THE JOURNAL OF FINANCE • VOL. LV, NO. 5 • OCT. 2000

Order Flow, Transaction Clock, and Normality of Asset Returns

THIERRY ANÉ and HÉLYETTE GEMAN*

ABSTRACT

The goal of this paper is to show that normality of asset returns can be recovered through a stochastic time change. Clark (1973) addressed this issue by representing the price process as a subordinated process with volume as the lognormally distributed subordinator. We extend Clark's results and find the following: (i) stochastic time changes are mathematically much less constraining than subordinators; (ii) the cumulative number of trades is a better stochastic clock than the volume for generating virtually perfect normality in returns; (iii) this clock can be modeled nonparametrically, allowing both the time-change and price processes to take the form of jump diffusions.

The relations among trading volume, stock prices, and price volatility, the subject of empirical and theoretical studies over many years, have lately received renewed attention with the increased availability of high frequency data. A vast amount of research has focused on issues such as news arrivals, volume, and price changes or volatility moves, usually outside any framework of general or even partial equilibrium. Is the normality of returns—a key issue, for example, in the mean-variance paradigm for portfolio choice, or the recent study of the problems of risk management (e.g., in Value at Risk)—verified at any time horizon? The evidence accumulated from a number of studies that document the presence of leptokurtosis and skewness in the distribution of returns of a wide variety of financial assets suggests that the answer is no. Studies as early as, for example, Fama (1965), showed that daily returns are more long tailed than the normal density, with the distribution of returns approaching normality as the holding period is extended to one month. In the same manner, volatility smiles and other observed

* Ané is from the University Paris Dauphine. Geman is from the University Paris Dauphine and ESSEC Graduate Business School. We thank the Center for International Securities and Derivative Markets of the University of Massachusetts for providing the data. We appreciate the comments made by participants of the Finance Seminar of the University of Chicago, the University of Washington, the University of Maryland, Carnegie Mellon University, the Workshop on "Stochastics, Information and Markets" at Humboldt University, and the European Finance Association meeting in Fontainebleau. The suggestions by René Stulz (the editor) and an anonymous referee were particularly helpful.

deviations from the Black and Scholes model seem to most directly contradict the assumption of normality in asset returns, which has obvious importance for the pricing and hedging of derivative instruments.

Many analytical approaches have been proposed to address and analyze the departure of returns from normality. Mandelbrot (1963) introduced a class of stable processes to account for the deviations of returns from Brownian motion. As an alternative explanation, Clark (1973) proposed linking the deviations from normality to the existence of variations in volume during different trading periods and introduced the use of subordinated processes in finance, that is, the idea that calendar time may not be the appropriate measure of time in financial markets. The investigation of specific subordinators that might better capture asset price has been sparse in the finance literature. Our goal is to update Clark's results from an economic viewpoint and to validate and generalize them from a modeling viewpoint using high frequency data series.

Since the 1980s, the relationship between trading volume and stock prices has been investigated in an impressive body of empirical and theoretical literature. Virtually all empirical studies establish a positive correlation between volatility—measured as absolute or squared price changes–and volume (see Karpoff (1987), Gallant, Rossi, and Tauchen (1992)). More recently, Jones, Kaul, and Lipson (1994) study daily prices of Nasdaq securities and conclude that it is the number of trades and not their size that generates volatility: "The average trade size has virtually no explanatory power when volatility is conditioned on the number of transactions."

Building on the result of Jones et al., we examine high frequency data on two major technology stocks, Cisco Systems and Intel, to determine whether it is the volume, as in Clark (1973), or the number of trades that best defines the business time. Introducing a general stochastic time change τ rather than a subordinator and making no a priori assumption on its distribution, we show that the clock that allows one to recover normality for asset returns is indeed defined by the number of trades. Using classical kernel density estimators, we construct empirical distributions of trades and stock returns, identify the distribution of the time change through its moments, and then reconstruct the density of the returns in the new "transaction time." Remarkably, this density is virtually normal. Last, because a stochastic clock naturally leads to stochastic volatility, we show how stochastic volatility models can be related to stochastic time changes.

The remainder of the paper is organized as follows. In Section I we recall some results on subordinating processes and explain the difference between these special processes and stochastic time changes, from both mathematical and modeling standpoints. In Section II, we describe our database and its statistical characteristics. In Section III we present our model and the main result of the paper, namely, the economic identification of the stochastic clock providing normality of asset returns. Section IV is dedicated to the relationship between stochastic time changes and stochastic volatility, and Section V contains some concluding comments.

I. Subordinators and Stochastic Time Changes

Subordinated processes in mathematics were first introduced in the domain of analysis—and not probability theory—in work on Fourier transforms and Laplace transforms (which, by today, happen to have become classical tools in mathematical finance). As defined by Bochner (1955), a subordinator $\tau(t)$ is a right-continuous increasing process that has independent and homogeneous increments. Besides the obvious degenerate case in which $\tau(t)$ is equal to t times a positive constant, the fundamental examples of subordinators are the Poisson process (simple or compound), the gamma process (see Madan and Seneta (1990)), and the stable process. For any process $X(t)$, the process $Y(t) = X(\tau(t))$ is called a subordinated process.

At the time of Clark's paper, the mathematical properties of subordination were well established and hence readily available for financial applications. However it is clear today that the condition of independent, identical increments imposed on $\tau(t)$ is far too constraining in finance and eventually undesirable (e.g., inconsistent with sampling a process at irregularly spaced dates). General *stochastic time changes*, defined by an increasing process τ and a given process X with certain desirable properties, provide all the flexibility necessary to represent any return process Y_t in the form $Y(t) = X(\tau(t))$. In Clark (1973), as in this paper, X is chosen to be the Brownian motion. Moreover, our choice for X is validated by the fact that any arbitrage-free return process Y_t can be written as a time-changed Brownian motion.[1] If the time change τ is continuous, the process $Y(t) = X(\tau(t))$ is also continuous, because the trajectories of the Brownian motion are everywhere continuous. If we want to account for periods of intense market activity, during crashes, for instance, τ may include a jump component that will in turn lead to a jump-diffusion representation for the process Y_t.

The idea of "business time" can be traced back to papers published by the NBER in the 1940s; Burns and Mitchell (1946) transform economic data to an alternative timescale based on stages of business cycles. This economic clock first appeared in finance in the work of Clark (1973) in the form of cumulative volume. Clark's stated goal was to "present and test an opposing hypothesis" to Mandelbrot's (1963) stable processes. But we can observe that the modeling processes of Clark and Mandelbrot are not inherently contradictory, because a stable process (of index α) is a Brownian motion subordinated to another stable process (with index $\alpha/2$). They only differ by the choice of the subordinator: log-normal for Clark, stable for Mandelbrot. Our claims in this paper are as follows: (a) All processes defining asset returns can be represented as time-changed Brownian motions; hence, identifying the asset return process is tantamount to identifying the time change; and (b) This time change is the key element accounting for information arrival

[1]The no arbitrage assumption implies the existence of a probability measure Q under which discounted stock prices are martingales. Hence, the stock returns processes have to be semimartingales under the original probability measure P. Monroe (1978) established that any semimartingale is a time-changed Brownian motion.

and market activity and must have the right modeling properties to fit the available data on asset returns. Because the properties of the Brownian motion are well known, we can capture in the transaction time τ the specific traits of a given market in addition to the information flow during the period under scrutiny. As such, we adopt a nonparametric approach and do not specify in advance any distribution for τ; rather we identify the distribution of τ through its moments.

II. Empirical Analysis

We analyze tick-by-tick data of two technology stocks, namely, Cisco Systems and Intel, which are both traded on the Nasdaq. The data were obtained from Reuters. We characterize an asset's returns in a standard manner as the change in the logarithm of the stock price over a given interval of time. We let P_t denote the stock price and define the return process in calendar time as

$$Y_t = Y(t) = \ln \frac{P_t}{P_{t-1}}, \tag{1}$$

where the return is observed at equally spaced calendar intervals. After introducing a time change τ that transforms the calendar time into the operational time through the bijective mapping $s = \tau(t)$, we define the return in operational time by

$$X(\tau(t)) = \ln \frac{P(\tau(t))}{P(\tau(t) - 1)} \tag{2}$$

where the length of the interval $[\tau(t) - 1, \tau(t)]$ represents one unit of operational time. To take an elementary example, if $\tau(t) = 2t$ (e.g., a deterministic acceleration of time), $X(\tau(t))$would represent the return over the calendar interval $[2t - 2, 2t]$.

Our task is to identify the economic proxy of the timescale providing normality of asset returns. Numerous empirical studies examine the contemporaneous behavior of volume and absolute price changes and document a positive correlation between the two; for a survey up to 1987, see Karpoff (1987). More recent empirical investigations, including Gerety and Mulherin (1989) and Stephan and Whaley (1990), focus on the intraday patterns in volume and price volatility and find similar correlations. Gallant et al. (1992) use a semi-nonparametric method to estimate the joint density of price change and volume. On the other hand, Easley and O' Hara (1992) show that the number of trades is a good indicator of the rate of information flow, whereas Blume, Easley, and O' Hara (1994) observe that volume provides information on the quality of market information. Hence, our purpose is to investigate two possible representations of the time change $\tau(t)$. In the first case, as in

Clark (1973), $\tau(t)$ is an affine function of the cumulative traded volume up to time t, which we denote V_t. In the second case, as in Jones et al. (1994), $\tau(t)$is an affine function of the number of trades cumulated up to time t, denoted T_t. Hence, the natural first step is to empirically test these two choices.

