

Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns

TORBEN G. ANDERSEN and TIM BOLLERSLEV*

ABSTRACT

Recent empirical evidence suggests that the interdaily volatility clustering for most speculative returns are best characterized by a slowly mean-reverting fractionally integrated process. Meanwhile, much shorter lived volatility dynamics are typically observed with high frequency intradaily returns. The present article demonstrates, that by interpreting the volatility as a mixture of numerous heterogeneous short-run information arrivals, the observed volatility process may exhibit long-run dependence. As such, the long-memory characteristics constitute an intrinsic feature of the return generating process, rather than the manifestation of occasional structural shifts. These ideas are confirmed by our analysis of a one-year time series of five-minute Deutschemark-U.S. Dollar exchange rates.

Among the most puzzling issues is the behavior of volatility. While the general properties of volatility remain elusive, perhaps the most intriguing feature revealed by empirical work on volatility is its long persistence. Such behavior has sparked a search, almost akin to that for the Holy Grail, for the perfect GARCH model, but the underlying question of why such volatility persistence endures remains unanswered. We conjecture that the ability to analyze higher frequency data may be particularly useful in pursuing this issue. (Goodhart and O'Hara (1997))

THE PRONOUNCED VOLATILITY CLUSTERING is arguably one of the most striking features of financial price series recorded at daily or weekly intervals, and a large body of literature seeking to characterize this aspect of speculative returns has emerged over the past decade.¹ At the same time, convincing theoretical explanations for the underlying sources of long-run persistence in volatility remain elusive. The recent advent of comprehensive high-frequency

* Andersen is from the J.L. Kellogg Graduate School of Management of Northwestern University. Bollerslev is from the University of Virginia and is a research associate at the National Bureau of Economic Research. We gratefully acknowledge the financial support provided by a research grant from the Institute for Quantitative Research in Finance (the Q-Group). Special thanks are due to Olsen and Associates for making the intradaily exchange rate quotes and Reuter's News Tape available. We have received valuable comments from Wake Epps, Clive W.J. Granger, Stephen F. Gray, J. Huston McCulloch, as well as seminar participants at the December 1996 Triangle Econometrics Workshop, the 1997 AFA Meetings in New Orleans, Academia Sinica, the Wharton School, USC, and Georgetown and York universities. We remain fully responsible for the content.

¹ See Bollerslev, Chou, and Kroner (1992) for a survey of this literature using ARCH type models.

data sets has added substantially to our knowledge of the higher frequency features of the returns process, but the findings have tended to deepen rather than resolve the puzzle surrounding the persistence in the volatility dynamics. In particular, while studies using daily or lower frequency returns generally point to a very high degree of intertemporal volatility dependence, the dynamic impact of identifiable intraday volatility shocks seem rather low, with any noticeable effects on the overall volatility typically gone in a matter of hours or less.²

These findings have led to a certain fragmentation of the literature, with most market microstructure studies focusing predominantly on the highly significant patterns in the intradaily price movements (time-of-day effects), ignoring the longer-run inter-daily volatility dependencies. This separation may appear entirely natural, if not inevitable. High-frequency returns generally cover only shorter time-spans, and given the evidence of near covariance non-stationarity of the return volatility process, as indicated by the findings of (near) integrated general autoregressive conditional heteroskedasticity (IGARCH) type models (Bollerslev and Engle (1993)), it may seem futile to search for insights regarding volatility persistence over a short horizon. This intuition is further bolstered by the lessons from the extant time series literature, which has documented that only very long time-spans, as opposed to more frequent observations, can provide genuine information about the presence of a unit root in a borderline stationary series.³ At the same time, the assumption of structural stability, or invariance of the return generating process, may be unreasonable for very long time-spans. Specifically, if a structural break is present, and ignored, then the appearance of strong persistence may emerge artificially. This issue has been hotly debated in the context of modeling the level of economic time series and has led to the development of an extensive literature on testing for structural breaks; see e.g., Perron (1989). The same idea carries over to the volatility setting; see e.g., Lamoureux and Lastrapes (1990).

This article advocates an approach that speaks directly to this important issue, while avoiding the analysis of excessively long time-spans of return data. The basic premise holds that shorter time-spans of high frequency returns—in contrast to widely, if implicitly, held notions—are in fact very informative about longer-run volatility dependencies. The rationale is closely related to the observation of Merton (1980), that whereas it is not possible to learn much about the drift (mean) of the price process by sampling more frequently, high frequency observations may greatly enhance estimates of the volatility (variance). In fact, in a strict diffusion setting volatility can be measured perfectly by sampling the price process continuously; see, e.g., Nelson (1990, 1992). In reality, market microstructure features such as the bid-ask spread and discrete price quotes, combined with the presence of strong intraday volatility patterns render this infeasible. Moreover, these short-run

² For a recent survey of this literature see Goodhart and O'Hara (1997).

³ Shiller and Perron (1985) provide some of the first simulation based evidence along these lines.

effects are generally large relative to the return component that provides information regarding the innovation in the long-run volatility. However, these effects may be controlled for in ways that allows an investigation of the long-run volatility features based on high-frequency return series covering only a relatively short time-span.

Our empirical findings along these lines clearly points towards the existence of long-run volatility persistence in high frequency returns.⁴ Meanwhile, the empirical analysis also provides distinct evidence for the existence of multiple volatility components at the intradaily frequencies. This evidence of multiple volatility components, along with the apparent long-memory characteristics, motivates our theoretical developments, in which we formulate a version of the mixture-of-distributions hypothesis (MDH) for returns that explicitly accommodates numerous heterogeneous information arrival processes. Even though each of the constituent information flow processes exhibit only short-run memory, we demonstrate how the aggregate volatility process may display the dependency associated with a long-memory process. In addition, we show that this degree of volatility persistence should be invariant to temporal aggregation.

The MDH interpretation of market volatility as resulting from the aggregation of numerous components type processes apply equally well across most financial markets and instruments. However, for concreteness the empirical analysis in the article is focused on the foreign exchange market, and a one-year time-span of five-minute Deutschemark–U.S. Dollar (DM–\$) returns. While our empirical analysis is performed predominantly in the frequency domain, we also develop a low-pass filtering technique that purges all intraday volatility components with periodicities of less than one day. As such, the low-pass filter is designed to annihilate the strong intraday patterns, while retaining all the low-frequency information that pertains to the interdaily frequencies. The resulting high-frequency series thus speaks directly to the long-run features of the return series, and affords a direct time-domain analysis of the volatility persistence from the high frequency data. We verify that this approach yields results that are consistent with more standard frequency-based procedures, and also illustrate the usefulness of the filtered series for direct analysis of the longer run volatility implications of different macroeconomic announcements. The fact that our estimates for the degree of long-run volatility persistence are very similar across all intradaily sampling frequencies, and strikingly close to the estimates for a daily time series covering a much longer time-span, lend further support to the notion of long-memory dependence as an inherent feature of the return generating process, rather than an artifact of infrequent structural shifts, or changes in regimes.

The plan for the remainder of the article is as follows. Section I briefly reviews some of the existing literature on the intradaily dependencies in the

⁴ This is consistent with the recent findings of long-memory in the volatility of different daily return series, as formally modeled by Baillie, Bollerslev, and Mikkelsen (1996) and Ding and Granger (1996), among others.

foreign exchange market. This section also presents preliminary estimation results for specific announcement effects that point to the existence of multiple volatility components and long-run dependencies. The mixture-of-distributions hypothesis and the temporal aggregation argument are presented in Section II, where we show that, under suitable assumptions concerning the underlying component structure, the spectrum for the overall volatility process implies an eventual slow hyperbolic rate of decay in the autocorrelations for the absolute returns, irrespective of the sampling frequency. The theoretical developments are validated by the empirical analysis in Section III, which reports the results obtained from two alternative spectral-domain estimators for the degree of fractional integration in the volatility process across a range of different intradaily sampling frequencies. These findings are entirely consistent with the time-domain estimate for the fractional order of integration determined by the hyperbolic rate of decay in the autocorrelations for the low-pass filtered absolute returns presented in Section IV. This section also illustrates how the low-pass filtered absolute returns allow for important new insights into the structure behind the determination of the observed volatility. Section V concludes.

I. Preliminary Data Analysis

In order to motivate the subsequent theoretical developments, the following section describes the salient empirical features of the volatility in the DM-\$ foreign exchange market. However, the general ideas and empirical results apply equally well to other financial markets and high frequency return series; see, e.g., Andersen and Bollerslev (1997a) for a discussion of the parallels between the empirical properties of high frequency foreign exchange and equity index returns. The DM-\$ exchange rate data under study consist of all the quotes that appeared on the interbank Reuters network during the October 1, 1992 through September 29, 1993 sample period. The data were collected and provided by Olsen and Associates. Each quote contains a bid and an ask price along with the time to the nearest even second. The exchange rate corresponding to the endpoint of a given five-minute interval was determined as the interpolated average between the preceding and immediately following quotes weighted linearly by their inverse relative distance to this endpoint. The n th five-minute return for day τ , $R_{\tau,n}$, is then simply defined as the difference between the midpoint of the logarithmic bid and ask at the appropriately spaced time intervals. All 288 five-minute returns during the 24-hour daily trading cycle are used, but in order to avoid confounding the evidence by the decidedly slower trading patterns over the weekends, all returns from Friday 21:00 Greenwich Mean Time (GMT) through Sunday 21:00 GMT were excluded; see Bollerslev and Domowitz (1993) for a detailed analysis of the quote activity in the DM-\$ interbank market and a justification for this "weekend" definition. Similarly, to preserve the number of returns associated with one week we make no corrections for any worldwide or country specific Holidays that occurred during the sample period. All in all, this leaves us with

a sample of 260 days, for a total of 74,880 five-minute intraday return observations; i.e., R_t , $t = 1, 2, \dots, 74,880$, where $R_{(\tau-1) \cdot 288+n} \equiv R_{\tau,n}$, for $n = 1, 2, \dots, 288$, and $\tau = 1, 2, \dots, 260$. For further discussion of the data construction we refer to Andersen and Bollerslev (1997a), where the same dataset has previously been analyzed from a different perspective.

Aside from the numerically small, but statistically significant, negative first order autocorrelation coefficient of -0.040 , the five-minute returns appear to be well approximated by a martingale process.⁵ However, the returns are clearly not independent, as the first order autocorrelation coefficient for $|R_t|$ equals 0.309 . Also, the Ljung–Box statistic for up to tenth order serial correlation in the absolute returns takes a value of $36,680$, which is highly significant in the corresponding asymptotic chi-square distribution.

