

The Distribution of Realized Exchange Rate Volatility

Torben G. ANDERSEN, Tim BOLLERSLEV, Francis X. DIEBOLD, and Paul LABYS

Using high-frequency data on deutschemark and yen returns against the dollar, we construct model-free estimates of daily exchange rate volatility and correlation that cover an entire decade. Our estimates, termed realized volatilities and correlations, are not only model-free, but also approximately free of measurement error under general conditions, which we discuss in detail. Hence, for practical purposes, we may 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 persistent dynamics in volatilities and correlations, evidence of long-memory dynamics in volatilities and correlations, and remarkably precise scaling laws under temporal aggregation.

KEY WORDS: Forecasting; High-frequency Data; Integrated volatility; Long-memory; Quadratic variation; Realized volatility; Risk management.

1. INTRODUCTION

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

In this article, we introduce a new and complementary volatility measure, termed realized volatility. The mechanics are simple—we compute daily realized volatility simply by summing intraday squared returns—but the theory is deep: By sampling intraday returns sufficiently frequently, the realized volatility can be made arbitrarily close to the underlying integrated volatility, the integral of instantaneous volatility over the interval of interest, which is a natural volatility measure. Hence for practical purposes, we may treat volatility as observed, which enables us to examine its properties directly, using much simpler techniques than the complicated econometric models required when volatility is latent.

Our analysis is in the spirit of and extends earlier contributions of French, Schwert, and Stambaugh (1987), Hsieh (1991), Schwert (1989, 1990), and, more recently, Taylor and Xu (1997). We progress, however, in important directions. First, we provide rigorous theoretical underpinnings for the volatility measures for the general case of a special semimartingale. Second, our analysis is explicitly 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 that consists of 10 years of continuously recorded 5-min returns on two major currencies. The high-frequency returns allow us to examine *daily* volatilities, which are of central concern in both academia and industry. In particular, the persistent volatility fluctuations of interest in risk management, asset pricing, portfolio allocation, and forecasting are very much present at the daily horizon.

We proceed as follows. In Section 2 we provide a formal and detailed justification for our realized volatility and correlation measures as highly accurate estimates of the underlying quadratic variation and covariation, assuming only that returns evolve as special semimartingales. Among other things, we relate our realized volatilities and correlations to the conditional variances and correlations common in the econometrics literature and to the notion of integrated variance common in the finance literature, and we show that they remain valid in the presence of jumps. Such background is needed for a serious understanding of our volatility and correlation measures, and it is lacking in the earlier literature on which we build.

Torben G. Andersen is Professor, Department of Finance, Kellogg Graduate School of Management, Northwestern University and National Bureau of Economic Research (E-mail: t-andersen@kellogg.northwestern.edu). Tim Bollerslev is Professor, Departments of Economics and Finance, Duke University and National Bureau of Economic Research (E-mail: boller@econ.duke.edu). Francis X. Diebold is Professor, Departments of Economics and Statistics, University of Pennsylvania and National Bureau of Economic Research (E-mail: fdiebold@mail.sas.upenn.edu). Paul Labys is Ph.D. Candidate, Department of Economics, University of Pennsylvania (E-mail: labys@ssc.sas.upenn.edu). This work was supported by the National Science Foundation. The authors are grateful to Olsen and Associates for making the intraday exchange rate quotations available. The authors thank participants at the 1999 North American and European Meetings of the Econometric Society, the Spring 1999 National Bureau of Economic Research Asset Pricing Meeting, the 2000 Western Finance Association Meeting, and seminars at Boston University, Chicago, Columbia, Georgetown, International Monetary Fund, Johns Hopkins, New York University, and London School of Economics (LSE), as well as the editor, associate editor, two anonymous referees, Dave Backus, Michael Brandt, Rohit Deo, Rob Engle, Clive Granger, Lars Hansen, Joel Hasbrouck, Ludger Hentschel, Cliff Hurvich, Blake LeBaron, Richard Levich, Bill Schwert, Rob Stambaugh, George Tauchen, and Stephen Taylor. Much of this paper was written while Diebold was visiting the Stern School of Business, New York University, whose hospitality is gratefully acknowledged.

In Section 3, we discuss the high-frequency deutschemark–U.S. dollar (DM/\$) and yen–U.S. dollar (yen/\$) returns that provide the basis for our empirical analysis, and we also detail the construction of our realized daily variances and covariances. In Sections 4 and 5, we characterize the unconditional and conditional distributions of the daily volatilities, respectively, including long-memory features. In Section 6, we explore issues related to temporal aggregation, with particular focus on the scaling laws implied by long memory, and we provide concluding remarks in Section 7.

2. VOLATILITY MEASUREMENT: THEORY

In this section we develop the foundations of our volatility and covariance measures. When markets are open, trades may occur at any instant. Therefore, returns as well as corresponding measures of volatility may, in principle, be obtained over arbitrarily short intervals. We, therefore, model the underlying price process in continuous time. We first introduce the relevant concepts, after which we show how the volatility measures may be approximated using high-frequency data, and we illustrate the concrete implications of our concepts for standard Itô and mixed jump–diffusion processes.

2.1 Financial Returns as a Special Semimartingale

Arbitrage-free price processes of practical relevance for financial economics belong to the class of special semimartingales. They allow for a unique decomposition of returns into a local martingale and a predictable finite variation process. The former represents the “unpredictable” innovation, whereas the latter has a locally deterministic drift that governs the instantaneous mean return, as discussed in Back (1991).

Formally, for a positive integer T and $t \in [0, T]$, let \mathcal{F}_t be the σ field that reflects the information at time t , so that $\mathcal{F}_s \subseteq \mathcal{F}_t$ for $0 \leq s \leq t \leq T$, and let P denote a probability measure on (Ω, P, \mathcal{F}) , where Ω represents the possible states of the world and $\mathcal{F} \equiv \mathcal{F}_T$ is the set of events that are distinguishable at time T . Also assume that the information filtration $(\mathcal{F}_t)_{t \in [0, T]}$ satisfies the usual conditions of P completeness and right continuity. The evolution of any arbitrage-free logarithmic price process, p_k , and the associated continuously compounded return over $[0, t]$ may then be represented as

$$p_k(t) - p_k(0) = M_k(t) + A_k(t), \quad (1)$$

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

The formulation (1) is very general and includes all specifications used in standard asset pricing theory. It includes, for example, Itô, jump, and mixed jump–diffusion processes, and it does not require a Markov assumption. It can also accommodate long memory, either in returns or in return volatility, as long as care is taken to eliminate the possibility of arbitrage first noted by Meheshwaran and Sims (1993), using, for example, the methods of Rogers (1997) or Comte and Renault (1998).

Without loss of generality, each component in Equation (1) may be assumed to be cadlag (right continuous with left limits). The corresponding caglad (left continuous with right limits) process is now p_{k-} , defined as $p_{k-}(t) \equiv \lim_{s \rightarrow t, s \leq t} p_k(s)$ for each $t \in [0, T]$, and the jumps are $\Delta p_k \equiv p_k - p_{k-}$ or

$$\Delta p_k(t) \equiv p_k(t) - \lim_{s \rightarrow t, s \leq t} p_k(s). \quad (2)$$

By no arbitrage, the occurrence and size of jumps are unpredictable, so M_k contains the (compensated) jump part of p_k along with any infinite variation components, whereas A_k has continuous paths. We may further decompose M_k into a pair of local martingales, one with continuous and infinite variation paths, M_k^c , and another of finite variation, ΔM_k , representing the compensated jump component so that $M_k = M_k^c + \Delta M_k$. Equation (1) becomes

$$p_k(t) - p_k(0) = M_k^c(t) + \Delta M_k(t) + A_k(t). \quad (3)$$

Finally, we introduce some formal notation for the returns. For concreteness, we normalize the unit interval to be one trading day. For $m \cdot T$ a positive integer, indicating the number of return observations obtained by sampling prices m times per day, the return on asset k over $[t - 1/m, t]$ is

$$r_{k, (m)}(t) \equiv p_k(t) - p_k(t - 1/m), \quad t = 1/m, 2/m, \dots, T. \quad (4)$$

Hence, $m > 1$ corresponds to high-frequency intraday returns, whereas $m < 1$ indicates interdaily returns.

