Assessing the Quality of a Security Market: A New Approach to Transaction-Cost Measurement

Joel Hasbrouck
New York University

I discuss a new method for measuring the deviations between actual transaction prices and implicit efficients prices. The approach decomposes security transaction prices into random-walk and stationary components. The random-walk component may be identified with the efficient price. The stationary component, the difference between the e-price and the actual transaction price, is termed the pricing error. Its dispersion is a natural measure of market quality. I describe practical strategies for estimating these quantities. For a sample of NYSE stocks, the average pricing error standard deviation estimate is roughly 0.33 percent of the stock price. If the pricing error is normally distributed and if it is always a positive cost incurred by the transaction initiators, the correspending average transaction cost for these traders is 0.26 percent of the stock price. The dispersion of the pricing error is also found to be elevated at the beginning and end of the trading session,

For comments and suggestions, 1 am indebted to the editors (Michael Gibbons and Chester Spatt), the two referees (James Stock and Lawrence Harris), Yakov Amihud, Fischer Black, Richard Green, Maureen O'Hara, Robert Schwartz, and workshop participants at Berkeley, Dartmouth, MIT. New York University, Rutgers University. Stanford University, and the Securities and Exchange Commission. I am especially grateful to John Campbell for bringing the techniques of random-walk decomposition to my attention. All errors are my own responsibility. I am indebted to the New York Stock Exchange and the Institute for Quantitative Research in Finance for financial support. The comments and opinions contained in this article are those of the author only. In particular, the views expressed here do not necessarily reflect those of the directors, members, or officers of the New York Stock Exchange, Inc.
Address correspondence to Joel Hasbrouck. New York University, Stem School of Business, 44 West Fourth Street, New York, NY 10012.

The Review of *Financial Studies* 1993 Volume 6, number 1, pp. 191-212

It is commonly held that in addition to the explicit costs of executing a security transaction (such as the commission), the trader incurs an implicit cost in the difference between the actual transaction price and a benchmark price that is considered to be in some sense fair or efficient. Measurement of this difference arises in financial markets analyses in two related contexts. First, transaction-cost measurement traditionally aims at estimation of this difference for a buyer or seller in a specific transaction, usually with the purpose of evaluating the broker's performance. Second, comparative market analysis seeks average transaction-cost measures, with a view toward determining markets or regulatory structures under which these costs are minimized. In this article, I discuss and implement a class of improved measures that are useful for both purposes.

The new measures are based on existing techniques for resolving a nonstationary time series into a random-walk component and a residual stationary component.[1] Applied to security transaction prices, it is natural to identify the random-walk component as the efficient price. The residual stationary component, termed here the *pricing* error, is taken as the implicit transaction cost. The dispersion of the pricing error measures how closely actual transaction prices track a random walk and is therefore a natural measure of market quality.[2]

Earlier studies of transaction costs have used various benchmarks for comparison with observable transaction prices: the volume-weighted daily average price [Berkowitz, Logue, and Noser (1988)], daily high-low midpoint prices and closing prices [Beebower (1989)], and the midpoint of the prevailing bid and ask quotes [Perold (1988)]. These benchmarks often possess the virtue of convenience, but generally lack strong interpretations as efficient prices.

Comparative studies of market performance and regulatory impact analyses usually are based on posted bid-ask spreads, liquidity ratios, or return variance ratios.[3] Grossman and Miller (1988) summarize the problems with using posted bid-ask spreads and liquidity ratios as measures of market quality. The posted bid-ask spread is twice the transaction cost for a market-order trader under numerous restrictive

[1] See Beveridge and Nelson (1981), Watson (1986). Campbell and Mankiw (1987). and Quah (1990. 1992). Stock and Watson (1988) provide a highly readable review.

[2] Hasbrouck (1991b) uses this approach to characterize the random-walk component implicit in the quotes. As a measure of information asymmetries, in my earlier article I describe a computation of the random walk variance that is explained by trades, In contrast, in the present article 1 focus on actual transaction prices, and the characterization of the stationary component. The common use of random-walk decompositions leads to methodological similarities, while the divergent aims of the two articles lead to substantial differences.

[3] Cooper, Groth, and Avera (1985), Hasbrouck and Schwartz (1988), and Tanner and Pritchett (1992). compare the liquidity in over-the-counter and exchange markets for equities. Neal (1989) contrasts options trading on the CBOE and AMEX. For an example of regulatory studies, see the Commodities Futures Trading Commission (1989) draft report on dual trading.

assumptions, including absence of asymmetric information. Liquidity ratios, which relate price changes to trading volume, fail to differentiate between transient price impacts (a sign of illiquidity) and the persistent impacts (a consequence of information inferred from trades in an efficient market). The variance ratio is defined as the return variance per unit time measured over long horizons divided by that measured over short horizons. The extent to which this ratio deviates from unity (the value associated with a random walk) has been used to measure market performance by Barnea (1974) and Hasbrouck and Schwartz (1988). As summary measures, variance ratios suffer from problems of sensitivity to the horizons used. Furthermore, there is no general connection between the variance ratio and conventional transaction-cost measures.

In the remainder of this article, I deal with the formal description and implementation of this approach. In Section 1, I introduce the basic security price decomposition and establish the economic significance of the pricing error and its dispersion. In the three following sections, I describe problems of inference in progressively more complex models. In Section 2, I illustrate the basic principles in situations where transaction prices constitute the only available data. In Sections 3 and 4, I discuss the expansion of the conditioning variable set to include other variables (prices and trades) with general lagged dependencies. I then turn to empirical implementation. In Section 5, I describe estimation strategies. In Section 6, I report results based on a sample of NYSE firms, and discuss intraday patterns. I conclude with a brief summary in Section 7.

1. The Transaction Price Decomposition

The models in this article take the logarithm of the observed transaction price at time t, p_t, as the sum of two components:

$$p_t = m_t + s_t. \tag{1}$$

The first component (m_t) is the efficient price, defined as the expectation of the final (end-of-trading) value of the security conditional on all public information available at time t, including whatever private information may be inferred from the published terms of the transaction. The second component (s_t) is the deviation between this efficient price and the actual transaction price, and is termed the pricing error. For present purposes, t is assumed to index either transactions or brief intervals of natural time.

The pricing error is central to this article. It is viewed as impounding diverse microstructure effects such as discreteness, inventory control, the non-information-based component of the bid-ask spread, the transient component of the price response to a block trade, etc., which are not explicitly modeled. The pricing error is related to the transaction cost in that, as defined, s_t is the cost for the buyer and $-s_t$ is the cost for the seller. The transaction cost for the buyer and that for the seller sum to zero for all trades. Even if the unconditional expectation of s_t is zero, however, it does not follow that all traders consider the pricing error to be a fair-game perturbation, diversifiable over many transactions. Conditional on trader identity and order placement strategy, the pricing error is generally not a fair game. A market-order trader, for example, must expect to incur a positive transaction cost on all transactions.[4]

The estimate of the pricing error for a particular transaction, denoted $\hat{s}_t(\cdot)$, i is a general function of conditioning data. If the estimate is sought as an ex post measure, it may incorporate posttrade information. If the goal is estimation of the cost of a future contemplated trade, the conditioning data must be restricted accordingly. Furthermore, S_t measures the deviation relative to an efficient price [m_t in (1)] that impounds the information that the market infers from the trade. Traders often view the transaction cost as the difference between their execution price and the price that prevailed before they entered the market ("slippage"). This leads to a pricing error different from the one considered here.