While the latter statistics are strongly suggestive of intraday volatility clustering, their overwhelming significance is in part attributable to the strong intradaily volatility pattern that is present in most financial markets. The importance of this phenomenon for the DM–\$ rates is obvious from Figure 1, which graphs an estimate of the spectrum for the intradaily absolute five-minute returns.⁶ The spectrum has a distinct peak at the daily frequency of approximately $2 \cdot \pi/288 \approx 0.0218$, along with well defined peaks at the corresponding seasonal harmonics of $4 \cdot \pi/288 \approx 0.0436$, $6 \cdot \pi/288 \approx 0.0654$, $8 \cdot \pi/288 \approx 0.0873$. This pronounced daily periodicity is also evident in the autocorrelogram for the absolute 5-minute returns depicted out to a lag of 2880, or ten days, in Figure 2. The intradaily pattern induces a distorted U-shape in the sample autocorrelations across each day, with the autocorrelations spaced half-a-day apart actually turning negative. The shape of the underlying intradaily volatility pattern over the 24-hour trading cycle in the foreign exchange market, and the close connection to the market activity in the various financial centers around the world, have previously been documented by Andersen and Bollerslev (1997a), Baillie and Bollerslev (1991), Dacorogna *et al.* (1993), Ito, Lyons, and Melvin (1996), among others.⁷

Taken together, these results clearly point to the importance of appropriately accounting for the strong intradaily volatility patterns when analyzing

⁵ The small negative autocorrelation may be explained by the asymmetric positioning of quotes by the foreign exchange dealers as a way to manage their inventory positions, thus causing the midpoint of the quoted prices to move around in a fashion analogous to the well documented bid–ask bounce effect on organized exchanges. This is also consistent with the results in Bollerslev and Domowitz (1993), who find similarly constructed five-minute DM–\$ returns over a 3-month period in 1989 to be negatively correlated, while the first order autocorrelation for artificial constructed five-minute pseudo transactions price returns is actually positive. Similarly, the foreign exchange transactions prices analyzed by Goodhart, Ito, and Payne (1996) show no negative correlation between subsequent returns.

⁶ The spectrum is estimated as the smoothed sample periodogram based on a triangular kernel with a bandwidth of ten; see, e.g., Hamilton (1994) or Priestley (1981). To allow for easier interpretation of the subsequent results, the plot is given on a double logarithmic scale.

⁷ Equally pronounced patterns have been documented for other financial markets, including the well known intradaily U-shape for equity market volatility; see e.g., Wood, McInish, and Ord (1985) and Harris (1986).




Figure 1. The figure graphs the spectrum for the five-minute absolute returns on an open position in the Deutschemark–U.S. Dollar spot exchange rate market from October 1, 1992 through September 29, 1993; i.e., $|R_t| \equiv |\ln(P_t) - \ln(P_{t-1})|$ where $t = 1, 2, \dots, 74,880$. The spectrum is estimated as the smoothed sample periodogram using a triangular kernel with bandwidth of ten.

high frequency dynamic dependencies, both within and across different markets. However, even after explicitly modeling the typical intradaily patterns, several recent studies have found that the estimation of standard GARCH type models for the nonperiodic intradaily volatility clustering tends to falter, in the sense that the estimated parameters obtained for different intradaily sampling frequencies are at odds with the temporal aggregation results developed by Drost and Nijman (1993) and the continuous record asymptotic in Nelson (1990); for existing evidence along these lines pertaining to the DM–\$ rates analyzed here, see Andersen and Bollerslev (1997a). Thus, it appears that multiple volatility components are necessarily required in order to fully explain the complex intradaily dependencies that are present in all the major financial markets.⁸

The attempt to associate each of the underlying volatility components with explicit economic factors seems destined to fail. Meanwhile, the regularly scheduled releases of macroeconomic figures have been shown to induce heightened overall volatility immediately following the “news” release for a number of different financial instruments; for studies on the impact of macroeconomic announcements on high-frequency foreign exchange rates see e.g.,

⁸ Multiple components ARCH models have been proposed by Müller *et al.* (1997) and Ghose and Kroner (1996) to capture this phenomenon in the modeling of high-frequency foreign exchange returns.




Figure 2. The figure graphs the sample autocorrelogram for the five-minute absolute returns on an open position in the Deutschemark–U.S. Dollar spot exchange rate market from October 1, 1992 through September 29, 1993; i.e., $|R_t| = |\ln(P_t) - \ln(P_{t-1})|$ where $t = 1, 2, \dots, 74,880$.

Andersen and Bollerslev (1997b), Ederington and Lee (1993), Goodhart *et al.* (1993), Hogan and Melvin (1994), and Ito and Roley (1987). Not surprisingly, the actual effect tends to vary according to the type of announcement. From the analysis in Andersen and Bollerslev (1997b) the three regularly scheduled macroeconomic announcements with the largest instantaneous impact on DM–\$ volatility during the present sample period are the U.S. Employment Report, the biweekly Bundesbank meetings, and the U.S. Durable Goods figures.⁹ To gauge the dynamic impact of each of these events, we consider the following simple AR(1) model for the standardized five-minute DM–\$ absolute returns, augmented by a separate short-lived AR(1) component,

$$|\tilde{R}_t| = \gamma + \gamma_N \cdot I_N(t) + \phi \cdot |\tilde{R}_{t-1}| + \phi_N \cdot D_N(t) \cdot |\tilde{R}_{t-1}| + \varepsilon_t, \quad (1)$$

where the $I_N(t)$ indicator variable equals unity if an announcement of the given type, N , occurred during the t th time interval and zero otherwise, while the “news” dummy, $D_N(t)$, equals unity for the two hours immediately following the event; i.e., $D_N(t) \equiv I_N(t) + I_N(t - 1) + \dots + I_N(t - 23)$. In order to avoid confounding the dynamic dependencies by the strong intradaily volatility patterns, the returns in equation (1) are standardized by the average absolute return for the particular five-minute interval; i.e., $|\tilde{R}_t| \equiv$

⁹ The U.S. Employment Report, Durable Goods figures, and Business Inventories analyzed below, are all announced on a monthly schedule.

$260 \cdot |R_{(\tau-1) \cdot 288+n}| \cdot (\sum_{\tau=1}^{260} |R_{(\tau-1) \cdot 288+n}|)^{-1}$ for $n = 1, 2, \dots, 288$, and $\tau = 1, 2, \dots, 260$.¹⁰ The first order sample autocorrelation for these standardized returns, i.e., $I_N(t)$ and $D_N(t) \equiv 0$ for all t , equals $\hat{\phi} = 0.273$. However, the estimates obtained from the simple descriptive model in equation (1) suggest that the three major announcements result in their own distinct volatility response pattern. Specifically, for the Employment Report $\hat{\gamma}_N = 5.346$ (1.424) and $\phi_N = -0.081$ (0.050), respectively, where the numbers in parentheses represent robust standard errors. These estimates therefore indicate an immediate increase in the volatility, but largely unaltered short-run dynamic dependencies in the two hours following the release of the report. Meanwhile, the estimates for the Bundesbank biweekly meetings are 1.394 (0.651) and 0.102 (0.054), suggestive of a slower rate of decay in the volatility component associated with this "news" release. The Durable Goods report also elevates volatility, $\hat{\gamma}_N = 1.801$ (0.687), but the estimated short-run decay is actually faster than average with $\hat{\phi}_N = -0.130$ (0.037), or $\hat{\phi} + \hat{\phi}_N = 0.144$. Many other announcements share this dynamic pattern. For instance, for the Business Inventory figures the same two parameters are 0.021 (0.110) and -0.094 (0.045), respectively.

This relatively rapid decay of the readily identifiable "news" effects at the intradaily frequencies is in sharp contrast to the well documented highly persistent volatility clustering that is present at the lower interdaily sampling frequencies. In fact, the daily dependencies have often been modeled by an integrated autoregressive conditional heteroskedasticity (ARCH) variance process; see Bollerslev, Chou, and Kroner (1992) for a survey of this voluminous literature. More recently however, Baillie, Bollerslev, and Mikkelsen (1996), Dacorogna *et al.* (1993), and Harvey (1994), have questioned these earlier findings, on the grounds that the long-run dependence in the volatility of interdaily foreign exchange returns is better characterized by a slowly mean-reverting fractionally integrated process.¹¹ While shocks to the daily volatility process are highly persistent, it appears that they eventually dissipate, albeit at a very slow hyperbolic rate. As discussed in Section III, the approximate linear behavior of the estimated spectrum for $|R_t|$, when plotted on the double logarithmic scale in Figure 1, is consistent with this long-memory type dependence. Similarly, the rapid initial decay followed by the decidedly slower rate of dissipation for the longer-run autocorrelations in Figure 2 also indicate the existence of long-memory type dependencies, although the marked intradaily periodicity severely complicates the interpretation of the overall pattern in the autocorrelogram for the raw absolute returns.

¹⁰ Much more elaborate standardization procedures for characterizing the average intradaily volatility patterns have recently been developed by Andersen and Bollerslev (1997a, 1997b) and Dacorogna *et al.* (1993).

¹¹ Similar long-run dependencies in daily stock market volatility have been observed by Bollerslev and Mikkelsen (1996), Breidt, Crato, and de Lima (1997), Ding, Granger, and Engle (1993), Granger and Ding (1996, 1997), among others.

II. The Theoretical Volatility Model

This section demonstrates that the mixture-of-distribution hypothesis, originally advocated by Clark (1973), Epps and Epps (1976), and Tauchen and Pitts (1983) provides a framework for rationalizing the complicated, and seemingly conflicting, behavior of the return volatility dynamics that exist at the inter- and intradaily frequencies; see also the analyses in Andersen (1996), Harris (1987), Ross (1989), and Taylor (1994). Intuitively, the mixture-of-distributions hypothesis stipulates that the return generating process reflects the impact of a large number of innovations to the information processes associated with the economic factors of relevance for the valuation of the asset in question. While the innovations, by definition, are serially uncorrelated, they are not likely to be independent since information of a particular kind tends to be positively autocorrelated, thus inducing the kind of dependence in the absolute, or second order, moments of the returns discussed above.

To formalize these ideas, consider the following representation for the intradaily returns,

$$R_t = r_t - m_t = V_t^{1/2} \cdot Z_t, \tag{2}$$

where m_t denotes the conditional mean of the raw returns, r_t , Z_t is an independent identical distribution (i.i.d.) stochastic process with mean zero and variance one, and the nonnegative, positively serially-correlated mixing variable, V_t , serves as a proxy for the aggregate amount of information flow to the market. Equation (2) takes the typical form associated with the discrete time ARCH class of models in which both m_t and V_t are assumed measurable with respect to the time $t - 1$ observable information set, so that the conditional variance of the return equals $\text{Var}_{t-1}(R_t) = \text{Var}_{t-1}(r_t) = V_t$. More generally, however, the information flow is not directly observable, and V_t is naturally modeled as a latent, or stochastic, volatility process. Of course, lacking additional assumptions regarding the temporal dependencies in the V_t process, the model in equation (2) is void of testable empirical implications in regards to the observed volatility dynamics.

