The Distribution of Exchange Rate Volatility*

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

November 1998
This Version: November 2, 1999

Abstract

Using high-frequency data on Deutschemark and Yen returns against the dollar, we construct model-free estimates of daily exchange rate volatility and correlation, covering an entire decade. In addition to being model-free, our estimates are also approximately free of measurement error under general conditions, which we delineate. Hence, for all practical purposes, we can treat the exchange rate volatilities and correlations as observed rather than latent. We do so, and we characterize their joint distribution, both unconditionally and conditionally. Noteworthy results include a simple normality-inducing volatility transformation, high contemporaneous correlation across volatilities, high correlation between correlation and volatilities, pronounced and highly persistent temporal variation in both volatilities and correlation, clear evidence of long-memory dynamics in both volatilities and correlation, and remarkably precise scaling laws under temporal aggregation.

Key Words: Financial Market Volatility; High-Frequency Data; Realized Volatility; Quadratic Variation; Exchange Rates; Long-Memory.

---

* This work was supported by the National Science Foundation. We are grateful to Olsen and Associates for making the intraday exchange rate quotations available. For useful comments we thank participants at the 1999 North American and European Meetings of the Econometric Society, the Spring 1999 NBER Asset Pricing Meeting, and seminars at Boston University, Chicago, Columbia, Georgetown, IMF, Johns Hopkins, NYU, and LSE. We are particularly grateful to the Editor and Associate Editor, two anonymous referees, Dave Backus, Michael Brandt, Rohit Deo, Clive Granger, Lars Hansen, Joel Hasbrouck, Ludger Hentschell, Cliff Hurvich, Blake LeBaron, Richard Levich, Bill Schwert, Rob Stambaugh, George Tauchen, and Stephen Taylor. The paper has been significantly abridged for publication; see the working paper version (Andersen, Bollerslev, Diebold and Labys, 1998) for additional discussion, variations, and extensions.

[a] Department of Finance, Kellogg Graduate School of Management, Northwestern University, phone: 847-467-1285, e-mail: t-andersen@nwu.edu

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

[c] Department of Finance, Stern School of Business, New York University, and NBER, phone: 610-585-4057, e-mail: fdiebold@stern.nyu.edu

[d] Graduate Group in Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104-6297, phone: 215-545-0450, e-mail: labys@ssc.sas.upenn.edu

1. Introduction

It is now widely agreed that, although daily and monthly financial asset returns are approximately unpredictable, return *volatility* is highly predictable, a phenomenon with sweeping implications for financial economics and risk management (e.g., Bollerslev, Engle and Nelson, 1994). Of course, volatility is inherently unobservable, and most of what we think we know about volatility has been learned either by fitting parametric econometric models such as GARCH, by studying volatilities implied by options prices in conjunction with specific option pricing models such as Black-Scholes, or by studying direct indicators of volatility such as ex-post squared or absolute returns. But all of those approaches, valuable as they are, have distinct weaknesses. For example, the existence of competing parametric volatility models with different properties (e.g., GARCH versus stochastic volatility models) suggests misspecification; after all, at most *one* of the models could be correct, and surely, *none* is strictly correct. Similarly, the well-known smiles and smirks in volatilities implied by Black-Scholes prices for options written at different strikes provide evidence of misspecification of the underlying model. Finally, direct indicators, such as ex-post squared returns, are contaminated by measurement error, and Andersen and Bollerslev (1998a) document that the variance of the "noise" typically is very large relative to the "signal."

In this paper, motivated by the drawbacks of the popular approaches, we provide new and complementary measures of daily asset return volatility. The mechanics are straightforward: we estimate daily volatility by summing high-frequency intraday squared returns. With sufficiently frequently sampled underlying returns, the resulting volatility estimates are largely free of measurement error. Hence, for practical purposes we can treat volatility as *observed*. We do so, and we proceed to examine its distribution directly, using much simpler techniques than those required when volatility is latent.

Our analysis is in the spirit of, and directly extends, the earlier contributions of French, Schwert and Stambaugh (1987), Hsieh (1991), and Schwert (1989, 1990), and more recently Taylor and Xu (1997). We progress, however, in a number of important ways. First, we provide rigorous theoretical underpinnings for the volatility measures for the general case of a special semimartingale. Second, much of our analysis is multivariate; we develop and examine measures not only of return variance but also of covariance. Finally, our empirical work is based on a unique high-frequency dataset consisting of ten years of continuously-recorded 5-minute returns on two major currencies. These high-frequency returns enable us to compute and examine *daily* volatilities, which are of central concern in both academia and industry. In particular, the persistent volatility fluctuations of interest in risk management, asset pricing, portfolio allocation, forecasting, and analysis of market microstructure effects are very much present in daily returns.

We proceed as follows. In Section 2 we provide a formal justification for our volatility measures. Readers who are primarily interested in the empirical results may skip the technical details in Sections 2.1 and 2.2. In Section 3, we discuss the high-frequency Deutschemark - U.S. Dollar (DM/\$) and Yen - U.S. Dollar (Yen/\$) returns that provide the basis for our empirical analysis, and we also detail the construction of our realized daily variances and covariances. In Sections 4 and 5, we characterize the unconditional and conditional distributions of the daily volatilities, including long-memory features. In Section 6, we explore issues related to temporal aggregation and scaling in relation to long memory. Finally, we conclude in Section 7 with a summary of our results and suggestions for future research.

2. Volatility Measurement: Theory

Here we develop the theoretical foundation for our realized volatility and covariance measures. We introduce the relevant concepts for the general semimartingale case, then detail how the measures may be approximated directly from high-frequency return observations. Finally, we explore the implications within the more familiar settings of Itô processes and mixed jump-diffusions.

2.1 Realized Volatility and Covariance Measures when Returns Follow a Special Semimartingale

The most general specification of asset return processes of practical relevance for financial economics is the special semimartigale. It allows for a unique, canonical decomposition of the returns into a local martingale and a predictable, integrable finite variation process. Hence, the "drift" is identified and represents the conditional mean of the instantaneous return; see Back (1991) for further discussion.

Formally, let $t \in [0,T]$, $\mathscr{F}_t$ be a σ-field reflecting the information at time t, so that $\mathscr{F}_s \subseteq \mathscr{F}_t$ for $0 \leq s \leq t \leq T$, and let P denote a probability measure on $(\Omega, P, \mathscr{F})$, where Ω represents the states of the world, so that $\mathscr{F} \equiv \mathscr{F}_T$ is the set of events that are distinguishable at the horizon T. Also, the information filtration $(\mathscr{F}_t)_{t \in [0,T]}$ satisfies the "usual conditions," i.e., it is P-complete and right continuous. Any logarithmic price process, p_k, and the associated return over the t-period horizon is then given as

$$p_k(t) \ - \ p_k(0) \ = \ M_k(t) + A_k(t) \, , \tag{1}$$

where $M_k(0) = A_k(0) = 0$, M_k is a local martingale, and A_k is a locally integrable and predictable process of finite variation. For full generality, we define p_k to be inclusive of any cash receipts such as dividends and coupons, but exclusive of required cash pay-outs associated with, for example, margin calls.

The formulation in (1) includes Itô, jump and mixed jump-diffusion processes, and it does not require a Markov assumption. Without loss of generality, each component may be assumed cadlag (right-continuous with left limits). The corresponding caglad (left-continuous with right limits) process is then defined by $p_{k-}(t) \equiv \lim_{s \to t, s<t} \ p_k(s)$, identifying the jumps as

$$\Delta \, p_k(t) \ \equiv \ p_k(t) - p_{k-}(t). \tag{2}$$

Because, by no arbitrage, jumps are not predictable, M_k contains the (compensated) jump part of p_k along with any infinite variation terms, while A_k has continuous paths. Thus, the conditional mean of the return at time t is given by the predictable "drift", $A_k(t)$, and the innovation by the local martingale $M_k(t)$.

For any semimartingale X, the stochastic integral $H \cdot X = \{\int_0^t H(s)\, dX(s)\}_{t\in[0,T]}$ is defined for any caglad integrand H (Protter, 1992, ch. 2), and may be extended to predictable processes H, i.e., H belonging to the smallest σ-algebra rendering all caglad processes measurable (Protter, ch. 4). Moreover, for $H \cdot X$ well defined, the stochastic integral constitutes a semimartingale. In particular, for any two semimartingales X and Y, the quadratic variation process, $[X,X] = ([X,X])_{t\in[0,T]}$, and the quadratic covariation process, $[X,Y] = ([X,Y])_{t\in[0,T]}$, are well defined and given by

$$[X,X] \;=\; X^2 - 2\int X_-\, dX\,, \qquad and \qquad [X,Y] \;=\; XY - \int X_-\, dY - \int Y_-\, dX. \tag{3}$$

These processes are not only semimartingales, but also of finite variation on $[0,T]$. The following properties show that they may be interpreted as the realized cumulative instantaneous variability of X and the realized cumulative instantaneous covariability between X and Y, respectively.

For a sequence of random partitions of $[0,T]$, $\tau_{m,0} \le \tau_{m,1} \le \dots$, such that $sup_{j\ge 1}(\tau_{m,j+1} - \tau_{m,j}) \to 0$, $\tau_{m,0} \to 0$, and $sup_{j\ge 1}\, \tau_{m,j} \to T$ for $m \to \infty$ with probability one,

$$lim_{m\to\infty}\{\, X^2(0) \;+\; \Sigma_{j\ge 1}\, [\, X(\, t\wedge\tau_{m,j}) - X(\, t\wedge\tau_{m,j-1})\,]^2 \,\} \;\to\; [X,X]_t\,, \tag{4}$$

and

$$lim_{m\to\infty}\{\, X(0)\, Y(0) \;+\; \Sigma_{j\ge 1}\, [X(\, t\wedge\tau_{m,j}) - X(\, t\wedge\tau_{m,j-1})]\, [Y(\, t\wedge\tau_{m,j}) - Y(\, t\wedge\tau_{m,j-1})]\,\} \;\to\; [X,Y]_t\,, \tag{5}$$

where $t\wedge\tau \equiv min(t,\tau)$, and the convergence is uniform in probability for any $t \in [0,T]$. Moreover, $[X,X]$ is a monotone increasing process,

$$[X,Y]_0 \;=\; X_0\, Y_0\,, \qquad \Delta[X,Y] \;=\; \Delta X\, \Delta Y\,, \tag{6}$$

and for H and K integrable w.r.t. X and Y, respectively,

$$[H{\cdot}X,K{\cdot}Y]_t \;=\; \int_0^t H(s)\, K(s)\, d[X,Y]_s\,, \tag{7}$$

for any $t \in [0,T]$. Finally, if X and Y are locally square integrable local martingales so that conditional variances and covariances are meaningful, then $B = [X,Y]$ is the unique adapted, cadlag process with paths of finite variation that satisfies the conditions in equation (6) and

$$X\,Y \;-\; B \;=\; X\,Y \;-\; [X,Y] \tag{8}$$

is a local martingale. Hence, $[X,Y]$ is a measure of the realized covariability between X and Y, and the covariance structure of X and Y (letting $X(0)=Y(0)=0$) is given by $E(\, X(t)Y(t)\,) \;=\; E([X,Y]_t\,)$.

