THE JOURNAL OF FINANCE • VOL. XLIII, NO. 2 • JUNE 1988

Estimating the Volatility of Discrete Stock Prices

D. CHINHYUNG CHO and EDWARD W. FREES*

ABSTRACT

This paper introduces an estimator of stock price volatility that eliminates, at least asymptotically, the biases that are caused by the discreteness of observed stock prices. Assuming that the observed stock prices are continuously monitored, an estimator is constructed using the notion of how quickly the price changes rather than how much the price changes. It is shown that this estimator has desirable asymptotic properties, including consistency and asymptotic normality. Also, through a simulation study, the authors show that it outperforms natural estimators for the low- and middle-priced stocks. Furthermore, the simulation study demonstrates that the proposed estimator is robust to certain misspecifications in measuring the time between price changes.

THE VOLATILITY OF STOCK returns plays an important role in many areas of finance. The most common way of estimating the volatility is to calculate the standard deviation of the changes in stock price observed over fixed time intervals. The estimator is often called a "natural" estimator.[1] It is well known that, if the stock price used in the natural estimator is the "true" equilibrium price, then the natural estimator will be nearly unbiased[2] and efficient.

However, the organized exchanges restrict the stock price quotations to be a multiple of some constant, e.g., one eighth of a dollar.[3] Also, the observed stock price is not the true equilibrium price but either a bid price or an ask price. Recently, several researchers noted that these frictions in the market cause the natural estimator to be severely biased. Gottlieb and Kalay [12] demonstrate that the discreteness of the observed stock prices causes the natural estimator to

* School of Business and Department of Statistics, University of Wisconsin-Madison. An earlier version of this paper has been presented at Boston University, Columbia University, Merrill Lynch, Northwestern University, Purdue University, Southern Methodist University, University of California-Riverside, University of Illinois, University of Southern California, University of Wisconsin-Madison, Vanderbilt University, the Actuarial Research Conference in Toronto, the European Finance Association Meetings in Madrid, and the American Finance Association Meetings in Chicago. We would like to thank participants of these workshops. We also would like to thank Lillyn Teh for her computational assistance.

[1] Christie [5] proposes an estimator that relates to the firm's characteristics. Parkinson [20], Garman and Klass [9], and Beckers [2] show that the efficiency of the natural estimator can be improved by using not only the opening and closing prices but also the high and low prices. Latane and Rendleman [16] introduced the so-called implied variance based on the option-pricing formula. However, these estimators are still prone to the biases that we address in this paper.

[2] Even though the sample variance is an unbiased estimator of the true variance, the sample standard deviation is not an unbiased estimator of the true standard deviation. In fact, by using Jensen's inequality, one can show that the sample standard deviation underestimates the true standard deviation.

[3] This is provided that the stock price is greater than one for U.S. domiciled stocks.

overestimate the true volatility of the stock returns.[4] French and Roll [7] argue that the natural estimator overestimates the true volatility of the stock returns in the presence of bid-ask spreads.

The purpose of this study is to construct an estimator that will eliminate the biases, at least asymptotically, that are caused by the discreteness in the observed stock prices. In this study, we do not directly address the biases that are caused by bid-ask spreads. The intuition behind our estimator is simple and appealing. Unlike the natural estimator, we partition the level of the price process into several intervals of a unit length, say, one eighth of a dollar. Then we examine the time required for the prices to move beyond a given unit price interval using the notion of the so-called "first-passage time." In short, the natural estimator focuses on how much the price changes, whereas our "temporal" estimator focuses on how quickly the price changes. Conceptually, the more volatile stock should move more quickly, and, hence, the first-passage time should be shorter than for the less volatile stock.

Based on this notion of the first-passage time, and by assuming that the true stock prices are lognormally distributed and that the observed stock prices are continuously monitored, we are able to construct a consistent estimator along with its asymptotic sampling distribution. This estimator will eliminate the biases that are caused by the discreteness in stock prices, at least asymptotically. Furthermore, since this temporal estimator is based on the time between price changes rather than the magnitude of price changes, the biases caused by bid-ask spreads may be relatively small.

As noted by Niederhoffer and Osborne [18], Roll [22], and Blume and Stambaugh [3], bid-ask spreads produce negative first-order autocovariances in observed price changes. The reason is as follows. Assume that the market is informationally efficient and that the probability distribution of price changes is stationary. Then, according to Roll, in the absence of new information, if the transaction is at the bid price, then the next price change cannot be negative, while, if the transaction is at the ask price, then the next price change cannot be positive. Hence, the distribution of the magnitude of price changes depends on whether the last transaction price is at the bid or ask, and this induces negative serial dependence in the magnitudes of successive price changes and biases the natural estimator.

In other words, the bias in the natural estimator is produced mainly because of its focus on the magnitude of price changes.[5] On the contrary, it is unlikely

[4] Marsh and Rosenfeld [17] claim that there should be no biases in the natural estimator that are caused by the discreteness. They argue that the discreteness problem can be treated as the nontrading problem. The fact that the observed closing price is discrete implies that the true price must have been equal to the observed closing price at some time before the closing. Hence, they argue that, according to Scholes and Williams [23], the bias in the natural estimator should be proportional to the logarithmic mean return, but in practice the daily mean return is very small compared with daily volatility of stock returns. However, our simulation study shows that the natural estimator overestimates the true volatility in most cases.

[5] As Marsh and Rosenfeld [17] and Glosten [10] point out, there are two components of bid-ask spreads; one is due to liquidity trading, and the other is due to information trading. The liquidity-trading component of bid-ask spreads would induce an upward bias in the natural estimator. On the other hand, the information-trading component of bid-ask spreads would induce downward bias

that the bid-ask spreads would have a systematic impact on the time between price changes. For example, there is no strong reason to suspect that, if the transaction is at the bid price, the next price change should occur sooner than if the transaction were at the ask price, or vice versa. Hence, the temporal estimator may be less biased than the natural estimator even in the presence of bid-ask spreads.[6]

The rest of this paper is organized as follows. Section I describes the stock price process assumed in this paper. Section II presents a consistent estimator of the return variance. Section III compares this estimator with natural estimators through a simulation study. The simulation is done using discrete-time measurement intervals rather than continuous time. It turns out that, even in this situation, which is a violation of the continuous-monitoring assumption made in Section I, the temporal estimator is preferable to the natural estimator for the low- and middle-priced stocks. Finally, concluding remarks are made in Section IV.