We examine the one-minute and 10-minute Cisco Systems returns and also the five-minute and 15-minute Intel returns over the period January 2, 1997, through December 31, 1997. For each stock, our data provide the prices, together with the time of the transaction and the trading volume; transactions that do not give rise to a change from the last recorded price are also included. These data are then used to create equally spaced series at one-minute and 10-minute intervals of stock returns, trading volumes, and number of trades for the Cisco Systems security and at five-minute and 15-minute intervals for the Intel security. The trading volume and the number of trades during a particular time interval represent the increments of the quantities defined as V_t and T_t. The descriptive statistics of the three series for each stock and time resolution are displayed in Table I.

We first observe that each return series exhibits a slight skewness and a kurtosis significantly greater than three—traditional evidence of non-normality. To compare the trading volume and the number of transactions in terms of their power to explain variance changes (i.e., of stochastic volatility), we follow a procedure introduced in Schwert (1990) and compute unbiased estimates of return standard deviations. We proceed in the following order:

1. We perform a regression of the return Y_t over 12 lagged returns (any number between 10 and 15 gives approximately the same results),

$$Y_t = \sum_{j=1}^{12} \delta_j Y_{t-j} + \varepsilon_t. \tag{3}$$

   The lagged returns are used as regressors to estimate short term movements in conditional expected returns, and the residuals in regression (3) represent unexpected returns.

2. We define as an estimate of the volatility change at time t the quantity[2]

$$\hat{\sigma}_t = \sqrt{\frac{\pi}{2}} |\hat{\varepsilon}_t|. \tag{4}$$

[2] An elementary result on the Gaussian distribution is that if $X \approx N(0,\sigma^2)$, then $E(|X|) = \sqrt{2/\pi}\sigma$.

Table I

Descriptive Statistics of the Data

Descriptive statistics of the two technology stock returns, trading volume, and number of trades for the different time intervals over the period January 2, 1997, to December 31, 1997, are presented in this table. The data have been collected from Reuters. The number of observations for the period of analysis is n = 101,707 and 10,171 for the Cisco Systems series on one-minute and 10-minute intervals and n = 20,352 and 6,784 for the Intel series on 5-minute and 15-minute intervals. Whereas population skewness and kurtosis for a normal distribution are, respectively, 0 and 3, the variance coefficients of skewness and kurtosis for a size n sample extracted from a normal population are $6/n$ and $24/n$, respectively. The number of observations we are analyzing is sufficiently large to invoke the central limit theorem and to assume that sample coefficients are normally distributed. Thus, the 95% confidence intervals for a test of stock returns normality are given by $\pm 1.96 \times \sqrt{6/n}$ for the sample skewness and $3 \pm 1.96 \times \sqrt{24/n}$ for the sample kurtosis. For statistics computed on the returns Y_t, sample skewness and kurtosis values outside these confidence intervals indicate significant departures from normality, which is clearly the case here.

		Mean	Variance	$m_3(Z)$	Skewness	$m_4(Z)$	Kurtosis	$m_5(Z)$	$m_6(Z)$
					Panel A: Cisco Systems				
					One-Minute Returns				
Returns	$Z = Y_t$	4.229270E−06	4.997559E−07	8.004071E−11	2.265552E−01	1.020305E−11	4.085210E+01	1.286759E−15	7.187722E−17
Traded volume	$Z = V_t$	8.764378	2.6148632	−3.2579457	−7.704960E−01	42.864665	6.26904645	493.5682	5899.51385
Number of trades	$Z = T_t$	2.176064578	1.201186697	−1.931194E−02	−1.466933E−02	17.96898447	12.45381774	231.9889886	2457.402228
					10-Minute Returns				
Returns	$Z = Y_t$	4.228270E−05	3.475961E−06	2.070135E−09	3.194379E−01	2.059287E−10	17.04383066	1.281678E−14	8.031443E−15
Traded volume	$Z = V_t$	85.2078034	15.458209	−4.59984065	−0.075683991	528.256079	2.210681383	2077.213637	11038.57328
Number of trades	$Z = T_t$	21.67736593	21.57279083	−27.87452855	−0.278194175	7655.118284	16.44899015	−13578.8581	122695.116
					Panel B: Intel				
					Five-Minute Returns				
Returns	$Z = Y_t$	−5.2863E−06	1.024171E−06	−9.871156E−11	−9.523783E−02	2.226697E−11	21.22835857	−1.062492E−15	2.061844E−16
Traded volume	$Z = V_t$	17.21842406	4.2508945	1.03457629	0.118043487	88.4210365	4.893222322	632.1749035	8013.035466
Number of trades	$Z = T_t$	7.74524156	2.4295205	−1.63127814	−0.4307717	43.575287	7.382426305	457.428425	5227.11304
					15-Minute Returns				
Returns	$Z = Y_t$	−1.5839E−05	7.096932E−07	−1.043905E−10	−1.746043E−01	6.179025E−12	12.26814006	−1.360838E−16	3.038698E−17
Traded volume	$Z = V_t$	51.85321567	11.0563972	2.6083671	0.070949325	876.606294	7.17096033	3150.304752	14027.7482
Number of trades	$Z = T_t$	23.687323	3.3259344	−5.459321	−0.900053059	331.7532738	29.99078678	1066.311298	16378.64235

To compare the performance of the volume and the number of trades in explaining volatility changes, we run the following three regressions:

$$|\hat{\sigma}_t| = \alpha + \beta \Delta V_t + \sum_{j=1}^{12} \rho_j |\hat{\sigma}_{t-j}| + \eta_t^1, \tag{5a}$$

$$|\hat{\sigma}_t| = \alpha + \gamma \Delta T_t + \sum_{j=1}^{12} \rho_j |\hat{\sigma}_{t-j}| + \eta_t^2, \tag{5b}$$

$$|\hat{\sigma}_t| = \alpha + \beta \Delta V_t + \gamma \Delta T_t + \sum_{j=1}^{12} \rho_j |\hat{\sigma}_{t-j}| + \eta_t^3, \tag{5c}$$

where $\Delta V(k) = V(k) - V(k-1)$ and $\Delta T(k) = T(k) - T(k-1)$. The lags of the estimated standard deviation series are included to accommodate any persistence in the futures price volatility.

For each security and under the various timescales, the results summarized in Table II indicate the superiority of the number of trades in explaining volatility changes, which supports the results of Jones et al. (1994), who also analyze Nasdaq stock returns. For instance, the value 0.214708 obtained for the adjusted-R^2 of regression (5b) in the case of one-minute Cisco Systems returns using the number of trades as the sole explanatory variable is strictly higher than the value 0.167614 obtained in regression (5a). Moreover, the small difference (0.215899 − 0.214708) between the adjusted-R^2 of regressions (5b) and (5c) supports the hypothesis that the trading volume adds virtually no explanatory power when returns are already conditioned on the number of transactions. This observation holds for every return series presented in this paper.

Last, we construct the empirical distributions of the stock returns, trading volume, and number of transactions. The density estimation is traditionally performed using the kernel method, that is, the estimator

$$\hat{f}(x) = \frac{1}{nh} \sum_{i=1}^{n} K\left(\frac{x - x_i}{h}\right), \tag{6}$$

where n = number of observations, x_i = observation i, h = window width (also called the smoothing parameter), and K = the kernel estimator.

We choose as a kernel $K(x) = (1/\sqrt{2\pi})e^{-x^2/2}$, which ensures that $\hat{f}$ is a smooth curve having derivatives of all orders and, hence, ensures the existence of the first six moments of Y_t, which play a key role below. Following the standard approach (see, e.g., Silverman (1986)), we use for the window width the quantity $h = \sigma(4/3)^{1/5} n^{1/5}$, where σ represents the standard deviation of the whole series. The estimated density distribution functions of the different return series are given in Figure 1. They are more peaked than the Gaussian distribution and exhibit thinner flanks and fatter tails.

Table II

Regression Estimates

Cisco Systems and Intel returns series are regressed, respectively, on trading volume and numbers of trades over the period January 2, 1997, to December 31, 1997. The data were collected from Reuters. Standard errors for the coefficients of the regressions are provided in parentheses. The higher adjusted-R^2 of regressions (5b) when compared to regressions (5a) indicate that the number of trades T_t has a higher explanatory power of volatility changes. The values of the coefficient γ in regressions (5c) are not statistically different from the values of this coefficient in regressions (5b). On the contrary, the values of the coefficient β are significantly lower than in regressions (5a). Moreover the R^2 in regressions (5b) and (5c) are not statistically different. This indicates that not only does the number of trades have a higher explanatory power but also that the informational content of the traded volume is included in the informational content of the number of trades.

