

(Understanding, Optimizing, Using and Forecasting)

Realized Volatility and Correlation*

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

September 1999

This print/draft: October 26, 1999

* This work was supported by the National Science Foundation. We are grateful to Olsen & Associates, Zurich, for data, advice and hospitality.

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

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

^c Department of Finance, Stern School of Business, New York University, Departments of
Economics and Statistics, University of Pennsylvania, and NBER, phone: 610-585-4057,
e-mail: fdiebold@stern.nyu.edu

^d Graduate Group in Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA
19104-6297, phone: 919-929-2549, e-mail: labys@ssc.sas.upenn.edu

Volatility is central to many applied issues in finance and financial engineering, ranging from asset pricing and asset allocation to risk management. Hence, at least since the seminal contribution by Merton (1980) and later work by Nelson (1992), financial economists have been intrigued by the very high precision with which volatility can be estimated under the diffusion assumption routinely invoked in theoretical work. The basic insight follows from the observation that precise estimation of diffusion volatility does not require a long calendar span of data; rather, volatility can be estimated arbitrarily well from an arbitrarily short span of data, provided that returns are sampled sufficiently frequently. This contrasts sharply with precise estimation of the drift, which generally requires a long calendar span of data, regardless of the frequency with which returns are sampled.

Consequently, the volatility literature has steadily progressed toward the use of higher-frequency data. This is true in the parametric ARCH and stochastic volatility literatures (see Bollerslev, Engle and Nelson, 1994, for a review), as well as in the more traditional empirical finance literature which is more in line with our work. For example, Officer (1973) constructs annual volatilities from monthly returns on an equity index, whereas Merton (1980) and French, Schwert, and Stambaugh (1987) use daily returns to estimate monthly volatilities. Even more recently, Schwert (1998) relies on 15-minute returns for construction of daily stock market volatilities, while Taylor and Xu (1997) and Andersen, Bollerslev, Diebold and Labys (1999a) exploit 5-minute returns in the measurement of daily exchange rate volatilities.

Concurrent with the move toward the use of higher frequency data, recent work has clarified the comparative desirability of alternative volatility estimators. This is important as, for example, several different estimators were employed even in the small subset of the literature

briefly reviewed above. This emerging theory emphasizes the advantages of the so-called *realized volatility* estimator. In particular, Andersen and Bollerslev (1998) show that, under the usual diffusion assumptions, realized volatility computed from high-frequency intraday returns, now increasingly available, is effectively an error-free volatility measure. Moreover, construction of realized volatility is trivial – one simply sums intra-period high-frequency squared returns (or cross products, for realized covariance), period by period. For example, for a 24-hour market, daily realized volatility based on 5-minute underlying returns is defined as the sum of the 288 intra-day squared 5-minute returns, taken day by day.

Because the recent work on realized volatility cum high-frequency data concludes that realized volatility is, in principle, error-free, it is natural to treat volatility as observable. Observable volatility creates entirely new opportunities: we can analyze it, optimize it, use it and forecast it with much simpler techniques than the complex econometric models required when volatility is latent. Our recent work – and this survey – exploits this insight intensively. In Section 1 we describe our recent attempts at *understanding* both the unconditional and conditional distributions of realized asset return volatility, and in Section 2 we describe tools for *optimizing* the construction of realized volatility measures. In section 3, we put realized volatilities to work, *using* them to draw sharp inferences about the conditional distributions of asset returns. In Section 4 we discuss explicit modeling and *forecasting* of realized volatility, and we conclude in Section 5.

1. Understanding Realized Volatility and Correlation

Much of our work has focused on understanding exchange rate volatility dynamics. In particular, in Andersen, Bollerslev, Diebold and Labys (1999a) we use high-frequency data on

Deutschemark and Yen returns against the U.S. Dollar to construct model-free estimates of daily exchange rate volatility and correlation, spanning an entire decade. Although the study examines foreign exchange rate series in particular, preliminary results indicate that the patterns observed apply more broadly to other types of assets, including the thirty individual stocks in the Dow Jones Industrial Average, as studied in Andersen, Bollerslev, Diebold and Ebens (1999).