For the special martingale p_k in equation (1), the finite variation part has $[A_k,A_k] = 0$, so

$$[p_k\,,p_k\,]_t \;=\; [M_k\,,\, M_k]_t \;=\; [M_k^c\,,\, M_k^c\,]_t \;+\; \Sigma_{0\le s\le t}\, (\,\Delta M_k(s)\,)^2, \tag{9}$$

where M_k has been decomposed into two local martingale components; a continuous term with infinite variation paths and a term representing the compensated jump part of the process. Now, exploiting the identical decomposition for any other logarithmic price process, p_j, $j \neq k$, we have

$$[p_k, p_j]_t \quad = \quad [M_k, M_j]_t \quad = \quad [M^c_k, M^c_j]_t \quad + \quad \Sigma_{0 \leq s \leq t} \; \Delta M_k(s) \Delta M_j(s) \, . \tag{10}$$

The formulas (9) and (10) define the *realized volatility* and *realized covariance* measures.

It is worth emphasizing the generality of the semimartingale formulation. It encompasses all processes used within the standard arbitrage-free asset pricing literature. Note, however, that it does rule out the fractional Brownian motion, $B_d(t)$, $0 < d < ½$. The latter is not a semimartingale and allows for arbitrage, as shown by Maheswaran and Sims (1993) and Rogers (1997). Formally, $B_d(t)$ is given as an infinite moving average (MA) of a standard Wiener process where the defining MA kernel has a singularity at zero. This feature is readily corrected by modifying the MA kernel at zero, thus generating a semimartingale that is consistent with the no-arbitrage condition, while retaining the basic long-memory characteristics (Rogers, 1997, provides a concrete example). However, it is arguably more relevant to allow for long-range dependence in return volatility. This may be done by positing a non-negative long-memory process for the volatility, which does not generate an arbitrage unless derivative claims written on the volatility process are traded. For instance, the option pricing model in Comte and Renault (1998) is based on a fractionally integrated log-volatility process, which violates the semimartingale property. An alternative, and perhaps preferable approach, is to modify the MA-kernel of the volatility process as suggested above, thus retaining both the long-memory in the volatility process and the validity of the standard integration theory for semimartingales.

2.2 Measurement of Realized Volatility and Realized Covariance

The local martingale formulation in (8) yields the key insight that the quadratic variation and covariation associated with the price processes provide measures of cumulative instantaneous return variability and covariability, respectively. Moreover, (4) and (5) suggest that we may approximate these quantities directly from high-frequency data. In particular, the measures are invariant to the specification of the conditional mean, since the squared mean return is an order of magnitude smaller than the squared return innovations.

Specifically, let $r_{k,(m)}(t) \equiv p_k(t) - p_k(t\text{-}1/m)$, $t = 1/m, 2/m, \ldots T$, denote the discretely sampled returns computed from m equally spaced observations per period. For concreteness, we normalize the unit time interval, or $m = 1$, to represent one trading day. It follows then directly from equations (4) and (5), that

$$plim_{m \to \infty} \quad \Sigma_{i=1,..,mt} \;\; r^2_{k,(m)}(i/m) \quad = \quad [p_k, p_k]_t \, , \tag{11}$$

and

$$plim_{m \to \infty} \quad \Sigma_{i=1,..,mt} \;\; r_{k,(m)}(i/m) \;\; r_{j,(m)}(i/m) \quad = \quad [p_k, p_j]_t \; . \tag{12}$$

Consequently, cumulative squares and cross-products of finely sampled high-frequency returns should provide a good approximation to the quadratic variation and covariation processes. The identical procedure may be used to approximate the corresponding h-period measures ($t = h, 2h, \ldots T$),

$$\sigma^2_{k,(h)}(t) \equiv [p_k, p_k]_t - [p_k, p_k]_{t-h}, \tag{13}$$

and

$$\sigma_{kj,(h)}(t) \equiv [p_k, p_j]_t - [p_k, p_j]_{t-h}. \tag{14}$$

These measures constitute time series of realized h-period volatilities and covariabilities.

It is important to recognize that $\sigma^2_{k,(h)}(t)$ and $\sigma_{kj,(h)}(t)$ are generally not measurable w.r.t. $\mathscr{F}_{t-h}$. Thus, we typically have $E[\sigma_{kj,(h)}(t)|\mathscr{F}_{t-h}] \neq \sigma_{kj,(h)}(t)$, simply because realizations differ from ex-ante expectations. An Itô process with constant volatility as is assumed for example in the Black-Scholes model, constitutes an important exception, as in that case $\sigma^2_{k,(h)}(t) \equiv \sigma^2_k h$. However, in general, the volatility measures represent *realizations* of variability and covariability rather than conditional variances and covariances. Nonetheless, $Cov[r_{k,(h)}(t), r_{j,(h)}(t)|\mathscr{F}_{t-h}] = E[\sigma_{kj,(h)}(t)|\mathscr{F}_{t-h}]$, so that the realized volatility measures do provide unbiased estimates of the ex-ante conditional variances and covariances.

Note also that there is generally no direct link between $\sigma^2_{k,(h)}(t)$ and $\sigma_{kj,(h)}(t)$, and the conditional variance and covariance at time t, $E[\sigma^2_{k,(h)}(t+h)|\mathscr{F}_t]$ and $E[\sigma_{kj,(h)}(t+h)|\mathscr{F}_t]$. Hence, the volatility measures are primarily tools for measuring realized volatilities, and not for forecasting, although they may be useful for that purpose as well. Specifically, if $\Im_{t,h}$ denotes the σ-algebra generated by the past realized variances and covariances, $\{\sigma^2_{k,(h)}(t-ih)\}$ and $\{\sigma_{kj,(h)}(t-ih)\}$ for all k,j and $i=0,1,\ldots (t-h)/h$, then, typically, $E[\sigma^2_{k,(h)}(t+h)|\mathscr{F}_t] \neq E[\sigma^2_{k,(h)}(t+h)|\Im_{t,h}]$. In particular, any parametric model allowing for an asymmetric relation between returns and volatility, or simply auxiliary state variables in the volatility dynamics, will imply that optimal forecasts should exploit information beyond $\Im_{t,h}$. Nonetheless, because of the ex-ante unbiasedness, the ex-post realizations, $\sigma^2_{k,(h)}(t+ih)$ and $\sigma_{kj,(h)}(t+ih)$, $i=1,2,\ldots (T-t)/h$, still provide the appropriate benchmark for volatility forecast evaluation.

Lastly, note that, because no financial market literally provides continuously recorded transaction prices, we cannot measure realized volatility via (11) and (12) without error. Moreover, it is not necessarily preferable to compute the measures from the highest possible frequency available, as bid-ask bounce, or dealer spread positioning, tend to induce negative autocorrelation, in turn violating the semimartingale assumption for ultra high-frequency returns. Thus, as discussed further in Section 3, some experimentation is required in practice to balance the pertinent microstructure biases against the accuracy of the continuous record asymptotics.

2.3 The Integrated Volatility Measure for Itô Processes

Much finance theory assumes that logarithmic asset prices follow a diffusion,

$$dp_k(t) \;=\; \mu_k(t)\,dt \;+\; \sigma_k(t)\,dW(t)\,, \tag{15}$$

where $W(t)$ denotes a Wiener process. Formally, in the terminology of the preceding section,

$$p_k(t) - p_k(0) \;\equiv\; r_{k,(t)}(t) \;=\; \int_0^t \mu_k(s)\,ds \;+\; \int_0^t \sigma_k(s)\,dW(s)\,. \tag{16}$$

which constitutes the canonical decomposition into a predictable, or "drift", term of finite variation, and a local martingale, or "Wiener", term. Since $[W,W]_t = t\ a.s.$, it follows from equation (7) that

$$[p_k\,,p_k\,]_t \;=\; \int_0^t \sigma_k^2(s)\,ds \;\equiv\; \sigma_{k,t}^2(t)\,. \tag{17}$$

Equation (17) defines the so-called *integrated volatility* that is central to the option pricing in Hull and White (1987), and further discussed in Andersen and Bollerslev (1998a) and Barndorff-Nielsen and Shephard (1998). The result implies that $r_{k,(t)}(t)$ conditional on $[p_k, p_k\,]_t$, is normally distributed with variance $\int_0^t \sigma_k^2(s)\,ds$.

This result extends directly to the multivariate setting. Let $W = (W_1, \ldots , W_n)$ denote an n-dimensional standard Brownian motion, and $(\mathscr{F}_t)_{t\in[0,T]}$ refer to its completed natural filtration. Then by martingale representation (Protter, Theorem 4.42) any locally square integrable Itô price process may be written as

$$p_k(t) \;-\; p_k(0) \;=\; \int_0^t \mu_k(s)\,ds \;+\; \Sigma_{i=1}^n \int_0^t \sigma_{k,i}(s)\,dW_i(s)\,, \tag{18}$$

so that in particular,

$$[p_k\,,p_k\,]_t \;\equiv\; \sigma_{k,(t)}^2(t) \;=\; \Sigma_{i=1}^n \int_0^t \sigma_{k,i}^2(s)\;ds \tag{19}$$

and

$$[p_k\,,p_j\,]_t \;\equiv\; \sigma_{kj,(t)}(t) \;=\; \Sigma_{i=1}^n \int_0^t \sigma_{k,i}(s)\;\sigma_{j,i}(s)\;ds\,. \tag{20}$$

The former provides a natural generalization of the scalar integrated volatility concept, while $\sigma_{kj,(t)}(t)$ is denoted the *integrated covariance*. As a special case, one may dedicate a few orthogonal Wiener components to be common factors while others serve as idiosyncratic error terms, providing a continuous-time analogue to the discrete-time latent factor volatility model in Diebold and Nerlove (1989).

Of course, integrated volatilities are inherently unobservable. Gallant, Hsu and Tauchen (1999) propose an intriguing reprojection method for estimating the distribution of $\sigma_{k,(h)}^2(t)$ (see also Chernov and Ghysels, 1998), but it relies on specific parametric assumptions. Motivated by (11) and (12) we, in contrast, take a direct nonparametric approach to measuring the daily integrated volatility and covariance by summing squares and cross-products of high-frequency intraday returns. The resulting realized volatility and covariance series allow us to characterize both the unconditional distribution and associated dynamic features of return volatility by standard statistical procedures.