	β	γ	R^2	β	γ	R^2
	Panel A: Cisco Systems			Panel B: Intel		
	One-Minute Returns			10-Minute Returns		
Regression (5a)	7.47216E−05 (0.00000253)	—	0.167614	2.90345E−04 (0.000132)	—	0.140566
Regression (5b)	—	3.82534E−04 (0.000108)	0.214708	—	2.28340E−03 (0.000507)	0.199388
Regression (5c)	1.09674E−05 (0.0000378)	3.54379E−04 (0.000098)	0.215899	1.97693E−04 (0.000065)	2.10976E−03 (0.000773)	0.201327
	Five-Minute Returns			15-Minute Returns		
Regression (5a)	4.39265E−05 (0.0000208)	—	0.203811	9.85257E−05 (0.0000465)	—	0.182301
Regression (5b)	—	3.07456E−04 (0.000102)	0.286968	—	7.30843E−04 (0.000111)	0.227614
Regression (5c)	2.97855E−06 (0.0000102)	3.32056E−04 (0.000099)	0.287031	3.20867E−06 (0.00000203)	7.08942E−04 (0.0000100)	0.229031

Figures 2 and 3 give, respectively, the empirical density distributions of the trading volume and the number of trades for different time resolutions, also reconstructed through the kernel method. Stephan and Whaley (1990) show there is an interesting U-shaped pattern when the trading volume is plotted as a function of the time of day. We cannot directly compare our Figures 2 and 3 to the results of these authors, because our analysis does not provide intraday patterns as a function of calendar time but rather shows return changes as functions of market activity.

III. The Model

We know at this point the empirical distribution—hence, the empirical moments—of the asset returns and of the cumulative number of trades. As explained in Section I, we now search for a stochastic clock $\tau(t)$ and the two parameters μ and σ of an arithmetic Brownian motion such that

$$Y(t) = X(\tau(t)) \tag{7}$$

and

$$X(s) \overset{law}{=} N(\mu s, \sigma^2 s). \tag{8}$$

The unconditional centered moments of the process $Y(t)$ can be expressed in terms of the parameters μ, σ and of the centered moments of the process $\tau(t)$. The computation of these theoretical moments is presented in Appendix A. However, a direct estimation procedure is not possible because any systems formed with these (nonlinear) equations will have more unknowns than equations—for instance, with the theoretical moments computed up to order six, we obtain a nonlinear system of six equations with eight unknowns.

To recover sufficient information in the estimation procedure, we use the moment generating function of the return process Y_t and define the following minimization program:

$$Min \quad U = \sum_{j=1}^{k} [E[\exp(\beta_j Y_t)]^{theoretical} - E[\exp(\beta_j Y_t)]^{empirical}]^2 \tag{9}$$

for an appropriate choice of the numbers $\beta_1, \beta_2, \ldots, \beta_k$ and under the equality constraints $m_i(Y_t)^{theoretical} = m_i(Y_t)^{empirical}$ for $i = 1, 2, \ldots$. Appendix B explains how the moment generating function of the return process in calendar time can be expressed in terms of the unknown parameters of our estimation problem. An empirical estimator of the moment generating function is also introduced to solve the minimization problem.

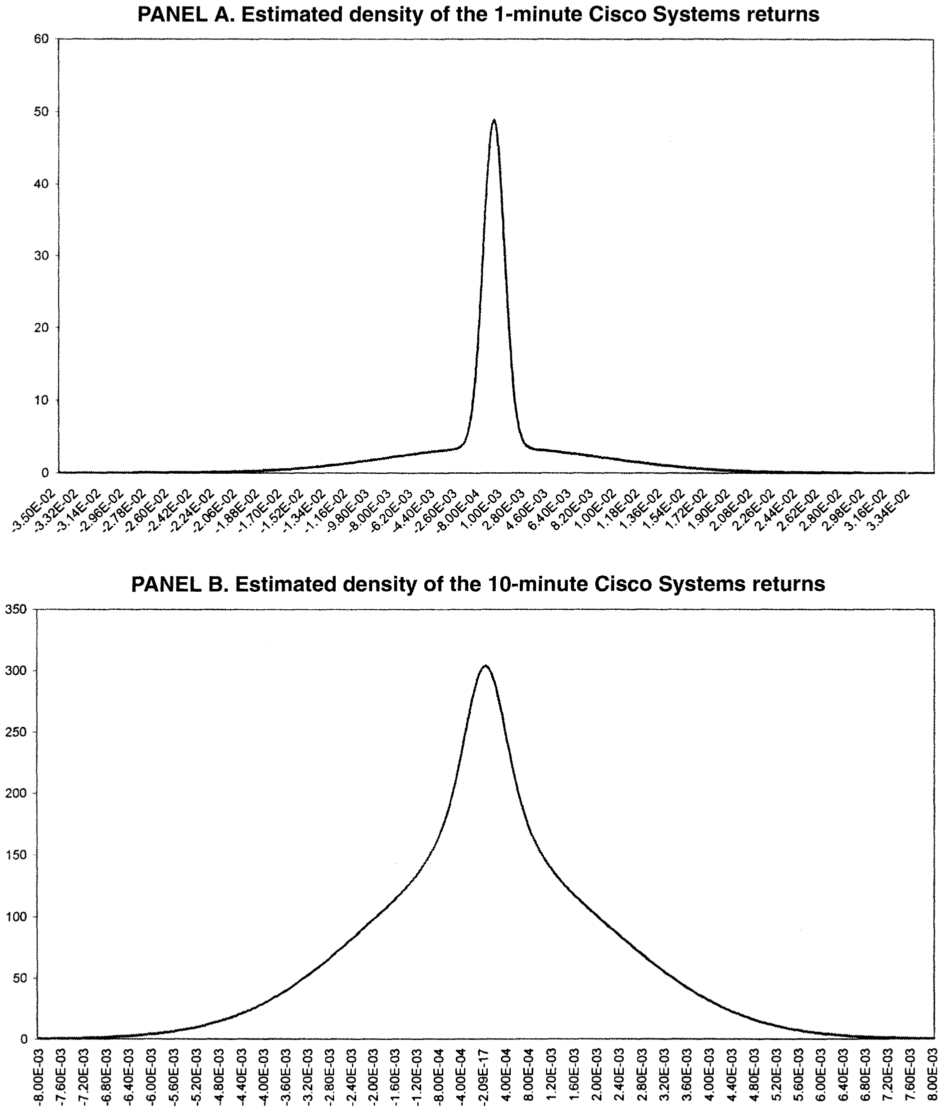

Figure 1. Estimated densities of Cisco Systems and Intel returns. This figure represents, for four different time resolutions, the density distributions of Cisco Systems and Intel returns. These are reconstructed through the kernel method, the horizontal axis representing the real values of asset returns. All distributions are more peaked than the Gaussian one and exhibit leptokurtosicity. (*Figure continues on facing page.*)

We perform our optimization procedure with several sets of values and obtain nearly identical solutions irrespective of the values of β used. The results of some of these minimizations are summarized in Table III.

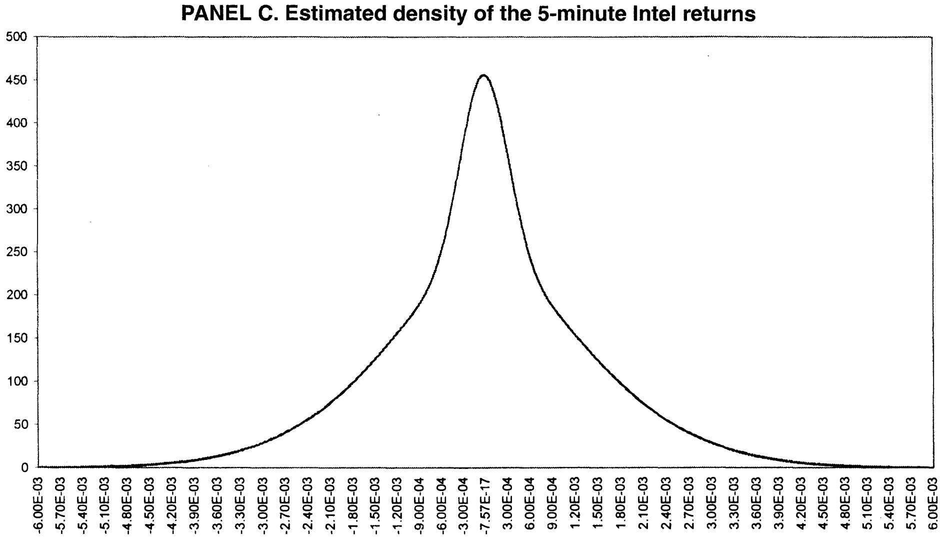

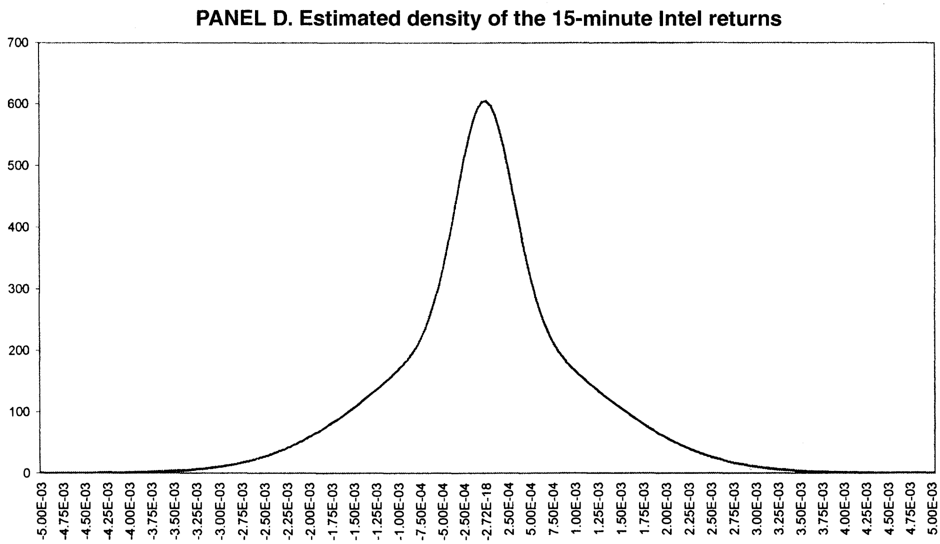

Figure 1. *Continued.*

A comparison of our findings with the empirical data presented in Table II establishes that the moments of the time change $\tau(t)$ greater than one are perfectly matched by the moments of the cumulative number of transactions T_t; only the mean differs significantly. If we make the simplifying (but standard in finance) assumption that a probability distribution is defined by the knowledge of its first several (usually four) moments, then it appears that the cumulated number of transactions (up to a constant) is a good