For purposes of assessing market design and regulation, I propose the standard deviation of the pricing error, σ_s, as a summary measure of market quality. Intuitively, this quantity measures how closely the transaction price tracks the efficient price. In the absence of a fully specified model of market operation, this dispersion does not have a direct transaction-cost interpretation. Its role as a proxy for market quality rests solely on the premise that as transaction costs and other barriers to trading are reduced, transaction prices should conform more closely to efficient prices.

From (1) the pricing error is defined only for trades that actually occur. Neither s_t nor σ_s directly reveals anything about the private or social cost of forgone trades (except in the sense that the pricing error impounds an immediacy cost). This shortcoming applies to most measures based on actual trade data. One can envision market regimes that achieve narrow bid-ask spreads, for example, by excluding traders who are likely to posses superior information. The apparent improvement in market quality implied by the smaller spread is offset

[4] Nevertheless, it cannot generally be assumed that $|s_t|$ is always a cost paid by a public trader to a market-maker. For example, if inventory control considerations or price continuity requirements force a dealer to place quotes sufficiently away from the efficient price, then a market-order trader may "pay" the spread and yet still have a negative transaction cost.

(and perhaps outweighed) by the social cost of forgone trades and reduced production of information.

Specification (1) is functionally similar to that used in Hasbrouck [1991b, equation (1)], except that in the earlier article I applied the decomposition to quote-midpoints. Prices are used here because the present goal is cost measurement for actual transactions. Quotes were used in the earlier article to avoid simultaneity problems involving trades. As in the earlier article, two assumptions are imposed on (1):

(i) The efficient price follows a random walk:

$$m_t = m_{t-1} + w_t, \tag{2}$$

where w_t are uncorrelated increments, $Ew_t = 0$, $Ew_t^2 = \sigma_w^2$, and $Ew_t w_\tau = 0$ for $t \neq \tau$.

(ii) The pricing error (s_t) is a zero-mean covariance-stationary stochastic process.

There are no requirements that s_t be serially uncorrelated or uncorrelated with w_t.

In the present article, I do not allow for heteroskedasticity in s_t or w_t. In the implementation, however, t indexes transactions rather than natural time. Periods of elevated return variance per unit time (the beginning and end of trading) are also periods of frequent transactions. Harris (1987) suggests that the behaviors ofvariance, skewness, kurtosis, and first-order autocorrelation of daily stock returns are consistent with a model in which prices and volumes evolve at a uniform rate in event time. Thus, using transaction time rather than natural time may mitigate return heteroskedasticity. Furthermore, because a primary source of return heteroskedasticity is time-of-day, estimations based on time-of-day subsamples will be presented.

The efficient price process specified in (2) does not include a drift term. Although the technique can easily be generalized to include a constant drift, practical econometric considerations favor suppression. The drift is the expected return (per increment in t). Most microstructure applications of this technique will involve data samples that are brief in calendar terms. In such situations, Merton (1980) notes that the standard error of the mean estimate is large relative to the expected return, and that the estimate is not improved by more frequent sampling. In consequence, he suggests that return variances using short data spans will have smaller estimation errors if they are centered around zero, rather than the sample mean return. Similar considerations are likely to apply in the present situation. This argument also suggests that the effect of ignoring time- and day-of-week variation in expected returns found by Harris (1986) is likely to be small.

Deviations from random-walk behavior in stock returns at daily and longer intervals have been found by Lo and MacKinlay (1988), Fama and French (1988), and Poterba and Summers (1988)) among others. If these deviations are attributed to low-frequency temporary components of stock prices, they should properly be impounded in the pricing error. The present analysis, which only allows for short-run (several-transaction) serial dependencies, is therefore probably misspecified. However, the present aim is the measurement of the microstructure-related component of s_t, which is presumably due to short-run phenomena. The misspecification will tend to impound the long-run components in the random-walk portion of the decomposition, leaving s_t to capture the short-run components of interest. In Section 5, I present a simulation that illustrates this point.

2. Inferences Based on Returns Only

In the preceding section, I motivated the economic interpretation and importance of the pricing error. I now turn to estimation of the pricing error and its dispersion. The exposition describes cases of increasing complexity. The model in this section uses returns only; that of the next section uses returns and signed trades; and the model of Section 4 uses broader and arbitrary sets of conditioning variables. With one exception (the lower-bound result), this discussion contains no new econometric results. It is a selective summary of the random-walk decomposition literature cited in the introduction, directed at microstructure issues.

To imbue (1) and (2) with specific economic content and to illustrate the econometric inferences, consider a special case with pricing error:

$$s_t = \alpha w_t + \eta_t, \tag{3}$$

where η_t is a disturbance uncorrelated with w_t. The Roll (1984) bid-ask spread model corresponds to $\alpha = 0$ and $\eta_t = \pm(\text{spread})/2$ (depending on the sign of the order). Alternatively, suppose that the spread is due in part to asymmetric information revealed in the trade and that the trade is the only update to the information set. This is a special case of Glosten (1987)) with no nontrade public information. In this case, $\alpha > 0$ will capture the cost-based component of the bid-ask spread and $\eta_t = 0$. A mixed case will result when the returns reflect public information supplementary to the trade.

The two terms in (3) correspond to components of the pricing error that may be thought of as information-correlated (αw_t) and information-uncorrelated (η_t). This distinction is a useful one for classifying microstructure effects. Information-uncorrelated pricing errors are likely to result from price discreteness, transient liquidity effects, inventory control effects, and "noise" trading. Information-correlated pricing errors arise from adverse-selection effects in the presence of fixed transaction costs and from lagged adjustment to information.

To estimate s_t and σ_s, it is necessary to connect the model to the observable data. Given (1)-(3), the return is

$$r_t = p_t - p_{t-1} = m_t - m_{t-1} + s_t - s_{t-1} = w_t + s_t - s_{t-1}. \tag{4}$$

Since w_t and s_t are serially uncorrelated (in the present example), r_t possesses nonzero autocovariances at the first lag only. This in turn implies that r_t may be represented as a first-order moving average process:

$$r_t = \epsilon_t - a\epsilon_{t-1}. \tag{5}$$

Two parameters, $\{a, \sigma_\epsilon^2\}$, fully characterize the mean and autocovariances of the return process. To deduce the return mean and autocovariances from the random-walk decomposition model requires specification of three parameters, $\{\sigma_u^2, \alpha, \sigma_\eta^2\}$. This model is therefore underidentified. (In this article, I make no assumptions and attempt no inferences concerning moments of higher order, such as skewness and kurtosis.)

