WHY DO SECURITY PRICES CHANGE? A TRANSACTION-LEVEL ANALYSIS OF NYSE STOCKS

by

Ananth Madhavan
Matthew Richardson
Mark Roomans

20-94

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

Why Do Security Prices Change? A Transaction-Level Analysis of NYSE Stocks

Ananth Madhavan, Matthew Richardson, and Mark Roomans*

Latest Revision: July, 1994

Abstract

This paper develops a structural model of intraday price formation that embodies both public information shocks and microstructure effects. Due to its structural nature, the model's underlying parameters provide summary measures to assess trading costs, the sources of short-run price volatility, and the speed of price discovery in an internally consistent, unified setting. We estimate the model using transaction level data for a cross-section of NYSE stocks. We find, for example, that the parameter estimates jointly explain the observed U-shaped pattern in quoted bid-ask spreads and in price volatility, the magnitude of transaction price volatility due to market frictions, and the autocorrelation patterns of transaction returns and quote revisions. Further, in contrast to bid-ask spread patterns, we find that execution costs of a trade are much smaller than the spread and increase monotonically over the course of the day. This may provide an explanation for why there is concentration in trade at the open.

*University of Southern California; Wharton School, University of Pennsylvania; and JP Morgan Investment Management Inc., respectively. Part of this research was completed while Madhavan was visiting the New York Stock Exchange. We thank Joel Hasbrouck, Dennis Sheehan, James Shapiro and seminar participants at the Western Finance Association Meetings, IESE Barcelona, UC Berkeley, UCLA, Duke University, Hebrew University, Iowa University, University of Maryland, New York University, University of North Carolina, Ohio State University, Penn State University, University of Southern California, Tel Aviv University, University of Utah, and the Wharton finance lunch group for helpful comments. The comments and opinions contained in this paper are those of the authors and do not necessarily reflect those of the directors, members or officers of the New York Stock Exchange, Inc. Research support from the Geewax-Terker Research Fund (Madhavan and Richardson) and the University of Pennsylvania Research Foundation (Richardson) is gratefully acknowledged.

- We provide an estimator of execution costs that takes into account the possibility that orders may execute within the bid-ask spread as well as information and inventory effects.

- The model provides insights into the determinants of the autocorrelations of quotes and returns, as well as other moments, such as the variance of quote changes. Of particular interest, we compute these serial covariance moments without estimating them from the data, and instead rely on parameters estimated from other restrictions of the model. Thus, the implied values of these moments can be compared to their corresponding actual (i.e., sample) moments, helping us better understand the economic factors underlying stock price dynamics.

We estimate the model using transaction-level data for 274 stocks listed on the New York Stock Exchange (NYSE). The results suggest that trading accounts for a significant fraction of individual stock price volatility at the transaction level. While new information arrivals are concentrated in the early part of the day, the effect of microstructure frictions increases over the day, so that intraday volatility exhibits the familiar U-shaped pattern.

Interestingly, the adverse selection component of the bid-ask spread decreases steadily throughout the day. This decrease in information asymmetry is consistent with models of price discovery (see, Schreiber and Schwartz (1985), Handa and Schwartz (1991), and Madhavan (1992)) where market makers learn from order flow. However, market maker costs increase over the day (possibly reflecting the costs of carrying inventory overnight), so that bid-ask spreads exhibit the U-shaped pattern noted in previous research (see, for example, Harris (1986), Jain and Joh (1988) and McInish and Wood (1992).

We also find the cost of transacting is significantly smaller than the bid-ask spread once the probability of executing within the quotes is considered. In contrast to the bid-ask spread, this measure of execution costs increases over the day. This result is consistent with concentrated trading at the open by discretionary liquidity traders who can selectively time their trades. Additionally, the model estimates imply particular behavior in return and quote autocorrelations, which closely resemble the actual autocorrelations of the data. For example, in our sample, autocorrelations of price changes average approximately -.219 on a transactions level basis, while our model implies -.211.

Before describing how quotes and transaction prices are determined, we first discuss the evolution of public beliefs. Changes in beliefs arise from two sources: (i) New public information announcements which are not associated with trading, or (ii) Order flow, which may provide a noisy signal about future asset value.

