

Wharton

**Financial
Institutions
Center**

*The Distribution of Exchange Rate
Volatility*

**by
Torben G. Andersen
Tim Bollerslev
Francis X. Diebold
Paul Labys**

99-08

THE WHARTON FINANCIAL INSTITUTIONS CENTER

The Wharton Financial Institutions Center provides a multi-disciplinary research approach to the problems and opportunities facing the financial services industry in its search for competitive excellence. The Center's research focuses on the issues related to managing risk at the firm level as well as ways to improve productivity and performance.

The Center fosters the development of a community of faculty, visiting scholars and Ph.D. candidates whose research interests complement and support the mission of the Center. The Center works closely with industry executives and practitioners to ensure that its research is informed by the operating realities and competitive demands facing industry participants as they pursue competitive excellence.

Copies of the working papers summarized here are available from the Center. If you would like to learn more about the Center or become a member of our research community, please let us know of your interest.

Anthony M. Santomero
Director

*The Working Paper Series is made possible by a generous
grant from the Alfred P. Sloan Foundation*

The Distribution of Exchange Rate Volatility ¹

January 27, 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.

¹This 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 the participants at the 1999 North American Winter Meetings of the Econometric Society, New York, as well as Dave Backus, Michael Brandt, Rohit Deo, Joel Hasbrouck, Cliff Hurvich, Richard Levich, Bill Schwert, Rob Stambaugh, and George Tauchen.

Torben G. Andersen is at the Department of Finance, Kellogg Graduate School of Management, Northwestern University, phone: 847-467-1285, e-mail: t-andersen@nwu.edu

Tim Bollerslev is at the Department of Economics, Duke University, and NBER, phone: 919-660-1846, e-mail: boller@econ.duke.edu

Francis X. Diebold is at the Departments of Economics and Statistics, University of Pennsylvania, and Departments of Finance, Economics and Statistics, Stern School of Business, New York University, and NBER, phone: 215-898-1507, e-mail: fdiebold@mail.sas.upenn.edu

Paul Labys is at the 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

Copyright © 1998, 1999 T.G. Andersen, T. Bollerslev, F.X. Diebold, and P. Labys.

1. Introduction

The seminal papers by Mandelbrot (1963) and Fama (1965) spurred a large, and still rapidly evolving, research program into the distributional characteristics of speculative returns. Consistent with the notion of efficient capital markets, the majority of the early work found that daily returns on most actively-traded instruments are approximately serially uncorrelated, a conclusion that is still hotly debated, but arguably a good first approximation to modern capital markets. Meanwhile, following the original contribution by Engle (1982), focus shifted toward return *volatility*, as opposed to returns themselves. It is now widely agreed that, although daily and monthly returns are approximately unpredictable, return volatility is highly predictable, a phenomenon with sweeping implications for financial economics and risk management.¹

Of course, asset return volatility is inherently unobservable. Most of what 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 econometric 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* of the models is strictly correct. Similarly, the well-known smiles and smirks in volatilities implied by Black-Scholes prices for options written at different strikes, as well as the non-constant term structure of implied volatility, 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 methods and models, we provide new and complementary measures of daily asset return volatilities. The mechanics of our methods are simple: we estimate daily volatility by summing high-frequency intraday squared returns. The resulting volatility estimates are valid under quite general conditions, which we delineate. Moreover, under those conditions, the volatility estimation error is in principle under our control and can be made arbitrarily small, by summing sufficiently finely sampled intraday returns. The resulting small amount

¹ The literature on the subject is huge; recent surveys include Bollerslev, Chou and Kroner (1992), Bollerslev, Engle and Nelson (1994), and Ghysels, Harvey and Renault (1996).

of measurement error in our daily volatility estimates means that for practical purposes we can treat the daily volatility as *observed*. We do so, and we proceed to examine directly the distributional characteristics of daily, weekly, and monthly volatility. The key insight is that because volatility is effectively observed, we can characterize both its marginal and conditional (dynamic) aspects using a variety of much simpler techniques than the complicated econometric models required when volatility is latent.

Our analysis is in the spirit of, and directly extends, the early contributions by French, Schwert and Stambaugh (1987) and Schwert (1989, 1990a, 1990b). We progress, however, in a number of important ways. First, we provide rigorous diffusion-theoretic underpinnings for the volatility measures. 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 five-minute returns on two major currencies. These high-frequency return series enable us to compute and examine the *daily* volatilities, which play a central role in the burgeoning volatility literature.² The daily frequency of observation is low enough such that the problems of intraday calendar effects and the coarseness of transactions time need not be dealt with explicitly, yet it is high enough such that the volatility dynamics are omnipresent. 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. The next section provides a formal justification for our volatility measures. In Section 3, we discuss the high-frequency Deutschemark - U.S. Dollar (DM/\$) and Yen - U.S. Dollar (Yen/\$) exchange rates that provide the basis for our empirical analysis, and we also detail the construction of our realized daily variances and covariances. In Section 4, we characterize the unconditional distributions of the daily volatilities, while Section 5 explores the conditional distributions, including long-memory features. In Section 6, we explore issues related to temporal aggregation and scaling in relation to long memory. Section 7 concludes with a summary of our main

² The literature has steadily progressed towards the use of higher-frequency data and better ex-post volatility measures. For instance, Officer (1973) uses monthly stock returns to compute a rolling 12-month standard deviation, whereas Merton (1980) employs monthly data and a rolling 12-month standard deviation over 1926-1978, and, importantly, also suggests the use of daily returns to estimate monthly standard deviations over a shorter sample from 1962-1978. Similarly, French, Schwert, and Stambaugh (1987) rely on daily stock returns to calculate monthly standard deviations over 1928-1984. Moreover, Schwert (1990b) uses 15-minute returns to estimate daily NYSE standard deviations from 1983-1989, although that analysis was not his primary focus. Finally, Hsieh (1991) and Fung and Hsieh (1991) analyze time series of daily standard deviations constructed from 15-minute futures returns. However, all these studies are strictly univariate and provide no formal justification for the approach.

empirical findings and brief remarks on directions for future research.

2. Volatility Measurement: Theory

Here we define and examine the notion of *integrated volatility*, and its estimation by *realized volatility*. In particular, we provide a rigorous diffusion-theoretic motivation for the use of realized volatility measures, at daily and lower frequencies (e.g., weekly, monthly), constructed using very high-frequency intraday returns. We discuss in detail the benchmark case of a univariate diffusion, after which we introduce jumps and multivariate aspects, all of which are accommodated by the theory.

Integrated and Realized Volatility for a Univariate Diffusion

Our empirical measures of volatility build upon the recent ideas and theoretical results for continuous-time stochastic processes in Andersen and Bollerslev (1998a) and Barndorff-Nielsen and Shephard (1998). In particular, consider the following representation for the continuous-time logarithmic price process p_t ,

$$dp_t = \sigma_t dW_t, \quad (1)$$

where $t \geq 0$, W_t denotes a standard Brownian motion, and σ_t is a strictly stationary process. Denote the corresponding discretely-sampled returns with m observations per period by

$$r_{(m),t} \equiv p_t - p_{t-1/m} = \int_0^{1/m} \sigma_{t-1/m+\tau} dW_{t-1/m+\tau}, \quad (2)$$

where $t = 1/m, 2/m, \dots$. In the empirical analysis we normalize the unit time interval, or $m = 1$, to represent one day. Note that, by definition, the expected returns are equal to zero for all return horizons, m . It is straightforward, albeit notationally more cumbersome, to allow for a drift or more general forms of conditional mean predictability. However, the assumption of no conditional mean dependence provides a very good first approximation for the high-frequency returns analyzed below.³

Assuming that σ_t and W_t are independent, it follows that the variance of the h -period returns, for $h > 0$, $r_{(1/h),t+h}$ conditional on the sample path $\{\sigma_{t+\tau}\}_{\tau=0}^h$ is

$$\sigma_{t,h}^2 \equiv \int_0^h \sigma_{t+\tau}^2 d\tau. \quad (3)$$

³ Moreover, Figlewski (1997) argues that even if a drift is operative, it may be preferable to set it to zero rather than to estimate it, because doing so will introduce only a small bias while achieving a large variance reduction.

This *integrated volatility* thus provides a natural definition of price variability, or volatility, in a continuous-time setting. Integrated volatility has previously been emphasized in the stochastic volatility option pricing literature by Hull and White (1987).

Of course, the integrated volatility is inherently unobservable. Gallant, Hsu and Tauchen (1998) and Chernov and Ghysels (1998) have recently pursued an indirect approach for estimating the distribution of $\sigma_{t,h}^2$, based on Gallant and Tauchen's (1998) re-projection method for simulation-based estimators. Their approach is elegant and intriguing, but it relies on number of auxiliary assumptions and can be challenging to implement. In the present paper we take a direct approach to measuring the daily integrated volatility by summing high-frequency intraday squared returns. The resulting *realized volatility* series in turn allows us to characterize both the unconditional and the conditional distribution of the volatility by standard statistical procedures.

More formally, suppose that the stochastic process for σ_t has cadlag sample paths, which in turn implies that the process for $\sigma_{t,h}^2$ has continuous sample paths.⁴ The quadratic variation of the returns defined in equation (2) then equals the integrated volatility over the relevant horizon as defined in equation (3); that is,⁵

$$plim_{m \rightarrow \infty} \sum_{j=1, \dots, mh} r_{(m),t+j/m}^2 = \sigma_{t,h}^2. \quad (4)$$

In words, the realized volatility is consistent (in m) for the integrated volatility; hence, by summing sufficiently many high-frequency discrete-time intraday returns one may approximate the integrated volatility arbitrarily well over any horizon.⁶ It is instructive to compare the realized volatility to the ex-post squared return over the relevant h -period horizon. Although the squared return is unbiased for the integrated volatility, so that $E_t(r_{(1/h),t+h}^2 - \sigma_{t,h}^2) = 0$, it is typically contaminated by substantial measurement error, which dwarfs the variation in the actual volatility process. For instance, as shown by Andersen and Bollerslev (1998a) for one-day volatility ($h = 1$) and empirically realistic parameter values, the variance of the measurement error in squared daily returns is easily twenty times the unconditional variance of $\sigma_{t,1}^2$. This contrasts sharply with the realized volatility defined by equation (4), which becomes *free* of measurement error as $m \rightarrow \infty$.

⁴ Continu à droite avec des limites à gauche (continuous from the left with limits from the right).

⁵ See Karatzas and Shreve (1988) or Barndorff-Nielsen and Shephard (1998).

⁶ This result, moreover, does not require that σ_t be independent of W_t .

Jumps

A number of authors, including Jorion (1988), Bates (1996), Andersen, Benzoni and Lund (1998), and Drost, Nijman and Werker (1998), have argued for the importance of including both time-varying volatility and jumps when modeling speculative returns over short horizons. The returns diffusion (1) readily accommodates such extensions. In particular, consider the jump diffusion

$$dp_t = \sigma_t dW_t + \kappa_t dN_t, \quad (5)$$

where N_t denotes an independent compound jump process.⁷ In this case,

$$plim_{m \rightarrow \infty} \sum_{j=1, \dots, mh} r_{(m), t+j/m}^2 = \sigma_{t,h}^2 + \kappa_{t,h}^2, \quad (6)$$

where, letting $I(\Delta N_s)$ denote the indicator function for a jump at a time $t = s$, we have

$$\kappa_{t,h}^2 \equiv \sum_{0 \leq \tau \leq h} \kappa_{t+\tau}^2 I(\Delta N_{t+\tau}),$$

so that the summation of the high-frequency returns approximates an additional set of jump terms. However, the interpretation of the right-hand-side of equation (6) as h -period return volatility remains intact. For notational simplicity, we shall therefore continue to refer to $\sigma_{t,h}^2$ as the realized volatility, omitting any further discussion of jump components.

Multivariate

Many interesting questions pertaining to the analysis of risk are inherently multivariate in nature. Hence, in addition to the above-discussed measures of variation, we also propose and examine new measures of covariation. In order to motivate our corresponding realized covariance measures, it is instructive to consider a simple multivariate extension of the model in equation (1). Again following Barndorff-Nielsen and Shephard (1998), let the N -dimensional diffusion for p_t have a single-factor representation

$$dp_t = \beta \sigma_t dW_t + \Omega_t dV_t, \quad (7)$$

where β is the N -dimensional vector of loadings on the common volatility factor $\sigma_t dW_t$, V_t is an N -dimensional standard Brownian motion with mutually-independent elements, and the diagonal matrix Ω_t

⁷ In the compound Poisson process most commonly employed in the literature, the jumps are assumed to be i.i.d and standard normal or lognormal, and they occur according to a Poisson process.

contains the N individual square-integrable mutually-independent asset-specific stochastic volatilities. This continuous-time latent-factor formulation directly parallels the idea behind the discrete-time latent-factor volatility models proposed by Diebold and Nerlove (1989) and King, Sentana and Wadhvani (1994). The generalization to multiple common factors would be conceptually straightforward, but notationally more cumbersome. Other more complicated multivariate continuous-time models could also be considered. However, the simple structure in equation (7) conveys the basic intuition.

Specifically, we consider the N -dimensional vector of h -period returns $r_{(1/h),t+h} \equiv p_{t+h} - p_t = \int_0^h \beta \sigma_t dW_t + \int_0^h \Omega_t dV_t$. Further, we let $\Sigma_{t,h}$ denote the corresponding $N \times N$ covariance matrix conditional on the sample path filtration generated by the latent volatility processes, $\{\sigma_{t+\tau}\}_{\tau=0}^h$ and $\{\Omega_{t+\tau}\}_{\tau=0}^h$. In parallel to the arguments underlying equation (3), the element of $\Sigma_{t,h}$ corresponding to the covariance between the i th and j th elements of $r_{(1/h),t+h}$, say $\{r_{(1/h),t+h}\}_i$ and $\{r_{(1/h),t+h}\}_j$ is then

$$\{\Sigma_{t,h}\}_{i,j} \equiv \{\beta \beta'\}_{i,j} \int_0^h \sigma_{t+\tau}^2 d\tau. \quad (8)$$

We call this the *integrated covariance*. By the theory of quadratic covariation,

$$plim_{m \rightarrow \infty} \Sigma_{k=1, \dots, mh} \{r_{(m),t+k/m}\}_i \{r_{(m),t+k/m}\}_j = \{\Sigma_{t,h}\}_{i,j}, \quad (9)$$

so that the integrated covariance may be estimated to any desired degree of accuracy by the *realized covariance* constructed by summing return cross products sampled sufficiently finely.