2.2 Quadratic Variation and Covariation

Development of formal volatility measures requires a bit of notation. For any semimartingale X and predictable integrand H , the stochastic integral $\int H dX = \{\int_0^t H(s) dX(s)\}_{t \in [0, T]}$ is well defined, and for two semimartingales X and Y , the quadratic variation and covariation processes, $[X, X] = ([X, X])_{t \in [0, T]}$ and $[X, Y] = ([X, Y])_{t \in [0, T]}$, are given by

$$[X, X] = X^2 - 2 \int X_- dX, \quad (5a)$$

$$[X, Y] = XY - \int X_- dY - \int Y_- dX, \quad (5b)$$

where the notation X_- means the process whose value at s is $\lim_{u \rightarrow s, u < s} X_u$; see Protter (1990, sections 2.4–2.6). These processes are semimartingales of finite variation on $[0, T]$. The following properties are important for our interpretation of these quantities as volatility measures. For an increasing sequence of random partitions of $[0, T]$, $0 = \tau_{m,0} \leq \tau_{m,1} \leq \dots$, so that $\sup_{j \geq 1} (\tau_{m,j+1} - \tau_{m,j}) \rightarrow 0$ and $\sup_{j \geq 1} \tau_{m,j} \rightarrow T$ for $m \rightarrow \infty$ with probability 1, we have for $t \wedge \tau \equiv \min(t, \tau)$ and $t \in [0, T]$,

$$\lim_{m \rightarrow \infty} \{X(0)Y(0) + \sum_{j \geq 1} [X(t \wedge \tau_{m,j}) - X(t \wedge \tau_{m,j-1})] \times [Y(t \wedge \tau_{m,j}) - Y(t \wedge \tau_{m,j-1})]\} \rightarrow [X, Y]_t, \quad (6)$$

where the convergence is uniform on $[0, T]$ in probability. In addition, we have that

$$[X, Y]_0 = X(0)Y(0), \tag{7a}$$

$$\Delta[X, Y] = \Delta X \Delta Y, \tag{7b}$$

$$[X, X] \text{ is an increasing process.} \tag{7c}$$

Finally, if X and Y are locally square integrable local martingales, the covariance of X and Y over $[t-h, t]$ is given by the expected increment to the quadratic covariation,

$$\text{Cov}(X(t), Y(t)/\mathcal{F}_{t-h}) = E([X, Y]_t/\mathcal{F}_{t-h}) - [X, Y]_{t-h}. \tag{8}$$

2.3 Quadratic Variation as a Volatility Measure

Here we derive specific expressions for the quadratic variation and covariation of arbitrage-free asset prices, and we discuss their use as volatility measures in light of the properties (6)–(8). The additive decomposition (3) and the fact that the predictable components satisfy $[A_k, A_j] = [A_k, M_j] = 0$, for all j and k , imply that

$$\begin{aligned} [p_k, p_j]_t &= [M_k, M_j]_t \\ &= [M_k^c, M_j^c]_t + \sum_{0 \leq s \leq t} \Delta M_k(s) \Delta M_j(s). \end{aligned} \tag{9}$$

We convert this cumulative volatility measure into a corresponding time series of incremental contributions. Letting the integer $h \geq 1$ denote the number of trading days over which the volatility measures are computed, we define the time series of h -period quadratic variation and covariation, for $t = h, 2h, \dots, T$, as

$$\text{Qvar}_{k,h}(t) \equiv [p_k, p_k]_t - [p_k, p_k]_{t-h}, \tag{10a}$$

$$\text{Qcov}_{kj,h}(t) \equiv [p_k, p_j]_t - [p_k, p_j]_{t-h}. \tag{10b}$$

Equation (9) implies that the quadratic variation and covariation for asset prices depend solely on the realization of the return innovations. In particular, the conditional mean is of no import. This renders these quantities model-free: regardless of the specific arbitrage-free price process, the quadratic variation and covariation are obtained by cumulating the instantaneous squares and cross-products of returns, as indicated by (6). Moreover, the measures are well defined even if the price paths contain jumps, as implied by (7), and the quadratic variation is increasing, as required of a cumulative volatility measure.

Equation (8) implies that the h -period quadratic variation and covariation are intimately related to, but distinct from, the conditional return variance and covariance. Specifically,

$$\text{Var}(p_k(t)/\mathcal{F}_{t-h}) = E[\text{Qvar}_{k,h}(t)/\mathcal{F}_{t-h}], \tag{11a}$$

$$\text{Cov}(p_k(t), p_j(t)/\mathcal{F}_{t-h}) = E[\text{Qcov}_{kj,h}(t)/\mathcal{F}_{t-h}]. \tag{11b}$$

Hence, the conditional variance and covariance diverge from the quadratic variation and covariation, respectively, by a zero-mean error. This is natural because the conditional variance and covariance are ex ante concepts, whereas the quadratic variation and covariation are ex post concepts. One can think

of the quadratic variation and covariation as unbiased for the conditional variance and covariance, or conversely. Either way, the key insight is that, unlike the conditional variance and covariance, the quadratic variation and covariation are in principle *observable* via high-frequency returns, which facilitates the analysis and forecasting of volatility using standard statistical tools. Shortly we exploit this insight extensively.

2.4 Approximating the Quadratic Variation and Covariation

Equation (6) implies that we may approximate the quadratic variation and covariation directly from high-frequency return data. In practice, we fix an appropriately high sampling frequency and cumulate the relevant intraday return products over the horizon of interest. Concretely, using the notation in Equation (4) for prices sampled m times per day, we define for $t = h, 2h, \dots, T$,

$$\text{var}_{k,h}(t; m) = \sum_{i=1, \dots, mh} r_{k,(m)}^2(t-h+(i/m)), \tag{12a}$$

$$\begin{aligned} \text{cov}_{kj,h}(t; m) &= \sum_{i=1, \dots, mh} r_{k,(m)}(t-h+(i/m)) \\ &\quad \times r_{j,(m)}(t-h+(i/m)). \end{aligned} \tag{12b}$$

We call the observed measures in (12) the time- t realized h -period volatility and covariance. Note that for any fixed sampling frequency m , the realized volatility and covariance are directly observable, in contrast to their underlying theoretical counterparts, the quadratic variation and covariation processes. For sufficiently large m , however, the realized volatility and covariance provide arbitrarily good approximations to the quadratic variation and covariation, because for all $t = h, 2h, \dots, T$ we have

$$\text{plim}_{m \rightarrow \infty} \text{var}_{k,h}(t; m) = \text{Qvar}_{k,h}(t), \tag{13a}$$

$$\text{plim}_{m \rightarrow \infty} \text{cov}_{kj,h}(t; m) = \text{Qcov}_{kj,h}(t). \tag{13b}$$

Note that the realized volatility measures $\text{var}_{k,h}(t; m)$ and $\text{cov}_{kj,h}(t; m)$ converge as $m \rightarrow \infty$ to $\text{Qvar}_{k,h}(t)$ and $\text{Qcov}_{kj,h}(t)$, but generally do not converge to the corresponding time $t-h$ conditional return volatility or covariance, $E[\text{Qvar}_{k,h}(t)/\mathcal{F}_{t-h}]$ and $E[\text{Qcov}_{kj,h}(t)/\mathcal{F}_{t-h}]$. Standard volatility models focus on the latter, which require a model for the return generating process. Our realized volatility and covariance, in contrast, provide unbiased estimators of the conditional variance and covariance, without taking a stand on any underlying model.

2.5 Integrated Volatility for Itô Processes

Much theoretical work assumes that logarithmic asset prices follow a univariate diffusion. Letting W be a standard Wiener process, we write $dp_k = \mu_k dt + \sigma_k dW$ or, more formally,

$$p_k(t) - p_k(t-1) \equiv r_k(t) = \int_{t-1}^t \mu_k(s) ds + \int_{t-1}^t \sigma_k(s) dW(s). \tag{14}$$

For notational convenience, we suppress the subscript m or h when we consider variables measured over the daily interval ($h = 1$). For example, we have $r_k(t) \equiv r_{k,(1)}(t)$ and $\text{Qcov}_{kj,1}(t) \equiv \text{Qcov}_{kj}(t)$.

Our volatility measure is the associated quadratic variation process. Standard calculations yield

$$\text{Qvar}_k(t) = [p_k, p_k]_t - [p_k, p_k]_{t-1} = \int_{t-1}^t \sigma^2(s) ds. \quad (15)$$

The expression $\int_{t-1}^t \sigma^2(s) ds$ defines the so-called integrated volatility, which is central to the option pricing theory of Hull and White (1987) and further discussed in Andersen and Bollerslev (1998a) and Barndorff-Nielsen and Shephard (1998). They note that, under the pure diffusion assumption, $r_k(t)$ conditional on $\text{Qvar}_k(t)$ is normally distributed with variance $\int_{t-1}^t \sigma^2(s) ds$.