Figure 1 shows daily realized volatility for a representative asset return series of 1000 days. (Unless otherwise noted, all of the graphics reported here are designed to be representative of daily returns.) It is clear that realized volatility changes from day to day, as one expects. Furthermore, its fluctuations display substantial persistence.

Figure 2 shows representative unconditional distributions of realized volatilities and correlations. Although the distribution of the realized variance is clearly skewed (first panel), transforming to realized standard deviation moves it toward symmetry (second panel), and transforming to log standard deviations renders it approximately Gaussian (third panel).

Similarly, moving to multivariate aspects of the unconditional distribution, we find that realized covariance tends to be highly skewed, but that a simple transformation to correlation delivers approximate normality, as shown in Figure 3. Realized correlation is almost always positive, often strongly so, and it displays substantial variation. We also find that realized correlation is itself highly correlated with realized volatility, which we call the “volatility effect in correlation.” In particular, return correlations tend to rise on high-volatility days, as we illustrate in Figure 4.

Let us now move from unconditional to conditional aspects of the distributions of realized volatilities and correlations. Although correlograms of realized volatilities tend to exhibit a slow

hyperbolic decay, as shown in Figure 5, we routinely and soundly reject the unit-root hypothesis. However, such autocorrelation behavior is also consistent with that of fractionally integrated long-memory processes. In fact, there is strong evidence to suggest that volatility is a long-memory process, an assertion we substantiate through a variety of analyses. First, we estimate the long-memory parameter directly; the estimates tend to be in the neighborhood of 0.4 for a variety of realized volatility and correlation series, and the associated standard error is very small, in the neighborhood of 0.02. Second, we verify that the degree of fractional integration is invariant to the horizon, which is a well-known property of long-memory processes, due to their self-similarity. Finally, we verify that our realized volatility and correlation series follow scaling laws, such that the logs of the variance of partial sums of the process are proportional to the logs of the horizon, which is also a well-known characteristic of long-memory processes.

Much has been learned, but much remains to be done. For example, although access to the high-frequency data necessary for construction of accurate realized volatilities is increasing rapidly, it is far from universal, and we need simple and operational ways of characterizing the measurement error remaining in realized volatilities constructed from insufficiently frequently sampled data. Moreover, even when high-frequency data is available, microstructure effects such as bid/ask bounce and asynchronous trading may distort associated realized volatilities. We now turn to a tool for identifying and mitigating such effects.

2. Optimizing Realized Volatility and Correlation

The appeal of realized volatility computed from high-frequency data relies at least partially on the assumption that log asset prices evolve as diffusions. This assumption, although adequate and useful in many contexts, becomes progressively less tolerable as transaction time is

approached and market microstructure effects emerge. Hence a tension arises: the optimal sampling frequency will likely not be the highest available, but rather some intermediate value, ideally high enough to produce a volatility estimate with negligible sampling variation, yet low enough to avoid microstructure bias. The choice of underlying return frequency is therefore critical, but the literature currently offers little guidance for making that decision.

In Andersen, Bollerslev, Diebold and Labys (1999b), we develop a tool designed to provide some guidance. On the one hand our motivation is highly pragmatic, as we seek to determine how one should choose the underlying sampling frequency for calculating the realized volatility of financial asset returns. Towards this end, we propose and illustrate a simple graphical diagnostic which we call a “volatility signature plot.” On the other hand our motivation is also rather lofty, as we attempt to characterize different market microstructures in terms of their volatility signatures – the patterns of bias injected in realized volatility as underlying returns are sampled progressively more frequently.

A key insight is that microstructure bias, if operative, will likely manifest itself as sampling frequency increases by distorting the average realized volatility. Consequently, a plot of average realized volatility against sampling frequency may help reveal the severity of microstructure bias as sampling frequency increases, and may therefore be useful in guiding the selection of sampling frequency. Interestingly, it turns out that the volatility signature is isomorphic to the variance-time function, which has been extensively studied in finance. However, notwithstanding the fact that there is no information in the volatility signature that is not also present in high-frequency return autocorrelations, the two are complements, not substitutes, as the information relevant for construction and interpretation of realized volatilities is more directly and transparently revealed

in the volatility signature plot.