The rest of this section develops a specific representation of equation (2) that affords an interpretation of volatility as governed by heterogeneous information arrivals, while retaining the capacity of reconciling the seemingly conflicting evidence regarding volatility at the various return frequencies, and producing interesting empirical implications that are directly testable.

A. Volatility as a Manifestation of Heterogeneous Information Arrivals

Motivated by the empirical observations in Section I, we assume that the volatility process reflects the aggregate impact of N distinct information arrival processes; say $V_{j,t} \geq 0$, where $j = 1, 2, \dots, N$. Also, following Andersen (1994), the temporal dependence of each constituent component is expressed in

terms of a standard log-normal stochastic volatility, or Exponential SARV, model,

$$\nu_{j,t} = \alpha_j \cdot \nu_{j,t-1} + \varepsilon_{j,t}, \quad (3)$$

where $\nu_{j,t} \equiv \ln(V_{j,t}) - \mu_j$, $\mu_j \equiv E[\ln(V_{j,t})]$, and the $\varepsilon_{j,t}$ s are assumed to be i.i.d. $N(0, \sigma_j^2)$ for all $j = 1, 2, \dots, N$. The logarithmic formulation in equation (3) ensures that the number of information arrivals dictated by the j th component process, $V_{j,t} \equiv \exp(\nu_{j,t} + \mu_j)$, is positive. The autoregressive coefficient, α_j , reflects the degree of persistence in the j th information arrival process. Consistent with the notion of positively serially-correlated but stationary "news" arrivals, we restrict these parameters to fall within the unit interval; i.e., $0 \leq \alpha_j < 1$.

According to the mixture-of-distributions hypothesis, each information arrival process has an effect on the aggregate latent volatility process. Assume that this combined effect may be represented by

$$V_t = \exp(\nu_t + \mu_\nu), \quad (4)$$

where $\nu_t \equiv \sum_{j=1}^N \nu_{j,t}$ and $\mu_\nu \equiv \sum_{j=1}^N \mu_j$. Note, the spectrum for the aggregate ν_t process at frequency ω is simply given by the sum of the spectra for the N underlying independent AR(1) processes, i.e.,

$$f_\nu(\omega) = (2\pi)^{-1} \cdot \sum_{j=1}^N \sigma_j^2 \cdot |1 - \alpha_j \cdot \exp(-i \cdot \omega)|^{-2}, \quad (5)$$

where $0 \leq \omega \leq \pi$. For small values of N , this spectral representation for ν_t can in principle be used to maximize Whittle's frequency domain approximation to the likelihood function for $\ln(R_t^2) = \mu_\nu + \nu_t + \ln(Z_t^2)$ as specified by equations (2), (3), and (4); see Breidt, Crato, and de Lima (1997) for an application of this frequency-domain based likelihood procedure in the estimation of a stochastic volatility model. Alternatively, the model may be estimated by any one of the Markov Chain Monte Carlo methods described in Jacquier, Polson, and Rossi (1994) and Shephard (1996), or by the Efficient Method of Moments developed by Gallant and Tauchen (1996). Yet, in the actual empirical implementation with high frequency data, a large number of heterogeneous information arrival processes is likely called for, rendering any of these direct estimation procedures for determining the α_j s as well as the remaining model parameters intractable.

In light of this, suppose that the heterogeneity in the individual component processes, as dictated by the AR(1) parameters, may be approximated by the normalized beta distribution,

$$dF_\alpha(\alpha) = 2 \cdot B(p, q)^{-1} \cdot \alpha^{2p-1} \cdot (1 - \alpha^2)^{q-1} d\alpha, \quad (6)$$

where $0 \leq \alpha < 1$, $0 < p$, $1 < q < 2$, and the beta function is given by $B(p, q) \equiv \int_0^1 x^{p-1}(1-x)^{q-1} dx = \Gamma(p) \cdot \Gamma(q) \cdot \Gamma(p+q)^{-1}$. This parametric density function allows for a wide variety of shapes in the bounded distribution of the α_j s, and thus provides a flexible way of characterizing different forms for dynamic dependencies in the latent information arrival processes; see, e.g., Johnson and Kotz (1970) for a discussion of the beta distribution. Moreover, this particular distributional assumption regarding the heterogeneity in the underlying $\nu_{j,t}$ processes has a number of testable implications in regard to the dynamic dependencies in the observable process for the absolute returns.

Specifically, consider the situation in which the number of component processes is arbitrarily large; i.e., $N \rightarrow \infty$, so that the sum of the individual AR(1) spectra in equation (5) converges to the probability weighted integral over the distribution for the α_j coefficients. Following Granger (1980), the resulting spectrum for the ν_t process may then be evaluated as

$$\begin{aligned}
 f_{\nu_t}(\omega) &= (2\pi)^{-1} \cdot \sigma_\varepsilon^2 \cdot \int_0^1 |1 - \alpha \cdot \exp(-i \cdot \omega)|^{-2} dF_\alpha(\alpha) \\
 &= (2\pi)^{-1} \cdot \sigma_\varepsilon^2 \cdot 2 \cdot B(p, q)^{-1} \cdot \sum_{j=-\infty}^{\infty} \left[\int_0^1 \alpha^{2p+|j|-1} \cdot (1-\alpha^2)^{q-2} d\alpha \right] \\
 &\quad \times \exp(-i \cdot \omega \cdot j) \\
 &= (2\pi)^{-1} \cdot \sigma_\varepsilon^2 \cdot \Gamma(p+q) \cdot \Gamma(q-1) \cdot \Gamma(p)^{-1} \cdot \Gamma(q)^{-1} \\
 &\quad \times \sum_{j=-\infty}^{\infty} \Gamma(p + 1/2|j|) \cdot \Gamma(p+q-1 + 1/2|j|)^{-1} \cdot \exp(-i \cdot \omega \cdot j), \quad (7)
 \end{aligned}$$

where $\sigma_\varepsilon^2 \equiv \lim_{N \rightarrow \infty} \sum_{j=1}^N \sigma_j^2$. Alternatively, the spectrum for the ν_t process may be written as $f_{\nu_t}(\omega) \equiv \sigma_\nu^2 \cdot (2\pi)^{-1} \cdot \sum_{j=-\infty}^{\infty} \rho(\nu_t, j) \cdot \exp(-i \cdot \omega \cdot j)$, where σ_ν^2 and $\rho(\nu_t, j)$ denote the variance and the j th order autocorrelation for ν_t , respectively. Thus, by equating terms in the identical powers of j , it follows that the autocorrelations for ν_t must be proportional to $\Gamma(p + 1/2|j|) \cdot \Gamma(p+q-1 + 1/2 \cdot |j|)^{-1}$. This implies that for large $|j|$, the autocorrelations for the aggregate latent logarithmic volatility process behaves like^{12,13}

$$\rho(\nu_t, j) \sim j^{1-q}. \quad (8)$$

¹² By Stirling's Formula $\Gamma(a+j) \cdot \Gamma(b+j)^{-1} \approx c \cdot j^{a-b}$ for large j , where c is a factor of proportionality.

¹³ This same idea has recently been employed by Ding and Granger (1996) in imposing a slowly hyperbolically decaying structure on the parameters in a GARCH type model for daily stock returns.

The dependence in ν_t will therefore dissipate at the slow hyperbolic rate of decay associated with the covariance stationary fractionally integrated, or $I(d)$, class of models, with $d = 1 - \frac{1}{2} \cdot q$.¹⁴

Of course, the mixture-of-distributions hypothesis in equation (2) is phrased in terms of the latent volatility process, V_t , and not ν_t . However, it is readily demonstrated that the long-run dependence induced in the autocorrelations for ν_t , for $0 < d = 1 - \frac{1}{2} \cdot q < \frac{1}{2}$, carries over to any positive power transforms of the fundamental V_t process; say $V_t^\theta = \exp^\theta(\nu_t + \mu_\nu)$, where $\theta > 0$. First, note that $E(V_t^\theta \cdot V_{t-j}^\theta) = [E(V_t^\theta)]^2 \cdot \exp[\theta^2 \cdot \rho(\nu_t, j) \cdot \sigma_\nu^2]$; see Granger and Newbold (1976) for a general discussion regarding the autocorrelation functions of power transforms of normals, and Andersen (1994) for a more detailed analysis of the particular case analyzed here. Consequently, the j th order autocorrelation for V_t^θ may be written as

$$\rho(V_t^\theta, j) = [\exp(\theta^2 \cdot \sigma_\nu^2) - 1]^{-1} \cdot [\exp\{\theta^2 \cdot \sigma_\nu^2 \cdot \rho(\nu_t, j)\} - 1]. \tag{9}$$

However, $\rho(\nu_t, j) \sim j^{2d-1} \rightarrow 0$ for $j \rightarrow \infty$, so that by a first order Taylor series expansion $\exp[\theta^2 \cdot \sigma_\nu^2 \cdot \rho(\nu_t, j)] \approx 1 + \theta^2 \cdot \sigma_\nu^2 \cdot \rho(\nu_t, j)$. From this it follows immediately that, for large j ,

$$\rho(V_t^\theta, j) \approx [\exp(\theta^2 \cdot \sigma_\nu^2) - 1]^{-1} \cdot \theta^2 \cdot \sigma_\nu^2 \cdot \rho(\nu_t, j) \sim j^{2d-1}. \tag{10}$$

Thus, the V_t^θ and the ν_t processes share the same long-run decay in their autocorrelation functions. More importantly, this long-run dependence in the V_t^θ process is propagated to the process for the absolute returns raised to an arbitrary power, $2 \cdot \theta$. To see this, note that from equation (2) any power transform of the absolute returns is directly related to the latent volatility process by $|R_t|^{2 \cdot \theta} = V_t^\theta \cdot |Z_t|^{2 \cdot \theta}$. Thus, because Z_t^θ is an i.i.d. process, it follows that $\rho(|R_t|^{2 \cdot \theta}, j)$ is proportional to $\rho(V_t^\theta, j)$. The return volatility will therefore exhibit the identical slow hyperbolic rate of decay whether measured in terms of the absolute returns, $|R_t|$, the squared returns, R_t^2 , or any other power transform, $|R_t|^{2 \cdot \theta}$; i.e., for large j ,

$$\rho(|R_t|^{2 \cdot \theta}, j) \sim j^{2d-1}. \tag{11}$$

This result is important because it demonstrates how the long-memory features of volatility may arise naturally through the interaction of a large number of diverse information processes. From a conceptual perspective, it implies that the long-memory characteristics reflect inherent properties of the return-generating process, rather than external shocks that induce a structural shift in the volatility process, as, e.g., suggested by Lamoureux and

¹⁴ The process y_t is integrated of order d , or $I(d)$, if $(1 - L)^d y_t$ is stationary and ergodic with a positive and bounded spectrum across all frequencies, where the fractional differencing operator $(1 - L)^d$ is defined by the binomial expansion, $(1 - L)^d \equiv \sum_{j=0, \dots, \infty} \Gamma(j - d) \cdot \Gamma(j + 1)^{-1} \cdot \Gamma(-d) \cdot L^j$. If $0 < d < \frac{1}{2}$ the autocorrelations for y_t are all eventually positive, and decay at a hyperbolic rate. For recent surveys of the relevant literature see, e.g., Baillie (1996), Beran (1994), and Robinson (1994b).