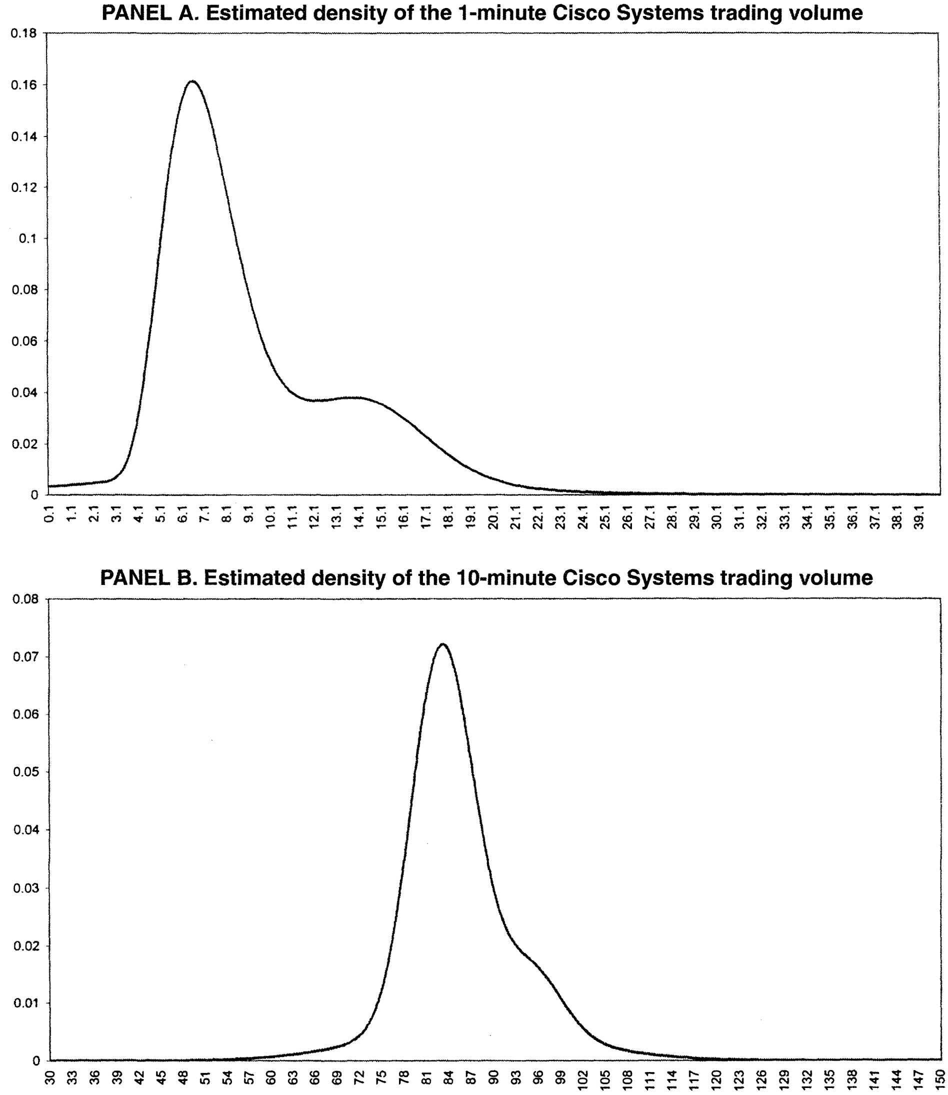

Figure 2. Estimated densities of Cisco Systems and Intel trading volumes. This figure displays the distributions of Cisco Systems and Intel trading volumes for the different time resolutions examined in Figure 1. Densities are reconstructed through the kernel method. The horizontal axis represents the trading volume in real values. (*Figure continues on facing page.*)

representation of the economic time. To ascertain whether the moments of the time change match the moments of the cumulative number of transactions, we recenter the time change so that its mean equals the mean of the number of transactions and then reconstruct the empirical density of the time change.

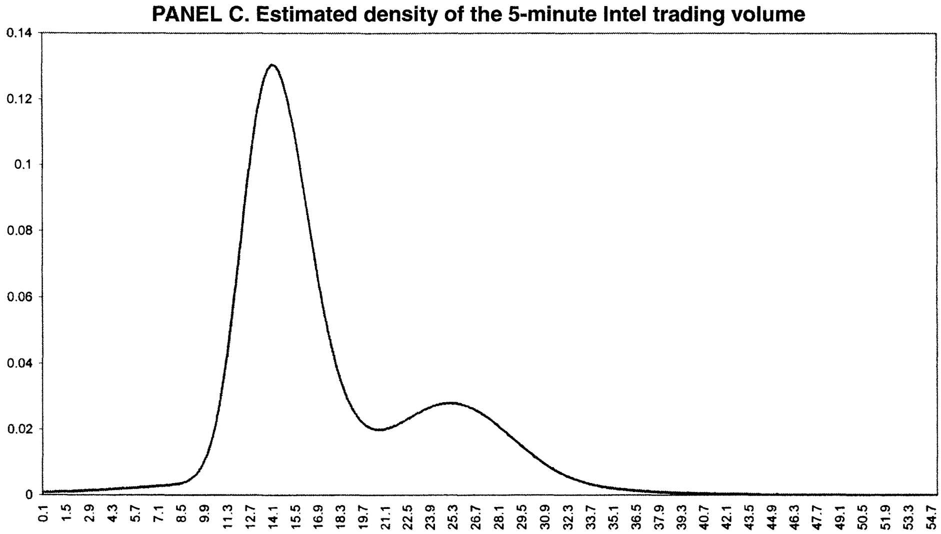

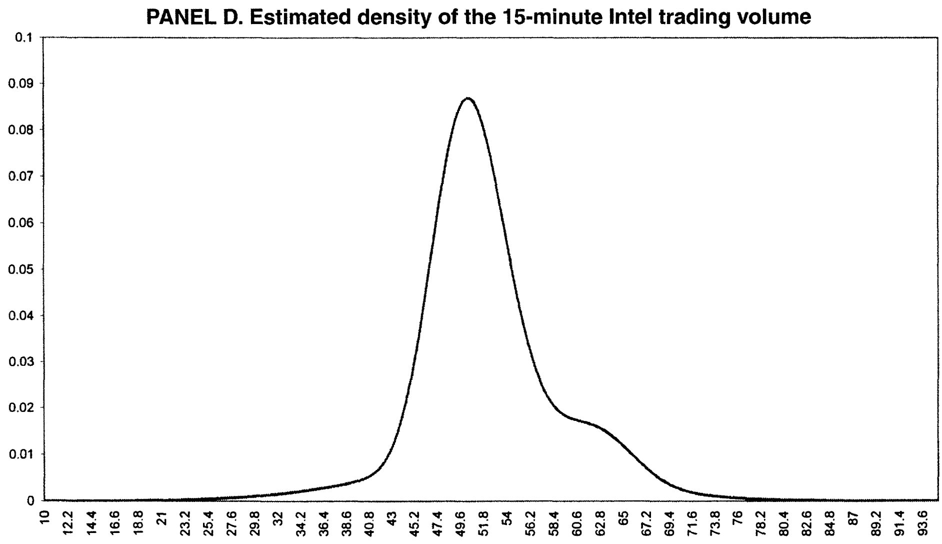

Figure 2. *Continued.*

We then observe that because $Y_t = X(\tau(t))$,

$$P(Y_t \in dy/\tau(t) = u) = P(X(u) \in dy/\tau(t) = u). \tag{10}$$

Assuming for simplicity, as does Clark (1973), the independence of X and τ, the previous relationship reduces to

