

The distribution of realized stock return volatility[☆]

Torben G. Andersen^{a,d,1}, Tim Bollerslev^{b,d,2},
Francis X. Diebold^{c,d,*}, Heiko Ebens^{e,3}

^a *Department of Finance, Kellogg Graduate School of Management, Northwestern University,
Evanston, IL 60208, USA*

^b *Department of Economics, Duke University, Durham, NC 27708, USA*

^c *Department of Economics, University of Pennsylvania, Philadelphia, PA 19104-6297, USA*

^d *NBER, USA*

^e *Department of Economics, Johns Hopkins University, Baltimore, MD 21218, USA*

Received 29 December 1999; received in revised form 9 July 2000; accepted 16 February 2001

Abstract

We examine “realized” daily equity return volatilities and correlations obtained from high-frequency intraday transaction prices on individual stocks in the Dow Jones

[☆]This work was supported by the National Science Foundation. We thank the editor and referee for several suggestions that distinctly improved this paper. Helpful comments were also provided by Dave Backus, Michael Brandt, Rohit Deo, Rob Engle, Clive Granger, Lars Hansen, Joel Hasbrouck, Ludger Hentschel, Cliff Hurvich, Pedro de Lima, Bill Schwert, Rob Stambaugh, George Tauchen, and Stephen Taylor, as well as seminar and conference participants at the 1999 North American Winter Meetings and European Summer Meetings of the Econometric Society, the May 1999 NBER Asset Pricing Meeting, Boston University, Columbia University, Johns Hopkins University, London School of Economics, New York University, Olsen & Associates, the Triangle Econometrics Workshop, and the University of Chicago. Much of this paper was written while Diebold visited the Stern School of Business, New York University, whose hospitality is gratefully acknowledged.

*Corresponding author. Tel.: +1-215-898-1507; fax: +1-215-573-4217.

E-mail addresses: t-andersen@nwu.edu (T.G. Andersen), boller@econ.duke.edu (T. Bollerslev), fdiebold@sas.upenn.edu (F.X. Diebold), ebens@jhu.edu (H. Ebens).

¹Tel.: +1-847-467-1285.

²Tel.: +1-919-660-1846.

³Tel.: +1-410-516-7601.

Industrial Average. We find that the unconditional distributions of realized variances and covariances are highly right-skewed, while the realized logarithmic standard deviations and correlations are approximately Gaussian, as are the distributions of the returns scaled by realized standard deviations. Realized volatilities and correlations show strong temporal dependence and appear to be well described by long-memory processes. Finally, there is strong evidence that realized volatilities and correlations move together in a manner broadly consistent with latent factor structure. © 2001 Elsevier Science S.A. All rights reserved.

JEL Classification: G10; C10; C20

Keywords: Integrated volatility; Correlation; Equity markets; High-frequency data; Long memory

1. Introduction

Financial market volatility is central to the theory and practice of asset pricing, asset allocation, and risk management. Although most textbook models assume volatilities and correlations to be constant, it is widely recognized among both finance academics and practitioners that they vary importantly over time. This recognition has spurred an extensive and vibrant research program into the distributional and dynamic properties of stock market volatility.⁴ Most of what we have learned from this burgeoning literature is based on the estimation of parametric ARCH or stochastic volatility models for the underlying returns, or on the analysis of implied volatilities from options or other derivatives prices. However, the validity of such volatility measures generally depends upon specific distributional assumptions, and in the case of implied volatilities, further assumptions concerning the market price of volatility risk. As such, the existence of multiple competing models immediately calls into question the robustness of previous findings. An alternative approach, based for example on squared returns over the relevant return horizon, provides model-free unbiased estimates of the ex post realized volatility. Unfortunately, however, squared returns are also a very

⁴For an early survey, see Bollerslev et al. (1992). A selective and incomplete list of studies since then includes Andersen (1996), Bekaert and Wu (2000), Bollerslev and Mikkelsen (1999), Braun et al. (1995), Breidt et al. (1998), Campbell and Hentschel (1992), Campbell et al. (2000), Canina and Figlewski (1993), Cheung and Ng (1992), Christensen and Prabhala (1998), Day and Lewis (1992), Ding et al. (1993), Duffee (1995), Engle and Ng (1993), Engle and Lee (1993), Gallant et al. (1992), Glosten et al. (1993), Hentschel (1995), Jacquier et al. (1994), Kim and Kon (1994), Kroner and Ng (1998), Kuwahara and Marsh (1992), Lamoureux and Lastrapes (1993), and Tauchen et al. (1996).

noisy volatility indicator and hence do not allow for reliable inference regarding the true underlying latent volatility.

The limitations of the traditional procedures motivate our different approach for measuring and analyzing the properties of stock market volatility. Using continuously recorded transactions prices, we construct estimates of ex post *realized* daily volatilities by summing squares and cross-products of intraday high-frequency returns. Volatility estimates so constructed are model-free, and as the sampling frequency of the returns approaches infinity, they are also, in theory, free from measurement error (Andersen, Bollerslev, Diebold and Labys, henceforth ABDL, 2001a).⁵ The need for reliable high-frequency return observations suggests, however, that our approach will work most effectively for actively traded stocks. We focus on the 30 stocks in the Dow Jones Industrial Average (DJIA), both for computational tractability and because of our intrinsic interest in the Dow, but the empirical findings carry over to a random sample of 30 other liquid stocks. In spite of restricting the analysis to actively traded stocks, market microstructure frictions, including price discreteness, infrequent trading, and bid–ask bounce effects, are still operative. In order to mitigate these effects, we use a five-minute return horizon as the effective “continuous time record”. Treating the resulting daily time series of realized variances and covariances constructed from a five-year sample of five-minute returns for the 30 DJIA stocks as being directly observable allows us to characterize the distributional features of the volatilities without attempting to fit multivariate ARCH or stochastic volatility models.

Our approach is directly in line with earlier work by French et al. (1987), Schwert (1989, 1990a,b), and Schwert and Seguin (1991), who rely primarily on daily return observations for the construction of monthly realized stock volatilities.⁶ The earlier studies, however, do not provide a formal justification for such measures, and the diffusion-theoretic underpinnings provided here explicitly hinge on the length of the return horizon approaching zero. Intuitively, following the work of Merton (1980) and Nelson (1992), for a continuous time diffusion process, the diffusion coefficient can be estimated arbitrarily well with sufficiently finely sampled observations, and by the theory of quadratic variation, this same idea carries over to estimates of the integrated volatility over

⁵Nelson (1990, 1992) and Nelson and Foster (1994) obtain a related but different result: misspecified ARCH models may work as consistent filters for the latent *instantaneous* volatility as the return horizon approaches zero. Similarly, Ledoit and Santa-Clara (1998) show that the Black-Scholes implied volatility for an at-the-money option provides a consistent estimate of the underlying latent instantaneous volatility as the time to maturity approaches zero.

⁶In a related analysis of monthly U.S. stock market volatility, Campbell et al. (2000) augment the time series of monthly sample standard deviations with various alternative volatility measures based on the dispersion of the returns on individual stocks in the market index.

fixed horizons. As such, the use of high-frequency returns plays a critical role in justifying our measurements. Moreover, our focus centers on *daily*, as opposed to monthly, volatility measures. This mirrors the focus of most of the extant academic and industry volatility literatures and more clearly highlights the important intertemporal volatility fluctuations.⁷ Finally, because our methods are trivial to implement, even in the high-dimensional situations relevant in practice, we are able to study the distributional and dynamic properties of *correlations* in much greater depth than is possible with traditional multivariate ARCH or stochastic volatility models, which rapidly become intractable as the number of assets grows.

Turning to the results, we find it useful to segment them into unconditional and conditional aspects of the distributions of volatilities and correlations. As regards the unconditional distributions, we find that the distributions of the realized daily variances are highly non-normal and skewed to the right, but that the logarithms of the realized variances are approximately normal. Similarly, although the unconditional distributions of the covariances are all skewed to the right, the realized daily correlations appear approximately normal. Finally, although the unconditional daily return distributions are leptokurtic, the daily returns normalized by the realized standard deviations are also close to normal. Rather remarkably, these results hold for the vast majority of the 30 volatilities and 435 covariances/correlations associated with the 30 Dow Jones stocks, as well as the 30 actively traded stocks in our randomly selected control sample.

Moving to conditional aspects of the distributions, all of the volatility measures fluctuate substantially over time, and all display strong dynamic dependence. Moreover, this dependence is well-characterized by slowly mean-reverting fractionally integrated processes with a degree of integration, d , around 0.35, as further underscored by the existence of very precise scaling laws under temporal aggregation. Although statistically significant, we find that the much debated leverage effect, or asymmetry in the relation between past negative and positive returns and future volatilities, is relatively unimportant from an economic perspective. Interestingly, the same type of asymmetry is also present in the realized correlations. Finally, there is a systematic tendency for the variances to move together, and for the correlations among the different stocks to be high/low when the variances for the underlying stocks are high/low, and when the correlations among the other stocks are also high/low.

Although several of these features have been documented previously for U.S. equity returns, the existing evidence relies almost exclusively on the estimation

⁷Schwert (1990a), Hsieh (1991), and Fung and Hsieh (1991) also study daily standard deviations based on 15-minute equity returns. However, their analysis is strictly univariate and decidedly less broad in scope than ours.

of specific parametric volatility models. In contrast, the stylized facts for the 30 DJIA stocks documented here are explicitly model-free. Moreover, the facts extend the existing results in important directions and both solidify and expand on the more limited set of results for the two exchange rates in ABDL (2001a,b) and the DJIA stock index in Ebens (1999a). As such, our findings set the stage for the development of improved volatility models—possibly involving a simple factor structure, which appears consistent with many of our empirical findings—and corresponding out-of-sample volatility forecasts, consistent with the distributional characteristics of the returns.⁸ Of course, the practical use of such models in turn should allow for better risk management, portfolio allocation, and asset pricing decisions.

The remainder of the paper is organized as follows. In Section 2 we provide a brief account of the diffusion-theoretic underpinnings of our realized volatility measures, along with a discussion of the actual data and volatility calculations. In Section 3 we discuss the unconditional univariate return volatility and correlation distributions, and we move to dynamic aspects, including long-memory effects and scaling laws, in Section 4. In Section 5 we assess the symmetry of responses of realized volatilities and correlations to unexpected shocks. We report on multivariate aspects of the volatility and correlation distributions in Section 6, and in Section 7 we illustrate the consistency of several of our empirical results with a simple model of factor structure in volatility. We conclude in Section 8 with a brief summary of our main findings and some suggestions for future research.