I. Model Description

We assume that the true stock price is $P(t) = P(0)\exp(\sigma B(t) + \mu t)$,[7] where $t \geq 0$. Here, $P(0)$ is a known constant, μ and σ are unknown parameters, and $\{B(t); t \geq 0\}$ is a standard Brownian motion over $[0, \infty)$. We assume that the observed stock price is

$$\hat{P}(t) = [P(t)/d]d, \tag{1}$$

where $[\cdot]$ is the greatest integer function and d is a known constant.[8] For example, $d = 1/8$ on the New York Stock Exchange. Alternatively, we could use the observed return process

$$\begin{aligned}\hat{R}(t) &= \hat{P}(t)/P(0) \\ &= [\exp(\sigma B(t) + \mu t)/(d/P(0))]d/P(0) \\ &= [\exp(\sigma B(t) + \mu t)/d^*]d^*, \end{aligned} \tag{2}$$

where $d^* = d/P(0)$. Note that the observed return process is discrete in jumps of $d/P(0)$. We will henceforth take $P(0) = 1$ without loss of generality and use $d = d^*$. With this convention, we also have $\hat{P}(t) = \hat{R}(t)$.

The purpose of this paper is to discuss estimation of $\sigma > 0$ from the observed process $\hat{P}(t)$. If one can observe the true pricing process, $P(t)$, estimation of σ is straightforward by using the logarithmic returns. For example, letting Δt be some

according to Marsh and Rosenfeld and no bias according to Glosten. The bias mentioned here is mainly the one due to the liquidity-trading component of bid-ask spreads.

[6] The possibility of systematic impacts of bid-ask spreads on the time between the price changes is left for future research. Recently, Harris [14] introduced a maximum-likelihood estimator that is supposed to take care of discreteness and bid-ask spread problems. See also Ball [1].

[7] This implies that the instantaneous expected rate of return is equal to $\mu + \sigma^2/2$.

[8] Alternatively, we could have assumed the observed stock price to be $\hat{P}(t) = kd$ if $kd - d/2 < P(t) \leq kd + d/2$ or $\hat{P}(t) = kd$ if $(k-1)d < P(t) \leq kd$. However, as pointed out by Gottlieb and Kalay [12], the size of bias would be almost identical and our discussion in this paper would not be affected.

fixed increment of time, say, one day, the logarithmic returns $X_i = \log(P((i + 1)\Delta t)/P(i\Delta t))$ are independent and identically distributed (i.i.d.) normal random variables with mean $\mu\Delta t$ and variance $\sigma^2\Delta t$. However, this is not true when one replaces P with $\hat{P}$ in the definition of X_i. This was pointed out by Gottlieb and Kalay [12], who also examined the magnitude of biases when σ was estimated from the observed pricing process $\hat{P}(i\Delta t)$, $i = 0, 1, 2, \cdots$, for some fixed Δt.

In this paper, we also assume that the discrete stock price, $\hat{P}(t)$, is continuously monitored. This assumption implies that we know when the true price becomes a multiple of d. Hence, there is little measurement error that is caused by the discreteness of observed stock prices if one uses the time rather than the magnitude of the observed price changes; that is what makes our temporal estimators unbiased in the limit.

To make the connection between variability of the process and average time between price changes more precise, we use some ideas of stopping-time random variables and basic probability calculations for a Brownian motion. As an important example of a stopping time, we usually look at random variables of the form $Z = \{\text{first time } t > 0 \text{ such that } B(t) + \theta t \notin (b, a)\}$, where a, b, and θ are constants such that $-\infty < b < 0 < a < \infty$. The random variable Z is the so-called first-passage time during which the process reached b or a for the first time.[9] Distributional properties of Z based on Brownian-motion probability calculations are well known. See Cox and Miller [6, chapter 5].

II. Temporal Volatility Estimator

We now construct an estimator that is based on the first time at which a stock price moves either up or down by a given unit price level. We remark that we can also construct an estimator based on the information provided by only upward movements of stock price. This latter estimator may be useful in situations where stock price changes are generated primarily by bid-ask movements. However, our earlier analysis shows that, in general, the temporal estimator based on both upward and downward movements of stock price outperforms the one using only upward movements. Hence, for brevity, we provide only the temporal estimator that uses both upward and downward movements of stock price.[10]

[9] Some good references are Cox and Miller [6] and Siegmund [24]. Siegmund provides some nice heuristic examples in his applications to sequential analysis. Cox and Miller supply the details of some well-known results that are used implicitly in the paper. In this paper, we use the "first-time" convention in a stopping-time distribution in lieu of the more precise infinum notation, i.e., $Z = \inf\{t > 0: B(t) + \theta t \notin (b, a)\}$.

[10] An estimator based on times to the next price advance can be constructed as follows. Recursively define

$$\tau_n^+ = \{\text{first time } t > 0 \text{ such that } P(t)/P(\tau_{n-1}^+) \geq 1 + d\},$$

where $\tau_0^+ = 0$. With the incremental time to the next price advance $\Delta\tau_i^+ = \tau_i^+ - \tau_{i-1}^+$, define the sample moments $\bar{\tau}_n^+ = n^{-1} \sum_{i=1}^n \Delta\tau_i^+ = \tau_n^+/n$ and $S_\tau^2 = n^{-1} \sum_{i=1}^n (\Delta\tau_i^+)^2 - (\bar{\tau}_n^{+2})$. Cho and Frees [4] developed the alternative temporal estimator $\hat{\sigma}_A^2 = \delta^2 S_\tau^2/(\bar{\tau}_n^+)^3$. Advantages of σ_A^2 are that it can be easily interpreted as the sample variance of stopping times, S_τ^2, times a scaling factor, and its distribution can be obtained in a straightforward fashion. This alternative temporal estimator may

We define a sequence of stopping-time random variables $\{\tau_n\}_{n=1}^{\infty}$ by

$$\tau_n = \{\text{first time } t > \tau_{n-1} \text{ such that } P(t)/P(\tau_{n-1}) \notin ((1+d)^{-1}, (1+d))\}, \tag{3}$$

where $\tau_0 = 0$. This is the first time at which a stock price increases or decreases by the multiple of $(1 + d)$ since the last price change. Because of the strong Markov property of the Brownian-motion process, a useful feature of the stopping-time sequence is that increments of time to change, $\Delta\tau_n = \tau_n - \tau_{n-1}$, are i.i.d. From (3), we can write

$$\tau \equiv \tau_1 = \{\text{first time } t > 0 \text{ such that } \sigma B(t) + \mu t \notin (-\delta, \delta)\}, \tag{4}$$

where $\delta = \log(1 + d)$.