$$P(Y_t \in dy/\tau(t) = u) = P(X(u) \in dy), \tag{11}$$

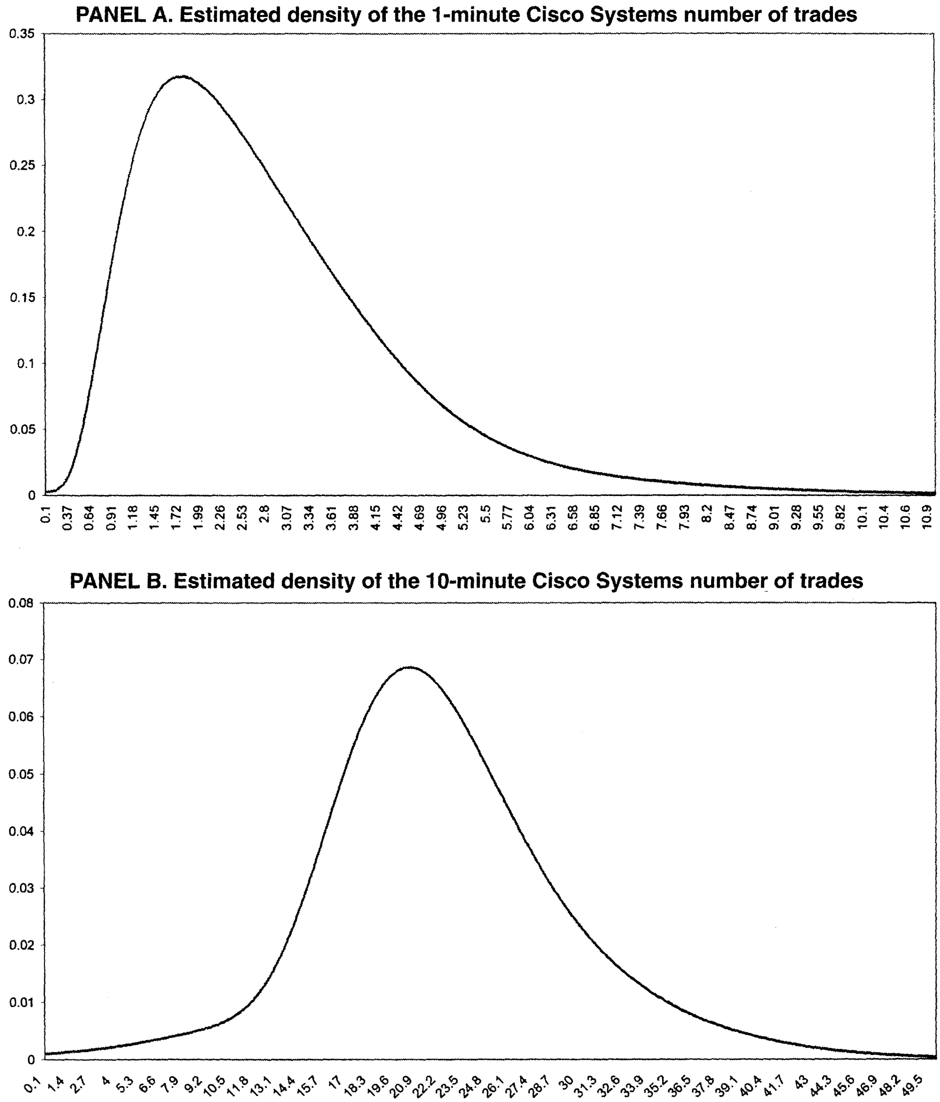

Figure 3. Estimated densities of Cisco Systems and Intel number of trades. This figure displays the distributions of Cisco Systems and Intel numbers of trades for the different time resolutions. The estimated density is reconstructed through the kernel method. The horizontal axis represents the real values of the number of trades. (*Figure continues on facing page.*)

where X is Brownian motion. Hence, if the time change has been properly chosen, the distribution of the return process Y, conditional on the time change τ, should be normal.

To test for this property in our data, we compute the series of the returns conditional on the (recentered) number of trades and use the kernel estimator to reconstruct the empirical conditional density. This density is plotted in

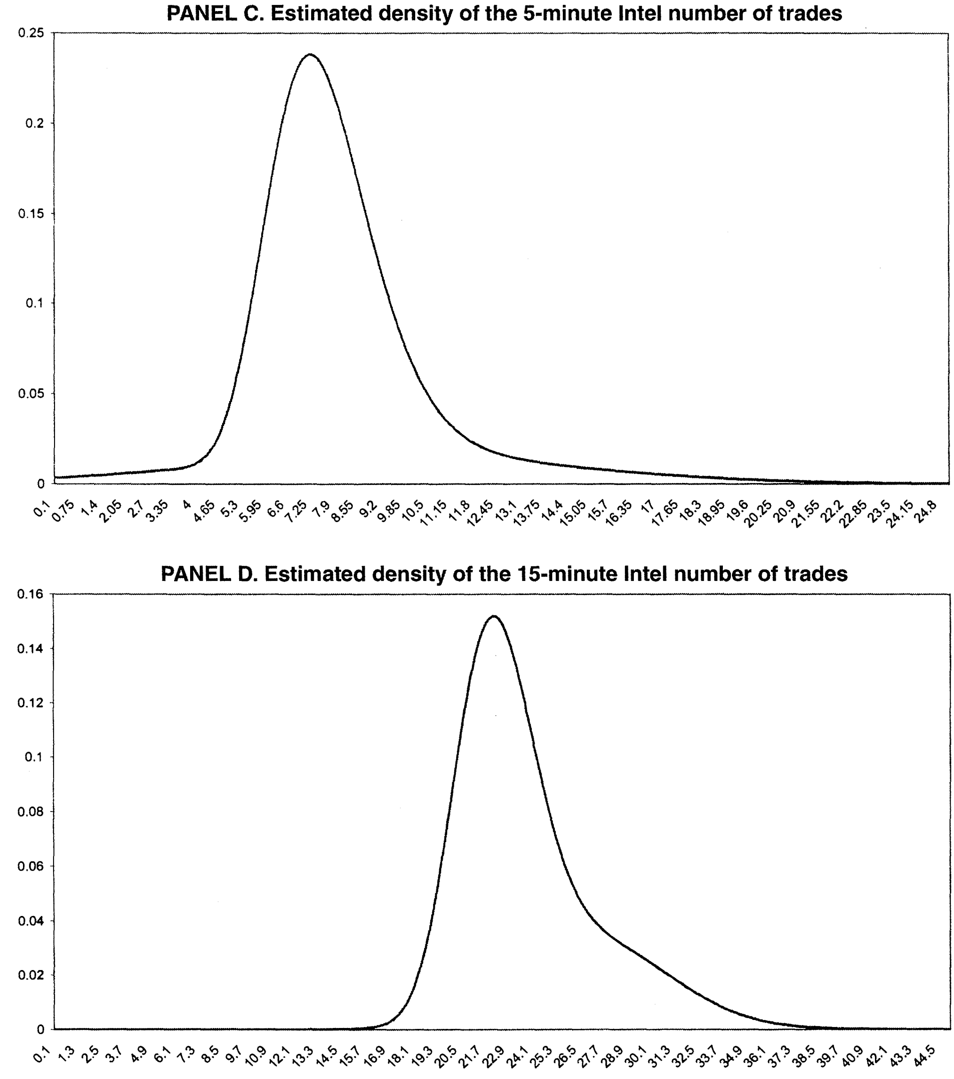

Figure 3. *Continued.*

Figure 4, together with a Gaussian distribution with the same mean and variance, for each time resolution and stock. The near-perfect normality strongly supports our conjecture—namely, that the time change generating conditional normality for the return process is properly represented by the number of trades, independent of any direct parametric representation of Y.

To confirm this result, we perform a chi-square test of goodness-of-fit on the conditional recentered series. The critical value at the one percent confidence level is $\chi^2 = 11.6$. The empirical statistics, $\hat{\chi}^2 = 5.21$, 7.09, 6.34, and

Table III
Moment Estimates

This table shows the estimated values of the unknown moments of the underlying Gaussian process and of the unobserved stochastic clock $\tau(t)$ obtained by the optimization program in equation (9). We minimize, under equality constraints on the centered theoretical and empirical moments, the sum of the square differences between the theoretical moment generating function of the return process (expressed in terms of the unknown parameters as discussed in Appendix B) and the empirical moment generating function of the return process. The minimization program has also been run with equality constraints on the uncentered theoretical and empirical moments of the return process, and the same values were found for the unknown parameters, exhibiting the robustness of the minimization procedure.

	μ	σ^2	$E(\tau_{(t)})$	$m_2(\tau_{(t)})$	$m_3(\tau_{(t)})$	$m_4(\tau_{(t)})$	$m_5(\tau_{(t)})$	$m_6(\tau_{(t)})$
Panel A: Cisco Systems								
One-Minute Returns								
Optimization value	1.371329E−05	1.619712E−06	3.084067E−01	1.2028836	−1.967611E−02	18.4472756	232.0183	2457.838915
Standard error	6.48208E−06	3.56967E−07	0.00237	0.57329	0.00474	2.06474	16.06836	30.68394
10-Minute Returns								
Optimization value	1.972051E−05	1.617121E−06	2.140840668	21.6328604	−27.6655439	7654.064486	−13577.01045	122697.2334
Standard error	7.20686E−06	8.90235E−07	0.48757	1.07998	2.08579	6.07687	23.86068	85.65343
Panel B: Intel								
Five-Minute Returns								
Optimization value	−8.3616E−06	1.6148E−06	6.3221E−01	2.432755	−1.66371814	43.8332287	457.337425	5229.06326
Standard error	2.46329E−07	2.07573E−07	3.30746E−03	0.10750	0.34524	1.53088	9.04525	23.96368
15-Minute Returns								
Optimization value	−1.5290E−05	6.8437E−07	1.0358626	3.314327	−5.48911321	332.614213	1068.01078	16380.33582
Standard error	7.86961E−06	7.6434E−08	0.10565	0.48460	0.79994	8.45207	11.96537	56.32110

8.25, respectively obtained for the one-minute and 10-minute Cisco Systems series and for the five-minute and 15-minute Intel series, indicate that the Gaussian hypothesis cannot be rejected. Applying the Jarque–Bera test of normality for each sampling frequency and each security also indicates that normality of the recentered conditional returns cannot be rejected at the one percent confidence level (with a critical Jarque–Bera statistic of 9.21 and test statistics of 0.337, 3.256, 1.440, and 3.907, respectively, for the one-minute and 10-minute Cisco Systems series and for the five-minute and 15-minute Intel series). Moreover, we provide in Figure 5 the quantile–quantile plots for each estimated density conditioned by the recentered number of trades. The perfect graphical fit is further strong evidence of the Gaussian nature of the series.