In Figure 6 we show two representative volatility signature plots – plots of average realized volatility against sampling interval k . The integer k represents multiples of the smallest sampling interval in the data; thus, if we have a series for which the smallest available sampling interval is one minute, for $k=1$ we construct average realized volatility using 1-minute returns, for $k=2$ we construct average realized volatility using 2-minute returns, and so forth. The top panel of Figure 6 represents a highly liquid asset for which the largest realized volatility estimates occur at the highest sampling rates, corresponding to the smallest values of k . This can be explained by negative serial correlation in the returns, most likely induced by bid-ask bounce: at the smallest sampling intervals, the volatility measures are very high, but as returns are aggregated across larger and larger sampling intervals the oscillating swings in the returns series tend to cancel, and overall volatility is lower. The volatility signature plot stabilizes at roughly $k=20$ (in this case corresponding to a 20-minute return sampling interval). Although high-frequency microstructural effects will also be small for sampling intervals larger than $k=20$, realized volatility estimates constructed from larger return intervals will begin to suffer from a higher sampling error. Thus, for this particular example, we would recommend the use of a sampling interval of $k=20$, which represents a reasonable tradeoff between minimizing microstructural bias and minimizing sampling error.

The bottom panel of Figure 6 represents a less liquid asset, whose volatility signature is quite different from that of the asset in the top panel. In this case, microstructural factors cause a positive serial correlation at high frequencies, resulting in a smaller estimate of realized volatility, which doesn't stabilize until the sampling interval reaches $k=15$, or 15 minutes. In this case, the

microstructure bias is likely induced by inactive trading.

Again, much remains to be done, including extensions of signature plots to multivariate and multimoment environments. Nevertheless, we feel confident that high-quality realized volatilities can be constructed in liquid markets, and we are optimistic regarding the potential for utilizing volatility signature plots to assist in the choice of underlying sampling frequency.

3. Using Realized Volatility and Correlation

Here we report on recent work that uses realized volatilities to study the conditional distribution of asset returns. In particular, Andersen, Bollerslev, Diebold, and Labys (1999c) characterize the distribution and temporal dependence of $\varepsilon_t = \frac{r_t}{\sigma_t}$, which we call the σ -standardized return. (σ denotes the realized standard deviation.)

There is a long tradition in the econometrics literature of needing and allowing for a fat-tailed conditional distribution of ε_t , as in Bollerslev (1987). But that literature typically works with returns standardized by volatilities obtained from discrete-time ARCH or stochastic volatility models, in which day-t volatility depends only on information at day t-1 and earlier. The situation is different with realized volatility; day-t realized volatility is based on information within day t, and the theoretical predictions for distributions of returns standardized by realized volatility are unambiguous: under the diffusion assumption they should be Gaussian.

This is, in fact, what we tend to find. The top panel of Figure 7 displays a representative QQ plot for unstandardized returns; because the points don't fall into a straight line, we conclude that the returns are not distributed normally. The bottom panel of Figure 7 displays a representative QQ plot for σ -standardized returns; it is close to linear, indicating that the σ -standardized returns are approximately Gaussian.

As always, much remains to be done. It is, for example, of practical importance to examine the distribution of returns standardized by *forecasts* of realized volatility, based on prior information only. On the theoretical side, it will be of interest to develop more formal tests for the presence of jumps from the distribution of the returns standardized by the realized volatility.

4. Forecasting Realized Volatility and Correlation

Our construction, optimization and use of realized volatilities has helped us reach two general conclusions: realized variances tend to be lognormally distributed, and asset returns standardized by realized standard deviations tend to be normally distributed. In turn, this suggests that a lognormal-normal mixture may be a good model for asset returns, an idea which inspires a simple modeling and forecasting strategy. This section briefly outlines an operational procedure for the univariate case, but extensions to the multivariate setting, exploiting realized covariances as well as volatilities, are straightforward and currently under investigation in Andersen, Bollerslev, Diebold, and Labys (1999d).