These results extend to the multivariate setting. If $W = (W_1, \dots, W_d)$ is a d -dimensional standard Brownian motion and $(\mathcal{F}_t)_{t \in [0, T]}$ denotes its completed natural filtration. Then, by martingale representation, any locally square integrable price process of the Itô form can be written as (Protter 1990, theorem 4.42),

$$p_k(t) - p_k(0) = \int_0^t \mu_k(s) ds + \sum_{i=1}^d \int_0^t \sigma_{k,i}(s) dW_i(s). \quad (16)$$

This result is related to the fact that any continuous local martingale, H , can be represented as a time change of a Brownian motion, that is, $H(t) = B([H, H]_t)$, a.s. (Protter 1990, theorem 2.41). That flexibility allows this particular specification to cover a large set of applications. Specifically, we obtain

$$\text{Qvar}_k(t) = \sum_{i=1}^d \int_{t-1}^t \sigma_{k,i}^2(s) ds, \quad (17a)$$

$$\text{Qcov}_{kj}(t) = \sum_{i=1}^d \int_{t-1}^t \sigma_{k,i}(s) \sigma_{j,i}(s) ds. \quad (17b)$$

The $\text{Qvar}_k(t)$ expression provides a natural multivariate concept of integrated volatility, and we may correspondingly denote $\text{Qcov}_{kj}(t)$ as the integrated covariance. As a special case of this framework, we may assign a few of the orthogonal Wiener components to be common factors and the remaining components to be pure idiosyncratic error terms. This produces a continuous-time analog to the discrete-time factor volatility models of Diebold and Nerlove (1989) and King, Sentana, and Wadhvani (1994).

Within this pure diffusion setting, stronger results may be obtained. Foster and Nelson (1996) constructed a volatility filter based on a weighted average of past squared returns that extracts the instantaneous volatility perfectly in the continuous record limit. There are two main differences between their approach and ours. From a theoretical perspective, their methods rely critically on the diffusion assumption and extract instantaneous volatility, whereas ours are valid for the entire class of arbitrage-free models, but extract only cumulative volatility over an interval. Second, from an empirical perspective, various market microstructure features limit the frequency at which returns can be productively sampled, which

renders infeasible a Foster–Nelson inspired strategy of extracting instantaneous volatility estimates for a large number of time points within each trading day. Consistent with this view, Foster and Nelson applied their theoretical insights only to the study of volatility filters based on daily data.

The distribution of integrated volatility also has been studied by previous authors, notably, Gallant, Hsu, and Tauchen (1999), who proposed an intriguing reprojection method for direct estimation of the relevant distribution given a specific parametric form for the underlying diffusion, whereas Chernov and Ghysels (2000) applied similar techniques, exploiting options data as well. Our high-frequency return methodology, in contrast, is simpler and more generally applicable, requiring only the special semi-Martingale assumption.

2.6 Volatility Measures for Pure Jump and Mixed Jump–Diffusion Processes

Jump processes have particularly simple quadratic covariation measures. The fundamental semimartingale decomposition (1) reduces to a compensated jump component and a finite variation term,

$$p_k(t) = p_k(0) + M_k(t) + \int_0^t \mu_k(s) ds, \quad (18)$$

where $\mu_k(t)$ denotes the instantaneous mean and the innovations in $M_k(t)$ are pure jumps. The specification covers a variety of scenarios in which the jump process is generated by distinct components,

$$M_k(t) = \sum_{i=1}^J \sum_{0 \leq s \leq t} \kappa_{k,i}(s) \Delta N_{k,i}(s) - \int_0^t \mu_k(s) ds, \quad (19)$$

where $\Delta N_{k,i}(t)$ is an indicator function for the occurrence of a jump in the i th component at time t , while the (random) $\kappa_{k,i}(t)$ term determines the jump size. From property (7),

$$\text{Qcov}_{kj}(t) = \sum_{i=1}^J \sum_{t-1 \leq s \leq t} \kappa_k(s) \kappa_j(s) \Delta N_k(s) \Delta N_j(s). \quad (20)$$

Andersen, Benzoni, and Lund (2000), among others, argued the importance of including both time-varying volatility and jumps when modeling speculative returns over short horizons, which can be accomplished by combining Itô and jump processes into a general jump–diffusion

$$p_k(t) - p_k(0) = \int_0^t \mu_k(s) ds + \int_0^t \sigma_k(s) dW(s) + \sum_{0 \leq s \leq t} k_k(s) \Delta N_k(s). \quad (21)$$

The jump–diffusion allows for a predictable stochastic volatility process $\sigma_k(t)$ and a jump processes, $k_k(t) N_k(t)$ with a finite conditional mean. The quadratic covariation follows directly from Equations (9) and (10):

$$\text{Qcov}_{kj}(t) = \int_{t-1}^t \sigma_k(s) \sigma_j(s) ds + \sum_{t-1 \leq s \leq t} k_k(s) k_j(s) \Delta N_k(s) \Delta N_j(s). \quad (22)$$

It is straightforward to allow for a d -dimensional Brownian motion, resulting in modifications along the lines of Equations (13)–(15), and the formulation readily accommodates multiple jump components, as in (19) and (20).

3. VOLATILITY MEASUREMENT: DATA

Our empirical analysis focuses on the bilateral DM/\$ and yen/\$ spot exchange rates, which are attractive candidates for examination because they represent the two main axes of the international financial system. We first discuss our choice of 5-min returns to construct realized volatilities, and then explain how we handle weekends and holidays. Finally, we detail the actual construction of the volatility measures.

3.1 On the Use of 5-Min Returns

In practice, the discreteness of actual securities prices can render continuous-time models poor approximations at very high sampling frequencies. Furthermore, tick-by-tick prices are generally only available at unevenly spaced time points, so the calculation of evenly spaced high-frequency returns necessarily relies on some form of interpolation among prices recorded around the endpoints of the given sampling intervals. It is well known that this nonsynchronous trading or quotation effect may induce negative autocorrelation in the interpolated return series. Moreover, such market microstructure biases may be exacerbated in the multivariate context if varying degrees of interpolation are employed in the calculation of the different returns.

Hence a tension arises in the calculation of realized volatility. On the one hand, the theory of quadratic variation of special semimartingales suggests the desirability of sampling at very high frequencies, striving to match the ideal of continuously observed frictionless prices. On the other hand, the reality of market microstructure suggests not sampling *too* frequently. Hence a good choice of sampling frequency must balance two competing factors; ultimately it is an empirical issue that hinges on market liquidity. Fortunately, the markets studied in this article are among the most active and liquid in the world, permitting high-frequency sampling without contamination by microstructure effects. We use a sampling frequency of 288 times per day ($m = 288$, or 5-min returns), which is high enough such that our daily realized volatilities are largely free of measurement error (see the calculations in Andersen and Bollerslev 1998a), yet low enough such that microstructure biases are not a major concern.

3.2 Construction of 5-Min DM/\$ and Yen/\$ Returns

The two raw 5-min DM/\$ and yen/\$ return series were obtained from Olsen and Associates. The full sample consists of 5-min returns covering December 1, 1986, through November 30, 1996, or 3,653 days, for a total of $3,653 \times 288 = 1,052,064$ high-frequency return observations. As in Müller et al. (1990) and Dacorogna, Müller, Nagler, Olsen, and Pictet (1993), the construction of the returns utilizes the interbank FX quotes that appeared on Reuter's FXFX page during the sample period. Each quote consists of a bid and an ask price together with a "time stamp" to the nearest even second. After filtering the data for outliers and other anomalies, the price

at each 5-min mark is obtained by linearly interpolating from the average of the log bid and the log ask for the two closest ticks. The continuously compounded returns are then simply the change in these 5-min average log bid and ask prices. Goodhart, Ito, and Payne (1996) and Danielsson and Payne (1999) found that the basic characteristics of 5-min FX returns constructed from quotes closely match those calculated from transactions prices, which are only available on a very limited basis.