2. Realized volatility measurement

2.1. Theory

Here we provide a discussion of the theoretical justification behind our volatility measurements. For a more thorough treatment of the pertinent issues within the context of special semimartingales we refer to ABDL (2001a) and the general discussion of stochastic integration in Protter (1992). To set out the basic idea and intuition, assume that the logarithmic $N \times 1$ vector price process, p_t , follows a multivariate continuous time stochastic volatility diffusion,

$$dp_t = \mu_t dt + \Omega_t dW_t, \quad (1)$$

where W_t denotes a standard N -dimensional Brownian motion, the process for the $N \times N$ positive definite diffusion matrix, Ω_t , is strictly stationary, and we

⁸Ebens (1999a), for example, makes an initial attempt at modeling univariate realized stock volatility for the DJIA index.

normalize the unit time interval, or $h=1$, to represent one trading day. Conditional on the sample path realization of μ_t and Ω_t , the distribution of the continuously compounded h -period returns, $r_{t+h,h} \equiv p_{t+h} - p_t$, is then

$$r_{t+h,h} | \sigma \{ \mu_{t+\tau}, \Omega_{t+\tau} \}_{\tau=0}^h \sim N \left(\int_0^h \mu_{t+\tau} d\tau, \int_0^h \Omega_{t+\tau} d\tau \right), \quad (2)$$

where $\sigma \{ \mu_{t+\tau}, \Omega_{t+\tau} \}_{\tau=0}^h$ denotes the σ -field generated by the sample paths of $\mu_{t+\tau}$ and $\Omega_{t+\tau}$ for $0 \leq \tau \leq h$. The integrated diffusion matrix thus provides a natural measure of the true latent h -period volatility. This notion of integrated volatility already plays a central role in the stochastic volatility option pricing literature, where the price of an option typically depends on the distribution of the integrated volatility process for the underlying asset over the life of the option.⁹

By the theory of quadratic variation, we have that under weak regularity conditions,

$$\sum_{j=1, \dots, [h/\Delta]} r_{t+j\Delta, \Delta} r'_{t+j\Delta, \Delta} - \int_0^h \Omega_{t+\tau} d\tau \rightarrow 0 \quad (3)$$

almost surely for all t as the sampling frequency of the returns increases, or $\Delta \rightarrow 0$. Thus, by summing sufficiently finely sampled high-frequency returns, it is possible to construct ex post *realized* volatility measures for the integrated latent volatilities that are asymptotically free of measurement error.¹⁰ This contrasts sharply with the common use of the cross-product of the h -period returns, $r_{t+h,h} r'_{t+h,h}$, as a simple ex post volatility measure. Although the squared return over the forecast horizon provides an unbiased estimate for the realized integrated volatility, it is an extremely noisy estimator, and predictable variation in the true latent volatility process is typically dwarfed by measurement error.¹¹ Moreover, for longer horizons any conditional mean dependence will tend to contaminate this variance measure. In contrast, as the length of the return horizon decreases the impact of the drift term vanishes, so that the mean is effectively annihilated.

These assertions remain valid if the underlying continuous time process in Eq. (1) contains jumps, so long as the price process is a special semimartingale, which will hold if it is arbitrage-free (see, for example, Back, 1991). Of course, in this case the limit of the summation of the high-frequency returns will

⁹ See, for example, the well-known contribution of Hull and White (1987).

¹⁰ Consider the simple case of univariate discretely sampled i.i.d. normal zero-mean returns; $N=1$, $\mu_t=0$, and $\Omega_t=\sigma^2$. It follows by standard arguments that $E(h^{-1} \sum_{j=1, \dots, [h/\Delta]} r_{t+j\Delta, \Delta}^2) = \sigma^2$, while $\text{Var}(h^{-1} \sum_{j=1, \dots, [h/\Delta]} r_{t+j\Delta, \Delta}^2) = (\Delta/h) 2\sigma^4 \rightarrow 0$, as $\Delta \rightarrow 0$.

¹¹ In empirically realistic situations, the variance of $r_{t+1,1} r'_{t+1,1}$ is easily 20 times the variance of the true daily integrated volatility, $\int_0^1 \Omega_{t+\tau} d\tau$. Andersen and Bollerslev (1998) provide numerical results along these lines.

involve an additional jump component, but the interpretation of the sum as the realized h -period return volatility remains intact; for further discussion along these lines see ABDL (2001a). Importantly, in the presence of jumps the conditional distribution of the returns in Eq. (2) is no longer Gaussian. As such, the corresponding empirical distribution of the standardized returns speaks directly to the relevance of allowing for jumps in the underlying continuous time process when analyzing the returns over longer h -period horizons. Of course, viewed as a non-parametric omnibus test for jumps, this may not be a very powerful procedure.¹²

2.2. Data

Our empirical analysis is based on data from the Trade And Quotation (TAQ) database. The TAQ data files contain continuously recorded information on the trades and quotations for the securities listed on the New York Stock Exchange (NYSE), American Stock Exchange (AMEX), and the National Association of Security Dealers Automated Quotation system (NASDAQ). The database is published monthly, and has been available on CD-ROM from the NYSE since January 1993; we refer the reader to the corresponding data manual for a more complete description of the actual data and the method of data capture. Our sample extends from January 2, 1993 until May 29, 1998, for a total of 1,366 trading days. A complete analysis based on all trades for all stocks, although straightforward conceptually, is infeasible in practice. We therefore restrict our analysis to the 30 DJIA firms, which also helps to ensure a reasonable degree of liquidity. A list of the relevant ticker symbols as of the reconfiguration of the DJIA index in March 1997 is contained in Andersen, Bollerslev, Diebold and Ebens (2001) (henceforth ABDE).

Although the DJIA stocks are among the most actively traded U.S. equities, the median inter-trade duration for all stocks across the full sample is 23.1 seconds, ranging from a low of 7 seconds for Merck & Co. Inc. (MRK) to a high of 54 seconds for United Technologies Corp. (UTX). As such, it is not practically feasible to push the continuous record asymptotics and the length of the observation interval Δ in Eq. (3) beyond this level. Moreover, because of the organizational structure of the market, the available quotes and transaction prices are subject to discrete clustering and bid–ask bounce effects. Such market microstructure features are generally not important when analyzing longer-horizon interdaily returns but can seriously distort the distributional properties of high-frequency intraday returns; see, for example, the textbook

¹²A similar idea underlies the test for jumps in Drost et al. (1998), based on a comparison of the sample kurtosis and the population kurtosis implied by a continuous time GARCH (1,1) model; see also ABDL (2001b).

treatment by Campbell et al. (1997). Thus, following the analysis in Andersen and Bollerslev (1997), we rely on artificially constructed five-minute returns.¹³ With the daily transaction record extending from 9:30 EST until 16:05 EST, there are a total of 79 five-minute returns for each day, corresponding to $\Delta = 1/79 \approx 0.0127$ in the notation above. The five-minute horizon is short enough that the accuracy of the continuous record asymptotics underlying our realized volatility measures work well, and long enough that the confounding influences from market microstructure frictions are not overwhelming; see ABDL (1999) for further discussion along these lines.¹⁴

2.3. Construction of realized equity volatilities

The five-minute return series are constructed from the logarithmic difference between the prices recorded at or immediately before the corresponding five-minute marks. Although the limiting result in Eq. (3) is independent of the value of the drift parameter, μ_t , the use of a fixed discrete time interval could allow dependence in the mean to systematically bias our volatility measures. Thus, in order to purge the high-frequency returns of the negative serial correlation induced by the uneven spacing of the observed prices and the inherent bid–ask spread, we first estimate an MA(1) model for each of the five-minute return series using the full five-year sample. Consistent with the spurious dependence that would be induced by non-synchronous trading and bid–ask bounce effects, all estimated moving-average coefficients are negative, with a median value of -0.214 across the 30 stocks. We denote the corresponding 30 de-meaned MA(1)-filtered return series of $79 \times 1,366 = 107,914$ five-minute returns by $r_{t+\Delta, \Delta}$.¹⁵ Finally, to avoid any confusion, we denote the daily unfiltered raw returns by a single time subscript, i.e., r_t where $t = 1, 2, \dots, 1,336$.

The realized daily covariance matrix is then

$$Cov_t \equiv \sum_{j=1, \dots, 1/\Delta} r_{t+j\Delta, \Delta} r'_{t+j\Delta, \Delta}, \quad (4)$$

where $t = 1, 2, \dots, 1,366$ and $\Delta = 1/79$. For notational simplicity we refer to realized daily variances given by the diagonal elements as $v_{j,t}^2 \equiv \{Cov_t\}_{jj}$, and the corresponding daily logarithmic standard deviations as $lv_{j,t} \equiv \log(v_{j,t})$.

¹³ An alternative and much more complicated approach would be to utilize all of the observations by explicitly modeling the high-frequency frictions.

¹⁴ As detailed below, the average daily variance of the “typical” DJIA stock equals 3.109. Thus, in the case of i.i.d. normally distributed returns, it follows that a five-minute sampling frequency translates into a variance for the daily variance estimates of 0.245.

¹⁵ We also experimented with the use of unfiltered and linearly interpolated five-minute returns, which produced very similar results.

Similarly, we denote the realized daily correlations by $Corr_{ij,t} \equiv \{Cov_t\}_{i,j} / (v_{i,t} \cdot v_{j,t})$. In addition to the daily measures, we also briefly consider the statistical properties of various multiday volatility measures, whose construction follows in straightforward fashion from Eq. (4) by extending the summation to cover h/Δ intervals, where $h > 1$ denotes the multiday horizon.

Because volatility is now effectively observable, we can rely on conventional statistical procedures for characterizing its distributional properties. In the next section we proceed to do so. Of course, it is possible that the 30 DJIA stocks analyzed here do not provide a representative picture of the return volatility for other actively traded stocks. As a robustness check we replicated the empirical analysis for a set of 30 randomly selected liquid stocks, picked from the 214 stocks with at least 158 trades per day at the beginning, middle and end of the sample period. Importantly, all of the results for this randomly selected sample closely match those reported below for the DJIA stocks, thus underscoring the general nature of our findings. However, for reasons of space conservation, we shall not discuss the parallel empirical findings here; instead, we refer the interested reader to ABDE (2001) for a detailed discussion and a full set of tables.