Public news announcements may cause revisions in beliefs without any trading volumes. Denote by ϵ_t the innovation in beliefs between times $t-1$ and t due to new public information. We assume that ϵ_t is an independent and identically distributed random variable with mean zero and variance σ_ϵ^2. In addition, if market makers' believe that some traders may possess private information about fundamental asset value, a buy (sell) order is associated with an upward (downward) revision of beliefs. We assume, following Glosten and Milgrom (1985), that the revision in beliefs is positively correlated with the *innovation* in order flow.[4] Formally, the change in beliefs due to order flow is $\theta(x_t - E[x_t|x_{t-1}])$, where $(x_t - E[x_t|x_{t-1}])$ is the surprise in order flow and $\theta \geq 0$ measures the degree of information asymmetry or the so-called permanent impact of the order flow innovation. Higher values of θ indicate larger revisions for a given innovation in order flow; in the absence of information asymmetry, the parameter $\theta = 0$.[5]

Let μ_t denote the post-trade expected value of the stock conditional upon public information and the trade initiation variable. The revision in beliefs is the sum of the change in beliefs due to new public information and order flow innovations, so that

$$\mu_t = \mu_{t-1} + \theta(x_t - E[x_t|x_{t-1}]) + \epsilon_t. \tag{1}$$

Market maker bid and ask quotations are *ex post* rational (see, e.g., Glosten and Milgrom (1985)) so that the ask (bid) price is conditioned on a trade being buyer-initiated (seller-initiated). Let p_t^a denote the (pre-trade) ask price at time t and similarly define the bid price, p_t^b. Market maker quotations also reflect their compensation for their service in providing

[4]In Glosten and Milgrom (1985), the revision in beliefs is directly proportional to the actual order flow because it is implicitly assumed that order flow is uncorrelated.

[5]In some models (e.g., Glosten and Harris (1988) and Madhavan and Smidt (1991)), the revision in beliefs is proportional to the net order imbalance in a particular period. It is possible to extend the model to incorporate such volume effects by having multiple indicator variables for various order size ranges. However, since previous studies conclude that the effect of order size is economically small relative to the indicator variables, we retain our simplifying assumption regarding volume. This assumption allows us to estimate a parsimonious model and compute closed-form solutions for the estimators of interest.

uncorrelated and $\rho = 0$.

Next, we need to compute the conditional expectation of the trade initiation variable given public information. Observe that if $x_{t-1} = 0$, $E[x_t|x_{t-1}] = 0$. If $x_{t-1} = 1$, $E[x_t|x_{t-1} = 1] = \Pr[x_t = 1|x_{t-1} = 1] - \Pr[x_t = -1|x_{t-1} = 1] = \gamma - (1 - \gamma - \lambda) = \rho$. Similarly, If $x_{t-1} = -1$, $E[x_t|x_{t-1} = -1] = -\rho$. Thus, the conditional expectation $E[x_t|x_{t-1}] = \rho x_{t-1}$.

To transform equation (2) into a testable equation, we need to substitute out the unobservable prior belief, μ_{t-1}. Using the fact that $\mu_{t-1} = p_{t-1} - \phi x_{t-1} - \xi_{t-1}$ and $E[x_t|x_{t-1}] = \rho x_{t-1}$, we can express equation (2) as

$$p_t - p_{t-1} = (\phi + \theta)x_t - (\phi + \rho\theta)x_{t-1} + \epsilon_t + \xi_t - \xi_{t-1}. \tag{3}$$

Equation (3) forms the basis for our investigation of intraday price movements. In the absence of market frictions, the model reduces to the classical description of an efficient market where prices follow a random walk. However, in the presence of frictions (i.e., transaction costs and information asymmetries), transaction price movements reflect order flow and noise induced by price discreteness, as well as public information shocks.

2.2 Model Estimation

The four parameters governing the behavior of transaction prices and quotes are: (i) θ, the asymmetric information parameter, (ii) ϕ, the cost of supplying liquidity, (iii) λ, the probability a transaction takes place inside the spread, and (iv) ρ, the autocorrelation of the order flow. Let $\beta = (\theta, \phi, \lambda, \rho)$ denote the vector of price and quote parameters.

Equation (3) expresses transaction price changes as a linear function of contemporaneous and past order flows. Thus, with adjustments to the standard errors for serial covariance of the errors induced by price discreteness (i.e., $\xi_t - \xi_{t-1}$), equation (3) can be estimated via ordinary least squares. Unfortunately, not all of the parameters in the vector β can be identified this way. However, using a time-series of T observations on transaction price changes and trade initiation, the model's parameters can be estimated using maximum likelihood or a similar nonlinear estimation procedure. The drawback to this approach is that it requires strong distributional assumptions on the processes generating public information which may be far from reality.