2.4 The Integrated Volatility Measure for Pure Jump Processes and mixed Jump-Diffusions

Special semimartingales of the pure jump variety have particularly simple quadratic variation and covariation processes. The process decomposes uniquely into a compensated local martingale jump component and a finite variation term with zero quadratic variation, i.e.,

$$p_k(t) \quad = \quad p_k(0) \quad + \quad M_k(t) \quad + \quad \int_0^t \mu_k(s)\, ds \,. \tag{21}$$

The innovations to $M_k(t)$ are pure jumps, so that from (6)

$$[p_k\,, p_k\,]_t \quad = \quad \Sigma_{0\leq s\leq t} \;(\,\Delta M_k(s)\,)^2. \tag{22}$$

This result covers a variety of complex scenarios, including multiple jump components as in

$$M_k(t) \quad = \quad \Sigma_{i=1}^{J}\; \Sigma_{0\leq s\leq t}\; \kappa_{k,i}(s)\, \Delta N_{k,i}(s) \;\; - \; \int_0^t \mu_k(s)\, ds \,, \tag{23}$$

where $\mu_k(s)$ denotes the conditional mean of the overall jump process, $\Delta N_{k,i}(s)$ is an indicator for the occurrence of a jump in the i'th component at time s, and the (random) $\kappa_{k,i}(s)$ determines the jump size. Hence, in this case

$$[p_k\,, p_k\,]_t \quad = \quad \Sigma_{0\leq s\leq t}\,(\,\Delta M_k(s)\,)^2 \quad = \quad \Sigma_{i=1}^{J}\; \Sigma_{0\leq s\leq t}\; \kappa_{k,i}^2(s)\, \Delta N_{k,i}(s). \tag{24}$$

Moreover, the quadratic covariation of a pure jump process with any other semimartingale is governed exclusively by their common jumps,

$$[p_k\,, p_j\,]_t \quad = \quad \Sigma_{0\leq s\leq t}\; \Delta M_k(s)\;\; \Delta M_j(s), \tag{25}$$

which equals zero unless the processes exhibit contemporaneous jumps.

Several authors (see for example Andersen, Benzoni and Lund, 1998, for evidence and references) argue for the importance of including both time-varying volatility and jumps when modeling short-horizon returns, as in

$$p_k(t) \;\; - \;\; p_k(0) \quad = \quad \int_0^t \mu_k(s)\, ds \quad + \quad \int_0^t \sigma_k(s)\, dW(s) \quad + \quad \Sigma_{0\leq s\leq t}\; \kappa_k(s)\, \Delta N_k(s). \tag{26}$$

Again, the quadratic variation follows directly from equation (9),

$$[p_k\;\, , p_k\,]_t \quad = \quad \int_0^t \sigma_k^2(s)\, ds \quad + \quad \Sigma_{0\leq s\leq t}\;\; \kappa_k^2(s)\, \Delta N_k(s)\,. \tag{27}$$

Extensions to a multivariate setting with an n-dimensional Brownian motion and multiple jump components are straightforward, resulting in modifications along the lines of equations (18)-(20) and (23)-(25).

3. Volatility Measurement: Data

Our empirical analysis focuses on the bilateral DM/$ and Yen/$ spot exchange rates, which are particularly attractive candidates for examination as they represent the two axes of the international financial system. They also represent the most actively traded and quoted foreign currencies, and hence they permit the construction of extremely accurate volatility measures. We first rationalize the use of underlying 5-minute returns to construct daily realized volatilities, and then detail our treatment of weekend and other holiday non-trading periods. Finally, we describe the actual construction of the realized volatility measures.

3.1 On the Use of 5-Minute Returns

In practice, the inherent discreteness of actual securities prices renders continuous-time models poor approximations at very high sampling frequencies. Furthermore, high-frequency, or tick-by-tick, prices are generally only available at unevenly-spaced discrete time points, so that calculation of evenly-spaced high-frequency returns must necessarily rely on some form of interpolation involving the recorded prices around the beginning and end of a given time interval. It is well known that this non-synchronous trading or quotation effect may induce negative autocorrelation in the interpolated return series. Moreover, such biases may be exacerbated in the multivariate context, if varying degrees of interpolation are employed in the calculation of the different returns.

The sampling frequency at which microstructure biases become a practical concern is largely an empirical question. For the actively quoted and traded foreign exchange rates analyzed here, a sampling frequency of 288 times per day (5-minute returns) represents a reasonable compromise between the accuracy of the theoretical approximations and the market microstructure considerations. That is, $m=288$ is high enough such that our daily realized volatilities are largely free of measurement error (see the calculations in Andersen and Bollerslev, 1998a), yet low enough such that microstructure biases are not a major concern. (Methods for diagnosing and avoiding microstructure biases are developed in Andersen, Bollerslev, Diebold and Labys, 1999a.)

3.2 Construction of 5-Minute DM/$ and Yen/$ Returns

The two raw 5-minute DM/$ and Yen/$ return series were obtained from Olsen and Associates. The full sample consists of continuously-recorded 5-minute returns from December 1, 1986 through November 30, 1996, or 3,653 days, for a total of 3,653·288 = 1,052,064 high-frequency return observations. As in Müller et al. (1990) and Dacorogna et al. (1993), the construction of the returns utilizes all of the interbank FX quotes that appeared on the Reuters screen during the sample period. Each quote consists of a bid and an ask price together with a "time stamp" to the nearest second. After filtering the data for outliers and other anomalies, the price at each 5-minute mark is obtained by linearly interpolating from the average of the log bid and the log ask for the two closest ticks. The continuously-compounded returns are then simply the change in these 5-minute average log bid and ask prices. Goodhart, Ito and Payne (1996) and Danielsson and Payne (1999) find that the basic characteristics of 5-minute FX returns constructed from quotes closely match those calculated from transactions prices (which are not generally available).

It is well known that the activity in the foreign exchange market slows decidedly over the weekend and certain holiday non-trading periods; see, e.g., Andersen and Bollerslev (1998b) and Müller et al. (1990). In order not to confound the distributional characteristics of the various volatility measures by these largely deterministic calendar effects, we explicitly excluded a number of days from the raw 5-minute return series. Whenever we did so, we always cut from 21:05 GMT the night before to 21:00 GMT that evening, to keep the daily periodicity intact. This particular definition of a "day" was motivated by the ebb and flow in the daily FX activity patterns documented in Bollerslev and Domowitz (1993). In addition to the thin weekend trading period from Friday 21:05 GMT until Sunday 21:00 GMT, we removed several fixed holidays, including Christmas (December 24 - 26), New Year's (December 31 - January 2), and July Fourth. We also cut the moving holidays of Good Friday, Easter Monday, Memorial Day, July Fourth (when it falls officially on July 3), and Labor Day, as well as Thanksgiving and the day after. Although our cuts do not account for all of the holiday market slowdowns, they capture the most important daily calendar effects.

Finally, we deleted some of the returns contaminated by brief lapses in the Reuters data feed. This problem, which occurred almost exclusively during the early part of the sample, manifested itself in the form of sequences of zero or constant 5-minute returns in places where the missing quotes had been interpolated. To remedy this, we simply removed the days containing the fifteen longest DM/$ zero runs, the fifteen longest DM/$ constant runs, the fifteen longest Yen/$ zero runs, and the fifteen longest Yen/$ constant runs. Because of the overlap among the four different sets of days defined by these criteria, we actually removed only 51 days. All in all, we were left with 2,449 complete days, or 2,449·288 = 705,312 5-minute return observations, for the construction of our daily realized variances and covariances.

3.3 Construction of DM/$ and Yen/$ Daily Realized Volatilities

In order to define our daily volatility measures formally, we denote the time series of 5-minute DM/$ and Yen/$ returns by $\Delta logD_{(288)}(\mathrm{t})$ and $\Delta logY_{(288)}(\mathrm{t})$, respectively, where $t = 1/288, 2/288, ..., 2,449$. We then form the corresponding 5-minute squared return and cross-product series $(\Delta logD_{(288)}(t))^2$, $(\Delta logY_{(288)}(t))^2$, and $\Delta logY_{(288)}(t)\cdot\Delta logD_{(288)}(\mathrm{t})$. The statistical properties of the squared return series closely resemble those found by Andersen and Bollerslev (1997a,b) with a much shorter one-year sample of 5-minute DM/$ returns. Interestingly, the basic properties of the 5-minute cross-product series, $\Delta logY_{(288)}(t)\cdot\Delta logD_{(288)}(\mathrm{t})$, are similar. In particular, all three series are highly persistent and display strong intraday calendar effects, the shape of which is driven by the opening and closing of the different financial markets around the globe during the 24-hour trading cycle.

Now, following the results in equations (11) and (12), we construct our estimates of the daily variances and covariances by summing the 288 5-minute observations within each day,

$$vard_t \equiv \Sigma_{j=1,..,288} \ (\Delta logD_{(288)}(t\text{-}1+j/288))^2 \tag{28}$$

$$vary_t \equiv \Sigma_{j=1,..,288} \ (\Delta logY_{(288)}(t\text{-}1+j/288))^2 \tag{29}$$

$$cov_t \equiv \Sigma_{j=1,..,288} \ \Delta logD_{(288)}(t\text{-}1+j/288)\cdot\Delta logY_{(288)}(t\text{-}1+j/288), \tag{30}$$

where $t = 1, 2, ..., T;$ here T = 2449. Our focus on the squared returns as a volatility measure, as opposed to say the absolute returns, is motivated by the diffusion theoretic foundation in Section 2. Of course, squared returns also have the closest link to the variance-covariance structures and standard notions of risk employed throughout the finance literature. However, in addition we shall also examine several alternative, but related, measures of realized variation and covariation derived from the realized variances and covariances in equations (28), (29) and (30), including realized standard deviations, $stdd_t \equiv vard_t^{1/2}$ and $stdy_t \equiv vary_t^{1/2}$, realized logarithmic standard deviations, $lstdd_t \equiv ½\cdot log(vard_t)$ and $lstdy_t \equiv ½\cdot log(vary_t)$, and realized correlations, $corr_t \equiv cov_t/(stdd_t \cdot stdy_t)$. In Section 4 we characterize the unconditional distribution of each of these realized volatilities, and in section 5 we characterize their conditional distributions.

In addition to daily volatilities, we also investigate the volatility of temporally aggregated returns. In particular, let $h \geq 1$ denote the length of the return horizon. We construct temporally aggregated realized variances and covariances for h-day returns as

$$vard_{t,h} \equiv \Sigma_{j=1,..,288\cdot h} \ (\Delta logD_{(288)}(t-h+j/288))^2 \tag{31}$$

$$vary_{t,h} \equiv \Sigma_{j=1,..,288\cdot h} \ (\Delta logY_{(288)}(t-h+j/288))^2 \tag{32}$$

$$cov_{t,h} \equiv \Sigma_{j=1,..,288\cdot h} \ \Delta logD_{(288)}(t-h+j/288) \cdot \Delta logY_{(288)}(t-h+j/288), \tag{33}$$

where $t = h, 2h, \ldots, [T/h]\cdot h$. We obtain the corresponding h-day standard deviations, $stdd_{t,h}$ and $stdy_{t,h}$, logarithmic standard deviations, $lstdd_{t,h}$ and $lstdy_{t,h}$, and correlations, $corr_{t,h}$ by appropriately transforming $vard_{t,h}$, $vary_{t,h}$ and $cov_{t,h}$. In Section 6 we analyze these temporally aggregated volatilities.

4. The Unconditional Distribution of Daily Realized FX Volatility

The unconditional distribution of volatility captures an important aspect of the return variance process, and as such it has immediate implications for risk measurement and management, asset pricing, and portfolio allocation. Here we provide a detailed characterization.

4.1 Univariate Unconditional Distributions

In the first two columns of the first panel of Table 1 we show a standard menu of moments (mean, variance, skewness, and kurtosis) summarizing the unconditional distributions of the daily realized variance series, $vard_t$ and $vary_t$, and in the top panel of Figure 1 we show kernel density estimates of the unconditional distributions. It is evident that the distributions are very similar and extremely right skewed. Thus, although the realized daily variances are constructed by summing 288 squared 5-minute returns, the strong heteroskedasticity in intraday returns renders the normal distribution a poor approximation.