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.⁸ They also represent the most actively traded and quoted foreign currencies, and hence permit the construction of extremely accurate volatility measures. We first rationalize the use of underlying five-minute returns to construct daily realized volatilities, and then detail our corrections for weekend and other holiday non-trading periods. Finally, we describe the actual construction of the realized volatility measures.

⁸ We use the standard terminology of the FX interbank spot market by measuring the prices and the corresponding rates of return as the prices of \$1 in terms of DM and Yen, i.e., DM/\$ and Yen/\$.

On the Use of Five-Minute Returns

In practice, the inherent discreteness of actual securities prices renders the continuous-time model in equation (1) a poor approximation for the very highest sampling frequencies. Furthermore, high-frequency, or tick-by-tick, prices are generally only available at unevenly-spaced discrete time points, so that the calculation of the evenly-spaced $1/m$ -period returns, $r_{(m),t} \equiv p_t - p_{t-1/m}$, must necessarily rely on some form of interpolation involving the recorded prices around the beginning and the 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; see, for example, Lo and MacKinlay (1990). Moreover, as recently illustrated by Andersen, Bollerslev and Das (1998), this same effect will also give rise to a bias in the estimated integrated volatility defined by equation (4).

The same market microstructure issues that may render the univariate model in equation (1) a poor approximation at the very highest sampling frequencies would obviously also invalidate the multivariate formulation in equation (7) at those same frequencies. In fact, the problems associated with unevenly-spaced price observations may be especially acute in the multivariate context, where different degrees of interpolation will typically have to be employed in the calculation of the two $1/m$ -period returns, $\{r_{(m),t+k/m}\}_i$ and $\{r_{(m),t+k/m}\}_j$. As recently shown by Lundin, Dacorogna and Müller (1998), with large values of m but relatively infrequent price observations, this non-synchronous trading or quotation effect can result in important downward biases in the standard sample correlations of two interpolated return series.⁹ A similar effect is likely to bias the realized covariance measure in equation (9) as $m \rightarrow \infty$.

The sampling frequency at which microstructure biases become a practical concern is largely an empirical question. For the very actively quoted and traded foreign exchange rates analyzed here, a sampling frequency of five-minutes, or $m = 288$, represents a reasonable compromise between the accuracy of the theoretical approximations in equations (4) and (9) on the one hand and the market microstructure and discreteness considerations on the other. That is, $m=288$ is high enough such that our daily realized volatilities are largely free of measurement error, see, e.g., the calculations in Andersen and Bollerslev (1998a), yet low enough such that microstructure biases are not a major concern.

⁹ The drop in high-frequency sample correlations as the length of the return interval approaches zero was first noted by Epps (1979). To circumvent this so-called Epps effect, Lundin, Dacorogna and Müller (1998) propose an alternative summary measure of short-run correlation based on a modified weighted average sample correlation in which the weights depend upon the observation frequency of the underlying price series.

Construction of Five-Minute DM/\$ and Yen/\$ Returns

The two raw five-minute DM/\$ and Yen/\$ return series were obtained from Olsen and Associates in Zürich, Switzerland. Our full sample consists of continuously-recorded 5-minute returns from December 1, 1986 through December 1, 1996, or 3,653 days, for a total of $3,653 \cdot 288 = 1,052,064$ high-frequency return observations. The actual 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 even second.¹⁰ After filtering the data for outliers and other anomalies, the price at each five-minute mark is obtained by linearly interpolating from the logarithmic average of the bid and the ask for the two closest ticks. The continuously-compounded returns are then simply the change in these logarithmic five-minute prices.¹¹ For a detailed account of the method of data capture and the outlier filters employed, see Müller et al. (1990) and Dacorogna et al. (1993).

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 five-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 systematic 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 (July 4!). We furthermore 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

¹⁰ During our ten year sample period, approximately 4,500 DM/\$ and 2,000 Yen/\$ quotes appeared on the Reuters screen on an average business day.

¹¹ Although the quoted bid-ask spreads in the interbank FX market are extremely narrow, see, e.g., Bollerslev and Melvin (1994), it would be preferable to compute the returns from actual transaction prices. Unfortunately, high-frequency FX transactions data are rarely available. However, in limited a study involving seven hours of firm quotes and prices from Reuters electronic broker system, D2000-2, Goodhart, Ito and Payne (1996) find that the time series characteristics of FX quotes closely match those of the actual transactions prices at the five-minute frequency.

of the holiday market slowdowns, they clearly capture the most important daily calendar effects.¹²

As a final adjustment, we also 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 five-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.¹³ All in all, we were left with 2,445 complete days, or $2,445 \cdot 288 = 704,160$ five-minute return observations, for the construction of our daily variance and covariance measures.

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

In order to define formally our daily volatility measures, let the two time series of five-minute DM/\$ and Yen/\$ returns be denoted by $\Delta \log D_{(288),t}$ and $\Delta \log Y_{(288),t}$, respectively, where $t = 1/288, 2/288, \dots, 2,445$. We then proceed by forming the corresponding five-minute squared return and cross-product series $(\Delta \log D_{(288),t})^2$, $(\Delta \log Y_{(288),t})^2$, and $\Delta \log Y_{(288),t} \cdot \Delta \log D_{(288),t}$, respectively. 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 five-minute DM/\$ returns. Interestingly, the basic properties of the five-minute cross-product series, $\Delta \log Y_{(288),t} \cdot \Delta \log D_{(288),t}$, are remarkably similar. In particular, all three series are highly persistent and display strong intraday calendar effects. The shape of the intraday volatility patterns are readily associated with the opening and closing of the different financial markets around the world during the 24-hour trading cycle. Given our main focus on measuring and analyzing daily and longer volatilities, we shall not dwell on these high-frequency characteristics any further here.

We construct our estimates of the daily variances and covariances by summing the 288 intraday observations within each day. Concretely, from the $288 \cdot T$ five-minute returns, we construct the daily realized variances and covariances

¹² The 5-minute returns also display distinct short-lived “announcement spikes” associated with the release of regularly scheduled macroeconomic news, see e.g., Andersen and Bollerslev (1998b). Because most of the important U.S. announcements occur around 1:00 GMT on Thursdays and Fridays, they may induce a day-of-the-week effect. In order to purge our daily volatility measures of this effect, we also experimented with the removal of all “announcement time” returns from 12:30 GMT - 1:45 GMT. Doing so did not materially alter any of our findings. Thus, although eliminating the announcement period returns may be important when actually modeling the temporal dependencies in high-frequency returns, because the corresponding price jumps may otherwise easily be confused with longer-lived volatility bursts, it appears unnecessary for characterizing the daily and lower-frequency returns that concern us here.

¹³ The removal of fewer or more days did not materially affect the results reported below.

$$vard_t \equiv \sum_{j=1, \dots, 288} (\Delta \log D_{(288), t-1+j/288})^2 \quad (10)$$

$$vary_t \equiv \sum_{j=1, \dots, 288} (\Delta \log Y_{(288), t-1+j/288})^2 \quad (11)$$

$$cov_t \equiv \sum_{j=1, \dots, 288} \Delta \log D_{(288), t-1+j/288} \cdot \Delta \log Y_{(288), t-1+j/288}, \quad (12)$$

where $t = 1, 2, \dots, T$; here $T = 2445$. Our focus on the squared returns as a volatility measure, as opposed to say the absolute returns, is motivated by the diffusion theoretic foundations 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 popular alternative measures of variation and covariation derived from the realized variances and covariances in equations (10), (11) and (12), including the standard deviations, $std_d_t \equiv vard_t^{1/2}$ and $std_y_t \equiv vary_t^{1/2}$, the logarithmic standard deviations, $lstd_d_t \equiv 1/2 \cdot \log(vard_t)$ and $lstd_y_t \equiv 1/2 \cdot \log(vary_t)$, and the correlation coefficient, $corr_t \equiv cov_t / (std_d_t \cdot std_y_t)$. 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 corresponding volatility measures for 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

$$vard_{t,h} \equiv \sum_{j=1, \dots, 288 \cdot h} (\Delta \log D_{(288), h \cdot (t-1) + j/288})^2 \quad (13)$$

$$vary_{t,h} \equiv \sum_{j=1, \dots, 288 \cdot h} (\Delta \log Y_{(288), h \cdot (t-1) + j/288})^2 \quad (14)$$

$$cov_{t,h} \equiv \sum_{j=1, \dots, 288 \cdot h} \Delta \log D_{(288), h \cdot (t-1) + j/288} \cdot \Delta \log Y_{(288), h \cdot (t-1) + j/288}, \quad (15)$$

where $t = 1, 2, \dots, [T/h]$.¹⁴ We obtain the corresponding h -day standard deviations, $std_{t,h}$ and $std_{y,t,h}$, logarithmic standard deviations, $lstd_{t,h}$ and $lstd_{y,t,h}$, and correlations, $corr_{t,h}$ by appropriately transforming $vard_{t,h}$, $vary_{t,h}$ and $cov_{t,h}$. Section 6 details our analysis of these temporally aggregated volatilities.

¹⁴ Of course, by definition $vard_{t,1} \equiv vard_t$, $vary_{t,1} \equiv vary_t$, and $cov_{t,1} \equiv cov_t$.

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.

Univariate Unconditional Distributions

A number of different volatility measures merit examination. We begin by analyzing the univariate unconditional distributions of the daily realized variance series, $vard_t$ and $vary_t$. We show a standard menu of statistics summarizing the unconditional distributions in the first two columns of Table 1 (mean, variance, skewness, and kurtosis, along with a number of representative fractiles), and we graph the estimated unconditional distributions in the top panel of Figure 1.¹⁵ 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 implied by standard central limit theorems a poor approximation.¹⁶

The standard deviation of returns is measured on the same scale as the returns, and as such provides a more readily interpretable measure of volatility than does the variance. We present summary statistics and density estimates for the two daily realized standard deviations, std_d_t and std_y_t , in columns three and four of Table 1 and in the middle 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 70 basis points.

Interestingly, the distributions of the two daily realized logarithmic standard deviations, $lstd_d_t$ and $lstd_y_t$, displayed in columns five and six of Table 1 and in the bottom panel of Figure 1, appear symmetric. Moreover, although one, under certain assumptions, can reject normality at conventional significance levels, it is obviously a much better approximation for the logarithmic standard deviations

¹⁵ All of the density estimates reported in the paper are based on an Epanechnikov kernel and Silverman's (1986) bandwidth selection.

¹⁶ For the same reason, the conventional variance ratio statistic with p-values based on an F-distribution, or its asymptotic normal approximation, commonly employed in testing equality of variances over different time periods, is likely to work equally poorly in the high-frequency data setting. The recent evidence in Andersen, Bollerslev and Das (1998) and Bollerslev and Wright (1998b) underscores this point.

than for the realized variances or standard deviations.¹⁷ This accords with the findings for aggregate equity index volatility in French, Schwert and Stambaugh (1987), who note that monthly realized logarithmic standard deviations appear normally distributed.¹⁸

Finally, we characterize the distribution of the daily realized covariances and correlations, cov_t and $corr_t$. 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.¹⁹ Remarkably, though, the distribution of the realized correlation is close to normal. Even so, the sample standard deviation of 0.167 indicates significant variation of the correlation around its mean of 0.435. Of course, the positive correlation is not surprising, as it may arise from a common dependency on U.S. macroeconomic fundamentals, but the strength and temporal variation of the effect is nonetheless crucial for short-term portfolio allocation and hedging decisions.

Multivariate Unconditional Distributions

The univariate distributions characterized above do not address any 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 realized volatilities and correlations.

The sample correlations in Table 2, along with the $lstdd_t$ - $lstdy_t$ scatterplot in Figure 3, clearly 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 also their volatilities are closely linked. Explanations in terms of volatility spillovers have been explored in a series of papers initiated by the work of Engle, Ito and Lin (1990) and Ito, Engle and Lin (1992). Of course, such a positive relation would be expected under the stylized multivariate continuous-time

¹⁷ Under the null hypothesis of iid normality, the sample skewness and kurtosis are both asymptotically normally distributed with means of 0 and 3, respectively, and standard errors equal to $(6/T)^{1/2}$ and $(24/T)^{1/2}$, or 0.050 and 0.099. However, the iid assumption is, of course, questionable.

¹⁸ In a related theoretical development consistent with our empirical findings of skewed variance but Gaussian log variance, Barndorff-Nielsen and Shephard (1998) argue that continuous-time models specified in terms of the instantaneous latent logarithmic volatility, $d\log(\sigma_t^2)$, may realistically be assumed to have Gaussian innovations, or Brownian motion driving processes, whereas empirically realistic volatility models specified in terms of $d\sigma_t^2$ necessitate the use of more general Lévy driving processes.

¹⁹ Of course, by the absence of triangular arbitrage, the covariance between the DM/\$ and Yen/\$ rates is simply equal to one-half times the sum of the two variances for the Yen/\$ and DM/\$ rates minus the variance for the Yen/DM rate.

volatility model in equation (7). More generally, however, the positive volatility correlation provides the empirical justification for the use of multivariate volatility models with factor structure, as suggested by Diebold and Nerlove (1989), Engle, Ng and Rothschild (1990), Bollerslev and Engle (1993), and King, Sentana and Wadhvani (1994).

The correlations in Table 2 and the $corr_t-lstd_t$ and $corr_t-lstdy_t$ scatterplots in Figure 4 also indicate positive association between correlation and volatility. Whereas some nonlinearity may be operative in the $corr_t-lstd_t$ relationship, with a flattened response for both very low and very high $lstd_t$ values, the $corr_t-lstdy_t$ relationship appears approximately linear. To quantify further this “volatility effect” in correlation, we show in the top panel of Figure 5 kernel density estimates of $corr_t$ when both $lstd_t$ and $lstdy_t$ are less than -0.46 (their median value, which happens to be the same for each) and when both $lstd_t$ and $lstdy_t$ are greater than -0.46. Similarly, we show in the bottom panel of Figure 5 the estimated $corr_t$ densities conditional on the more extreme volatility situation in which both $lstd_t$ and $lstdy_t$ are less than -0.880 (approximately the tenth percentile of each distribution) and when both $lstd_t$ and $lstdy_t$ are greater than 0.0 (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.²⁰ A similar correlation effect in volatility has been documented recently for international equity returns by Solnik, Boucrelle and Le Fur (1996).²¹

Some authors, including Erb, Harvey and Viskanta (1994) and Longin and Solnik (1998), have argued that the correlations among international equity markets also tend to be higher when the returns are negative. This “return effect” in equity correlation may be related to the leverage or volatility feedback effect, by which negative returns result in larger volatility bursts than do positive returns of the same magnitude.²² This asymmetric response in turn produces a larger increase in the correlation for negative than positive returns via the aforementioned volatility effect in correlation. This phenomenon is not operative in the foreign exchange market, however, which is to be expected as the rationalizations for the asymmetric volatility effect applies to equities only. To underscore this, we show in the top and middle panels of Figure 6 the distributions of $lstd_t$ and $lstdy_t$ conditional on the

²⁰ In order to quantify more formally this volatility effect in correlation, we also experimented with various non-parametric regressions. For instance, in the simple quadratic approximation to an arbitrary functional form obtained by the linear regression of $corr_t$ on an intercept, the levels, the squares and the cross-product of $lstd_t$ and $lstdy_t$, all of the individual coefficients were highly statistically significant, and the R^2 was close to twenty percent.