The following lemma provides some important relationships between the parameters of interest, μ and σ, and the expected time between price changes, $E\tau$.

LEMMA 1: *Consider the stopping time defined in* (4) *and assume that* $P(0) = 1$. *If* $\mu = 0$, *then the probability of the true price reaching the multiple of* $1 + d$ *before* $(1 + d)^{-1}$ *is*

$$\Pr(P(\tau) = 1 + d) \equiv p = \tfrac{1}{2}, \tag{5}$$

and the expected first time at which the true price reaches the multiple of $1 + d$ *or* $(1 + d)^{-1}$ *is*

$$E\tau = \sigma^{-2}\delta^2. \tag{6}$$

More generally, if $\mu \neq 0$, *then*

$$\Pr(P(\tau) = 1 + d) \equiv p = \{1 + (1 + d)^{-2\mu/\sigma^2}\}^{-1} \tag{7}$$

and

$$\begin{aligned} E\tau &= \mu^{-1}\delta\{2(1 + (1 + d)^{-2\mu/\sigma^2})^{-1} - 1\} \\ &= \mu^{-1}\delta(2p - 1). \end{aligned} \tag{8}$$

The proof of Lemma 1 is an immediate application of, for example, Theorem 3.6 of Siegmund [24, p. 36]. Note that, for a given $\mu > 0$, p increases and $E\tau$ decreases as σ increases and vice versa. That is, the more-volatile stock is more likely to increase and take less time for the price to change than the less-volatile stock. Also, it can be easily shown that (7) and (8) approach (5) and (6), respectively, as μ approaches zero.

Lemma 1 immediately suggests several estimators of σ. For example, by the strong law of large numbers, we have that, with probability one,

$$\lim_{n\to\infty} \bar{\tau}_n = E\tau,$$

where $\bar{\tau}_n = \tau_n/n$. Thus, in the special case of $\mu = 0$, $\delta^2/\bar{\tau}_n$ is a consistant estimator

be useful in periods of advancing periods for stocks with large bid-ask spreads. Heuristically, by measuring times between advances, we essentially measure times between ask prices. By restricting ourselves to ask prices, we increase the homogeneity of our observations and produce more reliable estimates.

of σ^2; i.e., with probability one,

$$\lim_{n\to\infty}\delta^2/\bar{\tau}_n=\sigma^2.$$

Similarly, from (8), we have

$$\sigma^2=2\mu\delta/(\log(\delta+\mu\mathrm{E}\tau)-\log(\delta-\mu\mathrm{E}\tau)). \tag{9}$$

Hence, we define the temporal estimator by

$$\hat{\sigma}^2=2\hat{\mu}\delta/(\log(\delta+\hat{\mu}\bar{\tau}_n)-\log(\delta-\hat{\mu}\bar{\tau}_n)), \tag{10}$$

where $\hat{\mu}=\bar{\tau}_n^{-1}\log(P(\tau_n))$ and $\bar{\tau}_n=\tau_n/n$. Note that $\hat{\mu}$ is the sample estimate of the price drift and $\bar{\tau}_n$ is the sample average of first-passage times.

We have the following asymptotic properties, the proofs of which are given in Cho and Frees [4].

LEMMA 2: *Consider the sequence* $\{\tau_n\}$ *defined in* (3). *Then,* $\hat{\mu}$ *is a consistent estimator of* μ.

THEOREM 1: *Assume that* $\mu\neq 0$. *Then,* $\hat{\sigma}$ *is a consistent estimator of* σ. *Furthermore,*

$$n^{1/2}(\hat{\sigma}-\sigma)\rightarrow_D N(0,\Sigma^2), \tag{11}$$

where $\Sigma^2=\sigma^4(1+d)^{4\mu/\sigma^2}\operatorname{var}\tau/\{64\delta^4p^4\}$ *and*

$$\operatorname{var}\tau=\frac{\sigma^2}{\mu^3}\delta(2p-1)-\frac{4\delta^2}{\mu^2}\left(\frac{1}{A+B}\right)^2, \tag{12}$$

where $A=\exp(-\mu\delta\sigma^2)$ *and* $B=\exp(\mu\delta/\sigma^2)$.

III. Comparison of Natural versus Temporal Volatility Estimators

In order to compare the temporal estimator with the natural estimator, we now examine the performances of the two estimators using simulated data. Since we do not know the "true" underlying parameters of stock price processes, we cannot compare estimators using the observed stock data. In carrying out the simulation study, we deliberately induce a certain model-misspecification error. The error we consider here is to assume that the price process is observed only at discrete time intervals rather than continuously.[11] We feel that this type of "error" is important to consider explicitly since it typically will be encountered by even the most careful stock analysts.

We use the following recursive relationship to generate the true prices at

[11] As suggested by an anonymous referee, an alternative viewpoint is to think of the actual price as observed only when a trade occurs. This could be modeled by assuming that the inter-arrival times between trades are distributed according to some exogenous stochastic process not necessarily related to the pricing process. However, by assuming that the parent geometric Brownian-motion process governs price, we implicitly assume that times to trades are i.i.d. random variables. It may be important to investigate whether the data follow this model or, if not, fit more general models than the ones introduced in this paper.

discrete time intervals:

$$P(t) = P(t-1)\exp(\sigma B(\Delta t) + \mu \Delta t), \quad t = 1, \cdots, T. \tag{13}$$