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

Finally, we deleted some returns contaminated by brief lapses in the Reuter's data feed. This problem manifests itself in long sequences of zero or constant 5-min returns in places where the missing quotes have been interpolated. To remedy this, we simply removed the days containing the 15 longest DM/\$ zero runs, the 15 longest DM/\$ constant runs, the 15 longest yen/\$ zero runs, and the 15 longest yen/\$ constant runs. Because of the overlap among the four different sets of days defined by these criteria, we actually removed only 51 days. All in all, we were left with 2,449 complete days, or $2,449 \times 288 = 705,312$ 5-min return observations, for the construction of our daily realized volatilities and covariances.

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

We denote the time series of 5-min DM/\$ and yen/\$ returns by $r_{D,(288)}(t)$ and $r_{Y,(288)}(t)$, respectively, where $t = 1/288, 2/288, \dots, 2,449$. We then form the corresponding 5-min squared return and cross-product series $(r_{D,(288)}(t))^2$, $(r_{Y,(288)}(t))^2$, and $r_{D,(288)}(t) \cdot r_{Y,(288)}(t)$. The statistical properties of the squared return series closely resemble those found by Andersen and Bollerslev (1997a,b) with a much shorter 1-year sample of 5-min DM/\$ returns. Interestingly, the basic properties of the 5-min cross-product series, $r_{D,(288)}(t) \cdot r_{Y,(288)}(t)$, are similar. In particular, all three series are highly persistent and display strong intraday calendar effects, the shape of which is driven by the opening and closing of the different financial markets around the globe during the 24-hour trading cycle.

Now, following (12), we construct the realized h -period variances and covariances by summing the corresponding 5-min observations across the h -day horizon. For notational simplicity, we suppress the dependence on the fixed sampling frequency ($m = 288$), and define $vard_{t,h} \equiv \text{var}_{D,h}(t; 288)$, $vary_{t,h} \equiv \text{var}_{y,h}(t; 288)$, and $cov_{t,h} \equiv \text{cov}_{Dy,h}(t; 288)$. Furthermore, for daily measures ($h = 1$), we suppress the subscript h , and simply write $vard_t$, $vary_t$, and cov_t . Concretely, we define for $t = 1, 2, \dots, [T/h]$,

$$vard_{t,h} \equiv \sum_{j=1, \dots, 288 \cdot h} (r_{D,(288)}(h \cdot (t-1) + j/288))^2, \quad (23a)$$

$$vary_{t,h} \equiv \sum_{j=1, \dots, 288 \cdot h} (r_{y,(288)}(h \cdot (t-1) + j/288))^2, \quad (23b)$$

$$cov_{t,h} \equiv \sum_{j=1, \dots, 288 \cdot h} r_{D,(288)}(h \cdot (t-1) + j/288) \cdot r_{y,(288)}(h \cdot (t-1) + j/288). \quad (23c)$$

In addition, we examine several alternative, but related, measures of realized volatility derived from those in (23), including realized standard deviations, $std_{t,h} \equiv vard_{t,h}^{1/2}$ and $stdy_{t,h} \equiv vary_{t,h}^{1/2}$, realized logarithmic standard deviations, $lstd_{t,h} \equiv 1/2 \cdot \log(vard_{t,h})$ and $lsty_{t,h} \equiv 1/2 \cdot \log(vary_{t,h})$, and realized correlations, $corr_{t,h} \equiv cov_{t,h} / (std_{t,h} \cdot stdy_{t,h})$. In Sections 4 and 5 we characterize the unconditional and conditional distributions of the daily realized volatility measures, whereas Section 6 details our analysis of the corresponding temporally aggregated measures ($h > 1$).

4. THE UNCONDITIONAL DISTRIBUTION OF DAILY REALIZED FX VOLATILITY

The unconditional distribution of volatility captures important aspects of the return process, with implications for risk management, asset pricing, and portfolio allocation. Here we provide a detailed characterization.

4.1 Univariate Unconditional Distributions

The first two columns of the first panel of Table 1 provide a standard menu of moments (mean, variance, skewness, and kurtosis) that summarizes the unconditional distributions of the daily realized volatility series, $vard_t$ and $vary_t$, and the top panel of Figure 1 displays kernel density estimates of the unconditional distributions. It is evident that the distributions are very similar and extremely right skewed. Evidently, although the realized daily volatilities are constructed by summing 288 squared 5-min returns, the pronounced heteroscedasticity in intraday returns renders the normal distribution a poor approximation.

The standard deviation of returns is measured on the same scale as the returns, and thus provides a more readily interpretable measure of volatility. We present summary statistics and density estimates for the two daily realized standard deviations, std_t and $stdy_t$, in columns three and four of the first panel of Table 1 and the second panel of Figure 1. The mean daily realized standard deviation is about 68 basis points, and although the right skewness of the distributions has been reduced, the realized standard deviations clearly remain non-normally distributed.

Interestingly, the distributions of the two daily realized logarithmic standard deviations, $lstd_t$ and $lsty_t$, in columns five

Table 1. Statistics That Summarize Unconditional Distributions of Realized DM/\$ and Yen/\$ Volatilities

	$vard_{t,h}$	$vary_{t,h}$	$std_{t,h}$	$stdy_{t,h}$	$lstd_{t,h}$	$lsty_{t,h}$	$cov_{t,h}$	$corr_{t,h}$
Daily, $h = 1$								
Mean	0.529	0.538	0.679	0.684	-0.449	-0.443	0.243	0.435
Variance	0.234	0.272	0.067	0.070	0.120	0.123	0.073	0.028
Skewness	3.711	5.576	1.681	1.867	0.345	0.264	3.784	-0.203
Kurtosis	24.090	66.750	7.781	10.380	3.263	3.525	25.250	2.722
Weekly, $h = 5$								
Mean	2.646	2.692	1.555	1.566	0.399	0.405	1.217	0.449
Variance	3.292	3.690	0.228	0.240	0.084	0.083	0.957	0.022
Skewness	2.628	2.769	1.252	1.410	0.215	0.382	2.284	-0.176
Kurtosis	14.200	14.710	5.696	6.110	3.226	3.290	10.020	2.464
Biweekly, $h = 10$								
Mean	5.297	5.386	2.216	2.233	0.759	0.767	2.437	0.453
Variance	10.440	11.740	0.389	0.403	0.072	0.070	2.939	0.019
Skewness	1.968	2.462	1.063	1.291	0.232	0.380	1.904	-0.147
Kurtosis	7.939	11.980	4.500	5.602	3.032	3.225	7.849	2.243
Triweekly, $h = 15$								
Mean	7.937	8.075	2.717	2.744	0.964	0.977	3.651	0.455
Variance	22.330	22.770	0.560	0.546	0.069	0.064	5.857	0.018
Skewness	2.046	2.043	1.033	1.177	0.208	0.400	1.633	-0.132
Kurtosis	9.408	8.322	4.621	4.756	2.999	3.123	6.139	2.247
Monthly, $h = 20$								
Mean	10.590	10.770	3.151	3.179	1.116	1.127	4.874	0.458
Variance	34.090	36.000	0.671	0.671	0.062	0.059	8.975	0.017
Skewness	1.561	1.750	0.906	1.078	0.295	0.452	1.369	-0.196
Kurtosis	5.768	6.528	3.632	4.069	2.686	2.898	4.436	2.196




Figure 1. Distribution of Daily Realized Exchange Rate Volatilities and Correlations.

and six of the first panel of Table 1 and in the third panel of Figure 1, appear symmetric, with skewness coefficients near zero. Moreover, normality is a much better approximation for these measures than for the realized volatilities or standard

deviations, because the kurtosis coefficients are near 3. This is in accord with the findings for monthly volatility aggregates of daily equity index returns in French, Schwert, and Stambaugh (1987), as well as evidence from Clark (1973) and Taylor (1986).

Finally, we characterize the distribution of the daily realized covariances and correlations, cov_t and $corr_t$, in the last columns of the first panel of Table 1 and the bottom panel of Figure 1. The basic characteristics of the unconditional distribution of the daily realized covariance are similar to those of the daily realized volatilities—they are extremely right skewed and leptokurtic. In contrast, the distribution of the realized correlation is approximately normal. The mean realized correlation is positive (.43), as expected, because both series respond to U.S. macroeconomic fundamentals. The standard deviation of the realized correlation (.17) indicates significant intertemporal variation in the correlation, which may be important for short-term portfolio allocation and hedging decisions.