The standard deviation of returns is measured on the same scale as the returns, and thus provides a more readily interpretable measure of volatility than the variance. We present summary statistics and density estimates for the two daily realized standard deviations, $stdd_t$ and $stdy_t$, in columns three and four of the first panel of Table 1 and in the second panel of Figure 1. The distributions of the standard deviations are clearly non-normal, but the right skewness has been significantly reduced relative to the distributions of the variances. The mean of each daily realized standard deviation is approximately 68 basis points.

Interestingly, the distributions of the two daily realized logarithmic standard deviations, $lstdd_t$ and $lstdy_t$, displayed in columns five and six of the first panel of Table 1 and in the third panel of Figure 1, appear symmetric. Moreover, normality is a much better approximation for the logarithmic standard deviations than for the realized variances or standard deviations. This accords with the findings for monthly volatility aggregates of daily equity index returns in French, Schwert and Stambaugh (1987), as well as the earlier findings in Clark (1973) and Taylor (1986).

Finally, we characterize the distribution of daily realized covariances and correlations, cov_t and $corr_t$, in the last two columns of the first panel of Table 1 and the bottom panel of Figure 1. The basic characteristic of the unconditional distribution of the covariance is similar to that of the two daily variances -- it is extremely right skewed and leptokurtic. Interestingly, however, the distribution of the realized correlation is close to normal. The mean realized correlation is positive (0.43), which is not surprising, as it may arise from common dependence on U.S. macroeconomic fundamentals. The standard deviation of realized correlation (0.17) indicates significant variation of the correlation around its mean, which may be important for short-term portfolio allocation and hedging decisions.

4.2 Multivariate Unconditional Distributions

The univariate distributions characterized above do not address relationships that may exist among the different measures of variation and covariation. Key financial and economic questions, for example, include whether the individual volatilities such as $lstdd_t$ and $lstdy_t$ move together, and whether they are positively correlated with movements in correlation. Although such questions are difficult to answer using conventional volatility models, they are relatively easy to address using our realized volatilities and correlations.

The sample correlations in the first panel of Table 2, along with the $lstdd_t$-$lstdy_t$ scatterplot in the top panel of Figure 2, indicate a strong positive association between the two exchange rate volatilities. Thus, not only do the two exchange rates tend to move together, as indicated by the positive means for cov_t and $corr_t$, but their volatilities are also closely linked. This provides empirical justification for the use of multivariate volatility models with a factor structure, as in Diebold and Nerlove (1989) and Bollerslev and Engle (1993).

The correlation figures in Table 2 along with the $corr_t$-$lstdd_t$ scatterplot in the second panel of Figure 2 also indicate a positive association between correlation and volatility. To quantify further this "volatility effect" in correlation, we show in the top panel of Figure 3 kernel density estimates of $corr_t$ when both $lstdd_t$ and $lstdy_t$ are less than -0.46 (their median value, which happens to be the same for each) and when both $lstdd_t$ and $lstdy_t$ are greater than -0.46. Similarly, we show in the bottom panel of Figure 3 the estimated $corr_t$

densities conditional on the more extreme volatility situation in which both $lstdd_t$ and $lstdy_t$ are less than -0.87 (approximately the tenth percentile of each distribution) and when both $lstdd_t$ and $lstdy_t$ are greater than 0.00 (approximately the ninetieth percentile of each distribution). In each case, the distribution of $corr_t$ conditional on being in the high volatility state is clearly shifted to the right. A similar correlation effect in volatility has been documented for international equity returns by Solnik, Boucrelle and Le Fur (1996) among others, while Ang and Bekaert (1999) have explored the optimal portfolio implications of such an effect. Of course, given that the high-frequency returns are positively correlated, some separation is to be expected (e.g., Ronn, 1998, and Forbes and Rigobon, 1999). However, the magnitude of the effect is nonetheless noteworthy.

5. The Conditional Distribution of Daily Realized FX Volatility

The value of a derivative security such as an option is closely linked to the expected volatility of the underlying asset until expiration. Hence improved volatility forecasts should, for example, lead to more accurate option pricing. The conditional dependence in volatility forms the basis for such forecasts. This feature is most easily identified in the daily realized correlations and logarithmic standard deviations which are approximately unconditionally normally distributed. To conserve space, we focus on those three series.

It is instructive first to consider the time series plots of the realized volatilities in Figure 4. The wide fluctuations and strong persistence evident in each of the univariate $lstdd_t$ and $lstdy_t$ series are of course manifestations of the widely documented return volatility clustering. It is striking that the time series plot for $corr_t$ shows equally pronounced persistence, with readily identifiable periods of high and low correlation.

This visual impression is borne out by the highly significant Ljung-Box tests reported in the first row of the first panel of Table 3. (The *0.001* critical value is 45.3.) The correlograms of $lstdd_t$, $lstdy_t$ and $corr_t$ in Figure 5 further underscore the point. The autocorrelations of the logarithmic standard deviations begin around 0.6 and decay very slowly to about 0.1 at a displacement of 100 days. Those of the realized daily correlations decay even more slowly, reaching 0.31 at the 100-day displacement. Similar results based on long series of daily absolute or squared returns from other markets have previously been obtained by a number of authors, including Ding, Granger and Engle (1993). The slow decay in Figure 5 is particularly noteworthy, however, in that the two realized daily volatility series span “only” ten years.

The findings of slow autocorrelation decay might indicate the presence of a unit root, as in the integrated GARCH model of Engle and Bollerslev (1986). However, Dickey-Fuller tests with ten augmentation lags soundly reject the unit root hypothesis for all the series, with test statistics ranging from -9.26 to -5.59, while the 0.01 and 0.05 critical values are -2.86 and -3.43. Although unit roots are soundly rejected, the very slow autocorrelation decay coupled with the negative and slowly decaying estimated augmentation lag coefficients in the Dickey-Fuller regressions still suggest that long-memory of a non unit-root variety is present. Hence we now turn to an investigation of fractional integration in the daily realized volatilities.

As noted by Granger and Joyeux (1980), a slow hyperbolic decay of the long-lag autocorrelations or, equivalently, the log-linear explosion of the low-frequency spectrum, are distinguishing features of a covariance stationary fractionally integrated, or I(d), process with $0 < d < ½$,. The low-frequency spectral behavior also forms the basis for the log-periodogram regression procedure of Geweke and Porter-Hudak (1983) and later refinements by Robinson (1994, 1995), Hurvich and Beltrao (1994) and Hurvich, Deo and Brodsky (1998). In particular, let $I(\omega_j)$ denote the sample periodogram at the jth Fourier frequency, $\omega_j = 2\pi j/T$, $j = 1, 2, ..., [T/2]$. The log-periodogram estimator of d is then based on the OLS regression,

$$log[\, I(\omega_j)\,] \quad = \quad \beta_0 \quad + \quad \beta_1 \cdot log(\omega_j) \quad + \quad u_j\,, \tag{34}$$

where $j = 1, 2, ..., m$, and $\hat{d} \equiv -½\cdot\hat{\beta}_1$. The least squares estimator of β_1, and hence $\hat{d}$, is asymptotically normal and the corresponding standard error, $\pi\cdot(24\cdot m)^{-½}$, depends only on the number of periodogram ordinates used. Although the earlier proofs for consistency and asymptotic normality of the log-periodogram regression estimator rely on normality, Deo and Hurvich (1998) and Robinson and Henry (1998) have recently shown that these same properties extend to non-Gaussian, possibly heteroskedastic, time series as well. Of course, the actual estimate of d depends upon the specific choice of m. While the theoretical standard error formula suggests choosing a large value of m in order to obtain a small standard error, doing so may induce a bias, because the relationship underlying equation (34) in general holds only for frequencies close to zero. Following Taqqu and Teverovsky (1996), we therefore graphed and examined $\hat{d}$ as a function of m, looking for a "flat region" in which we are plagued neither by high variance (m "too small") nor high bias (m "too large"). Our subsequent choice of $m = [\, T^{4/5}\,]$, or $m = 514$, is consistent with the optimal rate of $O(T^{4/5})$ established by Hurvich, Deo and Brodsky (1998).

We present the estimates of d in the second row of the first panel of Table 3. The estimates for all eight volatility series are highly statistically significant, and all are fairly close to the "typical value" of 0.4. These estimates for d are also directly in line with the estimates based on long time series of daily absolute and squared returns from other markets reported by Granger, Ding and Spear (1997), as well as the findings based on a much shorter one-year sample of intraday DM/$ returns reported in Andersen and Bollerslev (1997b). The results therefore suggest that the standard continuous-time models applied in much of the theoretical finance literature, in which the volatility is assumed to follow an Ornstein-Uhlenbeck (OU) process, are misspecified. Our results are also constructive, however, in that they indicate that simple and parsimonious long-memory models should accurately capture the long-lag autoregressive effects.

6. The Effects of Temporal Aggregation

The analysis in the preceding sections focused exclusively on the distributional properties of daily realized volatilities. However, many practical problems in asset pricing, portfolio allocation, and financial risk management invariably involve longer horizons. Here we examine the distributional aspects of the corresponding multi-day realized variances and correlations. As before, we begin with an analysis of unconditional distributions, followed by an analysis of the dynamic dependence, including a detailed examination of long-memory as it relates to temporal aggregation.

6.1 Univariate and Multivariate Unconditional Distributions

The lower panels of Table 1 provide summary statistics for the univariate unconditional distributions for the temporally aggregated volatility measures at weekly, bi-weekly, tri-weekly and monthly return horizons ($h = 5,\ 10, 15$, and 20, respectively, corresponding to sample sizes of 489, 244, 163 and 122 days). By construction, the means of the volatility series $vard_{t,h}$, $vary_{t,h}$, and $cov_{t,h}$ grow at the constant rate h, while the mean realized correlation, $corr_{t,h}$, is largely invariant to the level of aggregation. Of significantly, the growth of the variance of the realized variances and covariance adheres closely to h^{2d+1}, where d denotes the order of integration, a phenomenon that we discuss at length subsequently. Observe also that, even at the monthly level, the unconditional distributions of $vard_{t,h}$, $vary_{t,h}$, and $cov_{t,h}$ remain leptokurtic and highly rightward skewed. The basic characteristics of $sttd_{t,h}$ and $stdy_{t,h}$ are similar, with the mean increasing at the rate $h^{1/2}$. In contrast to previously, however, the unconditional variances of $lstdd_{t,h}$ and $lstdy_{t,h}$ now *decrease* with h, but again at a rate linked to the fractional integration parameter, as we document below.