3 Empirical Results

3.1 Data Sources and Methods

The data are drawn from a file of bid and ask quotations, transaction prices and volumes for equities in 1990, obtained from the Institute for the Study of Securities Markets (ISSM). Our initial sample is based on the first 750 stocks in the file. From the initial sample, we include only NYSE-listed common stocks trading in eighths which did not have any stock splits in the calendar year 1990.

As our objective is to understand the process by which security prices impound information, it is important to examine the evolution of the information parameters over the trading day.[9] Following Hasbrouck (1991b), we estimate the model for five intervals of the day: 9:30-10:00, 10:00-11:30, 11:30-2:00, 2:00-3:30, and 3:30-4:00. To ensure there are sufficient observations for model estimation, we consider only those stocks for which there are at least 250 observations per interval over the year 1990. These criteria reduce the sample to 274 stocks.

On a transaction basis, we impose filters on the data to eliminate possible recording errors.[10] For the opening period (9:30-10:00), overnight returns are eliminated since recent evidence (see, e.g., Amihud and Mendelson (1987)) indicates that they are likely to come from a different distribution. The opening transaction (which, in active stocks, is usually arranged in a batch or auction market), is eliminated.

To sign the trade initiation variable we must determine the quote prevailing at the time of the transaction. Although transactions and quotations in the ISSM data are time-stamped to the second, this presents a problem because there are often delays in the reporting of transactions resulting in a misaligned sequence of quotes and transactions. Thus, it is possible for a quote with the same time stamp as a transaction to represent the quote after the transaction. Accordingly, it has been common to use only those quotes which have been in

[9] This is especially relevant given empirical evidence (e.g., Harris (1989)) documenting temporal patterns in intraday returns and volatility.

[10] The filters are as follows: any trades below \$1 or above \$200 are excluded; and bid (ask) quote or transaction more than 50 percent away from the previous bid (ask) or transaction is eliminated; trades more than \$5 from the midquote are eliminated; and for stocks trading above \$10, any quote with a percentage spread above 20 percent is eliminated while for stocks below \$10 any quote implying spreads over \$2 is eliminated.

and transaction prices over the course of the day, as well as their respective magnitude, is suggestive of market frictions.

3.3 Parameter Estimates

Table 2 presents summary statistics on the individual parameter estimates governing the stochastic process for transaction price changes and quote revisions across the 274 stocks. The table presents the mean coefficient estimate, mean standard error, the standard deviation of the estimates and the median estimates of the parameter vector β for the 274 stocks in each of the five intraday trading intervals. The parameter estimates are, in general, highly significant and have economically reasonable values.[11]

The extent of information asymmetry (i.e., θ) is the main parameter of interest. From Table 2, it is clear that the degree of this asymmetry drops sharply after the opening half-hour interval. The mean value of θ falls by over a third from the opening to the middle of the day (from \$0.0415 to \$0.0275) and remains at this level until the final period where it increases slightly. With respect to formal tests, Table 2 also provides, for each stock, a Wald test that θ is equal over the course of the day. The average value of the Wald statistic over the sample of firms is 17.22. Since it is possible to show that the statistic follows a χ^2_4 distribution asymptotically, it is not surprising that 68.25% of the firms reject the restriction that θ does not change over the trading day.

The decline in θ has a clear economic interpretation. Recall that θ represents the magnitude of the revision in the market maker's beliefs concerning the security's value induced by order flow. A decline in θ, therefore, represents less reliance on the signal content of order flow. The greater reliance on prior beliefs is consistent with either (i) market makers learning about fundamental asset values (i.e., price discovery) through the trading process, or (ii) a larger percentage of liquidity traders (or equivalently, less information asymmetry) at the end of the day. However, the monotonicity in the parameter estimates suggests to us that the former interpretation is more reasonable. If this is the case, we would expect the average level of θ should decline over the week, a hypothesis that can be tested using the approach developed here. This may provide an explanation for the abnormal negative returns observed

[11]The constant α is close to zero and the estimates are not reported in the table.

The probability that a trade occurs within the quotes, λ, declines monotonically over the day. The mean estimates drop from 34 percent in the opening interval to 28 percent at the close. The joint Wald test rejects the equality restriction for 64.60% of the firms in the sample, with an average statistic value of 18.54. The steady decline in λ over the day may reflect the combined influence of (i) increased incentives to place limit orders when spreads are wide, (ii) a higher probability of a cross in intervals of high activity.

