, 105-108.

12. Andersen, T., T. Bollerslev, F.X. Diebold, P. Labys (2001). The Distribution of Realized Exchange Rate Volatility, *Journal of the American Statistical Association* 96, 42-55.

13. Andersen, T., T. Bollerslev, F.X. Diebold, P. Labys (2003). "Modeling and Forecasting Realized Volatility," *Econometrica*, 71(2), 579-626.

14. Andersen, T., T. Bollerslev, F.X. Diebold, and C. Vega (2004). "Real-Time Price Discovery in Stock, Bond and Foreign Exchange Markets," unpublished manuscript, Northwestern University, Duke University, University of Pennsylvania, and University of Rochester.

15. Andersen, T., T. Bollerslev, and N. Meddahi (2005). "Correcting the Errors: Volatility Forecast Evaluation Based on High Frequency Data and Realized Volatilities," *Econometrica*, 73(1), 279-296.

16. Andreou, E. and E. Ghysels (2002). "Rolling-Sample Volatility Estimators: Some New Theoretical, Simulation and Empirical Results," *Journal of Business and Economic Statistics*, 20, 363-376.

17. Andreou, E. and E. Ghysels (2002). "Detecting Multiple Breaks in Financial Market Volatility Dynamics," *Journal of Applied Econometrics*, Vol. 17, No. 5, 2002, pp. 579-600.

18. Bai, X. J.R. Russell, and G.C. Tiao (2000). "Beyond Merton's Utipia: Effects of Non-normality and Dependence on the Precision of Variance Estimates Using High-Frequency Financial Data," manuscript, Graduate School of Business, University of Chicago.

19. Bandi, F.M., and J.R. Russell (2003). "Microstructure Noise, Realized Volatility, and Optimal Sampling," manuscript, Graduate School of Business, University of Chicago.

20. Barndorff-Nielsen, O.E., and N. Shephard (2002a). "Estimating Quadratic Variation Using Realized Variance," *Journal of Applied Econometrics*, 17, 457-477.

21. Barndorff-Nielsen, O.E., and N. Shephard (2002b). "Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society, Series B*, 64, 253-280.

22. Barndorff-Nielsen, O.E., and N. Shephard (2004a). "Econometric Analysis of Realized Covariation: High Frequency Based Covariance, Regression, and Correlation in Financial Economics," *Econometrica*, 73(3), 885-926.

23. Barndorff-Nielsen, O.E., and N. Shephard (2004b). "How Accurate Is the Asymptotic Approximation to the distribution of Realized Volatility?," in *Identification and Inference for Econometric Models. a Festschrift in Honour of T.J. Rothenberg*, ed. by D.W.K. Andrews, J. Powell, P.A. Ruud, and J.H. Stock. Cambridge: Cambridge University Press.

24. Bollerslev, T. and H. Zhou (2001). "Estimating Stochastic Volatility Diffusion Using Conditional Moments of Integrated Volatility," *Journal of Econometrics*.

25. Geweke, J., and S. Porter-Hudak (1983). "The Estimation and Application of Long Memory Time-Series," *Journal of Time Series Analysis*, 4, 221-238.

26. Christensen, B.J., and N.R. Prabhala (1998). "The Relation Between Implied and Realized Volatility," *Journal of Financial Economics*, 37, 125-150.

27. Giot, P. and S. Laurent (2003). "Modelling Daily Value-at-Risk Using Realized Volatility and ARCH Type Models," unpublished manuscript, Universite Catholique de Louvain, Belgium.

28. Hansen, P.R., and A. Lunde (2005). "A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data," unpublished manuscript, Department of Economics, Stanford University.

29. Meddahi, N. (2002). "A Theoretical Comparison Between Integrated and Realized Volatilities," *Journal of Applied Econometrics*, 17, 479-508.

30. Meddahi, N. (2003). "ARMA Representation of Integrated and Realized Variances," *Econometrics Journal*, 6, 334-355.

31. Maheu, J.M. and T.H. McCurdy (2002). "Nonlinear Features of Realized FX Volatility," *Review of Economics and Statistics*, 84, 668-681.

32. Merton, R. (1980). "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323-361.

33. Schwert, G.W. (1989). "Why Does Stock Market Volatility Change Over Time?", *Journal of Finance*, 44, 1115-1153.

34. Taylor, S.J., and X. Xu (1997). "The Incremental Volatility Information in One Million Foreign Exchange Quotations," *Journal of Empirical Finance* 4, 317-340.

35. Zivot, E. and J. Wang (2005). *Modeling Financial Time Series with S-PLUS, Second Edition.* Springer-Verlag forthcoming.