²¹ Motivated by these findings, Bera and Kim (1998) have recently devised a formal test for constant conditional correlations in multivariate ARCH models.

²² See Nelson (1991), Campbell and Hentschel (1992), and Bekaert and Wu (1997).

signs of the corresponding two daily returns. The distributions on negative return days are almost indistinguishable from those on positive return days. Similarly, the distributions of $corr_t$ conditional on the signs of the two daily returns, which we show in the bottom panel of Figure 6, are also very close.

To summarize, the results in this section indicate a substantial amount of variation in volatilities and correlation, along with important contemporaneous dependence among the different measures. We now turn to a similar discussion of conditional dependence.

5. The Conditional Distribution of Daily Realized FX Volatility

The value of a derivative asset such as an option is closely linked to the expected volatility of the underlying over the time until expiration. Improved volatility forecasts should thus result in more accurate asset prices, and the conditional dependence in volatility forms the basis for such forecasts.²³ That dependence is most easily identified in the daily realized correlations and logarithmic standard deviations, which are approximately unconditionally normally distributed. Hence, in order to conserve space, we focus our discussion in the main text on those three series.

It is instructive first to consider the simple time series plots of the three realized volatilities in Figure 7. The wide fluctuations and strong persistence, that are evident in each of the univariate $lstd_t$ 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 an equally pronounced temporal dependence, 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 Table 3.²⁴ The correlograms of $lstd_t$, $lstdy_t$ and $corr_t$ in Figure 8 further underscore the point. The autocorrelation functions of the logarithmic standard deviations begin around 0.6 and decay very slowly to a value of about 0.1 at a displacement of 100 days. Those of the realized daily correlations dissipate even more slowly, reaching 0.35 at the 100 day displacement. Similar results based on long time series of daily absolute or squared returns from other markets have previously been documented by a number of authors including Ding, Granger and Engle (1993) and Granger and Marmol (1997). The slow

²³ Conversely, given a time series of option prices it is possible to infer a time series of volatilities implied by a particular option pricing model, see, e.g., Dumas, Fleming and Whaley (1998) for a recent discussion of S&P500 implied volatilities. More closely related to the present paper, Backus, Foresi, Li and Wu (1998) provide an intriguing characterization of the salient biases in the Black-Scholes pricing formula, along with an analysis of the corresponding implied volatilities for a set of foreign currency options.

²⁴ Equally significant short-run autocorrelations are evident in the implied daily FX volatilities analyzed by Backus, Foresi, Li and Wu (1998).

decay in Figure 8 is particularly noteworthy, however, in that the two realized daily volatility series “only” span ten years.

These findings of slow autocorrelation decay may seem to indicate the presence of a unit root, as in the integrated GARCH model proposed by Engle and Bollerslev (1986). However, the Dickey-Fuller tests with ten augmentation lags presented in the second row of Table 3 soundly reject the unit root hypothesis for all of the volatility series. This is also consistent with the findings of Wright (1999), who upon using more refined testing procedures reports strong rejections of the unit root hypothesis in daily log-squared FX and equity returns. Although unit roots may be formally rejected, the very slow autocorrelation decay coupled with the negative signs and slow decay of the estimated augmentation lag coefficients in the Dickey-Fuller tests suggest that long-memory of a non unit-root variety may be present. A number of authors have recently argued that this type of long-run dependence in financial market volatility may be conveniently modeled by a fractionally-integrated process, so that although volatility shocks are highly persistent, they eventually dissipate at a slow hyperbolic rate.²⁵ Hence, we now turn to an investigation of fractional integration in the daily realized volatility.

Fractionally integrated long-memory processes were introduced by Granger (1980, 1981) and Granger and Joyeux (1980); for a recent survey of their applications in economics see Baillie (1996). The slow hyperbolic decay of the long-lag sample autocorrelations and the log-linear explosion of the low-frequency spectrum are distinguishing features of a covariance stationary fractionally integrated, or $I(d)$, process with $0 < d < 1/2$.²⁶ The low-frequency spectral behavior also forms the basis for the log-periodogram regression estimation procedure proposed by Geweke and Porter-Hudak (1983) and later formalized by Robinson (1994a, 1995) and Hurvich, Deo and Brodsky (1998). In particular, let $I(\omega_j)$ denote the sample periodogram at the j th Fourier frequency, $\omega_j = 2\pi j/T$, $j = 1, 2, \dots, [T/2]$. The log-periodogram estimator of d is then based on the least squares regression,

$$\log[I(\omega_j)] = \beta_0 + \beta_1 \log(\omega_j) + u_j, \quad (16)$$

²⁵ See Baillie, Bollerslev and Mikkelsen (1996), Breidt, Crato and de Lima (1998), Granger and Ding (1996), Harvey (1998), and Robinson and Zaffaroni (1997, 1998) for various parametric formulations along those lines.

²⁶ Formal conditions for the equivalence between such low-frequency spectral behavior and long-lag autocorrelation behavior are given in Beran (1994) and Robinson (1994b).

where $j = 1, 2, \dots, m$, and $\hat{d} \equiv -1/2 \cdot \hat{\beta}_j$.²⁷ The least squares estimator of β_j , and hence \hat{d} , is asymptotically normal and the corresponding theoretical standard error, $\pi \cdot (24 \cdot m)^{-1/2}$, depends only on the number of periodogram ordinates used.²⁸

Of course, the actual value of the estimate of d also depends upon the particular choice of m . While the formula for the theoretical standard error suggests choosing a large value of m in order to obtain a small standard error, doing so may induce a bias in the estimator, because the relationship underlying equation (16) in general holds only for frequencies close to zero. Following the suggestion of Taquq and Teverovsky (1996), we therefore graph and examine \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”). The representative plots for $lstdd_t$ and $lstdy_t$ and $corr_t$, which we show in Figure 9, validate our subsequent use of $m = \lceil T^{4/5} \rceil$, or $m = 514$, as providing good estimates.²⁹

We present the estimates of d in the third row of Table 3. The estimates for all eight volatility series are highly statistically significant, and all are fairly close to the “typical value” of 0.4.³⁰ These estimates are directly in line with the estimates for d 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). Comparable results have also been obtained by Bollerslev and Wright (1998a) in their analysis of daily volatilities constructed from the aggregation of intraday returns. These 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. In that regard, our findings echo Fung and Hsieh (1991), who note that “... the stochastic process that best describes volatility of asset prices may well be more complex than the typical first-

²⁷ The calculations in Hurvich and Beltrao (1994) suggest that the estimator proposed by Robinson (1994a, 1995), which leaves out the very lowest frequencies in the regression in equation (16), has larger MSE than the original Geweke and Porter-Hudak (1983) estimator defined over all of the first m Fourier frequencies. For that reason, we include periodogram ordinates at all of the first m Fourier frequencies.

²⁸ While the earlier proofs for consistency and asymptotic normality of the log-periodogram regression estimator relied 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.

²⁹ This choice of m is consistent with the optimal rate of $O(T^{4/5})$ established by Hurvich, Deo and Brodsky (1998) and accords with the simulation results in Bollerslev and Wright (1998a).

³⁰ The commonality in the estimates for d across the different volatility series may be justified under auxiliary distributional assumptions following the arguments in Andersen and Bollerslev (1997b) and Bollerslev and Wright (1998a).

order mean-reverting models used in stochastic volatility models for option pricing ...”. 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.³¹

In contrast to the sample autocorrelations for the individual volatility series, the sample cross-correlations in Figure 10, $Corr(lstd_d_t, lstd_y_{t-j})$ and $Corr(lstd_{t-j}, lstd_y_t)$ for $j = 1, 2, \dots$, drop much more quickly and appear generally consistent with short-memory dependence. This suggests that a simple factor structure with a single common long-memory component would be too simplistic to fully describe the dynamic relationship between the two rates. It is also interesting to note that the cross-correlations at a displacement of one day are asymmetric, with the correlation between $lstd_{t-1}$ and $lstd_y_t$ being greater than the correlation between $lstd_d_t$ and $lstd_{y_{t-1}}$. We are not aware of any obvious economic reasons behind this asymmetric response, and we shall not pursue the issue further here.³² Instead, we turn to a characterization of the unconditional and conditional volatility distributions at horizons longer than one day.

6. The Effects of Temporal Aggregation

The analysis in the preceding sections focused exclusively on the distributional properties of daily realized volatilities and correlations. However, many practical problems in asset pricing, portfolio allocation, and financial risk management involve longer horizons.³³ Here, we provide a discussion of the distributional aspects of the corresponding multi-day realized volatilities and correlations. As before, we begin with an analysis of the unconditional distributional aspects, followed by an analysis of the dynamic dependence, including a detailed examination of long-memory as it relates to the temporal aggregation.

Univariate and Multivariate Unconditional Distributions

In Table 4, we summarize the univariate unconditional distributions for the eight different temporally aggregated volatility measures for weekly, bi-weekly, tri-weekly and monthly return horizons ($h = 5$,

³¹ Building on the earlier results of Cox (1991), Barndorff-Nielsen and Shephard (1998) have recently pointed out that, although OU volatility diffusions are short-memory, long-memory volatility diffusions can be constructed by appropriately superimposing an infinite number of OU processes. The continuous-time multifractal model of asset returns proposed by Mandelbrot, Fisher and Calvet (1997) also implies long-memory volatility dependence.

³² A similar asymmetry in the cross-correlations for the volatility of the same rate defined over different return horizons has recently been observed by Müller et al. (1997). They suggest that a multivariate extension of their “heterogeneous ARCH” volatility structure may be able to accommodate this asymmetry in the cross-correlations.

³³ See, for example, Diebold, Schuermann and Inoue (1998) and Christoffersen and Diebold (1998).

10, 15, and 20, respectively). Consistent with the notion of efficient capital markets and serially uncorrelated returns, the means of $vard_{t,h}$, $vary_{t,h}$, and $cov_{t,h}$ all grow at the constant rate h . In addition, the growth of the variance of the realized variances and covariance adheres closely to h^{2d+1} , where d denotes the order of integration of the series, a phenomenon that we discuss at length subsequently. It is also noteworthy that, even at the monthly level, the unconditional distributions of $vard_{t,h}$, $vary_{t,h}$, and $cov_{t,h}$ remain leptokurtic and highly skewed to the right. The basic characteristics of $stdd_{t,h}$ and $stdy_{t,h}$ are similar, with the sample mean increasing at a rate of $h^{1/2}$.³⁴ In contrast to the distributions of the other realized volatilities, 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. Finally, we note that the sample mean of the realized correlation, $corr_{t,h}$, is largely invariant to the level of temporal aggregation. Once again, this finding is consistent with fractional integration and will be explained shortly.

Next, turning to the multivariate unconditional distributions, we display in Table 5 the correlation matrices for all volatility measures for $h=5, 10, 15$, and 20 . While the correlation between the different measures of variation 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 0.606 at the daily horizon to a low of 0.531 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 0.391, whereas the correlation between $lstdd_{t,20}$ and $corr_{t,20}$ is 0.255. Similarly, the correlation between $lstdy_{t,h}$ and $corr_{t,h}$ drops from 0.295 for $h=1$ to 0.120 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.

The Conditional Distribution: Dynamic Dependence, Fractional Integration and Scaling

Our analysis of one-day volatilities reinforces earlier findings of strong volatility clustering. The Ljung-Box tests that we report in Table 6 indicate that the strong serial dependence is preserved under temporal aggregation. Even at the monthly level, or $h=20$, with “only” 122 observations, all of the test statistics are highly significant. This contrasts with previous evidence, which tends to show little or no

³⁴ The relatively slow decline of the sample kurtosis of the realized standard deviations is also consistent with the findings in Backus, Foresi, Li and Wu (1998) related to the distribution of implied FX volatilities.

significant evidence of volatility clustering by the time one aggregates to monthly returns; see, for example, Diebold (1988), Baillie and Bollerslev (1989), and Christoffersen and Diebold (1998). However, as previously noted, squared monthly returns are a very noisy proxy for the latent integrated volatility over the month. Using numerical techniques, Andersen, Bollerslev and Lange (1998) 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 Table 6 provide empirical confirmation.

The results in section 4 indicate that realized daily volatilities appear 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 of the process; see, for example, Beran (1994). This strong prediction is borne out by the estimates for d for the different levels of aggregation, which we report in Table 6. All of the estimates are within two asymptotic standard errors of the average estimate of 0.388 obtained for the daily series, and all are highly statistically significantly different from both zero and unity. These results are directly in line with the empirical evidence in Andersen and Bollerslev (1997b) and Bollerslev and Wright (1998a) pertaining to the estimates of d based on absolute returns at different levels of temporal aggregation (daily, weekly, etc.).

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

$$[x_t]_h \equiv \sum_{j=1, \dots, h} x_{h-(t-1)+j}, \quad (17)$$

denote the h -fold partial sum process for x_t , where $t = 1, 2, \dots, [T/h]$. It is then well known that, if the process for x_t is fractionally integrated, the partial sums obey a *scaling law* of the form

$$\text{Var}([x_t]_h) = c \cdot h^{2d+1}, \quad (18)$$

where d denotes the order of integration.³⁵ Thus, given d , if the unconditional variance is known for one aggregation level, it is possible to calculate the implied variance for any other aggregation level.³⁶

Of course, by definition $[var_d]_h \equiv vard_{t,h}$, $[vary_t]_h \equiv vary_{t,h}$, and $[cov_t]_h \equiv cov_{t,h}$, so that the

³⁵ See, for example, Beran (1994) and Diebold and Lindner (1996).