It is worth mentioning that the normality of the return process conditional on the number of trades was also exhibited in an empirical study we conducted independently on two high frequency databases of S&P 500 Futures prices (see Geman and Ané (1996)) and FTSE 100 index values. This indicates that the results presented here are not an accidental property of the data because they hold for various classes of equity instruments, each having clocks with their own tick rates. (The study of the multidimensional time-change problem is left for further research.)

Last, let us illuminate a property derived from our model. An elementary formula on conditional probabilities allows us to write

$$P(Y_t \in dy) = \sum_u P(Y_t \in dy/\tau(t) = u) f_\tau(u), \tag{12}$$

where f_τ denotes the distribution of τ that we assume, for the simplicity of argument, to be discrete. Hence,

$$P(Y_t \in dy) = \sum_u P(X_u \in dy) f_\tau(u) \tag{13}$$

from the independence of X and τ. Inspection of equation (13) reveals that the unconditional distribution of Y appears as a *mixture of normal distributions*. The mixture of distributions hypothesis (sometimes referred to as MDH) is a well-known representation of asset returns, and it has often been offered in the financial literature to model the observed leptokurtosis in returns. (For example, Richardson and Smith (1994) examine the MDH hypothesis empirically and propose a direct test of it using the generalized method of moments.) In our framework, the unconditional distribution of returns does behave as a mixture of normals, where the number of trades acts as the mixing variable. Obviously, the resulting process will crucially depend on the distribution of $\tau(t)$. We already mentioned that, in our view, in a number of markets (e.g., the newly deregulated electricity market), the process $(\tau(t))_{t\geq 0}$ should include jumps, which would lead to a mixing process quite different from the one classically presented in the literature.

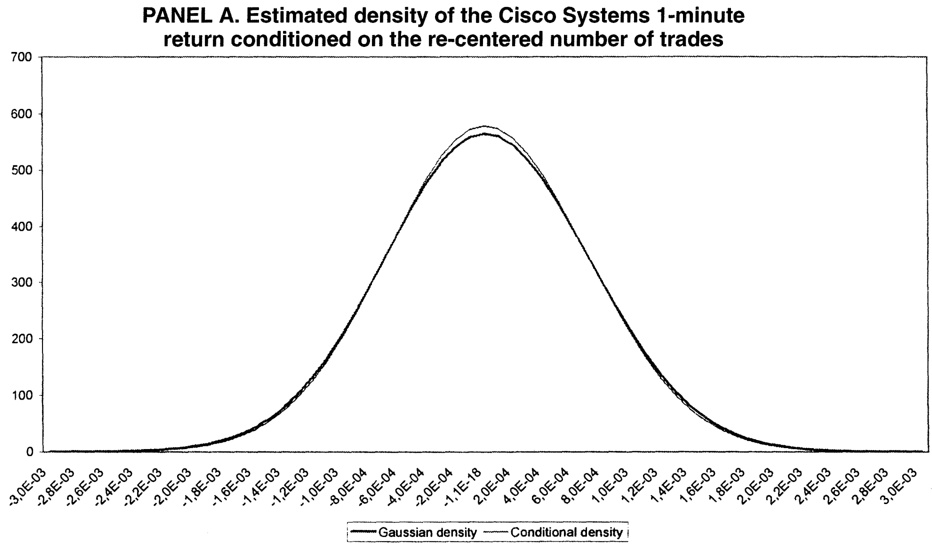

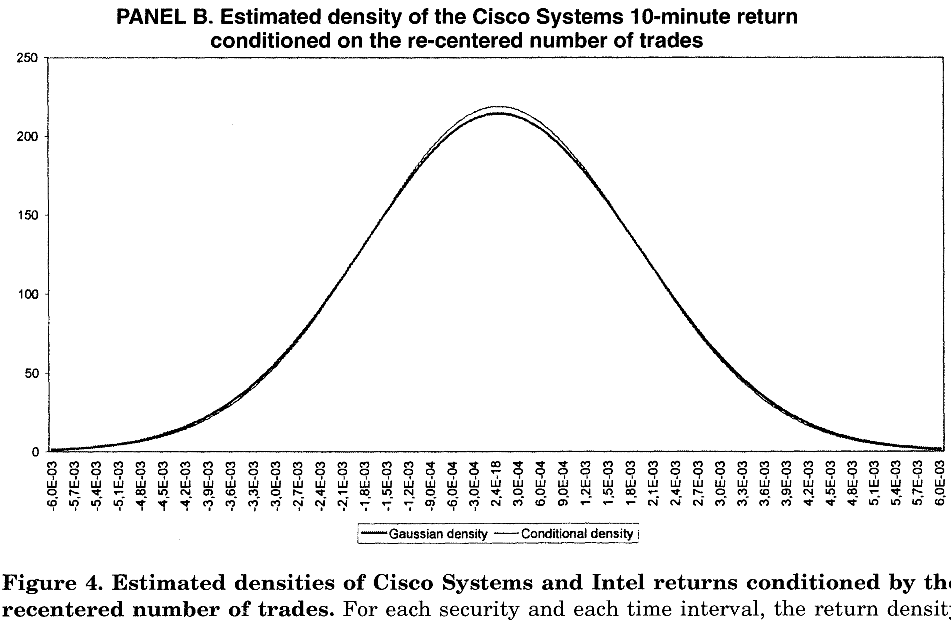

Figure 4. Estimated densities of Cisco Systems and Intel returns conditioned by the recentered number of trades. For each security and each time interval, the return density conditioned on the recentered number of trades is estimated using the kernel method and plotted here with the horizontal axis representing the real values of returns. On the same chart is represented the graph of the Gaussian distribution with the same mean and variance. The near-perfect identity of the two densities indicates that a time change that provides normality of returns is well proxied by an affine function of the number of trades. (*Figure continues on facing page.*)

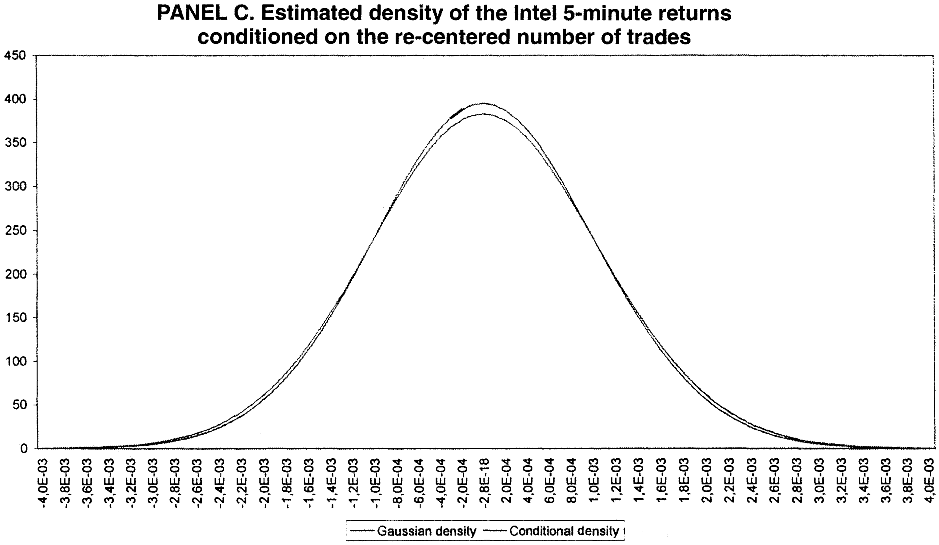

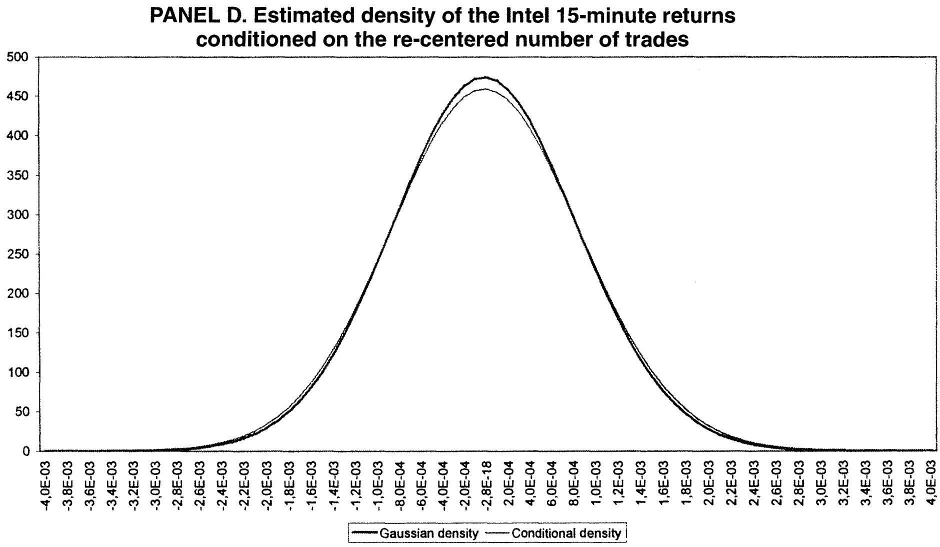

Figure 4. *Continued.*

IV. Stochastic Time Changes and Stochastic Volatility

The second line of equation (A1) in Appendix A shows that the variance of the unconditional return process (Y_t) evolves stochastically because the mean and the standard deviation of the number of trades are not constant over time. Hence, our representation of the return process in calendar time falls