Next, turning to the multivariate unconditional distributions, we display in the lower panels of Table 2 the correlation matrices of all volatility measures for $h = 5, 10, 15$, and 20. While the correlation between the different measures of volatility drops slightly under temporal aggregation, the strong positive association between the volatilities so apparent at the one-day return horizon is largely preserved under temporal aggregation. For instance, the correlation between $lstdd_{t,h}$ and $lstdy_{t,h}$ ranges from a high of 0.604 at the daily horizon to a low of 0.533 at the monthly horizon. Meanwhile, the volatility effect in correlation is reduced somewhat under temporal aggregation; the sample correlation between $lstdd_{t,1}$ and $corr_{t,1}$ equals 0.389, whereas the correlation between $lstdd_{t,20}$ and $corr_{t,20}$ is 0.245. Similarly, the correlation between $lstdy_{t,h}$ and $corr_{t,h}$ drops from 0.294 for $h = 1$ to 0.115 for $h = 20$. Thus, while the long-horizon correlations are still positively related to the overall level of volatility, the lower numerical values suggest that the benefits to international diversification may be the greatest over longer investment horizons.

6.2 The Conditional Distribution: Dynamic Dependence, Fractional Integration and Scaling

Andersen, Bollerslev and Lange (1999) have recently shown that, given the estimates typically obtained at the daily level, from a theoretical perspective the integrated volatility should remain strongly serially correlated and highly predictable under temporal aggregation, even at the monthly level. The Ljung-Box statistics for the realized volatilities presented in the lower panels of Table 3 provide strong empirical confirmation. Even at the monthly level, or $h = 20$, with only 122 observations, all of the test statistics are highly significant. This contrasts with other sorts of evidence, which tends to show little or no significant evidence of volatility clustering by the time one aggregates to monthly returns, as in Baillie and Bollerslev (1989) and Christoffersen and Diebold (2000).

The estimates for d in Section 4 all suggest that the realized daily volatilities are fractionally integrated. The class of fractionally integrated models is self-similar, so that the degree of fractional integration should be invariant to the sampling frequency; see, e.g., Beran (1994). This strong prediction is borne out by the estimates for d at the different levels of temporal aggregation, given in the lower panels of Table 3. All of the estimates are within two asymptotic standard errors of the average estimate of 0.391 obtained for the daily series, and all are highly statistically significantly different from both zero and unity.

Another implication of self-similarity concerns the variance of partial sums. In particular, let

$$[\, x_t \,]_h \;\equiv\; \Sigma_{j=1,..,h}\; x_{h\cdot(t-1)+j}\,, \tag{35}$$

denote the h-fold partial sum process for x_t, where $t = 1, 2, ..., [T/h]$. Then, as discussed by, e.g., Beran (1994) and Diebold and Lindner (1996), if x_t is fractionally integrated, the partial sums obey a *scaling law*,

$$Var(\,[\, x_t \,]_h\,) \;=\; c\cdot h^{2d+1}. \tag{36}$$

Of course, by definition $[\, vard_t \,]_h \equiv vard_{t,h}$ and $[\, vary_t \,]_h \equiv vary_{t,h}$, so the variance of the *realized* volatilities should grow at the rate h^{2d+1}. This theoretical implication is remarkably consistent with the unconditional sample variances and covariances in Table 1 for values of d around 0.35-0.40. Similar scaling laws for power transforms of absolute FX returns are reported in a series of papers initiated by Müller et al. (1990).

The striking accuracy of our scaling laws carries over to the partial sums of the alternative volatility series. The left panel of Figure 6 plots the logarithm of the sample variances of the partial sums for the realized logarithmic standard deviations versus the logarithm of the aggregation level; i.e., $log(\, Var([\, lstdd_t]_h\,))$ and $log(Var([\, lstdy_t\,]_h\,))$ against $log(h)$ for $h = 1, 2, ..., 30$. The linear fits implied by equation (36) are validated. Each of the slopes are very close to the theoretical value of $2d+1$ implied by the log-periodogram estimates for d, further solidifying the notion of long-memory volatility dependence. The estimated slopes in the top and bottom panels are 1.780 and 1.728, respectively, corresponding to d values of 0.390 and 0.364.

Because a non-linear function of a sum does not equal the sum of the non-linear function, it is not clear whether $lstdd_{t,h}$ and $lstdy_{t,h}$ will follow similar scaling laws. The estimates of d reported in Table 3 suggest that they should. The corresponding plots for the logarithm of the h-day logarithmic standard deviations

$log(Var(lstdd_{t,h}))$ and $log(Var(lstdy_{t,h}))$ against $log(h)$, for $h = 1, 2, ..., 30$, in the right panel of Figure 6, lend empirical support to this conjecture. Although the fits are not as perfect as those in Figure 8, the log-linear approximations are still remarkably accurate. Interestingly, however, the lines are downward sloped.

To understand why these slopes may be negative, assume that the returns are serially uncorrelated. The variance of the temporally aggregated return should then be proportional to the length of the return interval, that is, $E(var_{t,h}) = b \cdot h$, where $var_{t,h}$ refers to the temporally aggregated variance as defined above. Also, by the scaling law in (36), $Var(var_{t,h}) = c \cdot h^{2d+1}$. Furthermore, assume that the corresponding temporally aggregated logarithmic standard deviations, $lstd_{t,h} \equiv \frac{1}{2} \cdot log(var_{t,h})$, are normally distributed across all frequencies h with mean μ_h and variance σ_h^2. Note that these assumptions accord closely with the empirical distributions summarized in Table 1. It then follows from the properties of the lognormal distribution that

$$E(var_{t,h}) = exp(2\mu_h + 2\sigma_h^2) = b \cdot h \tag{37}$$

and

$$Var(var_{t,h}) = exp(4\mu_h)\, exp(4\sigma_h^2)\, [\, exp(4\sigma_h^2) - 1\,] = c \cdot h^{2d+1}, \tag{38}$$

so that solving for the variance of the log standard deviation yields

$$Var(lstd_{t,h}) \equiv \sigma_h^2 = log(c \cdot b^{-2} \cdot h^{2d-1} + 1). \tag{39}$$

With $2d\text{-}1$ slightly negative, this explains why the sample variances of $lstdd_{t,h}$ and $lstdy_{t,h}$ reported in Table 1 are decreasing with the level of temporal aggregation, h. Furthermore, by a log-linear approximation,

$$log[\, Var(lstd_{t,h})\,] \approx a + (2d\text{-}1) \cdot log(h), \tag{40}$$

which provides a justification for the apparent scaling law behind the two plots in the right panel of Figure 6, and the negative slopes of approximately $2d\text{-}1$. The slopes in the top and bottom panels are -0.222 and -0.270, respectively, and the implied d values of 0.389 and 0.365 are almost identical to the values implied by the scaling law in equation (36) and the two left panels of Figure 6.

7. Summary and Concluding Remarks

We have provided a theoretical basis for measuring and analyzing time series of realized volatilities constructed from high-frequency intraday returns. Utilizing a unique data set consisting of ten years of 5-minute DM/\$ and Yen/\$ returns, we find that the distributions of realized daily variances, standard deviations and covariances are skewed to the right and leptokurtic, but that the distributions of logarithmic standard deviations and correlations are approximately Gaussian. Volatility movements, moreover, are highly correlated across the two exchange rates, as would be implied by a factor structure induced by common dependence on U.S. fundamentals. We also find that the correlation between the exchange rates (as opposed to the correlation between their volatilities) increases with volatility, so that the benefits of portfolio diversification are reduced just when they are needed most. Finally, we confirm the wealth of existing evidence of strong volatility clustering effects in daily returns. However, in contrast to earlier work, which often indicates that volatility persistence decreases fairly quickly with the horizon, we find that even monthly realized volatilities remain highly persistent. Nonetheless, realized volatilities do not have unit roots; instead, they appear fractionally integrated and therefore very slowly mean-reverting. This finding is strengthened by our analysis of temporally aggregated volatility series, which appear to be governed by remarkably accurate scaling laws, as predicted by the structure of fractional integration.

A key conceptual distinction between this paper and the earlier work on which we build -- Andersen and Bollerslev (1998a) in particular -- is the recognition that realized volatility is usefully viewed as the object of intrinsic interest, rather than simply a post-modeling device to be used for evaluating parametric volatility models such as GARCH. As such, it is of interest to examine and model realized volatility directly. This paper is a first step in that direction, providing a nonparametric characterization of both the unconditional and conditional distributions of bivariate realized exchange rate volatility.

It will be of interest in future work to fit parametric models directly to the realized volatility, and in turn use them for forecasting in specific financial contexts. In particular, our findings suggest that modeling realized logarithmic daily standard deviations and correlations by a linear Gaussian multivariate long-memory model, could result in important improvements in the accuracy of long-term volatility forecasts and Value-at-Risk type calculations. This idea is pursued in Andersen, Bollerslev, Diebold and Labys (1999b).

REFERENCES

Andersen, T.G., Benzoni, L. and Lund, J. (1998), "Estimating Jump-Diffusions for Equity Returns," Manuscript, Department of Finance, Kellogg Graduate School of Management, Northwestern University.

Andersen, T.G. and Bollerslev, T. (1997a), "Intraday Periodicity and Volatility Persistence in Financial Markets," *Journal of Empirical Finance*, 4, 115-158.

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

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

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

Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P. (1998), "The Distribution of Exchange Rate Volatility," NBER Working Paper 6961.

Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P. (1999a), "Microstructure Bias and Volatility Signatures," work in progress.

Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P. (1999b), "Forecasting Realized Volatility: A VAR for VaR," work in progress.

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

Ang, A. and Bekaert, G. (1999), "International Asset Allocation with Time-Varying Correlations," NBER Working Paper No.7056.

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

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

Barndorff-Nielsen, O.E. and Shephard, N. (1998), "Aggregation and Model Construction for Volatility Models," Manuscript, Nuffield College, Oxford, UK.

Beran, J. (1994), *Statistics for Long-Memory Processes*, New York: Chapman and Hall.

Bollerslev, T. and Domowitz, I. (1993), "Trading Patterns and Prices in the Interbank Foreign Exchange Market," *Journal of Finance*, 48, 1421-1443.

Bollerslev, T. and Engle, R.F. (1993), "Common Persistence in Conditional Variances," *Econometrica*, 61, 166-187.

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

Chernov, M. and Ghysels, E. (1998), "What Data Should be Used to Price Options?," Manuscript, Department of Economics, Penn State University.

Christoffersen, P.F. and Diebold, F.X. (2000), "How Relevant is Volatility Forecasting for Financial Risk Management?," *Review of Economics and Statistics*, 82, forthcoming.

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

Comte, F. and Renault, E. (1998), "Long Memory in Continuous-Time Stochastic Volatility Models," *Mathematical Finance*, 8, 291-323.

Dacorogna, M.M., Müller, U.A., Nagler, R.J., Olsen, R.B., and Pictet, O.V. (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.

Danielsson, J. and Payne, R. (1999), "Real Trading Patterns and Prices in Spot Foreign Exchange Markets," Manuscript, Financial Markets Group, London School of Economics.

Deo, R. and Hurvich, C.M. (1998), "On the Log Periodogram Regression Estimator of the Memory Parameter in Long Memory Stochastic Volatility Models," Working Paper SOR-98-04, Stern School of Business, New York University.

Diebold, F.X. and Lindner, P. (1996), "Fractional Integration and Interval Prediction," *Economics Letters*, 50, 305-313.

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

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

Engle, R.F. and Bollerslev, T. (1986), "Modeling the Persistence of Conditional Variances," *Econometric Reviews*, 5, 1-50.

