MODELING AND FORECASTING REALIZED VOLATILITY*

by Torben G. Andersen[a], Tim Bollerslev[b], Francis X. Diebold[c] and Paul Labys[d]

First Draft: January 1999
This Version: January 2001

This paper provides a general framework for integration of high-frequency intraday data into the measurement, modeling, and forecasting of daily and lower frequency volatility and return distributions. Most procedures for modeling and forecasting financial asset return volatilities, correlations, and distributions rely on restrictive and complicated parametric multivariate ARCH or stochastic volatility models, which often perform poorly at intraday frequencies. Use of realized volatility constructed from high-frequency intraday returns, in contrast, permits the use of traditional time series procedures for modeling and forecasting. Building on the theory of continuous-time arbitrage-free price processes and the theory of quadratic variation, we formally develop the links between the conditional covariance matrix and the concept of realized volatility. Next, using continuously recorded observations for the Deutschemark / Dollar and Yen / Dollar spot exchange rates covering more than a decade, we find that forecasts from a simple long-memory Gaussian vector autoregression for the logarithmic daily realized volatilities perform admirably compared to popular daily ARCH and related models. Moreover, the vector autoregressive volatility forecast, coupled with a parametric lognormal-normal mixture distribution implied by the theoretically and empirically grounded assumption of normally distributed standardized returns, gives rise to well-calibrated density forecasts of future returns, and correspondingly accurate quantile estimates. Our results hold promise for practical modeling and forecasting of the large covariance matrices relevant in asset pricing, asset allocation and financial risk management applications.

KEYWORDS: Continuous-time methods, quadratic variation, realized volatility, realized correlation, high-frequency data, exchange rates, vector autoregression, long memory, volatility forecasting, correlation forecasting, density forecasting, risk management, value at risk.

---

* This paper supercedes the earlier manuscript "Forecasting Volatility: A VAR for VaR." The work reported in the paper was supported by the National Science Foundation. We are grateful to Olsen and Associates, who generously made available their intraday exchange rate quotation data. For insightful suggestions and comments we thank Rob Engle, Atsushi Inoue, Neil Shephard, Clara Vega, Sean Campbell, and seminar participants at Chicago, Michigan, Montreal/CIRANO, NYU, Rice, and the June 2000 Meeting of the Western Finance Association.

[a] Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, and NBER, phone: 847-467-1285, e-mail: t-andersen@kellogg.northwestern.edu

[b] Department of Economics, Duke University, Durham, NC 27708, and NBER, phone: 919-660-1846, e-mail: boller@econ.duke.edu

[c] Department of Economics, University of Pennsylvania, Philadelphia, PA 19104, and NBER, phone: 215-898-1507, e-mail: fdiebold@sas.upenn.edu

[d] Graduate Group in Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104, phone: 435-753-9671, e-mail: labys@ssc.sas.upenn.edu

1. INTRODUCTION

The joint distributional characteristics of asset returns are pivotal for many issues in financial economics. They are the key ingredients for the pricing of financial instruments, and they speak directly to the risk-return tradeoff critical for portfolio allocation, performance evaluation, and managerial decisions. Moreover, they are intimately related to the conditional portfolio return fractiles, which govern the likelihood of extreme shifts in portfolio value and therefore central to financial risk management, figuring prominently in both regulatory and private-sector initiatives.

The most critical feature of the conditional return distribution is arguably its second moment structure, which is empirically the dominant time-varying characteristic of the distribution. This fact has spurred an enormous literature on the modeling and forecasting of return volatility.[1] Over time, the availability of data for increasingly shorter return horizons has allowed the focus to shift from modeling at quarterly and monthly frequencies to the weekly and daily horizons. Along with the incorporation of more data has come definite improvements in performance, not only because the models now may produce forecasts at the higher frequencies, but also because they typically provide superior forecasts for the longer monthly and quarterly horizons than do the models exploiting only monthly data.

Progress in volatility modeling has, however, in some respects slowed over the last decade. First, the availability of truly high-frequency intraday data has made scant impact on the modeling of, say, daily return volatility. It has become apparent that standard volatility models used for forecasting at the daily level cannot readily accommodate the information in intraday data, and models specified directly for the intraday data generally fail to capture the longer interdaily volatility movements sufficiently well. As a result, standard practice is still to produce forecasts of daily volatility from daily return observations, even when higher-frequency data are available. Second, the focus of volatility modeling continues to be decidedly low-dimensional, if not universally univariate. Many multivariate ARCH and stochastic volatility models for time-varying return volatilities and conditional distributions have, of course, been proposed; see, Bollerslev, Engle and Nelson (1994), Ghysels, Harvey and Renault (1996), and Kroner and Ng (1998), but those models generally suffer from a curse-of-dimensionality problem that severely constrains their practical application. Consequently, it is rare to see practical applications of such procedures dealing with more than a few assets simultaneously.

In view of such difficulties, finance practitioners have largely eschewed formal volatility modeling and forecasting in the higher-dimensional situations of practical relevance, relying instead on

[1] Here and throughout, we use the generic term “volatilities” in reference both to variances (or standard deviations) and covariances (or correlations). When important, the precise meaning will be clear from context.

simple exponential smoothing methods for construction of volatility forecasts, coupled with an assumption of conditionally normally distributed returns. This approach is exemplified by J.P. Morgan's highly influential RiskMetrics, see J.P. Morgan (1997). Although such methods exploit outright counterfactual assumptions and almost certainly are suboptimal, such defects must be weighed against considerations of feasibility, simplicity and speed of implementation in high-dimensional environments.

Set against this background, we seek improvement along two important dimensions. First, we propose a new rigorous procedure for volatility forecasting and return fractile, value-at-risk (VaR), calculation that efficiently exploits the information in intraday return observations. In the process, we document significant improvements in predictive performance relative to the standard procedures that rely on daily data alone. Second, our methods achieve a simplicity and ease of implementation that allows for ready accommodation of higher-dimensional return systems. We achieve these dual objectives by focusing on an empirical measure of daily return variability termed *realized volatility*, which is easily computed from high-frequency intra-period returns. The theory of quadratic variation reveals that, under suitable conditions, realized volatility is not only an unbiased ex-post estimator of daily return volatility, but also asymptotically free of measurement error, as discussed in Andersen, Bollerslev, Diebold and Labys (2001a) (henceforth ABDL) as well as concurrent work by Barndorff-Nielsen and Shephard (2000, 2001). Building on the notion of continuous-time arbitrage-free price processes, we progress in several directions, including more rigorous theoretical foundations, multivariate emphasis, and links to modern risk management.

Empirically, by treating the volatility as observed rather than latent, our approach greatly facilitates modeling and forecasting using simple methods based directly on observable variables.[2] Although the basic ideas apply quite generally, we focus on the highly liquid U.S. dollar ($), Deutschemark (DM), and Japanese yen (¥) spot exchange rate markets in order to illustrate and evaluate our methods succinctly under conditions that allow for construction of good realized volatility measures. Our full sample consists of nearly thirteen years of continuously recorded spot quotations from 1986 through 1999. During this period, the dollar, Deutschemark and yen constituted the main axes of the international financial system, and thus spanned the majority of the systematic currency risk faced by most large institutional investors and international corporations.

---

[2] Earlier empirical work exploiting related univariate approaches from a heuristic perspective includes French, Schwert and Stambaugh (1987) and Schwert (1989), who rely on daily returns to estimate models for monthly realized U.S. equity volatility, and Hsieh (1991), who fits an AR(5) model to a time series of daily realized logarithmic volatilities constructed from 15-minute S&P500 returns.

We break the sample into a ten year "in-sample" estimation period, and a subsequent two and a half year "out-of-sample" forecasting period. The basic distributional and dynamic characteristics of the foreign exchange returns and realized volatilities during the in-sample period have been analyzed in detail by ABDL (2001a, 2001b).[3] Three pieces of their results form the foundation on which the empirical analysis of this paper is built. First, although raw returns are clearly leptokurtic, returns standardized by realized volatilities are approximately Gaussian. Second, although the distributions of realized volatilities are clearly right-skewed, the distributions of the logarithms of realized volatilities are approximately Gaussian. Finally, the long-run dynamics of realized logarithmic volatilities are well approximated by a fractionally-integrated long-memory process.

Motivated by the three ABDL empirical regularities, we proceed to estimate and evaluate a multivariate fractionally-integrated Gaussian vector autoregression (VAR) for the logarithmic realized volatilities. Comparing the resulting volatility forecasts to those obtained from daily ARCH and related models, we find our simple Gaussian VAR forecasts to be strikingly superior. Furthermore, we show that, given the theoretically motivated and empirically plausible assumption of normally distributed returns conditional on the realized volatilities, the resulting lognormal-normal mixture forecast distribution gives rise to well-calibrated density forecasts of returns, from which highly accurate estimates of return quantiles may be derived.

The rest of the paper is organized as follows. Section 2 develops the theory behind the notion of realized volatility. Section 3 focuses on measurement of realized volatilities using high-frequency foreign exchange returns. Next, Section 4 summarizes the salient distributional features of the returns and volatilities, which motivate the long-memory trivariate Gaussian VAR introduced in Section 5. Section 6 compares the resulting volatility forecasts to those obtained from traditional GARCH and related models, and Section 7 evaluates the success of density forecasts and corresponding VaR estimates generated from our long-memory Gaussian VAR in conjunction with a lognormal-normal mixture distribution. Section 8 concludes with suggestions for future research and discussion of issues related to the practical implementation of our approach for other financial instruments and markets.

2. QUADRATIC RETURN VARIATION AND REALIZED VOLATILITY

We consider a price process defined on a complete probability space, $(\Omega, \mathcal{F}, P)$, evolving in continuous

[3] Strikingly similar qualitative findings have been obtained from a separate sample consisting of individual U.S. stock returns in Andersen, Bollerslev, Diebold and Ebens (2001).

time over the interval $[0,T]$, where T denotes a positive integer. We further consider an information filtration, i.e., an increasing family of σ-fields, $(\mathscr{F}_t)_{t\in[0,T]} \subseteq \mathscr{F}$, which satisfies the usual conditions of P-completeness and right continuity. Finally, we assume that the asset prices through time t, including the relevant state variables, are known at time t and therefore included in the information set $\mathscr{F}_t$.

Under the standard assumptions that the return process does not allow for arbitrage and has a finite instantaneous mean, the asset price process, as well as smooth transformations thereof, belongs to the class of special semi-martingales, as detailed by Back (1991). A fundamental result of modern stochastic integration theory states that such processes permit a unique canonical decomposition.[4] In particular, we have the following characterization of the logarithmic asset price vector process, $p = (p(t))_{t\in[0,T]}$.

PROPOSITION 1: *For any n-dimensional arbitrage-free vector price process with finite mean, the associated logarithmic vector price process, p, may be written uniquely as the sum of a finite variation and predictable component, A, and a local martingale,* $M = (M_1 , \ldots , M_n)$. *The latter may be further decomposed into a continuous sample path local martingale,* M^c, *and a compensated jump martingale,* ΔM, *with the initial conditions* $M(0) = A(0) = 0$, *so that*

$$p(t) = p(0) + A(t) + M(t) = p(0) + A(t) + M^c(t) + \Delta M(t). \tag{1}$$

Proposition 1 provides a general qualitative characterization of the asset return process. We denote the (continuously compounded) return over $[t-h,t]$ by $r(t,h) = p(t) - p(t-h)$. The cumulative return process from $t=0$ onwards, $r = (r(t))_{t\in[0,T]}$, is then given as $r(t) \equiv r(t,t) = p(t) - p(0) = A(t) + M(t)$. Clearly, $r(t)$ inherits all the main properties of $p(t)$, and it may likewise be decomposed uniquely into the predictable and integrable mean component, A, and the local martingale, M.

Because the return process is a semi-martingale it has an associated quadratic variation process. This notion plays a critical role in our theoretical developments. The following proposition enumerates some essential properties of the quadratic return variation process.[5]

PROPOSITION 2: *For any n-dimensional arbitrage-free price process with finite mean, the quadratic*

[4] See, for example, Protter (1992), chapter 3.

[5] All of the properties in Proposition 2 follow, for example, from Protter (1992), chapter 2.

variation nxn matrix process of the associated return process, $[r,r] = \{ [r,r]_t \}_{t \in [0,T]}$, is well defined. The i'th diagonal element is called the quadratic variation process of the i'th asset return while the ij'th off-diagonal element, $[r_i , r_j]$, is termed the quadratic covariation process between asset returns i and j. Moreover, we have the following properties:

(i) *For an increasing sequence of random partitions of [0,T], $0 = \tau_{m,0} \leq \tau_{m,1} \leq \ldots$, such that $sup_{j\geq 1}(\tau_{m,j+1} - \tau_{m,j}) \rightarrow 0$ and $sup_{j\geq 1} \tau_{m,j} \rightarrow T$ for $m \rightarrow \infty$ with probability one, we have that*

$$lim_{m\rightarrow\infty} \{ \Sigma_{j\geq 1} [r(t\wedge\tau_{m,j}) - r(t\wedge\tau_{m,j-1})] [r(t\wedge\tau_{m,j}) - r(t\wedge\tau_{m,j-1})]' \} \rightarrow [r,r]_t , \qquad (2)$$

where $t \wedge \tau \equiv min(t,\tau)$, $t \in [0,T]$, and the convergence is uniform on [0,T] in probability.

(ii) $$[r_i,r_j]_t = [M_i ,M_j]_t = [M_i^c ,M_j^c]_t + \Sigma_{0\leq s\leq t} \Delta M_i(s) \Delta M_j(s) . \qquad (3)$$

The terminology of quadratic variation is justified by property (i) of Proposition 2. The quadratic variation process measures the realized sample-path variation of the squared return processes. Notice also that it suggests we may approximate the quadratic variation by cumulating cross-products of high-frequency returns. We refer to such measures, obtained from actual high-frequency data, as *realized volatility*. Property (ii) reflects the fact that quadratic variation of finite variation processes is zero, so the mean component is irrelevant for the quadratic variation. Moreover, jump components only contribute to the quadratic covariation if there are simultaneous jumps in the price path for the i'th and j'th asset, whereas the squared jump size contributes one-for-one to the quadratic variation.

The quadratic variation is the dominant determinant of the return covariance matrix, especially for shorter horizons. The reason is that the variation induced by the genuine return innovations, represented by the martingale component, locally is an order of magnitude larger than the return variation caused by changes in the conditional mean. Consequently, we have the following theorem which generalizes previous results in Andersen, Bollerslev, Diebold and Labys (2001a).

THEOREM 1: *Let an n-dimensional square-integrable arbitrage-free logarithmic price process with a unique canonical decomposition, as stated in equation (1), be given. The conditional return covariance matrix at time t for returns over [t, t+h], where $0 \leq t \leq t+h \leq T$, equals*

$$Cov(r(t+h,h)/\mathscr{F}_t) = E([r,r]_{t+h} - [r,r]_t /\mathscr{F}_t) + \Gamma_A (t+h,h) + \Gamma_{AM} (t+h,h) + \Gamma_{AM}'(t+h,h), \qquad (4)$$

where $\Gamma_A (t+h,h) = Cov(A(t+h) - A(t) /\mathscr{F}_t)$ and $\Gamma_{AM} (t+h,h) = E(A(t+h) [M(t+h) - M(t)]' /\mathscr{F}_t)$.

