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

Ananth Madhavan
University of Southern California

Matthew Richardson
New York University and NBER

Mark Roomans
JP Morgan Investment Management Inc.

This article develops and tests a structural model of intraday price formation that embodies public information shocks and microstructure effects. We use the model to analyze intraday patterns in bid-ask spreads, price volatility, transaction costs, and return and quote autocorrelations, and to construct metrics for price discovery and effective trading costs. Information asymmetry and uncertainty over fundamentals decrease over the day, although transaction costs increase. The results help explain the U-shaped pattern in intraday bid-ask spreads and volatility, and are also consistent with the intraday decline in the variance of ask price changes.

Part of this research was completed while A. Madhavan was visiting the New York Stock Exchange. We thank Larry Harris, Joel Hasbrouck, Dennis Sheehan, James Shapiro, Ingrid Werner, and an anonymous referee for their suggestions. Seminar participants at the Western Finance Association Meetings, the High Frequency Data Conference (Zurich), IESE Barcelona, UC Berkeley, UCLA, Duke University, Hebrew University, INSEAD, University of Iowa, London School of Economics, 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 Wharton provided many helpful comments. Research support from the Geewax-Terker Research Fund (Madhavan and Richardson) and the University of Pennsylvania Research Foundation (Richardson) is gratefully acknowledged. Address correspondence to Ananth Madhavan, Marshall School of Business, University of Southern California, Los Angeles, CA 90089-1421.

The Review of Financial Studies Winter 1997 Vol. 10, No. 4, pp. 1035–1064

Why do security prices change?

In the classical model of an efficient security market, prices move in response to new public information that causes traders to simultaneously revise their beliefs. Alternatively the process of trading itself may generate price movements because of various market imperfections and frictions. A realistic description of intraday security price movements must capture both of these elements.

Understanding the process of intraday price formation is important for several reasons. First, intraday prices and quotes exhibit several patterns that are not well understood. For example, it is well known [see, e.g., Harris (1986), Jain and Joh (1988), and McInish and Wood (1992)] that quoted bid-ask spreads and volume exhibit U-shaped patterns over the day. This finding is difficult to reconcile with evidence from theoretical models [Easley and O'Hara (1992), Madhavan (1992)] and laboratory experiments [Bloomfield (1996), Bloomfield and O'Hara (1996)] where information asymmetry and uncertainty over fundamentals, and hence bid-ask spreads, decline monotonically over the day as market participants learn from the trading process. Second, an examination of how much new public information flows or particular market frictions contribute to intraday price volatility is of great interest from a public policy viewpoint. Third, a better understanding of intraday price formation may also shed light on the magnitude, determinants, and composition of execution costs, a topic of considerable interest to portfolio managers, exchange officials, and traders.

Previous research has examined various microstructure phenomena (such as intraday bid-ask spreads, execution costs, and autocorrelation and volatility patterns of transaction prices and quotes) either individually or through generalized reduced-form investigations. By contrast, this article develops and estimates a structural model of price formation that captures many of these frictions in a unified setting. Our model incorporates both public information shocks and microstructure effects. The microstructure phenomena modeled include the possibility of crosses within the quoted bid-ask spread, autocorrelation of order flow, as well as trading frictions arising from asymmetric information, dealer costs, and price discreteness.

The structural model described in this article links two important areas of the literature. The first area examines the sources of intraday price volatility. While theoretical and laboratory experiments indicate that the trading process generates information that reduces pricing errors, the empirical evidence suggests otherwise. For example, French and Roll (1986) find that return volatility is significantly higher during trading hours than during nontrading hours, and attribute this to noise generated by the trading process. More recently, Hasbrouck

(1991) uses time-series techniques to decompose the random-walk component of returns using a VAR approach, while Hasbrouck (1993) decomposes the variance of the stationary component of returns. Like Hasbrouck, we decompose intraday volatility into components attributable to public information shocks and trading frictions, but our work differs in that our model is structural by nature. Thus we can relate the decomposition one-to-one with the underlying economic parameters of the model, although we lose the generality of Hasbrouck's approach.