3. Univariate unconditional return and volatility distributions

3.1. Returns

A voluminous literature, seeking to characterize the unconditional distribution of speculative returns, has evolved over the past three decades.¹⁶ Consistent with this literature, the summary statistics in Table 1 show that the daily DJIA returns, $r_{j,t}$, have fatter tails than the normal distribution and, for the majority of the stocks, are also skewed to the right.¹⁷

Quite remarkably, however, the next set of numbers in Table 1 indicates that all of the 30 standardized return series, $r_{j,t}/v_{j,t}$, are approximately unconditionally normally distributed. In particular, the median value of the sample kurtosis is reduced from 5.416 for the raw returns to only 3.129 for the standardized returns. This is also evident from Fig. 1, which plots the kernel density estimate for the mean-zero and unit-variance standardized returns for Alcoa Inc. (AA), the first of the 30 DJIA stocks, (alphabetically by ticker symbol). The close approximation afforded by the normal

¹⁶In early contributions, Mandelbrot (1963) and Fama (1965) argue that the Stable Paretian distributions provide a good approximation. Subsequently, however, Praetz (1972) and Blattberg and Gonedes (1974), among many others, find that finite variance-mixtures of normals, such as the student- t distribution, generally afford better characterizations.

¹⁷Under the null hypothesis of i.i.d. normally distributed returns, the sample skewness and kurtosis are asymptotically normal with means of 0 and 3, and variances of $6/T$ and $24/T$, where T denotes sample size. Thus for $T=1,366$ the two standard errors are 0.066 and 0.133.

Table 1
Unconditional daily return distributions^a

| Stock | $r_{i,t}$ | | | | $r_{i,t}/v_{i,t}$ | | | |
|---------|-----------|----------|--------|-------|-------------------|----------|--------|-------|
| | Mean | St. dev. | Skew. | Kurt. | Mean | St. dev. | Skew. | Kurt. |
| Min. | -0.059 | 1.149 | -0.221 | 3.810 | -0.033 | 0.623 | -0.054 | 2.734 |
| 0.10 | -0.024 | 1.222 | -0.022 | 3.964 | -0.013 | 0.697 | 0.037 | 2.821 |
| 0.25 | 0.007 | 1.275 | 0.035 | 4.236 | 0.002 | 0.772 | 0.081 | 3.005 |
| 0.50 | 0.041 | 1.419 | 0.159 | 5.416 | 0.024 | 0.806 | 0.113 | 3.129 |
| 0.75 | 0.071 | 1.538 | 0.231 | 6.587 | 0.048 | 0.852 | 0.164 | 3.302 |
| 0.90 | 0.084 | 1.704 | 0.487 | 8.462 | 0.060 | 0.928 | 0.228 | 3.414 |
| Max. | 0.140 | 1.833 | 0.564 | 11.98 | 0.099 | 0.960 | 0.322 | 3.848 |
| Mean | 0.036 | 1.438 | 0.172 | 5.908 | 0.025 | 0.808 | 0.125 | 3.156 |
| St.dev. | 0.046 | 0.181 | 0.192 | 2.016 | 0.030 | 0.080 | 0.081 | 0.251 |

^aThe table summarizes the daily return distributions for the 30 DJIA stocks, $r_{i,t}$. The sample covers the period from January 2, 1993 through May 29, 1998, for a total of 1,366 observations. The realized daily volatilities, $v_{i,t}$, are calculated from five-minute intraday returns, as detailed in the main text.




Fig. 1. Unconditional distribution of daily standardized returns. The figure shows the unconditional distributions of the standardized daily returns on Alcoa, $r_{AA,t}/v_{AA,t}$. The sample period extends from January 2, 1993 through May 29, 1998, for a total of 1,366 daily observations. The realized volatilities are calculated from five-minute intraday returns. The dotted line refers to the standard normal density.

reference density is striking.¹⁸ This result stands in sharp contrast to the leptokurtic distributions for the standardized daily returns that typically obtain when relying on an estimate of the one-day-ahead conditional variance from a parametric ARCH or stochastic volatility model; see Bollerslev et al. (1994) for

¹⁸The kernel density estimates are based on a Gaussian kernel and Silverman's (1986) bandwidth. Similar plots for all of the other stocks are available in ABDE (2001).

Table 2
Unconditional daily volatility distributions^a

| Stock | $v_{j,t}^2$ | | | | $lv_{i,t}$ | | | |
|---------|-------------|---------|-------|-------|------------|---------|--------|-------|
| | Mean | St.dev. | Skew. | Kurt. | Mean | St.dev. | Skew. | Kurt. |
| Min. | 1.899 | 1.159 | 1.451 | 5.789 | 0.239 | 0.225 | -0.537 | 2.282 |
| 0.10 | 2.009 | 1.348 | 2.306 | 11.98 | 0.280 | 0.228 | -0.308 | 3.245 |
| 0.25 | 2.539 | 1.665 | 3.516 | 27.84 | 0.403 | 0.238 | -0.015 | 3.475 |
| 0.50 | 3.108 | 1.988 | 5.609 | 66.16 | 0.476 | 0.264 | 0.192 | 3.885 |
| 0.75 | 3.390 | 2.458 | 8.322 | 142.6 | 0.544 | 0.280 | 0.465 | 4.758 |
| 0.90 | 4.315 | 4.346 | 18.89 | 518.0 | 0.664 | 0.294 | 0.777 | 5.136 |
| Max. | 6.854 | 6.319 | 20.70 | 567.8 | 0.894 | 0.353 | 1.023 | 6.620 |
| Mean | 3.178 | 2.355 | 7.433 | 143.6 | 0.478 | 0.263 | 0.222 | 4.101 |
| St.dev. | 1.146 | 1.203 | 5.664 | 176.8 | 0.150 | 0.029 | 0.388 | 0.870 |

^aThe table summarizes the distributions of the daily volatilities for the 30 DJIA stocks. The realized daily variances, $v_{j,t}^2$, and logarithmic standard deviations, $lv_{i,t} \equiv \log(v_{i,t})$, are calculated from five-minute intraday returns as detailed in the main text.

a general discussion, and Kim and Kon (1994) for explicit results related to the distributions of the DJIA stocks over an earlier time period. The results in Table 1 also imply that the unconditional distribution for the returns should be well-approximated by a continuous variance mixture of normals, as determined by the unconditional distribution for the mixing variable, $v_{j,t}^2$. The following section details this distribution.

3.2. Variances and logarithmic standard deviations

The first four columns in Table 2 provide the same set of summary statistics for the unconditional distribution of the realized daily variances. The median value for the sample means is 3.109, implying an annualized standard deviation for the typical stock of around 28%. However, there is considerable variation in the average volatility across the 30 stocks, ranging from a high of 42% for Walmart Stores Inc. (WMT) to a low of 22% for UTX.¹⁹ The standard deviations given in the second column also indicate that the realized daily volatilities fluctuate significantly through time. Finally, it is evident from the third and fourth columns that the distributions of the realized variances are extremely right-skewed and leptokurtic. This may seem surprising, as the realized daily variances are based on the sum of 79 five-minute return observations. However, as emphasized by ABDL (2000), intraday speculative returns are strongly dependent so that, even with much larger samples, standard Central Limit Theorem arguments often provide very poor approximations in the high-frequency data context.

¹⁹Details regarding the individual stocks are again available in ABDE (2001).




Fig. 2. Unconditional distribution of standardized daily logarithmic standard deviations. The figure shows the unconditional distribution of the standardized daily realized logarithmic standard deviations for Alcoa, $lv_{AA,t} \equiv \log(v_{AA,t})$. The realized volatilities are calculated from five-minute intraday returns. The dotted line refers to the standard normal density.

The next part of Table 2 refers to the realized logarithmic standard deviations, $lv_{j,t}$. Interestingly, the median value of the sample skewness across all of the 30 stocks is reduced to only 0.192, compared to 5.609 for the realized variances, and although the sample kurtosis for all but one of the stocks exceeds the normal value of three, the assumption of normality is obviously much better in this case. This is also illustrated by Fig. 2, in which we show estimates of the standardized unconditional density for $lv_{AA,t}$, along with the standard normal density. The normal approximation is very good.

This evidence is consistent with Taylor (1986) and French et al. (1987), who find that the distribution of logarithmic monthly standard deviations constructed from the daily returns within the month is close to Gaussian. It is also directly in line with the recent evidence in ABDL (2001a) and Zumbach et al. (1999), which indicates that realized daily foreign exchange rate volatilities constructed from high-frequency data are approximately log-normally distributed. Taken together, the results in Tables 1 and 2 imply that the unconditional distribution for the daily returns should be well-described by a continuous lognormal-normal mixture, as advocated by Clark (1973) in his seminal treatment of the Mixture-of-Distributions-Hypothesis (MDH). The foreign exchange rate forecasting results in ABDL (2000) corroborate this idea.

Our discussion thus far has centered on univariate return and volatility distributions. However, asset pricing, portfolio selection, and risk management decisions are invariably multivariate, involving many assets, with correlated returns. The next section summarizes the unconditional distributions of the pertinent realized covariances and correlations.