In essence, forecasting return volatility is equivalent to forecasting realized volatility (as long as high-quality intraday return data are available). Because realized volatility is effectively observed, it is amenable to direct analysis via standard time series methods. Given the stylized facts above, it is natural to assume that the log-volatility process falls within the usual Gaussian ARMA class of models. However, we noted the long-memory characteristics of the realized volatility series in section 1. Consequently, it is desirable to allow for fractional integration in the specification, leading to a so-called ARFIMA model.

Concretely, one may proceed as follows. First, determine the degree of fractional integration, d , in the realized log-volatility series – as noted above, the typical estimates suggest

a value of d around 0.4. Next, obtain the fractionally differenced series, say, $y_t = (1-L)^d \log \sigma_t$. This operation involves computing a long (in theory infinite, in practice long, but truncated) distributed lag of the underlying log-volatility series. This transformation ideally removes the long-run dependence in the series. For illustration, we display such a fractionally differenced log-volatility series in Figure 8 and the associated correlogram in Figure 9. They provide remarkable contrasts to Figures 1 and 5; any indication of long memory has indeed been annihilated. Hence, the final step of the modeling procedure is to obtain a parsimonious Gaussian ARMA representation for this fractionally differenced (residual) series.

Standard AR(FI)MA procedures may now be applied to generate predictions of future realized log-volatility based on the estimated model and the observed series. The result is a sequence of volatility forecasts with associated prediction errors that are (approximately) lognormally distributed. Moreover, because returns are normally distributed conditional on realized volatility, one may readily compute the fractiles of the conditional return distribution in closed form from the standard lognormal-normal mixture distribution.

The striking feature of this approach is that it builds directly on observed time series and utilizes only standard linear Gaussian modeling and forecasting techniques. Hence, it is simple to assess in-sample performance and evaluate model forecasts through well established out-of-sample procedures. It will be interesting in future work to investigate the actual performance of such an approach relative to popular frameworks such as ARCH, stochastic volatility and RiskMetrics. Because our approach exploits an arguably superior volatility measure along with more sound distributional assumptions, one may conjecture that it will outperform the standard procedures currently in use.

5. Concluding Remarks

Our findings have potentially wide-ranging implications for applied finance. The results on the unconditional and conditional distributions of asset return volatility are relevant for pricing derivative instruments. In fact, with the advent of volatility and covariance swaps, realized volatility itself is now the underlying! Such swaps are useful for, among others, holders of options who wish to vega hedge their holdings, i.e., offset the impact of changes in volatility on the value of their positions, as discussed for example in Demeterfi, Derman, Kamal and Zou (1999). Proper pricing of derivatives on volatility depends critically on how volatility itself varies over time (“the volatility of volatility”). Our methodology allows for a direct approach to this issue through the construction and analysis of historical realized volatility series.

Improved volatility and correlation forecasts will also be useful for portfolio allocation and management. Concrete indications that more traditional volatility forecasts can be of value in guiding portfolio allocation decisions are provided by Fleming, Kirby and Ostdiek (1999). To the extent that our procedures are able to improve on the volatility forecast performance, the implied economic benefits could be high.

Finally, our forecasting procedures for realized volatility and correlation lead directly to a characterization of the conditional return distribution (ignoring significant short-term variation in the conditional mean). The evaluation of fractiles of the conditional return distribution is, of course, a critical input into any active financial risk management program. Hence, extensions of our methodology to a richer multi-asset setting should provide potentially valuable inputs for practical risk management.