Lastrapes (1990). In other words, the mechanism responsible for the fractional integration in volatility is generic to the returns process, and thus ever present, so that with high frequency data it may be feasible to identify the manifestation of the phenomenon even over relatively short spans of calendar time. We pursue this possibility in Section III.

For now, we simply note that the result in equation (11) is consistent with the empirical behavior of the autocorrelograms for the various power transforms of daily equity returns reported in Ding, Granger, and Engle (1993). Although Harvey and Streibel (1996) argue that it is impossible, by theoretical means, to ascertain which value of θ will uniformly maximize the autocorrelations for the simple stochastic volatility model, corresponding to $N = 1$ in the current setup, it is noteworthy that the sample autocorrelations for the daily returns analyzed in Ding, Granger, and Engle (1993) attain their maxima for θ very close to one-half.¹⁵ Motivated by this observation, we concentrate on the correlation structure for the intradaily absolute returns in the empirical investigations below.¹⁶

B. Temporal Aggregation of Volatility

The long memory characterization obtained in equation (11) may appear incomplete, because the mixture-of-distributions hypothesis is silent regarding the proper discrete time sampling interval for the intradaily return series. While the complex volatility structure entertained renders a full distributional characterization infeasible, it is possible to show that the autocorrelations for the temporally aggregated squared returns eventually exhibit the identical long-run hyperbolic rate of decay irrespective of the sampling frequency.¹⁷ In particular, let $R_{\tau}^{(k)} \equiv R_{\tau \cdot k} + R_{\tau \cdot k-1} + \dots + R_{\tau \cdot k-k+1} = V_{\tau \cdot k}^{1/2} \cdot Z_{\tau \cdot k} + V_{\tau \cdot k-1}^{1/2} \cdot Z_{\tau \cdot k-1} + \dots + V_{\tau \cdot k-k+1}^{1/2} \cdot Z_{\tau \cdot k-k+1}$, where $\tau = 1, 2, \dots$, and $k = 1, 2, \dots$, denote the temporally aggregated returns.¹⁸ Since the Z_t s are i.i.d., it follows that the j th order autocorrelation for $(R_{\tau}^{(k)})^2$ is proportional to the correlation between $V_{\tau \cdot k} \cdot Z_{\tau \cdot k}^2 + V_{\tau \cdot k-1} \cdot Z_{\tau \cdot k-1}^2 + \dots + V_{\tau \cdot k-k+1} \cdot Z_{\tau \cdot k-k+1}^2$ and $V_{(\tau-j) \cdot k} \cdot Z_{(\tau-j) \cdot k}^2 + V_{(\tau-j) \cdot k-1} \cdot Z_{(\tau-j) \cdot k-1}^2 + \dots + V_{(\tau-j) \cdot k-k+1} \cdot Z_{(\tau-j) \cdot k-k+1}^2$. Thus, by collecting terms corresponding to the same lag length, the latter correlation is, in turn, proportional to the sum of the autocorrelations between

¹⁵ This phenomenon is referred to as the Taylor effect by Granger and Ding (1997).

¹⁶ Varying the power of the absolute returns effectively amounts to varying the importance of extreme events. For the five-minute DM-\$ returns analyzed here, the first order sample autocorrelations for $|R_t|$ and R_t^2 equal 0.309 and 0.201, respectively, consistent with the findings for other rates and time periods reported in Müller, Dacorogna, and Pictet (1996).

¹⁷ In a related context, assuming a Wold type representation, Chambers (1995) has recently shown that the degree of fractional integration in the levels is preserved under temporal aggregation of both stock and flow variables.

¹⁸ For the 74,880 five-minute returns analyzed here, an aggregation factor of $k = 2$ would correspond to the time series of 37,440 ten-minute returns, whereas $k = 3$ refers to the time series of 24,960 fifteen-minute returns, etc.

V_t^2 and $V_{t-j \cdot k-h}^2$, where the weights corresponding to $h = -k + 1, -k + 2, \dots, k - 1$ equal $(k - |h|)$, and zero otherwise. Consequently, for long lags, j ,

$$\begin{aligned} \rho([R_\tau^{(k)}]^2, [R_{\tau-j}^{(k)}]^2) &\sim k^{-2} \cdot \sum_{h=-k+1}^{k-1} (k - |h|) \cdot \rho(V_t^2, V_{t-j \cdot k-h}^2) \\ &\sim k^{-2} \cdot \sum_{h=-k+1}^{k-1} (k - |h|) \cdot (j \cdot k + h)^{2d-1} \approx (j \cdot k)^{2d-1} \sim j^{2d-1}. \end{aligned} \tag{12}$$

That is, the long-memory features of the squared returns are consistent with the characteristics of a self-similar process in the sense of Mandelbrot and van Ness (1968) and Mandelbrot and Wallis (1969). Moreover, by analogy to the $k = 1$ case detailed above, if the temporally aggregated latent volatility process is log-normally distributed, then the long-run autocorrelation structure for $|R_\tau^{(k)}|^{2 \cdot \theta}$ will be identical for all θ . Thus, in this case the autocorrelations for any power transform of the temporally aggregated absolute returns should again eventually decay at the identical hyperbolic rate of j^{2d-1} irrespective of the sampling frequency, k .¹⁹

C. Extensions

Before we describe the empirical findings pertaining to the propositions developed so far, we should emphasize that the theoretical model readily accommodates extensions in a number of directions that allow for added flexibility and realism in the portrayal of the volatility dynamics without affecting the salient long-run dependencies. First, the mixture-of-distributions hypothesis in equations (2), (3), and (4) obviously neglects the repetitive intradaily pattern in the volatility that is evident in the spectrum and the autocorrelogram for the 5-minute absolute DM-\$ returns in Figures 1 and 2. To incorporate this periodicity, let $s(t)$ denote the stage of the periodic cycle at time t . A slight modification of the model in equation (2), that explicitly allows for the pronounced intradaily pattern, takes the form,

$$R_t = V_t^{1/2} \cdot S_{t,s(t)} \cdot Z_t, \tag{13}$$

where the spectrum for the independent seasonal component $S_{t,s(t)}$ has no mass at frequencies lower than one day.²⁰ Since $\ln(R_t^2) = \mu_\nu + \nu_t + \ln(S_{t,s(t)}^2) + \ln(Z_t^2)$,

¹⁹ Previous empirical work with much simpler dependency structures have found the log-normal distribution to work remarkably well at various return frequencies; see, e.g., Andersen (1996), Shephard (1996), and Taylor (1994).

²⁰ A particularly simple representation that restricts $S_{t,s(t)}$ to depend only on the stage of the periodic cycle, $s(t)$, has been successfully employed by Andersen and Bollerslev (1997a) in modeling the periodicity in the DM-\$ return series analyzed here. Note also that periodicities at, say, the weekly frequency may be accommodated analogously.

it follows that the long-run behavior, as dictated by the spectrum for ν_t near frequency zero, is unaffected by this additional periodic component. Of course, the autocorrelation structure in the absolute returns may be severely influenced by the presence of the repetitive intradaily pattern. However, as demonstrated below, by explicitly filtering out the daily and higher frequency dynamics in the absolute returns process, it is possible to uncover the long-run hyperbolic rate of decay implied by the aggregation of the component processes.

Rather than being independent, the “news” arrivals might more realistically share a number of common factors related to the state of the overall economy, while still exhibiting their own separate influence on the valuation of the asset. To allow for such common effects, consider the following generalization of the simple logarithmic AR(1) processes in equation (3),

$$\nu_{j,t} = \alpha_j \cdot \nu_{j,t-1} + \sum_{k=1}^K \xi_{j,k} \cdot \kappa_{k,t} + \varepsilon_{j,t} \tag{14}$$

where $\kappa_{k,t}$ represents the de-measured k th common factor, and the factor loading for the j th arrival process associated with common factor k is denoted by $\xi_{j,k}$.²¹ Provided that each of the common factors is covariance stationary, so that their spectra are bounded across all frequencies, and that the heterogeneity in the individual persistence parameters, α_j , is determined by the beta distribution in equation (6), it follows immediately by analogy to the results in Granger (1980), that the hyperbolic decay in the autocorrelation function for ν_t is preserved under this more general specification. The independence assumption in equation (3) merely serves to simplify the exposition. The implied degree of long-run dependence in the absolute returns remains intact under the more realistic assumptions in equation (14).

Finally, while the derivation above explicitly relies on the beta distribution in equation (6) for characterizing the heterogeneity in the individual component processes, it is obvious that the degree of fractional integration, $d = 1 - \frac{1}{2} \cdot q$, is independent of the p parameter in the beta distribution. Only the shape of the distribution for α close to unity, as dictated by q , is important. Intuitively, if a sufficient number of the individual information arrival processes have high persistence, albeit $\alpha_j < 1$, the aggregate information arrival process will display long-memory characteristics. This suggests that the aggregation argument behind the long-run dependence in ν_t is somewhat more general than portrayed above. Indeed, Lin (1991) and Granger and Ding (1996) show that the so-called Generalized Integrated class of models, discussed by Granger (1987, 1988), may arise through the aggregation of a closely related component type structure.²²

²¹ This construction closely mirrors the spirit, if not the structure, of the arbitrage pricing theory of Ross (1976).

²² A closely related method for simulating fractional Gaussian noise has been proposed by Mandelbrot and Wallis (1969) and Mandelbrot (1971).

III. Estimating Long-Memory in Volatility

The theoretical framework developed in the previous section builds on the idea that the aggregate market volatility represents the manifestation of numerous heterogeneous information arrival processes; some with short-run volatility dependencies, others possessing more highly persistent volatility patterns. As time passes the short-run processes decay significantly, while the more highly persistent processes remain influential. Hence, while sudden bursts of volatility typically will possess both short-run and long-run components, the short-run decay stands out most clearly over the intradaily frequencies, whereas the highly persistent processes only will be noticeable over longer horizons. However, traditional correlation based measures for the degree of volatility persistence, obtained from high frequency intradaily data, will tend to pick up only the effects of the complex interaction between the short-run decay associated with the less persistent processes and the strong intradaily periodicity, thus missing the importance of the long-run volatility components.