Table 3
Unconditional daily covariance and correlation distributions^a

| Stock | $Cov_{i,j,t}$ | | | | $Corr_{i,j,t}$ | | | |
|---------|---------------|---------|--------|-------|----------------|---------|-------|-------|
| | Mean | St.dev. | Skew. | Kurt. | Mean | St.dev. | Skew. | Kurt. |
| Min. | 0.217 | 0.362 | -4.714 | 8.411 | 0.062 | 0.128 | 0.008 | 2.552 |
| 0.10 | 0.284 | 0.508 | 2.843 | 22.49 | 0.085 | 0.138 | 0.123 | 2.832 |
| 0.25 | 0.318 | 0.580 | 3.738 | 33.97 | 0.098 | 0.142 | 0.180 | 2.939 |
| 0.50 | 0.372 | 0.695 | 5.223 | 61.86 | 0.117 | 0.149 | 0.253 | 3.044 |
| 0.75 | 0.426 | 0.819 | 7.704 | 120.6 | 0.136 | 0.157 | 0.312 | 3.177 |
| 0.90 | 0.492 | 0.968 | 12.21 | 258.3 | 0.159 | 0.167 | 0.381 | 3.321 |
| Max. | 0.697 | 1.899 | 24.91 | 773.4 | 0.221 | 0.196 | 0.568 | 3.668 |
| Mean | 0.379 | 0.727 | 6.462 | 108.1 | 0.120 | 0.151 | 0.251 | 3.068 |
| St.dev. | 0.081 | 0.206 | 4.195 | 125.4 | 0.029 | 0.012 | 0.099 | 0.199 |

^aThe table summarizes the distributions of the 435 ($= 30 \cdot 29/2$) unique realized covariances and correlations for the 30 DJIA stocks. The realized daily covariances and correlations are calculated from five-minute intraday returns, as detailed in the main text.

3.3. Covariances and correlations

The realized covariance matrix for the 30 DJIA stocks contains a total of 435 unique elements. In Table 3 we report the median value of the sample mean, standard deviation, skewness, and kurtosis for the covariances and correlations for each of the 30 stocks with respect to all of the 29 other stocks, i.e., the median value of the particular sample statistic across the 29 time series for stock i as defined by $Cov_{i,j,t}$ and $Corr_{i,j,t}$, where $j = 1, 2, \dots, 30$, and $j \neq i$.

The median of the mean covariance across all of the stocks equals 0.373, while the typical correlation among the DJIA stocks is around 0.113. However, the realized covariances and correlations exhibit considerable variation across the different stocks and across time. For instance, the median of the average correlations for Union Carbide Corp. (UK) equals 0.080, whereas the median for General Electric (GE) is as high as 0.150.

As with the realized variances, the distributions of the realized covariances are extremely right skewed and leptokurtic. Interestingly, however, the realized correlations appear approximately normally distributed. In particular, the median kurtosis for all of the 435 realized covariances equals 61.86, whereas the median kurtosis for the realized correlations equals 3.037. To illustrate this result, Fig. 3 graphs the unconditional distribution of the standardized realized correlations for Alcoa with respect to Exxon Corp. (XON), the alphabetically last ticker symbol of the 30 DJIA stocks.²⁰ It is obvious that the standard normal reference density affords a close approximation.

²⁰Similar graphs for all of the other correlations with respect to XON are available in ABDE (2001).




Fig. 3. Unconditional distribution of standardized daily correlations. The figure shows the unconditional distribution of the standardized daily realized correlations between Alcoa and Exxon, $Corr_{AA, XON,t}$. The realized correlations are calculated from five-minute intraday returns. The dotted line refers to the standard normal density.

The unconditional distributions detailed above capture important aspects of the return generating process, and they indicate that all of the realized volatilities vary importantly through time. In the next section, we explore the associated dynamic dependence. Again, the use of realized volatilities allows us to do so in a model-free environment, without reliance on complicated and intractable parametric latent volatility models.

4. Temporal dependence, long memory and scaling

The conditional distribution of stock market volatility has been the subject of extensive research effort during the past decade. Here we solidify and extend the findings in that literature; in particular, we reinforce the existence of pronounced long-run dependence in volatility and show that this effect is also present in correlations. Motivated by the results of the previous section, we focus on the logarithmic volatilities and correlations.

4.1. Logarithmic standard deviations

It is instructive first to consider the representative time-series plot for $lv_{AA,t}$ in Fig. 4. It is evident that the series is positively serially correlated, with distinct periods of high and low volatility readily identifiable. This is, of course, a manifestation of the well-documented volatility clustering effect, and directly in line with the results reported in the extant ARCH and stochastic volatility




Fig. 4. Time series of daily logarithmic standard deviations. The figure shows the time series of the daily realized logarithmic standard deviations for Alcoa, $lv_{AA,t} \equiv \log(v_{AA,t})$. The realized volatilities are calculated from five-minute intraday returns.

Table 4
Dynamic volatility dependence^a

| Stock | $lv_{i,t}$ | | | | $Corr_{i,j,t}$ | | | |
|---------|------------|--------|-----------|-------|----------------|--------|-----------|-------|
| | Q_{22} | ADF | d_{GPH} | d_S | Q_{22} | ADF | d_{GPH} | d_S |
| Min. | 982 | -4.850 | 0.263 | 0.286 | 155 | -5.351 | 0.117 | 0.177 |
| 0.10 | 2080 | -4.466 | 0.284 | 0.334 | 395 | -4.566 | 0.278 | 0.240 |
| 0.25 | 2966 | -3.918 | 0.317 | 0.359 | 660 | -4.065 | 0.326 | 0.271 |
| 0.50 | 4715 | -3.327 | 0.349 | 0.386 | 1169 | -3.542 | 0.380 | 0.308 |
| 0.75 | 6075 | -2.992 | 0.392 | 0.400 | 2167 | -2.983 | 0.439 | 0.347 |
| 0.90 | 6921 | -2.676 | 0.409 | 0.412 | 3431 | -2.571 | 0.486 | 0.376 |
| Max. | 14254 | -2.178 | 0.416 | 0.463 | 7209 | -1.917 | 0.600 | 0.422 |
| Mean | 4729 | -3.450 | 0.350 | 0.377 | 1267 | -3.548 | 0.381 | 0.308 |
| St.dev. | 2556 | 0.665 | 0.046 | 0.038 | 1267 | 0.746 | 0.081 | 0.051 |

^aThe table summarizes the time-series dependence in the 30 realized logarithmic standard deviations and 435 realized correlations for the DJIA stocks. The table reports the Ljung-Box portmanteau test for up to 22nd order autocorrelation, Q_{22} , the augmented Dickey-Fuller test for a unit root involving 22 augmentation lags, ADF , the Geweke-Porter-Hudak estimate for the degree of fractional integration, d_{GPH} , and the estimate for the degree of fractional integration based on the scaling law, d_S .

literatures; see Lamoureux and Lastrapes (1990) or Kim and Kon (1994) for estimation of GARCH models for individual daily stock returns.

To underscore the significance of this effect more generally, the first column in Table 4 summarizes the values of the standard Ljung-Box portmanteau test for the joint significance of the first 22 autocorrelations of $lv_{i,t}$ (about one month of trading days). The white noise hypothesis is overwhelmingly rejected

for all 30 stocks. The correlogram for Alcoa in Fig. 5 shows why. The autocorrelations are systematically above the conventional Bartlett 95% confidence band, the upper range of which is given by the flat dashed line, even at the longest displacement of 120 days (approximately half a year). Similarly slow decay rates have been documented in the literature with daily time series of absolute or squared returns spanning several decades (see, for example, Crato and de Lima, 1993; and Ding et al., 1993), but the results in Fig. 5 are noteworthy in that the sample “only” spans five-and-a-half years. In spite of this slow decay, the augmented Dickey-Fuller tests, reported in the second column in Table 4, reject the null hypothesis of a unit root for all but four of the stocks when judged by the conventional -2.86 5% critical value.²¹

In response to such findings, a number of recent studies have argued that the long-run dependence in financial market volatility can be conveniently modeled by fractional integrated ARCH or stochastic volatility models; see, Baillie et al. (1996), Breidt et al. (1998) and Robinson and Zaffaroni (1998). The log-periodogram regression estimates of the degree of fractional integration, or d , for the realized logarithmic volatilities, given in the third column in Table 4, are directly in line with these studies, and all 30 estimates are very close to the median value of 0.349, and highly statistically significantly different from both 0 and 1.²² It is also evident that the implied hyperbolic decay rate, $j^{2 \cdot d - 1}$, superimposed in Fig. 5, affords a close approximation to the correlogram for $lv_{AA,t}$, and equally good fits obtain for each of the 29 other stocks.

Another implication of the long memory associated with fractional integration concerns the behavior of the variance of partial sums. In particular, let $[x_t]_h \equiv \sum_{j=1, \dots, h} x_{h(t-1)+j}$ denote the h -fold partial sum process for x_t . If x_t is fractionally integrated, the partial sums will obey a *scaling law* of the form $Var([x_t]_h) = ch^{2d+1}$. Thus, given d and the unconditional variance at one aggregation level, it is possible to calculate the implied variance for any other

²¹ It is well-known, however, that the outcome of standard unit root tests should be carefully interpreted with slowly decaying processes; see Schwert (1987).

²² The theory behind the log-periodogram regression is developed in Geweke and Porter-Hudak (GPH, 1983) and Robinson (1995). The slow hyperbolic decay of the long-lag autocorrelations, or equivalently the log-linear explosion of the low-frequency spectrum, are both distinguishing features of a covariance stationary fractionally integrated, or I(d), process with $0 < d < 1/2$. Accordingly, let $I(\omega_j)$ denote the sample periodogram estimate for the spectrum at the j th Fourier frequency, $\omega_j = 2\pi j/T$, $j = 1, 2, \dots, [T/2]$. The GPH estimator for d is then based on the least squares regression,

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

where $j = 1, 2, \dots, m$, and $\hat{d} \equiv -1/2 \cdot \hat{\beta}_1$ is asymptotically normal with a standard error of $\pi(24 \cdot m)^{-1/2}$. For the estimates in Table 4 we take $m = [1,366]^{3/5} = 76$, thus implying an asymptotic standard error of 0.074. This particular choice was motivated by Deo and Hurvich (1999), who show that the GPH estimator is consistent and asymptotically normal provided that $m = O(T^{-\delta})$, where $\delta < 4d(1 + 4d)^{-1}$.




Fig. 5. Sample autocorrelations for daily logarithmic standard deviations. The figure shows the sample autocorrelations for the daily realized logarithmic standard deviations for Alcoa, $lv_{AA,t} \equiv \log(v_{AA,t})$, out to a displacement of 100 days. The realized volatilities are calculated from five-minute intraday returns. The dotted line gives the minimum-distance estimates of the hyperbolic decay rate, ch^{2d-1} . The dashed line give the upper range of the conventional Bartlett 95% confidence band.

aggregation level. To explore whether this implication of fractional integration is satisfied by our equity volatilities, we plot in Fig. 6 the logarithm of the variance of the partial sum of the daily realized logarithmic standard deviations, $\log(\text{Var}[lv_{AA,t}]_h)$, against the logarithm of the aggregation level, $\log(h)$, for $h = 1, 2, \dots, 30$. The accuracy of the fitted line, $c + (2d + 1) \log(h)$, is striking.²³ Moreover, the corresponding regression estimates for d for all of the stocks reported in the fourth column in Table 4, are generally very close to the GPH estimates.