PROOF: From equation (1), $r(t+h,h) = [A(t+h) - A(t)] + [M(t+h) - M(t)]$. The martingale property implies $E(M(t+h) - M(t) / \mathscr{F}_t) = E([M(t+h) - M(t)] A(t) / \mathscr{F}_t) = 0$, so, for $i,j \in \{1, ..., n\}$, $Cov([A_i(t+h) - A_i(t)], [M_j(t+h) - M_j(t)] / \mathscr{F}_t) = E(A_i(t+h) [M_j(t+h) - M_j(t)] / \mathscr{F}_t)$. Exploiting these results, it follows that $Cov(r(t+h,h) / \mathscr{F}_t) = Cov(M(t+h) - M(t) / \mathscr{F}_t) + \Gamma_A(t+h,h) + \Gamma_{AM}(t+h,h) + \Gamma_{AM}'(t+h,h)$. Hence, it only remains to show that the conditional covariance of the martingale term equals the expected value of the quadratic variation. We proceed by verifying the equality for an arbitrary element of the covariance matrix. If this is the i'th diagonal element, we are studying a univariate square-integrable martingale and by Protter (1992), chapter II.6, corollary 3, we have $E[M_i^2(t+h)] = E([M_i ,M_i]_{t+h})$, so $Var(M_i(t+h) - M_i(t) / \mathscr{F}_t) = E([M_i ,M_i]_{t+h} - [M_i ,M_i]_t / \mathscr{F}_t) = E([r_i ,r_i]_{t+h} - [r_i ,r_i]_t / \mathscr{F}_t)$, where the second equality follows from equation (3) of Proposition 2. This confirms the result for the diagonal elements of the covariance matrix. An identical argument works for the off-diagonal terms by noting that the sum of two square-integrable martingales remains a square-integrable martingale and then applying the reasoning to each component of the polarization identity, $[M_i ,M_j]_t = ½ ([M_i+M_j , M_i+M_j]_t - [M_i ,M_i]_t - [M_j ,M_j]_t)$. In particular, it follows as above that $E([M_i ,M_j]_{t+h} - [M_i ,M_j]_t / \mathscr{F}_t) = ½ [Var([M_i(t+h)+M_j(t+h)] - [(M_i(t)+M_j(t)] / \mathscr{F}_t) - Var(M_i(t+h) - M_i(t) / \mathscr{F}_t) - Var(M_j(t+h) - M_j(t) / \mathscr{F}_t)]= Cov([M_i(t+h) - M_i(t)], [M_j(t+h) - M_j(t)] / \mathscr{F}_t)$. Equation (3) of Proposition 2 again ensures that this equals $E([r_i ,r_j]_{t+h} - [r_i ,r_j]_t / \mathscr{F}_t)$. □

A couple of scenarios highlight the role of the quadratic variation in driving the return volatility process. These important special cases are collected in a corollary which follows immediately from Theorem 1.

COROLLARY: *Let an n-dimensional square-integrable arbitrage-free logarithmic price process, as described in Theorem 1, be given. If the mean process,* $\{A(s) - A(t)\}_{s\in[t,t+h]}$ *, conditional on information at time t is independent of the return innovation process,* $\{M(u)\}_{u\in[t,t+h]}$ *, then the conditional return covariance matrix reduces to the conditional expectation of the quadratic return variation plus the conditional variance of the mean component, i.e.,*

$$Cov(r(t+h,h) / \mathscr{F}_t) = E([r,r]_{t+h} - [r,r]_t / \mathscr{F}_t) + \Gamma_A(t+h,h), \qquad (5)$$

where $0 \le t \le t+h \le T$. *If the mean process,* $\{A(s) - A(t)\}_{s\in[t,t+h]}$*, conditional on information at time t is a predetermined function over [t, t+h], then the conditional return covariance matrix equals the conditional expectation of the quadratic return variation process, i.e.,*

$$Cov(r(t+h,h) / \mathscr{F}_t) = E([r,r]_{t+h} - [r,r]_t / \mathscr{F}_t), \qquad (6)$$

where $0 \le t \le t+h \le T$.

It is apparent, under the conditions leading to equation (6), that the quadratic variation is *the* critical ingredient in volatility measurement and forecasting. The conditional covariance matrix is simply given by the conditional expectation of the quadratic variation. Moreover, it follows that the ex-post realized quadratic variation is an unbiased estimator for the return covariance matrix *conditional on information at time t*. Although these conclusions may appear to hinge on restrictive assumptions, they apply to a wide set of models used in the literature. For example, a constant mean is frequently invoked in models for daily or weekly asset returns. Equation (6) further allows for deterministic intra-period variation in the conditional mean process, induced, e.g., by time-of-day or other calendar type effects. Of course, the specification in (6) also accommodates a stochastic evolution of the mean process as long as it remains a function, over the interval $[t\text{-}h, t]$, of variables that belong to the information set at time $t\text{-}h$. What is precluded are feedback effects from the random intra-period evolution of the system to the instantaneous mean. Although this may be counter-factual, such effects are likely trivial in magnitude, as discussed below. It is also worth stressing that equation (6) is compatible with the existence of a so-called leverage, or asymmetric return-volatility, relation. The latter arises from a correlation between the return innovations - measured as deviations from the conditional mean - and the innovations to the volatility process. Hence, the leverage effect does not require contemporaneous correlation between the return innovations and the instantaneous mean return. And, as emphasized above, the formulation (6) does allow for the return innovations over $[t\text{-}h, t]$ to impact the conditional mean over $[t, t+h]$ and onwards, so that the intra-period evolution of the system still may impact the future expected returns. In fact, this is how potential interaction between risk and return is captured within discrete-time ARCH or stochastic volatility models that incorporate leverage effects.

In contrast to equation (6), equation (5) does accommodate continually evolving random variation in the conditional mean process, although it must be independent of the return innovations. However, even with this feature present, the quadratic variation is likely an order of magnitude larger than the mean variation, and hence the former remains the critical determinant of the return volatility over shorter horizons. This observation follows from the fact that, locally, over horizons of length h, with h small, the mean return is of order h, and the variance of the mean return thus of order h^2, while the quadratic variation is of order h. It is obviously an empirical question whether these results are a good guide for volatility measurement at practically relevant frequencies.[6] To illustrate the likely implications

[6] Merton (1982) provides a similar intuitive account of the continuous record *h-asymptotics*. These limiting results are also closely related to the theory rationalizing the quadratic variation formulas in Proposition 2 and Theorem 1.

at a daily horizon, consider an asset return with (typical) standard deviation of 1% daily, or 15.8% annually, and a (large) mean return of 0.1%, or about 25% annually. The squared mean return is still only one-hundredth of the variance. The expected daily variation of the mean return is obviously smaller yet, unless the required daily return is expected to behave truly erratically within the day. In fact, we would generally expect the within-day variance of the expected daily return to be much smaller than the expected daily return itself. Hence, the daily return fluctuations induced by within-day variations in the required mean return are almost certainly trivial. Even for a weekly horizon, similar calculations suggest that the identical conclusion applies.

The general case, covered by Theorem 1, allows for direct intra-period interaction between the return innovations and the instantaneous mean. This occurs, for example, in the arguably empirically relevant scenario where there is a leverage effect, or asymmetry, by which the volatility impacts the contemporaneous mean drift. In this setting, a string of negative within-period return innovations will be associated with an increase in return volatility and this may in turn raise the risk premium and induce a larger return drift. Relative to the corollary, the theorem involves an additional set of Γ_{AM} terms. Nonetheless, the results and intuition discussed above survive. It is readily established that the *ik*'th component of these terms may be bounded as, $\{\Gamma_{AM}(t+h,h)\}_{i,k} \leq \{\mathrm{Var}(A_i(t+h) - A_i(t) \,/\, \mathscr{F}_t)\}^{1/2}$ $\{\mathrm{Var}(M_k(t+h) - M_k(t) \,/\, \mathscr{F}_t)\}^{1/2}$, but the latter terms are of order h and $h^{1/2}$ respectively, so the Γ_{AM} terms are at most of order $h^{3/2}$, which again is dominated by the corresponding quadratic variation of order h. Moreover, this upper bound is quite conservative, since it allows for a correlation of unity, whereas the typical correlation estimated from daily or weekly returns is much lower, de facto implying that the quadratic variation process is the main driving force behind the high-frequency return volatility.

We now turn towards an even more ambitious goal. Accepting that the above results carry implications for the measurement and modeling of return volatility, it is natural to ask whether we also can infer something about the appropriate specification of the return generating process that builds on the realized volatility measures. Obviously, at the level of generality that we are operating at so far - requiring only square integrability and absence of arbitrage - we cannot derive specific distributional results. However, it turns out that we may obtain a useful benchmark under somewhat restrictive conditions, including a continuous price process, i.e., no jumps or $\Delta M \equiv 0$. We first recall the martingale representation theorem.[7]

[7] See, for example, Karatzas and Shreve (1991), chapter 3.

PROPOSITION 3: *For any n-dimensional square-integrable arbitrage-free logarithmic price process, p, with continuous sample path and a full rank of the associated nxn quadratic variation process,* $[r,r]_t$ *we have a.s.(P) for all* $0 \leq t \leq T$,

$$r(t+h,h) = p(t+h) - p(t) = \int_0^h \mu_{t+s}\, ds + \int_0^h \sigma_{t+s}\, dW(s) \,, \tag{7}$$

where μ_s *denotes an integrable predictable nx1 dimensional vector,* $\sigma_s = (\sigma_{(i,j),s})_{i,j=1,\ldots,n}$ *is a nxn matrix, W(s) is a nx1 dimensional standard Brownian motion, and integration of a matrix (vector) w.r.t. a scalar denotes component-wise integration, e.g., the mean component is the nx1 vector,*

$$\int_0^h \mu_{t+s}\, ds = \left(\int_0^h \mu_{1,t+s}\, ds, \ldots, \int_0^h \mu_{n,t+s}\, ds \right)',$$

and integration of a matrix w.r.t. a vector denotes component-wise integration of the associated vector,

$$\int_0^h \sigma_{t+s}\, dW(s) = \left(\int_0^h \Sigma_{j=1,\ldots,n}\ \sigma_{(1,j),t+s}\, dW_j(s), \ldots, \int_0^h \Sigma_{j=1,\ldots,n}\ \sigma_{(n,j),t+s}\, dW_j(s) \right)'. \tag{8}$$

Moreover, we have

$$P\left[\int_0^h (\sigma_{(i,j),t+s})^2\, ds < \infty \right] = 1, \quad 1 \leq i, j \leq n. \tag{9}$$

Finally, letting $\Omega_s = \sigma_s \sigma_s'$, *the increments to the quadratic return variation process take the form*

$$[r,r]_{t+h} - [r,r]_t = \int_0^h \Omega_{t+s}\, ds \,. \tag{10}$$

The condition of Proposition 3 that the *nxn* matrix $[r,r]_t$ is of full rank for all *t*, implies that no asset is redundant at any time, so that no individual asset return can be spanned by a portfolio created by the remaining assets. This condition is not restrictive; if it fails, a parallel representation may be achieved on an extended probability space.[8]

We are now in position to state a distributional result that inspires our empirical modeling of the full return generating process in Section 7. It extends a result recently noted by Barndorff-Nielsen and Shephard (2000) by allowing for a more general specification of the conditional mean process and, more importantly, accommodating a multivariate setting. It should be noted that if we only focus on volatility forecasting, as in Sections 5 and 6 below, we do not need the auxiliary assumptions invoked here.

THEOREM 2: *For a n-dimensional square-integrable arbitrage-free price process with continuous sample paths satisfying Proposition 3, and thus representation (7), with conditional mean and volatility processes,* μ_s *and* σ_s, *that are independent of the innovation process, W(s), over [t,t+h], we have*

$$r(t+h,h) \,/\, \sigma\{ \mu_{t+s}, \sigma_{t+s} \}_{s\in[0,h]} \sim N\left(\int_0^h \mu_{t+s}\, ds, \int_0^h \Omega_{t+s}\, ds \right), \tag{11}$$

where $\sigma\{ \mu_{t+s}, \sigma_{t+s} \}_{s\in[0,h]}$ *denotes the* σ*-field generated by* $(\mu_{t+s}, \sigma_{t+s})_{s\in[0,h]}$.

---

[8] See Karatzas and Shreve (1991), section 3.4.

PROOF: Clearly, $r(t+h,h) - \int_0^h \mu_{t+s}\, ds = \int_0^h \sigma_{t+s}\, dW(s)$ and $E(\int_0^h \sigma_{t+s}\, dW(s) / \sigma\{ \mu_{t+s}, \sigma_{t+s} \}_{s\in[0,h]}) = 0$. We proceed by establishing the normality of $\int_0^h \sigma_{t+s}\, dW(s)$ conditional on the volatility path $\{ \sigma_{t+s} \}_{s\in[0,h]}$. The integral is n-dimensional, and we define $\int_0^h \sigma_{t+s}\, dW(s) = (\int_0^h (\sigma_{(1),t+s})' dW(s), .., \int_0^h (\sigma_{(n),t+s})' dW(s))'$, where $\sigma_{(i),s} = (\sigma_{(i,1),s}, ... , \sigma_{(i,n),s})'$, so that $\int_0^h (\sigma_{(i),t+s})' dW(s)$ denotes the i'th element of the $nx1$ vector in equation (8). The vector is multivariate normal, if and only if any linear combination of the elements are univariate normal but this follows readily if each element of the vector is univariate normal. From equation (8), each element of the vector is a sum of integrals and hence will be normally distributed if each component of the sum is univariate normal conditional on the volatility path. This is what we establish next. A typical element of the sums in equation (8), representing the j'th volatility factor loading of asset i over $[t,t+h]$, takes the form, $I_{i,j}(t+h,h) = \int_0^h \sigma_{(i,j),t+s}\, dW_j(s)$, for $1 \leq i, j \leq n$. Obviously, $I_{i,j}(t) \equiv I_{i,j}(t,t)$ is a continuous local martingale, and then by the "change of time" result, see, e.g., Protter (1992), Chapter II, Theorem 41, it follows that $I_{i,j}(t) = B([I_{i,j}, I_{i,j}]_t)$, where $B(t)$ denotes a standard univariate Brownian motion. Further, we have $I_{i,j}(t+h,h) = I_{i,j}(t+h) - I_{i,j}(t) = B([I_{i,j}, I_{i,j}]_{t+h}) - B([I_{i,j}, I_{i,j}]_t)$, and this increment to the Brownian motion is distributed $N(0, [I_{i,j}, I_{i,j}]_{t+h} - [I_{i,j}, I_{i,j}]_t)$. Finally, the quadratic variation governing the variance of the Gaussian distribution above is readily determined to be $[I_{i,j}, I_{i,j}]_{t+h} - [I_{i,j}, I_{i,j}]_t = \int_0^h (\sigma_{(i,j),t+s})^2\, ds$, see, e.g., Protter (1992), Chapter II.6, which is finite by equation (9) of Proposition 3. Conditional on the ex-post realization of the volatility path, the quadratic variation is given (measurable), and the conditional normality of $I_{i,j}(t+h,h)$ follows. Since both the mean and the volatility paths are independent of the return innovations over [t,t+h], the mean is readily determined from the first line of the proof. This verifies the conditional normality asserted in equation (11). The only remaining issue is to identify the conditional return covariance matrix. For the ik'th element of the matrix we have

$$
\begin{aligned}
Cov[\int_0^h (\sigma_{(i),t+s})' dW(s) , \int_0^h (\sigma_{(k),t+s})' dW(s) / \sigma\{ \mu_{t+s}, \sigma_{t+s} \}_{s\in[0,h]}] \\
&= E[\Sigma_{j=1,..,n} \int_0^h \sigma_{(i,j),t+s}\, dW_j(s) \cdot \Sigma_{j=1,..,n} \int_0^h \sigma_{(k,j),t+s}\, dW_j(s) / \sigma\{ \mu_{t+s}, \sigma_{t+s} \}_{s\in[0,h]}] \\
&= \Sigma_{j=1,..,n} E[\int_0^h \sigma_{(i,j),t+s}\, \sigma_{(k,j),t+s}\, ds / \sigma\{ \mu_{t+s}, \sigma_{t+s} \}_{s\in[0,h]}] \\
&= \Sigma_{j=1,..,n} \int_0^h \sigma_{(i,j),t+s}\, \sigma_{(k,j),t+s}\, ds \\
&= \int_0^h (\sigma_{(i),t+s})' \sigma_{(k),t+s}\, ds \\
&= (\int_0^h \sigma_{t+s} (\sigma_{t+s})' ds)_{ik} \\
&= (\int_0^h \Omega_{t+s}\, ds)_{ik} .
\end{aligned}
$$

This confirms that each element of the conditional return covariance matrix equals the corresponding element of the variance term indicated in equation (11). □

Notice that the distributional characterization in Theorem 2 is conditional on the ex-post sample-path realization of (μ_s, σ_s). Theorem 2 may thus appear to be of little practical relevance, because such realizations typically are unobservable. However, Proposition 2 and equation (10) suggest that we may construct approximate measures of the realized quadratic variation, and hence of the conditional return variance, directly from high-frequency return observations. In addition, as discussed previously, for daily or weekly returns, the conditional mean is largely negligible relative to the return innovations. Consequently, ignoring the time variation of the conditional mean, the daily returns, say, follow a Gaussian mixture distribution with the realized daily quadratic return variation governing the mixture.