The second related area concerns the measurement of the components of the bid-ask spread. An extensive theoretical literature shows that the costs of trading, as represented by the effective bid-ask spread, consist of three components: asymmetric information, inventory carrying costs, and order processing costs. The relative importance of these components is a topic of considerable practical and academic interest, especially since it may provide insights into the why execution costs vary across markets and stocks [see, e.g., Choi, Salandro, and Shastri (1988), George, Kaul, and Nimalendran (1991), Glosten (1987), Glosten and Harris (1988), Huang and Stoll (1994), Roll (1984), and Stoll (1989)]. More recently Huang and Stoll (1996) develop a general model to investigate the components of the bid-ask spreads. Although their model (which was independently derived) is closely related to ours, it decomposes the noninformation part of the spread into the inventory and order processing components. By contrast, our focus is more on explaining the effect of information flows on stock prices over the day, that is, complementary. In this sense, our work is also related to Huang and Stoll (1994), who examine the predictability of returns over very short horizons while controlling for various microstructure effects.

We estimate the model using transaction-level data for a number of stocks listed on the New York Stock Exchange (NYSE). Because the model is so simple, the parameters are estimated in a setting that provides a high comfort level in terms of estimation error, that is, OLS normal equations plus some augmented direct estimates of order flow behavior and crossing probabilities.

Several interesting results emerge from our analysis. First, both information flows and trading frictions are important factors in explaining intraday price volatility in individual stocks. Second, information asymmetry decreases steadily throughout the day, consistent with theoretical models [Handa and Schwartz (1991) and Madhavan (1992)] where market makers learn from order flow, as well as with evidence from experimental markets [Bloomfield (1996), Bloomfield and O'Hara (1996)]. However, dealer 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. Third, 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 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.

Finally, 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. In particular the autocorrelation of quote returns implied by our model may be positive or negative, even though beliefs follow a martingale and there are no lagged inventory effects. Additionally the pattern in return and quote autocorrelations implied by our model closely resemble the actual autocorrelations of the data. For example, in our sample, autocorrelations of price changes average approximately -0.219 on a transactions level basis, while our model implies an average of -0.211.

The article proceeds as follows: In Section 1, we develop a structural model of price formation, and describe the procedure for estimating this model. Section 2 provides a brief description of the data and estimates of the model's underlying parameters over the day. Section 3 discusses the economic implications of the model's parameters for execution costs, price discovery, and the autocorrelation of price changes and quote revisions. Finally, Section 4 summarizes our findings and describes some possible extensions to the model for future research.

1. A Structural Model of Price Formation

1.1 Bid-ask quotes and transaction prices

We begin by examining a prototypical microstructure model of the quote and return generating mechanism.[1] Consider the market for a risky security whose fundamental value (which can be thought of as the present value of future dividends) evolves through time. The security is traded in an auction-dealer mechanism where liquidity providers (who may be traders using limit orders or designated market makers such as NYSE specialists) quote bid and ask prices at which

[1] The model incorporates a variety of microstructure effects discussed individually. In particular, the models of Choi, Salandro, and Shastri (1988), Garbade and Silber (1979), Glosten and Milgrom (1985), Roll (1984), and Stoll (1989) can be viewed as special cases of our model.

they are willing to trade. An order also may be executed within the quotes.[2]

Let p_t denote the transaction price of the security at time t. Denote by x_t an indicator variable for trade initiation, where $x_t = +1$ if trade t is buyer initiated and -1 if the trade is seller initiated. Some trades (such as prenegotiated crosses within the prevailing bid-ask spread) can be viewed as both buyer and seller initiated, and in this case $x_t = 0$. Let λ denote the unconditional probability the transaction occurs within the quoted spread, that is, $\lambda = \Pr[x_t = 0]$. We assume that buys and sells are (unconditionally) equally likely, so that $E[x_t] = 0$ and $\text{var}[x_t] = (1 - \lambda)$.

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, and (ii) order flow, which may provide a noisy signal about future asset values. 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 the order flow.[3] Formally the change in beliefs due to order flow is $\theta(x_t - E[x_t \mid x_{t-1}])$, where $(x_t - E[x_t \mid x_{t-1}])$ 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.

Our assumption of a fixed order size is consistent with much of the previous literature [see, e.g., Choi, Salandro, and Shastri (1988), George, Kaul, and Nimalendran (1991), Glosten and Milgrom (1985), Huang and Stoll (1994), Roll (1984), and Stoll (1989)]. Alternatively, following Glosten and Harris (1988) and Madhavan and Smidt (1991), the revision in beliefs could be modeled as proportional to the net or-

[2] On the NYSE, such price improvement may occur because of limit orders or through the actions of floor traders or the specialist. Although we do not explicitly model the limit order book, we can interpret a limit order trader as another market maker. To this extent the bid and ask quotes may be thought of as arising not necessarily from the exchange floor. Such an approach is pursued by Greene (1996), who extends our model in this direction.