We simulate 1000 price series, each of which contains T = 1000 observations. We use the initial prices $P(0)$ of \$1, \$25, and \$100 to represent the low, intermediate and high stock prices, respectively. The daily mean returns μ of 0.03 percent, 0.08 percent, and 0.13 percent and the daily standard deviations σ of 0.8 percent, 1.2 percent, and 1.6 percent are used. These numbers correspond to the annual returns of 7.5 percent, 20.0 percent, and 32.5 percent and the annual standard deviations of 12.6 percent, 19.0 percent, and 25.3 percent, respectively, if 250 trading days are assumed in a year. We use d = ⅛ as our discretizing constant. Hence, after standardization, the constant is d^* = $0.125/P(0)$. Different values of Δt are used depending on the initial prices $P(0)$. The expected time to price change, $E\tau$, depends on the parameters $P(0)$, μ, and σ. For example, in the case of μ = 0.03 percent and σ = 0.8 percent, the \$1 stock takes more than twenty times longer to change by the amount of one eighth of a dollar than the \$25 stock. Hence, for the \$1 stock, we do not have to generate as many prices as for the \$25 stock if we want to observe a similar number of price changes. Thus, in order to reduce the computational costs, we let Δt be one day for the \$1 stock, 0.02 days for the \$25 stock, and 0.002 days for the \$100 stock.

These true prices are then discretized according to (1) to generate observed prices, and these observed prices are used to calculate various estimators. Results are summarized in Tables I through III for the \$1, \$25, and \$100 stocks, respectively. Each of these tables contains the average point estimates of $\hat{\mu}$ and $\hat{\sigma}$, along with their root mean squares, for three different types of estimators. These estimators are the natural estimators based on the true and observed prices, respectively, and the temporal estimator (denoted as *True*, *Nat*, and *Temp*, respectively). Tables II and III contain additional natural estimators based on

Table I

Comparison of Estimators on Simulated Data: \$1 (Days = 1000, Δt = 1)[a]

		σ = 0.8%				σ = 1.2%				σ = 1.6%			
	μ(%)	$\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$	$\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$	$\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$
True	0.03	0.030	0.025	0.800	0.018	0.030	0.037	1.200	0.026	0.031	0.050	1.600	0.035
Nat	0.03	0.025	0.027	2.840	2.080†	0.025	0.040	3.470	2.330†	0.025	0.053	4.050	2.550†
Temp	0.03	0.038	0.038	0.805	0.270*	0.038	0.050	1.180	0.276*	0.036	0.059	1.500	0.355*
True	0.08	0.080	0.025	0.800	0.018	0.080	0.037	1.200	0.026	0.081	0.050	1.600	0.035
Nat	0.08	0.077	0.026	2.520	1.750†	0.077	0.039	3.070	1.910†	0.077	0.052	3.570	2.040†
Temp	0.08	0.088	0.028	1.140	0.421*	0.092	0.044	1.410	0.365*	0.091	0.057	1.740	0.400*
True	0.13	0.130	0.025	0.800	0.018	0.130	0.037	1.200	0.026	0.131	0.050	1.600	0.035
Nat	0.13	0.129	0.025	2.270	1.490†	0.129	0.038	2.770	1.610†	0.129	0.051	3.200	1.660†
Temp	0.13	0.135	0.026	1.660	0.914*	0.140	0.040	1.800	0.729*	0.142	0.054	2.080	0.705*

[a] *True*, natural estimator based on true prices; *Nat*, natural estimator based on observed prices; *Temp*, temporal estimator; μ, σ, (daily) mean and standard deviations of true price process (%); $\hat{\mu}$, $\hat{\sigma}$, estimates of mean and standard deviations (%)—average of 1000 simulation results; *RMS*, root mean squares of estimates (%).

† Largest root mean square for a given μ and σ except "*True*".

* Smallest root mean square for a given μ and σ except "*True*".

Table II

Comparison of Estimators on Simulated Data: $25 (Days = 20, $\Delta t = 0.02$)[a]

	μ (%)	$\sigma = 0.8\%$ $\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$	$\sigma = 1.2\%$ $\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$	$\sigma = 1.6\%$ $\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$
True	0.03	0.032	0.176	0.800	0.018	0.033	0.264	1.200	0.026	0.034	0.352	1.600	0.035
Nat 50	0.03	0.020	0.177	1.500	0.707†	0.021	0.264	1.840	0.646†	0.022	0.353	2.150	0.556†
Nat 25	0.03	0.020	0.177	1.270	0.470	0.021	0.264	1.580	0.381	0.022	0.353	1.900	0.308
Nat 5	0.03	0.020	0.177	0.923	0.141	0.021	0.264	1.290	0.128*	0.022	0.353	1.670	0.141*
Nat 1	0.03	0.020	0.177	0.820	0.134	0.021	0.264	1.210	0.187	0.022	0.353	1.600	0.251
Temp	0.03	0.032	0.180	0.712	0.101*	0.033	0.267	1.020	0.189	0.035	0.355	1.300	0.310
True	0.08	0.082	0.176	0.800	0.018	0.083	0.264	1.200	0.026	0.084	0.352	1.600	0.035
Nat 50	0.08	0.070	0.177	1.500	0.704†	0.071	0.265	1.840	0.642†	0.072	0.353	2.150	0.551†
Nat 25	0.08	0.070	0.177	1.260	0.467	0.071	0.265	1.570	0.378	0.072	0.353	1.900	0.308
Nat 5	0.08	0.070	0.177	0.922	0.141	0.071	0.265	1.280	0.126*	0.072	0.353	1.670	0.140*
Nat 1	0.08	0.070	0.177	0.820	0.128	0.071	0.265	1.200	0.189	0.072	0.353	1.600	0.251
Temp	0.08	0.083	0.180	0.716	0.098*	0.084	0.266	1.020	0.186	0.085	0.354	1.300	0.304
True	0.13	0.132	0.176	0.800	0.018	0.133	0.264	1.200	0.026	0.134	0.352	1.600	0.035
Nat 50	0.13	0.120	0.177	1.500	0.701†	0.121	0.265	1.840	0.639†	0.122	0.353	2.140	0.548†
Nat 25	0.13	0.120	0.177	1.260	0.464	0.121	0.265	1.570	0.375	0.122	0.353	1.900	0.307
Nat 5	0.13	0.120	0.177	0.922	0.139	0.121	0.265	1.290	0.126*	0.122	0.353	1.670	0.139*
Nat 1	0.13	0.120	0.177	0.820	0.129	0.121	0.265	1.200	0.190	0.122	0.353	1.600	0.252
Temp	0.13	0.134	0.179	0.721	0.092*	0.134	0.267	1.030	0.183	0.134	0.354	1.310	0.299

[a] *True*, natural estimator based on true prices; *Nat n*, natural estimator based on observed prices using n equally spaced observations per day; *Temp*, temporal estimator; μ, σ, (daily) mean and standard deviations of true price process (%); $\hat{\mu}$, $\hat{\sigma}$, estimates of mean and standard deviations (%)—average of 1000 simulation results; *RMS*, root mean squares of estimates (%).