4.2 Multivariate Unconditional Distributions

The univariate distributions characterized in the foregoing text do not address relationships that may exist among the different measures of variation and covariation. Key issues relevant in financial economic applications include, for example, whether and how $lstdd_t$, $lstdy_t$, and $corr_t$ move together. Such questions are difficult to answer using conventional volatility models, but they are relatively easy to address using our realized volatilities and correlations.

The sample correlations in the first panel of Table 2, along with the $lstdd_t$ - $lstdy_t$ scatterplot in the top panel of Figure 2, 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 so too do their volatilities. This suggests factor structure, as in Diebold and Nerlove (1989) and Bollerslev and Engle (1993).

The correlations in the first panel of Table 2 and the $corr_t$ - $lstdd_t$ and $corr_t$ - $lstdy_t$ scatterplots in the second and third panels of Figure 2 also indicate positive association between correlation and volatility. Whereas some nonlinearity may be operative in the $corr_t$ - $lstdd_t$ relationship, with a flattened response for both very low and very high $lstdd_t$ 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 3 kernel density estimates of $corr_t$ when both $lstdd_t$ and $lstdy_t$ are less than $-.46$ (their median value) and when both $lstdd_t$ and $lstdy_t$ are greater than $-.46$. Similarly, we show in the bottom panel of Figure 3 the estimated $corr_t$ densities conditional on the more extreme volatility situation in which both $lstdd_t$ and $lstdy_t$ are less than $-.87$ (their 10th percentile) and when both $lstdd_t$ and $lstdy_t$

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

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




Figure 2. Bivariate Scatterplots of Realized Volatilities and Correlations.

are greater than .00 (their 90th percentile). It is clear that the distribution of realized correlation shifts rightward when volatility increases. A similar correlation effect in volatility was documented for international equity returns by Solnik, Boucelle, and Le Fur (1996). Of course, given that the high-frequency returns are positively correlated, some such effect is to be expected, as argued by Ronn (1998), Boyer, Gibson, and Loretan (1999), and Forbes and Rigobon (1999). However, the magnitude of the effect nonetheless appears noteworthy.

To summarize, we have documented a substantial amount of variation in volatilities and correlation, as well as important contemporaneous dependence measures. We now turn to dynamics and dependence, which characterize the conditional, as opposed to unconditional, distribution of realized volatility and correlation.

5. THE CONDITIONAL DISTRIBUTION OF DAILY REALIZED FX VOLATILITY

The value of financial derivatives such as options is closely linked to the expected volatility of the underlying asset over the time until expiration. Improved volatility forecasts should, therefore, yield more accurate derivative prices. The conditional dependence in volatility forms the basis for such forecasts. That dependence is most easily identified in the

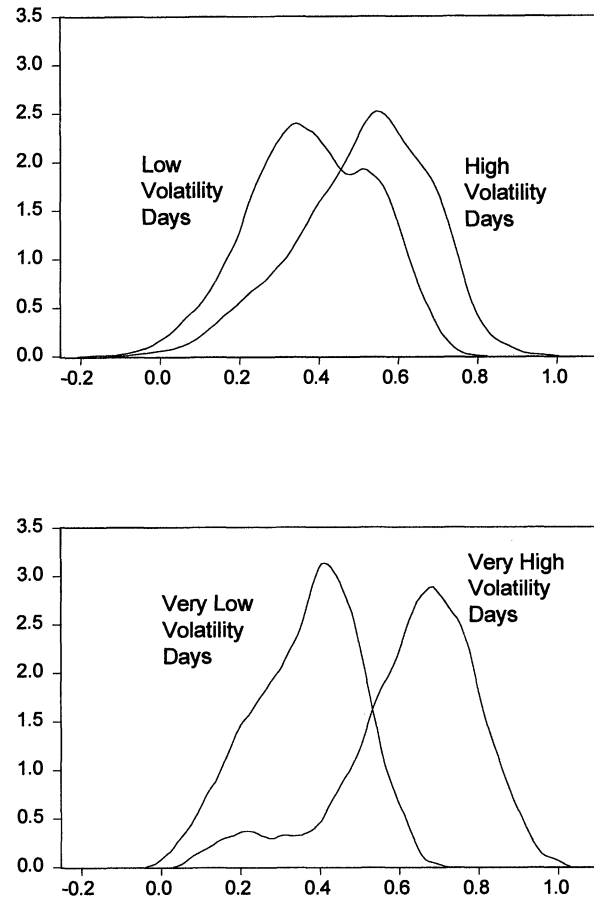


Figure 3. Distributions of Realized Correlations: low volatility versus high volatility days.

daily realized correlations and logarithmic standard deviations, which we have shown to be approximately unconditionally normally distributed. To conserve space, we focus our discussion on those three series.

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

The visual impression of strong persistence in the volatility measures is confirmed by the highly significant Ljung-Box tests reported in the first panel of Table 3. (The .001 critical value is 45.3.) The correlograms of $lstdd_t$, $lstdy_t$, and $corr_t$ in Figure 5 further underscore the point. The autocorrelations of the realized logarithmic standard deviations begin around .6 and decay very slowly to about .1 at a displacement of 100 days. Those of the realized correlations decay even more slowly, reaching just .31 at the 100-day displacement. Similar results based on long series of daily absolute or squared returns from other markets have been obtained previously by a number of authors, including Ding, Granger, and Engle (1993). The slow decay in Figure 5 is particularly noteworthy, however, in that the two realized daily volatility series span "only" 10 years.

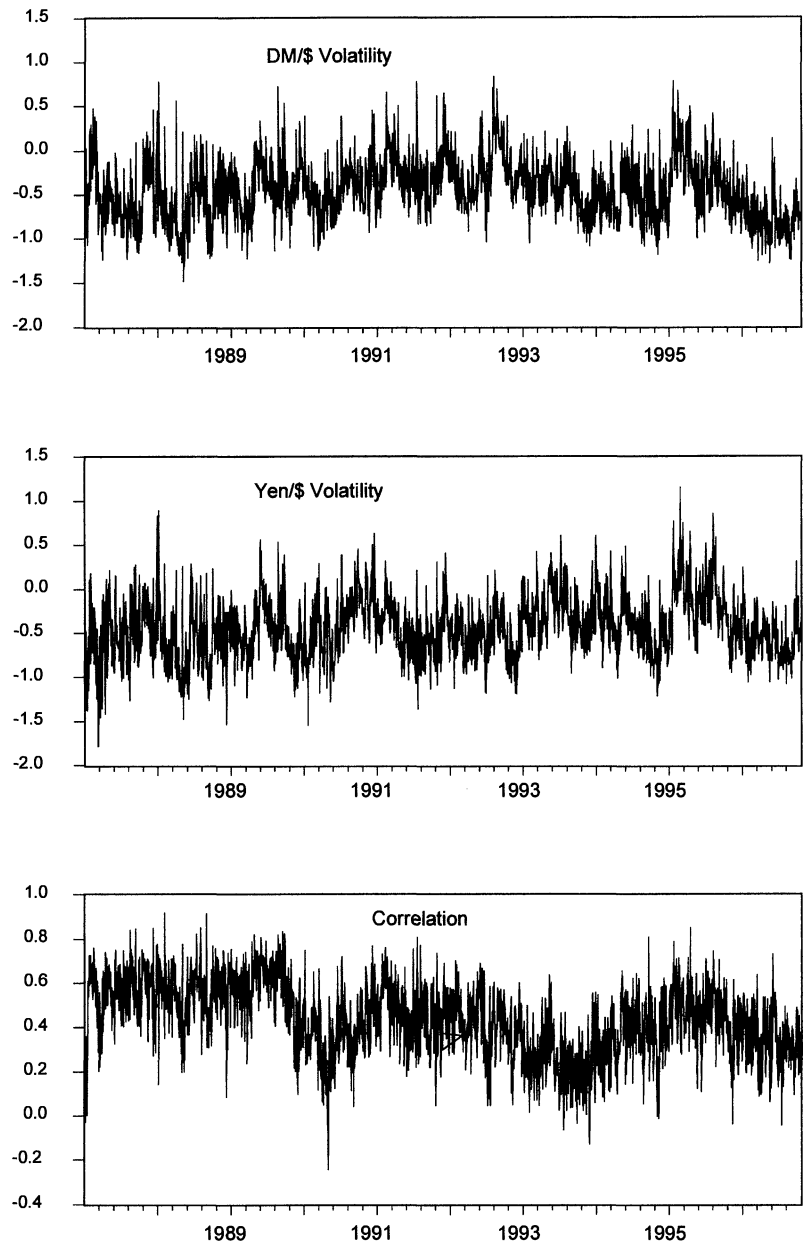


Figure 4. Time Series of Daily Realized Volatilities and Correlation.