To illustrate, consider the third row in Table I, which reports the first order sample autocorrelation, $\hat{\rho}^{(k)}$, for the temporally aggregated absolute five-minute returns; i.e., $|R_\tau^{(k)}|$ for $\tau = 1, 2, \dots, 74,880/k$. The first order sample autocorrelations for the highest intradaily frequencies all indicate very significant positive serial correlation in the absolute returns. However, there is a sharp dropoff in the value of the sample autocorrelations at the three-hour sampling frequency, i.e., $k = 36$. In fact, the first order autocorrelations for the eight and twelve hours absolute returns, i.e., $k = 96$ and 144 , are both negative, suggestive of an anti-persistent volatility process.²³

The latter conclusion is, of course, grossly misleading, being driven by the combined effects of the intradaily periodic pattern and the complex multiple volatility component structure, thus creating an overall system that conforms to the law of motion outlined in equation (13). In order to justify this interpretation, however, we must develop an alternative estimation approach. In particular, while strong intraday periodicity and pronounced volatility clustering have been extensively documented at the high frequency level, there is little direct evidence for the existence of longer-run volatility components that induce the type of long-memory behavior in the high frequency volatility process implied by our information aggregation rendition of the mixture-of-distributions hypothesis.

Specifically, from the model detailed in the previous section, the absolute returns should exhibit the identical long-run dependence for all sampling frequencies, say $I(d)$, where $0 < d < 1/2$. Representing the process for the absolute returns as $(1 - L)^d |R_\tau^{(k)}| = \eta_\tau^{(k)}$, where $\eta_\tau^{(k)}$ is a stationary and ergodic

²³ Müller *et al.* (1997) attribute this negative correlation for the half-day lag to a heterogeneous market with different groups of traders participating during their regular business hours.

Table I

Persistence Measures for Temporally Aggregated Absolute Intraday Returns

The percentage returns are based on interpolated five-minute logarithmic average bid-ask quotes for the Deutschemark-U.S. Dollar spot exchange rate from October 1, 1992 through September 29, 1993. Quotes from Friday 21:00 Greenwich Mean Time (GMT) through Sunday 21:00 GMT have been excluded, resulting in a total of 74,880 five-minute return observations. The length of the different intraday sampling intervals equal 5-k minutes. Each absolute return series consists of a total of 74,880/k nonoverlapping observations; i.e., $|R_\tau^{(k)}| = |R_{\tau \cdot k} + R_{\tau \cdot k - 1} + \dots + R_{\tau \cdot k - k + 1}|$ where $\tau = 1, 2, \dots, T/k$ and $T = 74,880$. The $\hat{\rho}^{(k)}$ column gives the first order autocorrelations for $|R_\tau^{(k)}|$. The Log-Periodogram estimates for d from equation (18) are denoted by $\hat{d}_{GH}^{(k)}$. The truncation and trimming parameters for $k = 2, 3, \dots, 144$ are determined by $m = \lceil (T/k)^{1/2} \rceil$ and $\ell = \lfloor m^{1/4} \rfloor$, respectively, where $\lceil \cdot \rceil$ denotes the integer value. In order to avoid any confounding effects from frequencies corresponding to less than 1 day, the estimate for $\hat{d}_{GH}^{(1)}$ is based on $m = 255$ and $\ell = 3$. The $\hat{d}_{AP}^{(k)}$ column gives the average periodogram estimates for d defined in equation (21). The truncation parameters are determined by $m = \lceil (T/k)^{1/2} \rceil$ for $k = 2, 3, \dots, 144$, while $\hat{d}_{AP}^{(1)}$ is based on $m = 255$. The scalar q is fixed at $1/4$ across all sampling frequencies.

k	T/k	$\hat{\rho}^{(k)}$	$\hat{d}_{GH}^{(k)}$	$\hat{d}_{AP}^{(k)}$
1	74,880	0.309	0.321	0.385
2	37,440	0.313	0.367	0.375
3	24,960	0.307	0.337	0.376
4	18,720	0.287	0.339	0.364
6	12,480	0.268	0.288	0.351
8	9,360	0.272	0.281	0.340
9	8,320	0.251	0.214	0.333
12	6,240	0.229	0.309	0.317
16	4,680	0.246	0.269	0.303
18	4,160	0.193	0.457	0.310
24	3,120	0.164	0.220	0.291
32	2,340	0.171	0.221	0.302
36	2,080	0.097	0.178	0.275
48	1,560	0.075	0.236	0.311
72	1,040	0.007	0.214	0.270
96	780	-0.025	0.382	0.274
144	520	-0.033	0.382	0.286

process with a bounded spectrum, $f_{\eta,k}(\omega)$, for all frequencies, ω , the spectrum for $|R_\tau^{(k)}|$ may therefore be written as²⁴

$$f_{|R|,k}(\omega) = |[1 - \exp(-i \cdot \omega)^{-d}]|^2 \cdot f_{\eta,k}(\omega) = |2 \cdot \sin(1/2 \cdot \omega)|^{-2d} \cdot f_{\eta,k}(\omega). \quad (15)$$

Since $\lim_{\omega \rightarrow 0} \omega^{-1} \cdot \sin(1/2 \cdot \omega) = 1/2$, it follows that for the frequencies close to zero, i.e., $\omega \approx 0$,

$$f_{|R|,k}(\omega) \approx f_{\eta,k}(0) \cdot |\omega|^{-2d}, \quad (16)$$

²⁴ Formal conditions for the equivalence between this spectral definition of long-memory and the hyperbolic decay rate in the autocorrelation function are discussed in Beran (1994) and Robinson (1994b); see also Granger and Ding (1996).

or,

$$\ln[f_{|R|,k}(\omega)] \approx \ln[f_{\eta,k}(0)] - 2 \cdot d \cdot \ln(\omega). \tag{17}$$

Hence, the spectrum should be approximately log-linear for the long-run frequencies. Indeed, when viewed on the double logarithmic scale in Figure 1, the spectrum for the five-minute absolute returns, i.e., $k = 1$, is very close to a straight line over the interdaily frequencies, $0 < \omega < 2 \cdot \pi/288 \approx 0.0218$.

The Geweke and Porter-Hudak (1983) log-periodogram estimate for the fractional order of integration is based directly on this relationship. These estimates for the dependence in $|R_\tau^{(k)}|$, obtained across all of the intradaily frequencies, are reported in the column labeled $\hat{d}_{GH}^{(k)}$ in Table I. Formally,

$$\hat{d}_{GH}^{(k)} = -\frac{1}{2} \cdot \left\{ \sum_{j=\ell+1}^m [\ln(\omega_{k,j}) - \bar{\omega}_{k,\ell,m}] \cdot \ln(I_{|R|,k}(\omega_{j,k})) \right\} \cdot \left\{ \sum_{j=\ell+1}^m [\ln(\omega_{k,j}) - \bar{\omega}_{k,\ell,m}]^2 \right\}^{-1}, \tag{18}$$

where $\bar{\omega}_{k,\ell,m} \equiv (m - \ell)^{-1} \cdot \sum_{j=\ell+1}^m \omega_{j,k}$, and $I_{|R|,k}(\omega_{j,k})$ denotes the sample periodogram for $|R_\tau^{(k)}|$ at the j th Fourier frequency, i.e., $\omega_{j,k} = 2 \cdot \pi \cdot j \cdot k/T$. Although this estimator for d has been fairly widely used in the literature, consistency for $0 < d < 1/2$ has only recently been established by Robinson (1995) under regularity conditions that include the truncation and trimming parameters both tending to infinity, albeit at a slower rate than the sample size; i.e., $m \rightarrow \infty$, $\ell \rightarrow \infty$, $\ell/m \rightarrow 0$, and $m \cdot k/T \rightarrow 0$. However, the regularity conditions also require that $|R_\tau^{(k)}|$ be normally distributed, which in turn implies that the estimator itself, $\hat{d}_{GH}^{(k)}$, is asymptotically normal with a variance of $m^{-1} \cdot (\pi^2/24)$, independent of the sample size, T/k .

For the results reported in Table I, we took $m = [(T/k)^{1/2}]$ and $\ell = [m^{1/4}]$, respectively, where $[\cdot]$ denotes the integer value. However, in order to avoid any confounding effects from frequencies corresponding to less than one day, the estimate for $\hat{d}_{GH}^{(1)}$ is based on $m = 255$ and $\ell = 3$. Note that, in contrast to the sample autocorrelations reported in the $\hat{\rho}^{(k)}$ column, the $\hat{d}_{GH}^{(k)}$ estimates for the degree of long-run volatility dependence are remarkably stable across the different intradaily return intervals. In fact, when judged by the asymptotic normal distributions, all of the estimates are within less than one asymptotic standard error of $d = 0.359$.²⁵ To illustrate, consider the estimate for $\hat{d}_{GH}^{(1)} = 0.321$, corresponding to the average slope of the spectrum for the five-minute absolute returns in Figure 1 over the frequencies $0 < \omega \leq 255 \cdot 2 \cdot \pi/74,888 \approx$

²⁵ This value of d corresponds to the estimated hyperbolic decay rate in the low-pass filtered five-minute absolute returns described further in Section IV below.

0.0214.²⁶ The theoretical standard error for this estimate equals $\pi \cdot (24 \cdot 255)^{-1/2} \approx 0.040$. As such, the estimates for $k = 1, 2, \dots, 144$ confirm the proposition that the degree of fractional integration in the absolute returns is invariant with respect to the sampling frequency.²⁷

Unfortunately, the assumption of normality underlying the formal statistical justification for the $\hat{d}_{GH}^{(k)}$ estimates is clearly violated in the present context. For instance, the sample skewness and kurtosis for the absolute five-minute returns equal 0.367 and 21.5, respectively. The last column in Table I therefore reports the results from a less restrictive semiparametric estimation procedure for determining d , based on the ratio of the periodogram for two frequencies close to zero.²⁸ To motivate this estimator, let

$$G_{m,k} = 2 \cdot \pi \cdot (T/k)^{-1} \cdot \sum_{j=1}^m I_{|R|,k}(\omega_{j,k}) \tag{19}$$

denote the average periodogram for the frequencies $j = 1, 2, \dots, m$. Then, following Robinson (1994a), for $0 < d < 1/2$ and $m \rightarrow \infty$, but $m \cdot k/T \rightarrow 0$,

$$\ln(G_{m,k}) - \ln(G_{[qm],k}) + (1 - 2 \cdot d) \cdot \ln(q) \approx 0, \tag{20}$$

where $0 < q < 1$. Thus, upon rearranging the terms in equation (20), the following estimator for d becomes apparent,

$$\hat{d}_{AP}^{(k)} = 1/2 + 1/2 \cdot \ln(q)^{-1} \cdot [\ln(G_{m,k}) - \ln(G_{[qm],k})]. \tag{21}$$

Consistency of this frequency-domain estimator for d has been established by Robinson (1994a) under much weaker regularity conditions than those available for $\hat{d}_{GH}^{(k)}$. Furthermore, given the assumption of normality underlying the existing consistency proof for the log-periodogram estimator, $\hat{d}_{GH}^{(k)}$, Lobato and Robinson (1996) have recently shown that the alternative $\hat{d}_{AP}^{(k)}$ estimator is asymptotically normal for $0 < d < 1/4$, but nonnormally distributed for $1/4 \leq d < 1/2$.