These parameter estimates are of interest not only in themselves but also because of their implications for price formation and the cost of trading. In particular, due to the structural nature of the model, these parameters will have implications for many intraday pricing phenomena. In the next section, we discuss and investigate the more interesting of these issues.

4 Applications of the Model

4.1 Application I: The Cost of Trading

4.1.1 The Bid-Ask Spread

The bid-ask spread is a common measure of transaction costs. In our model, the implied bid-ask spread at time t (i.e., $p_t^a - p_t^b$) is a random variable with mean $2(\theta + \phi)$. Let s denote the expected implied bid-ask spread. As s is a function of identifiable parameters, and the GMM estimators of these parameters have well-known asymptotic distributions, s can be estimated in a straightforward manner from the data. Specifically, it can be shown that the estimator of s, $2(\hat{\phi} + \hat{\theta})$, is consistent and asymptotically normal with variance

$$[2,\ 2] V_{\hat{\theta},\hat{\phi}} \begin{bmatrix} 2 \\ 2 \end{bmatrix},$$

where $\hat{\theta}$ and $\hat{\phi}$ denote the sample estimates of the parameters θ and ϕ and $V_{\hat{\theta},\hat{\phi}}$ is the estimated covariance matrix of the GMM estimators.[13]

Table 3 presents the mean coefficient estimate, mean standard error, the standard deviation of the coefficient estimate, and the median estimate of the bid-ask spread measure, as

[13]In our discussion below, we use this notation to describe all parameter estimates and associated covariance matrices.

bid-ask spread, we can also estimate the fraction of the implied spread attributable to asymmetric information. Define by r the ratio of the information component of the spread (i.e., 2θ) to the total implied spread. If $r = 0$, the spread is entirely attributable to the costs of supplying liquidity, and if $r = 1$, direct liquidity costs are negligible and adverse selection costs constitute the entire bid-ask spread. Then, the estimate of r has mean

$$\frac{\theta}{\phi + \theta},$$

and asymptotic variance

$$\left[\frac{\theta + \phi - 1}{(\theta + \phi)^2}, \frac{-\theta}{(\theta + \phi)^2}\right] V_{\hat{\theta},\hat{\phi}} \left[\begin{array}{c} \frac{\theta+\phi-1}{(\theta+\phi)^2} \\ \frac{-\theta}{(\theta+\phi)^2} \end{array}\right].$$

An examination of the proportion of the spread due to asymmetric information r over the course of the day supports this hypothesis for the U-shaped pattern in spreads. As shown in Table 3, r is 51 percent in the initial period, and falls steadily to 36 percent in the third period and remains at about this level for the rest of the day. Note that the average standard errors range between 1%-3%. These results then are strongly significant for the majority of firms (i.e., 91.54% of the sample imply Wald statistics which are significant at the 5% level).

4.1.2 The Effective Cost of Trading

An alternative measure of trading costs is the *effective bid-ask spread* measured by the expected price difference between a notional purchase at time t and a notional sale at some future time $t + k$. Recognizing the potential for a cross at either time, the potential changes are from ask to bid, ask to the midpoint, midpoint to bid, and midpoint to midpoint. If the notional sale takes place several transactions after the notional purchase (i.e., if k is sufficiently large), we can ignore the effect of the autocorrelation in order flow (which is of the order ρ^k). Accordingly, we assume that at the time of the notional sale market makers, on average, expect buys and sells to be equally likely. In this case, the round-trip expected costs associated with each of the four possibilities are $2(\phi+\theta)$, ϕ, $(\phi+\theta)$, and 0, respectively.

Observe that the round-trip costs from an ask to a midquote are not equal to the costs from midquote to bid. In the former case, the initial trade results in a permanent upward shock to price which improves the selling price, whereas in the latter case there is no such

Unlike the implied spread, the effective spread does not exhibit a U-shaped pattern. Indeed, the effective spread increases monotonically over the day, from \$0.073 in the opening interval to \$0.086 at the close. The average statistic across the firms is 21.81, with 71.90% of the firms displaying a statistically significant pattern. Thus, there is reasonable evidence to suggest that the U-shaped pattern in implied spreads does not carry through to the effective spread measure.

In fact, our result seems surprising because the implied spread S is actually highest in the opening period. This result reflects two factors. First, the effective spread takes into account the probability of execution within the quotes, and this probability decreases over the day (see Table 2). Second, the asymmetric information parameter is largest in the opening period, and this parameter has more impact on the implied spread than on the effective spread. This is because the effective spread takes into account the systematic tendency for prices to rise (fall) following a transaction at the ask (bid). In contrast, the market maker's transaction costs increase over the day, and this has relatively more impact on the effective spread.