† Largest root mean square for a given μ and σ except "*True*".

* Smallest root mean square for a given μ and σ except "*True*".

Table III

Comparison of Estimators on Simulated Data: $100 (Days = 2, $\Delta t = 0.002$)[a]

		$\sigma = 0.8\%$				$\sigma = 1.2\%$				$\sigma = 1.6\%$			
	μ (%)	$\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$	$\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$	$\hat{\mu}$	$RMS_{\hat{\mu}}$	$\hat{\sigma}$	$RMS_{\hat{\sigma}}$
True	0.03	0.020	0.581	0.800	0.018	0.015	0.872	1.200	0.027	0.010	1.160	1.600	0.035
Nat 500	0.03	−0.012	0.582	1.340	0.538†	−0.015	0.874	1.650	0.455	−0.021	1.160	1.970	0.368
Nat 50	0.03	−0.012	0.582	0.877	0.100	−0.015	0.874	1.250	0.104*	−0.021	1.160	1.640	0.125*
Nat 25	0.03	−0.012	0.582	0.837	0.092*	−0.015	0.874	1.220	0.124	−0.021	1.160	1.610	0.165
Nat 5	0.03	−0.012	0.582	0.787	0.190	−0.015	0.874	1.180	0.286	−0.021	1.160	1.560	0.376
Nat 1	0.03	−0.012	0.582	0.630	0.526	−0.015	0.874	0.946	0.787†	−0.021	1.160	1.260	1.050†
Temp	0.03	0.019	0.589	0.691	0.115	0.015	0.879	0.979	0.225	0.010	1.170	1.240	0.365
True	0.08	0.070	0.581	0.800	0.018	0.065	0.872	1.200	0.027	0.060	1.160	1.600	0.035
Nat 500	0.08	0.039	0.584	1.340	0.537†	0.034	0.873	1.650	0.453	0.030	1.170	1.970	0.369
Nat 50	0.08	0.039	0.584	0.877	0.100	0.034	0.873	1.250	0.106*	0.030	1.170	1.640	0.125*
Nat 25	0.08	0.039	0.584	0.838	0.094*	0.034	0.873	1.220	0.127	0.030	1.170	1.610	0.164
Nat 5	0.08	0.039	0.584	0.787	0.190	0.034	0.873	1.170	0.282	0.030	1.170	1.570	0.378
Nat 1	0.08	0.039	0.584	0.634	0.524	0.034	0.873	0.946	0.785†	0.030	1.170	1.260	1.050†
Temp	0.08	0.072	0.588	0.692	0.115	0.066	0.877	0.980	0.224	0.060	1.170	1.240	0.364
True	0.13	0.120	0.581	0.800	0.018	0.115	0.872	1.200	0.027	0.110	1.160	1.600	0.035
Nat 500	0.13	0.089	0.584	1.340	0.540†	0.084	0.873	1.650	0.451	0.079	1.170	1.970	0.369
Nat 500	0.13	0.089	0.584	0.880	0.103	0.084	0.873	1.250	0.103*	0.079	1.170	1.640	0.125*
Nat 25	0.13	0.089	0.584	0.837	0.093*	0.084	0.873	1.220	0.126	0.079	1.170	1.610	0.163
Nat 5	0.13	0.089	0.584	0.790	0.188	0.084	0.873	1.180	0.283	0.079	1.170	1.570	0.375
Nat 1	0.13	0.089	0.584	0.628	0.525	0.084	0.873	0.947	0.785†	0.079	1.170	1.260	1.050†
Temp	0.13	0.122	0.589	0.693	0.114	0.115	0.879	0.978	0.226	0.110	1.170	1.240	0.363

[a] *True*, natural estimator based on true prices; *Nat n*, natural estimator based on observed prices using *n* equally spaced observations per day; *Temp*, temporal estimator; μ, σ, (daily) mean and standard deviations of true price process (%); *RMS*, root mean squares of estimates (%).

† Largest root mean square for a given μ and σ except "*True*".

* Smallest root mean square for a given μ and σ except "*True*".

observed prices that use a different number of observations per day.[12] As noted by Gottlieb and Kalay [12], with observed prices the natural estimator should get worse as one uses more observations per day, unlike the case of true prices. This is so because the bias due to discreteness should pile up with more observations of observed prices. Hence, it is of interest to examine how the performance of natural estimators on the observed prices behaves as a function of the number of observations.

In estimating the stopping times for the temporal estimator, we have to make some adjustments to (3). The stopping times defined in (3) are based on the true prices that cannot be observed. Note that (3) can be expressed as

$$\tau_n = \left\{ \text{first time } t > \tau_{n-1} \text{ such that } P(t) - P(\tau_{n-1}) \notin \left(-\frac{d}{1+d} P(\tau_{n-1}), dP(\tau_{n-1}) \right) \right\}.$$

Since the differences between two observed prices must be multiples of d, we must adjust the price interval in the above equation. Also, recall that we assume the observed price to be the largest multiple of d that is less than or equal to the true price. That is, $\hat{P}(t) = kd$ if $kd \leq P(t) < (k+1)d$ for some integer k. This implies that, if the true price has just crossed kd in an upward direction, then the observed price at the moment after is kd, whereas, if the true price crossed kd in a downward direction, then the observed price at the moment after would be $(k-1)d$.