The estimates reported in the last column of Table I are based on a truncation parameter of $m = \lceil (T/k)^{1/2} \rceil$ for $k = 2, 3, \dots, 144$, whereas $m = 255$ for $k = 1$ in order to avoid any confounding effects from the intradaily dependencies in the estimation of $\hat{d}_{AP}^{(1)}$. The value for the scalar q was fixed at 0.25 across all the

²⁶ Including all of the frequencies up to $j = \lceil 74,880^{1/2} \rceil = 273$, the estimate for $\hat{d}_{GH}^{(1)}$ drops to a value of only 0.242, highlighting the importance of explicitly excluding the intradaily effects in the estimation.

²⁷ Similar estimates for d based on 10-minute absolute exchange rate returns, corresponding to $k = 2$, have been reported independently by Henry and Payne (1996). These results are also consistent with the notion of an intrinsic time scale in the foreign exchange market as discussed by Müller *et al.* (1993).

²⁸ The same log-periodogram and average periodogram estimators for d implemented here have previously been employed by Chambers (1995) and Delgado and Robinson (1994) in the analysis of various U.K. macroeconomic time series and monthly Spanish inflation rates, respectively.

sampling frequencies. In line with the simulation evidence reported in Lobato and Robinson (1996), some informal sensitivity analysis revealed the results to be fairly robust with respect to this choice. Turning to the actual estimates, the similarities across the different values of k are even more striking than for the log-periodogram estimates. The average value of $\hat{d}_{AP}^{(k)}$ equals 0.321, while ranging from a low of 0.270 for $k = 72$ to a high of only 0.385 for $k = 1$.²⁹

Consistent with the notion of a heterogeneous component structure and the invariance under temporal aggregation, these estimates for d , based on a single year of intradaily returns, correspond very closely to the estimates reported in the extant literature with longer time-spans of daily data. For instance, the same log-periodogram regression and average periodogram estimates for d based on the 3,649 daily DM-\$ absolute returns from March 14, 1979 through September 29, 1993, equal 0.344 and 0.301, respectively.³⁰ Thus, the relatively simple semiparametric frequency-domain estimators in equations (18) and (21) are both capable of uncovering the inherent long-run volatility dependencies in the time series of high-frequency intradaily returns, without having to impose any specific structure on the short-run behavior of the system in order to accommodate the complex intradaily dynamics and repetitive periodic patterns that corrupt conventional correlation based measures.

To further appreciate the notion of long-memory volatility dynamics, consider the properties of the $(1 - L)^d$ filter designed to annihilate the long-run dependence in the absolute returns. Expressing the fractional differencing operator in terms of its binomial expansion, the gain of the filter may be written as

$$| [1 - \exp(-i \cdot \omega)]^d | = \left\{ \left[\sum_{j=0}^{\infty} \delta_j \cdot \cos(\omega \cdot j) \right]^2 + \left[\sum_{j=0}^{\infty} \delta_j \cdot \sin(\omega \cdot j) \right]^2 \right\}^{1/2}, \quad (22)$$

where $0 \leq \omega \leq \pi$, and $\delta_j \equiv \Gamma(j - d) \cdot \Gamma(j + 1)^{-1} \cdot \Gamma(-d)$ for $j = 0, 1, \dots$. Intuitively, the gain at frequency ω represents the magnitude by which the filter multiplies the component of the time series with a repetitive cycle of $2 \cdot \pi/\omega$ periods. The enhanced flexibility provided by allowing for fractional orders of integration is evident from Figure 3, which plots the gains of $(1 - L)^d$ for $d = 0, 0.25, 0.359, 0.5$, and 1. Although the fractional differencing operators with $d > 0$ completely eliminate the zero frequency component, the filters

²⁹ The estimate for $\hat{d}_{AP}^{(1)}$ corrupted by the intradaily frequencies, $j = 256, 257, \dots, 273$, equals 0.172.

³⁰ Both estimates are based on $m = 168$ corresponding to a period of $3,648/168 \approx 21.7$ trading days, or approximately one month, along with $\ell = 4$ and $q = 0.25$.




Figure 3. The figure graphs the filter gains for the fractional differencing operators, $(1 - L)^d$, corresponding to $d = 0, 0.25, 0.359, 0.5$, and 1 .

differ greatly in terms of their gains over the finite frequencies, $0 < \omega \leq \pi$.³¹ The filter gain of $[2 - 2 \cdot \cos(\omega)]^{1/2}$ for the first difference operator, $(1 - L)$, associated with the Integrated GARCH, or IGARCH, class of models for σ_t^2 , implies a much greater down-weighting of the long-run dependencies than the fractional differencing filter with $d = 0.359$.

The justification of filtering the five-minute DM-\$ absolute returns by $(1 - L)^{0.359}$ is also evident from the plot of the estimated spectrum for $(1 - L)^{0.359}|R_t|$ given in Figure 4.³² The flat spectrum for $0 < \omega < 2 \cdot \pi/288 \approx 0.0218$ reveals that the fractional differencing operator is, indeed, successful in eliminating the longer-run dependencies. At the same time, the pronounced peaks associated with the daily periodicity remain very similar across Figures 1 and 4. These results are further underscored by the time-domain autocorrelogram for $(1 - L)^{0.359}|R_t|$ depicted in Figure 5. Although more noisy, the autocorrelations exhibit the same repetitive daily cycles as the autocorrelations for the raw absolute returns in Figure 2.³³ However, in contrast to the autocorrelogram for

³¹ Since the autocorrelations of a long-memory process with $d > 0$ is not summable, all of the $(1 - L)^d$ filters necessarily eliminate the zero frequency component in order to achieve a bounded spectrum for the fractionally differenced series.

³² The binomial expansion for $(1 - L)^{0.359}$ was truncated at a lag length of 1440, corresponding to one week. Also, following Baillie, Bollerslev, and Mikkelsen (1996) all of the presample values for $|R_t|$ were fixed at their unconditional sample analogues.

³³ Note that the scales for the autocorrelograms differ across Figures 2 and 5. In particular, the first order sample autocorrelation of -0.160 for $(1 - L)^{0.359}|R_t|$ does not fit on the scale in Figure 5.




Figure 4. The figure graphs the spectrum for the fractionally differenced five-minute absolute returns on an open position in the Deutschemark–U.S. Dollar spot exchange rate market from October 1, 1992 through September 29, 1993; i.e., $(1 - L)^{0.359} |R_t|$ where $t = 1, 2, \dots, 74,880$. The spectrum is estimated as the smoothed sample periodogram using a triangular kernel with bandwidth of ten.

the raw absolute returns, which display an overall slow hyperbolic rate of decay, the daily periodic patterns in the autocorrelogram for the fractionally differenced five-minute absolute returns are centered on zero. The long-run dependence in the series has been eliminated.

IV. Low-Pass Filtering and Long-Run Volatility Dynamics

The frequency domain based estimators outlined in the previous section allow for the determination of the degree of long-run volatility dependence across the different sampling frequencies by explicitly focusing on the shape of the spectrum around the origin. Alternatively, a nonstructural time-domain estimation procedure could be based on the eventual hyperbolic decay of the autocorrelation function implied by the presence of long-memory. However, the implementation of this idea is obscured by the strong intradaily periodicity, as exemplified by the sample autocorrelogram for the five-minute absolute returns depicted in Figure 2. Various standardization procedures have previously been proposed for modeling the intradaily patterns in the volatility of high frequency foreign exchange rates, including the time-deformation approach advocated by Dacorogna *et al.* (1993) and Ghysels and Jasiak (1995), and the flexible Fourier functional form utilized by Andersen and Bollerslev (1997a, 1997b). These procedures are directly applicable in a forecasting context. In contrast, the low-pass filtering technique developed below is based on




Figure 5. The figure graphs the sample autocorrelogram for the fractionally differenced five-minute absolute returns on an open position in the Deutschemark–U.S. Dollar spot exchange rate market from October 1, 1992 through September 29, 1993; i.e., $(1 - L)^{0.359} |R_t| \equiv |\ln(P_t) - \ln(P_{t-1})|$ where $t = 1, 2, \dots, 74,880$.

a two-sided weighted average of both past and future absolute returns. By explicitly annihilating the dependencies with a periodicity of less than one day, the in-sample low-pass filtered absolute returns are designed to be void of any short-run intradaily dynamics, and as such provide a framework for the ex-post analysis of the long-run volatility determinants based on conventional time-domain methods.

Restricting the attention to daily and longer run dynamics, the ideal low-pass filter would have a gain, or a frequency response function, of zero for all of the intradaily frequencies, and a gain of unity for the interdaily frequencies; i.e., $\beta(\omega) = 1$ for $0 \leq \omega < \omega_D$ and $\beta(\omega) = 0$ for $\omega_D \leq \omega \leq \pi$, where ω_D denotes the daily frequency. By standard filter theory the weights in the corresponding infinite two-sided time-domain filter, $b_\infty(L) = \sum_{j=-\infty}^{\infty} b_{\infty,j} L^j$, is readily found by the inverse Fourier transform; i.e., $b_{\infty,j} = \int_0^\pi \beta(\omega) \cdot \exp(i \cdot \omega \cdot j) d\omega = \sin(j \cdot \omega_D) / (j \cdot \pi)$ for $j \neq 0$ and $b_{\infty,0} = \int_0^\pi \beta(\omega) d\omega = \omega_D / \pi$. Of course, in practice, with a finite number of data points, this infinite filter is not applicable. However, the weights in the finite dimensional two-sided approximate filter,

$$b_p(L) = \sum_{j=-p}^p b_{p,j} L^j, \tag{23}$$

that achieves the minimum squared approximation error, subject to the constraint that the filter weights sum to unity, is given by $b_{p,j} = b_{\infty,j} - (b_{\infty,0} +$




Figure 6. The figure graphs the filter weights for the approximate two-sided low-pass filter; i.e., $b_{p,j} = b_{\infty,j} - (b_{\infty,0} + 2 \cdot \sum_{h=1}^p b_{\infty,h} - 1)/(2 \cdot p + 1)$, where $p = 1440$, $\omega_D = 0.021$, $b_{\infty,0} = \omega_D/\pi$, and $b_{\infty,j} = \sin(j \cdot \omega_D)/(j \cdot \pi)$ for $j = -1440, \dots, -1, 0, 1, 2, \dots, 1440$.