The fact that the effective spread is smallest at the open has interesting implications for theoretical models (e.g., Admati and Pfleiderer (1988)) which predict that trading should concentrate at certain periods of the day. Our results suggest a natural explanation for why this concentration should occur at the beginning of the day rather than at other times. Intuitively, discretionary liquidity traders will migrate to periods where their effective costs of trading are lowest. Discretionary traders find it cheapest to trade at the start of the day, even though the degree of information asymmetry is highest; the high value of θ is balanced against the market maker's transaction costs which increase over the course of the day. In turn, the concentration of trading by such traders increases the probability that a transaction occurs within the spread, which also sustains a small effective spread.

4.2 Application II: The Determinants of Price Volatility

The model can be used to decompose transaction price volatility into its components. Using equation (3), the variance of stock price changes is

$$\mathrm{Var}[\Delta p_t] = \sigma_\epsilon^2 + 2\sigma_\xi^2 + (1-\lambda)[(\theta+\phi)^2 + (\theta\rho+\phi)^2 - 2(\theta+\phi)(\theta\rho+\phi)\rho]. \tag{4}$$

moments

$$E\begin{pmatrix} (u_t-\alpha)^2-(\sigma_\epsilon^2+2\sigma_\xi^2) \\ (u_t-\alpha)(u_{t-1}-\alpha)+\sigma_\xi^2 \end{pmatrix}=0,$$

where u_t is defined as $\Delta p_t-(\theta+\phi)x_t+(\theta\rho+\phi)x_{t-1}$ and α is the estimated drift term.[15]

We can further decompose π into four parts: (i) the effect of price discreteness, (ii) the asymmetric information effect, (iii) the trading cost effect, and (iv) the interaction between these effects, as measured by the ratio of the terms $2\sigma_\xi^2$, A, B, and C, respectively, to the variance of price changes. These measures allow us to assess the relative contributions of the microstructure frictions to price volatility.

Table 4 provides summary statistics on the individual parameter estimates and the percentage of the variance in price changes attributable to (i) public information, (ii) price discreteness, (iii) asymmetric information, (iv) trading costs, and (v) the interaction between the asymmetric information and cost components.

For each component, the table displays the mean coefficient estimate, the mean standard error of the estimate, the standard deviation of the coefficient estimates, and the median estimate over the 274 stocks, by time of day.

The public information component of volatility, σ_ϵ^2, declines by about a third over the day. The decline is monotonic except for the last half hour interval where there is a small increase in the variance. Although these results are highly significant (as measured by the mean standard error of the estimate), there is a wide range in the estimates of σ_ϵ^2 (measured by the standard deviation of the parameter estimates) across stocks. As the variance decreases, the fraction of variance attributable to market frictions (i.e., π) increases steadily from 54 percent at the open to 65 percent at the close. This result is consistent with evidence provided by French and Roll (1986) who find that prices are more variable during trading hours than during non-trading hours. This can be explained if public information events are more likely to occur during business hours or if the process of trading creates volatility. Our estimates suggest that both public information shocks and order flow are important sources of intraday volatility, but that the relative importance of public information declines over the day.

The decline in σ_ϵ^2 over the day may reflect more frequent occurrences of public information

[15] The estimates of the other parameters are unaffected by the addition of these moment conditions.

(which in fact exist), this will no longer be the case. Given the estimated frictions (i.e., $\hat{\theta}$, $\hat{\phi}$,...), what are the implications for time variation in transaction returns and quote revisions?

4.3.1 Transaction Prices

Using equation (3), we obtain

$$\text{Cov}(\Delta p_t, \Delta p_{t-1}) = -\sigma_\xi^2 + \rho(1-\lambda)[(\theta+\phi)^2 + (\theta\rho+\phi)^2] - (1-\lambda)(\theta\rho+\phi)(\theta+\phi)(1+\rho^2). \quad (6)$$

Simplifying this expression, we can show that $\text{Cov}(\Delta p_t, \Delta p_{t-1}) < 0$. Stock price changes are negatively autocorrelated if there are costs to providing liquidity (i.e., $\phi > 0$) or if there are rounding errors induced by price discreteness (i.e., σ_ξ^2), as these frictions generate bid-ask bounce. Larger frictions and greater information asymmetry increase the absolute magnitude of the serial covariance term. The absolute size of the covariance term is a decreasing function of the probability of executing within the spread, λ, because this mitigates the bid-ask bounce. The autocorrelation in order flow, however, has an ambiguous effect on the serial covariance term.