³⁶ The log-linear relationship in equation (18) also underlies the formal estimators for d recently analyzed by Giraitis, Robinson and Surgailis (1998), Taqqu, Teverovsky and Willinger (1995), and Teverovsky and Taqqu (1997).

variance of the *realized* variances and covariance should grow at the rate h^{2d+1} . This theoretical implication accords extremely well with the values for the unconditional sample variances and covariances reported in Tables 1 and 4 and a value of d around 0.35-0.40. Similar scaling laws for various power transforms of absolute FX returns have previously been emphasized in a series of papers by Müller et al. (1990), Guillaume et al. (1997), Fisher, Calvet and Mandelbrot (1997), and Mandelbrot, Fisher and Calvet (1997). If the d values for $vard_t$, $vary_t$, and cov_t are approximately the same, the scaling law also explains why the sample means for $corr_{t,h}$ reported in Table 4 should be invariant to the aggregation level, h .

The striking accuracy of scaling carries over to the partial sums of the other volatility series. To illustrate, we graph in Figure 11 the logarithm of the sample variances of the partial sums for the realized logarithmic standard deviations and correlation against the logarithm of the aggregation level; i.e., $\log(\text{Var}([lstd_d]_h))$, $\log(\text{Var}([lstd_y]_h))$, and $\log(\text{Var}([corr]_h))$, against $\log(h)$ for $h = 1, 2, \dots, 30$. The linear fits implied by equation (18) are obviously extremely good. Each of the three slopes accords very closely with the theoretical value of $2d+1$ implied by the log-periodogram estimates for d , further solidifying the notion of long-memory volatility dependence.³⁷

Because a non-linear function of a sum is not the sum of the non-linear function, it is not clear whether $lstd_{t,h}$ and $lstd_{y,t,h}$ will follow similar scaling laws. The estimates for d that we report in Table 5 suggest that they should. The corresponding plots for the logarithm of the logarithmic standard deviations $\log(\text{Var}(lstd_{t,h}))$ and $\log(\text{Var}(lstd_{y,t,h}))$ against $\log(h)$, for $h = 1, 2, \dots, 30$, which we show in Figure 12, lend empirical confirmation to this conjecture. Although the fits are not as impressive as those in Figure 11, the log-linear approximations are still very accurate. Interestingly, however, the slopes of the lines are negative.

To understand why these slopes are 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 $E(\text{var}_{t,h}) = b \cdot h$, where $\text{var}_{t,h}$ refers to the temporally aggregated variance as defined above. Also, by the scaling law in equation (18), $\text{Var}(\text{var}_{t,h}) = c \cdot h^{2d+1}$. Furthermore, assume that the corresponding temporally aggregated logarithmic standard deviations, $lstd_{t,h} \equiv 1/2 \cdot \log(\text{var}_{t,h})$, are normally distributed across all frequencies h with mean and variance denoted by μ_h and σ_h^2 , respectively. Of course, as discussed above, these three assumptions accord closely with the actual empirical distributions reported in Tables 1 and 4. It follows then from the properties of the lognormal

³⁷ The estimated slopes in the top, middle and bottom panels are 1.779, 1.726, and, 1.819, respectively, corresponding to d values of 0.389, 0.363, and, 0.409.

distribution that

$$E(\text{var}_{t,h}) = \exp(2\mu_h + 2\sigma_h^2) = b \cdot h$$

and

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

so that solving for the variance of the logarithmic standard deviation,

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

With $2d-1$ marginally less than zero, this explains why the sample variances of $\text{lstd}_{t,h}$ and $\text{lstdy}_{t,h}$ reported in Table 4 are decreasing with the level of temporal aggregation, h . Furthermore, by a log-linear approximation,

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

which provides a justification for the apparent scaling law behind the two plots in Figure 12, and the negative slopes approximately equal to $2d-1$.³⁸ Also, solving for the mean yields

$$E(\text{lstd}_{t,h}) \equiv \mu_h = \frac{1}{2} \cdot \log(h) + \frac{1}{2} \cdot \log(c) + \frac{1}{4} \cdot \log(c \cdot b^{-2} \cdot h^{2d-1} + 1).$$

To a first order approximation the sample means of $\text{lstd}_{t,h}$ and $\text{lstdy}_{t,h}$ should therefore increase at the rate $\frac{1}{2} \cdot \log(h)$, which again accords very well with the summary statistics reported in Tables 1 and 4.

7. Summary and Concluding Remarks

We have provided a formal diffusion-theoretic justification 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

³⁸ The slopes of the lines in the top and bottom panels are -0.227 and -0.275, respectively, and the corresponding d values of 0.387 and 0.363 are almost identical to the values implied by the scaling law in equation (18) underlying Figure 11.

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. There is, however, no evidence that the correlation is related to the sign of the returns. Finally, we confirm the wealth of existing evidence of strong volatility clustering effects in daily returns. However, in contrast to earlier work such as Christoffersen and Diebold (1998), which often indicates that volatility persistence decreases fairly quickly with the horizon, we find that even monthly realized volatilities remain highly persistent. But 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 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 realized volatility, and to use them for forecasting, risk management, asset allocation, and asset pricing. The results of this paper suggest modeling the joint conditional dependence in the realized logarithmic daily standard deviations by a standard linear Gaussian vector autoregressive (VAR) model, and that doing so may, for example, result in important improvements in the accuracy of long-term volatility forecasts. This idea is pursued in Andersen, Bollerslev, Diebold and Labys (1999).

References

- Andersen, T.G., L. Benzoni and J. Lund (1998), "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 (1997a), "Intraday Periodicity and Volatility Persistence in Financial Markets," *Journal of Empirical Finance*, 4, 115-158.
- Andersen, T.G. and T. Bollerslev (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 T. Bollerslev (1998a), "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885-905.
- Andersen, T.G. and T. Bollerslev (1998b), "DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer-Run Dependencies," *Journal of Finance*, 53, 219-265.
- Andersen, T.G., T. Bollerslev and A. Das (1998), "Testing for Market Microstructure Effects in Intraday Volatility: A Reassessment of the Tokyo FX Experiment," Manuscript, Department of Finance, J.L. Kellogg Graduate School of Management, Northwestern University, and Department of Economics, Duke University.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (1999), "Forecasting Realized Volatility," work in progress.
- Andersen, T.G., T. Bollerslev and S. Lange (1998), "Forecasting Financial Market Volatility: Sample Frequency vis-à-vis Forecast Horizon," Manuscript, Department of Finance, J.L. Kellogg Graduate School of Management, Northwestern University, and Departments of Economics, Duke University and University of Virginia.
- Backus, D., S. Foresi, K. Li, and L. Wu (1998), "Accounting for Biases in Black-Scholes," Manuscript, Stern School, New York University.
- Baillie, R.T. (1996), "Long-Memory Processes and Fractional Integration in Econometrics," *Journal of Econometrics*, 73, 5-59.
- 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., T. Bollerslev and H.O. Mikkelsen (1996), "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 74, 3-30.
- Barndorff-Nielsen, O.E. and N. Shephard (1998), "Aggregation and Model Construction for Volatility Models," Manuscript, Department of Mathematical Sciences, University of Aarhus, and Nuffield College, Oxford.

- Bates, D.S. (1996), "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options," *Review of Financial Studies*, 9, 69-107.
- Bekaert, G. and G. Wu (1997), "Asymmetric Volatility and Risk in Equity Markets," Manuscript, Departments of Finance, Stanford University and University of Michigan.
- Bera, A. and A. Kim (1998), "Testing Constancy of Correlation with an Application to International Equity Returns," Manuscript, Department of Economics, University of Illinois, Urbana-Champaign.
- Beran, J. (1994), *Statistics for Long-Memory Processes*. New York: Chapman and Hall.
- Bollerslev, T. and I. Domowitz (1993), "Trading Patterns and Prices in the Interbank Foreign Exchange Market," *Journal of Finance*, 48, 1421-1443.
- Bollerslev, T., R.Y. Chou and K.F. Kroner (1992), "ARCH Modeling in Finance: A Selective Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5-59.
- Bollerslev, T. and R.F. Engle (1993), "Common Persistence in Conditional Variances," *Econometrica*, 61, 166-187.
- 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 M. Melvin (1994), "Bid-Ask Spreads and Volatility in the Foreign Exchange Market: An Empirical Analysis," *Journal of International Economics*, 36, 355-372.
- Bollerslev, T. and J.H. Wright (1998a), "Estimating Long-Memory Volatility Dependencies: The Role of High Frequency Data," Manuscript, Departments of Economics, Duke University and University of Virginia.
- Bollerslev, T. and J.H. Wright (1998b), "Frequency Domain Inference for Speculative Return Volatility Dynamics," Manuscript, Departments of Economics, Duke University and University of Virginia.
- Breidt, F.J., N. Crato, and P. de Lima (1998), "On the Detection and Estimation of Long-Memory in Stochastic Volatility," *Journal of Econometrics*, 83, 325-348.
- Campbell, J.Y. and L. Hentschel (1992), "No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns," *Journal of Financial Economics*, 31, 281-318.
- Chernov, M. and E. Ghysels (1998), "What Data Should be Used to Price Options?," Manuscript, Department of Economics, The Pennsylvania State University.
- Christoffersen, P.F. and F.X. Diebold (1998), "How Relevant is Volatility Forecasting for Financial Risk Management?," Wharton Financial Institutions Center Working Paper 97-45 and NBER Working Paper 6844.
- Cox, D.R. (1991), "Long-Range Dependence, Non-Linearity and Time Irreversibility," *Journal of Time Series Analysis*, 12, 329-335.

- 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.
- Deo, R. and C.H. Hurvich (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. (1988), *Empirical Modeling of Exchange Rate Dynamics*. New York: Springer Verlag.
- Diebold, F.X. and P. Lindner (1996), "Fractional Integration and Interval Prediction," *Economics Letters*, 50, 305-313.
- 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. Schuermann and A. Inoue (1998), "Scale Models," *Risk*, 11, 104-107.
- 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.
- Drost, F.C., T.E. Nijman and B.J.M. Werker (1998), "Estimation and Testing in Models Containing Both Jumps and Conditional Heteroskedasticity," *Journal of Business and Economic Statistics*, 16, 237-243.
- Dumas, B., J. Fleming and R.E. Whaley (1998), "Implied Volatility Functions: Empirical Tests," *Journal of Finance*, 53, 2059-2106.
- Engle, R.F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation," *Econometrica*, 50, 987-1008.
- Engle, R.F. and T. Bollerslev (1986), "Modeling the Persistence of Conditional Variances," *Econometric Reviews*, 5, 1-50.
- Engle, R.F., T. Ito and W.L. Lin (1990), "Meteor Showers or Heat Waves? Heteroskedastic Intraday Volatility in the Foreign Exchange Market," *Econometrica*, 58, 525-542.
- 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.
- Epps, T.W. (1979), "Comovements in Stock Prices in the Very Short Run," *Journal of the American Statistical Association*, 74, 291-298.
- Erb, C.B., C.R. Harvey and T.E. Viskanta (1994), "Forecasting International Equity Correlations," *Financial Analysts Journal*, November/December, 32-45.
- Fama, E.F. (1965), "The Behavior of Stock Market Prices," *Journal of Business*, 38, 34-105.

- Figlewski, S. (1997), "Forecasting Volatility," *Financial Markets, Institutions and Instruments* (New York University Salomon Center), 6, 1-88.
- Fisher, A., L. Calvet and B. Mandelbrot (1997), "Multifractality of Deutschemark/U.S. Dollar Exchange Rates," Manuscript, Cowles Foundation, Yale University.
- French, K.R., G.W. Schwert and R.F. Stambaugh (1987), "Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3-29.
- Fung, W.K.H. and D.A. Hsieh (1991), "Empirical Analysis of Implied Volatility: Stocks, Bonds and Currencies," Manuscript, Fuqua School of Business, Duke University.
- Gallant, A.R. and Tauchen, G. (1998), "Reprojecting Partially Observed Systems with Application to Interest Rate Diffusions," *Journal of the American Statistical Association*, forthcoming.
- Gallant, A.R., C.-T. Hsu and G.E. Tauchen (1998), "Calculating Volatility Diffusions and Extracting Integrated Volatility," Manuscript, Departments of Economics, University of North Carolina, University of Vienna, and Duke University.
- 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.
- Giraitis, L., P.M. Robinson and D. Surgailis (1998), "Variance-Type Estimation of Long Memory," Manuscript, London School of Economics.
- Goodhart, C., T. Ito and R. Payne (1996), "One Day in June 1993: A Study of the Working of the Reuters 2000-2 Electronic Foreign Exchange Trading System," in J.A. Frankel, G. Galli and A. Giovannini (eds.), *The Microstructure of Foreign Exchange Markets*, 107-179. Chicago: University of Chicago Press.
- Granger, C.W.J. (1980), "Long Memory Relationships and the Aggregation of Dynamic Models," *Journal of Econometrics*, 14, 227-238.
- Granger, C.W.J. (1981), "Some Properties of Time Series Data and Their Use in Econometric Model Specification," *Journal of Econometrics*, 16, 121-130.
- Granger, C.W.J. and Z. Ding (1996), "Varieties of Long Memory Models," *Journal of Econometrics*, 73, 61-77.
- Granger, C.W.J., Z. Ding and S. Spear (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 R. Joyeux (1980), "An Introduction to Long Memory Time Series Models and Fractional Differencing," *Journal of Time Series Analysis*, 1, 15-39.