$2 \cdot \sum_{h=1}^p b_{\infty,h} - 1)/(2 \cdot p + 1)$ for $j = -p, -p + 1, \dots, p$; see, e.g., Baxter and King (1995).³⁴ Since the weights are symmetric, the gain of this approximate low-pass filter may be conveniently written as,³⁵

$$|b_p[\exp(i \cdot \omega)]| = b_{p,0} + 2 \cdot \sum_{j=1}^p b_{p,j} \cdot \cos(j \cdot \omega). \tag{24}$$

The higher the value of p , the more accurate this gain approximates the ideal gain of $\beta(\omega)$. Of course, in practice a tradeoff is necessarily called for in terms of the shape of the gain given by equation (24) and the number of observations that have to be sacrificed at the beginning and end of the sample in the implementation of $b_p(L)$.

For the five-minute returns analyzed here, we took $\omega_D = 0.021$, corresponding to roughly 299 periods or close to 25 hours, along with $p = 1440$, or one week. The corresponding filter weights, $p_{-1440}, p_{-1439}, \dots, p_{1440}$, are given in Figure 6. The accuracy afforded by this choice of p is illustrated in Figure 7,

³⁴ Formally, the weights minimize the squared approximation error, $\int_0^\pi [b_p(\exp(i \cdot \omega)) - \beta(\omega)]^2 d\omega$. The constraint that the weights sum to unity ensures that the long-run zero frequency behavior of the series is unaltered by the filtering.

³⁵ The symmetry of the filter also guarantees that the phase is equal to zero across all frequencies; i.e., $\phi(\omega) = \tan^{-1}\{\text{Im}[b_p(\exp(i \cdot \omega))]/\text{Re}[b_p(\exp(i \cdot \omega))]\} = 0$. Intuitively, $\phi(\omega)/\omega$ represents the amount by which the filter shifts the series back in time at frequency ω .




Figure 7. The figure graphs the ideal and the approximate low-pass filter gains for $\omega_D = 0.021$ and $p = 1440$. The gain for the approximate two-sided filter, $|b_{1440}[\exp(i \cdot \omega)]|$, is given by equation (24) in the text.

which graphs the gain of the ideal low-pass filter along with this two-sided approximation. The overall coherence between the gain of the two filters is generally very good. Only for the frequencies close to ω_D is there some evidence that the frequencies greater than ω_D do not receive a zero weight and that the frequency gains for $\omega < \omega_D$ are different from unity. Such “leakage” and “compression” is inevitable with a finite value of p .

The effectiveness of this approximate low-pass filter in eliminating the short-run intradaily volatility components is clearly seen from Figure 8. The estimated spectrum for the filtered five-minute absolute returns, $b_{1440}(L)|R_t|$, where $t = 1441, 1442, \dots, 73440$, has virtually no-mass at the frequencies higher than ω_D .³⁶ Of course, the log-linear relation for the interdaily frequencies, implied by the presence of long-memory, remains intact.

This long-run dependence is also evident in the autocorrelation function for the low-pass filtered absolute returns. In contrast to the overall decay of the autocorrelogram for the raw absolute returns in Figure 2, which is masked by the strong recurring daily pattern, the autocorrelation function for $b_{1440}(L)|R_t|$,

³⁶ Due to the loss of one week of observations at the beginning and the end of the sample, the filtered time series consist of “only” $74,880 - 2 \cdot 1,440 = 72,000$ observations. Of course, the logarithmic scale in Figure 7 may be slightly misleading as the spectrum is not identically equal to zero for $\omega > \omega_D$.




Figure 8. The figure graphs the spectrum for the low-pass filtered five-minute absolute returns on an open position in the Deutschemark–U.S. Dollar spot exchange rate market from October 1, 1992 through September 29, 1993; i.e., $b_{1440}(L)|R_t|$ where $t = 1441, 1442, \dots, 73440$, and the coefficients in the filter are defined in the text. The spectrum is estimated as the smoothed sample periodogram using a triangular kernel with bandwidth of ten.

depicted in Figure 9, display a distinct hyperbolic rate of decay.³⁷ The actual magnitude of the correlations have also increased substantially. Even out to lag 2,880, or ten days, the autocorrelations all remain above 0.2. Matching the sample autocorrelations for the low-pass filtered absolute returns with the hyperbolic decay implied by the presence of long-memory thus provides an alternative time-domain procedure for estimating d .

In general, since the autocorrelations of a long-memory process are eventually all positive, it follows that, for large j ,

$$\ln(\rho_j) \approx \ln(c) + (2 \cdot d - 1) \cdot \ln(j), \quad (25)$$

where ρ_j denotes the j th order autocorrelation, and c is just a factor of proportionality. Replacing the autocorrelations by their sample analogues, $\hat{\rho}_j$, therefore suggests the following least squares estimator,

$$\hat{d}_{AC} = \frac{1}{2} + \frac{1}{2} \cdot \left\{ \sum_{j=r+1}^{r+n} [\ln(j) - \bar{j}_{r,n}] \cdot \ln(\hat{\rho}_j) \right\} \cdot \left\{ \sum_{j=r+1}^{r+n} [\ln(j) - \bar{j}_{r,n}]^2 \right\}^{-1} \quad (26)$$

³⁷ A similarly shaped autocorrelogram for twenty-minute absolute DM–\$ returns standardized by a measure of the degree of market activity, or “theta-time,” has been reported by Dacorogna *et al.* (1993).




Figure 9. The figure graphs the sample autocorrelogram for the low-pass filtered five-minute absolute returns on an open position in the Deutschemark–U.S. Dollar spot exchange rate market from October 1, 1992 through September 29, 1993; i.e., $b_{1440}(L)|R_t|$ where $t = 1441, 1442, \dots, 73440$, and the coefficients in the filter are defined in the text. The dotted line gives the curve for the estimated hyperbolic decay rate; i.e., $j^{2 \cdot 0.359 - 1} = j^{-0.282}$.

where $\bar{j}_{r,n} \equiv n^{-1} \cdot \sum_{j=r+1}^{r+n} \ln(j)$, and $r \rightarrow \infty$, but $r/T \rightarrow 0$. This semi-parametric time-domain estimator for the fractional degree of integration, proposed by Robinson (1994b), is closely related to the minimum distance estimators recently analyzed by Tieslau, Schmidt, and Baillie (1996). No formal asymptotic distribution theory is yet available for the estimator in equation (26), although it seems likely that \hat{d}_{AC} will be consistent under rather weak regularity conditions.

Applying this estimator to the sample autocorrelations for the low-pass filtered absolute returns, $\hat{\rho}(b_{1440}(L)|R_t|, j)$ for $j = 5, 6, \dots, 2880$, yields $\hat{d}_{AC} = 0.359$. This estimate is thus fully consistent with the results from the frequency based procedures reported in the previous section. It is furthermore evident from the plot in Figure 9, that the implied hyperbolic rate of decay, $j^{2 \cdot 0.359 - 1} = j^{-0.282}$, is in close accordance with the actual shape of the autocorrelogram. It is worth stressing that, due to the strong intradaily volatility patterns and the associated recurring negative sample autocorrelations at the half-day lags, this estimation procedure for determining the fractional order of integration simply is not applicable with the raw absolute returns. Only by explicitly eliminating the intradaily dependencies do the sample autocorrelations all become positive and the hyperbolic decay stand out clearly.

The low-pass filtered returns also set the stage for a more structural investigation of the determinants behind the important volatility components. A detailed analysis along these lines is beyond the scope of the present article. To

simply illustrate the idea consider the announcement effects associated with the Employment report, the Bundesbank biweekly meeting, the Durable Goods figures, and the report on Business Inventories discussed in Section I above. On calculating the average increase in the absolute five-minute DM-\$ returns in the two hours immediately following these announcements, the four effects, with robust standard errors in parentheses, are 0.0091 (0.0052), 0.0133 (0.0025), 0.0067 (0.0032), and 0.0143 (0.0023), respectively.³⁸ Thus, the increase in the volatility associated with the Employment report is apparently not significant at the usual five percent level, while the Durable Goods report is only marginally significant. Also, the figures on Business Inventories appear to be the most significant of the four “news” events. These estimates should be carefully interpreted, however, as the intradaily pattern is prone to obscure the fundamental relationships. In fact, on calculating the average two-hour increase in the volatility for the low-pass filtered returns, $b_{1440}(L)|R_t|$, the four effects are 0.0076 (0.0005), 0.0078 (0.0004), 0.0032 (0.0004), and -0.0005 (0.0005), respectively. The t -statistics for the three former announcements are now all overwhelmingly significant, while the release of the new figures for Business Inventories do not result in any increase in the volatility once the intradaily pattern in the absolute returns are filtered out.³⁹ As such, these results clearly demonstrate how the low-pass filtered absolute returns provide a valuable framework for the further study of the structural determinants behind financial market volatility clustering.

V. Conclusion

The temporal dependence in the volatility of speculative returns is of the utmost importance for the pricing and hedging of financial contracts. Yet, the empirical analysis of low frequency interdaily and high frequency intradaily returns have hitherto given rise to very different conclusions regarding the degree of volatility persistence for any particular asset. The mixture-of-distributions hypothesis developed here provides a justification for these conflicting empirical findings by interpreting the volatility as resulting from the aggregation of numerous constituent component processes; some with very short-run decay rates and others possessing much longer-run dependencies. When analyzing intradaily returns the short-run components will tend to dominate the estimates obtained with traditional time series models, whereas for daily or longer-run return intervals the estimates will be driven by the more persistent components. Nonetheless, under suitable conditions the aggregation of these multiple components implies, that the process for the volatility should exhibit the identical form of long-memory dependence irrespective of the

³⁸ In order to be compatible with the results for the low-pass filtered returns, the estimates are based on observations $t = 1441, 1442, \dots, 73440$, only.

³⁹ These findings are directly in line with the evidence reported in Andersen and Bollerslev (1997b), who rely on a flexible Fourier functional form in explicitly modeling the periodicity in the intradaily volatility along with the imposition of a declining weight structure for the announcement effects.

sampling intervals. This proposition is confirmed by our empirical analysis, which also demonstrates that, by annihilating the intradaily dependencies in time series of intradaily returns, it is possible to uncover this inherent long-memory dependence in relatively short calendar time-spans of high frequency data using simple correlation based procedures. As such the techniques discussed here set the stage for the development of improved long-run interdaily volatility forecasts based on the large samples of intradaily prices that have recently become available for a wide variety of different instruments. Only future research will reveal the extent to which these techniques will support the development of new and improved empirical pricing relationships.