The above theoretical results provide explicit representations for the serial covariances of transaction price changes. It is interesting to relate the implied correlations with the actual correlations present in the data to see how well the model explains the short-horizon time-variation of returns. Table 5 presents, for the five intraday time intervals, the mean implied autocorrelation (from the model and estimated parameters), and the mean sample autocorrelation of transaction price changes for the 274 stocks in the sample. The implied estimates are computed on a stock-by-stock basis using equations (6) and (8). The table also reports the standard deviation of the actual estimates and the standard deviation of the difference between the actual and implied estimates.

There is a close correspondence between the actual and implied autocorrelation of transaction price changes, especially after the 9:30-10:00 period. For example, the implied and sample autocorrelations after this period equal (-.2026,-.2501,-.2588,-.2525) and (-.2166,-.2433,-.2484,-.2197) respectively. This is especially interesting because the autocorrelation moments were not used to estimate the underlying parameters of the model. Thus, the similarity between the actual and implied estimates suggests that the model is well-specified, order flow and the relative magnitudes of the parameters.[16] Intuitively, if order flow is positively correlated, successive transactions at the bid or the ask are more likely than reversals. Market makers take this effect into account in forming their beliefs, so that the expected revision in beliefs is zero. However, successive transactions at the bid or the ask will still lead to quote revisions unless they are fully anticipated, and this creates positive serial correlation in ask price revisions. The larger the information asymmetry component and the lower the probability of a cross within the quotes, the stronger this effect. By contrast, if the autocorrelation in order flow is negative, the covariance term in equation (8) is unambiguously negative, because reversals are more likely than continuations and quotes are revised in the direction of order flow. Thus, ask price changes may contain important information about the underlying determinants of security price movements.

Table 5 shows that the actual serial correlation of successive ask price revisions is negative, but is small in magnitude. The model's implied serial correlation of quote revisions is of similar magnitude. However, our implied estimates are, on average, closer to zero than the actual estimates and become more negative over the day, so that the deviations become steadily smaller over the day. Recall that the theoretical autocovariance of ask price revisions is negative if the variance in the rounding error, σ_ξ^2 is sufficiently large. The fact that our implied estimates are, on average, slightly larger than the actual estimates, suggests that the variance of the rounding error due to price discreteness is underestimated, especially in the early part of the day. Alternatively, this finding may reflect possible autocorrelation of these errors.

Table 5 also provides a comparison between the intraday patterns in the sample variance of ask price changes and the implied variance of ask price changes (from the model). Of particular interest, the actual variances and the implied variances both drop after the beginning of the day, and level off around midday. For example, the implied variances drop from .0042 to .0035, and then level off at .0031. Similarly, the sample variances drop from .0061 to .0048 and then level off at .0040. To the extent that the implied variances of ask price changes are calculated using parameters estimated from transactions data, the similarity in

[16] This is consistent with recent empirical studies which find ask-to-ask returns and mid-quote returns do not follow martingales. See, e.g., Hasbrouck and Ho (1987) and Handa (1991). The autocorrelation may also represent time-varying expected returns as in Conrad, Kaul, and Nimalendran (1991).

significantly smaller than the cost implied by the bid-ask spread. Interestingly, our measure of execution costs increases over the day as the probability of trading within the quotes declines monotonically. Thus, although the opening period is associated with the greatest information asymmetry and widest quoted spreads, it is in reality the least costly period in which to trade. This result may explain the observed pattern of trade concentration early in the day, but may also be a consequence of this phenomenon.

On a more general note, we believe there is a definitive advantage of the structural approach to modeling intraday price formation. Given the structural model and its underlying economics, the parameters of that model imply microstructure price behavior in a unified framework. Thus, trading costs, bid-ask spreads, time variation, and volatility patterns can all be examined in an internally consistent environment. Our approach has the additional advantage that we use only transactions data and a given set of model restrictions to estimate these phenomena. In this paper, we find that, with a relatively simple structural model, many stylized facts about intraday price dynamics can be jointly explained.