From the auxiliary assumptions invoked in Theorem 2, the Gaussian mixture distribution is strictly only applicable if the price process has continuous sample paths and the volatility and mean processes are independent of the within-period return innovations. This raises two main concerns. First, some recent evidence suggests the possibility of discrete jumps in asset prices, rendering sample paths discontinuous.[9] On the other hand, the findings also tend to indicate that jumps are infrequent and have a jump size distribution about which there is little consensus. Second, for some asset classes there is evidence of leverage effects that may indicate a correlation between concurrent return and volatility innovations. However, as argued above, such contemporaneous correlation effects are likely quantitatively insignificant. Indeed, the theorem does allow for the more critical impact leading from the current return innovations to the volatility in subsequent periods, corresponding exactly to the effect captured in the related discrete-time literature. We thus retain the Gaussian mixture distribution as a natural starting point for empirical work. Obviously, if the realized volatility-standardized returns fail to be normally distributed, it may speak to the importance of incorporating jumps and/or contemporaneous return innovation-volatility interactions into the data generating process.

In summary, the arbitrage-free setting imposes a semi-martingale structure that leads directly to the representation in Proposition 1 and the associated quadratic variation in Proposition 2. In addition, equation (2) in Proposition 2 suggests a practical way to approximate the quadratic variation. Theorem 1 and the associated corollary reveal the intimate relation between the quadratic variation and the return volatility process. For the continuous sample path case, we further obtain the representation in equation (7), and the quadratic variation reduces by equation (10) to $\int_0^h \Omega_{t+s}\, ds$, which is often referred to as the *integrated volatility*. Theorem 2 consequently strengthens Theorem 1 by showing that the realized quadratic variation is not only a useful estimator of the ex-ante conditional volatility, but also, under

[9] See, for example, Andersen, Benzoni and Lund (2000), Bates (2000), Bakshi, Cao and Chen (1997), Pan (1999), and Eraker, Johannes and Polson (2000).

auxiliary assumptions, identical to the realized integrated return volatility over the relevant horizon. Moreover, it delivers a reference distribution for appropriately standardized returns. Combined these results provide a general framework for integration of high-frequency intraday data into the measurement and estimation of daily and lower frequency volatility and return distributions.

3. MEASURING REALIZED FOREIGN EXCHANGE VOLATILITY

Practical implementation of the procedures suggested by the theory in Section 2 must confront the fact that no financial market provides a frictionless trading environment with continuous price recording. Consequently, the notion of quadratic return variation is an abstraction that, strictly speaking, cannot be observed. Nevertheless, we may use the continuous-time arbitrage-free framework to motivate and guide the creation of return series and associated volatility measures from high-frequency data. We do not claim that this provides exact counterparts to the (non-existing) corresponding continuous-time quantities. Instead, we assess the usefulness of the theory through the lens of predictive accuracy. Specifically, theory guides the nature of the data collected, the way the data are transformed into volatility measures, and the model used to construct conditional return volatility and density forecasts.

3.1 *Data*

Our empirical analysis focuses on the spot exchange rates for the U.S. dollar, the Deutschemark and the Japanese yen.[10] The raw data consists of all interbank DM/$ and ¥/$ bid/ask quotes displayed on the Reuters FXFX screen during the sample period from December 1, 1986 through June 30, 1999.[11] These quotes are merely indicative (that is, non-binding) and subject to various market microstructure "frictions," including strategic quote positioning and standardization of the size of the quoted bid/ask spread. Such features are generally immaterial when analyzing longer horizon returns, but may distort the statistical properties of the underlying "equilibrium" high-frequency intraday returns. The sampling frequency at which such considerations become a concern is intimately related to market activity. For our exchange rate series, preliminary analysis based on the methods of ABDL (2000) suggests that the use of equally-spaced thirty-minute returns strikes a satisfactory balance between the accuracy of the

[10] Before the advent of the Euro, the dollar, Deutschemark and yen were the most actively traded currencies in the foreign exchange market, with the DM/$ and ¥/$ accounting for nearly fifty percent of the daily trading volume, according to a 1996 survey by the Bank for International Settlements.

[11] The data comprise several million quotes kindly supplied by Olsen & Associates. Average daily quotes number approximately 4,500 for the Deutschemark and 2,000 for the Yen.

continuous-record asymptotics underlying the construction of our realized volatility measures on the one hand, and the confounding influences from the market microstructure frictions on the other.[12]

The actual construction of the returns follows Müller et al. (1990) and Dacorogna et al. (1993). First, we calculate thirty-minute prices from the linearly interpolated logarithmic average of the bid and ask quotes for the two ticks immediately before and after the thirty-minute time stamps throughout the global 24-hour trading day. Second, we obtain thirty-minute returns as the first difference of the logarithmic prices.[13] In order to avoid modeling specific weekend effects, we exclude all of the returns from Friday 21:00 GMT until Sunday 21:00 GMT. Similarly, to avoid complicating the inference by the decidedly slower trading activity during certain holiday periods, we delete a number of other inactive days from the sample. We are left with a bivariate series of thirty-minute DM/$ and ¥/$ returns spanning a total of *3,045* days. In order to explicitly distinguish the empirically constructed continuously compounded discretely sampled returns, and corresponding volatility measures, from the theoretical counterparts in Section 2, we will refer to the former by time subscripts. Specifically for the half-hour returns $r_{t+\Delta,\Delta}$, $t = \Delta, 2\Delta, 3\Delta, ..., 3{,}045$, where $\Delta = 1/48 \approx 0.0208$. Also, for notational simplicity we label the corresponding daily returns by a single time subscript, so that $r_{t+1} \equiv r_{t+1,1} \equiv r_{t+\Delta,\Delta} + r_{t+2\Delta,\Delta} + ...\ r_{t+1,\Delta}$ for $t = 1, 2, ..., 3{,}045$. Finally, we partition the full sample period into an in-sample estimation period covering the *2,449* days from December 1, 1986 through December 1, 1996, and an out-of-sample forecast evaluation period covering the *596* days from December 2, 1996 through June 30, 1999.[14]

3.2 *Construction of Realized Volatilities*

The preceding discussion suggests that meaningful ex-post interdaily volatility measures may be constructed by cumulating cross-products of intraday returns sampled at an appropriate frequency, such as thirty minutes. In particular, based on the bivariate vector of thirty-minute DM/$ and ¥/$ returns, i.e., with $n = 2$, we define the *h-day realized volatility,* for $t = 1, 2, ..., 3045$, $\Delta = 1/48$, by

$$V_{t,h} \equiv \Sigma_{j=1,..,h/\Delta}\ r_{t-h+j\cdot\Delta,\Delta} \cdot r'_{t-h+j\cdot\Delta,\Delta} = R_{t,h}{'} R_{t,h}, \tag{12}$$

where the $(h/\Delta)xn$ matrix, $R_{t,h}$, is defined by $R_{t,h}{'} \equiv (r_{t-h+\Delta,\Delta}, r_{t-h+2\cdot\Delta,\Delta}, ... , r_{t,\Delta})$. As before, we simplify the

---

[12] An alternative approach would be to utilize all of the observations by explicitly modeling the high-frequency market microstructure. That approach, however, is much more complicated and subject to numerous pitfalls of its own.

[13] We follow the standard terminology of the interbank market by measuring the exchange rates and computing the corresponding rates of return from the prices of $1 expressed in terms of DM and ¥, i.e., DM/$ and ¥/$. Similarly, we express the cross rate as the price of one DM in terms of ¥, i.e., ¥/DM.

[14] All of the empirical results in ABDL (2001a, 2001b), which in part motivate our approach, are based on data for the in-sample period, justifying the claim that our forecast evaluation is "out-of-sample."

notation for the daily horizon by defining $V_t \equiv V_{t,1}$. The $V_{t,h}$ measure constitutes the empirical counterpart to the h-period quadratic return variation and, for the continuous sample path case, the integrated volatility. In fact, by Proposition 2, as the sampling frequency of intraday returns increases, or $\Delta \to 0$, $V_{t,h}$ converges almost surely to the quadratic variation. The same intuition underlies the continuous record asymptotics for the estimation of a time-invariant diffusion in Merton (1980) and the filtering results for continuous-time stochastic volatility models in Nelson and Foster (1995).

One obstacle frequently encountered when constructing conditional covariance matrix estimates from a finite set of return observations is that the estimator becomes non-positive definite. In fact, even for relatively low-dimensional cases, such as three or four assets, imposition and verification of conditions that guarantee positive definiteness of conditional covariance matrices can be challenging; see, e.g., the treatment of multivariate GARCH processes in Engle and Kroner (1995). Interestingly, it is straightforward to establish positive definiteness of our $V_{t,h}$ measure. The following proposition follows from the correspondence between our realized volatility measures and standard unconditional sample covariance matrix estimators which, of course, are positive semi-definite by construction.

PROPOSITION 4: *If the columns of $R_{t,h}$ are linearly independent, then $V_{t,h}$ is positive definite.*

PROOF: It suffices to show that $a' V_{t,h}\, a > 0$ for all non-zero a. Linear independence of the columns of $R_{t,h}$ ensures that $b_{t,h} = R_{t,h}\, a \neq 0,\ \forall a \in \mathbb{R}^n \backslash \{0\}$, and in particular that at least one of the elements of $b_{t,h}$ is non-zero. Hence $a' V_{t,h}\, a = a' R_{t,h}' R_{t,h}\, a = b_{t,h}'\, b_{t,h} = \Sigma_{j=1,\ldots,h/\Delta}\, (b_{t,h})_j^2 > 0,\ \forall a \in \mathbb{R}^n \backslash \{0\}$. □

The fact that positive definiteness of the volatility measure is virtually assured within high-dimensional applications is encouraging. However, the theorem also points to a problem that will arise for very high-dimensional systems. The assumption of linear independence of the columns of $R_{t,h}$, although weak, will ultimately be violated as the dimension of the price vector increases relative to the sampling frequency of the intraday returns. Specifically, for $n > h/\Delta$ the rank of the $R_{t,h}$ matrix is obviously less than n, so $R_t' R_t = V_t$ will not have full rank and it will fail to be positive definite. Hence, although the use of V_t facilitates rigorous measurement of conditional volatility in much higher dimensions than is feasible with most alternative approaches, it does not allow the dimensionality to become arbitrarily large. Concretely, the use of thirty-minute returns, corresponding to $1/\Delta = 48$ intraday observations, for construction of daily realized volatility measures, implies that positive definiteness of V_t requires n, the number of assets, to be no larger than 48.

The construction of an observable series for the realized volatility allows us to model the daily conditional volatility measure, V_t, using standard and relatively straightforward time series techniques. The diagonal elements of V_t, say $v_{t,1}$ and $v_{t,2}$, correspond to the daily DM/$ and ¥/$ realized variances, while the off-diagonal element, say $v_{t,12}$, represents the daily realized covariance between the two rates. We could then model vech(V_t) = $(v_{t,1}, v_{t,12}, v_{t,2})'$ directly but, for reasons of symmetry, we replace the realized covariance with the realized variance of the ¥/DM cross rate which may be done, without loss of generality, in the absence of triangular arbitrage, resulting in a system of three realized volatilities.

To appreciate the implication of precluding triangular arbitrage, note that this constraint requires the continuously compounded return on the ¥/DM cross rate to equal the difference between the ¥/$ and DM/$ returns, which has two key consequences. First, it implies that, even absent direct data on the ¥/DM cross rate, we can calculate it using our DM/$ and ¥/$ data. We can then calculate the realized cross-rate variance, $v_{t,3}$, by summing the implied thirty-minute squared cross-rate returns,

$$v_{t,3} \;=\; \Sigma_{j=1,..,1/\Delta} \left[(-1\,,1)' \cdot r_{t-1+j\cdot\Delta,\Delta} \right]^2 . \tag{13}$$

Second, because it implies that $v_{t,3} = v_{t,1} + v_{t,2} - 2\, v_{t,12}$, we can infer the realized covariance from the three realized volatilities,[15]

$$v_{t,12} \;=\; ½\cdot(v_{t,1} + v_{t,2} - v_{t,3}) . \tag{14}$$

Building on this insight, we infer the covariance from the three variances, $v_t \equiv (v_{t,1}, v_{t,2}, v_{t,3})'$, and the identity in equation (14), instead of directly modeling vech(V_t).

We now turn to a discussion of the pertinent empirical regularities that guide our specification of a trivariate forecasting model for the three DM/$, ¥/$, and ¥/DM volatility series.

4. PROPERTIES OF EXCHANGE RATE RETURNS AND REALIZED VOLATILITIES

The in-sample distributional features of the DM/$ and ¥/$ returns and the corresponding realized volatilities have been characterized previously by ABDL (2001a, 2001b).[16] Here we briefly summarize those parts of the ABDL results that are relevant for the present inquiry. We also provide new results for the ¥/DM cross rate volatility and an equally-weighted portfolio that explicitly incorporates the realized covariance measure discussed above.

4.1 *Returns*

---

[15] Similarly, the realized correlation between the two dollar rates is given by $\rho_{t,12} = ½\cdot(v_{t,1} + v_{t,2} - v_{t,3}) / (v_{t,1} \cdot v_{t,2})^{½}$.

[16] For a prescient early contribution along these lines, see also Zhou (1996).

The statistics in the top panel of Table 1 refer to the two daily dollar denominated returns, $r_{t,1}$ and $r_{t,2}$, and the equally-weighted portfolio $½\cdot(r_{t,1}+r_{t,2})$. As regards unconditional distributions, all three return series are approximately symmetric with zero mean. However, the sample kurtoses indicate fat tails relative to the normal, which is confirmed by the kernel density estimates shown in Figure 1. As regards conditional distributions, the Ljung-Box test statistics indicate no serial correlation in returns, but strong serial correlation in squared returns. These results are entirely consistent with the extensive literature documenting fat tails and volatility clustering in asset returns, dating at least to Mandelbrot (1963) and Fama (1965).

The statistics in the bottom panel of Table 1 refer to the distribution of the standardized daily returns $r_{t,1}\cdot v_{t,1}^{-1/2}$ and $r_{t,2}\cdot v_{t,2}^{-1/2}$, along with the standardized daily equally-weighted portfolio returns $½\cdot(r_{t,1}+r_{t,2})\cdot(¼\cdot v_{t,1}+¼\cdot v_{t,2}+½\cdot v_{t,12})^{-½}$, or equivalently by equation (14), $½\cdot(r_{t,1}+r_{t,2})\cdot(½\cdot v_{t,1}+½\cdot v_{t,2}-¼\cdot v_{t,3})^{-½}$. The results are striking. Although the kurtosis for all of the three standardized returns are less than the normal value of three, the returns are obviously close to being Gaussian. Also, in contrast to the raw returns in the top panel, the standardized returns display no evidence of volatility clustering.[17] This impression is reinforced by the kernel density estimates in Figure 1, which visually convey the approximate normality.

Of course, the realized volatility used for standardizing the returns is only observable ex post. Nonetheless, the result is in stark contrast to the typical finding that, when standardizing daily returns by the one-day-ahead forecasted variance from ARCH or stochastic volatility models, the resulting distributions are invariably leptokurtic, albeit less so than for the raw returns; see, e.g., Baillie and Bollerslev (1989) and Hsieh (1989). In turn, this has motivated the widespread adoption of volatility models with non-Gaussian conditional densities, as suggested by Bollerslev (1987).[18] The normality of the standardized returns in Table 1 and Figure 1 suggests a different approach: a fat-tailed Gaussian mixture distribution governed by the realized volatilities. Of course, this is also consistent with the results of Theorem 2. We now turn to a discussion of the distribution of the realized volatilities.

4.2 *Realized Volatilities*

The statistics in the top panel of Table 2 summarize the distribution of the realized volatilities, $v_{t,i}^{1/2}$, for

---

[17] Similar results obtain for the multivariate standardization $V_t^{-1/2} r_t$, where $[\cdot]^{-1/2}$ refers to the Cholesky factor of the inverse matrix, as documented in ABDL (2001b).