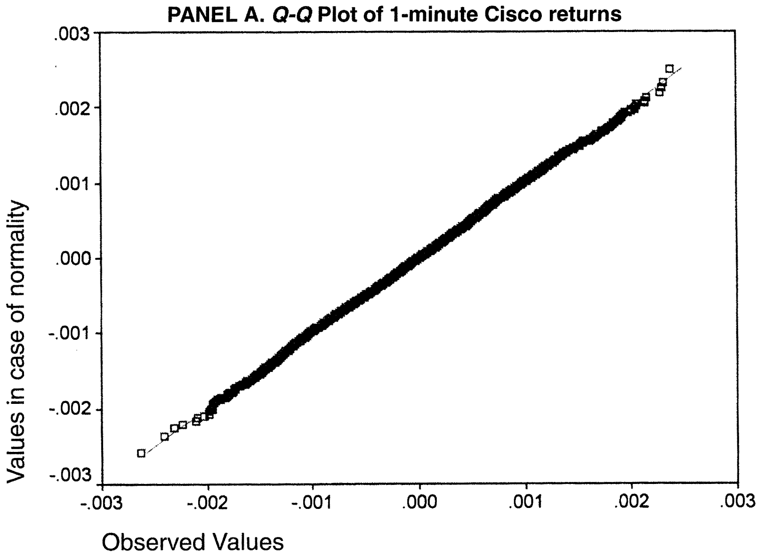

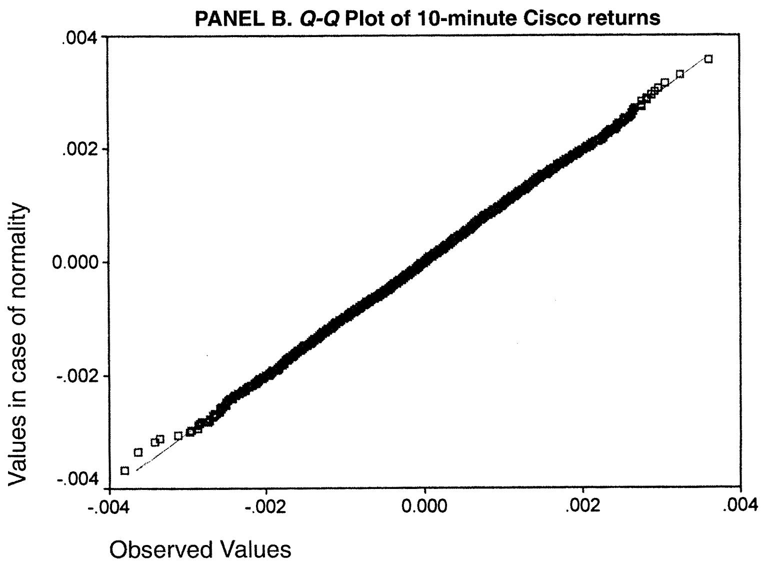

Figure 5. *Q-Q* plots of Cisco Systems and Intel conditional returns. This figure provides the quantile-quantile plots for return densities conditioned on the recentered number of trades, for the four time resolutions analyzed all along. The *Q-Q* plot is a well-known criterion of normality when the set of points is a 45° line. The excellent fits we obtained are a strong evidence of return normality in transaction time. (*Figure continues on facing page.*)

into the group of *stochastic volatility models*. Moreover, the above formula shows that the volatility is monotonically related to the average of the number of trades but also to the unexpected component of the trading activity

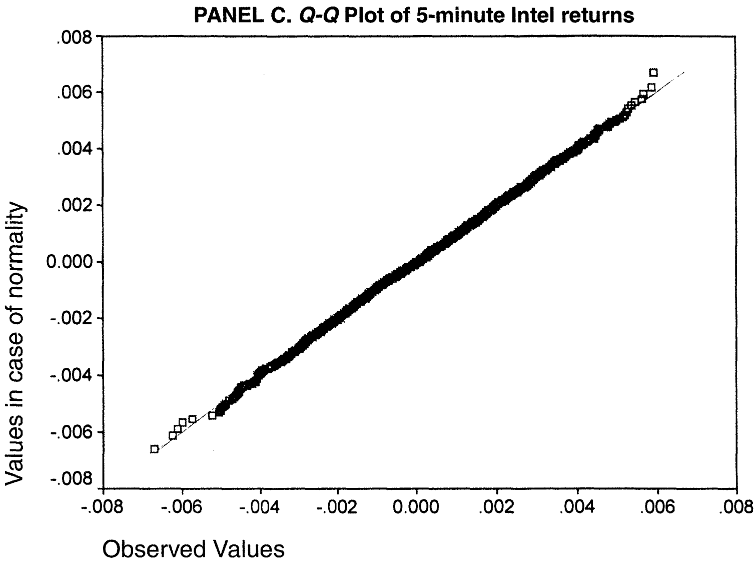

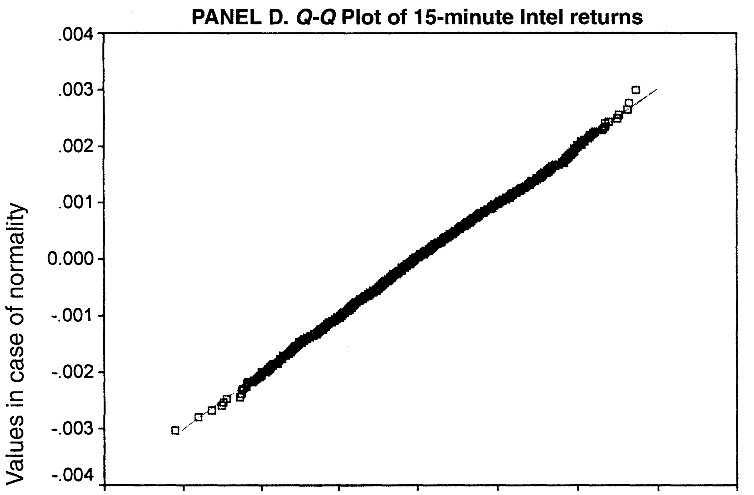

Figure 5. *Continued.*

(represented by $\mathrm{Var}(\tau(t))$). This result is in agreement with the findings of Bessembinder and Seguin (1993), who specifically examine the effects on the volatility of the expected and unexpected transacted volumes.

More generally, our setting can accommodate any type of stochastic volatility model for the return process, the changes in volatility being translated by contraction or dilatation of time (the mathematics associated with the construction of τ are described in the Appendixes). In one of his last papers, Black (1992) comments on the limits of the Black–Scholes formula and its extensions. In particular, he argues that an alternative to the addition of a jump component in the stock price (to model the effects of the arrival of news) is to incorporate a higher volatility in the formula (hence, the importance of the problem of proper volatility estimation). Our modest claim is that the introduction of a business clock τ allows one to reconcile these two possible representations of non-normal returns.

Last, let us illustrate by an example that in option pricing and hedging, the stochastic volatility of the underlying asset price dynamics can be transformed into constant volatility through a stochastic time change. Consider for instance a financial institution that has sold standard European options at time zero on the basis of a Black–Scholes–Merton volatility equal to b, where b denotes a positive constant. Let the asset price dynamics be described by

$$\frac{dS_t}{S_t} = \mu\, dt + \sigma\, dW_t, \tag{14}$$

where $(W_t)_{t\geq 0}$ is a Brownian motion on the probability space (Ω, F, F_t, P) representing the randomness of the economy, P is the statistical probability and F_t is the filtration of information available at time t. If σ is itself stochastic on the interval $[0, T]$, the value of the self-financing portfolio built at time zero with the premium recovered from the sale of the option and dynamically readjusted over time to replicate the option will not equal at maturity T the payoff of the option. Instead of experiencing a deviation from the target, one may be interested in finding the first time τ_b at which perfect replication is achieved, that is, when

$$\int_0^{\tau_b} [\sigma(s)]^2 ds = b^2 T. \tag{15}$$

Obviously, τ_b is a stopping time whose probability distribution is extremely valuable to option traders. As an example, assume the Hull and White (1987) framework, that is, asset price dynamics given by the equations

$$\frac{dS_t}{S_t} = \mu\, dt + \sigma_t\, dW_t^1, \tag{16}$$

$$\frac{dY_t}{Y_t} = \theta\, dt + \zeta\, dW_t^2, \tag{17}$$

where Y_t is equal to σ_t^2 and ρ is the correlation between the Brownian motions W^1 and W^2. Geman and Yor (1993) provide the Laplace transform of the density f_b of the stopping time τ_b, namely, the quantity $\int_0^{+\infty} f_b(x) e^{-\lambda x} dx$, in terms of the parameters θ, ζ defining the volatility process dynamics; as usual, the drift of the stock price plays no role. By inversion of the Laplace transform, it is possible to obtain the complete density $f_b(x)$; if one only wishes to compute the moments of the "perfect replication time" τ_b, one can, using standard properties of Laplace transforms, repeatedly differentiate the transform the appropriate number of times (n times for the nth uncentered moment) and evaluate at zero. In particular, one can obtain the average time between issuance at a volatility level b and optimal replication of the option. (Strictly speaking, the replication process need not be unique, because the presence of stochastic volatility that does not depend solely on S_t means that markets are incomplete.)