Forbes, K. and Rigobon, R. (1999), "No Contagion, Only Interdependence: Measuring Stock Market Co-Movements," NBER Working Paper No.7267.

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

Gallant, A.R., Hsu, C.-T., and Tauchen, G.E. (1999), "Calculating Volatility Diffusions and Extracting Integrated Volatility," *Review of Economics and Statistics*, forthcoming.

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

Goodhart, C., Ito, T., and Payne, R. (1996), "One Day in June 1993: A Study of the Working of the Reuters 2000-2 Electronic Foreign Exchange Trading System," in *The Microstructure of Foreign Exchange Markets*, 107-179, eds. J.A. Frankel, G. Galli and A. Giovannini), Chicago: Univ. of Chicago Press.

Granger, C.W.J., Ding, Z., and Spear, S. (1997), "Stylized Facts on the Temporal and Distributional Properties of Daily Data From Speculative Markets," Manuscript, Department of Economics, University of California, San Diego.

Granger, C.W.J. and Joyeux, R. (1980), "An Introduction to Long Memory Time Series Models and Fractional Differencing," *Journal of Time Series Analysis*, 1, 15-39.

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

Hull, J. and White, A. (1987), "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42, 381-400.

Hurvich, C.M. and Beltrao, K.I. (1994), "Automatic Semiparametric Estimation of the Memory Parameter of a Long-Memory Time Series," *Journal of Time Series Analysis*, 15, 285-302.

Hurvich, C.M., Deo, R., and Brodsky, J. (1998), "The Mean Squared Error of Geweke and Porter Hudak's Estimator of the Memory Parameter of a Long-Memory Time Series," *Journal of Time Series Analysis*, 19, 19-46.

Maheswaran, S. and Sims, C.A.. (1993), "Empirical Implications of Arbitrage-Free Asset Markets," in *Models, Methods and Applications of Econometrics: Essays in Honor of A.R. Bergstrom*, ed. P.C.B. Phillips, Cambridge, MA: Blackwell Publishers.

Müller, U.A., Dacorogna, M.M., Olsen, R.B., Pictet, O.V., Schwarz, M., and Morgenegg, C. (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.

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

Robinson, P.M. (1994), "Semiparametric Analysis of Long-Memory Time Series," *Annals of Statistics*, 22, 515-539.

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

Robinson, P.M. and Henry, M. (1998), "Long and Short Memory Conditional Heteroskedasticity in Estimating the Memory Parameter of Levels," Manuscript, London School of Economics.

Rogers, L.C.G. (1997), "Arbitrage with Fractional Brownian Motion," *Mathematical Finance*, 7, 95-105.

Ronn, E. (1998), "The Impact of Large Changes in Asset Prices on Intra-Market Correlations in the Stock and Bond Markets," Manuscript, Department of Finance, University of Texas, Austin.

Schwert, G.W. (1989), "Why does Stock Market Volatility Change over Time?," *Journal of Finance*, 44, 1115-1154.

Schwert, G.W. (1990), "Stock Market Volatility," *Financial Analysts Journal*, May-June, 23-34.

Solnik, B., Boucrelle, C., and Le Fur, Y. (1996), "International Market Correlation and Volatility," *Financial Analysts Journal*, September-October, 17-34.

Taqqu, M.S. and Teverovsky, V. (1996), "Semi-Parametric Graphical Estimation Techniques for Long-Memory Data," in *Time Series Analysis in Memory of E.J. Hannan*, 420-432, eds. P.M. Robinson and M. Rosenblatt, New York: Springer-Verlag.

Taylor, S. J. (1986), *Modelling Financial Time Series*, Chichester: John Wiley and Sons.

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

Table 1
Statistics Summarizing Unconditional Distributions of Realized DM/$ and Yen/$ Volatilities

	$vard_{t,h}$	$vary_{t,h}$	$stdd_{t,h}$	$stdy_{t,h}$	$lstdd_{t,h}$	$lstdy_{t,h}$	$cov_{t,h}$	$corr_{t,h}$
Daily, h=1								
Mean	0.529	0.538	0.679	0.684	-0.449	-0.443	0.243	0.435
Variance	0.234	0.272	0.067	0.070	0.120	0.123	0.073	0.028
Skewness	3.711	5.576	1.681	1.867	0.345	0.264	3.784	-0.203
Kurtosis	24.09	66.75	7.781	10.38	3.263	3.525	25.25	2.722
Weekly, h=5								
Mean	2.646	2.692	1.555	1.566	0.399	0.405	1.217	0.449
Variance	3.292	3.690	0.228	0.240	0.084	0.083	0.957	0.022
Skewness	2.628	2.769	1.252	1.410	0.215	0.382	2.284	-0.176
Kurtosis	14.20	14.71	5.696	6.110	3.226	3.290	10.02	2.464
Bi-Weekly, h=10								
Mean	5.297	5.386	2.216	2.233	0.759	0.767	2.437	0.453
Variance	10.44	11.74	0.389	0.403	0.072	0.070	2.939	0.019
Skewness	1.968	2.462	1.063	1.291	0.232	0.380	1.904	-0.147
Kurtosis	7.939	11.98	4.500	5.602	3.032	3.225	7.849	2.243
Tri-Weekly, h=15								
Mean	7.937	8.075	2.717	2.744	0.964	0.977	3.651	0.455
Variance	22.33	22.77	0.560	0.546	0.069	0.064	5.857	0.018
Skewness	2.046	2.043	1.033	1.177	0.208	0.400	1.633	-0.132
Kurtosis	9.408	8.322	4.621	4.756	2.999	3.123	6.139	2.247
Monthly, h=20								
Mean	10.59	10.77	3.151	3.179	1.116	1.127	4.874	0.458
Variance	34.09	36.00	0.671	0.671	0.062	0.059	8.975	0.017
Skewness	1.561	1.750	0.906	1.078	0.295	0.452	1.369	-0.196
Kurtosis	5.768	6.528	3.632	4.069	2.686	2.898	4.436	2.196

Table 2
Correlation Matrices of Realized DM/$ and Yen/$ Volatilities

	$vary_{t,h}$	$stdd_{t,h}$	$stdy_{t,h}$	$lstdd_{t,h}$	$lstdy_{t,h}$	$cov_{t,h}$	$corr_{t,h}$
Daily, h=1							
$vard_t$	0.539	0.961	0.552	0.860	0.512	0.806	0.341
$vary_t$	1.000	0.546	0.945	0.514	0.825	0.757	0.234
$stdd_t$	-	1.000	0.592	0.965	0.578	0.793	0.383
$stdy_t$	-	-	1.000	0.589	0.959	0.760	0.281
$lstdd_t$	-	-	-	1.000	0.604	0.720	0.389
$lstdy_t$	-	-	-	-	1.000	0.684	0.294
cov_t	-	-	-	-	-	1.000	0.590
Weekly, h=5							
$vard_{t,h}$	0.494	0.975	0.507	0.907	0.495	0.787	0.311
$vary_{t,h}$	1.000	0.519	0.975	0.514	0.908	0.761	0.197
$stdd_{t,h}$	-	1.000	0.545	0.977	0.545	0.789	0.334
$stdy_{t,h}$	-	-	1.000	0.555	0.977	0.757	0.220
$lstdd_{t,h}$	-	-	-	1.000	0.571	0.748	0.336
$lstdy_{t,h}$	-	-	-	-	1.000	0.718	0.235
$cov_{t,h}$	-	-	-	-	-	1.000	0.617
Bi-weekly, h=10							
$vard_{t,h}$	0.500	0.983	0.503	0.931	0.490	0.776	0.274
$vary_{t,h}$	1.000	0.516	0.980	0.514	0.923	0.772	0.170
$stdd_{t,h}$	-	1.000	0.533	0.982	0.531	0.780	0.293
$stdy_{t,h}$	-	-	1.000	0.544	0.981	0.762	0.188
$lstdd_{t,h}$	-	-	-	1.000	0.556	0.753	0.300
$lstdy_{t,h}$		-	-	-	1.000	0.726	0.202
$cov_{t,h}$	-	-	-	-	-	1.000	0.609
Tri-weekly, h=15							
$vard_{t,h}$	0.498	0.982	0.505	0.931	0.497	0.775	0.255
$vary_{t,h}$	1.000	0.522	0.984	0.525	0.939	0.763	0.146
$stdd_{t,h}$	-	1.000	0.538	0.983	0.539	0.787	0.277
$stdy_{t,h}$	-	-	1.000	0.551	0.984	0.756	0.155
$lstdd_{t,h}$	-	-	-	1.000	0.564	0.765	0.285
$lstdy_{t,h}$	-	-	-	-	1.000	0.727	0.162
$cov_{t,h}$	-	-	-	-	-	1.000	0.605
Monthly, h=20							
$vard_{t,h}$	0.479	0.988	0.484	0.952	0.479	0.764	0.227
$vary_{t,h}$	1.000	0.501	0.988	0.509	0.953	0.747	0.109
$stdd_{t,h}$	-	1.000	0.512	0.988	0.511	0.775	0.241
$stdy_{t,h}$	-	-	1.000	0.527	0.988	0.741	0.112
$lstdd_{t,h}$	-	-	-	1.000	0.533	0.763	0.245
$lstdy_{t,h}$		-	-	-	1.000	0.719	0.115
$cov_{t,h}$	-	-	-	-	-	1.000	0.596

Table 3
Dynamic Dependency Measures for Realized DM/$ and Yen/$ Volatilities

		$vard_{t,h}$	$vary_{t,h}$	$stdd_{t,h}$	$stdy_{t,h}$	$lstdd_{t,h}$	$lstdy_{t,h}$	$cov_{t,h}$	$corr_{t,h}$
Daily, h=1									
	LB	4539.3	3257.2	7213.7	5664.7	9220.7	6814.1	2855.2	12197
	$\hat{d}$	0.356	0.339	0.381	0.428	0.420	0.455	0.334	0.413
Weekly, h=5									
	LB	592.7	493.9	786.2	609.6	930.0	636.3	426.1	2743.3
	$\hat{d}$	0.457	0.429	0.446	0.473	0.485	0.496	0.368	0.519
Bi-weekly, h=10									
	LB	221.2	181.0	267.9	206.7	305.3	203.8	155.4	1155.6
	$\hat{d}$	0.511	0.490	0.470	0.501	0.515	0.507	0.436	0.494
Tri-weekly, h=15									
	LB	100.7	108.0	122.6	117.3	138.3	112.5	101.6	647.0
	$\hat{d}$	0.400	0.426	0.384	0.440	0.421	0.440	0.319	0.600
Monthly, h=20									
	LB	71.8	69.9	83.1	70.9	94.5	66.0	78.5	427.3
	$\hat{d}$	0.455	0.488	0.440	0.509	0.496	0.479	0.439	0.630

Figure 1
Distributions of Daily Realized Exchange Rate Volatilities and Correlations