[3] 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.

der imbalance in a particular period. Although it is possible to extend the model to incorporate such volume effects (as discussed below), there are several arguments in favor of a constant order size model. First, and most importantly, the simplifying assumptions regarding volume permit us to estimate a parsimonious model and compute closed-form solutions for the estimators of interest. Second, large-block trades often originate in the upstairs market [see, e.g., Keim and Madhavan (1996), Madhavan and Cheng (1997) and Seppi (1990)] where intermediaries or block brokers search for counter-parties to the trade to cushion the price impact of the trade. Further, if the floor market is truly anonymous, informed traders will break up their large trades. These factors suggest that trade direction may have more explanatory power than trade size. Third, the assumption facilitates comparisons of our results with much of the previous literature on intraday price movements.

Let μ_t denote the posttrade 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

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

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 (pretrade) 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 liquidity on demand. Let $\phi \geq 0$ represent market-makers' cost per share for supplying liquidity. We can interpret ϕ as dealers' compensation for transaction costs, inventory costs, risk bearing, and possibly the return to their unique position. It follows that the ask price (i.e., the price conditional upon $x_t = +1$) is $p_t^a = \mu_{t-1} + \theta(1 - E[x_t|x_{t-1}]) + \phi + \epsilon_t$. Similarly, the bid price is $p_t^b = \mu_{t-1} - \theta(1 + E[x_t|x_{t-1}]) - \phi + \epsilon_t$. Thus ϕ captures the temporary (or transitory) effect of order flow on prices. Not all orders are executed at the quoted bid or ask prices. For example, floor brokers or upstairs market makers may negotiate trades within the bid-ask spread. In this case, transactions are assumed to execute at the midquote, that is, $(p_t^a + p_t^b)/2$.

The transaction price can then expressed as

$$p_t = \mu_t + \phi x_t + \xi_t, \tag{2}$$

where ξ_t is an independent and identically distributed random variable with mean zero. The term ξ_t captures the effect of stochastic rounding errors induced by price discreteness or possibly time-varying returns.[4] Observe that any systematic deviation from zero in the rounding error (arising perhaps from a tendency to round up on buys and down on sells) is captured in the dealer cost component ϕ.

In interpreting Equation (2), note that μ_t is the conditional mean given that x_t is observed. Even though market makers fix their quotes before the trade is observed, they can condition on the incoming trade's sign because buys execute at the ask price while sells execute at the bid. Under rational expectations, dealers set regret-free prices, recognizing the effect of trades on their ex post beliefs. In setting such ex post rational prices, the dealer must also incorporate the transaction cost ϕ into the quotes. For this reason, the transaction price at time t already impounds the effect of the trade direction on the incoming trade.

Using Equations (1) and (2), we obtain

$$p_t = \mu_{t-1} + \theta(x_t - E[x_t|x_{t-1}]) + \phi x_t + \epsilon_t + \xi_t. \tag{3}$$

To estimate Equation (3), we must describe the temporal behavior of order flow. We assume a general Markov process for the trade initiation variable. Let γ denote the probability that a transaction at the ask (bid) follows a transaction at the ask (bid), that is, $\gamma = \Pr[x_t = x_{t-1}|x_{t-1} \neq 0]$. As large traders typically break up their orders into smaller components for easier execution, continuations are more likely than reversals, so that $\gamma > \frac{1}{2}$. This may also occur because of price continuity rules, trade reporting practices, and other institutional factors.

Let ρ denote the first-order autocorrelation of the trade initiation variable, that is, $\rho = \frac{E[x_t x_{t-1}]}{\mathrm{Var}[x_{t-1}]}$. It is straightforward to prove that $\rho = 2\gamma - (1 - \lambda)$, so that the autocorrelation of order flow is an increasing function of the probability of a continuation, γ, and the probability of a midquote execution, λ. Observe that when there is no possibility of transacting within the quotes (i.e., $\lambda = 0$) and order flow is independent (i.e., $\gamma = \frac{1}{2}$), order flow is serially uncorrelated, and $\rho = 0$.[5]

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 $E[x_t|x_{t-1}] = \rho x_{t-1}$.

[4] The assumption that the rounding error is serially uncorrelated is made for simplicity; since the rounding error is, on average, one-sixteenth of a dollar (i.e., half the minimum variation), the effects on our estimates of any serial correlation are unlikely to be significant.

[5] In our model, midquote transactions are informative because these transactions lead to innovations in market-maker beliefs when order flow is autocorrelated.