[18] This same observation also underlies the ad hoc multiplication factors often employed by practioners in the construction of VaR forecasts.

each of the three exchange rates: DM/$, ¥/$ and ¥/DM. All the volatilities are severely skewed to the right and highly leptokurtic. In contrast, the skewness and kurtosis for the three logarithmic standard deviations, $y_{t,i} \equiv ½ \cdot log(v_{t,i})$, shown in the bottom panel of Table 2, appear remarkably Gaussian. Figure 2 confirms these impressions. In fact, the kernel density estimates for the logarithmic volatilities nearly coincide with the normal reference densities, rendering the curves effectively indistinguishable.

The log-normality of realized volatility suggests the use of standard Gaussian distribution theory and associated critical values for modeling and forecasting realized logarithmic volatilities. Moreover, combining the results for the returns in Table 1, suggesting $r_t \cdot V_t^{-1/2} \sim N(0, I)$, with the results for the volatility in Table 2, $y_t \equiv (y_{t,1}, y_{t,2}, y_{t,3})' \sim N(\mu, \Omega)$, we should expect the overall return distribution (not conditioned on the realized volatility) to be well approximated by a lognormal-normal mixture.[19] We present density forecasts and VaR calculations below that explicitly build on this insight.

Turning again to Table 2, the Ljung-Box statistics indicate strong serial correlation in the realized daily volatilities, in accord with the significant Ljung-Box statistics for the squared unstandardized returns in the top panel of Table 1. It is noteworthy, however, that the $Q^2(20)$ statistics in Table 1 are orders of magnitude smaller than those in Table 2. This reflects the fact that the daily squared returns constitute very noisy volatility proxies relative to the daily realized volatilities.[20] Consequently, the strong persistence of the underlying volatility dynamics is masked by the measurement error in the daily squared returns.

Several recent studies have suggested that the strong serial dependence in financial asset return volatility may be conveniently captured by a long-memory, or fractionally-integrated, process (e.g., Ding, Granger and Engle, 1993, and Andersen and Bollerslev, 1997). Hence in the last column of Table 2 we report estimates of the degree of fractional integration, obtained using the Geweke and Porter-Hudak (1983) (GPH) log-periodogram regression estimator as formally developed by Robinson (1995). The three estimates of d are all significantly greater than zero and less than one half when judged by the

---

[19] The density function for the lognormal-normal mixture is formally given by

$$f(r) = (2 \cdot \pi \cdot \sigma^2)^{-1} \int_0^\infty y^{-3/2} \cdot exp\{-½ \cdot [\, r^2 y^{-1} + \sigma^{-2}(log(y) - \mu)^2 \,]\}\, dy \,,$$

where μ and σ denote the mean and the variance of the logarithmic volatility. This same mixture distribution has previously been advocated by Clark (1973), without any of the direct empirical justification provided here.

[20] See Andersen and Bollerslev (1998) for a detailed efficiency comparison of various volatility proxies.

standard error of 0.028 in the asymptotic normal distribution.[21] Moreover, the three estimates are very close, indicative of a common degree of long-run dependence in the logarithmic volatilities. The multivariate extension of the GPH estimator developed by Robinson (1995) provides a formal framework for testing this hypothesis. On implementing Robinson's estimator we obtain $\hat{d} = 0.401$, and the corresponding test statistic for a common d across the three volatilities has a p-value of 0.510 in the asymptotic chi-square distribution with three degrees of freedom.

Figure 3 provides graphical confirmation and elaboration of the long-memory results. It displays the sample autocorrelations of the realized logarithmic volatilities out to a displacement of 70 days, or about one quarter. The slow hyperbolic autocorrelation decay symptomatic of long memory is evident, and the qualitatively identical autocorrelation values across the three volatilities supports the assertion of a common degree of fractional integration. Figure 3 also shows the sample autocorrelations of the logarithmic volatilities, fractionally differenced by applying the filter $(1-L)^{0.401}$. It is evident that this single fractional differencing operator eliminates the bulk of the serial dependence in each of the three realized logarithmic volatilities, although Ljung-Box portmanteau tests (not reported here) do reject the hypothesis of white noise fractionally-differenced volatilities.

It is possible that the three series are fractionally cointegrated, so that a linear combination will exhibit a degree of fractional integration less than 0.401. On heuristically testing for this by regressing each of the logarithmic volatilities on the two other logarithmic volatilities and a constant, and then estimating the degree of fractional integration in the residuals, the three estimates for d are 0.356, 0.424, and 0.393, respectively, all of which are very close to the value of d for the original series in Table 2.[22]

Although the realized logarithmic volatility series do not appear to be fractionally cointegrated, they are nevertheless strongly contemporaneously correlated. In particular, the sample correlations between $y_{t,1}$ and $y_{t,2}$ and $y_{t,3}$ are respectively 0.591 and 0.665, while the correlation between $y_{t,2}$ and $y_{t,3}$ is 0.648. Again, this is entirely consistent with the extant ARCH and stochastic volatility literatures. In the next section, we propose a simple multivariate model capable of accommodating both the strong dynamics and contemporaneous correlations in the realized volatilities.

5. A VAR FOR MODELING AND FORECASTING REALIZED VOLATILITY

[21] Consistent with the simulation evidence in Bollerslev and Wright (2000), the corresponding estimates of d for v_t, not reported here, are about 0.15 less than those for y_t reported in the bottom panel of Table 2.

[22] More complicated formal testing procedures for fractional cointegration have recently been developed by Robinson and Marinucci (1999) and Davidson (2000).

The distributional features highlighted in the previous section suggest that a long-memory Gaussian VAR for the realized logarithmic volatilities should provide a good description of the volatility dynamics. We therefore consider the simple tri-variate VAR,

$$A(L)(1-L)^d(y_t - \mu) = \epsilon_t \,, \tag{15}$$

where ϵ_t is a vector white noise process.

The impulse response functions for the model (15) are easily calculated and lend insight into its implied dynamics. On rearranging terms and multiplying by (1-L), we obtain,

$$(1-L)\,(y_t - \mu) = A(L)^{-1}\,(1-L)^{1-d}\,\epsilon_t \,. \tag{16}$$

Hence the impulse response coefficients associated with the lag k shocks are simply given by the powers in the matrix lag polynomial $\Gamma(L) \equiv A(L)^{-1}\,(1-L)^{1-d}$, say Γ_k. If all of the roots of the estimated polynomial $|A(z)|=0$ are outside the unit circle, the model is stationary and $\Gamma_k \to 0$. Moreover, all of the cumulative impulse response coefficients associated with the volatility shocks, $\Gamma_1+\Gamma_2+...+\Gamma_k$, eventually dissipate at the slow hyperbolic rate k^{d-1}.

We now estimate the unrestricted VAR in equation (15) by applying OLS equation-by-equation. In so doing, we impose the normalization $A(0) = I$, and fix the value of d at the earlier-reported common estimate of 0.401. We also assume that the orders of the lag polynomials in $A(L)$ are all equal to five days, or one week. This choice is somewhat arbitrary, and the model could easily be refined through a more detailed specification search explicitly allowing for zero parameter restrictions and/or different autoregressive lag lengths.[23] Additional explanatory variables, such as interest rate differentials, daily trading activity measures, etc., could also easily be included. However, in order to facilitate the comparison with the daily volatility models in common use, for which the mechanics of including additional explanatory variables are much more complicated and typically not entertained, we restrict our attention to the simple unrestricted VAR in equation (15).

Many of the estimated VAR coefficients (not shown) are statistically significant, and all the roots of the estimate of the matrix lag polynomial A(L) are outside the unit circle, consistent with covariance stationarity. Moreover, Ljung-Box tests for serial correlation in the VAR residuals reveal no evidence against the white noise hypothesis, indicating that the VAR has successfully accommodated all volatility dynamics not already captured by the first-stage long memory modeling.[24]

---

[23] Both the Akaike and Schwarz criteria select a first-order VAR. Degrees of freedom are plentiful, however, so we included a week's worth of lags to maintain conservatism.

[24] Details concerning model estimates and related test statistics are available upon request.

It is striking to note how much volatility variation is explained by the univariate long-memory models (the R^2 values are in the neighborhood of 50%), and how little of the variation of the residuals from the long-memory models is explained by the VAR (the R^2 values are in the neighborhood of 2%). Effectively, the univariate one-parameter long-memory models are so successful at explaining realized volatility dynamics that only little is left for the VAR. This is also evident from the previously discussed plots of the autocorrelations in Figure 3. Nevertheless, the Ljung-Box statistics for the three univariate fractionally differenced volatility series all indicate significant serial correlation, while those for the residuals from the VAR do not. Moreover, the VAR captures important cross-rate linkages. In particular, Granger causality tests reveal strong evidence of predictive enhancement from including lags of the logarithmic DM/$ and ¥/$ volatility in the realized logarithmic ¥/DM volatility equation.

It is natural to conjecture that vector autoregressive forecasts of realized volatility based on equation (15) will outperform more traditional daily ARCH and related volatility forecasts. Our forecasts are based on explicitly-modeled long-memory dynamics, which seem to be a crucial feature of the data. Long-memory may, of course, also be incorporated in otherwise standard ARCH models, as in Baillie, Bollerslev and Mikkelsen (1996). As such, the genuinely distinctive feature of our approach is instead the effective incorporation of information contained in the high-frequency data, which enables the realized volatilities and their forecasts to quickly adapt to changes in the underlying latent volatility. The next section explores these conjectures in detail.

6. EVALUATING AND COMPARING ALTERNATIVE VOLATILITY FORECASTS

Volatility forecasts play a central role in the financial decision making process. In this section we assess the performance of the realized volatility forecasts generated from our simple vector autoregressive model. In Figure 4, to convey some initial feel, we plot the daily realized DM/$, ¥/$, and ¥/DM standard deviations, along with the corresponding one-day-ahead VAR forecasts for the out-of-sample period, December 2, 1996, through June 30, 1999. It appears that the VAR does a good job of capturing both the low-frequency and the high-frequency movements in the realized volatilities. We next proceed to a more thorough statistical evaluation of the forecasts along with a comparison to several alternative volatility forecasting procedures currently in widespread use.

6.1 *Forecast Evaluation Regressions*

Many methods have been proposed for modeling and forecasting financial market volatility. The most widespread procedure used by practitioners is arguably J.P. Morgan's (1997) RiskMetrics. The

RiskMetrics daily variances and covariances are calculated as exponentially weighted averages of cross products of daily returns, using a smoothing factor of λ*=0.94*. This corresponds to an IGARCH(1,1) model in which the intercept is fixed at zero and the moving average coefficient in the ARIMA representation for the squared returns equals *-0.94*.

The most widespread procedure used by academics is arguably the GARCH model of Engle (1982) and Bollerslev (1986), with the GARCH(1,1) model constituting the leading case. As with the VAR model discussed in the previous section, we will base the GARCH(1,1) model estimates on the *2,449* daily in-sample returns from December 1, 1996, through December 1, 1996. Consistent with previous results reported in the literature, the quasi-maximum likelihood parameter estimates indicate a strong degree of volatility persistence, with the autoregressive roots for each of the three rates equal to *0.986*, *0.968*, and *0.990*, respectively.[25]

Finally, we will also compare the forecasts from our fifth-order VAR for the long-memory filtered realized volatility (henceforth VAR-RV) to those obtained from a fifth-order VAR for the long-memory filtered daily logarithmic absolute returns (henceforth VAR-ABS). This makes for a controlled comparison, as the model structures are identical in all respects except the volatility proxy: one uses daily realized volatility, while the other uses daily absolute returns.

No universally acceptable loss function exists for the ex-post evaluation and comparison of nonlinear model forecasts, and in the context of volatility modeling, several statistical procedures have been used for assessing the quality of competing forecasts (see, e.g. the discussion in Andersen, Bollerslev, and Lange, 1999, and Christoffersen and Diebold, 1999). Following the analysis in Andersen and Bollerslev (1998), and in the tradition of Mincer-Zarnowitz (1969) and Chong and Hendry (1986), we focus our evaluations on regressions of the realized volatilities on a constant and the various model forecasts.

For the one-day-ahead in-sample and out-of-sample forecasts reported in Tables 3A and 3B, the evaluation regressions take the form[26]

---

[25] Further details regarding the actual estimates are available on request. To explicitly accommodate long-memory in the volatilities, we also estimated a set of FIGARCH(1,d,0) models (see, Baillie, Bollerslev, and Mikkelsen, 1996). However, the one- and ten-day-ahead predictions for these models were generally very close to the results for the GARCH(1,1) models discussed below, and are therefore omitted to conserve space. Again, further details are available upon request.

[26] Following standard practice, we focus on forecasts of the volatility, $v_t^{1/2}$, based on the square root of the unbiased GARCH(1,1) and VAR forecasts for v_t. Of course, these forecasts are not formally unbiased, so we also experimented with a first order Taylor series expansion of the square root operator to adjust for this bias. The resulting regression estimates and R^2's were almost identical to the ones reported in Table 3. Similarly, the evaluation regressions for the

$$(\{v_{t+1}\}_i)^{½} = b_0 + b_1 \cdot (\{v_{t+1/t,VAR\text{-}RV}\}_i)^{½} + b_2 \cdot (\{v_{t+1/t,\,Model}\}_i)^{½} + u_{t+1,i}. \quad (17)$$

The results are striking. For the regressions including just one volatility forecast, the regression R^2 is always the highest for the VAR-RV model, and for none of the VAR-RV forecasts can we reject the hypothesis that $b_0 = 0$ and $b_1 = 1$ using the corresponding t tests. In contrast, we reject the hypothesis that $b_0 = 0$ and/or $b_2 = 1$ for most RiskMetrics, GARCH and VAR-ABS in-sample and out-of-sample forecasts. Moreover, on including both the VAR-RV and the RiskMetrics, GARCH or VAR-ABS forecast in the same regression, the coefficient estimates for b_1 and b_2 are generally not different from unity and zero, respectively. Finally, inclusion of the RiskMetrics, GARCH or VAR-ABS forecasts improves the R^2s very little relative to those based solely on the VAR-RV forecasts.[27]

Turning to the ten-day-ahead forecasts, the results are qualitatively identical. The evaluation regressions are

$$(\{\textstyle\sum_{j=1,..,10} v_{t+j}\}_i)^{½} = b_0 + b_1 \cdot (\{\textstyle\sum_{j=1,..,10} v_{t+j/t,VAR\text{-}RV}\}_i)^{½}$$
$$+ b_2 \cdot (\{\textstyle\sum_{j=1,..,10} v_{t+j/t,\,Model}\}_i)^{½} + u_{t+10,i}\,. \quad (18)$$

Both the in-sample and out-of-sample results in Tables 3C and 3D, respectively, systematically favor the VAR-RV forecasts for all three rates.[28] As with the one-day-ahead regressions discussed above, the estimates for b_1 are generally not significantly different from unity, while very few of the estimates for b_0 and b_2 in the multiple regressions including both the VAR-RV forecasts and the other forecasts are statistically significantly different from zero. These results are especially noteworthy insofar as several previous studies have found it difficult to outperform simple daily GARCH(1,1) formulations for exchange rate volatility using more complicated multivariate models (e.g., Sheedy, 1998), or ARCH models estimated directly from high-frequency data (e.g., Beltratti and Morana, 1999).

6.2 *On the Superiority of VAR-RV Forecasts*

Why does the VAR-RV produce superior forecasts? We have identified the quadratic variation and its

---

realized variances, v_t, and logarithmic standard deviations, $y_t \equiv ½ \cdot log(v_t)$, produced very similar results to the ones reported here. Details are available upon request.

[27] These results are consistent with Engle (2000), who reports that the inclusion of the lagged daily realized variance in the conditional variance equation of a GARCH(1,1) model for the daily DM/$ returns analyzed here renders the coefficient associated with the lagged daily squared returns insignificant. In contrast, the results of Taylor and Xu (1997), based on a limited one-year sample of *253* daily DM/$ returns, suggest that the lagged daily realized variance offers little incremental explanatory power over a univariate GARCH(1,1) model.

[28] To account for the overlap in the multi-day forecasts, we use robust standard errors, calculated using an unweighted covariance matrix estimator allowing for up to ninth order serial correlation in $u_{t+10,i}$.

empirical counterpart, the *realized volatility*, as the key objects of interest for volatility measurement, and consequently assess our various volatility forecast relative to this measure. It is therefore maybe not surprising that models built directly for the realized volatility produce forecasts superior to those obtained from more indirect methods, a conjecture consistent with the literature on forecasting under the relevant loss function, e.g., Christoffersen and Diebold (1997).