- Granger, C.W.J. and F. Marmol (1997), "The Correlogram of a Long Memory Process Plus a Simple Noise," Manuscript, Department of Economics, University of California, San Diego.
- Guillaume, D.M., M.M. Dacorogna, R.R. Davé, U.A. Müller, R.B. Olsen, and O.V. Pictet (1997), "From the Bird's Eye to the Microscope: A Survey of New Stylized Facts of the Intraday Foreign Exchange Markets," *Finance and Stochastics*, 1, 95-129.
- Harvey, A.C. (1998), "Long-Memory in Stochastic Volatility," in J. Knight and S. Satchell (eds.), *Forecasting Volatility in Financial Markets*. London: Butterworth-Heinemann.
- Hsieh, D.A. (1991), "Chaos and Nonlinear Dynamics: Application to Financial Markets," *Journal of Finance*, 46, 1839-1877.
- Hull, J. and A. White (1987), "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42, 381-400.
- 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.
- Hurvich, C.M. and K.I. Beltrao (1994), "Automatic Semiparametric Estimation of the Memory Parameter of a Long-Memory Time Series," *Journal of Time Series Analysis*, 15, 285-302.
- Ito, T., R.F. Engle and W.-L. Lin (1992), "Where Does the Meteor Shower Come From? The Role of Stochastic Policy Coordination," *Journal of International Economics*, 32, 221-240.
- Jorion, P. (1988), "On Jump Processes in the Foreign Exchange and Stock Markets," *Review of Financial Studies*, 1, 427-445.
- Karatzas, I. and S.E. Shreve (1988). *Brownian Motion and Stochastic Calculus*. New York: Springer Verlag.
- King, M., E. Sentana and S. Wadhvani (1994), "Volatility and Links Between National Stock Markets," *Econometrica*, 62, 901-933.
- Lo, A. and A.C. MacKinlay (1990), "An Econometric Analysis of Nonsynchronous Trading," *Journal of Econometrics*, 45, 181-212.
- Longin, F. and B. Solnik (1998), "Correlation Structure of International Equity Markets During Extremely Volatile Periods," ESSEC Graduate Business School and HEC School of Management.
- Lundin, M., M.M. Dacorogna and U.A. Müller (1998), "Correlation of High Frequency Financial Time Series," in P. Lequeux (ed.), *The Financial Markets Tick by Tick*. London: John Wiley & Sons.
- Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices," *Journal of Business*, 36, 394-419.
- Mandelbrot, B., A. Fisher and L. Calvet (1997), "A Multifractal Model of Asset Returns," Manuscript, Cowles Foundation, Yale University.

- Merton, R.C. (1980), "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323-361.
- Müller, U.A., M.M. Dacorogna, R.B. Olsen, O.V. Pictet, M. Schwarz and C. Morgengegg (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.
- Müller, U.A., M.M. Dacorogna, R.D. Davé, R.B. Olsen, O.V. Pictet, and J.E. von Weizsäcker (1997), "Volatilities of Different Time Resolutions - Analyzing the Dynamics of Market Components," *Journal of Empirical Finance*, 4, 213-239.
- Nelson, D.B. (1991), "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59, 347-370.
- Officer, R.R. (1973), "The Variability of the Market Factor of the NYSE," *Journal of Business*, 46, 434-453.
- Robinson, P.M. (1994a), "Semiparametric Analysis of Long-Memory Time Series," *Annals of Statistics*, 22, 515-539.
- Robinson, P.M. (1994b), "Time Series with Strong Dependence," in C.A. Sims (ed.), *Advances in Econometrics: Sixth World Congress Vol.1*. Cambridge, UK: Cambridge University Press.
- Robinson, P.M. (1995), "Log-Periodogram Regression of Time Series with Long-Range Dependence," *Annals of Statistics*, 23, 1048-1072.
- Robinson, P.M. and M. Henry (1998), "Long and Short Memory Conditional Heteroskedasticity in Estimating the Memory Parameter of Levels," Manuscript, London School of Economics.
- Robinson, P.M. and P. Zaffaroni (1997), "Modelling Nonlinearity and Long Memory in Time Series," in C.D. Cutler and D.T. Kaplan (eds.), *Nonlinear Dynamics and Time Series*, 161-170. Fields Institute Communications.
- Robinson, P.M. and P. Zaffaroni (1998), "Nonlinear Time Series with Long Memory: A Model for Stochastic Volatility," *Journal of Statistical Planning and Inference*, 68, 359-371.
- Schwert, G.W. (1989), "Why does Stock Market Volatility Change over Time?," *Journal of Finance*, 44, 1115-1154.
- Schwert, G.W. (1990a), "Stock Volatility and the Crash of '87," *Review of Financial Studies*, 3, 77-102.
- Schwert, G.W. (1990b), "Stock Market Volatility," *Financial Analysts Journal*, May-June, 23-34.
- Silverman, B.W. (1986). *Density Estimation for Statistics and Data Analysis*. New York: Chapman and Hall.
- Solnik, B., C. Boucrelle and Y. Le Fur (1996), "International Market Correlation and Volatility," *Financial Analysts Journal*, September-October, 17-34.

Taqqu, M.S., V. Teverovsky and W. Willinger (1995), "Estimators for Long-Range Dependence: An Empirical Study," *Fractals*, 3, 785-798.

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

Teverovsky, V. and M.S. Taqqu (1997), "Testing for Long-Range Dependence in the Presence of Shifting Mean, or a Slowly Declining Trend, Using a Variance-Type Estimator," *Journal of Time Series Analysis*, 18, 279-304.

Wright, J.H. (1999), "Testing for a Unit Root in the Volatility of Asset Returns," *Journal of Applied Econometrics*, forthcoming.

Table 1
Statistics Summarizing Unconditional Univariate Distributions
Daily Realized DM/\$ and Yen/\$ Volatilities

	$vard_t$	$vary_t$	$stdd_t$	$stdy_t$	$lstdd_t$	$lstdy_t$	cov_t	$corr_t$
<i>Mean</i>	0.529	0.538	0.679	0.684	-0.450	-0.444	0.243	0.435
<i>Variance</i>	0.238	0.272	0.068	0.071	0.121	0.124	0.075	0.028
<i>Skewness</i>	3.750	5.438	1.696	1.852	0.348	0.283	3.876	-0.203
<i>Kurtosis</i>	24.53	63.61	7.867	10.13	3.265	3.470	26.33	2.716
<i>Quantiles</i>								
<i>Min.</i>	0.061	0.029	0.247	0.171	-1.397	-1.767	-0.096	-0.244
<i>0.001</i>	0.068	0.044	0.261	0.209	-1.345	-1.568	-0.050	-0.108
<i>0.005</i>	0.084	0.067	0.290	0.260	-1.238	-1.349	-0.002	-0.009
<i>0.01</i>	0.094	0.084	0.306	0.290	-1.183	-1.238	0.008	0.033
<i>0.05</i>	0.136	0.138	0.368	0.371	-0.999	-0.992	0.034	0.152
<i>0.10</i>	0.169	0.176	0.411	0.420	-0.889	-0.869	0.052	0.214
<i>0.50</i>	0.397	0.399	0.630	0.632	-0.462	-0.459	0.166	0.437
<i>0.90</i>	0.985	1.012	0.992	1.006	-0.008	0.006	0.514	0.649
<i>0.95</i>	1.373	1.365	1.172	1.168	0.158	0.155	0.710	0.697
<i>0.99</i>	2.520	2.570	1.587	1.603	0.462	0.472	1.444	0.767
<i>0.995</i>	3.620	3.189	1.903	1.786	0.643	0.580	1.654	0.801
<i>0.999</i>	4.646	5.301	2.156	2.300	0.768	0.832	2.737	0.851
<i>Max.</i>	5.503	9.965	2.346	3.157	0.853	1.150	3.087	0.918

Notes: The daily volatility measures are constructed by summation of the corresponding intraday 5-minute returns, as detailed in the text. The sample period is December 1, 1986 until December 1, 1996.

Table 2
Correlation Matrix
Daily Realized DM/\$ and Yen/\$ Volatilities

	$vary_t$	std_t	$stdy_t$	$lstd_t$	$lstdy_t$	cov_t	$corr_t$
$vard_t$	0.538	0.961	0.553	0.858	0.513	0.809	0.344
$vary_t$	1.000	0.547	0.946	0.515	0.828	0.760	0.238
std_t	-	1.000	0.594	0.965	0.580	0.794	0.385
$stdy_t$	-	-	1.000	0.591	0.960	0.761	0.283
$lstd_t$	-	-	-	1.000	0.606	0.719	0.391
$lstdy_t$	-	-	-	-	1.000	0.683	0.295
cov_t	-	-	-	-	-	1.000	0.587

Notes: The daily volatility measures are constructed by summation of the corresponding intraday 5-minute returns, as detailed in the text. The sample period is December 1, 1986 until December 1, 1996.

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

	$vard_t$	$vary_t$	$stdd_t$	$stdy_t$	$lstdd_t$	$lstdy_t$	cov_t	$corr_t$
<i>LB</i>	3199.0	3230.2	6314.7	5086.4	9051.1	6706.4	1125.8	12291
<i>ADF</i>	-7.495	-9.083	-6.941	-8.008	-6.370	-7.817	-7.565	-5.589
<i>d</i>	0.346	0.353	0.388	0.424	0.421	0.448	0.304	0.423

Notes: The realized daily volatilities are constructed by the summation of the DM/\$ and Yen/\$ 5-minute returns from December 1, 1986 until December 1, 1996, as detailed in the text. The first row reports Ljung-Box tests using sample autocorrelations through displacement 20. The 0.001 critical value is 45.3. The second row reports Augmented Dickey-Fuller tests for unit roots in the autoregressive lag-operator polynomials of univariate ARMA representations, using ten augmentation lags. The corresponding 0.01 and 0.05 critical values are -2.86 and -3.43. The third row reports the log-periodogram regression estimate for the fractional integration parameter, d . The estimates are based on the $m = \lceil T^{4/5} \rceil = 514$ lowest-frequency periodogram ordinates. The asymptotic standard error of the estimator depends only on m ; it is $\pi \cdot (24 \cdot m)^{-1/2} = 0.028$.

Table 4
Statistics Summarizing Unconditional Marginal Distributions
Realized DM/\$ and Yen/\$ Volatilities at Different Levels of Temporal Aggregation

	$vard_{t,h}$	$vary_{t,h}$	$std_{t,h}$	$stdy_{t,h}$	$lstd_{t,h}$	$lsty_{t,h}$	$cov_{t,h}$	$corr_{t,h}$
<i>Weekly, h=5</i>								
<i>Mean</i>	2.645	2.691	1.555	1.567	0.399	0.406	1.217	0.448
<i>Variance</i>	3.250	3.643	0.227	0.238	0.083	0.083	0.947	0.021
<i>Skewness</i>	2.483	2.815	1.230	1.404	0.242	0.350	2.281	-0.127
<i>Kurtosis</i>	12.40	15.33	5.422	6.204	3.146	3.334	10.30	2.393
<i>Bi-Weekly, h=10</i>								
<i>Mean</i>	5.296	5.387	2.216	2.234	0.759	0.767	2.438	0.454
<i>Variance</i>	10.52	11.27	0.387	0.398	0.071	0.070	2.840	0.019
<i>Skewness</i>	2.109	2.260	1.108	1.195	0.262	0.350	1.715	-0.141
<i>Kurtosis</i>	9.378	10.87	4.771	5.050	3.050	3.129	6.533	2.242
<i>Tri-Weekly, h=15</i>								
<i>Mean</i>	7.935	8.074	2.718	2.746	0.965	0.978	3.652	0.456
<i>Variance</i>	21.78	23.03	0.548	0.537	0.068	0.061	5.698	0.018
<i>Skewness</i>	2.036	2.289	1.023	1.297	0.216	0.496	1.633	-0.144
<i>Kurtosis</i>	9.636	10.43	4.578	5.341	2.993	3.229	6.019	2.231
<i>Monthly, h=20</i>								
<i>Mean</i>	10.59	10.77	3.150	3.179	1.116	1.127	4.875	0.458
<i>Variance</i>	34.35	36.46	0.668	0.668	0.061	0.058	8.992	0.017
<i>Skewness</i>	1.619	1.893	0.959	1.147	0.311	0.472	1.428	-0.188
<i>Kurtosis</i>	5.816	7.457	3.792	4.422	2.818	3.042	4.855	2.130

Note: See notes to Table 1. The results for the weekly, bi-weekly, tri-weekly and monthly distributions are based on 489, 244, 163 and 122 observations, respectively.

Table 5
Correlation Matrices
Realized DM/\$ and Yen/\$ Volatilities at Different Levels of Temporal Aggregation

	$vary_{t,h}$	$std_{t,h}$	$stdy_{t,h}$	$lstdd_{t,h}$	$lstdy_{t,h}$	$cov_{t,h}$	$corr_{t,h}$
<i>Weekly, h=5</i>							
$vard_{t,h}$	0.496	0.977	0.507	0.911	0.494	0.780	0.314
$vary_{t,h}$	1.000	0.522	0.975	0.518	0.905	0.771	0.207
$std_{t,h}$	-	1.000	0.547	0.978	0.545	0.786	0.338
$stdy_{t,h}$	-	-	1.000	0.558	0.976	0.764	0.232
$lstdd_{t,h}$	-	-	-	1.000	0.572	0.750	0.342
$lstdy_{t,h}$	-	-	-	-	1.000	0.721	0.249
$cov_{t,h}$	-	-	-	-	-	1.000	0.621
<i>Bi-weekly, h=10</i>							
$vard_{t,h}$	0.469	0.982	0.480	0.929	0.475	0.768	0.279
$vary_{t,h}$	1.000	0.499	0.981	0.507	0.929	0.749	0.165
$std_{t,h}$	-	1.000	0.520	0.982	0.522	0.782	0.298
$stdy_{t,h}$	-	-	1.000	0.538	0.982	0.748	0.184
$lstdd_{t,h}$	-	-	-	1.000	0.550	0.760	0.302
$lstdy_{t,h}$	-	-	-	-	1.000	0.720	0.199
$cov_{t,h}$	-	-	-	-	-	1.000	0.622
<i>Tri-weekly, h=15</i>							
$vard_{t,h}$	0.501	0.982	0.505	0.933	0.495	0.770	0.245
$vary_{t,h}$	1.000	0.518	0.983	0.517	0.937	0.764	0.134
$std_{t,h}$	-	1.000	0.533	0.983	0.532	0.782	0.268
$stdy_{t,h}$	-	-	1.000	0.543	0.984	0.757	0.146
$lstdd_{t,h}$	-	-	-	1.000	0.554	0.761	0.275
$lstdy_{t,h}$	-	-	-	-	1.000	0.728	0.155
$cov_{t,h}$	-	-	-	-	-	1.000	0.599
<i>Monthly, h=20</i>							
$vard_{t,h}$	0.477	0.988	0.481	0.950	0.474	0.767	0.233
$vary_{t,h}$	1.000	0.499	0.987	0.507	0.949	0.745	0.107
$std_{t,h}$	-	1.000	0.510	0.987	0.508	0.777	0.248
$stdy_{t,h}$	-	-	1.000	0.525	0.987	0.740	0.113
$lstdd_{t,h}$	-	-	-	1.000	0.531	0.764	0.255
$lstdy_{t,h}$	-	-	-	-	1.000	0.717	0.120
$cov_{t,h}$	-	-	-	-	-	1.000	0.595