References

- Andersen, T. and T. Bollerslev (1998), "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885-905.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and H. Ebens (1999), "The Distribution of U.S. Stock Return Volatility," Manuscript in progress.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (1999a), "The Distribution of Exchange Rate Volatility," Wharton Financial Institutions Center Working Paper 99-08 and NBER Working Paper 6961.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (1999b), "Microstructure Bias and Volatility Signatures," Manuscript in progress.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (1999c), "Exchange Rate Returns Standardized by Realized Volatility are (Nearly) Gaussian," Manuscript in progress.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (1999d), "Forecasting Realized Volatility: A VAR for VaR," Manuscript in progress.
- Bollerslev, T. (1987), "A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *Review of Economics and Statistics*, 69, 542-547.
- Bollerslev, T., R.F. Engle and D.B. Nelson (1994), "ARCH Models," in R.F. Engle and D. McFadden (eds.), *Handbook of Econometrics, Volume IV*, 2959-3038. Amsterdam: North-Holland.
- Demeterfi, K., E. Derman, M. Kamal and J. Zou (1999), "More Than You Ever Wanted To Know About Volatility Swaps (but Less Than can be Said)," Goldman Sachs & Co., Quantitative Strategies Research Notes.
- Fleming, J., C. Kirby and B. Ostdiek (1999), "The Economic Value of Volatility Timing," Manuscript, Department of Finance, Rice University.
- French, K.R., G.W. Schwert and R.F. Stambaugh (1987), "Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3-29.
- Merton, R.C. (1980), "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323-361.
- Nelson, D.B. (1992), "Filtering and Forecasting with Misspecified ARCH Models I: Getting the Right Variance with the Wrong Model," *Journal of Econometrics*, 52, 61-90.

Officer, R.R. (1973), "The Variability of the Market Factor of the NYSE," *Journal of Business*, 46, 434-453.

Schwert, G.W. (1998), "Stock Market Volatility: Ten Years After the Crash," *Brookings-Wharton Papers on Financial Services*, 1998. Washington, DC: Brookings Institution.

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

Figure 1
Representative Realized Volatility Series




Figure 2
Representative Distributions
of Realized Volatility


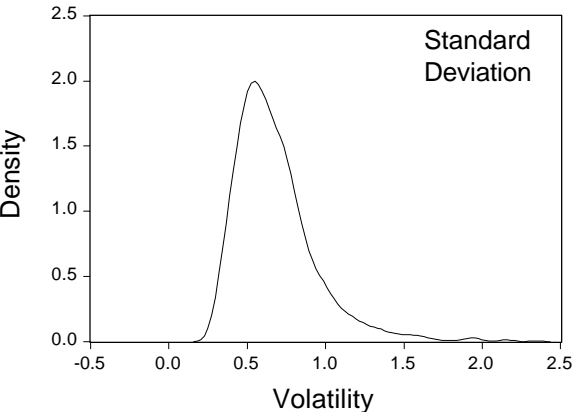
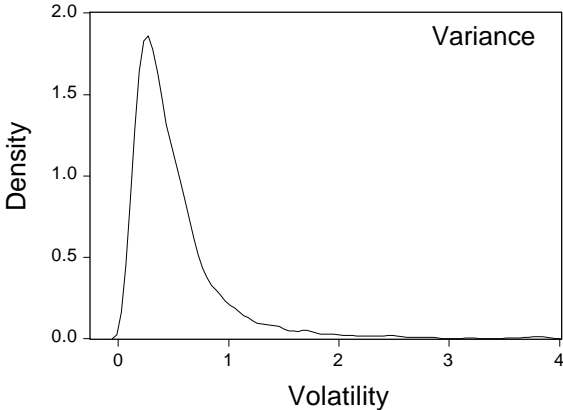