Hence, we use the following approximation of the stopping time τ_n:

$$\hat{\tau}_n = \left\{ \text{first time } t > \hat{\tau}_{n-1} \text{ such that} \right.$$
$$\left. \hat{P}(t) - \hat{P}(\hat{\tau}_{n-1}) \notin \left(-2\frac{d}{1+d} IP(\hat{\tau}_{n-1}), dIP(\hat{\tau}_{n-1}) \right) \right\},$$

with $\hat{\tau}_0 = 0$ and $\hat{P}(\hat{\tau}_n) = \hat{P}(\hat{\tau}_{n-1}) + dIP(\hat{\tau}_{n-1})$ or $\hat{P}(\hat{\tau}_n) = \hat{P}(\hat{\tau}_{n-1}) - dIP(\hat{\tau}_{n-1})$, depending on whether the price moves up or down, respectively.[13] $IP(\cdot)$ is the nearest integer of $\hat{P}(\cdot)$. The intuition is as follows. At time $\hat{\tau}_{n-1}$, $\hat{P}(\hat{\tau}_{n-1})$ is a multiple of d for each n. As before, if the true price goes up, then the first-passage time $\hat{\tau}_n$ (after $\hat{\tau}_{n-1}$) satisfies $\hat{P}(\hat{\tau}_n) = \hat{P}(\hat{\tau}_{n-1}) + dIP(\hat{\tau}_{n-1})$. However, if the true price goes down, then the first-passage time $\hat{\tau}_n$ must satisfy $\hat{P}(\hat{\tau}_n + h) = \hat{P}(\hat{\tau}_{n-1}) - 2\frac{d}{1+d} IP(\hat{\tau}_{n-1})$ for small h. This is due to the discretization process assumed in Section I. See Figure 1 for a sequence of stopping times.[14]

[12] We could have compared the high-low estimator of Garman and Klass [9], which has recently gained some popularity due to its high relative efficiency compared with the natural estimator using only closing prices. We did not do this because the high-low estimator assumes μ to be zero.

[13] Special considerations are necessary if the true price descends to a multiple of $d/(1+d)$ but does not drop below the multiple. However, this happens with probability zero and can be ignored. Also, note that $d/(1+d) \simeq d$.

[14] Since the observed process is discrete in multiples of d, we use an approximation to the stopping time by substituting the integer part of $\hat{P}(\hat{\tau}_{n-1})$ for $P(\tau_{n-1})$. This is effectively a Taylor-series linearization due to the exponential scaling of the process. Indeed, in the random-walk case without this scaling, no approximation is necessary and the observed stopping times can be measured without

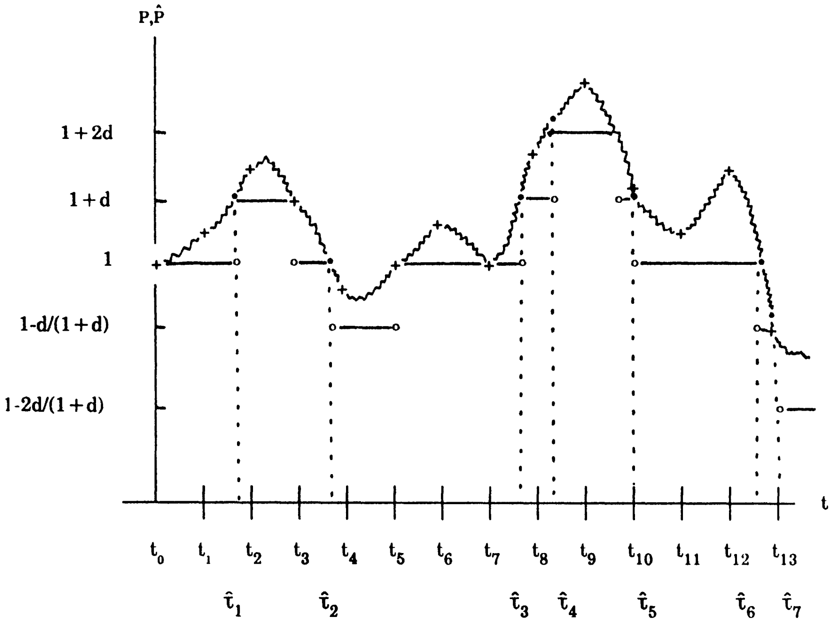

Figure 1. Stopping times of stock price changes. t_i, time at which price is generated; $\hat{\tau}_i$, estimated stopping time of price changes; ∿, true prices; ——, observed prices; +, point at which true price is generated; •, point at which true price changes to the multiple of $(1 + d)$ or $(1 + d)^{-1}$; ∘, point of discontinuity.

In comparing estimators, we examine the root mean square of the point estimate. We use the root mean square in the usual sense, which equals the square root of the variance plus bias squared, to simultaneously discuss the statistic's variability and deviation from the parameter of interest. Even though the temporal estimator is a consistent estimator, there will be some biases mainly due to errors in measuring times between price changes in finite samples. Hence, it is proper to compare the estimators using the root mean squares rather than the standard errors. We note that the comparison of estimators reduces to comparing the biases due to measurement errors in the magnitude of price changes with the biases due to measurement errors in the time between price changes.

In Table I, the $1 case, we see that the natural estimator using true prices always gives the most desirable estimate. On the other hand, the natural estimator on the observed prices performs the worst. This is not surprising because the impact of discreteness (one eighth of a dollar) is significant for low-priced stocks. Also, we note that the natural estimator using observed prices performs better as μ increases. This is due to the fact that there will be less measurement error from the discreteness when the level of price increases. With the temporal estimator, the performance gets worse as μ increases. This may be due to the fact

error. Heuristically, the temporal estimator is robust to this measurement error since it relies only on the sum of the stopping times rather than on precise measurement of each one.

that there will be more measurement errors in the time between price changes by not being able to detect the time properly for downward movements of stock price. On the other hand, there is no consistent pattern with regard to σ for the temporal estimator.

Table II, the \$25 case, also shows that the natural estimator based on the true prices gives the most desirable estimates, as expected, and that the natural estimator based on the entire set of observed prices (fifty observations per day) gives the worst estimates. Note that, as the number of observations per day decreases, the performance of the natural estimator based on the observed prices actually improves. In fact, the natural estimator based on five observations per day performs the best among all estimators except the natural estimator based on the true prices for $\sigma = 1.2$ percent and $\sigma = 1.6$ percent. For $\sigma = 0.8$ percent, the temporal estimator performs the best.