There is a more direct reason for the superior performance of the VAR-RV forecasts, however. The essence of forecasting is quantification of the mapping from the past and, in particular, the present into the future. Hence, quite generally, superior estimates of present conditions translate into superior forecasts of the future. Realized volatility excels in this dimension: it provides a superior and quickly-adapting estimate of current volatility. Standard models based on daily data such as GARCH and RiskMetrics rely on long and slowly decaying weighted moving averages of past squared returns and therefore adapt only gradually to volatility movements.[29]

Figure 5 confirms the above assertions graphically. We display the realized standard deviations for DM/$ returns, ¥/$ returns and ¥/DM returns, along with the corresponding one-day-ahead GARCH forecasts for the out-of-sample period, December 2, 1996, through June 30, 1999. The GARCH forecasts appear to track the low-frequency variation adequately, matching the broad temporal movements in volatility, but they track much less well at higher frequencies. Note the striking contrast with Figure 4 which, as discussed earlier, reveals a close coherence between daily realized volatility and VAR-RV forecasts at high as well as low frequencies.

We provide a more detailed illustration of the comparative superiority of the VAR-RV forecasts in Figure 6, which depicts four illustrative DM/$ episodes of thirty-five days each.[30] First, for days one through twenty-five (the non-shaded region) we show daily realized volatility together with one-day-ahead forecasts made on the previous day using VAR-RV and GARCH models. The accuracy of the VAR-RV forecasts is striking, as is the *in*accuracy of the GARCH forecasts, due to their inability to adapt rapidly to high-frequency movements. Second, for days twenty-six through thirty-five (the shaded

---

[29] Suppose, for example, that integrated volatility has been low for many days, $t = 1, ..., T-1$, so that both realized and GARCH volatilities are presently low as well. Now suppose that integrated volatility increases sharply on day T and that the effect is highly persistent as is typical. Realized volatility for day T, which makes extensive use of the day-T information, will increase sharply as well, as is appropriate. GARCH or RiskMetrics volatility, in contrast, won't move at all on day T, as they depend only on squared returns from days T-1, T-2, ..., and they will increase only gradually on subsequent days, as they approximate volatility via a long and slowly decaying exponentially weighted moving average.

[30] The actual start dates for the forecasts in each of the four panels are 11/16/1988, 3/6/1989, 5/13/1991, and 8/28/1998, respectively.

region), we continue to display the daily realized volatility, but we show one- through ten-day-ahead VAR-RV and GARCH forecasts based on information available on day twenty-five. Hence the forecasts for day twenty-six are one-day-ahead, the forecasts for day twenty-seven are two-day-ahead, and so on. Examination of these multi-step trajectories makes clear the forecasting benefit of having a superior estimate of current volatility: in each case the VAR-RV forecasts "take off" from a good estimate of current (day twenty-five) volatility, and then revert to the unconditional mean of approximately 0.6 in a fashion consistent with historical experience. In contrast, the lack of an equally accurate assessment of current volatility within the GARCH setting results in markedly inferior longer-run volatility forecasts.

The above findings do not reflect a failure of the GARCH model per se, but rather the efficacy of exploiting volatility measures based on intraday return observations. In fact, Andersen and Bollerslev (1998) find that the GARCH model explains about as much of the future variation in volatility as is theoretically feasible, and thus exploits the information in daily foreign exchange data very well. The point is that the noise in daily squared returns necessarily renders the measurement of the current innovation to volatility imprecise, independent of the correct model for daily volatility.

In light of the excellent performance of the VAR-RV volatility forecasts, we now turn to an investigation of the corresponding return density forecasts and VaR calculations, in which we explicitly incorporate the theoretically motivated distributional features of the realized volatilities and standardized returns highlighted in section 4.

7. DENSITY FORECASTS AND QUANTILE CALCULATIONS: A VAR FOR VaR

Measuring and forecasting portfolio Value-at-Risk, or VaR, and fluctuations in VaR due to changing market conditions and/or portfolio shares, is an important part of modern financial risk management, see, e.g., Gouriéroux, Laurent and Scaillet (2000). Our results suggest that accurate return density forecasts and associated VaR estimates may be obtained from a long-memory VAR for realized volatility, coupled with the assumption of normally distributed standardized returns.[31] We assess this conjecture using the methods of Diebold, Guther and Tay (1998). The basic idea is as follows. Suppose that the daily returns $\{r_t\}$ is generated from the series of one-day-ahead conditional densities $\{f(r_t|\mathscr{F}_{t-1})\}$, where $\mathscr{F}_{t-1}$ denotes the full information set available at time t-1. If the series of one-day-ahead conditional density

[31] Portfolio VaR at confidence level ρ percent and horizon k is simply the ρth percentile of the k-step-ahead portfolio return density forecast. For an overview, see Duffie and Pan (1997). When calculating VaR at confidence level ρ and horizon k, the appropriate values of ρ and k are generally situation-specific, although the Basel Committee on Banking Supervision has advocated the systematic use of five- and one-percent VaR, and one- and ten-day horizons.

forecasts $\{f_{t|t-1}(r_t)\}$ coincides with $\{f(r_t|\mathscr{F}_{t-1})\}$, it follows under weak conditions, that the sequence of probability integral transforms of r_t with respect to $f_{t|t-1}(\cdot)$ should be *iid* uniformly distributed on the unit interval. That is, $\{z_t\} \overset{iid}{\sim} U(0,1)$, where the probability integral transform z_t is simply the cumulative density function corresponding to $f_{t|t-1}(\cdot)$, evaluated at r_t, i.e., $z_t = \int_{-\infty}^{r_t} f_{t|t-1}(u)\,du$. Hence the adequacy of the VAR-RV based volatility forecast and the lognormal-normal mixture distribution may be assessed by checking whether the corresponding distribution of $\{z_t\}$ is *iid* $U(0,1)$.

To this end, Table 4 reports the percentage of the realized DM/\$, ¥/\$ and equally-weighted portfolio returns that are less than various quantiles forecast by the long-memory lognormal-normal mixture model. The close correspondence between the percentages in each column and the implied quantiles by the model, both in-sample and out-of-sample, is striking. It is evident that the VAR-RV lognormal-normal mixture model affords a very close fit for all of the relevant VaRs.[32]

Although uniformity of the z_t sequences is necessary for adequacy of the density forecasts, it is not sufficient. The z_t's must be independent as well, to guarantee, for example, that violation of a particular quantile forecast on day *t* conveys no information regarding its likely violation on day *t+1*. In general, dependence in z_t would indicate that dynamics have been inadequately modeled and captured by the forecasts. To assess independence, Figure 7 therefore plots the sample autocorrelation functions for $(z_t - \bar{z})$ and $(z_t - \bar{z})^2$ corresponding to the density forecasts for the DM/\$, the ¥/\$, and the equally-weighted portfolio returns. All told, there is no evidence of serial correlation in any of the six series, indicating that the model's density forecasts are not only correctly unconditionally calibrated, but also correctly conditionally calibrated.[33]

8. CONCLUSIONS AND DIRECTIONS FOR FUTURE RESEARCH

Guided by a general theory for continuous-time arbitrage-free price processes, we develop a framework for the direct modeling and forecasting of *realized* volatility and correlation. Our methods are simple to implement empirically, even in multivariate situations. We illustrate the idea in the context of the foreign exchange market, specifying and estimating a long-memory Gaussian VAR for a set of daily logarithmic realized volatilities. The model produces strikingly successful volatility forecasts, which dominate those

[32] Standard qq-plots, available on request, reveal an equally close fit throughout the full support of the distributions.

[33] This apparent lack of dependence is confirmed by formal Ljung-Box portmanteau tests for the joint significance of the depicted autocorrelations. One could, of course, examine the sample autocorrelations of higher powers of $(z_t - \bar{z})$ as well, but the second moment is the one crucial for our application, which centers on volatility dynamics.

from conventional GARCH and related approaches. It also generates well-calibrated density forecasts and associated quantile, or VaR, estimates. Numerous interesting directions for future research remain.

First, the realized volatility measures used in this paper do not distinguish between variability originating from continuous price movements or jumps. However, as discussed in Section 2, the dynamic impact may differ across the two sources of variability. Hence, it is possible that improved volatility forecasting models may be constructed by explicitly modeling the jump component, if present.

Second, although the lognormal-normal mixture distribution works very well in the present context, the predictive distribution could be refined and adapted to more challenging contexts using the numerical simulation methods of Geweke (1989), the Cornish-Fisher expansion of Baillie and Bollerslev (1992), or the recalibration methods of Diebold, Hahn and Tay (1999).

Third, although we focused on using density forecasts to obtain VaR estimates (quantiles), the same density forecasts could of course be used to calculate other objects of interest in financial risk management. Examples include the probability of loss exceeding a specified threshold (shortfall probabilities), and the expected loss conditional upon loss exceeding a pre-specified threshold (expected shortfall), as discussed for example in Heath, Delbaen, Eber and Artzner (1999) and Basak and Shapiro (1999). In contrast, the quantile regression methods of Granger, White and Kamstra (1989) and Engle and Manganelli (1999) do not readily lend themselves to such extensions.

Fourth, our approach to exchange rate density forecasting could be extended to other classes of financial assets. Although the structure of our proposed modeling framework builds directly on the empirical regularities for the foreign exchange markets documented in Section 4, the empirical features characterizing other asset markets appear remarkably similar, as shown for example by Andersen, Bollerslev, Diebold and Ebens (2001) for U.S. equities.

Fifth, and perhaps most importantly, volatility forecasts figure prominently in many practical financial decisions, such as asset allocation (e.g., Fleming, Kirby and Ostdiek, 2001), market timing (e.g., Busse, 1999), and derivative pricing (e.g., Bollerslev and Mikkelsen, 1999). It will be of interest to explore the gains afforded by the simple volatility modeling and forecasting procedures developed here, particularly in high-dimensional environments. In this regard, a couple of issues merit particular attention. One critical task is to develop realized volatility forecasting models that are parameterized in ways that guarantee positive definiteness of forecasted covariance matrices within high-dimensional settings. Because the in-sample realized covariance matrix is positive definite under quite general conditions, one approach would be to model the Cholesky factors rather than the realized covariance matrix itself. The corresponding forecasts for the Cholesky factors are then readily transformed into forecasts for the future variances and covariances by simple matrix multiplications.

Finally, and also of particular relevance in high-dimensional situations, allowing for factor structure in the modeling and forecasting of realized volatility may prove useful, as factor structure is central to both empirical and theoretical financial economics. Previous research on factor volatility models has typically relied on complex procedures involving a latent volatility factor, as for example in Diebold and Nerlove (1989), Engle, Ng and Rothschild (1990), and King, Sentana and Wadhwani (1994). In contrast, factor analysis of realized volatility should be relatively straightforward, even in high-dimensional environments. Moreover, the identification of explicit volatility factors, and associated market-wide variables that underlie the systematic volatility movements, may help to provide an important step towards a better understanding of "the economics of volatility."

REFERENCES

Andersen, T.G., L. Benzoni, and J. Lund (2000), "Estimating Jump-Diffusions for Equity Returns," Manuscript, Department of Finance, J.L. Kellogg Graduate School of Management, Northwestern University.

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

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

Andersen, T.G., T. Bollerslev, F.X. Diebold and H. Ebens (2001), "The Distribution of Realized Stock Return Volatility," *Journal of Financial Economics*, in press.

Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2000), "Market Microstructure Effects and the Estimation of Integrated Volatility," Work in Progress, Northwestern University, Duke University, and University of Pennsylvania.

Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2001a), "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association*, in press.

Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2001b), "Exchange Rate Returns Standardized by Realized Volatility are (Nearly) Gaussian," *Multinational Finance Journal*, in press.

Andersen, T.G., T. Bollerslev and S. Lange (1999), "Forecasting Financial Market Volatility: Sample Frequency vis-a-vis Forecast Horizon," *Journal of Empirical Finance*, 6, 457-477.

Back, K. (1991), "Asset Prices for General Processes," *Journal of Mathematical Economics*, 20, 317-395.

Baillie, R.T. and T. Bollerslev (1989), "The Message in Daily Exchange Rates: A Conditional Variance Tale," *Journal of Business and Economic Statistics*, 7, 297-305.

Baillie, R.T. and T. Bollerslev (1992), "Prediction in Dynamic Models with Time-Dependent Conditional Variances," *Journal of Econometrics*, 52, 91-113.

Baillie, R.T., T. Bollerslev and H.O. Mikkelsen (1996), "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 74, 3-30.

Bakshi, G, C. Cao and Z. Chen (1997), “Empirical Performance of Alternative Option Pricing Models,” *Journal of Finance*, 52, 2003-2049.

Barndorff-Nielsen, O.E. and N. Shephard (2000), "Econometric Analysis of Realised Volatility and its Use in Estimating Levy Based Non-Gaussian OU Type Stochastic Volatility Models," Manuscript, Nuffield College, Oxford, United Kingdom.

Barndorff-Nielsen, O.E. and N. Shephard (2001), "Non-Gaussian Ornstein-Uhlenbeck-Based Models and Some of Their Uses in Financial Economics," *Journal of the Royal Statistical Society, Series B*, in press.

Basak, S. and A. Shapiro (1999), "Value-at-Risk Based Risk Management: Optimal Policies and Asset Prices," Manuscript, University of Pennsylvania.

Bates, D. (2000), "Post-'87 Crash Fears in S&P 500 Futures Options," *Journal of Econometrics*, 94, 181-238.

Beltratti, A. and C. Morana (1999), "Computing Value-at-Risk with High-Frequency Data," *Journal of Empirical Finance*, 6, 431-455.

Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307-327.

Bollerslev, T. (1987), "A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *Review of Economics and Statistics*, 69, 542-547.

Bollerslev, T., R.F. Engle and D.B. Nelson (1994), "ARCH Models," in R.F. Engle and D. McFadden (eds.), *Handbook of Econometrics, Volume IV*, 2959-3038. Amsterdam: North-Holland.

Bollerslev, T. and H.O. Mikkelsen (1999), "Long-Term Equity Anticipation Securities and Stock Market Volatility Dynamics," *Journal of Econometrics*, 92, 75-99.

Bollerslev, T. and J.H. Wright (2000), "Semiparametric Estimation of Long-Memory Volatility Dependencies: The Role of High Frequency Data," *Journal of Econometrics*, 98, 81-106.

Busse, J.A. (1999), "Volatility Timing in Mutual Funds: Evidence from Daily Returns," *Review of Financial Studies*, 12, 1009-1041.

Chong, Y.Y. and D.F. Hendry (1986), "Econometric Evaluation of Linear Macroeconomic Models," *Review of Economic Studies*, 53, 671-690.

Christoffersen, P.F. and F.X. Diebold (1997), "Optimal Prediction Under Asymmetric Loss," *Econometric Theory*, 13, 808-817.

Christoffersen, P.F. and F.X. Diebold (1999), "How Relevant is Volatility Forecasting for Financial Risk Management," *Review of Economics and Statistics*, 82, 12-22.

Clark, P.K. (1973), "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices," *Econometrica*, 41, 135-155.

Dacorogna, M.M., U.A. Müller, R.J. Nagler, R.B. Olsen and O.V. Pictet (1993), "A Geographical Model for the Daily and Weekly Seasonal Volatility in the Foreign Exchange Market," *Journal of International Money and Finance*, 12, 413-438.

Davidson, J. (2000), "A Model of Fractional Cointegration, and Tests for Fractional Cointegration Using the Bootstrap," Manuscript, Cardiff University, Wales.

Diebold, F.X. and M. Nerlove (1989), "The Dynamics of Exchange Rate Volatility: A Multivariate Latent Factor ARCH Model," *Journal of Applied Econometrics*, 4, 1-22.

Diebold, F.X., T. Gunther and A.S. Tay (1998), "Evaluating Density Forecasts, with Applications to Financial Risk Management," *International Economic Review*, 39, 863-883.

Diebold, F.X., J. Hahn and A.S. Tay (1999), "Multivariate Density Forecast Evaluation and Calibration in Financial Risk Management: High-Frequency Returns on Foreign Exchange," *Review of Economics and Statistics*, 81, 661-673.

Ding, Z., C.W.J. Granger and R.F. Engle (1993), "A Long Memory Property of Stock Market Returns and a New Model," *Journal of Empirical Finance*, 1, 83-106.