4.2. Correlations

The estimation of parametric multivariate volatility models is notoriously difficult and, as a result, relatively little is known about the temporal behavior of individual stock return correlations.²⁴ The last four columns of Table 4 provide our standard menu of summary statistics for the 435 series of daily realized correlations. In accordance with our convention in Section 3.3 above, each entry gives the median value of that particular statistic across the 30 stocks.

²³ LeBaron (1999) has recently demonstrated that apparent scaling laws may arise for short-memory, but highly persistent, processes. In the present context, the hyperbolic decay in Fig. 5 further buttresses the long-memory argument.

²⁴ In a recent paper, Campbell et al. (2000) argue that although the number of stocks required to achieve a given level of diversification has increased noticeably over the past two decades, firm-specific volatility has also gone up, so that individual stock return correlations have actually decreased over the same time period.




Fig. 6. Volatility scaling plots for daily logarithmic standard deviations. The figure shows the logarithm of the variance of the partial sum of the daily realized logarithmic standard deviations for Alcoa, $\log(\text{Var}[lv_{AA,t}]_h)$, plotted against the logarithm of the aggregation level, $\log(h)$, for $h = 1, 2, \dots, 30$. The sample period extends from January 2, 1993 through May 29, 1998, for a total of 1,366 observations at the daily level. The daily realized volatilities are calculated from five-minute intraday returns. The dotted line refers to the least-squares estimates of the regression line $c + (2d + 1)\log(h)$.




Fig. 7. Time series of daily correlations. The figure shows the time series of daily realized correlations between Alcoa and Exxon, $\text{Corr}_{AA,XON,t}$. The sample period extends from January 2, 1993 through May 29, 1998, for a total of 1,366 daily observations. The realized correlations are calculated from five-minute intraday returns.

Turning to the results, the time-series plot of the correlation between Alcoa and Exxon, $\text{Corr}_{AA,XON,t}$, in Fig. 7 suggests important dependence and hence predictability in the correlations. This impression is confirmed by the correlogram in Fig. 8 and the Ljung-Box portmanteau statistics for up to 22nd-order serial correlation reported in Column 5 of Table 4.²⁵ Moreover, as

²⁵ As in Fig. 5, the flat dashed line denotes the upper range of the 95% Bartlett confidence band.




Fig. 8. Sample autocorrelations of daily correlations. The figure shows the sample autocorrelations for the daily realized correlations between Alcoa and Exxon, $Corr_{AA,XON,t}$. The realized correlations are calculated from five-minute intraday returns. The dotted line gives the minimum-distance estimates of the hyperbolic decay rate, ch^{2d-1} . The dashed line give the upper range of the conventional Bartlett 95% confidence band.




Fig. 9. Volatility scaling plots for daily correlations. The figure shows the logarithm of the variance of the partial sum of the daily realized correlations between Alcoa and Exxon, i.e., $\log(\text{Var}[Corr_{AA,XON,t}]_h)$, plotted against the logarithm of the aggregation level, $\log(h)$, for $h = 1, 2, \dots, 30$. The sample period extends from January 2, 1993 through May 29, 1998, for a total of 1,366 observations at the daily level. The daily realized correlations are calculated from five-minute intraday returns. The dotted line refers to the least-squares estimates of the regression line $c + (2d + 1)\log(h)$.

with the ADF tests for $lv_{j,t}$, the tests for $Corr_{i,j,t}$ reported in the sixth column systematically reject the unit root hypothesis. Accordingly, the GPH estimates for d are significantly different from zero (and unity), with typical values around 0.35. The corresponding hyperbolic decay rate for $Corr_{AA,XON,t}$ superimposed in Fig. 8 and the scaling law in Fig. 9, in which we plot $\log(\text{Var}[Corr_{AA,XON,t}]_h)$ against $\log(h)$, for $h = 1, 2, \dots, 30$, also reveal highly accurate fits.

Overall, our results thus far suggest that the *univariate* unconditional and conditional distributions for the realized correlations closely mimic the qualitative characteristics of the realized volatilities discussed earlier. We now turn to multivariate aspects of the distributions, focusing first on issues related to asymmetry in the distributions of the volatilities.

5. Asymmetric responses of volatilities and correlations

A number of previous studies document an asymmetry in the relation between equity volatility and returns, i.e., positive returns have a smaller impact on future volatility than do negative returns of the same absolute magnitude. Two competing explanations have been put forth to rationalize this phenomenon. According to the so-called leverage effect, a large negative return increases financial and operating leverage, in turn raising equity return volatility (Black, 1976; Christie, 1982). Alternatively, if the market risk premium is an increasing function of volatility, large negative returns increase the future volatility by more than positive returns due to a volatility feedback effect (Campbell and Hentschel, 1992). We now reevaluate the underlying empirical evidence on the basis of our realized volatility measures.

5.1. Logarithmic standard deviations

The use of realized volatilities allows for direct tests of asymmetries in the impact of past returns. However, in order to avoid confusing such effects with the strong serial correlation documented in the previous section, it is imperative that dynamic dependence be modeled properly. The first four columns in Table 5 report the regression estimates based on the fractionally differenced series,

$$(1 - L)^{d_i} l v_{i,t} = \omega_i + \gamma_i l v_{i,t-1} + \phi_i l v_{i,t-1} I(r_{i,t-1} < 0) + u_{i,t}, \quad (5)$$

where $I(\cdot)$ refers to the indicator function, and the values for d_i are fixed at the d_{GPH} estimates reported in Table 4. Also, to accommodate any additional short-run dynamics, the t -statistics are based on a Newey-West robust covariance matrix estimator using 22 lags.

The median estimated value of γ_i equals -0.023 , and only one of the 30 t -statistics for γ_i is statistically significantly greater than zero when judged by the standard 95% critical value of 1.645.²⁶ Simultaneously, the median estimate for

²⁶Note that as long as $-\gamma_i < d_i$, a negative value for γ_i is fully consistent with the strong volatility clustering effect documented above, as a series expansion of the fractional differencing operator in Eq. (5) would imply that terms of the form $d_i l v_{i,t-1}, \frac{1}{2} \cdot d_i(1 - d_i) l v_{i,t-2}, \dots$ also enter the right-hand side in the corresponding equation for $l v_{i,t}$.

Table 5
News impact functions^a

| Stock | $lv_{i,t}$ | | | | $Corr_{i,j,t}$ | | |
|---------|------------|------------|-----------|-------------|----------------|------------|-------------|
| | γ | T_γ | φ | t_φ | t_γ | t_θ | t_φ |
| Min. | -0.092 | -4.098 | 0.007 | 0.236 | -5.198 | -2.060 | -1.299 |
| 0.10 | -0.067 | -2.930 | 0.022 | 0.908 | -3.015 | -0.911 | 0.162 |
| 0.25 | -0.049 | -1.737 | 0.028 | 1.224 | -2.451 | -0.288 | 0.840 |
| 0.50 | -0.023 | -0.754 | 0.053 | 2.277 | -1.660 | 0.333 | 1.406 |
| 0.75 | 0.000 | 0.008 | 0.067 | 2.874 | -0.929 | 1.051 | 2.061 |
| 0.90 | 0.026 | 1.039 | 0.081 | 3.753 | -0.118 | 1.646 | 2.815 |
| Max. | 0.051 | 1.833 | 0.130 | 4.314 | 2.444 | 3.727 | 4.395 |
| Mean | -0.021 | -0.825 | 0.051 | 2.246 | -1.636 | 0.387 | 1.450 |
| St.dev. | 0.035 | 1.455 | 0.027 | 1.067 | 1.180 | 0.988 | 1.026 |

^aThe table reports the OLS regression estimates for the news impact functions for the fractionally differenced logarithmic standard deviations, $(1-L)^{d_i}lv_{i,t} = \omega_i + \gamma_i lv_{i,t-1} + \varphi_i lv_{i,t-1} I \times (r_{i,t-1} < 0) + u_{i,t}$, and correlations, $(1-L)^{d_{i,j}}Corr_{i,j,t} = \omega_{i,j} + \gamma_{i,j}(lv_{i,t-1} + lv_{j,t-1}) + \theta_{i,j} \times (lv_{i,t-1} + lv_{j,t-1})I(r_{i,t-1}r_{j,t-1} > 0) + \varphi_{i,j}(lv_{i,t-1} + lv_{j,t-1})I(r_{i,t-1} < 0, r_{j,t-1} < 0) + u_{i,j,t}$, where the values for d_i and $d_{i,j}$ are fixed at the d_{GPH} estimates reported in Table 4. The t -statistics are based on a Newey-West HAC covariance matrix estimator with 22 lags.

φ_i equals 0.053, and 22 of the 30 t -statistics exceed the 5% critical value. These results are broadly consistent with the EGARCH model estimates for daily individual stock returns reported by Cheung and Ng (1992) and Kim and Kon (1994), indicating a differential impact, or asymmetry, in the influence of past negative and positive returns.

However, although statistically significant for most of the stocks, the economic importance of this effect is marginal. Consider Fig. 10, which displays the scatterplots for the logarithmic standard deviation for Alcoa, $lv_{AA,t}$, against the lagged standardized returns, $r_{AA,t-1}/v_{AA,t-1}$. For visual reference, we have superimposed the two regression lines corresponding to negative and positive returns. This figure provides a direct analogy to the news impact curves for parametric ARCH models previously studied by Pagan and Schwert (1990) and Engle and Ng (1993). Although the news impact curve is more steeply sloped to the left of the origin, the systematic effect is obviously not very strong; similar plots for each of the 29 other stocks are available in ABDE (2001). This parallels the findings for the four individual stocks in Tauchen et al. (1996), who note that while asymmetry is a characteristic of the point estimates, the magnitude is quite small. In contrast, the parametric volatility model estimates reported in Nelson (1991), Glosten et al. (1993) and Hentschel (1995), among others, all point toward important asymmetries in market-wide equity index returns, which calls into question the leverage explanation and instead suggests that the significant asymmetries for the




Fig. 10. News impact functions for daily logarithmic standard deviations. The figure shows the scatterplot of the daily realized logarithmic volatilities for Alcoa, $lv_{AA,t}$, against the lagged standardized returns, $r_{AA,t-1}/lv_{AA,t-1}$. The solid lines refer to the estimated regression lines for the lagged returns being negative or positive. The realized volatilities are calculated from five-minute intraday returns.

aggregate market returns reported in these studies are most likely due to a volatility feedback effect (see also the recent discussion of Bekaert and Wu, 2000).