In the macroeconomic literature, there are two common identification restrictions. The first, due to Beveridge and Nelson (BN) (1981)) implies for the present model that $\eta_t = 0$ in (3) [i.e., that the pricing error is entirely information-correlated ("Glosten model")]. The correspondence between (4) and (5) then establishes that $\epsilon_t - a\epsilon_{t-1} = w_t + \alpha w_t - \alpha w_{t-1}$. This further implies $\alpha = a/(1 - a)$, $w_t = (1 - a)\epsilon_t$, $s_t = \alpha w_t$, $\sigma_u^2 = (1 - a)^2\sigma_\epsilon^2$, and, $\sigma_s^2 = a^2\sigma_\epsilon^2$. By substitution and recursion, the estimate of the pricing error for a particular transaction is

$$\hat{s}_t(r_t, r_{t-1}, \ldots) = -a\epsilon_t = -a(r_t + ar_{t-1} + a^2 r_{t-2} + \cdots). \tag{6}$$

Under the BN restriction, this estimate is exact. In the following discussion, however, it will find application in more general situations.

The second identification restriction is due to Watson (1986), and implies for the model (3) that $\alpha = 0$ [i.e., that the pricing error is completely information-uncorrelated ("Roll model") and $\sigma_s^2 = \sigma_\eta^2$. The analysis (slightly more difficult than the BN case) shows[5] that $\sigma_w^2 = (1 - a)^2\sigma_\epsilon^2$ and $\sigma_s^2 = a\sigma_\epsilon^2$. In contrast with the BN case, it is not possible to express the pricing error as an exact function of the current

[5]The correspondence between (4) and (5) is $\epsilon_t - a\epsilon_{t-1} = w_t + \eta_t - \eta_{t-1}$. Taking the variance of each side establishes $(1 + a^2)\sigma_\epsilon^2 = \sigma_w^2 + 2\sigma_\eta^2$. Taking the first-order autocovariance of each side yields $a\sigma_\epsilon^2 = \sigma_\eta^2 = \sigma_s^2$. These may be solved to give the indicated results.

disturbance t,. Watson shows that the best linear estimate of s_t based on current and lagged returns is identical to that obtained under the BN restriction (6), but that estimates that incorporate future returns will be superior.

There are several interesting features of these two cases. First, note that the random-walk variance σ_w^2 is invariant. This is a general property of these decompositions. Intuitively, observed returns computed over progressively longer horizons are dominated to a greater degree by the random-walk component, irrespective of the form of the pricing error. Since both identification structures must be consistent with the same observed return process (5), it follows that the random-walk variance must also be the same.

Next, note from the expression for σ_s^2 that feasibility of the Watson identification requires $a \geq 0$. This implies that the first-order autocovariance of r_t is nonpositive. Recall that the Watson identification restriction $(\alpha = 0, \sigma_\eta^2 \neq 0)$ characterizes the simple Roll bid-ask model, which is incompatible with a positive return autocorrelation. In data samples, of course, a positive estimated autocorrelation can arise from estimation error. [See Harris (1990).]

Between the two identification restrictions, the dispersion of the pricing error differs considerably. To get a feel for this difference, note that for a first-order autocorrelation coefficient ρ near zero, the moving-average parameter a is approximately equal to ρ. For daily stock returns, French and Roll (1986) estimate the first-order autocorrelation coefficient to be ρ = –.0115 (their Table 2). The values for σ_s^2 obtained under the two identification restrictions differ by a factor of *a*, a factor of 100 in this case. The situation is generally better with transactions data, where a may reach 0.5. Nevertheless, these computations suggest that the identification issue is of great practical concern. To make matters worse, the economics of the problem do not argue conclusively in favor of either identification restriction. The earlier discussion noted that certain plausible microstructure effects are likely to be information-correlated, while others are likely to be information-uncorrelated.

The general lack of identification impairs estimation of s_t and σ_s, but some results are still attainable. Watson shows that the estimate of s_t given in (6) is the best linear estimate based on current and lagged returns not only for the two special cases considered here, but for all feasible identifications [i.e., all values of σ and σ_η^2 in (3) that imply the given $\{a, \sigma_\epsilon^2\}$].

The identification-invariance of this particular form of $\hat{s}_t$ leads to a lower bound for σ_s^2. To see this, consider the precision (mean-square error) of a general estimate of s_t, $\hat{s}_t$: $0 \leq E(s_t - \hat{s}_t)^2 = Es_t^2 - E\hat{s}_t^2 - 2E\hat{s}_t(s_t - \hat{s}_t)$. If $\hat{s}_t$ is the BN form given in (6), then the last term vanishes, since $E\hat{s}_t(s_t - \hat{s}_t) = E[E[\hat{s}_t(s_t - \hat{s}_t) \mid r_t, r_{t-1}, \ldots]] = 0$ This establishes the lower bound: $\sigma_s^2 = Es_t^2 \leq E\hat{s}_t^2$. It should be emphasized that the needed orthogonality condition is guaranteed to hold for the BN form of $\hat{s}_t$ only if the model (4) is correctly specified.

Now since this restriction rules out any information-uncorrelated effects, it is at this point a very weak lower bound. It will be shown in the next section, however, that the lower bound can be strengthened substantially by considering variables besides the transaction prices. The model eventually developed in this article is one that permits both information-correlated and uncorrelated effects.

Finally, the Watson identification restriction does not provide an upper bound for the pricing error. In fact, no meaningful upper bound for σ_s^2 exists: zero autocorrelations in r_t do not necessarily imply zero pricing errors.[6]

3. Inferences Based on Returns and Trades

In this section, I describe how the price error framework may be strengthened by including explanatory variables in addition to prices. The discussion is based on a simple bivariate model of prices and trades that can be viewed as the reduced form of a reasonable microstructure model.

The trade variable at time t is the trade volume x_t, signed to be positive if the agent who initiates the trade is buying and negative if the agent is selling. Trades are assumed to be symmetric: $Ex_t = 0$. It is provisionally assumed that trades and the pricing error are not autocorrelated. (These restrictions will be dropped in the next section.) With the availability of the trade variable, the random-walk component in (1) may itself be decomposed into two parts:

Here, γx_t reflects the information the market infers from the trade. The second part (u_t) is an innovation uncorrelated with x_t that stems from nontrade public information. The pricing error may also be written to reflect a trade component:

$$s_t = \alpha x_t + \beta u_t + \eta_t, \tag{8}$$

[6]Suppose that $\mathrm{Var}(r_t) = 1$, with zero autocovariances at ail lags. This is consistent with the "natutal" model: a random-walk decomposition in which the pricing error is completely absent ($s_t = 0$) and the efficient price changes have unit variance ($\sigma_w^2 = 1$). The autocovariances are also consistent, however, with a model in which the pricing error is perfectly negatively correlated with the changes in the efficient price ($s_t = -w_t$), in which case $r_t = w_t - w_t + w_{t-1}$. This implies $\mathrm{Var}(r_t) = 1$ (as required), but the pricing error variance is $\sigma_s^2 = 1$. The example may be extended by allowing longer lags on the pricing error. If $s_t = -w_t - w_{t-1}$, then although $\mathrm{Var}(r_t) = 1$ and the return autocovariances remain at zero, $\sigma_s^2 = 2$. In general, the pricing error variance may be made arbitrarily large by considering longer lags. Quah (1992) discusses related points.