Figure 5. Sample Autocorrelations of Realized Volatilities Correlation.

Table 3. Dynamic Dependence Measures for Realized DM/\$ and Yen/\$ Volatilities

	$vard_{t,h}$	$vary_{t,h}$	$stdd_{t,h}$	$stdy_{t,h}$	$lstdd_{t,h}$	$lstdy_{t,h}$	$cov_{t,h}$	$corr_{t,h}$
Daily, $h = 1$								
LB	4539.300	3257.200	7213.700	5664.700	9220.700	6814.100	2855.200	12197.000
\hat{d}	0.356	0.339	0.381	0.428	0.420	0.455	0.334	0.413
Weekly, $h = 5$								
LB	592.700	493.900	786.200	609.600	930.000	636.300	426.100	2743.300
\hat{d}	0.457	0.429	0.446	0.473	0.485	0.496	0.368	0.519
Biweekly, $h = 10$								
LB	221.200	181.000	267.900	206.700	305.300	203.800	155.400	1155.600
\hat{d}	0.511	0.490	0.470	0.501	0.515	0.507	0.436	0.494
Triweekly, $h = 15$								
LB	100.700	108.000	122.600	117.300	138.300	112.500	101.600	647.000
\hat{d}	0.400	0.426	0.384	0.440	0.421	0.440	0.319	0.600
Monthly, $h = 20$								
LB	71.800	69.900	83.100	70.900	94.500	66.000	78.500	427.300
\hat{d}	0.455	0.488	0.440	0.509	0.496	0.479	0.439	0.630

The findings of slow autocorrelation decay may seem to indicate the presence of a unit root, as in the integrated GARCH model of Engle and Bollerslev (1986). However, Dickey–Fuller tests with 10 augmentation lags soundly reject this hypothesis for all of the volatility series. (The test statistics range from -9.26 to -5.59 , and the .01 and .05 critical values are -2.86 and -3.43 .) 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 regressions suggest that long memory of a non-unit-root variety may be present. Hence, we now turn to an investigation of fractional integration in the daily realized volatilities.

As noted by Granger and Joyeux (1980), the slow hyperbolic decay of the long-lag autocorrelations or, equivalently, the log-linear explosion of the low-frequency spectrum are distinguishing features of a covariance stationary fractionally integrated, or $I(d)$, process with $0 < d < \frac{1}{2}$. The low-frequency spectral behavior also forms the basis for the log-periodogram regression estimation procedure proposed by Geweke and Porter-Hudak (1983) and refined by Robinson (1994, 1995), Hurvich and Beltrao (1994) and Hurvich, Deo, and Brodsky (1998). In particular, let $I(\omega_j)$ denote the sample periodogram at the 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

$$\log[I(\omega_j)] = \beta_0 + \beta_1 \cdot \log(\omega_j) + u_j, \quad (24)$$

where $j = 1, 2, \dots, n$, and $\hat{d} \equiv -\hat{\beta}_1/2$. The least squares estimator of β_1 , and hence \hat{d} , is asymptotically normal and the corresponding standard error, $\pi \cdot (24 \cdot n)^{-1/2}$, depends only on the number of periodogram ordinates used. Although the earlier proofs for consistency and asymptotic normality of the log-periodogram regression estimator rely on normality, Deo and Hurvich (1998) and Robinson and Henry (1999) showed that these properties extend to non-Gaussian, possibly heteroscedastic, time series as well. Of course, the actual value of the estimate of d depends upon the choice of n . Although the formula for the theoretical standard error suggests choosing a large n to obtain a small standard error, doing so produces bias in the estimator, because the relationship underlying (24), in general, holds only for frequencies close to

zero. Following Taqqu and Teverovsky (1996), we therefore graphed and examined \hat{d} as a function of n , looking for a flat region in which we are plagued by neither high variance (n too small) nor high bias (n too large). Our subsequent choice of $n = [T^{4/5}]$, or $n = 514$, is consistent with the optimal rate of $O(T^{4/5})$ established by Hurvich, Deo, and Brodsky (1998).

The estimates of d are given in the first panel of Table 3. The estimates are highly statistically significant for all eight volatility series, and all are fairly close to the “typical value” of .4. These estimates for d are also in line with the estimates based on longer time series of daily absolute and squared returns from other markets reported by Granger, Ding, and Spear (1997), and the findings based on a much shorter 1-year sample of intraday DM/\$ returns reported in Andersen and Bollerslev (1997b). This suggests that the continuous-time models used in much of theoretical finance, where volatility is assumed to follow an Ornstein–Uhlenbeck type process, are misspecified. Nonetheless, our results are constructive, in that they also indicate that parsimonious long-memory models should be able to accommodate the long-lag autoregressive effects.

Having characterized the distributions of the daily realized volatilities and correlations, we now turn to longer horizons.

6. TEMPORAL AGGREGATION AND SCALING LAWS

The analysis in the preceding sections focused on the distributional properties of daily realized volatility measures. However, many practical financial problems invariably involve longer horizons. Here we examine the distributional aspects of the corresponding multiday realized variances and correlations. As before, we begin with an analysis of unconditional distributions, followed by an analysis of dynamics and dependence, including a detailed examination of long memory as it relates to temporal aggregation.

6.1 Univariate and Multivariate Unconditional Distributions

In the lower panels of Table 1 we summarize the univariate unconditional distributions of realized volatilities and correlations at weekly, biweekly, triweekly, and monthly horizons

($h = 5, 10, 15$, and 20 , respectively), implying samples of length 489, 244, 163, and 122. Consistent with the notion of efficient capital markets and serially uncorrelated returns, the means of $vard_{t,h}$, $vary_{t,h}$, and $cov_{t,h}$ grow at the rate h , while the mean of the realized correlation, $corr_{t,h}$ is largely invariant to the level of temporal aggregation. 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 we discuss at length subsequently. We also note that even at the monthly level, the unconditional distributions of $vard_{t,h}$, $vary_{t,h}$, and $cov_{t,h}$ remain leptokurtic and highly right skewed. The basic characteristics of $std_{t,h}$ and $stdy_{t,h}$ are similar, with the means increasing at rate $h^{1/2}$. The unconditional variances of $lstdd_{t,h}$ and $lstdy_{t,h}$, however, decrease with h , but again at a rate linked to the fractional integration parameter, as we document subsequently.

Next, turning to the multivariate unconditional distributions, we display in the lower panels of Table 2 the correlation matrices of all volatility measures for $h = 5, 10, 15$, and 20 . Although the correlation between the different measures drops slightly under temporal aggregation, the positive association between the volatilities, so apparent at the 1-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 .604 at the daily horizon to a low of .533 at the monthly horizon. Meanwhile, the volatility effect in correlation is somewhat reduced by temporal aggregation; the sample correlation between $lstdd_{t,1}$ and $corr_{t,1}$ equals .389, whereas the one between $lstdd_{t,20}$ and $corr_{t,20}$ is .245. Similarly, the correlation between $lstdy_{t,h}$ and $corr_{t,h}$ drops from .294 for $h = 1$ to .115 for $h = 20$. Thus, whereas the long-horizon correlations remain positively related to the level of volatility, the lower values suggest that the benefits to international diversification may be the greatest over longer investment horizons.

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

Andersen, Bollerslev, and Lange (1999) showed that, given the estimates obtained at the daily level, the integrated volatility should, in theory, remain strongly serially correlated and highly predictable, even at the monthly level. The Ljung–Box statistics for the realized volatilities in the lower panels of Table 3 provide strong empirical backing. 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 that finds little evidence of volatility clustering for monthly returns, such as Baillie and Bollerslev (1989) and Christoffersen and Diebold (2000). However, the methods and/or data used in the earlier studies may produce tests with low power.

The estimates of d reported in Section 4 suggest that the realized daily volatilities are fractionally integrated. The class of fractionally integrated models is self-similar, so that the degree of fractional integration is invariant to the sampling frequency (see, e.g., Beran 1994). This strong prediction is borne out by the estimates for d for the different levels of temporal aggregation reported in the lower panels of Table 3. All of the estimates are within two asymptotic standard errors of the average estimate of .391 obtained for the daily series,

and all are highly statistically significantly different from both zero and unity.