This seems a little surprising. However, if one allows the natural estimator to have different time intervals, then there can be some natural estimator that can dominate our temporal estimator. This is so because, in this case, the natural estimator is in fact using the information provided not only by the magnitude of price changes but also by the time between price changes. Unfortunately, one does not know a priori how many observations are optimal to use. Also, if we generate more observations per day by reducing Δt so that we have less errors in estimating the time between price changes, then the temporal estimator will improve.[15] Note that, from Figure 1, if we use coarse partitions of times in generating prices, then we will not be able to continuously monotor the time at which the observed price changes. This induces additional errors in estimating the time between price changes. Thus, the results of the simulation study are biased *against* the temporal estimator in that the simulation study uses observations taken at small time increments Δt, whereas we have been assuming continuous monitoring of all the data. In practice, continuous monitoring may not be available, but these results indicate that the temporal estimator will perform well in situations approximating continuous monitoring.

Table III, the \$100 case, again shows that the natural estimator based on the true prices gives the most desirable estimates, the natural estimator based on the entire set of observed prices (500 observations per day) performs the worst for $\sigma = 0.8$ percent, and the natural estimator based on one observation performs the worst for $\sigma = 1.2$ percent and $\sigma = 1.6$ percent. On the other hand, excluding the natural estimator based on the true prices, the natural estimator based on twenty-five observations performs the best for $\sigma = 0.8$ percent, and the natural estimator based on fifty observations performs the best for $\sigma = 1.2$ percent and $\sigma = 1.6$ percent. This may be explained by the fact that the impact of discreteness is minor for high-priced stocks, and so there will be smaller biases in the natural estimators. (Recall that $d = 0.125/P(0)$.)

[15] In order to see this, we generated prices using $\Delta t = 0.002$ days for the \$25 stock with $\mu = 0.03$ percent and $\sigma = 0.12$ percent. The root mean square of the temporal estimator was 0.087. The root mean squares of natural estimators were 2.080, 0.647, 0.380, 0.127, and 0.200 for the daily observations of 500, fifty, twenty-five, five, and one, respectively. Thus, in this case, the temporal estimator performed the best and the natural estimator based on 500 observations performed the worst.

Overall, our simulation study indicates that the temporal estimator is the preferred estimator for the case of low- and middle-priced stocks. The natural estimator with a moderate number of observations per day is the preferred estimator for the case of high-priced stocks. Also, it demonstrates that, contrary to common belief, the natural estimator does not become better as we add more observations per day as long as the estimator is based on the observed discrete stock prices.

Finally, we compare estimators by examining the actual transactions data of Advanced Micro Devices (AMD) stock for each of three days, December 3 through 5, 1984. A graphical representation of this data set can be found in Figure 2. The calculations are summarized in Table IV. The temporal estimator agrees with the number of transactions and price changes and indicates that AMD was more volatile on December 5 than on December 3. On the other hand, the natural

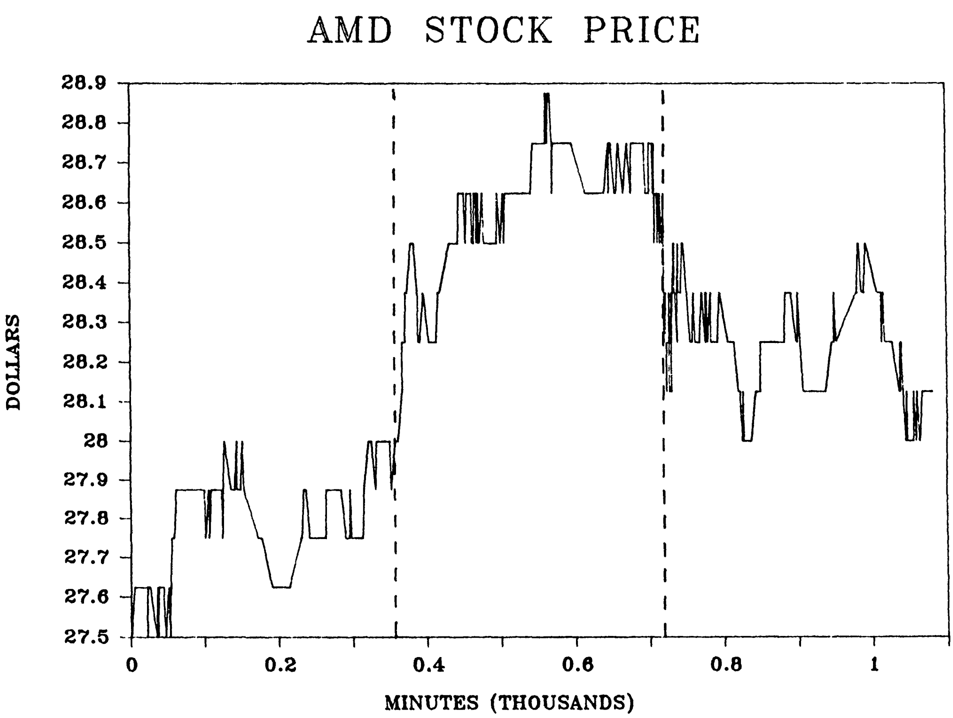

Figure 2. Transaction Stock Price of AMD (12/3/84–12/5/84).

Table IV

Comparison of Estimators on AMD Data[a]

Date	*NT*	*NP*	*Nat*	*Temp*
12/3/84	102	40	1.506	1.300
12/4/84	145	57	1.761	1.451
12/5/84	159	66	1.447	2.133

[a] *NT*, number of transactions; *NP*, number of price changes; *Nat*, natural estimator based on observed hourly prices (%); *Temp*, temporal estimator (%).

estimator based on hourly observations indicates that the stock was more volatile on December 3 than on December 5. Naturally, this is a real data set, and no general conclusions can be ascertained by studying it. However, a visual inspection of Figure 2 leads us to conclude that the temporal estimator is in greater concordance with the data than the natural estimator. To the extent that the volatility should reflect trading activity, the temporal estimator captures more of this activity than the natural estimator. By fixing observation times, the natural estimator fails to reflect trading activity between the observation times.

IV. Conclusion

In this paper, we addressed the issue of constructing a volatility estimator that eliminates biases that are caused by the discreteness of observed stock prices. Assuming that the stock prices are lognormally distributed and that the observed stock prices are continuously monitored, we derived a consistent estimator using the notion of first-passage-time random variables. Also, we identified the asymptotic sampling distribution of this estimator that would enable us to carry out statistical inferences.