where η_t i is a disturbance that is uncorrelated with both x_t and u_t. This specification assumes that both w, and s, are linear in x_t, a restriction that will be relaxed in the empirical analysis. The behavior of the return series implied by the random-walk decomposition model is

$$r_t = b_0 x_t + b_1 x_{t-1} + \epsilon_t - a\epsilon_{t-1}, \tag{9}$$

that is, a regression with a moving-average error term. The correspondence between (4) [substituting in from (8)] and (9) shows

$$\begin{aligned} &b_0 x_t + b_1 x_{t-1} + \epsilon_t - a\epsilon_{t-1} \\ &\quad = (\alpha + \gamma)x_t - \alpha x_{t-1} + (1 + \beta)u_t + \eta_t - \beta u_{t-1} - \eta_{t-1}. \end{aligned} \tag{10}$$

The random-walk variance is identified: $\sigma_w^2 = (b_0 + b_1)^2\sigma_x^2 + (1 - a)^2\sigma_\epsilon^2$. Without additional restrictions, the remaining parameters of the random-walk decomposition are not identified.

The BN restriction sets $\eta_t = 0$ in (8). This implies $\alpha = -b_1, \gamma = b_0 + b_1, \beta = a/(1 - a)$, and $u_t = (1 - a)\epsilon_t$. In addition, $\sigma_u^2 = (1 - a)^2\sigma_\epsilon^2$, and therefore, $\sigma_s^2 = b_1^2\sigma_x^2 + a^2\sigma_\epsilon^2$. The properties of the BN representation in the univariate case apply here as well. A minor generalization of Watson (1986) shows that for all feasible identifications, the best linear estimate of the pricing error based on current and lagged returns and signed trades is that obtained under the BN identification restriction. [This estimate is computed by using the BN values for a and β in (8) and dropping the η_t term.] Since this estimate is identification-invariant, the associated σ_s is a lower bound by the same logic as in the univariate case.

This simple bivariate specification may be viewed as a representation of a microstructure model in the spirit of Glosten and Milgrom (1985). The market consists of a risk-neutral dealer, and informed and uninformed traders. At the close of period t - 1, the trading history and the efficient price m_{t-1} are common knowledge. At the beginning of period t, public information (u_t) arrives, making the expected security value $m_{t-1} + u_t$. About this value, the dealer posts bid and ask quotes: $q_t^b = m_{t-1} + u_t - c$ and $q_t^a = m_{t-1} + u_t + c$, where c is the half-spread. The arriving trader faces these quotes and may buy or sell one unit of the security $x_t = \{-1, +1\}$. The transaction price is $p_t = q_t^b$ if $x_t = -1$ and $p_t = q_t^a$ if $x_t = +1$. From the relations $p_t = m_{t-1} + u_t + cx_t$ and $m_t = m_{t-1} + u_t + \gamma x_t$, it is apparent that $s_t = p_t - m_t = (c - \gamma)x_t$. This is identical with (8) with $\alpha = (c - \gamma)$, $\beta = 0$, and $\eta_t = 0$ (i.e., the BN restriction). Therefore, s_t may be computed directly, and the lower bound is exact. In this example, it is interesting to contrast the estimated pricing error variance with the bid-ask spread. The half-spread is, of course, c. This overstates the pricing error, because a portion of this (γ) reflects the update to the efficient price conditional on the trade. The pricing error reflects only the remaining component, $(c - \gamma)$.

4. Inference in the General Multivariate Case

In the previous section, I generalized the basic model of Section 1 by permitting inclusion of other variables. In this section, I expand the model further to allow general serial correlations in the returns (and other explanatory variables). From a microstructure perspective, this is important because many market imperfections lead to lagged effects. These include inventory control mechanisms, lagged price adjustment, and price discreteness. The precise nature of these effects is unknown, however, and there is in consequence no strong prior specification.

As in Hasbrouck (1991a, 1991b), vector autoregression (VAR) provides a useful framework that is general enough to capture the unspecified lagged dependencies and is also amenable to computation. A representative bivariate VAR involving trades and price changes is

$$\begin{aligned} r_t &= a_1 r_{t-1} + a_2 r_{t-2} + \cdots + b_1 x_{t-1} + b_2 x_{t-2} + \cdots + v_{1,t}, \\ x_t &= c_1 r_{t-1} + c_2 r_{t-2} + \cdots + d_1 x_{t-1} + d_2 x_{t-2} + \cdots + v_{2,t}. \end{aligned} \tag{11}$$

In the present discussion, x_t may be taken as the signed trade variable defined in the previous section. More generally, x_t is a column vector of explanatory variables, and b_i and di are conformable coefficient matrices. The innovations are zero-mean, serially uncorrelated disturbances. The VAR may be transformed to obtain a vector moving average (VMA) representation that expresses the variables in terms of current and lagged disturbances [see Judge et al. (1985, p. 657)]. A VMA corresponding to (11) is

$$\begin{aligned} r_t = {} & a_0^* v_{1,t} + a_1^* v_{1,t-1} + a_2^* v_{1,t-2} + \cdots + b_0^* v_{2,t} + b_1^* v_{2,t-1} \\ & + b_2^* v_{2,t-2} + \cdots, \\ x_t = {} & c_0^* v_{1,t} + c_1^* v_{1,t-1} + c_2^* v_{1,t-2} + \cdots + d_0^* v_{2,t} + d_1^* v_{2,t-1} \\ & + d_2^* v_{2,t-2} + \cdots. \end{aligned} \tag{12}$$

For the present calculations, only the r_t equation in (12) is used. The underlying random-walk decomposition model is (1), but with an expanded representation for the pricing error:

$$\begin{aligned} s_t = {} & \alpha_0 v_{1,t} + \alpha_1 v_{1,t-1} + \cdots + \beta_0 v_{2,t} + \beta_1 v_{2,t-1} \\ & + \cdots + \eta_t + \gamma_1 \eta_{t-1} + \cdots, \end{aligned} \tag{13}$$

where η_t is a disturbance orthogonal to all components of v_t.

By a slight modification of the computation in the appendix to Hasbrouck (1991b), the variance of the random-walk component may be computed as

$$\sigma_w^2 = \left[\sum a_i^* \sum b_i^*\right] \mathrm{Cov}(v) \begin{bmatrix} \sum a_i^{*} \\ \sum b_i^{*\prime} \end{bmatrix}. \tag{14}$$

[In Hasbrouck (1991b) the disturbance covariance matrix *Cov(v)* is block-diagonal, a structure that is not imposed here.] By a minor generalization of the demonstration in Beveridge and Nelson (1981), the α's and β's in (13) may be computed under the BN identification restriction $(\eta_t = \gamma_1 = \cdots = 0)$:

$$\alpha_j = -\sum_{k=j+1}^{\infty} a_k^*, \qquad \beta_j = -\sum_{k=j+1}^{\infty} b_k^*. \tag{15}$$

When these coefficients are used in (13) and the η_t terms are dropped, the result is an identification-invariant best-linear estimate of s_t conditional on current and lagged v_t. The pricing error variance may be computed as

$$\sigma_s^2 = \sum_{j=0}^{\infty} [\alpha_j \quad \beta_j] \mathrm{Cov}(v) \begin{bmatrix} \alpha_j \\ \beta_j' \end{bmatrix}. \tag{16}$$

By the same logic as that used in the discussion of the univariate model, this constitutes a lower bound over the more general cases.