Figure 3
Representative Distributions
of Realized Covariance and Correlation

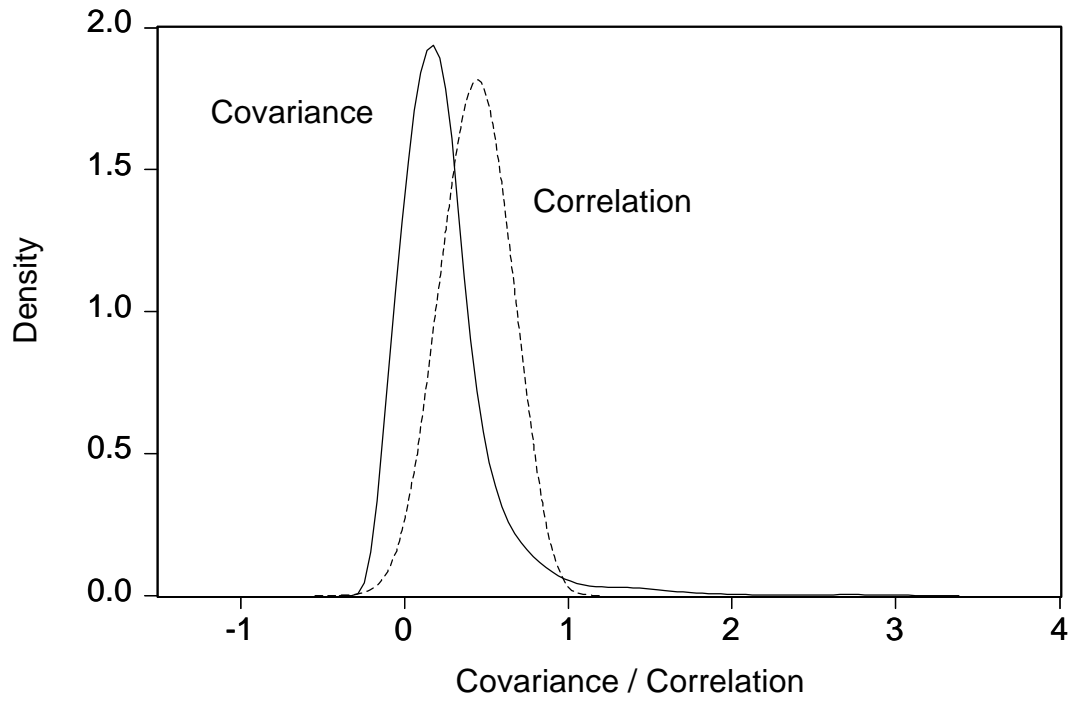


Figure 4
Representative Distributions of Realized Correlation
in Low Volatility vs. High Volatility Periods