Obviously, the use of time changes can be extended to any stochastic volatility model and provides a general and powerful technique to handle problems related to stochastic volatility. In particular, the economic identification of the appropriate time change τ and its density may facilitate the recovery of option prices as the Black–Scholes price for a fixed time integrated against the density of τ. More specifically, the European fixed maturity option price in the case of stochastic volatility will depend on the distribution of the cumulative quadratic variation associated with the underlying asset, which will depend on the nature of the time change. We leave these extensions to future research as our goal here was to provide a precise analysis of the methodology and empirics of the recovery of normality of stock returns.

V. Conclusion

This paper looks at the distribution of a high frequency database of technology stock returns using stochastic time changes rather than subordinators. It is shown that, to recover normality in asset returns, the number of trades is a better time change than the traditionally used trading volume. The near-perfect normality in transaction time is exhibited through the reconstruction of the time-changed return process. No particular distribution is assumed for the time change itself. Instead, the time change is characterized through its moments. Due to this flexibility, the time change can be represented by a lognormal or another continuous distribution; it can also include jumps to account for periods of high market activity, leading in turn to a jump-diffusion model for the asset return process.

Appendix A

In this Appendix, we give the theoretical form of the first six centered moments of the return process in terms of the corresponding moments of the unobservable stochastic time change $\tau(t)$ and the two moments (μ, σ^2) of the underlying Gaussian distribution. Denoting by $m_i(Y)$ the centered moment of order i of the variable Y, and using the assumption of independence of X and τ, we can express the theoretical moments of the return process Y_t as follows:

$$
\begin{aligned}
E(Y_t) &= \mu E(\tau(t)),\\
Var(Y_t) &= m_2(Y_t) = \sigma^2 E(\tau(t)) + \mu^2 Var(\tau(t)),\\
m_3(Y_t) &= 3\mu\sigma^2 Var(\tau(t)) + \mu^3 m_3(\tau(t)),\\
m_4(Y_t) &= \mu^4 m_4(\tau(t)) + 6\sigma^2\mu^2 m_3(\tau(t)) + 6\sigma^2\mu^2 E(\tau(t))Var(\tau(t))\\
&\quad + 3\sigma^4[Var(\tau(t)) + [E(\tau(t))]^2],\\
m_5(Y_t) &= \mu^5 m_5(\tau(t)) + 10\sigma^2\mu^3 m_4(\tau(t)) + (10\sigma^2\mu^3 E(\tau(t))\\
&\quad + 15\mu\sigma^4)m_3(\tau(t)) + 30\mu\sigma^4 Var(\tau(t)),\\
m_6(Y_t) &= \mu^6 m_6(\tau(t)) + 15\sigma^6(m_3(\tau(t)) + 3m_2(\tau(t))E(\tau(t)) + (E(\tau(t)))^3)\\
&\quad + 15\mu^4\sigma^2(m_5(\tau(t)) + m_4(\tau(t))E(\tau(t)))) + 45\mu^2\sigma^4(m_4(\tau(t))\\
&\quad + 2m_3(\tau(t))E(\tau(t)) + m_2(\tau(t))(E(\tau(t)))^2).
\end{aligned}
\tag{A1}
$$

Appendix B

Using the properties of conditional expectations and the assumption of Brownian motion for the process $X(s)$, the moment generating function of the return process Y_t can be expressed in terms of the moment generating function of the time change process $\tau(t)$,

$$
\begin{aligned}
E[\exp(\beta Y_t)] &= E[E(\exp(\beta Y_t)/\tau(t)]\\
&= E[\exp(\beta\mu + \tfrac{1}{2}\beta^2\sigma^2)\tau(t)]\\
&= E[\exp(A\tau(t))],
\end{aligned}
\tag{B1}
$$

where $A = \beta\mu + \frac{1}{2}\beta^2\sigma^2$.

The nonparametric setting we adopt in this study gives us no direct information on the distribution of the time change process or on its moment generating function. To obtain the moment generating function of the return process Y_t in terms of μ, σ and the moments of τ, we write a series expansion of $\exp(A\tau(t))$ around the mean of the time change $\tau(t)$:

$$\exp(A\tau(t)) \cong \exp(AE(\tau(t))) * \left[1 + A(\tau(t) - E(\tau(t)) + \frac{A^2}{2}(\tau(t) - E(\tau(t))^2 + \frac{A^3}{6}(\tau(t) - E(\tau(t))^3 + \frac{A^4}{24}(\tau(t) - E(\tau(t))^4\right]. \tag{B2}$$

The expectation of this series expansion yields a theoretical expression of the moment generating function of the directing process $\tau(t)$ and hence of the return process Y_t,

$$E[\exp(\beta_j Y_t)]^{theoretical} \cong e^{AE(\tau(t)} * \left[1 + \frac{A^2}{2} Var(\tau(t)) + \frac{A^3}{6} m_3(\tau(t)) + \frac{A^4}{24} m_4(\tau(t))\right]. \tag{B3}$$

For any given value of β, say, β_j, the moment generating function of the return process Y_t may also be approximated empirically by

$$E[\exp(\beta_j Y_t)]^{empirical} = \frac{1}{n} \sum_{i=1}^{n} \exp(\beta_j y_{t_i}). \tag{B4}$$

The unknown parameters can then be obtained by minimizing the expression

$$Min \quad U = \sum_{j=1}^{k} [E[\exp(\beta_j Y_t)]^{theoretical} - E[\exp(\beta_j Y_t)]^{empirical}]^2 \tag{B5}$$

under the constraints on equality $m_i(Y_t)^{theoretical} = m_i(Y_t)^{empirical}$ for $i = 1, 2, \ldots$. We minimize the sum of squared differences between the sample data moment generating function and the theoretical moment generating function computed for the values $\beta_1, \beta_2, \ldots \beta_k$ of β. Obviously the choice of the values $\beta_1, \beta_2, \ldots \beta_k$ is crucial. Simulations we ran suggest that very large and very small values of β should be avoided.

REFERENCES

Bessembinder, Hendrik, and Paul J. Seguin, 1993, Price volatility, trading volume and market depth: Evidence from futures markets, *Journal of Financial and Quantitative Analysis* 28, 21–35.

Black, Fisher, 1992, Living up to the model; in *From Black-Scholes to Black Holes: New Frontiers in Options* (Risk Publications, London).

Blume, Lawrence, David Easley, and Maureen O' Hara, 1994, Market statistics and technical analysis: The role of volume, *Journal of Finance* 49, 153–181.

Bochner, Salomon, 1955, *Harmonic Analysis and the Theory of Probability* (University of California Press, Berkeley).

Burns, Arthur F., and Wesley C. Mitchell, 1946, *Measuring Business Cycles* (National Bureau of Economic Research, New York).

Clark, Peter K., 1973, A subordinated stochastic process with finite variance for speculative prices, *Econometrica* 41, 135–156.

Easley, David, and Maureen O' Hara, 1992, Time and the process of security price adjustment, *Journal of Finance* 47, 577–605.

Fama, Eugene F., 1965, The behavior of stock market prices, *Journal of Business* 38, 34–105.

Gallant, A. Ronald, Peter E. Rossi, and George Tauchen, 1992, Stock prices and volume, *Review of Financial Studies* 5, 199–242.

Geman, Hélyette, and Thierry Ané, 1996, Stochastic subordination, *Risk* 9, 145–149.

Geman, Hélyette, and Marc Yor, 1993, Bessel processes, Asian options, and perpetuities, *Mathematical Finance* 3, 349–375.

Gerety, Mason S., and Harold J. Mulherin, 1989, Intraday trading behavior in securities markets: Hourly NYSE volume and returns, 1933–1988, Working paper, U.S. Securities and Exchange Commission.

Hull, John, and Alan White, 1987, The pricing of options with stochastic volatilities, *Journal of Finance* 42, 281–300.

Jones, Charles, Gautam Kaul, and Marc L. Lipson, 1994, Transactions, volume and volatility, *Review of Financial Studies* 7, 631–651.

Karpoff, Jonathan M., 1987, The relation between price changes and trading volume: A survey, *Journal of Financial and Quantitative Analysis* 22, 109–126.

Madan, Dilip, and Eugene Seneta, 1990, The variance gamma model for share market returns, *Journal of Business* 63, 511–524.

Mandelbrot, Benoit, 1963, The variation of certain speculative prices, *Journal of Business* 36, 394–413.

Monroe, Itrel, 1978, Processes that can be embedded in Brownian motion, *Annals of Probability* 6, 42–56.

Richardson, Matthew, and Tom Smith, 1994, A direct test of the mixture of distributions hypothesis: Measuring the daily flow of information, *Journal of Financial and Quantitative Analysis* 29, 101–116.

Schwert, William G., 1990, Stock volatility and the crash of '87, *Review of Financial Studies* 3, 77–102.

Silverman, Bernie W., 1986, *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, London).

Stephan, Jens E., and Robert E. Whaley, 1990, Intraday price change and trading volume relations in the stock and stock option markets, *Journal of Finance* 45, 191–220.