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

$$[x_t]_h \equiv \sum_{j=1, \dots, h} x_{h \cdot (t-1)+j} \quad (25)$$

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

$$\text{Var}([x_t]_h) = c \cdot h^{2d+1}. \quad (26)$$

Of course, by definition $[vard_t]_h \equiv vard_{t,h}$ and $[vary_t]_h \equiv vary_{t,h}$, so the variance of the realized volatilities should grow at rate h^{2d+1} . This implication is remarkably consistent with the values for the unconditional sample (co)variances reported in Table 1 and a value of d around .35–.40. Similar scaling laws for power transforms of absolute FX returns have been reported in a series of articles initiated by Müller et al. (1990).

The striking accuracy of our scaling laws carries over to the partial sums of the alternative volatility series. The left panel of Figure 6 plots the logarithm of the sample variances of the partial sums of the realized logarithmic standard deviations versus the log of the aggregation level; that is, $\log(\text{Var}([lstdd_t]_h))$ and $\log(\text{Var}([lstdy_t]_h))$ against $\log(h)$ for $h = 1, 2, \dots, 30$. The linear fits implied by (26) are validated. Each of the slopes is very close to the theoretical value of $2d+1$ implied by the log-periodogram estimates for d , further solidifying the notion of long-memory volatility dependence. The estimated slopes in the top and bottom panels are 1.780 and 1.728, respectively, corresponding to values of \hat{d} of .390 and .364.

Because a nonlinear function of a sum is not the sum of the nonlinear function, it is not clear whether $lstdd_{t,h}$ and $lstdy_{t,h}$ will follow similar scaling laws. The estimates of d reported in Table 3 suggest that they should. The corresponding plots for the logarithm of the h -day logarithmic standard deviations $\log(\text{Var}(lstdd_{t,h}))$ and $\log(\text{Var}(lstdy_{t,h}))$ against $\log(h)$, for $h = 1, 2, \dots, 30$, in the right panel of Figure 6, lend empirical support to this conjecture. Interestingly, however, the lines are downward sloped.

To understand why these slopes should be negative, assume that the returns are serially uncorrelated. The variance of the temporally aggregated return should then be proportional to the length of the return interval, that is, $E(\text{var}_{t,h}) = b \cdot h$, where $\text{var}_{t,h}$ refers to the temporally aggregated variance as defined in the preceding text. Also, by the scaling law (26), $\text{Var}(\text{var}_{t,h}) = c \cdot h^{2d+1}$. 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 at all aggregation horizons h with mean μ_h and variance σ_h^2 . Of course, these assumptions agree closely with the actual empirical distributions summarized in Table 1. It then follows from the properties of the log-normal distribution that

$$E(\text{var}_{t,h}) = \exp(2\mu_h + 2\sigma_h^2) = b \cdot h, \quad (27a)$$

$$\begin{aligned} \text{Var}(\text{var}_{t,h}) &= \exp(4\mu_h) \exp(4\sigma_h^2) [\exp(4\sigma_h^2) - 1] \\ &= c \cdot h^{2d+1}, \end{aligned} \quad (27b)$$

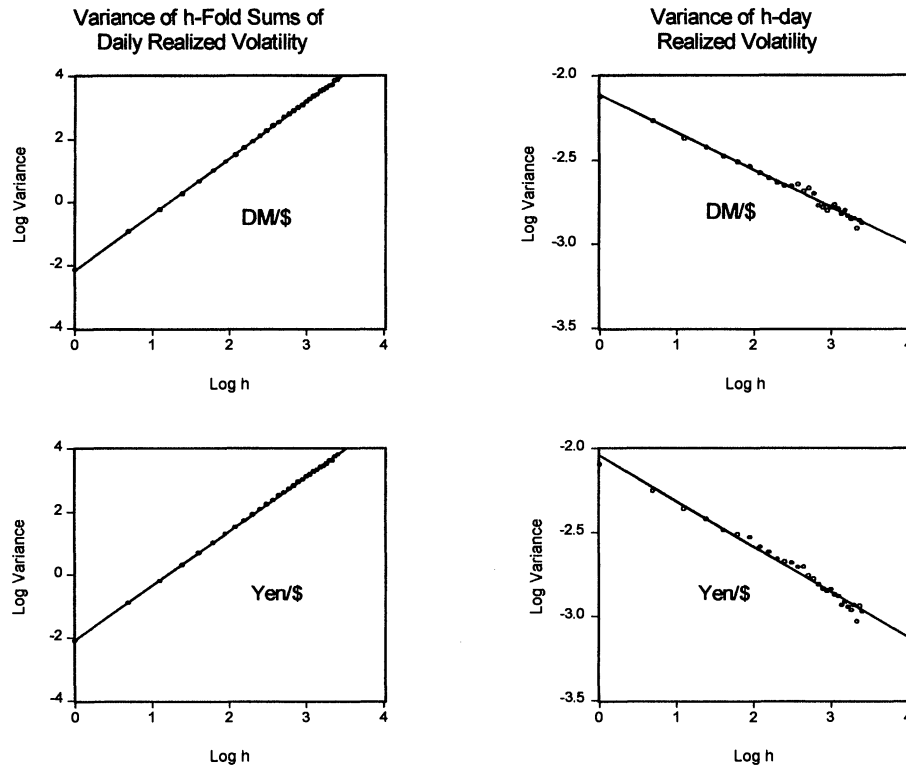


Figure 6. Scaling Laws Under Temporal Aggregation.

and solving for the variance of the log standard deviation yields

$$\text{Var}(\text{lstd}_{t,h}) \equiv \sigma_h^2 = \log(c \cdot b^{-2} \cdot h^{2d-1} + 1). \quad (28)$$

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

$$\log[\text{Var}(\text{lstd}_{t,h})] \approx a + (2d - 1) \cdot \log(h), \quad (29)$$

which explains the apparent scaling law behind the two plots in the right panel of Figure 6, and the negative slopes of approximately $2d - 1$. The slopes in the top and bottom panels are $-.222$ and $-.270$, respectively, and the implied d values of $.389$ and $.365$ are almost identical to the values implied by the scaling law (26) and the two left panels of Figure 6.

7. SUMMARY AND CONCLUDING REMARKS

We first strengthened the theoretical basis for measuring and analyzing time series of realized volatilities constructed from high-frequency intraday returns, and then we put the theory to work, examining a unique dataset that consists of 10 years of 5-min DM/\$ and yen/\$ returns. We found 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. We also found that the correlation between the exchange rates (as opposed to the correlation between their volatilities) increases with volatility. Finally, we confirmed the wealth of existing

evidence of strong volatility clustering effects in daily returns. However, in contrast to earlier work, which often indicates that volatility persistence decreases quickly with the horizon, we find that even monthly realized volatilities remain highly persistent. Nonetheless, realized volatilities do not have unit roots; instead, they appear to be fractionally integrated and, therefore, very slowly mean-reverting. This finding is strengthened by our analysis of temporally aggregated volatility series, whose properties adhere closely to the scaling laws implied by the structure of fractional integration.

A key conceptual distinction between this article 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 postmodeling device to be used to evaluate parametric conditional variance models such as GARCH. As such, it is of interest to examine and model realized volatility directly. This article 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 in turn to use them for forecasting in specific financial contexts. In particular, our findings suggest that a multivariate linear Gaussian long-memory model is appropriate for daily realized logarithmic standard deviations and correlations. Such a model could result in important improvements in the accuracy of volatility and correlation forecasts and related value-at-risk type calculations. This idea is pursued in Andersen, Bollerslev, Diebold, and Labys (2000).

[Received February 1999. Revised October 2000.]