To analyze the practical usefulness of the lower bound, it is necessary to consider what economic forces might cause it to be exceeded. Intuitively, the present technique regresses (projects) the actual pricing error on a set of known variables. Under the BN model is the one-lag univariate case described in the introduction, the pricing error is given by (3). Under the BN restriction, w_t may be computed from past returns. Viewing (3) as a linear regression, the explained variance is $\alpha^2\sigma_w^2$. The true value for σ_s^2 is $\alpha^2\sigma_w^2 + \sigma_\eta^2$. The excess is due to the "residual variance" σ_η^2. This is a general property of random-walk decompositions. In the most general multivariate model, the "residual" in the pricing error is constituted by the η_t components of (13). and these are the terms dropped under the BN restriction.

The residuals in a linear regression are uncorrelated with the explanatory variables, and this is true of the projection residual in (3) as well. This immediately leads to the basic principle: any factor that causes σ_s^2 to exceed its lower bound cannot be perfectly correlated with any linear combination of the explanatory variables. The residual η_t in the univariate pricing error (3) is uncorrelated with current and lagged returns. It is not too difficult to conceive of a pricing error component that is uncorrelated with lagged returns. In the univariate model, a simple fixed-cost (i.e., non-information-based) bid-ask spread satisfies the requirement.

In the bivariate model of Section 2, the residual η_t in the pricing error (8) must be uncorrelated with current and lagged returns *and* trades. When trades are included, the BN specification picks up the fixed-cost portion of the spread. In seeking factors that might cause the actual value of σ_s^2 to exceed the lower bound, one is limited to components uncorrelated with current and lagged returns and trades. This is a significant restriction. It must be admitted, however, that such components may arise from discreteness in reported prices. A simulation dealing with this issue is presented in Section 5.

The distinction between the use of transaction prices in the present article and the use of quote midpoints in Hasbrouck (1991b) can be explained as follows. The analysis of the earlier article required that the contemporaneous causality running from trades to quotes be one-way. This was justified economically by noting that in actual market operations the quote is adjusted subsequently to the trade. In the econometric specifications, this leads to a recursive structure: the earlier article employs a VAR that is very similar to (11), except that the return specification includes the contemporaneous trade, x_t. The disturbance covariance matrix is consequently block-diagonal, and the impact of a trade on the quotes may be determined relatively precisely. In the present framework, however, no one-way contemporaneous causality can be asserted a priori for trades and actual transaction prices. The two are determined simultaneously, often as the result of negotiation. Trade-price simultaneity also implies that it is generally inappropriate to regard the $\beta_i v_{2,t-i}$ terms in (13) as measuring the pricing error "caused" by a given set of trade innovations.[7]

5. Estimation

An estimated VMA (12) may be obtained by inverting an estimated VAR similar to (11) that is truncated at some lag beyond which all serial dependencies are assumed negligible. Equations (13) (with all $\eta_t = 0$) and (15) may be used to obtain the BN estimate of s_t. The BN (lower-bound) estimate of σ_s is computed from (16). Standard errors for the BN σ_s estimate and the s_t estimate [under the assumption that

[7]The BN estimate of σ_s does not change when the contemporaneous trade is included in the return specification. More generally, the left-hand side of the r_t specification in (12) does not change when it is written in terms of the transformed disturbances $v_t^\dagger = Av_t$ and coefficients $[a_i^\dagger\ b_i^\dagger] = [a_i\ b_i]A^{-1}$, for an orthonormal rotation matrix A. Inclusion of the contemporaneous trade simply diagonalizes $Cov(v_t)$.

the η_t in (13) equal zero] can be obtained using generalized method of moments techniques. The present article, however, is concerned only with summary averages of σ_s over broad classes of stocks. In consequence, the standard error of the mean estimates are easily computed, while the standard errors for the individual stocks are not required. Furthermore, the techniques are applied to stocks with numerous observations.

As an example, consider a simulation of the model of Section 3. The trades are generated as $x_t = \{+1, -1\}$ with equal probability (implying $\sigma_x^2 = 1$). The half-spread is $c = 0.005$. The other parameters are $\sigma_u = 0.00346$ and $\gamma = 0.002$.[8] Since $s_t = \alpha x_t$ and $\alpha = c - \gamma$, $\sigma_s =$ 0.003 (roughly 0.3 percent of the stock price). From (7), $\sigma_w = 0.004$. The generated sample consisted of 1000 observations on r_t and x_t.

Consider first an application of the simple univariate return analysis discussed in Section 2. The estimated variance and first-order autocovariance for the generated data were $\mathrm{Var}(r_t) = 0.0000468$ and $Cov(r_t, r_{t-1}) = 0.0000155$, implying a first-order autocorrelation of $\rho = -.331$. The corresponding parameters of the first-order moving-average model are $a = 0.378$ and $\sigma_\epsilon^2 = 0.0000410$. Using the BN identification restriction, $\sigma_s = 0.0024$ (i.e., a 20 percent underestimate). Using the Watson restriction, $\sigma_s = 0.0039$ (a 30 percent overestimate).

Consider next a joint bivariate analysis of r_t and x_t. Since these data are generated from a known model, one could estimate the reduced-form specification (9) using maximum-likelihood methods. The method used here, however, is the vector autoregression method described in the last section. The estimated bivariate VAR truncated at one lag is

$$\begin{aligned} r_t &= \underset{(-0.88)}{-0.375 r_{t-1}} - \underset{(-9.27)}{0.00271 x_{t-1}} + \hat{v}_{1,t}, \\ x_t &= \underset{(-1.54)}{-10.6 r_{t-1}} + \underset{(1.55)}{0.0729 x_{t-1}} + \hat{v}_{2,t}, \end{aligned}$$

$$\mathrm{Cov}(\hat{v}) = \begin{bmatrix} 0.0000384 & 0.00511 \\ 0.00511 & 0.997 \end{bmatrix}.$$

The corresponding VMA representation (also truncated at one lag) is

$$r_t = v_{1,t} - 0.0375 v_{1,t-1} - 0.00271 v_{2,t-1},$$

$$x_t = -10.6 v_{1,t-1} + v_{2,t} + 0.0729 v_{2,t-1}.$$

[8] These parameters arise as follows. A quarter-spread on a $25 security is *a* proportional spread of 0.01. Of the 0.005 half-spread, $\gamma = 0.002$ is assumed due to adverse selection. The intertransaction random-walk variance is therefore $\sigma_w^2 = \gamma^2\sigma_x^2 + \sigma_u^2 = 0.002^2(1) + 0.000012 = 0.000016$ On an annual basis, assuming 250 trading days and 10 transactions per day, this is 0.04, an annual return standard deviation of approximately 20 percent.