Duffie, D. and J. Pan (1997), "An Overview of Value-at-Risk," *Journal of Derivatives*, 4, 9-49.

Engle, R.F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation," *Econometrica*, 50, 987-1008.

Engle, R.F. (2000), "Discussion of 'The Distribution of Exchange Rate Volatility,' by Andersen, Bollerslev, Diebold and Labys," Western Finance Association Meetings, Sun Valley, Idaho.

Engle, R.F. and Kroner, K. (1995), "Multivariate Simultaneous Generalized ARCH," *Econometric Theory*, 11, 122-150.

Engle, R.F., V. Ng and M. Rothschild (1990), "Asset Pricing with a Factor ARCH Covariance Structure: Empirical Estimates for Treasury Bills," *Journal of Econometrics*, 45, 213-238.

Engle, R.F. and S. Manganelli (1999), "CAViaR: Conditional Autoregressive Value-at-Risk by Regression Quantiles," Manuscript, University of California, San Diego.

Eraker, B., M. Johannes and N. Polson (2000), "Evidence for and the Impact of Jumps in Volatility and Returns," Manuscript, Graduate School of Business, University of Chicago.

Fama, E.F. (1965), "The Behavior of Stock Market Prices," *Journal of Business*, 38, 34-105.

Fleming, J., C. Kirby, and B. Ostdiek (2001), "The Economic Value of Volatility Timing," *Journal of Finance*, in press.

French, K.R., G.W. Schwert and R.F. Stambaugh (1987), "Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3-29.

Geweke, J., (1989), "Exact Predictive Densities for Linear Models with ARCH Disturbances," *Journal of Econometrics*, 40, 63-86.

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

Ghysels, E., A. Harvey and E. Renault (1996), "Stochastic Volatility," in G.S. Maddala (ed.), *Handbook of Statistics Vol. 14, Statistical Methods in Finance*, 119-191. Amsterdam: North-Holland.

Gouriéroux, C., J.P. Laurent and O. Scaillet (2000), "Sensitivity analysis of Values at Risk," *Journal of Empirical Finance*, forthcoming.

Granger, C.W.J., H. White and M. Kamstra (1989), "Interval Forecasting: An Analysis Based Upon ARCH-Quantile Estimators," *Journal of Econometrics*, 40, 87-96.

Heath, D., F. Delbaen, J.M. Eber and P. Artzner (1999), "Coherent Measures of Risk," *Mathematical Finance*, 9, 203-228.

Hsieh, D.A. (1989), "Modeling Heteroskedasticity in Daily Foreign Exchange Rates," *Journal of Business and Economic Statistics*, 7, 307-317.

Hsieh, D.A. (1991), "Chaos and Nonlinear Dynamics: Application to Financial Markets," *Journal of Finance*, 46, 1839-1877.

J.P. Morgan (1997), *RiskMetrics Technical Documents*, 4th edition, New York.

Karatzas, I. and S.E. Shreve (1991), *Brownian Motion and Stochastic Calculus*, (Second Edition). New York: Springer-Verlag.

King, M., E. Sentana and S. Wadhwani (1994), "Volatility and Links Between National Stock Markets," *Econometrica*, 62, 901-933.

Kroner, K.F. and V.K. Ng (1998), "Modeling Asymmetric Comovements of Asset Returns," *Review of Financial Studies*, 11, 817-844.

Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices," *Journal of Business*, 36, 394-419.

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

Merton, R.C., (1982), “On the Mathematics and Economics Assumptions of Continuous-Time Models,” in W.F. Sharpe and C.M. Cootner (eds.), *Financial Economics: Essays in Honor of Paul Cootner*. Englewood Cliffs, New Jersey.

Mincer, J. and V. Zarnowitz (1969), "The Evaluation of Economic Forecasts," in J. Mincer (ed.), *Economic Forecasts and Expectations*. New York: National Bureau of Economic Research.

Müller, U.A., M.M. Dacorogna, R.B. Olsen, O.V. Pictet, M. Schwarz and C. Morgenegg (1990), "Statistical Study of Foreign Exchange Rates, Empirical Evidence of a Price Change Scaling Law, and Intraday Analysis," *Journal of Banking and Finance*, 14, 1189-1208.

Nelson, D.B. and D. Foster (1995), “Filtering and Forecasting with Misspecified ARCH Models II: Making the Right Forecast with the Wrong Model,” *Journal of Econometrics*, 67, 303-335.

Pan, J. (1999), "Integrated Time-Series Analysis of Spot and Option Prices," Manuscript, Graduate School of Business, Stanford University.

Protter, P. (1992), *Stochastic Integration and Differential Equations: A New Approach* (Second Edition). New York: Springer Verlag.

Robinson, P.M. (1995), "Log-Periodogram Regression of Time Series with Long-Range Dependence," *Annals of Statistics*, 23, 1048-1072.

Robinson, P.M. and D. Marinucci (1999), "Semiparametric Frequency Domain Analysis of Fractional Cointegration," Manuscript, London School of Economics.

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

Sheedy, E., (1998), "Correlation in Currency Markets: A Risk Adjusted Perspective," *Journal of International Financial Markets, Institutions and Money*, 8, 59-82.

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

Zhou, B. (1996), "High-Frequency Data and Volatility in Foreign Exchange Rates," *Journal of Business and Economic Statistics*, 14, 45-52.

TABLE 1

Daily Return Distributions

	Mean	*Std.Dev.*	*Skewness*	*Kurtosis*	*Q(20)*	*$Q^2(20)$*
<u>Returns</u>						
DM/$	-0.007	0.700	0.003	5.28	14.13	140.3
¥/$	-0.010	0.692	-0.129	6.64	27.88	142.6
Portfolio	-0.008	0.630	-0.046	5.81	20.30	111.5
<u>Standardized Returns</u>						
DM/$	-0.007	0.993	0.032	2.57	19.42	23.32
¥/$	0.007	0.964	-0.053	2.66	32.13	24.83
Portfolio	-0.005	0.993	0.028	2.61	25.13	29.20

Notes: The daily returns cover the period from December 1, 1986 through December 1, 1996. The top panel refers to the distribution of raw daily returns, while the bottom panel refers to the distribution of daily returns standardized by realized volatility. The rows labeled "Portfolio" refer to returns on an equally-weighted portfolio. The columns labeled *Q(20)* and *$Q^2(20)$* contain Ljung-Box test statistics for up to twentieth order serial correlation in the returns and the squared returns.

TABLE 2

Daily Realized Volatility Distributions

	Mean	*St.Dev.*	*Skewness*	*Kurtosis*	*Q(20)*	*d*
<u>Volatility</u>						
DM/$	0.626	0.283	1.99	10.49	5249.2	-
¥/$	0.618	0.290	2.20	12.94	4155.1	-
¥/DM	0.571	0.234	1.49	7.86	10074.2-	
<u>Logarithmic Volatility</u>						
DM/$	-0.554	0.405	0.251	3.29	7659.6	0.387
¥/$	-0.572	0.419	0.191	3.44	5630.0	0.413
¥/DM	-0.637	0.388	0.058	3.04	12983.20.430	

Notes: The sample covers the period from December 1, 1986 until December 1, 1996. The top panel refers to the distribution of realized standard deviations, $v_t^{1/2}$. The daily realized volatilities are constructed from 30-minute squared returns, as detailed in the text. The bottom panel refers to the distribution of logarithmic realized standard deviations, $y_t \equiv \frac{1}{2} \cdot \log(v_t)$. The column labeled *Q(20)* contains Ljung-Box test statistics for up to twentieth order serial correlation. The last column gives the log-periodogram regression estimate of the fractional integration parameter, *d*, based on the $m = [T^{4/5}] = 514$ lowest-frequency periodogram ordinates. The asymptotic standard error for all the *d* estimates is $\pi \cdot (24 \cdot m)^{-1/2} = 0.028$.

TABLE 3A
Mincer-Zarnowitz Regressions for Realized Volatilities
In-Sample, One-Day-Ahead

	b_0	b_1	b_2	R^2
DM/$				
VAR-RV	-0.008 (0.020)	1.065 (0.035)	-	0.353
RiskMetrics	0.110 (0.019)	-	0.767 (0.030)	0.262
GARCH	-0.071 (0.026)	-	1.014 (0.040)	0.261
VAR-ABS	0.242 (0.019)	-	1.176 (0.061)	0.152
VAR-RV+RiskMetrics	-0.019 (0.020)	0.948 (0.059)	0.120 (0.046)	0.356
VAR-RV+GARCH	-0.047 (0.024)	0.950 (0.059)	0.157 (0.062)	0.356
VAR-RV+VAR-ABS	-0.018 (0.020)	1.030 (0.046)	0.093 (0.072)	0.354
¥/$				
VAR-RV	-0.010 (0.021)	1.069 (0.039)	-	0.374
RiskMetrics	0.142 (0.021)	-	0.725 (0.035)	0.263
GARCH	-0.102 (0.033)	-	1.064 (0.052)	0.295
VAR-ABS	0.258 (0.019)	-	1.186 (0.065)	0.153
VAR-RV+RiskMetrics	-0.021 (0.022)	0.947 (0.054)	0.126 (0.046)	0.377
VAR-RV+GARCH	-0.080 (0.032)	0.869 (0.060)	0.277 (0.079)	0.381
VAR-RV+VAR-ABS	-0.022 (0.021)	1.029 (0.047)	0.117 (0.066)	0.375
¥/DM				
VAR-RV	-0.022 (0.013)	1.079 (0.025)	-	0.505
RiskMetrics	0.162 (0.012)	-	0.735 (0.023)	0.368
GARCH	0.107 (0.015)	-	0.816 (0.027)	0.389
VAR-ABS	0.192 (0.014)	-	1.331 (0.051)	0.299
VAR-RV+RiskMetrics	-0.022 (0.013)	0.989 (0.046)	0.087 (0.037)	0.507
VAR-RV+GARCH	-0.029 (0.013)	0.941 (0.048)	0.144 (0.043)	0.509
VAR-RV+VAR-ABS	-0.028 (0.013)	1.012 (0.037)	0.146 (0.060)	0.507

Notes: We report OLS parameter estimates for Mincer-Zarnowitz regressions of realized volatility on a constant and forecasts from different models. The regression is $(\{v_{t+1}\}_i)^{½} = b_0 + b_1\cdot(\{v_{t+1/t,VAR\text{-}RV}\}_i)^{½} + b_2\cdot(\{v_{t+1/t,Model}\}_i)^{½} + u_{t+1,i}$, with heteroskedasticity robust standard errors in parentheses. The forecast evaluation period is December 1, 1987 through December 1, 1996, for a total of *2,223* daily observations. All model parameter estimates are based on data from December 1, 1986 through December 1, 1996. VAR-RV denotes forecasts from a long-memory vector autoregression for daily realized volatility, VAR-ABS denotes forecasts from a long-memory vector autoregression for daily log absolute returns, RiskMetrics denotes forecasts from an exponential smoothing model applied to squared daily returns, and GARCH denotes forecasts from a univariate GARCH(1,1) model. See the main text for details.

TABLE 3B
Mincer-Zarnowitz Regressions for Realized Volatilities
Out-of-Sample, One-Day-Ahead

	b_0	b_1	b_2	R^2
DM/$				
VAR-RV	0.021 (0.049)	0.987 (0.092)	-	0.249
RiskMetrics	0.219 (0.042)	-	0.618 (0.075)	0.097
GARCH	0.051 (0.063)	-	0.854 (0.105)	0.096
VAR-ABS	0.439 (0.028)	-	0.450 (0.089)	0.028
VAR-RV+RiskMetrics	0.017 (0.046)	0.979 (0.133)	0.014 (0.112)	0.249
VAR-RV+GARCH	0.015 (0.060)	0.980 (0.134)	0.016 (0.156)	0.249
VAR-RV+VAR-ABS	0.035 (0.046)	1.018 (0.107)	-0.106 (0.103)	0.250
¥/$				
VAR-RV	-0.006 (0.110)	1.085 (0.151)	-	0.329
RiskMetrics	0.164 (0.108)	-	0.767 (0.131)	0.266
GARCH	-0.002 (0.147)	-	1.020 (0.187)	0.297
VAR-ABS	0.349 (0.086)	-	1.256 (0.241)	0.115
VAR-RV+RiskMetrics	-0.022 (0.115)	0.859 (0.113)	0.219 (0.123)	0.336
VAR-RV+GARCH	-0.081 (0.144)	0.733 (0.121)	0.424 (0.247)	0.346
VAR-RV+VAR-ABS	-0.038 (0.111)	1.037 (0.159)	0.179 (0.132)	0.331
¥/DM				
VAR-RV	-0.047 (0.101)	1.146 (0.143)	-	0.355
RiskMetrics	0.248 (0.084)	-	0.668 (0.107)	0.286
GARCH	0.243 (0.092)	-	0.692 (0.119)	0.300
VAR-ABS	0.405 (0.062)	-	1.063 (0.175)	0.119
VAR-RV+RiskMetrics	-0.029 (0.093)	0.941 (0.130)	0.160 (0.117)	0.360
VAR-RV+GARCH	-0.020 (0.086)	0.873 (0.137)	0.215 (0.159)	0.364
VAR-RV+VAR-ABS	-0.039 (0.097)	1.176 (0.167)	-0.080 (0.146)	0.355

Notes: We report OLS parameter estimates for Mincer-Zarnowitz regressions of realized volatility on a constant and forecasts from different models. The regression is $(\{v_{t+1}\}_i)^{½} = b_0 + b_1 \cdot (\{v_{t+1/t,VAR\text{-}RV}\}_i)^{½} + b_2 \cdot (\{v_{t+1/t,Model}\}_i)^{½} + u_{t+1,i}$, with heteroskedasticity robust standard errors in parentheses. The forecast evaluation period covers December 2, 1996 through June 30, 1999, for a total of *596* daily observations. All model parameter estimates are based on data from December 1, 1986 through December 1, 1996. VAR-RV denotes forecasts from a long-memory vector autoregression for daily realized volatility, VAR-ABS denotes forecasts from a long-memory vector autoregression for daily log absolute returns, RiskMetrics denotes forecasts from an exponential smoothing model applied to squared daily returns, and GARCH denotes forecasts from a univariate GARCH(1,1) model. See the main text for details.

TABLE 3C
Mincer-Zarnowitz Regressions for Realized Volatilities
In-Sample, Ten-Days-Ahead

	b_0	b_1	b_2	R^2
DM/$				
VAR-RV	-0.046 (0.135)	1.036 (0.068)	-	0.436
RiskMetrics	0.712 (0.134)	-	0.638 (0.066)	0.337
GARCH	0.026 (0.189)	-	0.938 (0.090)	0.374
VAR-ABS	0.692 (0.153)	-	0.224 (0.026)	0.269
VAR-RV+RiskMetrics	-0.029 (0.139)	0.980 (0.137)	0.046 (0.097)	0.437
VAR-RV+GARCH	-0.152 (0.149)	0.801 (0.128)	0.269 (0.130)	0.445
VAR-RV+VAR-ABS	-0.055 (0.137)	0.989 (0.102)	0.017 (0.028)	0.437
¥/$				
VAR-RV	-0.012 (0.152)	1.020 (0.079)	-	0.387
RiskMetrics	0.848 (0.134)	-	0.578 (0.069)	0.299
GARCH	-0.086 (0.205)	-	0.991 (0.100)	0.366
VAR-ABS	0.784 (0.130)	-	0.097 (0.011)	0.230
VAR-RV+RiskMetrics	0.030 (0.153)	0.922 (0.131)	0.075 (0.093)	0.389
VAR-RV+GARCH	-0.239 (0.176)	0.634 (0.130)	0.467 (0.146)	0.413
VAR-RV+VAR-ABS	-0.028 (0.155)	0.944 (0.100)	0.013 (0.011)	0.389
¥/DM				
VAR-RV	-0.163 (0.102)	1.102 (0.058)	-	0.584
RiskMetrics	0.705 (0.089)	-	0.659 (0.049)	0.461
GARCH	0.444 (0.103)	-	0.782 (0.057)	0.513
VAR-ABS	0.520 (0.101)	-	0.299 (0.022)	0.436
VAR-RV+RiskMetrics	-0.131 (0.124)	1.033 (0.134)	0.054 (0.082)	0.585
VAR-RV+GARCH	-0.100 (0.105)	0.813 (0.112)	0.257 (0.078)	0.599
VAR-RV+VAR-ABS	-0.151 (0.104)	0.976 (0.106)	0.049 (0.029)	0.588

Notes: We report OLS parameter estimates for Mincer-Zarnowitz regressions of realized volatility on a constant and forecasts from different models. The regression is $(\{\sum_{j=1,..,10} v_{t+j}\}_i)^{1/2} = b_0 + b_1 \cdot (\{\sum_{j=1,..,10} v_{t+j/t,VAR\text{-}RV}\}_i)^{1/2} + b_2 \cdot (\{\sum_{j=1,..,10} v_{t+j/t,Model}\}_i)^{1/2} + u_{t+10,i}$. We report robust standard errors calculated using an unweighted covariance matrix estimator allowing for up to ninth order serial correlation in $u_{t+10,i}$. The forecast evaluation period covers December 1, 1987 through December 1, 1996, for a total of *2,223* daily observations. All model parameter estimates are based on data from December 1, 1986 through December 1, 1996. VAR-RV denotes forecasts from a long-memory vector autoregression for daily realized volatility, VAR-ABS denotes forecasts from a long-memory vector autoregression for daily log absolute returns, RiskMetrics denotes forecasts from an exponential smoothing model applied to squared daily returns, and GARCH denotes forecasts from a univariate GARCH(1,1) model. See the main text for details.