5.2. Correlations

As noted above, little is known about the distributions of individual stock return correlations. If the volatility asymmetry at the individual stock level is caused by a leverage effect, then a change in financial leverage is likely to also affect the covariances between different stocks, which in turn can affect the correlations. In this regard it is interesting to note that the different multivariate ARCH models estimated in Kroner and Ng (1998) generally result in statistically significant asymmetries in the conditional covariance matrices for the weekly returns on a pair of well-diversified small- and large-stock portfolios.²⁷ Similarly, Ang and Chen (2000) have recently demonstrated significant asymmetries in the correlations between the market and various industry, size, and book-to-market sorted portfolios. At the same time, the bivariate EGARCH models in Braun et al. (1995) indicate that while the overall market volatility responds asymmetrically to positive and negative shocks, monthly conditional (time-varying) betas for size- and industry-sorted portfolios are mostly symmetric. More recently, however, Cho and Engle (1999) report statistically significant asymmetries in daily EGARCH betas for

²⁷ In the context of international equity markets, Erb, et al. (1994) and Longin and Solnik (1998) have also argued that the cross-country correlations tend to be higher when the returns are negative.




Fig. 11. News impact functions for daily correlations. The figure shows the scatterplots for the daily realized correlations between Alcoa and Exxon, $Corr_{AA, XON,t}$, against the average lagged standardized returns, $\frac{1}{2}(r_{XON,t-1}/v_{XON,t-1} + r_{j,t-1}/v_{j,t-1})$. The solid lines refer to the estimated regression lines for the sum of the lagged returns being negative or positive. The realized correlations are calculated from five-minute intraday returns.

a small set of individual stocks, suggesting that the apparent symmetry in the monthly portfolio betas in Braun et al. (1995) are due to cross-sectional and/or temporal aggregation effects.

In light of these findings, we now extend the analysis above to test for asymmetries in the realized daily correlations. In particular, the last three columns in Table 5 report the results from the regressions

$$\begin{aligned}
 (1 - L)^{d_{i,j}} Corri_{i,j,t} = & \omega_{i,j} + \gamma_{i,j}(lv_{i,t-1} + lv_{j,t-1}) \\
 & + \theta_{i,j}(lv_{i,t-1} + lv_{j,t-1})I(r_{i,t-1} \cdot r_{j,t-1} > 0) \\
 & + \varphi_{i,j}(lv_{i,t-1} + lv_{j,t-1})I(r_{i,t-1} < 0, r_{j,t-1} < 0) + u_{i,j,t}, \quad (6)
 \end{aligned}$$

where, as before, the $d_{i,j}$ are fixed at the d_{GPH} estimates reported in Table 4, and the t -statistics are based on a Newey-West HAC covariance matrix estimator using 22 lags. Note that $\gamma_{i,j}$ captures the impact of the past realized volatilities on the correlations, $\theta_{i,j}$ gives the additional influence when the past returns are of the same sign, while the overall impact of the past volatility if both of the returns are negative is measured by $\gamma_{i,j} + \theta_{i,j} + \varphi_{i,j}$. This particular formulation therefore facilitates a direct test of asymmetry based on the t -statistic for φ_{ij} .²⁸

Turning to the results, most of the 435 estimates for φ_{ij} are indeed positive. However, fewer than half are significant at the usual 95% level when judged by the t_φ -statistics. This relatively weak asymmetry is underscored by Fig. 11, which plots the daily realized correlations for Alcoa and Exxon,

²⁸ Because of the fractional differencing operator on the left side of the equation, the actual coefficient values should be carefully interpreted.

$Corr_{AA,XON,t}$, against the average lagged standardized returns, $1/2(r_{AA,t-1}/v_{AA,t-1} + r_{XON,t-1}/v_{XON,t-1})$. As with the realized volatility news impact curve in Fig. 10, the line corresponding to the sum of the two lagged returns being negative is slightly more steeply sloped than the line corresponding to the sum of the lagged returns being positive. However, the systematic influence of this effect is clearly not very important. Similar graphs obtain for all of the other stocks (see ABDE 2001).

6. Multivariate unconditional volatility distributions

Here we investigate various aspects of the multivariate unconditional volatility distributions. Many key economic and financial, as well as regulatory, questions depend upon the perceived commonality in volatility movements across assets and markets. Most of the existing evidence concerning the extent of such comovements relies on very specific parametric volatility models. The realized volatility measures, in contrast, allow for a direct assessment of the relation between the individual standard deviations and correlations.

We begin in Fig. 12 with a scatterplot of the realized daily logarithmic standard deviation of Alcoa, $lv_{AA,t}$, against the logarithmic standard deviation of Exxon, $lv_{XON,t}$. It is evident that the two volatilities move together. This feature also holds for the other stocks. From the first column in Table 6, the median correlation between $lv_{i,t}$ and $lv_{j,t}$ across the 435 unique pairwise combinations equals 0.205.²⁹ As discussed further below, this tendency of return volatility to vary in tandem across individual stocks is consistent with factor structure, as in Diebold and Nerlove (1989), Tauchen and Tauchen (1999), and others.

Next, in Fig. 13 and the second column of Table 6, we document the presence of what might be termed a volatility-in-correlation effect. In particular, in Fig. 13 we plot the average realized daily correlations for Alcoa, $(1/29)\sum_i Corr_{AA,i,t}$ for $i \neq AA$, against the logarithmic standard deviation for Alcoa, $lv_{AA,t}$. As with the foreign exchange rates analyzed in ABDL (2001a), a strong positive association is evident. This is further underscored by the results in the second column in Table 6. The median correlation between all of the individual correlations, $Corr_{i,i,t}$, and the corresponding logarithmic standard deviations, $lv_{h,t}$ for $i = h$ or $j = h$, equals 0.150. While similar volatility-in-correlation effects have been documented for other broadly defined market

²⁹ Embrechts et al. (1999) have recently advocated the use of copulas and rank statistics when measuring dependence in non-normally distributed financial data. However, because the unconditional distributions that we explore in Table 6 are all approximately Gaussian, the linear correlation affords the most natural measure in the present context.




Fig. 12. Scatterplot of daily realized logarithmic standard deviations. The figure shows the scatterplot of realized daily logarithmic standard deviations for Alcoa, $v_{AA,t}$, against the logarithmic standard deviations for Exxon, $v_{XON,t}$. The realized volatilities are calculated from five-minute intraday returns.

indexes, our direct model-free measurements of realized correlations and volatilities are very different from the procedures previously entertained in the literature, and as such our findings provide additional empirical support for the phenomenon.³⁰ As shown below, a volatility effect in correlation is also to be expected within a factor structure, just as with the positive correlation across volatilities (see also Ronn et al., 1999). At the same time, the specific manifestation of the effect is model dependent, which renders direct predictions about magnitudes impossible within our nonparametric setting. Nonetheless, the strength of the effect is noteworthy and provides a benchmark measure that candidate models should be able to accommodate. At the least, it suggests that standard mean-variance efficiency calculations based on constant correlations are misguided.

Our final look at the multivariate volatility distributions in Fig. 14 shows the scatter plot of average realized daily correlations for Alcoa against the average realized correlations for Exxon, i.e., $(1/28) \sum_i Corr_{AA,i,t}$ versus $(1/28) \sum_i Cor r_{XON,i,t}$ for $i \neq AA$ and $i \neq XON$. The strong association between the realized daily correlations is truly striking. Clearly, there is a powerful commonality in the comovements across the individual stocks. The last column of Table 6 tells

³⁰ Similar observations have recently been made in the context of international equity index returns by Solnik et al. (1996). This also motivates the switching ARCH model estimated by Ramchand and Susmel (1998), who argue that the correlations between the U.S. and other world markets are on average 2 to 3.5 times higher when the U.S. market is in a high-variance state as compared to a low-variance state.

Table 6
Multivariate unconditional volatility distributions^a

| Stock | $Corr(lv_{i,t}, lv_{j,t})$ | $Corr(Corr_{i,j,t}, lv_{i,t})$ | $Corr(Corr_{i,j,t}, Corr_{h,k,t})$ |
|---------|----------------------------|--------------------------------|------------------------------------|
| Min | -0.327 | -0.209 | 0.093 |
| 0.10 | -0.016 | -0.086 | 0.230 |
| 0.25 | 0.081 | 0.032 | 0.265 |
| 0.50 | 0.205 | 0.150 | 0.308 |
| 0.75 | 0.321 | 0.236 | 0.358 |
| 0.90 | 0.439 | 0.296 | 0.407 |
| Max. | 0.641 | 0.536 | 0.601 |
| Mean | 0.206 | 0.130 | 0.314 |
| St.dev. | 0.172 | 0.148 | 0.069 |

^aThe column labeled $Corr(lv_{i,t}, lv_{j,t})$ gives the distribution of the 435 ($=30 \cdot 29/2$) unique correlations between the 30 daily realized logarithmic volatilities for the DJIA stocks. The second column, $Corr(Corr_{i,j,t}, lv_{i,t})$, refers to the distribution of the 870 ($=30 \cdot 29$) unique correlations between the daily realized correlations and the corresponding logarithmic standard deviations. The last column denoted $Corr(Corr_{i,j,t}, Corr_{h,k,t})$ gives the distribution of the 82,215 ($=30 \cdot 29 \cdot 28 \cdot 27/8$) unique correlations between the realized daily correlations.




Fig. 13. Scatterplot of average daily realized correlations versus logarithmic standard deviations. The figure shows the scatterplot of the average realized daily correlations for Alcoa, $(1/29) \sum_i Corr_{AA,i,t}$ and $i \neq AA$, against the logarithmic standard deviations for Alcoa, $lv_{AA,t}$. The realized volatilities and correlations are calculated from five-minute intraday returns.

the same story. The smallest correlation among the 82,215 ($=30 \cdot 29 \cdot 28 \cdot 27/8$) unique correlations is as high as 0.093, and the median correlation between the daily time series of realized correlations equals 0.308. Again, this seems to suggest that there is a lower-dimensional factor structure driving the second-moment characteristics of the joint distribution, to which we now turn.