From (14), the variance of the random-walk component may be computed as

$$\hat{\sigma}_w^2 = [0.963 \quad -0.00271]\begin{bmatrix} 0.000038 & 0.00511 \\ 0.00511 & 0.997 \end{bmatrix}\begin{bmatrix} 0.963 \\ -0.00271 \end{bmatrix}$$
$$= 0.0000162.$$

Under the BN restriction, the estimated pricing error is [from (6) and (15)]

$$\hat{s}_t = 0.00375\hat{v}_{1,t} + 0.00271\hat{v}_{2,t}. \tag{17}$$

Given the relative magnitudes of $\mathrm{Var}(\hat{v}_{1,t})$ and $\mathrm{Var}(\hat{v}_{2,t})$, $\hat{s}_t \approx 0.00271\hat{v}_{1,t} \approx 0.00271x_t$. (By construction, the coefficient of x_t is $\alpha = c - \gamma = 0.003$.) The estimated standard deviation is $\sigma_s = 0.0029$, which is close to the assumed value (0.003). The computations cited to this point are based on one sample draw. Over 30 sample draws, the average value of the σ_s estimates was 0.00303, and the standard deviation of these estimates was 0.000027.

Simulations are also useful for investigating the consequences of misspecification. As noted earlier, two particularly problematic features of return data involve fads and discreteness. To investigate the first of these, the prices simulated in the first study were perturbed by the addition of a constructed fad of the form $f_t = 0.99835f_{t-1} + \eta_t$, where $\mathrm{Var}(\eta_t) = 8.24 \times 10^{-6}$ and η_t is independent of x_t and u_t. These parameters imply a fad with a standard deviation of 0.05 (approximately 5 percent of the stock price) and a half-life of 420 transactions. With the inclusion of this fad, the pricing error variance may be written as $\sigma_s^2 = (0.003)^2 + (0.05)^2$, a value that is strongly dominated by the fad. Fads are generally viewed, however, as arising from considerations beyond the short-run operations of the market. In the present context, the relevant question is, to what extent does the σ_s estimated using a truncated VAR capture the microstructure-based component of the pricing error? As a large-sample illustration, when a five-lag VAR analysis was implemented for a single draw of 50,000 observations, the estimated σ_s was found to be 0.00298. In 30 draws of 1000 observations each, the mean and standard deviation of the estimates were 0.00308 and 0.00027. These results support the conjecture, advanced in Section 1, that the misspecification is likely to limit the estimated σ_s to short-run microstructure effects.

To study the effects of discreteness, the generated log transaction prices were rounded to the nearest 0.005 (1/8 for a $25 security). The rounding elevates the standard deviation of the pricing error. In a single draw of 50,000 observations, the standard deviation of the generated s_t was found to be 0.00334. The value for σ_s estimated using the VAR procedure was 0.00299. Thus, presumably because the discreteness disturbance is uncorrelated with either trades or movements in the efficient price, it is not captured by the BN σ_s measure. The contribution of discreteness remains, however, a relatively small part of the total.

6. A Profile of the NYSE

In this section, I present the computation of the lower-bound pricing-error variance calculations for a sample of NYSE issues. The aims of this analysis are illustration of the technique, evaluation of patterns in the pricing-error variance across market-value subsamples, and a preliminary analysis of intradaily patterns. Estimates of individual pricing errors are not computed, since in the absence of knowledge of trader identity, they are not highly meaningful.

The transactions data were collected from the Institute for the Study of Security Markets (ISSM) tape for the first quarter of 1989. Supplementary data were obtained from the daily CRSP file. First, for all firms present on the ISSM and CRSP tapes, I computed equity capitalizations as of the close of 1988, and formed four subsamples based on these ranked capitalizations. I then applied the analysis to the first 50 firms in each subsample that had at least 500 transactions over the quarter. (For the lowest-value subsample, only 27 firms satisfied this criterion.)

In constructing the time series of returns and trades, natural time was ignored. The data were viewed as an untimed sequence of observations, and the time subscript t was incremented each time a transaction occurred. Trade classifications were made by reference to the quotes. As noted by Hasbrouck (1988) and Lee and Ready (1991), the NYSE reporting process delays most trades relative to quotes, leading to a spurious reversal in the transaction record. In consequence, trade classification was made on the basis of quotes prevailing as of five seconds prior to the trade print time. For trades that occurred at the midpoint of these quotes, the trade variable x_t, was set to zero. Overnight returns were not used. The trade/returns process was assumed to start afresh each morning, at which time lagged values of trades and returns were set to zero.

Summary statistics are reported in Table 1. As expected, beginning share price and number of transactions are positively related to market value, and the average log spread is negatively related to market value. The proportion of unclassified trades (roughly 42 percent) is slightly higher than the 30 percent found by Lee and Ready (in a 1988 non-censored sample). The proportion of unclassified volume is lower than the proportion of unclassified trades, suggesting that the mid-point-trade problem is less significant for larger trades. Furthermore, the classification proportions exhibit a pattern across market-value subsamples. The lowest market-value subsample has the smallest midpoint proportions, and with increasing market value, the proportions generally increase.

In Table 1, I present two estimates of the BN pricing error dispersion. The first, $\sigma_{s,r}$, is based on a five-lag univariate autoregression (i.e., r_2, was regressed against $\{r_{t-1}, \ldots, r_{t-5}\}$). The average for the total sample is 0.243 (i.e., approximately 0.243 percent of the stock price). A striking feature of the $\sigma_{s,r}$ averages is that they are substantially lower than the corresponding average spreads. For the total sample, the average log spread is 1.52 percent. If the spread were entirely due to liquidity and non-information-related transaction costs (cf. the Roll model), the analysis of Section 1 suggests that the value of σ_s should be (1.52 percent)/2 = 0.76 percent. At 0.243 percent, the average value of $\sigma_{s,r}$ is only about a third of this. The estimate values are in some respects sensitive to choice of lag length: Generally, σ_s increases with the cutoff.[9] The cross-sectional patterns (across firms and across time) remain unaffected, however.

Of course, $\sigma_{s,r}$ is a lower bound, and from the earlier remarks, it would be expected to understate the true value of σ_s in the presence of a pricing error component that is uncorrelated with information, but correlated with trades. Addition of trade variables to the explanatory variable set would therefore be expected to strengthen this lower bound. This enhancement was implemented as follows. Following Hasbrouck (1991a, 1991b), the signed trade variable of power k is defined by $x_t^k = \text{sign}(x_t)|x_t|^k$. Thus, x_t^0 is an indicator variable that takes on values in $\{-1, 0, +1\}$, $x_t^1 = x_t$, and $x_t^{1/2}$ is a signed square-root trade variable (included to allow for concavity in the trade-price relation). The VAR was then estimated over the four-variable set $\{r_t, x_t^0, x_t, x_t^{1/2}\}$: each variable was regressed against the full set through the fifth lag. The simultaneous use of various powers of the trade is intended to allow for nonlinear dependencies in both the random-walk and pricing error components.