TABLE 3D
Mincer-Zarnowitz Regressions for Realized Volatilities
Out-of-Sample, Ten-Days-Ahead

	b_0	b_1	b_2	R^2
DM/$				
VAR-RV	0.417 (0.329)	0.753 (0.176)	-	0.197
RiskMetrics	1.304 (0.265)	-	0.312 (0.144)	0.056
GARCH	0.780 (0.397)	-	0.559 (0.205)	0.082
VAR-ABS	1.633 (0.241)	-	0.042 (0.037)	0.012
VAR-RV+RiskMetrics	0.441 (0.318)	0.880 (0.250)	-0.147 (0.173)	0.204
VAR-RV+GARCH	0.443 (0.340)	0.772 (0.266)	-0.031 (0.266)	0.198
VAR-RV+VAR-ABS	0.471 (0.302)	0.981 (0.234)	-0.087 (0.049)	0.232
¥/$				
VAR-RV	0.188 (0.530)	1.010 (0.227)	-	0.278
RiskMetrics	1.329 (0.322)	-	0.526 (0.133)	0.211
GARCH	0.706 (0.403)	-	0.810 (0.175)	0.263
VAR-ABS	0.935 (0.581)	-	0.113 (0.040)	0.186
VAR-RV+RiskMetrics	0.202 (0.551)	0.989 (0.292)	0.015 (0.104)	0.278
VAR-RV+GARCH	0.194 (0.502)	0.622 (0.229)	0.387 (0.196)	0.297
VAR-RV+VAR-ABS	0.059 (0.571)	0.834 (0.178)	0.035 (0.022)	0.288
¥/DM				
VAR-RV	0.154 (0.438)	1.048 (0.201)	-	0.293
RiskMetrics	1.258 (0.329)	-	0.529 (0.136)	0.299
GARCH	1.173 (0.324)	-	0.578 (0.135)	0.328
VAR-ABS	0.992 (0.533)	-	0.286 (0.097)	0.218
VAR-RV+RiskMetrics	0.614 (0.289)	0.522 (0.269)	0.301 (0.227)	0.316
VAR-RV+GARCH	0.667 (0.367)	0.399 (0.331)	0.404 (0.252)	0.341
VAR-RV+VAR-ABS	0.139 (0.452)	0.848 (0.195)	0.084 (0.118)	0.302

Notes: We report OLS parameter estimates for Mincer-Zarnowitz regressions of realized volatility on a constant and forecasts from different models. The regression is $(\{\sum_{j=1,..,10} v_{t+j}\}_i)^{½} = b_0 + b_1 \cdot (\{\sum_{j=1,..,10} v_{t+j/t,VAR\text{-}RV}\}_i)^{½} + b_2 \cdot (\{\sum_{j=1,..,10} v_{t+j/t,\,Model}\}_i)^{½} + u_{t+10,i}$. We report robust standard errors calculated using an unweighted covariance matrix estimator allowing for up to ninth order serial correlation in $u_{t+10,i}$. The forecast evaluation period covers December 2, 1996 through June 30, 1999, for a total of *596* daily observations. All model parameter estimates are based on data from December 1, 1986 through December 1, 1996. VAR-RV denotes forecasts from a long-memory vector autoregression for daily realized volatility, VAR-ABS denotes forecasts from a long-memory vector autoregression for daily log absolute returns, RiskMetrics denotes forecasts from an exponential smoothing model applied to squared daily returns, and GARCH denotes forecasts from a univariate GARCH(1,1) model. See the main text for details.

TABLE 4
Distributions of One-Day-Ahead Probability Integral Transforms
Density Forecasts From Long-Memory Lognormal-Normal Mixture Model

Quantile:	*1%*	*5%*	*10%*	*90%*	*95%*	*99%*
In-Sample						
DM/$	0.016	0.059	0.105	0.895	0.943	0.987
¥/$	0.016	0.061	0.103	0.901	0.951	0.990
Portfolio	0.010	0.052	0.091	0.912	0.958	0.990
Out-of-Sample						
DM/$	0.005	0.045	0.092	0.893	0.941	0.990
¥/$	0.019	0.055	0.099	0.884	0.956	0.993
Portfolio	0.010	0.042	0.079	0.909	0.968	1.000

Notes: We report selected points of the empirical c.d.f. for the probability integral transform of returns with respect to the density forecasts from our VAR-RV long-memory lognormal-normal mixture model; see the main text for details. The rows labeled "portfolio" give the results for an equally-weighted DM/$, ¥/$ portfolio. The in-sample forecast evaluation period is December 1, 1987 through December 1, 1996, for a total of *2,223* daily observations. The out-of-sample period is December 2, 1996 through June 30, 1999, for a total of *596* daily observations. All model parameter estimates used for out-of-sample forecasting are based on data from December 1, 1986 through December 1, 1996.

Figure 1
Return Distributions

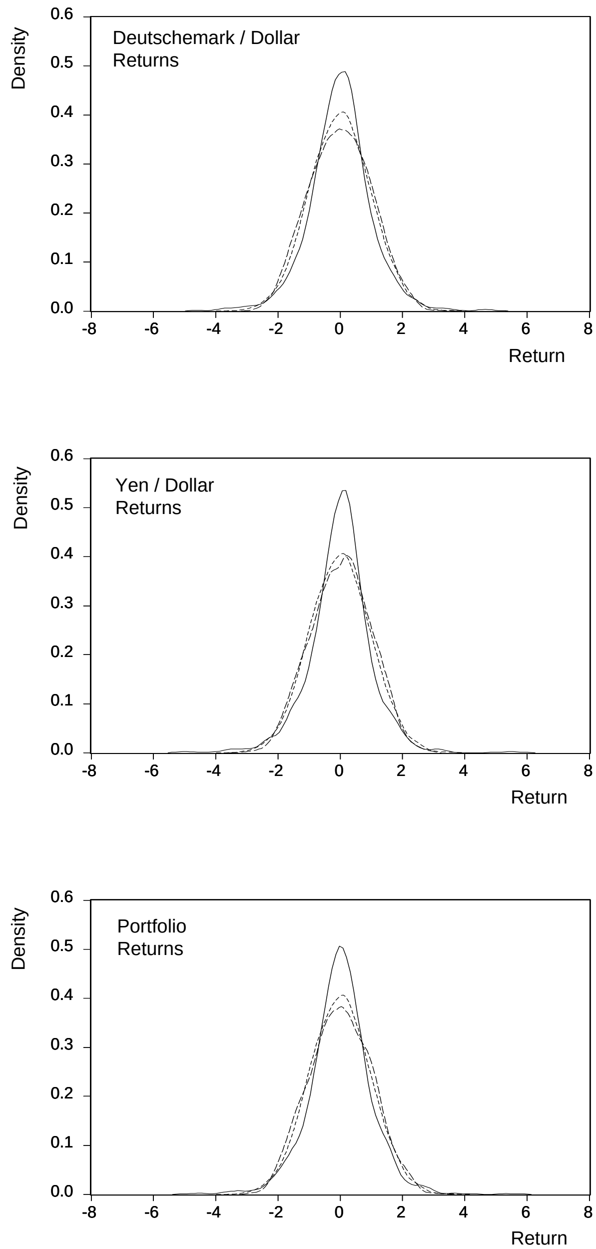

Notes: We show kernel estimates of the density of daily returns on the DM/$ rate, the ¥/$ rate, and an equally-weighted portfolio. The sample period extends from December 1, 1986 through December 1, 1996. The solid line is the estimated density of raw returns, standardized using its (constant) sample mean and sample standard deviation. The dashed line is the estimated density of returns standardized using its constant sample mean and time-varying realized standard deviation. The dotted line is a N(0,1) density for visual reference.

Figure 2
Realized Volatility Distributions

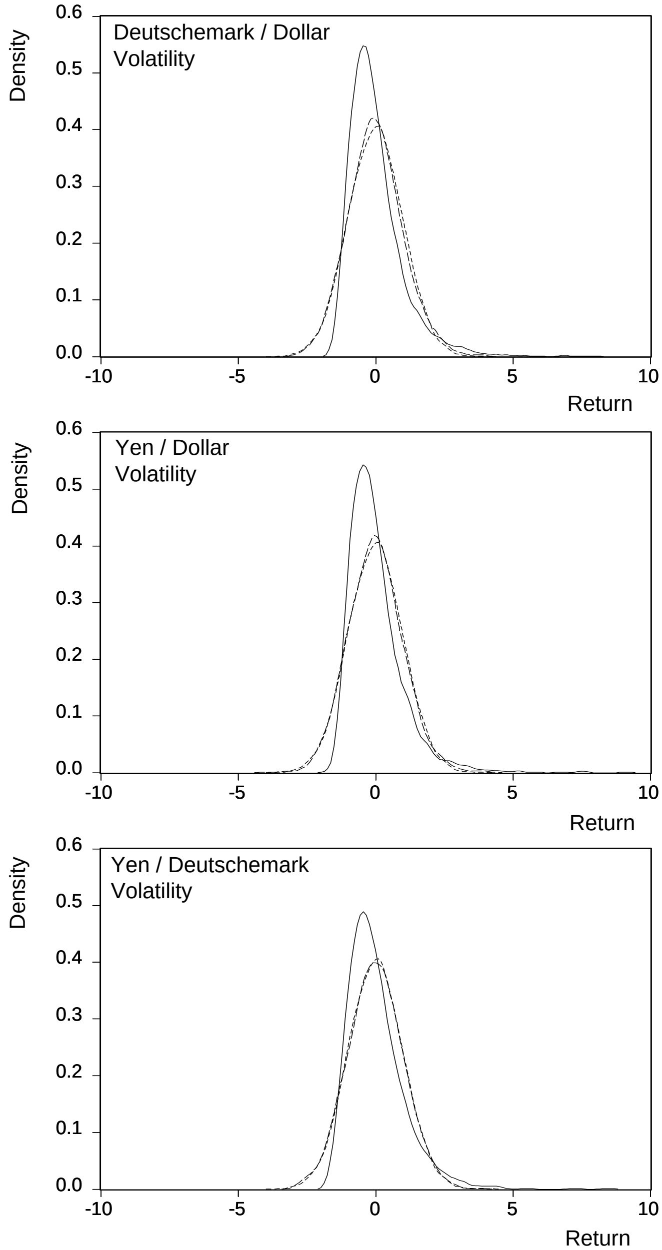

Notes: We show kernel estimates of the density of daily realized DM/$, ¥/$ and ¥/DM volatility. The sample period extends from December 1, 1986 through December 1, 1996. The solid line is the estimated density of the realized standard deviation, standardized to have zero mean and unit variance. The dashed line is the estimated density of the realized logarithmic standard deviation, standardized to have zero mean and unit variance. The dotted line is a N(0,1) density for visual reference.

Figure 3
Realized Volatility Autocorrelations

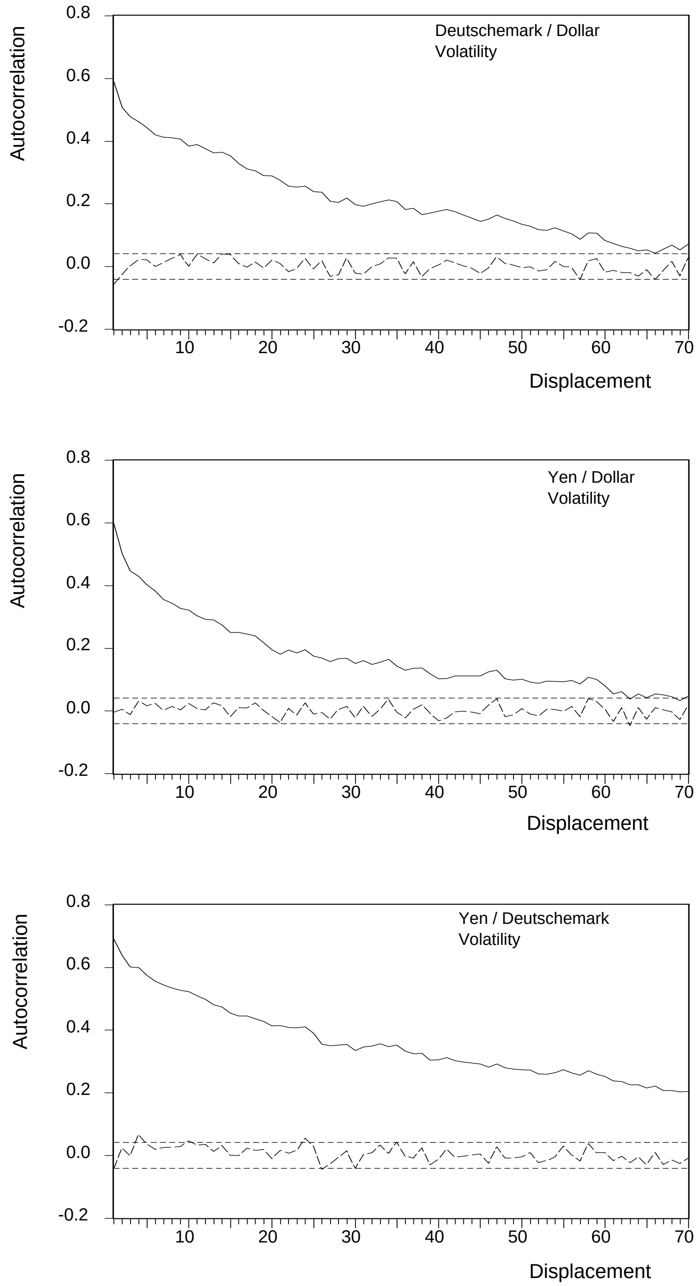

Notes: We show the sample autocorrelation functions for daily DM/$, ¥/$ and ¥/DM realized volatility. The sample period extends from December 1, 1986 through December 1, 1996. The solid line gives the autocorrelation function of realized logarithmic standard deviation, while the dashed line refers to the autocorrelation function of realized logarithmic standard deviation fractionally differenced by $(1\text{-}L)^{0.401}$. The dotted lines are the Bartlett two standard error bands.