Fig. 14. Scatterplot of average daily realized correlations. The figure shows the scatterplot of the average realized daily correlations for Alcoa against the average realized correlations for Exxon, i.e., $(1/28) \sum_i \text{Corr}_{AA,i,t}$ against $(1/28) \sum_i \text{Corr}_{XON,i,t}$ for $i \neq AA$ and $i \neq XON$. The realized correlations are calculated from five-minute intraday returns.

7. Latent factor structure in volatility

The notion of a low-dimensional factor structure is central to modern asset pricing theory (see, for example, Cochrane, 2000). We briefly explore the properties of realized volatility in the context of a simple multivariate model with an explicit factor structure. We focus on three of the empirical results noted above: the tendency for volatilities to move together, the tendency for correlations to be high when the corresponding volatilities are high, and the tendency for an arbitrary correlation to be high when other correlations are also high.

Consider an N -dimensional diffusion for log price p_t with the single-factor representation

$$dp_t = \lambda \sigma_t dW_t + \Omega dV_t, \quad (7)$$

where λ is an N -dimensional vector of loadings on the common volatility factor $\sigma_t dW_t$, V_t is an N -dimensional standard Brownian motion with mutually independent elements, and the diagonal matrix Ω contains N individual asset-specific volatilities. Note that each element of the N -vector of returns dp_t is driven in part by a single latent factor with stochastic volatility and in part by an orthogonal idiosyncratic noise.

Given the simple model (7), the N -dimensional vector of daily returns is

$$r_{t+1} \equiv p_{t+1} - p_t = \int \lambda \sigma_t dW_t + \int \Omega dV_t. \quad (8)$$

Letting Σ_t denote the corresponding $N \times N$ covariance matrix conditional on the sample path filtration generated by the latent volatility process, $\{\sigma_{t+\tau}\}_{\tau=0}^1$, the element of Σ_t corresponding to the covariance between the i th and j th elements of r_{t+1} , say $\{r_{t+1}\}_i$ and $\{r_{t+1}\}_j$, is

$$\{\Sigma_t\}_{i,j} \equiv \{\lambda\lambda'\}_{i,j} \int \sigma_{t+\tau}^2 d\tau + \Omega_{i,j}. \quad (9)$$

Hence, the conditional variances and covariances inherit their dynamics from σ_t , a fact with important implications for comovements among volatilities and correlations.

In order to relate this factor model directly to daily realized volatilities and correlations, it is convenient to restate the system in discrete time. The continuous time latent factor volatility model (7)–(9) maps directly into a discrete time model that has been studied by a number of authors, including Diebold and Nerlove (1989), Harvey et al. (1994), King et al. (1994), Fiorentini et al. (1998), and Jacquier and Marcus (2000):

$$\begin{aligned} r_{it} &= \lambda_i f_t + v_{it} \\ f_t | I_t &\sim (0, h_t) \\ v_{it}^{\text{iid}} &\sim (0, \omega_i^2) \\ \text{cov}(v_{it} v_{jt'}) &= 0, \quad \forall i \neq j, t \neq t', \end{aligned} \quad (10)$$

where $i, j = 1, \dots, N$, and $t = 1, \dots, T$.

It is readily established that volatilities tend to move together in such a factor model. Concretely, the i th and j th time- t conditional variances, for arbitrary i and j , are

$$\begin{aligned} h_{it} &= \lambda_i^2 h_t + \omega_i^2, \\ h_{jt} &= \lambda_j^2 h_t + \omega_j^2. \end{aligned} \quad (11)$$

Note in particular that the conditional variances, which are themselves covariance stationary stochastic processes, are linear functions of latent volatility h_t and are therefore driven entirely by movements in volatility. The unconditional covariance between h_{it} and h_{jt} is

$$\begin{aligned} \text{cov}(h_{it}, h_{jt}) &= \text{E} \left[\left((\lambda_i^2 h_t + \omega_i^2) - (\lambda_i^2 \text{E}(h_t) + \omega_i^2) \right) \right. \\ &\quad \left. \times \left((\lambda_j^2 h_t + \omega_j^2) - (\lambda_j^2 \text{E}(h_t) + \omega_j^2) \right) \right] = \lambda_i^2 \lambda_j^2 (h_t - \text{E}(h_t))^2, \end{aligned} \quad (12)$$

which is unambiguously positive. Hence the unconditional correlation between h_{it} and h_{jt} is also unambiguously positive.

It is also readily seen why a factor structure induces high correlations in situations of high volatility. The ij th time- t conditional covariance is

$$\text{cov}_{ijt} = \lambda_i \lambda_j h_t, \quad (13)$$

so the conditional correlation is

$$corr_{ijt} = \frac{\lambda_i \lambda_j h_t}{\sqrt{\lambda_i^2 h_t + \omega_i^2} \sqrt{\lambda_j^2 h_t + \omega_j^2}}. \quad (14)$$

Note the conditional variance effect in conditional correlation: if, $\lambda_i, \lambda_j, h_t > 0$, then $\partial corr_{ijt} / \partial h_t > 0$. Moreover, $\lim_{h_t \rightarrow \infty} corr_{ijt} = 1$ and $\lim_{h_t \rightarrow 0} corr_{ijt} = 0$.

Finally, it is straightforward to verify that a factor structure implies that if the correlation between an arbitrary pair of stocks is high, the correlations between other stocks tend to be high also. In fact, Eq. (14) makes clear that, so long as all stocks load positively off the common factor, *all* pairwise correlations are increasing in volatility. Hence, as volatility moves, the pairwise correlations all move as well, and in the same direction.

In closing, we note that it is easy to extend these results to richer factor structures, including models with dynamics in Ω and models with multiple factors, as in recent work by Lo and Wang (2000) on modeling volume, which is intimately related to the modeling of volatility.

8. Conclusions

We exploit direct model-free measures of realized daily volatility and correlation obtained from high-frequency intraday stock prices to confirm, solidify and extend existing characterizations. Our findings are remarkably consistent with existing work such as ABDL (2001a, b) and Ebens (1999a). This is true of the right-skewed distributions of the variances and covariances, the normal distributions of the logarithmic standard deviations and correlations, the normal distributions of daily returns standardized by realized standard deviations, and the strongly persistent dynamics of the realized volatilities and correlations, well-described by a stationary fractionally integrated process and conforming to scaling laws under temporal aggregation. The striking congruence of all findings across asset classes (equity vs. foreign exchange) and underlying method of price recording (transaction prices vs. averages of logarithmic bid and ask quotes) suggests that the results reflect fundamental attributes of speculative returns.

Our analysis is noteworthy not only for confirming and checking the robustness of existing results, but also for achieving significant extensions, facilitated throughout by the model-free measurement of realized volatility and correlation afforded by high-frequency data, and the simplicity of our methods, which enable straightforward high-dimensional correlation estimation. We shed new light on some distinct properties of equity return dynamics and illustrate them, for example, via the news impact curve. We confirm the existence of an asymmetric relation between returns and volatility, with

negative returns being associated with higher volatility innovations than positive returns of the same magnitude. However, the effect is much weaker at the individual stock level than at the aggregate market level, thus lending support to a volatility risk premium feedback explanation rather than a financial leverage effect. Finally, we find a pronounced volatility-in-correlation effect. The volatility-in-correlation effect, the strong positive relations between individual stock volatilities, and the corresponding strong positive relations between contemporaneous stock correlations should motivate additional work on the development of parsimonious factor models for the covariance structure of stock returns.

We envision several applications of the approach adopted in this paper. For example, the direct measurement of volatilities and correlations should alleviate the errors-in-variables problem that plagues much work on the implementation and testing of the CAPM, because realized betas can be constructed directly from the corresponding realized covariances and standard deviations. Multifactor models based on factor-replicating portfolios are similarly amenable to direct analysis. As a second example, the effective observability of volatilities and correlations facilitates direct time-series modeling of portfolio choice and risk management problems under realistic and testable distributional assumptions. Work along those lines is pursued in Andersen et al. (2000). Finally, our methods will also facilitate direct comparisons of volatility forecasts generated by alternative models and procedures. Such explorations are underway in Ebens (1999b) and Ebens and de Lima (1999).

References

- Andersen, T.G., 1996. Return volatility and trading volume: an information flow interpretation of stochastic volatility. *Journal of Finance* 51, 169–204.
- Andersen, T.G., Bollerslev, T., 1997. Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns. *Journal of Finance* 52, 975–1005.
- Andersen, T.G., Bollerslev, T., 1998. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.
- Andersen, T.G., Bollerslev, T., Das, A., 2000. Variance ratios and high-frequency data: testing for changes in intraday volatility patterns. *Journal of Finance*, forthcoming.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2001. The distribution of stock return volatility: appendix. Working Paper, Northwestern University, Duke University, University of Pennsylvania and Johns Hopkins University.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 1999. Microstructure bias and volatility signatures. Working Paper, Northwestern University, Duke University and University of Pennsylvania.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2000. Modeling and forecasting realized volatility. NBER Working Paper 8160.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2001a. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* 96, 42–55.

- Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2001b. Exchange rate returns standardized by realized volatility are (nearly) Gaussian. *Multinational Finance Journal*, forthcoming.
- Ang, A., Chen, J., 2000. Asymmetric correlations of equity portfolios. Working paper, Graduate School of Business Columbia University.
- Back, K., 1991. Asset prices for general processes. *Journal of Mathematical Economics* 20, 317–395.
- Baillie, R.T., Bollerslev, T., Mikkelsen, H.O., 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 74, 3–30.
- Bekaert, G., Wu, G., 2000. Asymmetric volatility and risk in equity markets. *Review of Financial Studies* 13, 1–42.
- Black F., 1976. Studies of stock market volatility changes. *Proceedings of the American Statistical Association, Business and Economic Statistics Section* 177–181.
- Blattberg, R.C., Gonedes, N.J., 1974. Comparison of the stable and student distributions as statistical models for stock prices. *Journal of Business* 47, 244–280.
- Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH modeling in finance: a selective review of the theory and empirical evidence. *Journal of Econometrics* 52, 5–59.
- Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH models. *Handbook of Econometrics* 4, 2959–3038.
- Bollerslev, T., Mikkelsen, H.O., 1999. Long-term equity anticipation securities and stock market volatility dynamics. *Journal of Econometrics* 92, 75–99.
- Braun, P.A., Nelson, D.B., Sunier, A.M., 1995. Good news, bad news, volatility and betas. *Journal of Finance* 50, 1575–1603.
- Breidt, F.J., Crato, N., de Lima, P., 1998. On the detection and estimation of long-memory in stochastic volatility. *Journal of Econometrics* 83, 325–348.
- Campbell, J.Y., Hentschel, L., 1992. No news is good news: an asymmetric model of changing volatility in stock returns. *Journal of Financial Economics* 31, 281–318.
- Campbell, J.Y., Lettau, M., Malkiel, B.G., Xu, Y., 2000. Have individual stocks become more volatile? an empirical exploration of idiosyncratic risk. *Journal of Finance*, forthcoming.
- Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. *The Economics of Financial Markets*. Princeton University Press, Princeton.
- Canina, L., Figlewski, S., 1993. The informational content of implied volatility. *Review of Financial Studies* 6, 659–681.
- Cheung, Y.-W., Ng, L., 1992. Stock price dynamics and firm size: an empirical investigation. *Journal of Finance* 47, 1985–1997.
- Cho, Y.H., Engle, R.F., 1999. Time-varying betas and asymmetric effects of news: empirical analysis of blue chip stocks. NBER working paper No.7330.
- Christensen, B.J., Prabhala, N.R., 1998. The relation between implied and realized volatility. *Journal of Financial Economics* 50, 125–150.
- Christie, A.A., 1982. The stochastic behavior of common stock variances: value, leverage and interest rate effects. *Journal of Financial Economics* 10, 407–432.
- Clark, P.K., 1973. A Subordinated stochastic process model with finite variance for speculative prices. *Econometrica* 41, 135–155.
- Cochrane, J.H., 2000. Asset pricing. Manuscript, University of Chicago.
- Crato, N., de Lima, P.J., 1993. Long-range dependence in the conditional variance of stock returns. *Economic Letters* 25, 281–285.
- Day, T.E., Lewis, C.M., 1992. Stock market volatility and the information content of stock index options. *Journal of Econometrics* 52, 267–287.
- Deo, R.S., Hurvich, C.M., 1999. On the log-periodogram regression estimator of the memory parameter in long-memory stochastic volatility models. Working paper, New York University.
- Diebold, F.X., Nerlove, M., 1989. The dynamics of exchange rate volatility: a multivariate latent factor ARCH model. *Journal of Applied Econometrics* 4, 1–22.

- Ding, Z., Granger, C.W.J., Engle, R.F., 1993. A long memory property of stock market returns and a new model. *Journal of Empirical Finance* 1, 83–106.
- Drost, F.C., Nijman, T.E., Werker, B.J.M., 1998. Estimation and testing in models containing both jumps and conditional heteroskedasticity. *Journal of Business and Economic Statistics*. 16, 237–243.
- Duffee, G.R., 1995. Stock returns and volatility: a firm level analysis. *Journal of Financial Analysis* 37, 399–420.
- Ebens, H., 1999a. Realized stock index volatility. Working paper No. 420, Department of Economics Johns Hopkins University.
- Ebens, H., 1999b. Forecasting stock volatility. Work in progress.
- Ebens, H., de Lima, P., 1999. The volatility implied by options and realized volatility. Work in progress.
- Embrechts, P., McNeil, A., Straumann, D., 1999. Correlation and dependency in risk management: properties and pitfalls. Working paper, ETH Zürich, Switzerland.
- Engle, R.F., Lee, G.G.J., 1993. Long run volatility forecasting for individual stocks in a one factor model. Working paper, University of California, San Diego.
- Engle, R.F., Ng, V.K., 1993. Measuring and testing the impact of news on volatility. *Journal of Finance* 48, 1749–1778.
- Erb, C.B., Harvey, C.R., Viskanta, T.E., 1994. Forecasting international equity correlations. *Financial Analysts Journal*, November/December, 32–45.
- Fama, E.F., 1965. The behavior of stock market prices. *Journal of Business* 38, 34–105.
- Fiorotini, G., Sentana, E., Shephard, N., 1998. Exact likelihood-based estimation of conditionally heteroskedastic factor models. Working paper, University of Alcantara, CEMFI and Oxford University.
- French, K.R., Schwert, G.W., Stambaugh, R.F., 1987. Expected stock returns and volatility. *Journal of Financial Economics* 19, 3–29.
- Fung, W.K.H., Hsieh, D.A., 1991. Empirical analysis of implied volatility: stocks, bonds and currencies. Working paper, Duke University.
- Gallant, A.R., Rossi, P.E., Tauchen, G.E., 1992. Stock prices and volume. *Review of Financial Studies* 5, 199–242.
- Geweke, J., Porter-Hudak, S., 1983. The estimation and application of long memory time-series models. *Journal of Time Series Analysis* 4, 221–238.
- Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779–1801.
- Harvey, A., Ruiz, E., Shephard, N., 1994. Multivariate stochastic variance models. *Review of Economic Studies* 61, 247–264.
- Hentschel, L., 1995. All in the family: nesting symmetric and asymmetric GARCH models. *Journal of Financial Economics* 39, 71–104.
- Hsieh, D.A., 1991. Chaos and nonlinear dynamics: application to financial markets. *Journal of Finance* 46, 1839–1877.
- Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42, 381–400.
- Jacquier, E., Marcus, A.J., 2000. Market volatility and asset correlation structure. Working paper, Boston College.
- Jacquier, E., Polson, N.G., Rossi, P.E., 1994. Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* 12, 371–417.
- Kim, D., Kon, S.J., 1994. Alternative models for the conditional heteroskedasticity of stock returns. *Journal of Business* 67, 563–598.
- King, M., Sentana, E., Wadhvani, S., 1994. Volatility and links between national stock markets. *Econometrica* 62, 901–933.

- Kroner, K.F., Ng, V.K., 1998. Modeling asymmetric co-movements of asset returns. *Review of Financial Studies* 11, 817–844.
- Kuwahara, H., Marsh, T.A., 1992. The pricing of Japanese equity warrants. *Management Science* 38, 1610–1641.
- Lamoureux, C.G., Lastrapes, W.D., 1990. Heteroskedasticity in stock return data: volume versus GARCH effects. *Journal of Finance* 45, 221–229.
- Lamoureux, C.G., Lastrapes, W.D., 1993. Forecasting stock returns variances: towards understanding stochastic implied volatility. *Review of Financial Studies* 6, 293–326.
- LeBaron, B., 1999. Volatility persistence and apparent scaling laws in finance. Working paper, Brandeis University.
- Ledoit, O., Santa-Clara, P., 1998. Relative pricing of options with stochastic volatility. Working paper, UCLA.
- Lo, A.W., Wang, J., 2000. Trading volume: definitions, data analysis, and implications of portfolio theory. *Review of Financial Studies* 13, 257–300.
- Longin, F., Solnik, B., 1998. Correlation structure of international equity markets during extremely volatile periods. Working paper, ESSEC Graduate Business School and HEC School of Management.
- Mandelbrot, B., 1963. The variation of certain speculative prices. *Journal of Business* 36, 394–419.
- Merton, R.C., 1980. On estimating the expected return on the market. *Journal of Financial Economics* 8, 323–361.
- Nelson, D.B., 1990. ARCH models as diffusion approximations. *Journal of Econometrics* 45, 7–38.
- Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* 59, 347–370.
- 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.
- Nelson, D.B., Foster, D.P., 1994. Asymptotic filtering theory for univariate ARCH models. *Econometrica* 62, 1–41.
- Pagan, A.R., Schwert, G.W., 1990. Alternative models for conditional stock volatility. *Journal of Econometrics* 45, 267–290.
- Praetz, P.D., 1972. The distribution of share price changes. *Journal of Business* 45, 49–55.
- Protter, P., 1992. *Stochastic Integration and Differential Equations*. Springer, New York.
- Ramchand, L., Susmel, R., 1998. Volatility and cross correlation across major stock markets. *Journal of Empirical Finance* 5, 397–416.
- Robinson, P.M., 1995. Log-periodogram regression of time series with long-range dependence. *Annals of Statistics* 23, 1048–1072.
- Robinson, P.M., Zaffaroni, P., 1998. Nonlinear time series with long memory: a model for stochastic volatility. *Journal of Statistical Planning and Inference* 68, 359–371.
- Ronn, E.I., Sayrak, A., Tompaidis, S., 1999. The impact of large changes in asset prices on intra-market correlations in the stock, bond, and international markets. Working paper, University of Texas at Austin.
- Schwert, G.W., 1987. Effects of model specification on tests for unit roots in macroeconomic data. *Journal of Monetary Economics* 20, 73–103.
- Schwert, G.W., 1989. Why does stock market volatility change over time? *Journal of Finance* 44, 1115–1153.
- Schwert, G.W., 1990a. Stock market volatility. *Financial Analysts Journal* 46, 23–34.
- Schwert, G.W., 1990b. Stock volatility and the crash of '87. *Review of Financial Studies* 3, 77–102.
- Schwert, G.W., Seguin, P.J., 1991. Heteroskedasticity in stock returns. *Journal of Finance* 45, 1129–1155.
- Silverman, B.W., 1986. *Density Estimation*. Chapman & Hall, London, UK.
- Solnik, B., Boucrelle, C., Le Fur, Y., 1996. International market correlation and volatility. *Financial Analysts Journal* 52, 17–34.

- Tauchen, G.E., Tauchen, H., 1999. The volatility structure of the dow, work in progress.
- Tauchen, G.E., Zhang, H., Liu, M., 1996. Volume, volatility and leverage: a dynamic analysis. *Journal of Econometrics* 74, 177–208.
- Taylor, S.J., 1986. *Modeling Financial Time Series*. Wiley, Chichester.
- Zumbach, G.O., Dacorogna, M.M., Olsen, J.L., Olsen, R.B., 1999. Introducing a scale of market shocks. Working paper, Olsen and Associates, Zurich.