The methodology suggests several avenues of future research. First, on a structural basis, it would be interesting to extend the model to incorporate volume. This issue is perhaps not so relevant for the frequently traded stocks in our sample where almost all the trades take place either at the quotes or within the quotes for this data. For series involving inactive securities, however, large trades may have a substantial effect on prices (see, e.g., Keim and Madhavan (1993)) and this issue may be very important. One way to incorporate trade size is to allow θ to vary over different trade sizes (e.g., θ_1 for small trades, θ_2 for medium-size trades, etc.) or be a function of trade size. Another structural extension is to explicitly model the intraday evolution of the parameters as a function of the time of day. Here, we approximate the within-day variation in parameters by allowing discrete jumps at various times of the day, but an alternative may be to allow the parameters to exponentially decay as a function of elapsed time. Second, on an empirical basis, a cross-sectional analysis of the parameters and resulting model implications is of interest. For example, what fundamentals determine the rate of price discovery for a particular firm and the level of information asymmetry? Is the temporal variation in the model's parameters very different for internationally cross-listed stocks (where there may be no overnight non-trading period), for stocks that are often used in

TABLE 1

Descriptive Statistics for the Sample of Stocks (1990)

Panel A provides summary statistics on the variance of transaction price changes, variance of ask price changes, average number of transactions per day, market capitalization, and price for 274 NYSE-listed stocks in 1990. Panel B provides mean estimates of the variance of transaction and ask price changes, number of transactions per hour, the mean hourly volume (in round lots of 100 shares), and the dollar spread for five time intervals during the day.

Panel A	Mean	Std.Dev.	75%	Median	25%
Variance of ΔP	.0067	.0025	.0079	.0062	.0051
Variance of ΔP^{ask}	.0044	.0030	.0059	.0037	.0025
Transactions/Day	95	86	107	66	44
Market Cap. ($ bn.)	4.36	6.95	4.42	2.21	1.02
Price ($)	38.85	21.82	49.13	36.63	22.25
Panel B	**9:30-10:00**	**10:00-11:30**	**11:30-2:00**	**2:00-3:30**	**3:30-4:00**
Variance of ΔP	.0073	.0068	.0065	.0066	.0070
Variance of ΔP^{ask}	.0061	.0048	.0042	.0040	.0040
Transactions/hour	17	16	12	13	17
Volume/hour (100s)	385.8	357.4	235.6	252.3	272.5
Spread ($)	.228	.211	.204	.205	.210

TABLE 3
Summary Statistics of Estimated Trading Costs

Table 3 presents summary statistics of estimates of trading costs for 274 NYSE-listed stocks in the 1990 sample period over five intraday trading intervals. Specifically, the mean coefficient estimate across the stocks, the mean standard error of the mean estimates, the standard deviation of the estimates across the 274 stocks, and the median estimate are provided for various parameters of interest: s, the implied spread; r, the fraction of the implied spread attributable to asymmetric information; s^E, the effective bid-ask spread; and r^E, the ratio of the effective to the implied spread. Joint tests of the restriction that each parameter is constant throughout the day are also given. A Wald test of this restriction, i.e., $s_{9:30-10:00} = s_{10:00-11:30} = \ldots = s_{3:30-4:00}$, follows an asymptotic χ^2_4 distribution. The estimates of these trading costs parameters are based on the GMM coefficient estimates described in Table 2. The final column of the table reports the mean and median value of the χ^2_4 statistic and the percentage of stocks for which the null hypothesis is rejected at the 5 percent level.

	9:30-10:00	10:00-11:30	11:30-2:00	2:00-3:30	3:30-4:00	Joint Tests
s						
Mean	0.1518	0.1440	0.1425	0.1448	0.1496	Av. χ^2_4 - 14.25
(Av. Std.Er.)	(0.0066)	(0.0027)	(0.0024)	(0.0029)	(0.0048)	Median χ^2_4 - 9.77
Std. Dev.	0.0331	0.0252	0.0233	0.0238	0.0246	% Reject (5%) - 51.46%
Median	0.1467	0.1389	0.1380	0.1419	0.1461	
s^E						
Mean	0.0728	0.0773	0.0814	0.0834	0.0864	Av. χ^2_4 - 21.81
(Av. Std.Er.)	(0.0050)	(0.0022)	(0.0019)	(0.0024)	(0.0040)	Median χ^2_4 - 16.42
Std. Dev.	0.0142	0.0129	0.0125	0.0125	0.0123	% Reject (5%) - 71.90%
Median	0.0735	0.0768	0.0808	0.0838	0.0863	
r						
Mean	0.5107	0.4149	0.3630	0.3553	0.3601	Av. χ^2_4 - 58.43
(Av. Std.Er.)	(0.0378)	(0.0167)	(0.0138)	(0.0165)	(0.0270)	Median χ^2_4 - 37.00
Std. Dev.	0.2527	0.2153	0.1977	0.1943	0.1994	% Reject (5%) - 91.54%
Median	0.4812	0.3923	0.3345	0.3302	0.3210	
r^E						
Mean	0.5019	0.5552	0.5888	0.5927	0.5947	Av. χ^2_4 - 11.13
(Av. Std.Er.)	(0.0469)	(0.0220)	(0.0196)	(0.0240)	(0.0381)	Median χ^2_4 - 8.21
Std. Dev.	0.1435	0.1414	0.1409	0.1375	0.1360	% Reject (5%) - 43.06%
Median	0.4975	0.5392	0.5747	0.5842	0.5899	