The major reason why the temporal estimator is unbiased (in the limit) is that, with the assumption of continuous monitoring, we can accurately calculate the time during which the true price process moves by the amount of the discretization constant. On the other hand, the natural estimator is biased because the magnitude of price changes for a fixed time interval cannot be calculated without an error. In practice, the assumption of continuous monitoring may not be reasonable, in which case we can measure the time between price changes only with error. However, our simulation study shows that measurement errors in the time between price changes induce less bias than measurement errors in the magnitude of price changes.

We also examined the relative performance of the temporal estimator compared with natural estimators using a simulation study. Our results show that the temporal estimator outperforms natural estimators for the low- and middle-priced stocks. Finally, we demonstrated that, contrary to common belief, the natural estimator does not become better as one adds more observations per day.

We remark that the temporal estimator works even in a situation where there is no discreteness problem. In this case, we could partition the level of price process by any amount d. As d gets smaller, the estimator should get better. In fact, in the limit, the estimator approaches the true volatility with probability one.[16]

[16] While the temporal estimator is intuitively appealing and has desirable properties, we conjecture that it is not optimal in some sense. For example, by examining Figure 1, we see that the estimator does not use the information contained in each observed price movement but only in changes wherein the underlying true price has moved by a unit d. In other words, by the sample function properties of the Brownian-motion path, $P(t)$ will cross 1 infinitely many times during the interval $[0, \Delta t)$, where Δt is arbitrarily small. Hence, during this time interval, $\hat{P}(t)$ will oscillate between 1 and $1 - d$ $(1 + d)$

Recently, transactions data have become more readily available and have been frequently used in many studies.[17] Hence, the temporal estimator is potentially beneficial and should be used in the event-type studies or the option-related studies. Especially for the low- and middle-priced stocks, one should seriously consider using the temporal estimator. Currently, we are investigating the performance of the temporal estimator in relation to the implied volatilities of stocks that have options. We hope to report this in a later paper using transactions data.

---

infinitely often, and it may be possible to construct an estimator of σ based on this oscillation. However, if we defined the observed process to be

$$\hat{P}(t) = P(\tau_{n-1}), \quad \tau_{n-1} \leq t < \tau_n,$$

instead of (1), then the temporal estimator would use the information in each jump. We intend to explore this application in a later paper.

[17] For examples, there are studies by Niederhoffer and Osborne [18], Garbade and Lieber [8], Oldfield and Rogalski [19], Wood, McInish, and Ord [25], Patell and Wolfson [21], Harris [13, 14], Hasbrouck and Ho [15], and Glosten and Harris [11], to name a few.

REFERENCES

1. C. Ball. "Security Price Estimation Bias Induced by Discrete Observations." Working paper, University of Michigan, 1987.
2. S. Beckers. "Variance of Security Price Returns Based on High, Low and Closing Prices." *Journal of Business* 56 (January 1983), 97–112.
3. M. Blume and S. Stambaugh. "Biases in Computed Returns." *Journal of Financial Economics* 12 (November 1983), 387–404.
4. D. Cho and E. Frees. "Estimating the Volatility of Discrete Stock Prices." Working paper, University of Wisconsin-Madison, 1986.
5. A. Christie. "The Stochastic Behavior of Common Stock Variances." *Journal of Financial Economics* 10 (December 1982), 407–32.
6. D. Cox and H. Miller. *The Theory of Stochastic Processes*. London: Methuen, 1965.
7. K. French and R. Roll. "Stock Return Variances: The Arrival of Information and the Reaction of Traders." *Journal of Financial Economics* 17 (September 1986), 5–26.
8. K. Garbade and Z. Lieber. "On the Independence of Transactions on the New York Stock Exchange." *Journal of Banking and Finance* 1 (October 1977), 151–72.
9. M. Garman and M. Klass. "On the Estimation of Security Price Volatilities from Historical Data." *Journal of Business* 53 (January 1980), 67–78.
10. L. Glosten. "Components of the Bid-Ask Spread and the Statistical Properties of Transaction Prices." Working paper, Northwestern University, 1985.
11. ——— and L. Harris. "Estimating the Components of the Bid/Ask Spread." Working paper, University of Southern California, 1986.
12. G. Gottlieb and A. Kalay. "Implications of the Discreteness of Observed Stock Prices." *Journal of Finance* 40 (March 1985), 135–53.
13. L. Harris. "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns." *Journal of Financial Economics* 16 (May 1986), 99–117.
14. ———. "Estimation of 'True' Stock Price Variances and Bid-Ask Spreads from Discrete Observations." Working paper, University of Southern California, 1986.
15. J. Hasbrouck and T. Ho. "Intraday Stock Returns: Empirical Evidence of Lagged Adjustment." Working paper, New York University, 1986.
16. H. Latane and R. Rendleman. "Standard Deviation of Stock Price Return Implied in Option Prices." *Journal of Finance* 31 (May 1976), 369–81.
17. T. Marsh and E. Rosenfeld. "Non-Trading, Market Making, and Estimates of Stock Price Volatility." *Journal of Financial Economics* 15 (March 1986), 359–72.

18. V. Niederhoffer and M. F. M. Osborne. "Market Making and Reversal on the Stock Exchange." *Journal of the American Statistical Association* 61 (December 1966), 897–916.
19. G. Oldfield, Jr. and R. Rogalski. "A Theory of Common Stock Returns over Trading and Non-Trading Periods." *Journal of Finance* 35 (June 1980), 729–51.
20. M. Parkinson. "The Extreme Value Method for Estimating the Variance of the Rate of Return." *Journal of Business* 53 (January 1980), 61–65.
21. J. Patell and M. Wolfson. "The Intraday Speed of Adjustment of Stock Prices to Earnings and Dividend Announcements." *Journal of Financial Economics* 13 (June 1984), 223–52.
22. R. Roll. "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market." *Journal of Finance* 39 (September 1984), 1127–39.
23. M. Scholes and J. Williams. "Estimating Betas from Nonsynchronous Data." *Journal of Financial Economics* 5 (December 1977), 309–27.
24. D. Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. New York: Springer-Verlag, 1985.
25. R. Wood, T. McInish, and J. Ord. "An Investigation of Transactions Data for NYSE Stocks." *Journal of Finance* 40 (July 1985), 723–39.