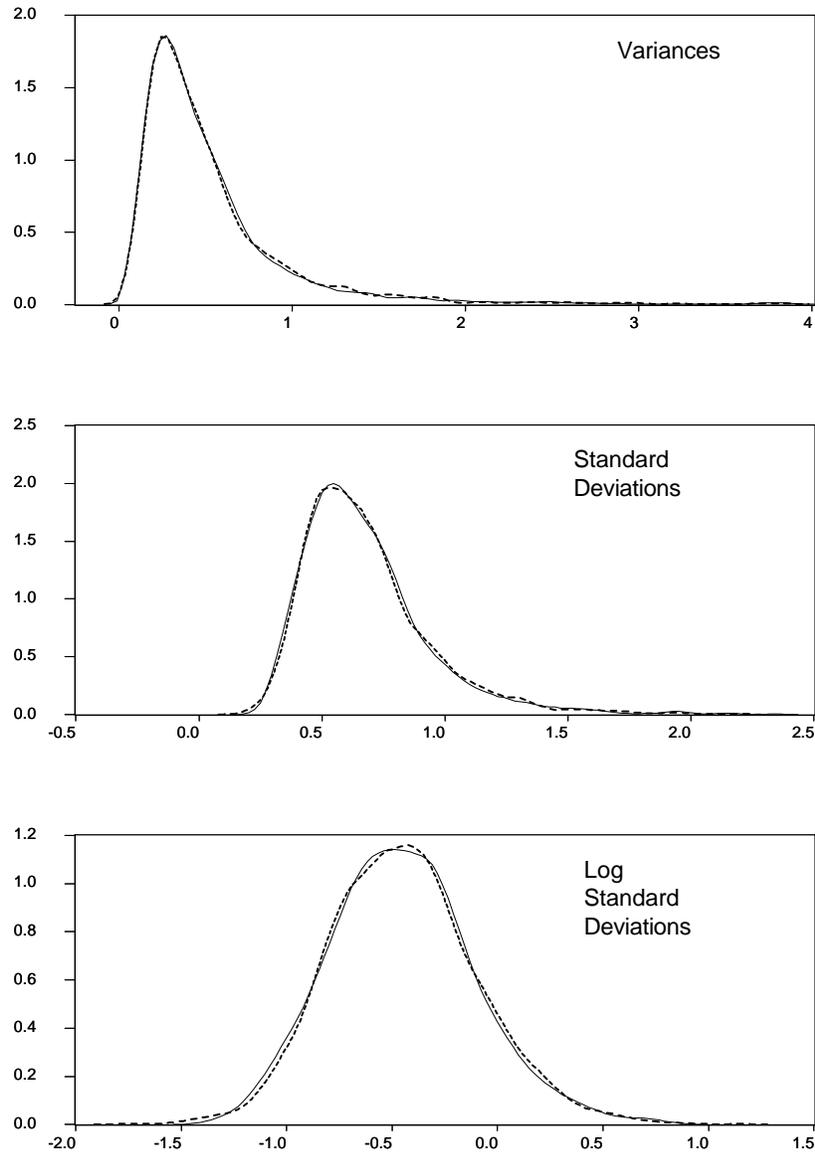
Note: See notes to Table 2. The weekly, bi-weekly, tri-weekly and monthly correlations are based on 489, 244, 163 and 122 observations, respectively.

Table 6
Dynamic Dependency Measures
Realized DM/\$ and Yen/\$ Volatilities Under Temporal Aggregation

	$vard_{t,h}$	$vary_{t,h}$	$stdd_{t,h}$	$stdy_{t,h}$	$lstdd_{t,h}$	$lstdy_{t,h}$	$cov_{t,h}$	$corr_{t,h}$
<i>Weekly, h=5</i>								
<i>LB</i>	621.5	503.3	799.1	612.0	938.2	627.4	447.7	2857
<i>d</i>	0.485	0.455	0.494	0.511	0.528	0.533	0.425	0.488
<i>Bi-weekly, h=10</i>								
<i>LB</i>	220.3	200.9	270.3	219.0	311.7	207.5	176.5	1246
<i>d</i>	0.510	0.501	0.501	0.500	0.479	0.485	0.480	0.499
<i>Tri-weekly, h=15</i>								
<i>LB</i>	107.6	103.7	130.1	117.9	147.8	116.6	113.6	672.0
<i>d</i>	0.430	0.416	0.432	0.430	0.424	0.429	0.357	0.551
<i>Monthly, h=20</i>								
<i>LB</i>	68.7	67.2	81.2	71.4	95.0	68.7	77.9	412.5
<i>d</i>	0.433	0.474	0.426	0.485	0.463	0.482	0.435	0.628

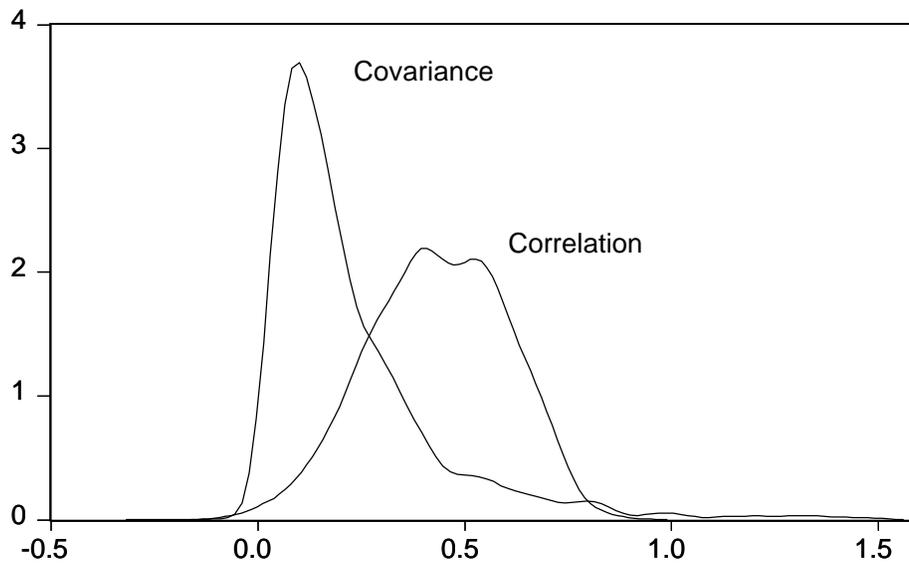
Notes: See notes to Table 3. The Ljung-Box tests, LB, using sample autocorrelations through displacement 20, are based on 489, 244, 163, and, 122 observations for the weekly, bi-weekly, tri-weekly and monthly estimates, respectively. The 0.001 critical value is 45.3. The reported log periodogram estimates of d are calculated from the first $m = \lfloor (T/h)^{4/5} \rfloor$, or 141, 81, 58, and 46, periodogram ordinates, respectively. The corresponding asymptotic standard errors are 0.054, 0.071, 0.084 and 0.095, respectively.

Figure 1
Distributions of Daily Realized Exchange Rate Volatilities:
Variances, Standard Deviations, and Log Standard Deviations



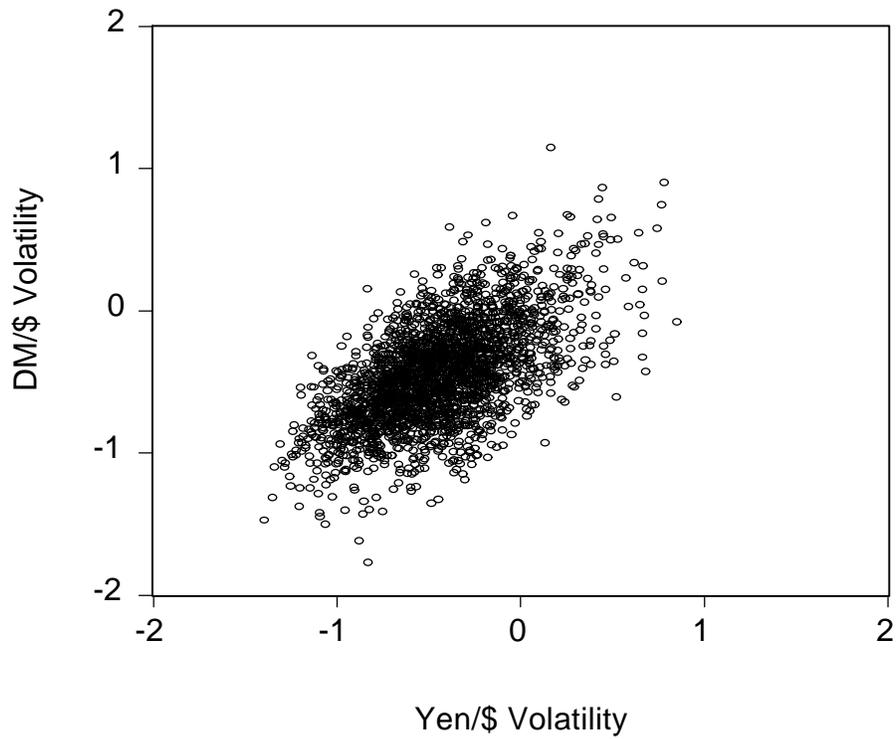
Notes: The figure shows estimates of the unconditional densities of daily realized DM/\$ and Yen/\$ volatility, measured in three ways. The top panel shows estimated distributions for the DM/\$ variance (solid line, *vard* in the notation of the text) and Yen/\$ variance (dashed line, *vary* in the notation of the text). The middle panel shows estimated distributions for the DM/\$ standard deviation (solid line, *std*) and Yen/\$ standard deviation (dashed line, *stdy*). The bottom panel shows estimated distributions for the DM/\$ and Yen/\$ logarithmic standard deviations (*lstd*, solid line, and *lstdy*, dashed line). The sample period runs from December 1, 1986, until December 1, 1996.

Figure 2
Distributions of Daily Realized Exchange Rate Comovement:
Covariances and Correlations



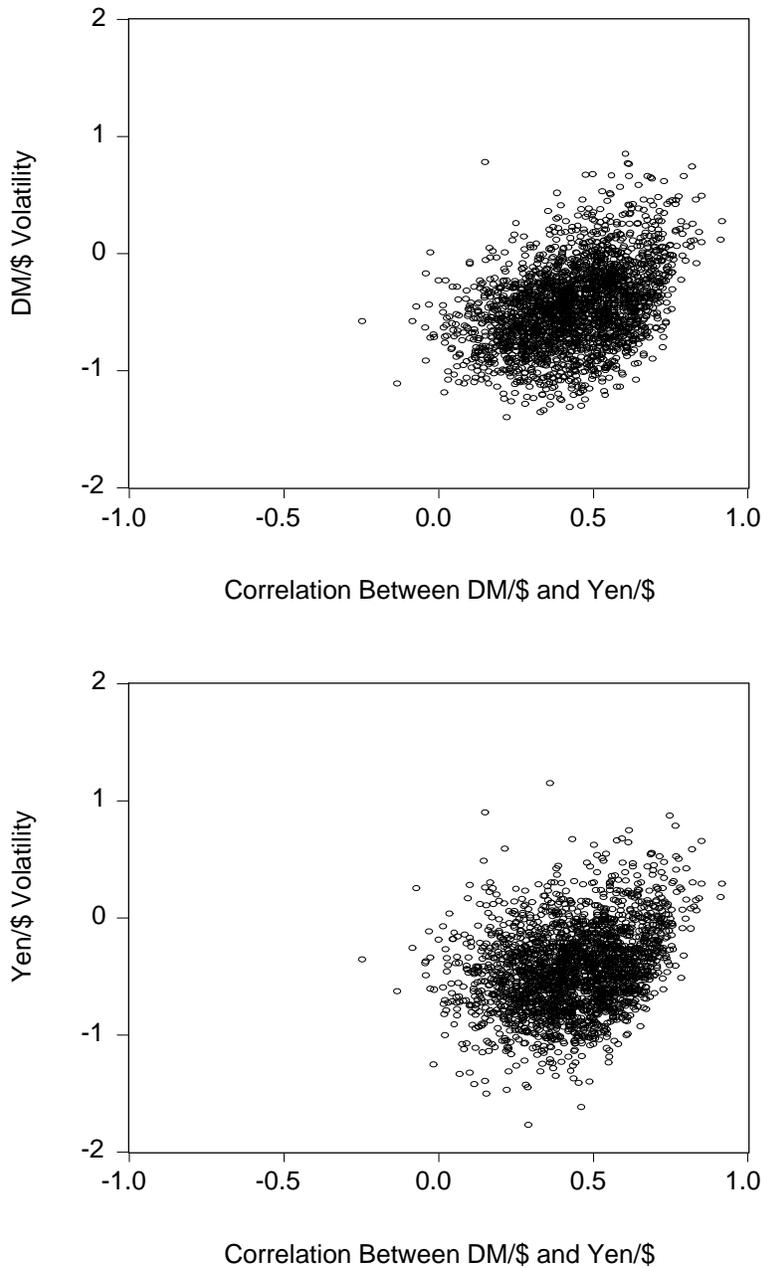
Notes: The figure shows estimates of the unconditional densities of daily realized DM/\$ and Yen/\$ comovement, measured in two ways: covariance and correlation (*cov* and *corr*, in the notation of the text). The sample period runs from December 1, 1986, until December 1, 1996.