Figure 4

Realized Volatility and Out-of-Sample VAR-RV Forecasts

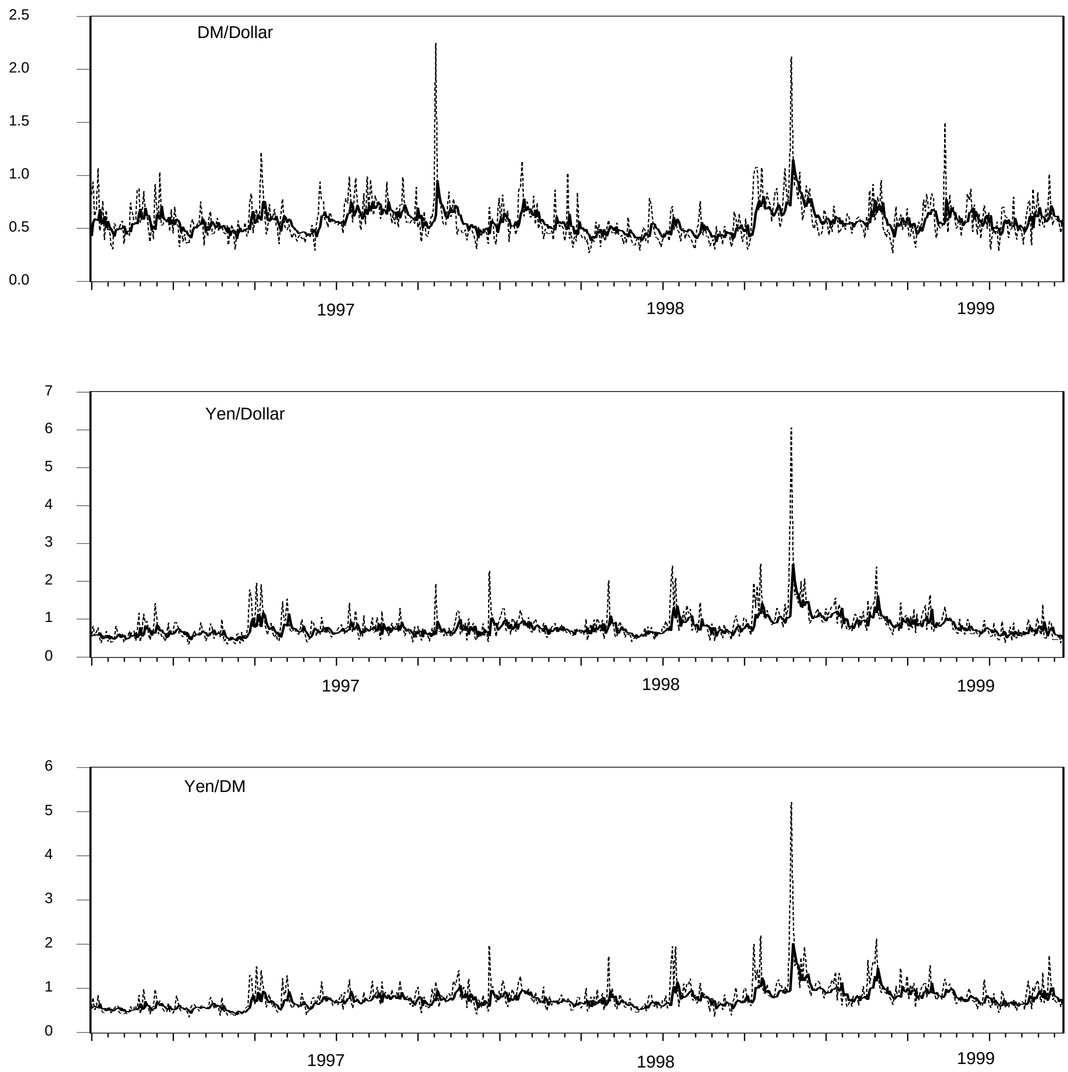

Notes: We show time series of daily realized volatility for DM/$, ¥/$ and ¥/DM, along with one-day-ahead VAR-RV forecasts. The plot spans the out-of-sample period from December 2, 1996 through June 30, 1999. The dotted line is realized volatility, while the solid line gives the corresponding one-day-ahead VAR-RV forecast from a long-memory vector autoregression for the daily realized volatility. See the main text for details.

Figure 5
Realized Volatility and Out-of-Sample GARCH Forecasts

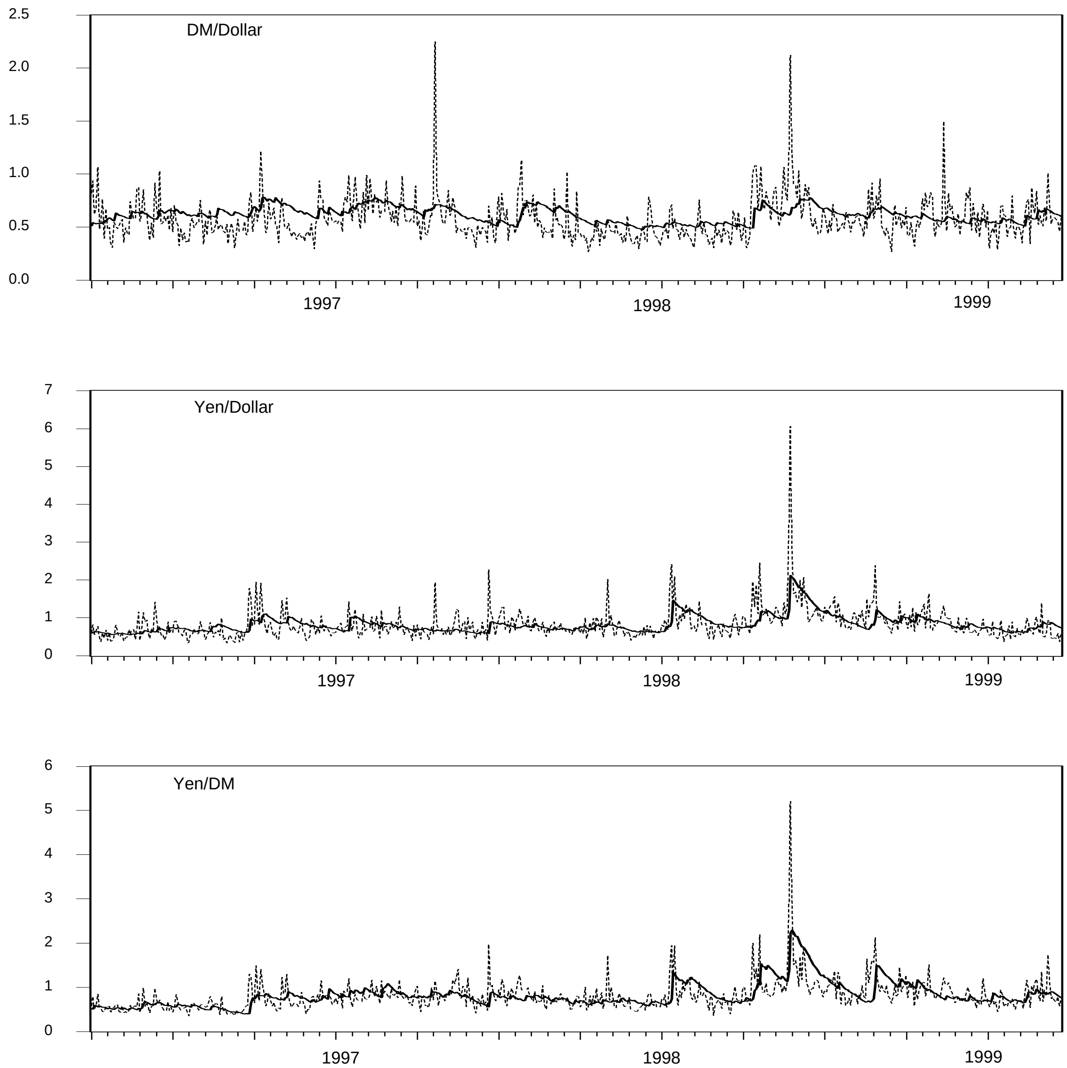

Notes: We show time series of daily realized volatility for DM/$, ¥/$ and ¥/DM, along with one-day-ahead GARCH(1,1) forecasts. The plot spans the out-of-sample period, running from December 2, 1996 through June 30, 1999. The dotted line is realized volatility, while the solid line gives the corresponding one-day-ahead GARCH forecast. See the main text for details.

Figure 6
Realized Volatility and Two Out-of-Sample Forecasts

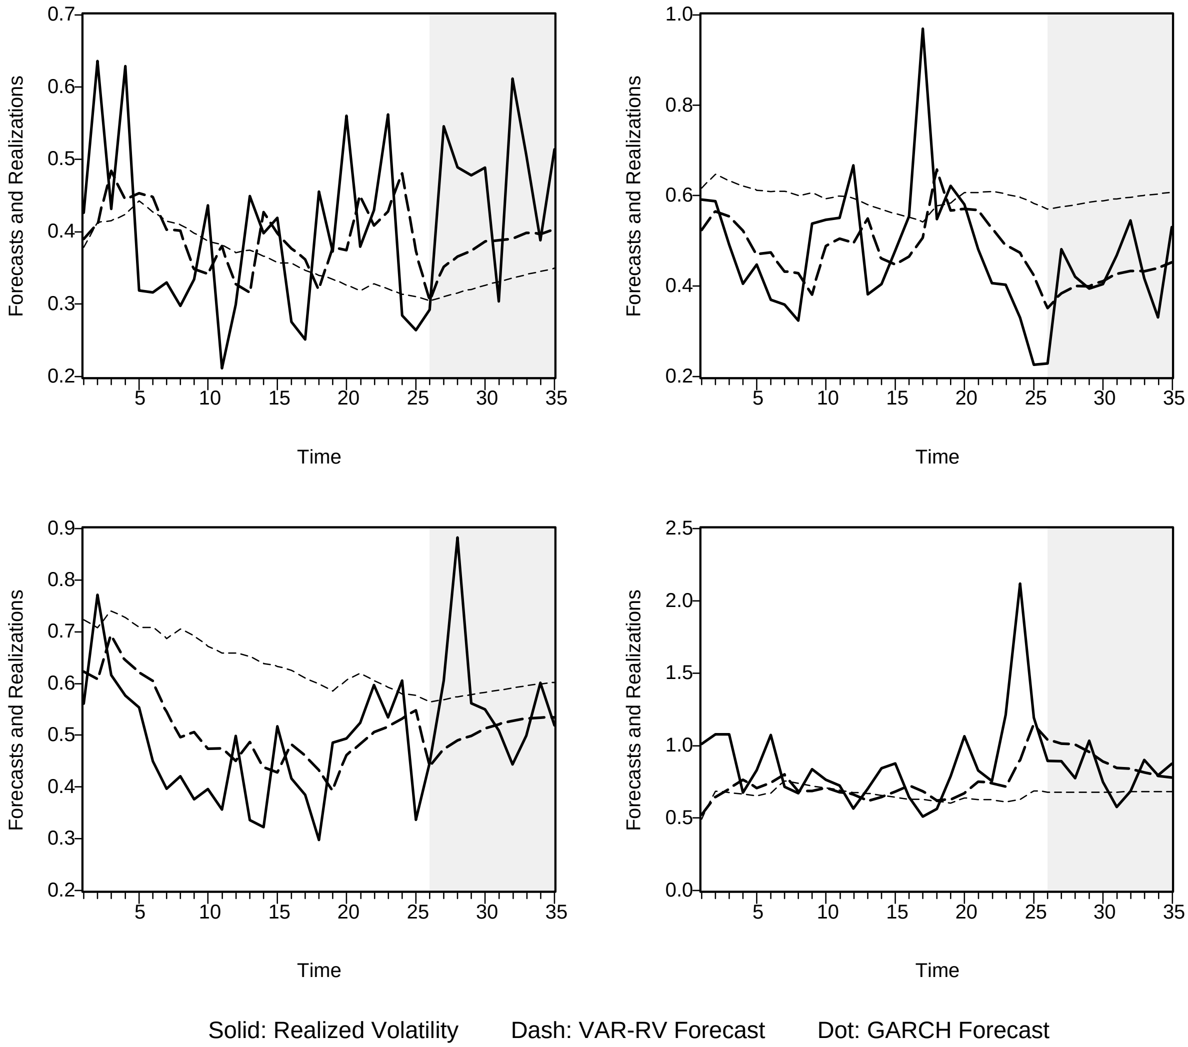

Notes: We show four out-of-sample episodes of thirty-five days each. For each of the first twenty-five days, we show the daily realized volatility together with the one-day-ahead forecasts made on the previous day using the VAR-RV and GARCH(1,1) models. Then, for days twenty-six through thirty-five (shaded), we continue to show daily realized volatility, but we show multi-step VAR-RV and GARCH forecasts based on information available on day twenty-five. Hence the forecasts for day twenty-six are one-day-ahead, the forecasts for day twenty-seven are two-day-ahead, and so on. See the main text for details.

Figure 7

Dependence Structure of Probability Integral Transforms
Out-of-Sample One-Day-Ahead Density Forecasts
Long-Memory Lognormal-Normal Mixture Model

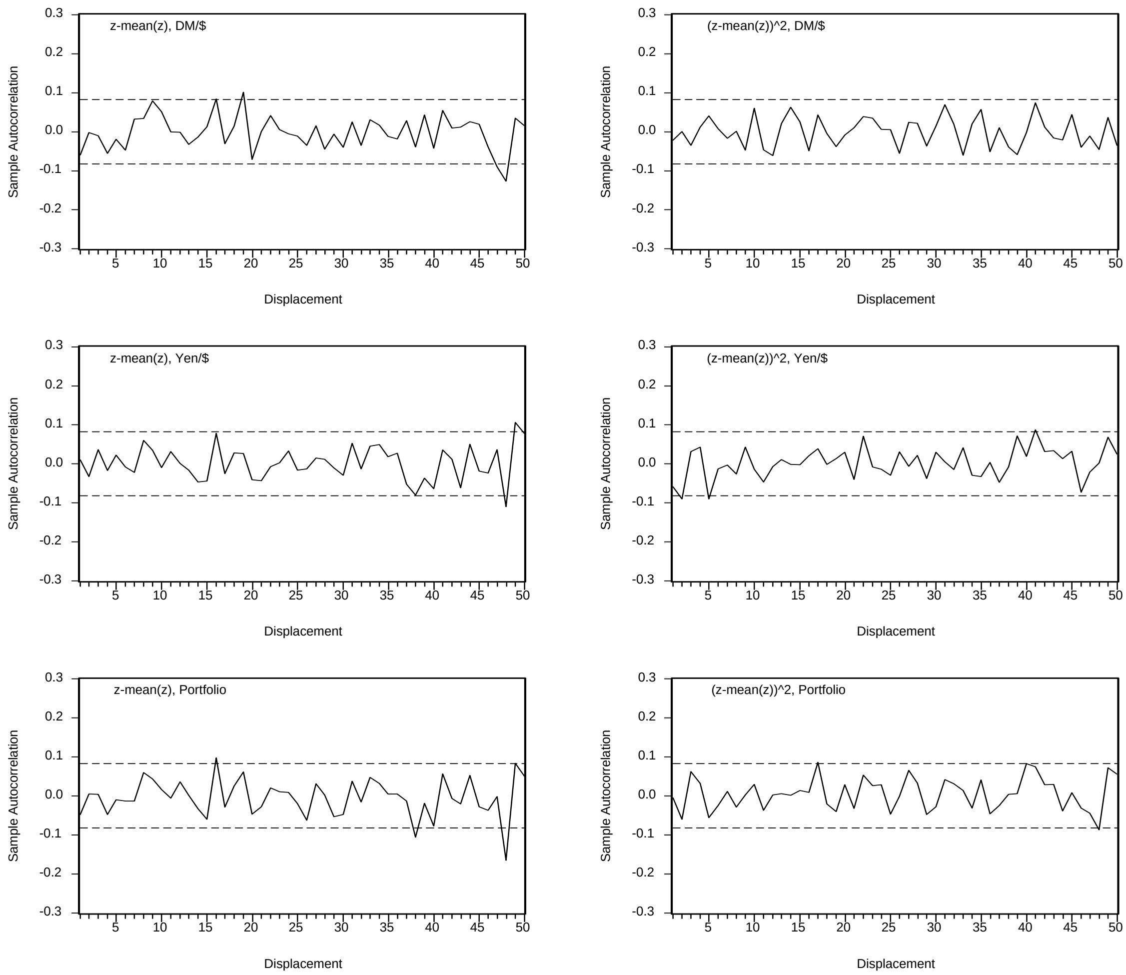

Notes: We graph the sample autocorrelation functions of $\left(z_t - \bar{z}\right)$ and $\left(z_t - \bar{z}\right)^2$, where z_t denotes the probability integral transform of returns with respect to the one-day-ahead density forecasts from our long-memory lognormal-normal mixture model; see the main text for details. Other things the same, the closer are $\left(z_t - \bar{z}\right)$ and $\left(z_t - \bar{z}\right)^2$ to white noise, the better. The three subplots correspond to DM/$ returns, ¥/$ returns, and equally-weighted portfolio returns. The dashed lines are Bartlett two standard error bands. The out-of-sample period is December 2, 1996 through June 30, 1999.