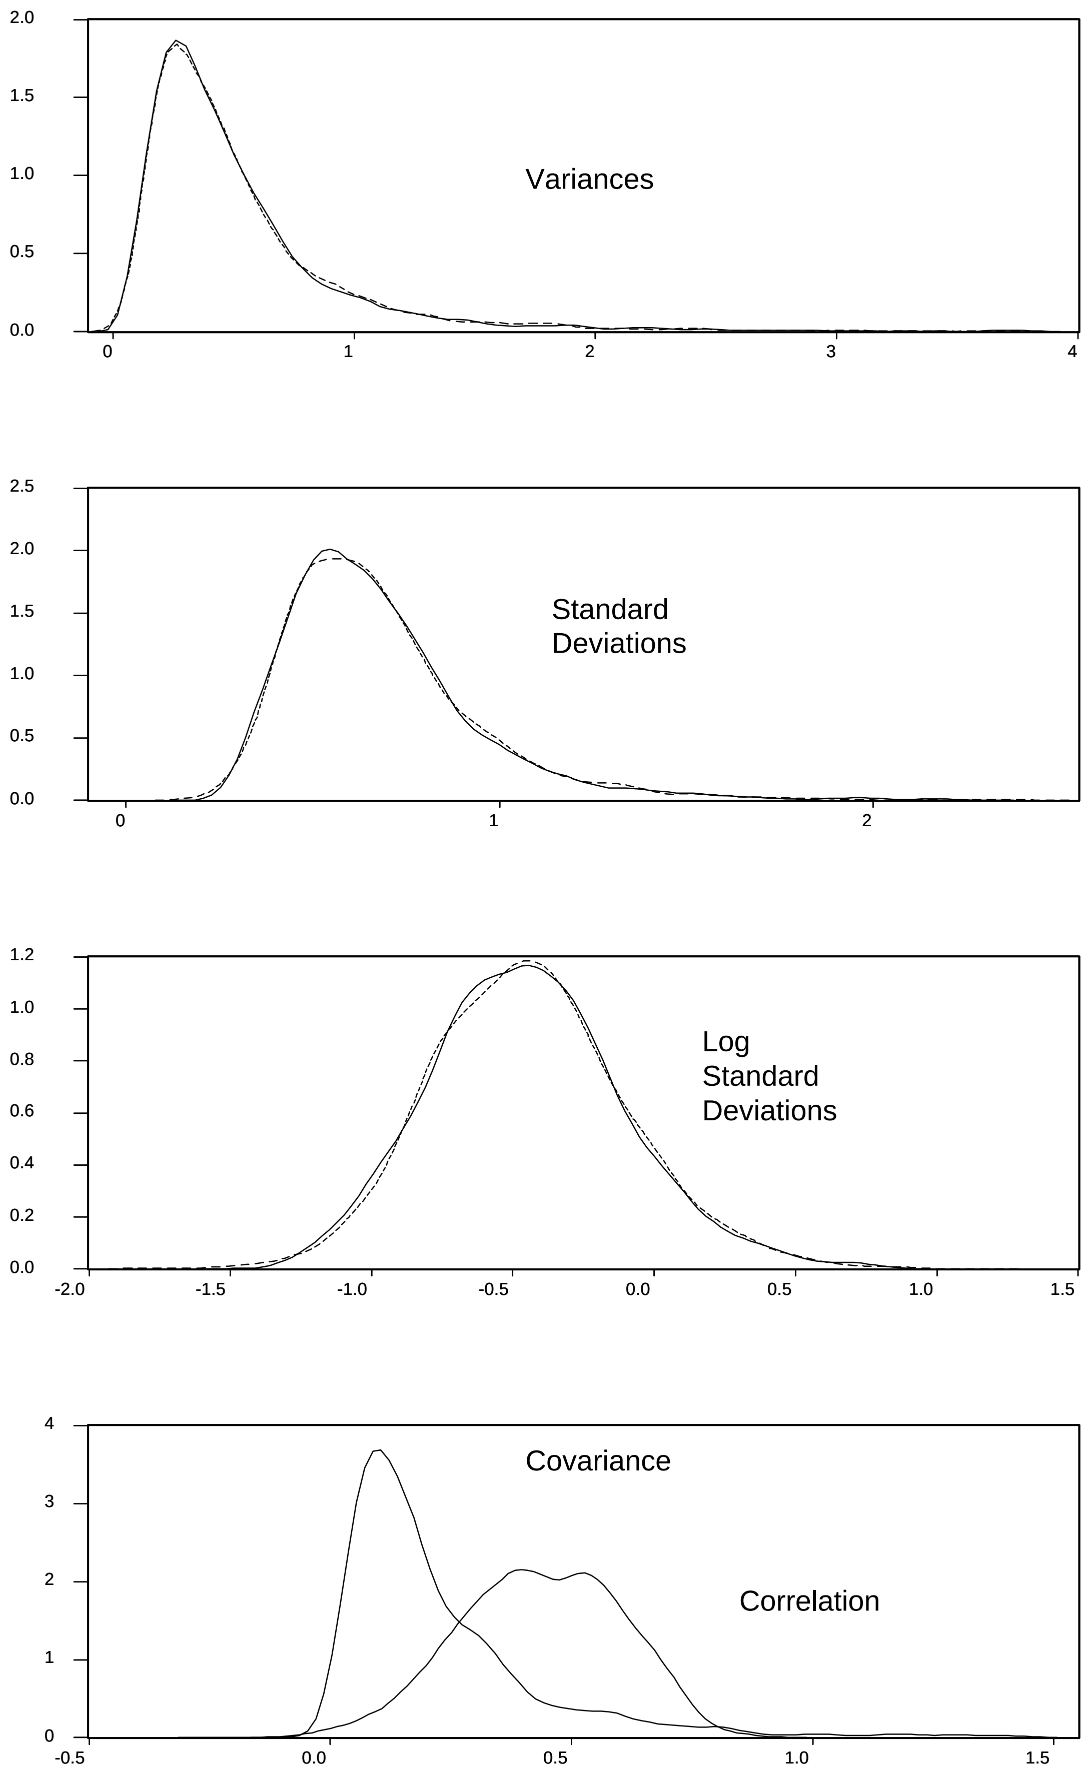

Figure 2
Bivariate Scatterplots of Realized Volatilities and Correlations

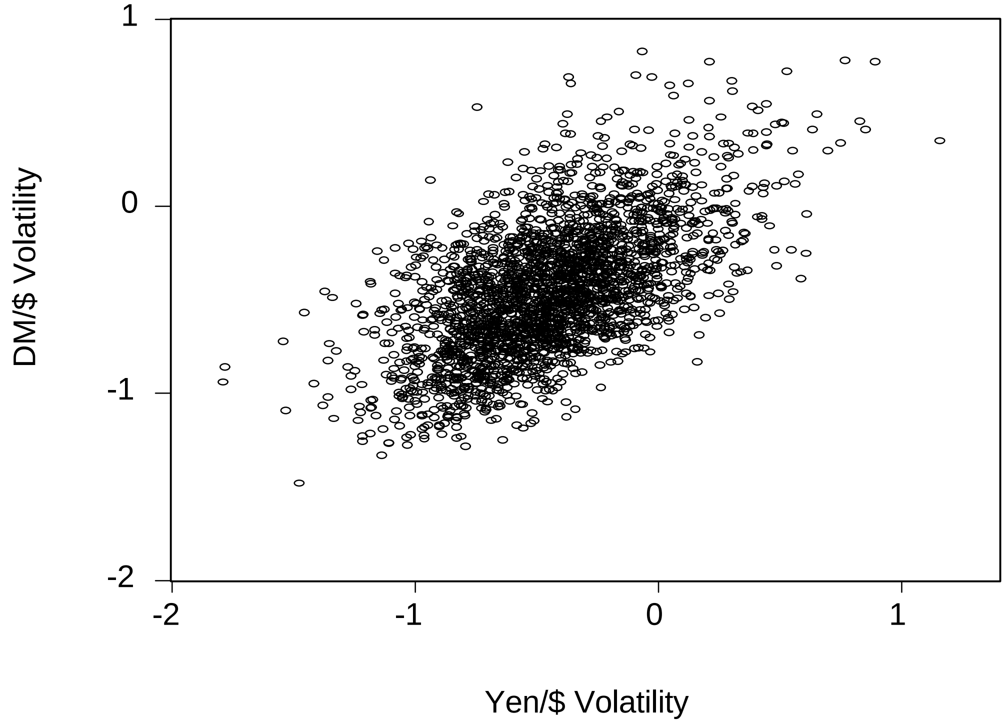

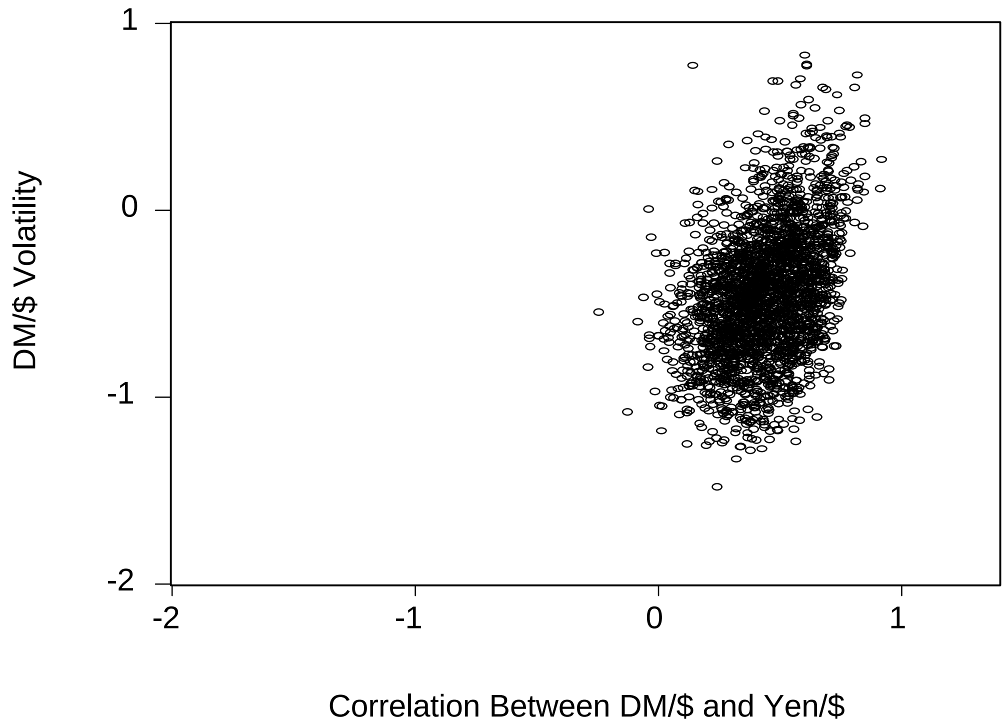

Figure 3
Distributions of Daily Realized Exchange Rate Correlation:
Low Volatility vs. High Volatility Days