Figure 3
Daily Realized Exchange Rate Volatilities: DM/\$ vs. Yen/\$



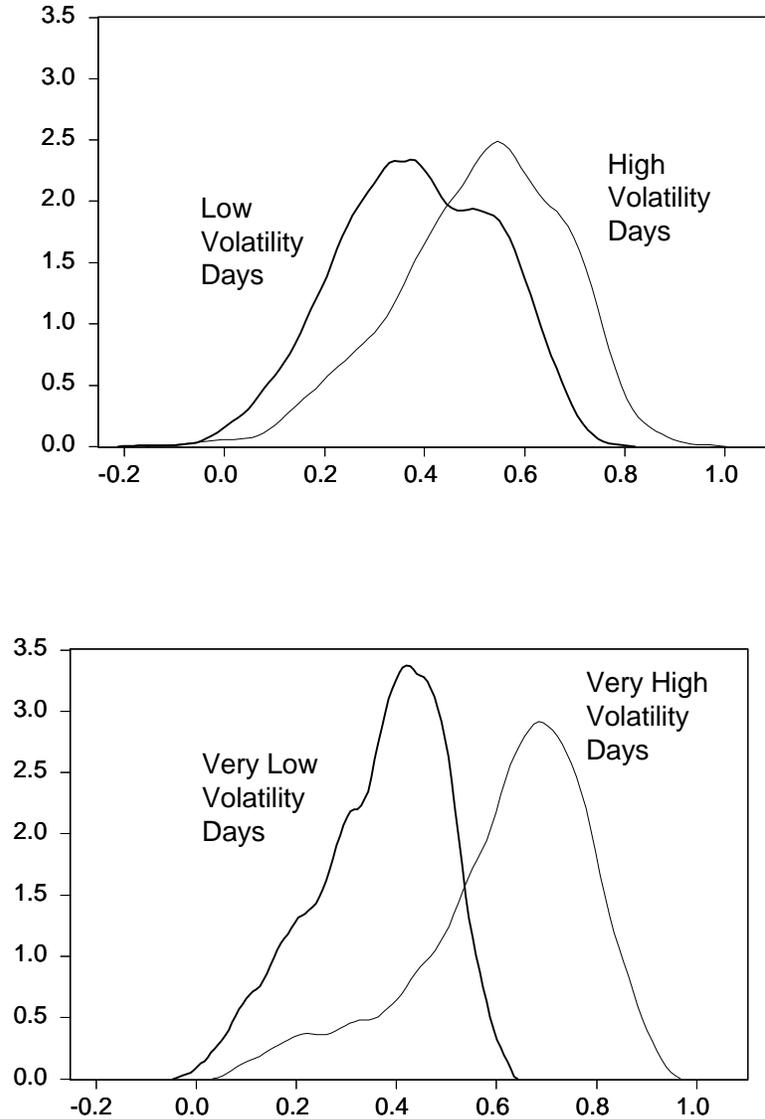
Notes: The figure shows a scatterplot of daily realized DM/\$ and Yen/\$ volatilities, measured as logarithmic standard deviations (*lstd* vs. *lstdy*, in the notation of the text). The sample period runs from December 1, 1986, until December 1, 1996.

Figure 4
Daily Realized Exchange Rate Volatilities vs. Correlation



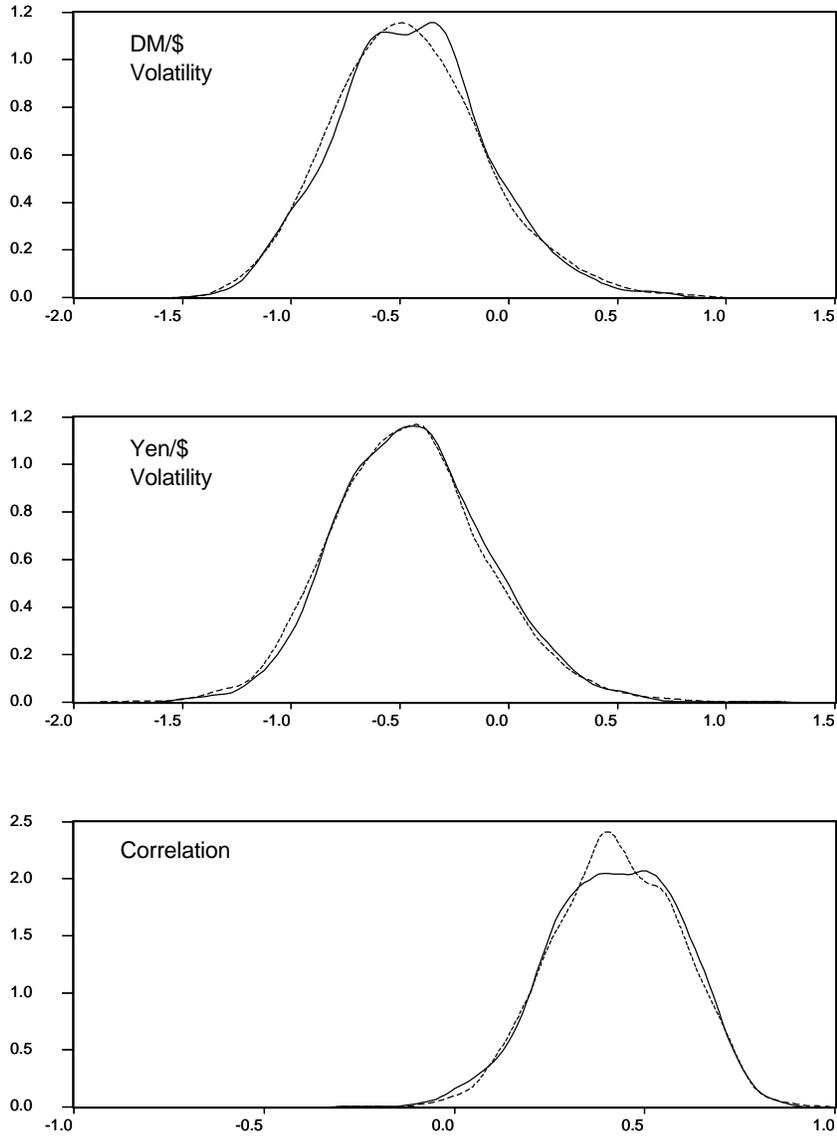
Notes: The top panel shows a scatterplot of daily realized DM/\$ volatility measured as logarithmic standard deviation (*lstd*, in the notation of the text) vs. daily realized (DM/\$, Yen/\$) correlation (*corr*, in the notation of the text). The bottom panel shows a scatterplot of daily realized Yen/\$ volatility measured as logarithmic standard deviation (*lstdy*) vs. daily realized (DM/\$, Yen/\$) correlation (*corr*). The sample period runs from December 1, 1986, until December 1, 1996.

Figure 5
Distributions of Daily Realized Exchange Rate Correlation:
Low Volatility vs. High Volatility Days



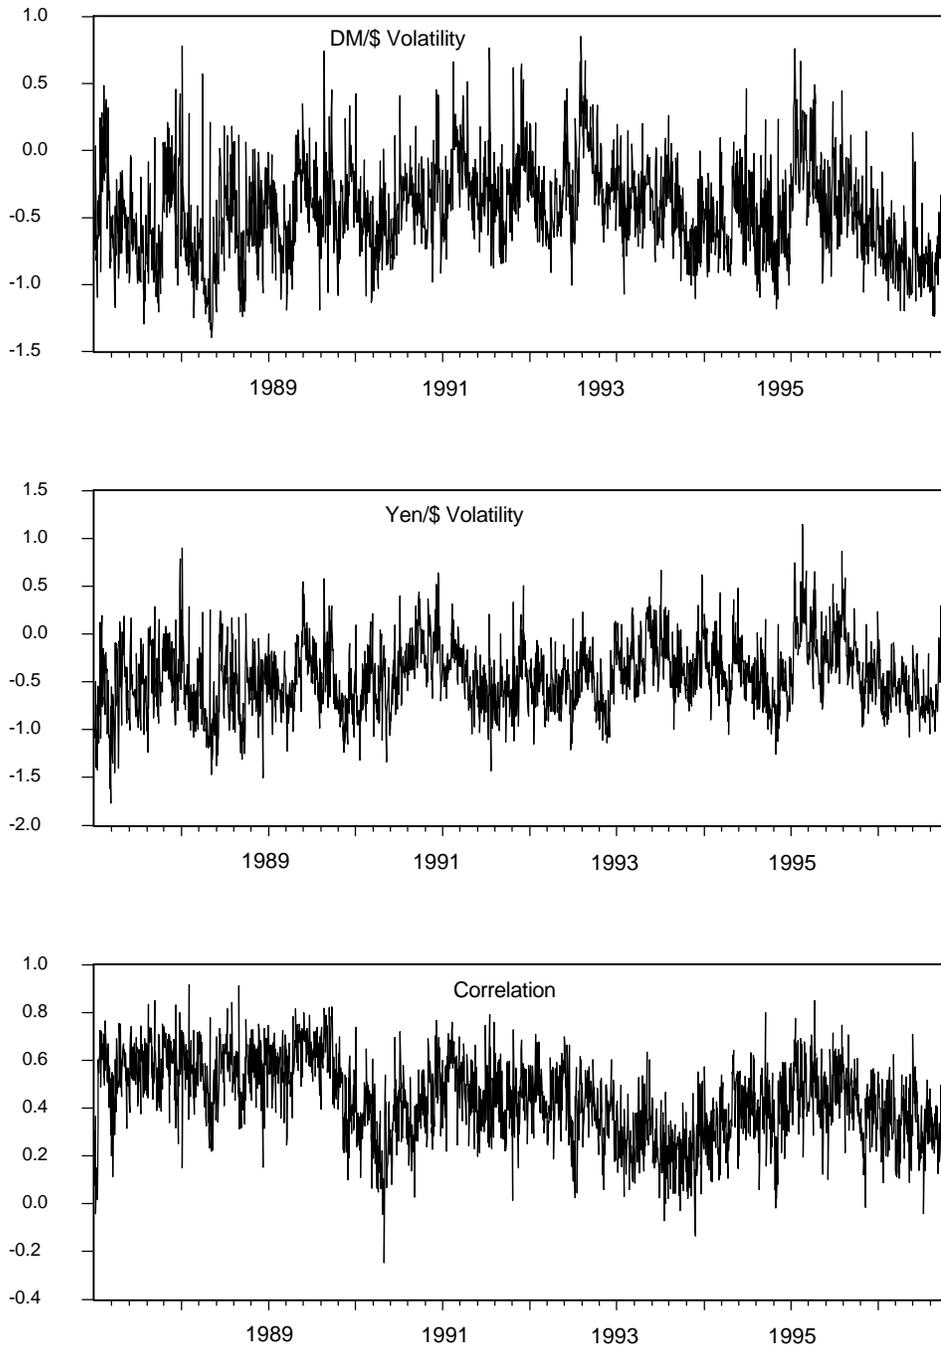
Notes: The figure shows estimates of the unconditional densities of daily realized (DM/\$, Yen/\$) correlation (*corr*, in the notation of the text) in low-volatility and high-volatility periods, where volatility is measured as logarithmic standard deviation (*lstd* and *lstdy*). The top panel shows the estimated distribution of *corr* when both *lstd* and *lstdy* are less than their median (“low volatility”) and when both are greater than their median (“high volatility”). The bottom panel shows the estimated distribution of *corr* when both *lstd* and *lstdy* are less than their tenth percentile (“very low volatility”) and when both are greater than their ninetieth percentile (“very high volatility”). The sample period runs from December 1, 1986, until December 1, 1996.

Figure 6
Distribution of Realized Volatilities and Correlation:
Positive- vs. Negative-Return Days



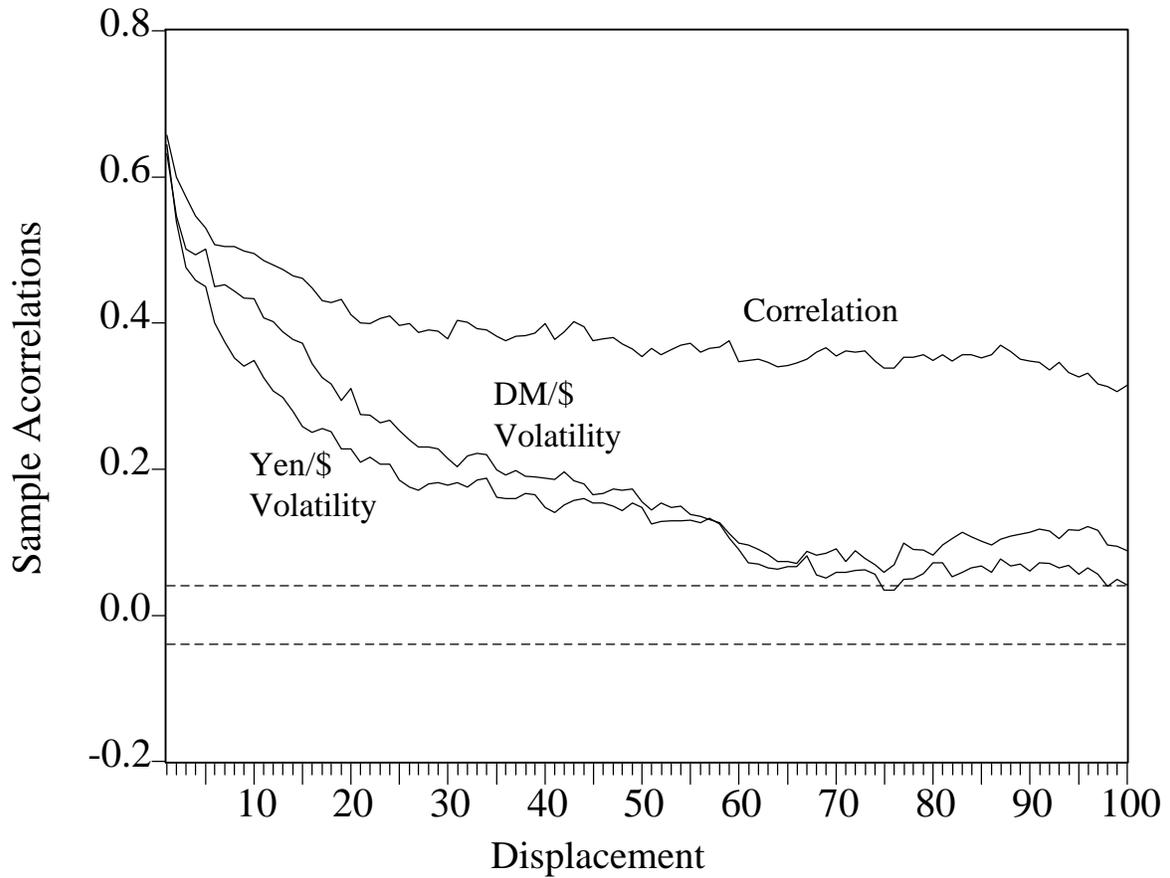
Notes: The figure shows estimates of the unconditional densities of daily realized DM/\$ and Yen/\$ volatilities (*lstd* and *lsty*, in the notation of the text) and correlation (*corr*) on positive and negative return days. The top panel shows the distribution of *lstd* on positive-return days (solid line) and negative-return days (dashed line). The middle panel shows the distribution of *lsty* on positive-return days (solid line) and negative-return days (dashed line). The bottom panel shows the distribution of *corr* on positive-return days (solid line) and negative-return days (dashed line). The sample period runs from December 1, 1986, until December 1, 1996.