Average values for these trade-based estimates of the standard deviation of the pricing error, denoted $\sigma_{s,r,x}$, are presented in Table 1. These values are generally higher than the corresponding $\sigma_{s,r}$ values. This is most striking for the low market-value subsample. In moving to the higher market-value subsamples, the discrepancy declines, and the highest market-value subsample, the trade-based σ_s's, are only

[9] This increase may be occurring for methodological reasons. As discussed in Section 4, the pricing error variance is essentially an "explained variance" in an implicit linear regression. In any given data sample, such a quantity will rend to increase whenever explanatory variables are added. The increase may also represent, however, contributions to the pricing error from economic forces operating over horizons longer than the brief spans generally considered in microstructure analysis.

Table 1
Summary statistics by market capitalization

	Total sample	Market-value subsamples 1 (Lowest)	2	3	4 (Highest)
No. of firms	175	27	50	48	50
Equity capitalization ($MM)	1145	26	93	416	3502
Beginning share price ($)	20.05	7.56	11.96	20.37	34.57
Average log spread (×100)	1.52	2.63	1.98	1.29	.70
No. of transactions	3489	761	1221	2792	7899
Proportion of unclassified trades	.420	.297	.409	.448	.451
Proportion of unclassified volume	.342	.283	.352	.347	.349
$\sigma_{s,r}$ (×100)	0.243	0.305	0.329	0.230	0.136
	[0.160]	[0.182]	[0.164]	[0.148]	[0.064]
$\sigma_{s,r,x}$ (×100)	0.330	0.552	0.411	0.307	0.153
	[0.235]	[0.285]	[0.161]	[0.238]	[0.078]

The sample is drawn from a sample of NYSE firms stratified by market value of equity. The number of firms in the total sample and in the subsamples is reported in the first line. All other figures in the table are sample averages. Equity capitalization and beginning share price are taken from the CRSP file. The average log spread for a given firm is an average of the difference between the log of the bid and the log of the ask quotes, weighted by the time over which the quote was in effect (from the ISSM tape, 1989 first quarter). The number of transactions is the number reported on the ISSM tape for the first quarter of 1989. The pricing-error standard deviations are computed under the Beveridge- Nelson identification restriction: $\sigma_{s,r}$ is based on a univariate autoregressive model of the stock return. $\sigma_{s,r,x}$ is based on a VAR of returns and trades. Standard deviations of the variable are given in brackets.

slightly higher than $\sigma_{s,r}$. The trade-based σ_s's remain, however, substantially lower than the average half-spreads.

In discussing the magnitudes of these estimates, it is well to emphasize that identification of σ_s with an average transaction cost for a particular trader or class of traders requires further assumptions about the operation of the market. Empirical implementation generally also requires identification of those transactions in which the traders participated. Approximate inferences can be made under simplifying assumptions. If it is assumed that the initiator of a trade (such as a market-order trader) always incurs a positive cost, then the transaction cost is $|s_t|$. (The receipt of this amount by the contraparty is presumably compensation for market-making services.) If s_t were normally distributed, then the expected transaction cost for these traders would $E|s_t| = (\sqrt{2/\pi})\sigma_s \approx 0.8\sigma_s$, roughly 0.8(0.33 percent) = 0.26 percent for the NYSE total sample. This should be viewed as very approximate because of the assumptions of one-sidedness in the cost, symmetry of buys and sells, and normality.

In a pure liquidity model in which all trades occurred at the bid or the ask, and trades conveyed no information, one would expect to see σ_s equal to one half the spread. In actuality, the average values

Table 2
Summary statistics by time of day

	Morning (9:30–10:00 A.M.)	Midday (10:00 A.M.–3:30 P.M.)	Afternoon (3:30–4:00 P.M.)
$\sigma_{s,r} \times 100$	0.165**	0.197	0.229**
	[0.117]	[0.149]	[0.232]
$\sigma_{s,r,x} \times 100$	0.395*	0.218	0.258**
	[0.618]	[0.166]	[0.241]
Average log spread × 100	1.000**	0.905	0.913
	[0.623]	[0.593]	[0.597]

The values in the table are sample means and (in brackets) standard deviations of the pricing-error standard deviations estimated for each firm in the sample. $\sigma_{s,r}$ is based on a univariate autoregressive model of the stock return. $\sigma_{s,r,x}$ is based on a VAR of returns and trades. The average log spread is time weighted.

*,** The estimate differs from the corresponding midday estimate at significance levels of .10 and .05. respectively (using a sign-rank test on the differences).

for the trade-based σ_s's are roughly one quarter of the corresponding average spreads. There are at least three considerations bearing on this discrepancy. First, a substantial proportion of the trades occur at the midpoint of the bid and ask quotes. This decreases the effective spread. As an offsetting effect, however, the quoted spread is only valid for trades of relatively small size. Large trades may take place outside of the quoted spread. Finally, besides the fixed transaction costs, the quoted spread impounds an asymmetric information charge not included in the pricing error. In short, the discrepancy between the estimated pricing-error standard deviation and the half-spread is easily accounted for by considerations that illustrate and reinforce the conclusion that posted spreads are suspect measures of market quality.

Previous research has characterized beginning- and end-of-trading elevations in returns, variances, and volumes. Spreads have also been found to be higher at the beginning of trading [see Harris (1986, 1989), Foster and Viswanathan (1990, 1993), Jain and Joh (1988), McInish and Wood (1992), Mulherin and Gerety (1989), and Wood, McInish, and Ord (1985)]. Given that the pricing error impounds microstructure effects that are likely to depend on these factors, there are grounds for expecting beginning- and end-of-trading patterns in the pricing error as well. To explore this hypothesis, separate estimates of the pricing-error standard deviation were computed for the beginning-of-day (defined as the first half hour of trading), end-of-day (last half hour of trading), and midday. When a 250-transaction minimum was imposed on the subperiods, the sample declined to 36 firms. The estimates are reported in Table 2. The trade-based estimate of σ_s is higher at the beginning and, to a lesser extent, at the end of the trading session. (This pattern is only partially consistent with that of the average log spread, which is not substantially elevated at the end of trading.)

7. Conclusions

In this article, I define the pricing error (s_t) as the difference between the actual transaction price and an implicit unobservable efficient price assumed to follow a random walk. The standard deviation of this pricing error (σ_s) naturally arises as a measure of market quality. In establishing the correspondence between this model and the observed transactions data, it is shown that generally neither s_t nor σ_s is econometrically identified. It is nevertheless possible to compute certain estimates of these values: a best-linear estimate of s, based on current and lagged data and a lower bound for σ_s. Both estimates may be improved by the addition of explanatory variables. If signed trade variables are included in the estimation set, the estimates are likely to be quite close to the true values.