TABLE 5
Actual and Implied Moments of Price Changes and Quote Revisions

Table 5 presents summary statistics for the difference between actual and implied moments of price changes and quote revisions for the 274 NYSE-listed stocks in the 1990 sample period over five intraday trading intervals. The table presents the mean estimate of both the implied moment (from the model) and corresponding sample moment, the standard deviation of the sample moment across the 274 stocks, and the standard deviation of the difference between the implied and sample moment across the stocks. The particular moments computed are: $corr(\Delta P_t, \Delta P_{t-1})$, the autocorrelation of price changes; $corr(\Delta P_t^{ask}, \Delta P_{t-1}^{ask})$, the autocorrelation of ask price changes; and $var(\Delta P_t^{ask})$, the variance of ask price changes.

	9:30-10:00	10:00-11:30	11:30-2:00	2:00-3:30	3:30-4:00
$corr(\Delta P_t, \Delta P_{t-1})$					
Implied	-0.0972	-0.2026	-0.2501	-0.2588	-0.2525
Sample	-0.1719	-0.2166	-0.2433	-0.2484	-0.2197
Std. Dev.	0.1168	0.1114	0.1069	0.1054	0.1139
Std. Dev. of Diff.	0.1831	0.0975	0.0875	0.0971	0.1278
$corr(\Delta P_t^{ask}, \Delta P_{t-1}^{ask})$					
Implied	0.0117	-0.0123	-0.0285	-0.0333	-0.0315
Sample	-0.0358	-0.0573	-0.0626	-0.0597	-0.0458
Std. Dev.	0.0659	0.0464	0.0401	0.0428	0.0542
Std. Dev. of Diff.	0.0648	0.0512	0.0444	0.0513	0.0842
$var(\Delta P_t^{ask})$					
Implied	0.0042	0.0035	0.0031	0.0031	0.0032
Sample	0.0061	0.0048	0.0042	0.0040	0.0040
Std. Dev.	0.0028	0.0025	0.0022	0.0022	0.0022
Std. Dev. of Diff.	0.0017	0.0020	0.0022	0.0022	0.0022

Glosten, L., and L. Harris, 1988, "Estimating the Components of the Bid-Ask Spread," *Journal of Financial Economics*, 21, 123–142.

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

Handa, P., 1991, "Order flow and Bid-Ask Dynamics: An Empirical Investigation," Working paper, New York University.

Handa, P., and R. Schwartz, 1991, "The Dynamics of Price Discovery in a Securities Market," Working paper, New York University.

Hansen, L., 1982, "Large Sample Properties of Generalized Method of Moment Estimators," *Econometrica*, 50, 1029–1084.

Harris, L., 1986, "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns," *Journal of Financial Economics*, 16, 99–118.

Harris, L., 1989, "A Day-End Transaction Price Anomaly," *Journal of Financial and Quantitative Analysis*, 24, 29–46.

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

Hasbrouck, J., 1991a, "Measuring the Information Content of Stock Trades," *Journal of Finance*, 46, 178–208.

Hasbrouck, J., 1991b, "The Summary Informativeness of Stock Trades: An Econometric Analysis," *Review of Financial Studies*, 4(3), 571–595.

Hasbrouck, J., and T. Ho, 1987, "Order Arrival, Quote Behavior, and the Return Generating Process," *Journal of Finance*, 42, 1035–1048.

Hasbrouck, J., and G. Sofianos, 1992, "The Trades of Market-Makers: An Analysis of NYSE Specialists," Working paper, New York Stock Exchange.

Ho, T., and H. Stoll, 1983, "The Dynamics of Dealer Markets under Competition," *Journal of Finance*, 38, 1053–1074.

Stoll, H., 1989, "Inferring the Components of the Bid-Ask Spread: Theory and Empirical Tests," *Journal of Finance*, 44, 115–134.