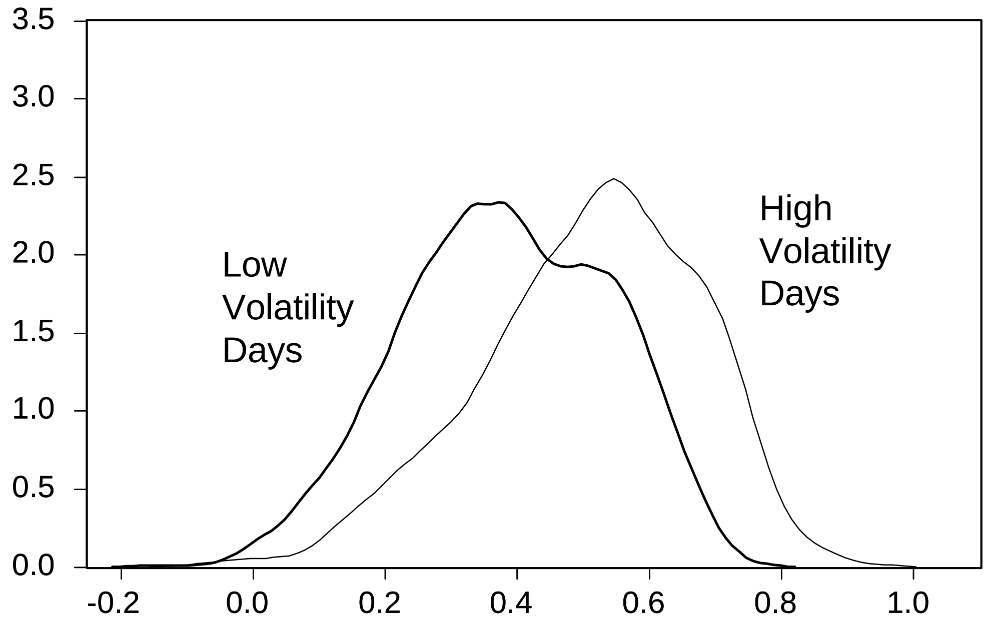

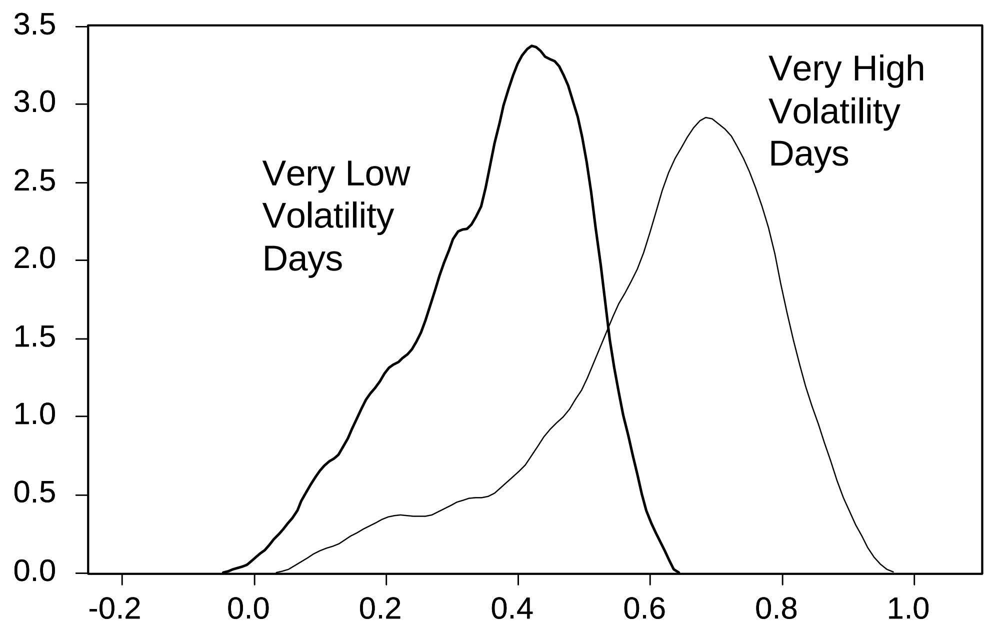

Figure 4
Time Series of Daily Realized Volatilities and Correlation

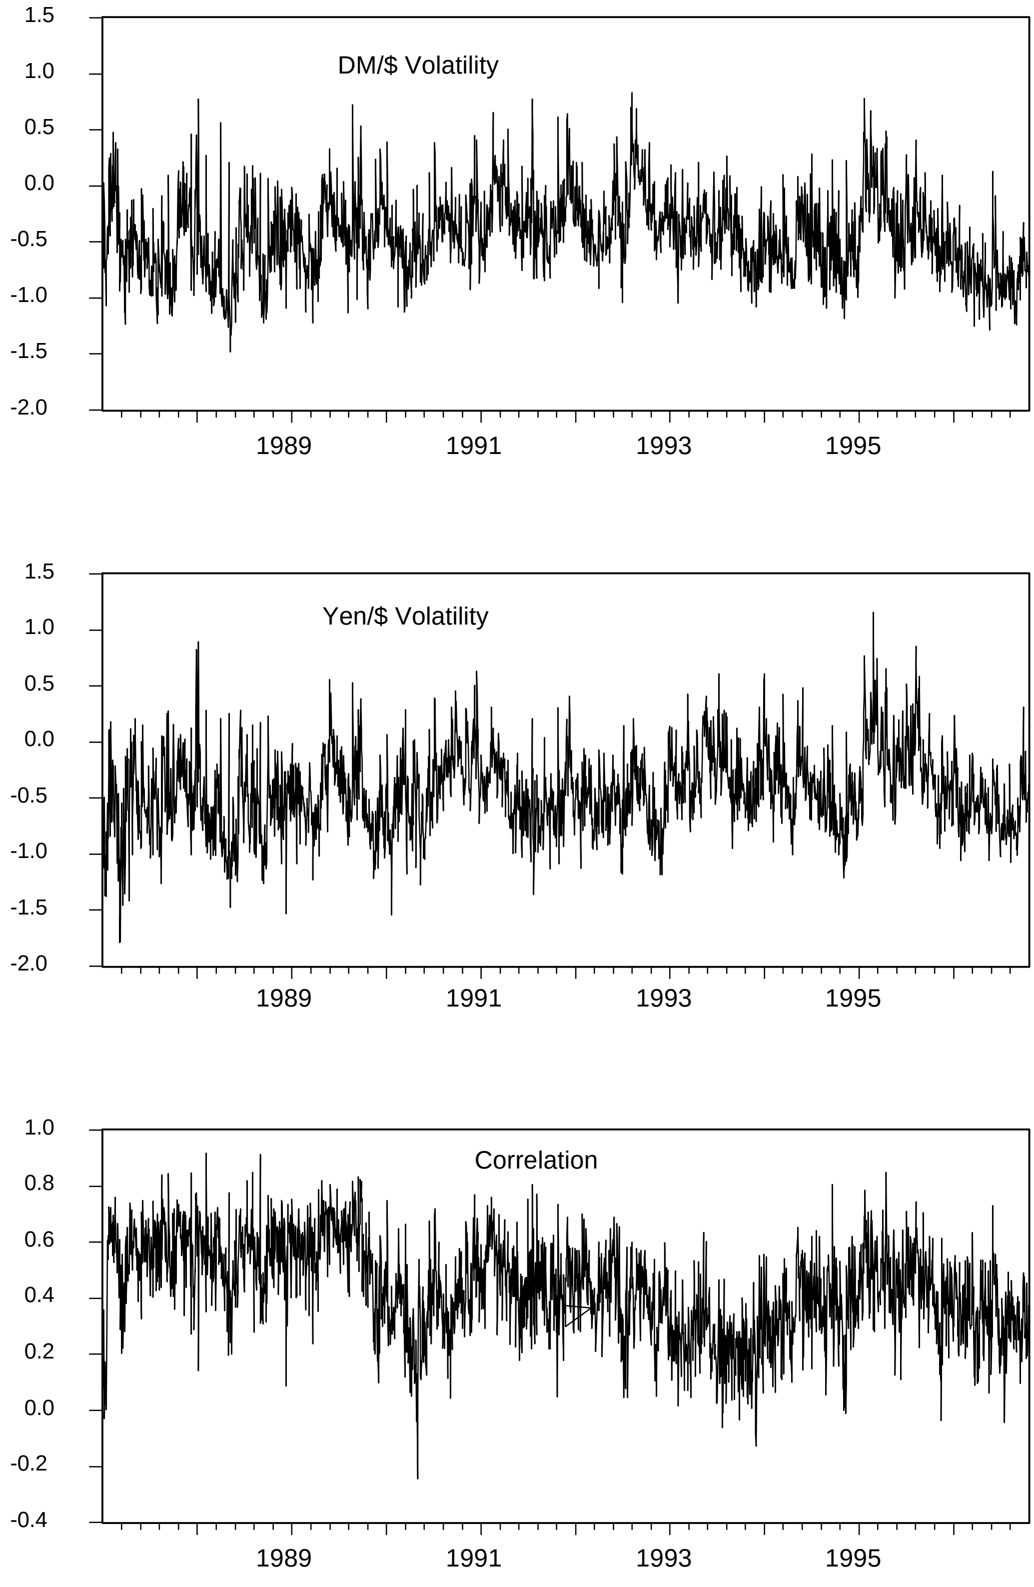

Figure 5
Sample Autocorrelations of Realized Volatilities and Correlation

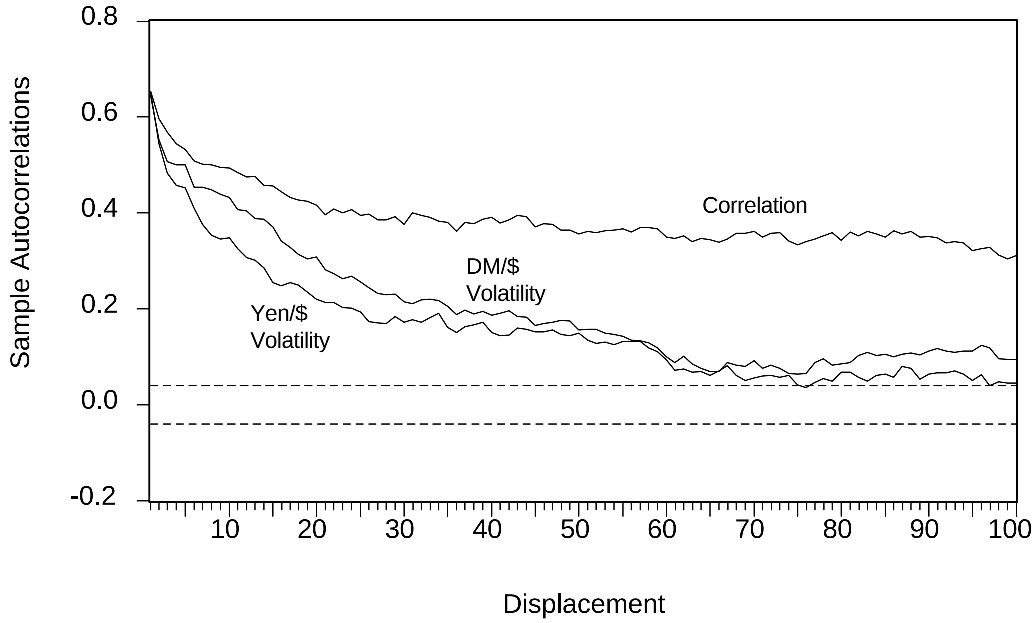

Figure 6
Scaling Laws Under Temporal Aggregation