REFERENCES

- Andersen, Torben G., 1994, Stochastic autoregressive volatility: A framework for volatility modeling, *Mathematical Finance* 4, 75–102.
- Andersen, Torben G., 1996, Return volatility and trading volume: An information flow interpretation of stochastic volatility, *Journal of Finance* 51, 169–204.
- Andersen, Torben G., and Tim Bollerslev, 1997a, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance*, forthcoming.
- Andersen, Torben G., and Tim Bollerslev, 1997b, DM-Dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies, *Journal of Finance*, forthcoming.
- Baillie, Richard T., 1996, Long-memory processes and fractional integration in econometrics, *Journal of Econometrics* 73, 5–59.
- Baillie, Richard T., and Tim Bollerslev, 1991, Intra-day and inter-market volatility in foreign exchange rates, *Review of Economic Studies* 58, 565–585.
- Baillie, Richard T., Tim Bollerslev, and Hans-Ole Mikkelsen, 1996, Fractionally integrated generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 74, 3–30.
- Baxter, Marianne, and Robert G. King, 1995, Measuring business cycles: Approximate band-pass filters for economic time series, Working paper, University of Virginia.
- Beran, Jan, 1994, *Statistics for Long-Memory Processes* (Chapman and Hall, New York).
- Bollerslev, Tim, Ray Y. Chou, and Kenneth F. Kroner, 1992, ARCH modeling in finance, *Journal of Econometrics* 52, 5–59.
- Bollerslev, Tim, and Ian Domowitz, 1993, Trading patterns and prices in the interbank foreign exchange market, *Journal of Finance* 48, 1421–1443.
- Bollerslev, Tim, and Robert F. Engle, 1993, Common persistence in conditional variances, *Econometrica* 61, 166–187.
- Bollerslev, Tim, and Hans-Ole Mikkelsen, 1996, Modeling and pricing long-memory in stock market volatility, *Journal of Econometrics* 73, 151–184.
- Breidt, F. Jay, Nuno Crato, and Pedro J. F. de Lima, 1997, On the detection and estimation of long memory in stochastic volatility, *Journal of Econometrics*, forthcoming.
- Chambers, Marcus J., 1995, Long memory and aggregation in macroeconomic time series, Working paper, University of Essex.
- Clark, Peter K., 1973, A subordinated stochastic process model with finite variance for speculative prices, *Econometrica* 41, 135–155.
- Dacorogna, Michel M., Ulrich A. Müller, Robert J. Nagler, Richard B. Olsen, and Olivier 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.
- Delgado, Miguel A., and Peter M. Robinson, 1994, New methods for the analysis of long memory time series: application to spanish inflation, *Journal of Forecasting* 13, 97–107.
- Ding, Zhuanxin, and Clive W. J. Granger, 1996, Modeling volatility persistence of speculative returns: A new approach, *Journal of Econometrics* 73, 185–215.
- Ding, Zhuanxin, Clive W. J. Granger, and Robert F. Engle, 1993, A long memory property of stock market returns and a new model, *Journal of Empirical Finance* 1, 83–106.

- Drost, Feike C., and Theo E. Nijman, 1993, Temporal aggregation of GARCH processes, *Econometrica* 61, 909–927.
- Ederington, Louis H., and Jae Ha Lee, 1993, How markets process information: News releases and volatility, *Journal of Finance* 48, 1161–1191.
- Epps, Thomas W., and Mary L. Epps, 1976, The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis, *Econometrica* 44, 305–321.
- Gallant, A. Ronald, and George E. Tauchen, 1996, Which moments to match, *Econometric Theory* 12, 657–681.
- Geweke, John, and Susan Porter-Hudak, 1983, The estimation and application of long-memory time series models, *Journal of Time Series Analysis* 4, 221–238.
- Ghose, Devajyoti, and Kenneth F. Kroner, 1996, Components of volatility in foreign exchange markets: An empirical analysis of high frequency data, Working paper, University of Arizona.
- Ghysels, Eric, and Joanna Jasiak, 1995, Trading patterns, time deformation, and stochastic volatility in foreign exchange markets, Working paper, University of Montreal.
- Goodhart, Charles A. E., S. G. Hall, S. G. B. Henry, and B. Pesaran, 1993, News effects in a high-frequency model of the sterling-dollar exchange rate, *Journal of Applied Econometrics* 8, 1–13.
- Goodhart, Charles A. E., Takatoshi Ito, and Richard Payne, 1996, One day in June 1993: A study of the working of Reuters' Dealing 2000-2 electronic foreign exchange trading system, in Jeffrey A. Frankel, Giampaolo Galli, and Alberto Giovannini, Eds.: *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago).
- Goodhart, Charles A. E., and Maureen O'Hara (1997), High frequency data in financial markets: Issues and applications, *Journal of Empirical Finance*, forthcoming.
- Granger, Clive W. J., 1980, Long memory relationships and the aggregation of dynamic models, *Journal of Econometrics* 14, 227–238.
- Granger, Clive W. J., 1988, Models that generate trends, *Journal of Time Series Analysis* 9, 329–343.
- Granger, Clive W. J., 1987, Generalized integrated processes, Working paper, University of California, San Diego.
- Granger, Clive W. J., and Zhuanxin Ding, 1996, Varieties of long memory models, *Journal of Econometrics* 73, 61–77.
- Granger, Clive W. J., and Zhuanxin Ding, 1997, Some properties of absolute returns: An alternative measure of risk, *Annales d'Economie et de Statistique*, forthcoming.
- Granger, Clive W. J., and Paul Newbold, 1976, Forecasting transformed series, *Journal of the Royal Statistical Society, Series B* 38, 189–203.
- Hamilton, James D., 1994, *Time Series Analysis* (Princeton University Press: Princeton, N.J.).
- Harris, Lawrence, 1986, A transaction data study of weekly and intraday patterns in stock returns, *Journal of Financial Economics* 16, 99–117.
- Harris, Lawrence, 1987, Transaction data tests of the mixture-of-distributions hypothesis, *Journal of Financial and Quantitative Analysis* 22, 127–141.
- Harvey, Andrew C., 1994, Long-memory in stochastic volatility, Working paper, London School of Economics.
- Harvey, Andrew C., and Mariane Streibel, 1996, Testing for a slowly changing level with special reference to stochastic volatility, Working paper, London School of Economics.
- Henry, Marc, and Richard Payne, 1996, An investigation of long range dependence in intra-day foreign exchange rate volatility, Working paper, London School of Economics.
- Hogan, Kedreth C., Jr., and Michael T. Melvin, 1994, Sources of heat waves and meteor showers in the foreign exchange markets, *Journal of International Economics* 37, 239–247.
- Ito, Takatoshi, and V. Vance Roley, 1987, News from the U.S. and Japan: Which moves the yen/dollar exchange rate?, *Journal of Monetary Economics* 19, 255–277.
- Ito, Takatoshi, Richard K. Lyons, and Michael T. Melvin, 1996, Is there private information in the FX market? The Tokyo experiment, Working paper, University of California at Berkeley.
- Jacquier, Eric, Nicholas G. Polson, and Peter E. Rossi, 1994, Bayesian analysis of stochastic volatility models, *Journal of Business and Economic Statistics* 12, 371–417.

- Johnson, Norman L., and Samuel Kotz, 1970, *Continuous Univariate Distributions*—2 (John Wiley and Sons, New York).
- Lamoureux, Christopher G., and William D. Lastrapes, 1990, Persistence in variance, structural change, and the GARCH model, *Journal of Business and Economic Statistics* 8, 225–234.
- Lin, Jin-Lung, 1991, Generalized integrated processes and the aggregation of dynamic time series, *Academia Economic Papers* 19, 207–226.
- Lobato, Ignacio N., and Peter M. Robinson, 1996, Averaged periodogram estimation of long-memory, *Journal of Econometrics* 73, 303–324.
- Mandelbrot, Benoit B., 1971, A fast fractional Gaussian noise generator, *Water Resources Research* 7, 543–553.
- Mandelbrot, Benoit B., and J. R. Wallis, 1969, Computer experiments with fractional Gaussian noises, *Water Resources Research* 5, 321–340.
- Mandelbrot, Benoit B., and J. W. van Ness, 1968, Fractional Brownian motions, fractional noises and applications, *SIAM Review* 10, 422–437.
- Merton, Robert C., 1980, On estimating the expected return on the market, *Journal of Financial Economics* 8, 323–361.
- Müller, Ulrich A., Michel M. Dacorogna, Rakhil D. Davé, Richard B. Olsen, Olivier V. Pictet, and Jakob E. von Weizsäcker, 1997, Volatilities of different time resolutions—Analyzing the dynamics of market components, *Journal of Empirical Finance*, forthcoming.
- Müller, Ulrich A., Michel M. Dacorogna, and Olivier V. Pictet, 1996, Heavy tails in high-frequency financial data, Working paper, Olsen & Associates, Research Institute for Applied Economics, Zürich, Switzerland.
- Müller, Ulrich A., Michel M. Dacorogna, Rakhil D. Davé, Olivier V. Pictet, Richard B. Olsen, and Robert Ward, 1993, Fractals and intrinsic time—A challenge to econometricians, Working paper, Olsen & Associates, Research Institute for Applied Economics, Zürich, Switzerland.
- Nelson, Daniel B., 1990, ARCH models as diffusion approximations, *Journal of Econometrics* 45, 7–38.
- Nelson, Daniel B., 1992, Filtering and forecasting with misspecified ARCH models: Getting the right variance with the wrong model, *Journal of Econometrics* 52, 61–90.
- Perron, Pierre, 1989, The great crash, the oil price shock, and the unit root hypothesis, *Econometrica* 57, 1361–1401.
- Priestley, Maurice B., 1981, *Spectral Analysis and Time Series* (Academic Press, London, UK).
- Robinson, Peter M., 1994a, Semiparametric analysis of long-memory time series, *Annals of Statistics* 22, 515–539.
- Robinson, Peter M., 1994b, Time series with strong dependence, in Christopher A. Sims, ed.: *Advances in Econometrics: Sixth World Congress* (Cambridge University Press, Cambridge, UK).
- Robinson, Peter M., 1995, Log-periodogram regression of time series with long-range dependence, *Annals of Statistics* 23, 1048–1072.
- Ross, Stephen A., 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–360.
- Ross, Stephen A., 1989, Information and volatility: The no-arbitrage martingale approach to timing and resolution irrelevancy, *Journal of Finance* 44, 1–18.
- Shephard, Neil, 1996, Statistical aspects of ARCH and stochastic volatility, in David R. Cox, David V. Hinkley, and Ole E. Barndorff-Nielsen, Eds.: *Likelihood, Time Series with Econometric and Other Applications* (Chapman and Hall, London, UK).
- Shiller, Robert J., and Pierre Perron, 1985, Testing the random walk hypothesis: Power versus frequency of observations, *Economics Letters* 18, 381–386.
- Tauchen, George E., and Mark Pitts, 1983, The price variability-volume relationship on speculative markets, *Econometrica* 51, 485–505.
- Taylor, Stephen J., 1994, Modeling stochastic volatility, *Mathematical Finance* 4, 183–204.
- Tieslau, Margie A., Peter Schmidt, and Richard T. Baillie, 1996, A minimum distance estimator for long-memory processes, *Journal of Econometrics* 71, 249–264.
- Wood, Robert A., Thomas H. McInish, and J. Keith Ord, 1985, An investigation of transaction data for NYSE stocks, *Journal of Finance* 25, 723–739.