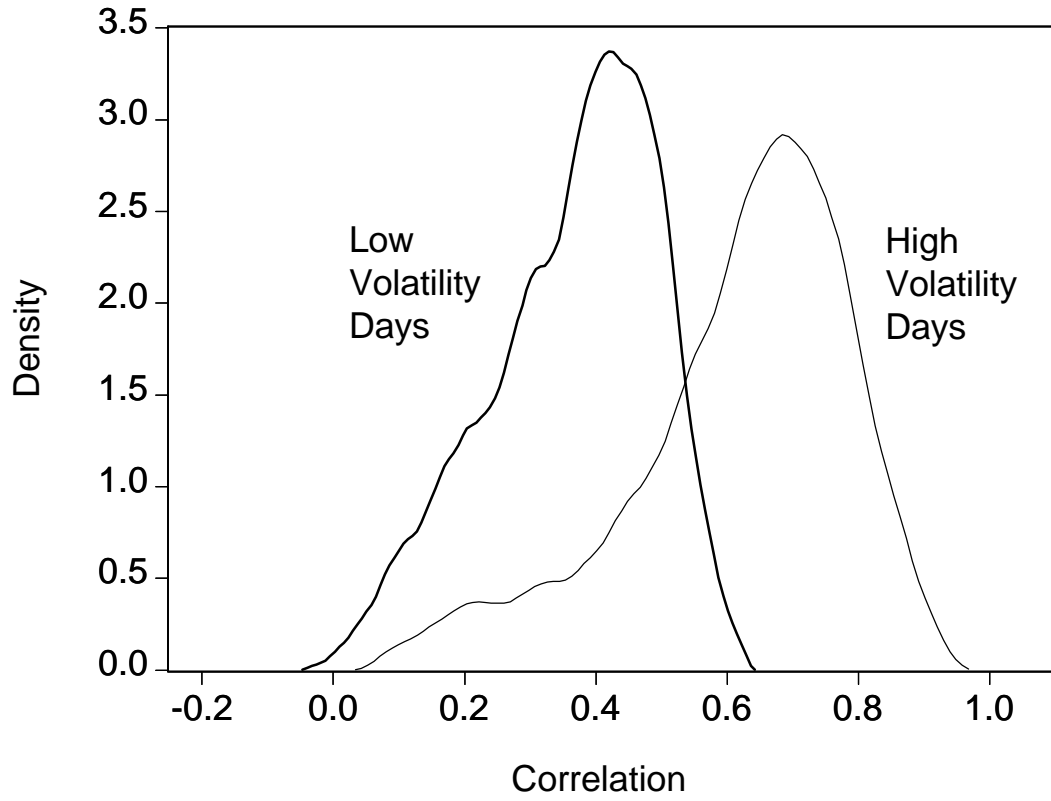


Figure 5
Representative Sample Autocorrelations
of Realized Volatility or Correlation

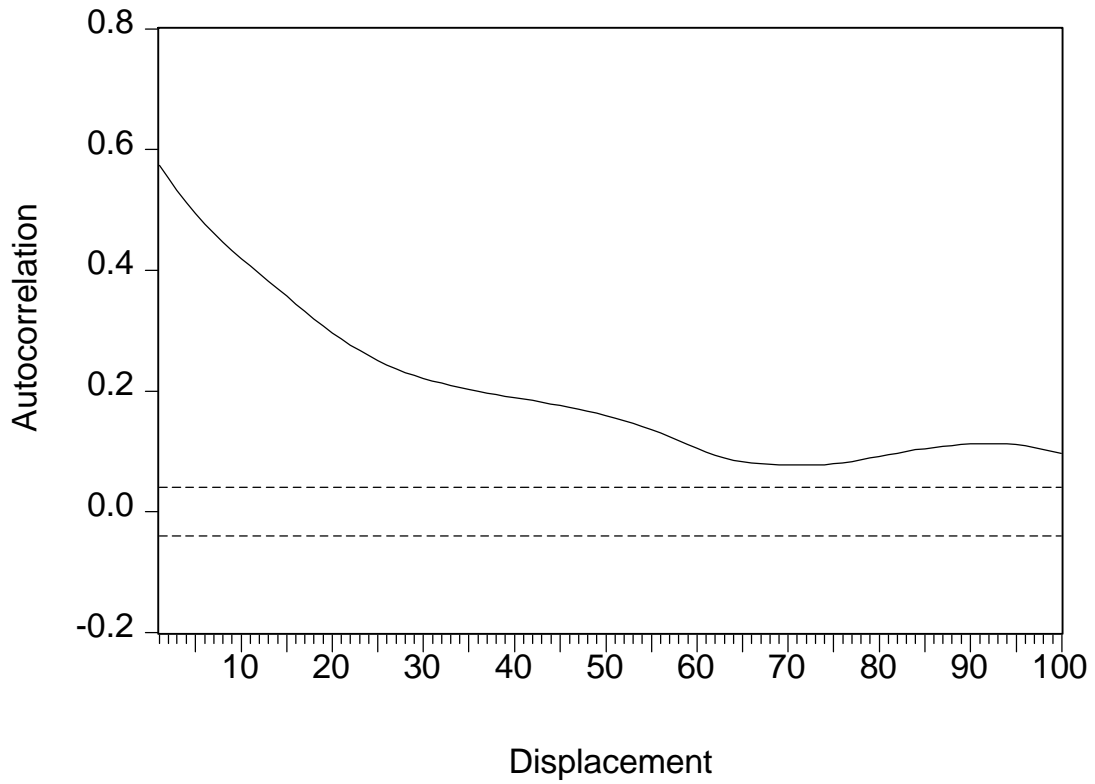


Figure 6
Representative Volatility Signature Plots
Liquid and Illiquid Assets





Figure 7
Representative QQ Plots
Unstandardized and σ -Standardized Returns





Figure 8
Representative Fractionally-Differenced Realized Volatility Series




Figure 9
Representative Sample Autocorrelations
of Fractionally-Differenced Realized Volatility or Correlation