REFERENCES

- Andersen, T. G., Benzoni, L., and Lund, J. (2000), "Estimating Jump-Diffusions for Equity Returns," unpublished manuscript, Northwestern University, Dept. of Finance, J.L. Kellogg Graduate School of Management.
- Andersen, T. G., and Bollerslev, T. (1997a), "Intraday Periodicity and Volatility Persistence in Financial Markets," *Journal of Empirical Finance*, 4, 115–158.
- (1997b), "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns," *Journal of Finance*, 52, 975–1005.
- (1998a), "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885–905.
- (1998b), "DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer-Run Dependencies," *Journal of Finance*, 53, 219–265.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2000), "Modeling and Forecasting Realized Volatility," unpublished manuscript, Northwestern University, Duke University, and University of Pennsylvania.
- Andersen, T. G., Bollerslev, T., and Lange, S. (1999), "Forecasting Financial Market Volatility: Sample Frequency vis-à-vis Forecast Horizon," *Journal of Empirical Finance*, 6, 457–477.
- Back, K. (1991), "Asset Prices for General Processes," *Journal of Mathematical Economics*, 20, 317–395.
- Baillie, R. T., and Bollerslev, T. (1989), "The Message in Daily Exchange Rates: A Conditional Variance Tale," *Journal of Business and Economic Statistics*, 7, 297–305.
- Barndorff-Nielsen, O. E., and Shephard, N. (1998), "Aggregation and Model Construction for Volatility Models," unpublished manuscript, Nuffield College, Oxford, UK.
- Beran, J. (1994), *Statistics for Long-Memory Processes*, New York: Chapman and Hall.
- Bollerslev, T., and Domowitz, I. (1993), "Trading Patterns and Prices in the Interbank Foreign Exchange Market," *Journal of Finance*, 48, 1421–1443.
- Bollerslev, T., and Engle, R. F. (1993), "Common Persistence in Conditional Variances," *Econometrica*, 61, 166–187.
- Bollerslev, T., Engle, R. F., and Nelson, D. B. (1994), "ARCH Models," in *Handbook of Econometrics* (Vol. IV), eds. R. F. Engle and D. McFadden, Amsterdam: North-Holland, pp. 2959–3038.
- Boyer, B. H., Gibson, M. S., and Loretan, M. (1999), "Pitfalls in Tests for Changes in Correlations," IFS Discussion Paper No. 597, Federal Reserve Board.
- Chernov, M., and Ghysels, E. (2000), "A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation," *Journal of Financial Economics*, 56, 407–458.
- Christoffersen, P. F., and Diebold, F. X. (2000), "How Relevant is Volatility Forecasting for Risk Management?," *Review of Economics and Statistics*, 82, 12–22.
- Clark, P. K. (1973), "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices," *Econometrica*, 41, 135–155.
- Comte, F., and Renault, E. (1998), "Long Memory in Continuous-Time Stochastic Volatility Models," *Mathematical Finance*, 8, 291–323.
- Dacorogna, M. M., Müller, U. A., Nagler, R. J., Olsen, R. B., and Pictet, O. V. (1993), "A Geographical Model for the Daily and Weekly Seasonal Volatility in the Foreign Exchange Market," *Journal of International Money and Finance*, 12, 413–438.
- Danielsson, J., and Payne, R. (1999), "Real Trading Patterns and Prices in Spot Foreign Exchange Markets," unpublished manuscript, London School of Economics, Financial Markets Group.
- Deo, R., and Hurvich, C. M. (1998), "On the Log Periodogram Regression Estimator of the Memory Parameter in Long Memory Stochastic Volatility Models," *Econometric Theory*, forthcoming.
- Diebold, F. X., and Lindner, P. (1996), "Fractional Integration and Interval Prediction," *Economics Letters*, 50, 305–313.
- Diebold, F. X., and Nerlove, M. (1989), "The Dynamics of Exchange Rate Volatility: A Multivariate Latent Factor ARCH Model," *Journal of Applied Econometrics*, 4, 1–22.
- Ding, Z., Granger, C. W. J., and Engle, R. F. (1993), "A Long Memory Property of Stock Market Returns and a New Model," *Journal of Empirical Finance*, 1, 83–106.
- Engle, R. F., and Bollerslev, T. (1986), "Modeling the Persistence of Conditional Variances," *Econometric Reviews*, 5, 1–50.
- Forbes, K., and Rigobon, R. (1999), "No Contagion, Only Interdependence: Measuring Stock Market Co-Movements," Working Paper 7267, National Bureau of Economic Research.
- Foster, D. P., and Nelson, D. B. (1996), "Continuous Record Asymptotics for Rolling Sample Variance Estimators," *Econometrica*, 64, 139–174.
- French, K. R., Schwert, G. W., and Stambaugh, R. F. (1987), "Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3–29.
- Gallant, A. R., Hsu, C.-T., and Tauchen, G. E. (1999), "Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance," *Review of Economics and Statistics*, 81, 617–631.
- Geweke, J., and Porter-Hudak, S. (1983), "The Estimation and Application of Long Memory Time Series Models," *Journal of Time Series Analysis*, 4, 221–238.
- Goodhart, C., Ito, T., and Payne, R. (1996), "One Day in June 1993: A Study of the Working of the Reuters 2000–2 Electronic Foreign Exchange Trading System," in *The Microstructure of Foreign Exchange Markets*, eds. J. A. Frankel, G. Galli, and A. Giovannini, Chicago: University of Chicago Press, pp. 107–179.
- Granger, C. W. J., Ding, Z., and Spear, S. (1997), "Stylized Facts on the Temporal and Distributional Properties of Daily Data From Speculative Markets," unpublished manuscript, University of California, San Diego, Dept. of Economics.
- Granger, C. W. J., and Joyeux, R. (1980), "An Introduction to Long Memory Time Series Models and Fractional Differencing," *Journal of Time Series Analysis*, 1, 15–39.
- Hsieh, D. A. (1991), "Chaos and Nonlinear Dynamics: Application to Financial Markets," *Journal of Finance*, 46, 1839–1877.
- Hull, J., and White, A. (1987), "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42, 381–400.
- Hurvich, C. M., and Beltrao, K. I. (1994), "Automatic Semiparametric Estimation of the Memory Parameter of a Long-Memory Time Series," *Journal of Time Series Analysis*, 15, 285–302.
- Hurvich, C. M., Deo, R., and Brodsky, J. (1998), "The Mean Squared Error of Geweke and Porter Hudak's Estimator of the Memory Parameter of a Long-Memory Time Series," *Journal of Time Series Analysis*, 19, 19–46.
- King, M., Sentana, E., and Wadhvani, S. (1994), "Volatility and Links Between National Markets," *Econometrica*, 62, 901–933.
- Meheswaran, S., and Sims, C. A. (1993), "Empirical Implications of Arbitrage-Free Asset Markets," in *Models, Methods and Applications of Econometrics: Essays in Honor of A. R. Bergstrom*, ed. P. C. B. Phillips, Cambridge, MA: Blackwell.
- Müller, U. A., Dacorogna, M. M., Olsen, R. B., Pictet, O. V., Schwarz, M., and Morgenege, C. (1990), "Statistical Study of Foreign Exchange Rates, Empirical Evidence of a Price Change Scaling Law, and Intraday Analysis," *Journal of Banking and Finance*, 14, 1189–1208.
- Protter, P. (1990), *Stochastic Integration and Differential Equations: A New Approach*, New York: Springer-Verlag (third printing, 1995).
- Robinson, P. M. (1994), "Semiparametric Analysis of Long-Memory Time Series," *The Annals of Statistics*, 22, 515–539.
- (1995), "Log-Periodogram Regression of Time Series with Long-Range Dependence," *The Annals of Statistics*, 23, 1048–1072.
- Robinson, P. M., and Henry, M. (1999), "Long and Short Memory Conditional Heteroskedasticity in Estimating the Memory Parameter of Levels," *Econometric Theory*, 19, 299–336.
- Rogers, L. C. G. (1997), "Arbitrage with Fractional Brownian Motion," *Mathematical Finance*, 7, 95–105.
- Ronn, E. (1998), "The Impact of Large Changes in Asset Prices on Intra-Market Correlations in the Stock and Bond Markets," unpublished manuscript, University of Texas, Austin, Dept. of Finance.
- Schwert, G. W. (1989), "Why does Stock Market Volatility Change Over Time?," *Journal of Finance*, 44, 1115–1154.
- (1990), "Stock Market Volatility," *Financial Analysts Journal*, May–June, 23–34.
- Solnik, B., Boucrelle, C., and Le Fur, Y. (1996), "International Market Correlation and Volatility," *Financial Analysts Journal*, September–October, 17–34.
- Taqqu, M. S., and Teverovsky, V. (1996), "Semi-Parametric Graphical Estimation Techniques for Long-Memory Data," in *Time Series Analysis in Memory of E. J. Hannan*, eds. P. M. Robinson and M. Rosenblatt, New York: Springer-Verlag, pp. 420–432.
- Taylor, S. J. (1986), *Modelling Financial Time Series*, Chichester: Wiley.
- Taylor, S. J., and Xu, X. (1997), "The Incremental Volatility Information in One Million Foreign Exchange Quotations," *Journal of Empirical Finance*, 4, 317–340.