Figure 7
Times Series of Daily Realized Volatilities and Correlation



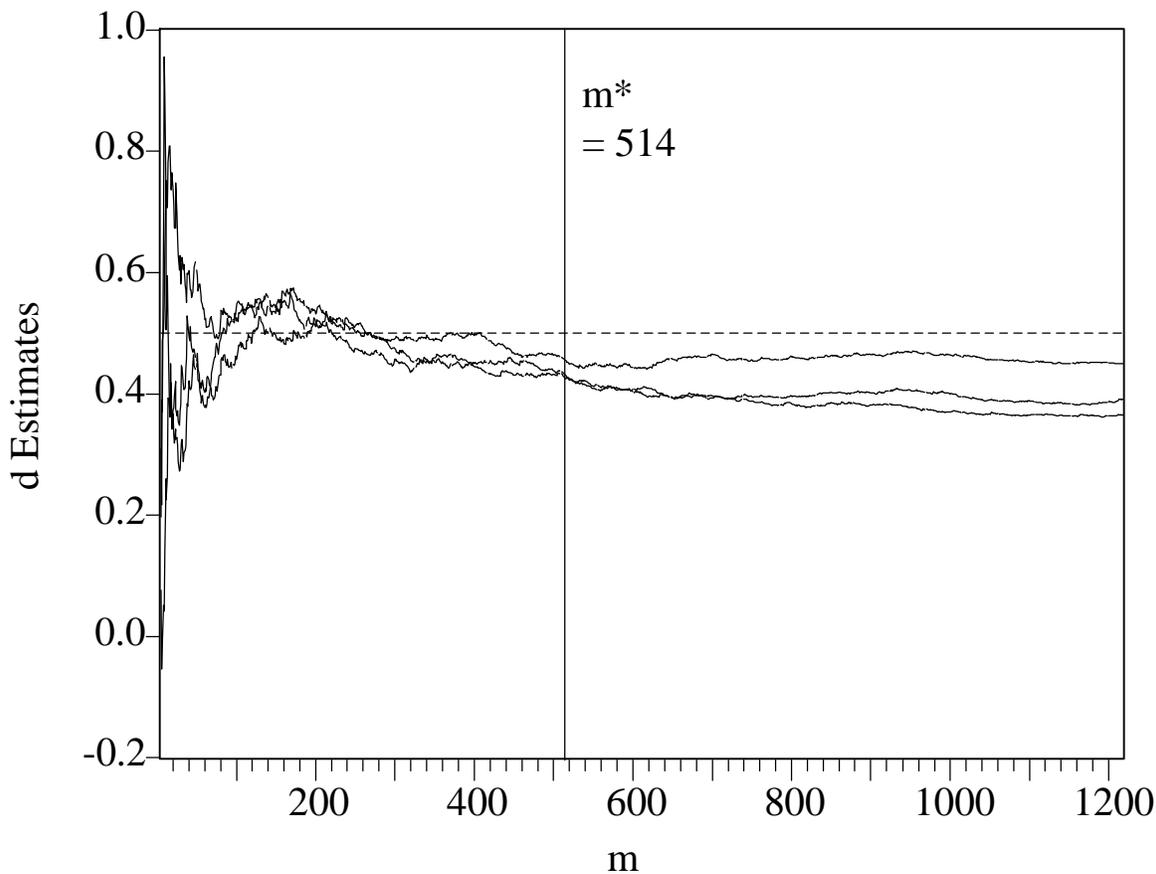
Notes: The figure shows the time series of daily realized DM/\$ and Yen/\$ volatilities (*lstd* and *lstdy*, in the notation of the text) and correlation (*corr*). The top panel shows *lstd*, the middle panel shows *lstdy*, and the bottom panel shows *corr*. The sample period runs from December 1, 1986, until December 1, 1996.

Figure 8
Sample Autocorrelations of Realized Volatilities and Correlation



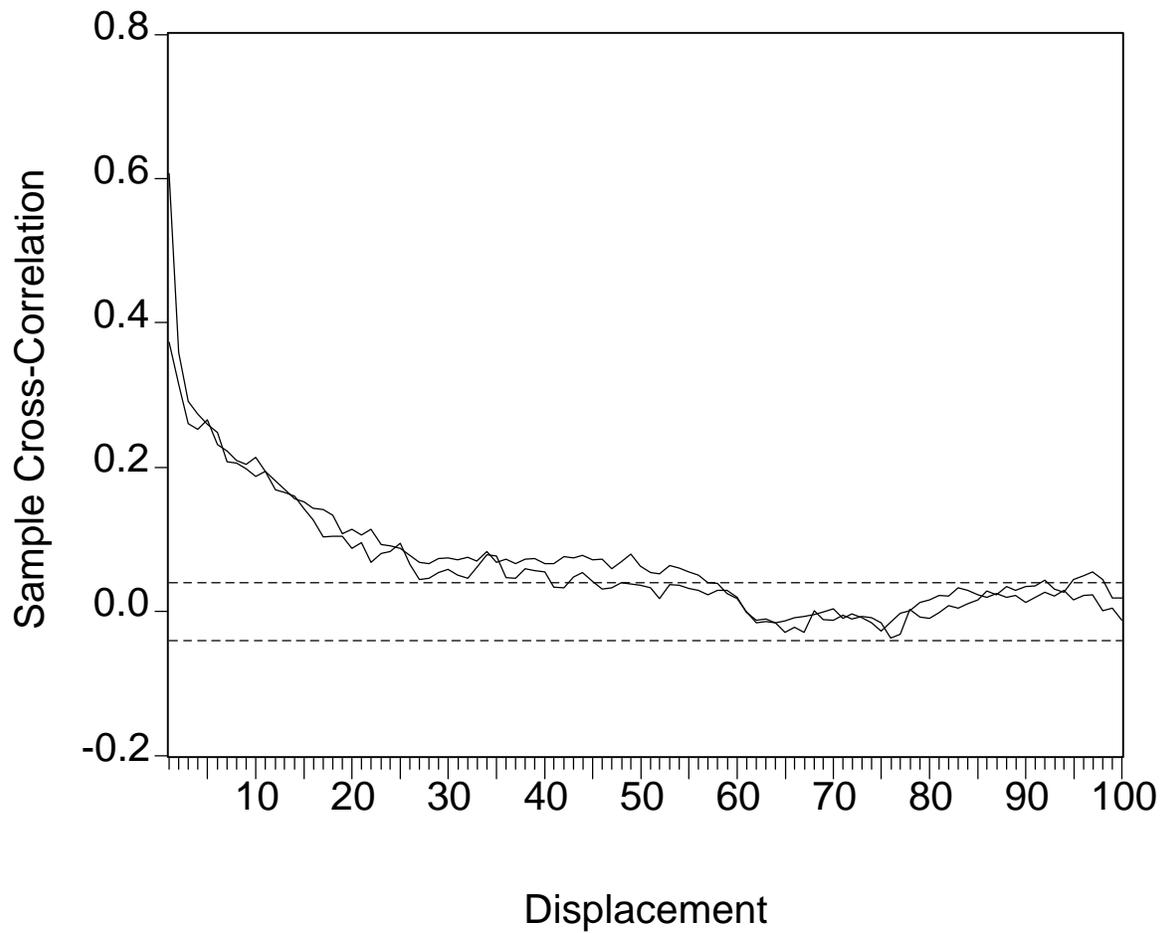
Notes: The figure shows estimates of the autocorrelation functions of daily realized DM/\$ and Yen/\$ volatilities (*lstd* and *lstdy*, in the notation of the text) and correlation (*corr*), through displacement 100. The parallel dashed lines are the 95% Bartlett intervals. The sample period runs from December 1, 1986, until December 1, 1996.

Figure 9
Estimates of Parameters of Fractional Integration
Realized Volatilities and Correlation



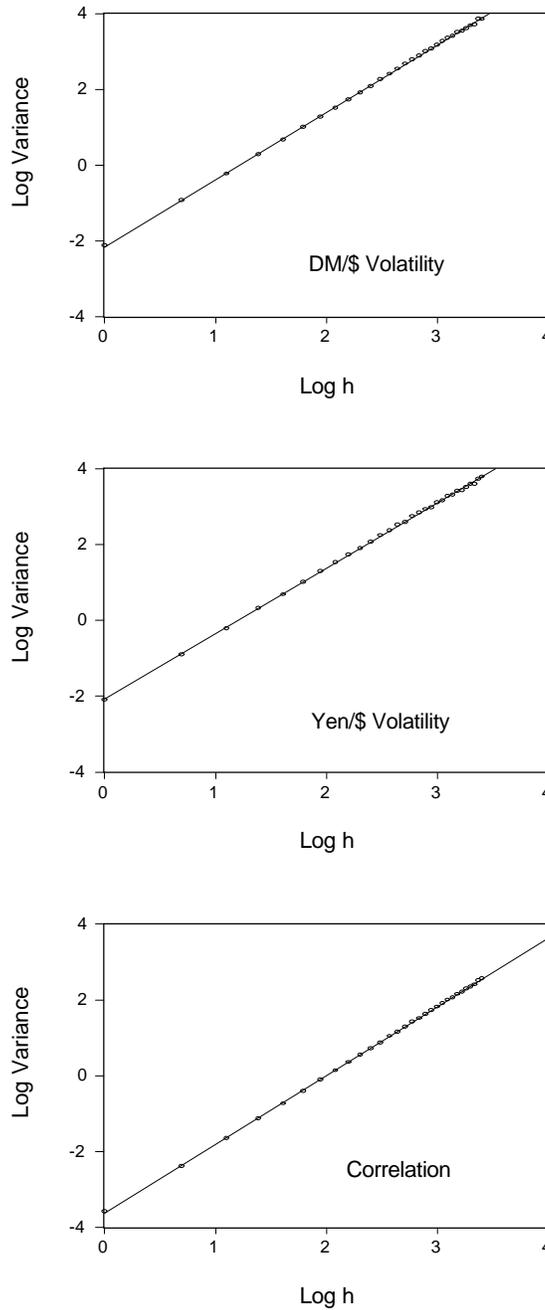
Notes: The figure shows log-periodogram estimates of the degree of fractional integration, d , as a function of the number of periodogram ordinates, m , used for their computation. m^* is a benchmark choice of m motivated by the results of Hurvich, Deo and Brodsky (1998). We show estimates of d for realized DM/\$ volatility (*lstdd*), realized Yen/\$ volatility (*lstdy*), and realized correlation (*corr*). The sample period runs from December 1, 1986, until December 1, 1996.

Figure 10
Sample Cross-Correlations of Realized DM/\$ and Yen/\$ Volatilities



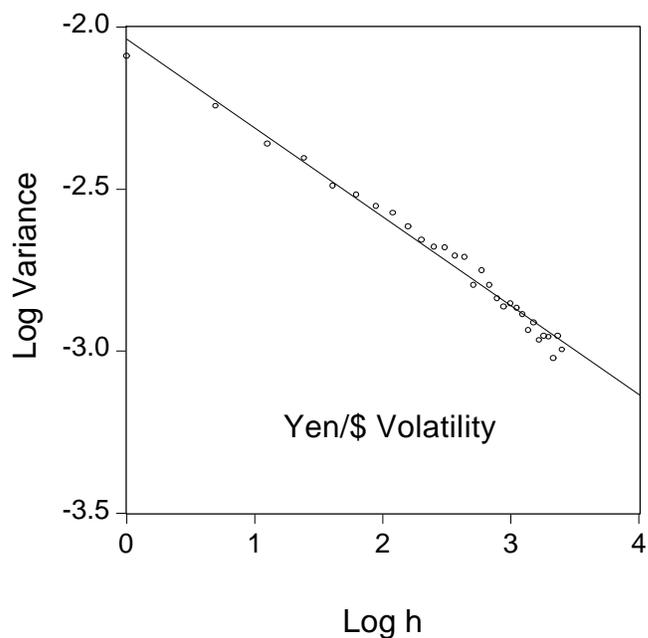
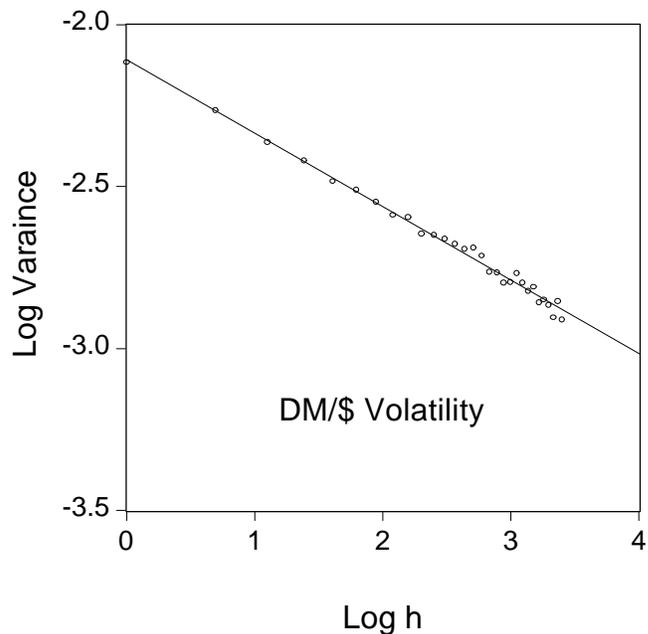
Notes: The figure shows estimates of the cross-correlation functions of daily realized DM/\$ and Yen/\$ volatilities ($lstd$ and $lstdy$, in the notation of the text). That is, it shows $Corr(lstd_t, lstdy_{t-j})$ and $Corr(lstd_{t-j}, lstdy_t)$, $j = 1, 2, \dots, 100$. The parallel dashed lines are the 95% Bartlett intervals. The sample period runs from December 1, 1986, until December 1, 1996.

Figure 11
Realized Daily Volatilities and Correlation:
Variance of h-fold Partial sums as a function of h



Notes: The figure shows the logarithm of the variance of partial sums against the logarithm of the aggregation levels, $h = 1, 2, \dots, 30$, for the DM/\$ and Yen/\$ daily realized volatilities and correlation. In the notation of the text, the top panel shows $\log(\text{Var}([lstdd_t]_h))$, the middle panel shows $\log(\text{Var}[lstdy_t]_h)$, and the bottom panel shows $\log(\text{Var}([corr_t]_h))$. The sample period runs from December 1, 1986, until December 1, 1996.

Figure 12
Realized h-Day Volatilities:
Variance as a function of h



Notes: The figure shows the logarithm of the variance of h-day volatilities against the logarithm of the aggregation levels, $h = 1, 2, \dots, 30$, for the DM/\$ and Yen/\$ h-day realized volatilities. In the notation of the text, the top panel shows $\log(\text{Var}(\text{lstd}_{t,h}))$, and the bottom panel shows $\log(\text{Var}(\text{lstd}_{y,t,h}))$.