The empirical implementation surveys estimates of σ_s for a sample of NYSE firms. For the full sample, the average value of the estimated σ_s is approximately 0.33 percent of the stock price. The estimates of σ_s are negatively related to market value, and also exhibit elevation at the beginning and end of the trading session.

There are many directions for further research. In the area of transaction-cost measurement, it would be useful to compute average values of the pricing error for particular traders or classes of traders. Along the lines of comparative market assessment, the financial community is currently experiencing a period of expansion and experimentation in alternative market regimes. The σ_s statistic presented in this article may be a useful tool for assessing the fitness of different market structures.

References

Admati. A. R.. and P. Pfleiderer. 1988, “A Theory of Intraday Patterns: Volume and Price Variability,” *Review of Financial Studies,* 1, 3-40.

Barnes, A., 1974, “Performance Evaluation of New York Stock Exchange Specialists,” *Journal of Financial and Quantitative Analysis,* 9, 511 -535.

Beebower. G., 1989, “Evaluating Transaction Cost,” in W. H. Wagner (ed.), *The Complete Guide to Securities Transactions: Enhancing Investment Performance and Controlling Cost,* Wiley, New York.

Berkowitz, S. A., D. E. Logue, and E. A. Noser. 1988, “The Total Cost of Transactions on the NYSE,” *Journal of Finance,* 41, 97-112.

Beveridge, S., and C. Nelson, 1981, “A New Approach to the Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to the Measurement of the ‘Business Cycle’,” *Journal of Monetary Economics,* 7, 151-174.

Campbell, J. Y., and N. G. Mankiw, 1987, "Permanent and Transitory Components in Macroeconomic Fluctuations," American Economic Review, 77, 111-117.

Commodities Futures Trading Commission, 1989, "Economic Analysis of Dual Trading on Commodity Exchanges," working paper, Division of Economic Analysis. Washington, DC.

Cooper, K.. J. Groth, and W. Avera. 1985. "Liquidity, Exchange Listing, and Common Stock Performance," *Journal of Economics and Business,* 37, 21-33.

Fama, E. F.. and K. R. French, 1988, "Permanent and Temporary Components of Stock Prices," *Journal* of Political Economy, 96, 246-273.

Foster, D. F., and S. Viswanathan, 1990, "A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Securities Markets," *Review of Financial Studies, 3, 593-624.*

Foster, D. F.. and S. Viswanathan, 1993, "Variations in Trading Volume, Return Volatility and Trading Costs: Evidence on Recent Price Formation Models," *Journal of Finance,* forthcoming.

French, K. R., and R. Roll, 1986, "Stock Return Variances: The Arrival of Information and the Reaction of Traders," *Journal of Financial Economics,* 17, 5-26.

Glosten, L. R., 1987, "Components of the Bid-Ask Spread and the Statistical Properties of Transaction Prices," *Journal of Finance,* 42, 1293-1307.

Glosten, L. R., and P. R. Milgrom, 1985, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics,* 14, 71-100.

Grossman, S. J.. and M. H. Miller, 1988, "Liquidity and Market Structure," *Journal of Finance, 43,* 617-633.

Harris, L. E., 1986, "A Transactions Data Study of Weekly and Intradaily Patterns in Stock Returns," *Journal of Financial Economics,* 16, 99-117.

Harris, L. E., 1987, "Transaction Data Tests of the Mixture of Distributions Hypothesis," *Journal of Financial and Quantitative Analysis,* 22, 127-141.

Harris, L. E., 1989, "A Day-End Transaction Price Anomaly," *Journal of Financial and Quantitative Analysis,* 24. 29-45.

Harris, L. E.. 1990. "Statistical Properties of the Roll Serial Covariance Bid/Ask Spread Estimator," *Journal of Finance,* 45,579-590.

Hasbrouck. J., 1988, "Trades. Quotes, Inventories and Information," *Journal of Financia1 Economics,* 22. 229-252.

Hasbrouck. J., 1991a, "Measuring the Information Content of Stock Trades," *Journal of Finance,* 46, 179-207.

Hasbrouck. J., 1991b, "The Summary Informativeness of Stock Trades: An Econometric Analysis," *Review of Financial Studies.* 4, 571-595.

Hasbrouck. J., and R. A. Schwartz, 1988, "An Assessment of Stock Exchange and Over-the-Counter Markets," *Journal of Portfolio Management,* 14, 10-16.

Jam, P., and G. Job. 1988, "The Dependence Between Hourly Prices and Trading Volume." *Journal of Financial and Quantitative Analysis.* 23, 269-283.

Judge, G. G., W. E. Griffiths, R. C. Hall. H. Lutkepol, and T-C. Lee, 1985, *The Theory and Practice of Econometrics* (2d ed.), Wiley, New York.

Lee. C., and M. Ready, 1991, "Inferring Trade Direction from Intradaily Data," *Journal of Finance,* 46, 733-746.

Lo, A. W.. and A. C. MacKhtlay, 1988. "Stock Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies,* 1, 41-66.

McInish. T. H., and R. A. Wood, 1992, "An Analysis of Intraday Patterns in Bid/Ask Spreads for NYSE Stocks," *Journal of Finance,* 47, 753-764.

Merton. R. C., 1980, "On Estimating the Expected Return on the Market: An Exploratory Investigation." *Journal of Financial Economics,* 8, 323-361.

Mulherin, J. H., and M. S. Gerety, 1989, "Intraday Trading Behavior in Securities Markets: Hourly NYSE Volume and Returns, 1933-1988," working paper, US. Securities and Exchange Commission.

Neal. R., 1989, "Market Structure and Transaction Costs," working paper, University of Washington

Perold, A., 1988, "The Implementation Shortfall: Paper versus Reality," *Journal of Portfolio Management,* 14, 4-9.

Poterba, J. M., and L. H. Summers, 1988, "Mean Reversion in Stock Prices: Evidence and Implications." *Journal of Financial Economics,* 22, 27-60.

Quah, D., 1990, "Permanent and Transitory Movements in Labor Income: An Explanation for 'Excess Smoothness' in Consumption," *Journal of Political Economy,* 98, 449-475.

Quah, D., 1992, "The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds," *Econometrica,* 60, 107-118.

Roll, R., 1984, "A Simple Model of the Implicit Bid-Ask Spread in an Efficient Market," *Journal of Finance, 39. 1127-1139.*

Stock, J. H., and M. W. Watson, 1988, "Variable Trends in Economic Time Series," *Journal of Economic Perspectives,* 2, 147-174.

Tanner, J. E., and J. B. Pritchett, 1992, "The Pricing of Market Maker Services Under Siege: Nasdaq vs. NYSE on Black Monday," in *The NASDAQ Handbook,* Probus. Chicago.

Watson, M. W., 1986, "Univariate Detrending Methods with Stochastic Trends," *Journal of Monetary Economics, 18. 49-75.*

Wood, R. A., T. H. Mclnish, and J. K. Ord, 1985, "An Investigation of Transactions Data for NYSE Stocks," *Journal of Finance, 40, 723-739.*