To transform Equation (3) 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 (3) as

$$p_t - p_{t-1} = (\phi + \theta)x_t - (\phi + \rho\theta)x_{t-1} + \epsilon_t + \xi_t - \xi_{t-1}. \tag{4}$$

Equation (4) 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.

1.2 Model estimation

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

Equation (4) expresses transaction price changes as a linear function of contemporaneous and past order flows. We use the generalized method of moments (GMM) to estimate the model. This technique is more appropriate than alternatives such as maximum likelihood because it does not require strong assumptions for the stochastic process generating the data and it allows for adjustment for general forms of autocorrelation and conditional heteroskedasticity. Indeed, Huang and Stoll (1994, 1996) use GMM to test their models which closely resemble ours.

The GMM procedure consists of choosing parameter values for β that minimize a criterion function based on the orthogonality restrictions (or moment conditions) implied by the model. Because of the model's implied linearity between the observable variables, the moment conditions correspond to the normal equations of OLS, with some additional restrictions which help decompose the OLS coefficients. Specifically, let $u_t = p_t - p_{t-1} - (\phi + \theta)x_t + (\phi + \rho\theta)x_{t-1}$. Then the following population moments implied by the model exactly identify the parameter vector β and a constant (drift) α:

$$E \begin{pmatrix} x_t x_{t-1} - x_t^2 \rho \\ |x_t| - (1 - \lambda) \\ u_t - \alpha \\ (u_t - \alpha)x_t \\ (u_t - \alpha)x_{t-1} \end{pmatrix} = 0. \tag{5}$$

The first equation is simply the definition of the autocorrelation in trade initiation, the second equation defines the crossing probability, the third equation defines the drift term as the average pricing error, and the last two equations are the OLS normal equations. Although the crossing probability λ can be estimated independently of the structural model, it does affect our derived measures of transaction costs.

The idea behind GMM is to choose parameter values for β such that the sample moments, denoted by $g_T(\beta)$, closely approximate these underlying population moments. Hansen (1982) proves that the GMM estimates, $\hat{\beta}$, are consistent and asymptotically normally distributed. In particular, the variance-covariance matrix of $\hat{\beta}$ equals $V_{\hat{\beta}} = [D_0^{'} S_0^{-1} D_0]^{-1}$, where $D_0 = E[\frac{\partial g_T(\beta)}{\partial \beta}]$ and $S_0 = \sum_{l=-\infty}^{l=+\infty} E[f_t f_{t-l}]$, where f_t is the vector of arguments in the expectation that defines the moment conditions. In practice, we estimate S_0 using the Newey–West procedure to obtain a heteroskedasticity consistent covariance matrix.

2. Empirical Results

2.1 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 that did not have any stock splits in the calendar year 1990. 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 outliers or recording errors.[6] 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.

Although our estimation uses only transaction data, we use quotation data to sign the trade initiation variable. Quotes and transac-

[6] The filters are as follows: any trades below $1 or above $200 are excluded; and a bid (ask) quote or transaction more than 50% 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 greater than 20% is eliminated while for stocks below $10 any quote implying spreads over $2 is eliminated.

Table 1
Descriptive statistics for the sample of stocks (1990)

Panel A

	Mean	Std.Dev.	75%	Median	25%
Variance of ΔP	0.0067	0.0025	0.0079	0.0062	0.0051
Variance of ΔP^{ask}	0.0044	0.0030	0.0059	0.0037	.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	0.0073	0.0068	0.0065	0.0066	0.0070
Variance of ΔP^{ask}	0.0061	0.0048	0.0042	0.0040	0.0040
Transactions/hour	17	16	12	13	17
Volume/hour (100s)	385.8	357.4	235.6	252.3	272.5
Volume/transaction	22.7	22.3	19.6	19.4	16.0
Spread ($)	0.228	0.211	0.204	0.205	0.210

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), the volume per transaction (in round lots of 100 shares), and the dollar spread for five time intervals during the day.

tions, although time stamped to the second, may be systematically misaligned because different reporting systems are used on the floor, as shown by Lee and Ready (1991). We adopt the Lee and Ready procedure to determine trade initiation using the 16-second time lag recommended by Blume and Goldstein (1992) to determine the quote prevailing at the time of the transaction. Finally, because dealers may provide multiple quotations even if the security's price is unchanged, we measure quote revisions in transaction time.

2.2 Descriptive statistics

Table 1 presents descriptive statistics for the sample of stocks. Panel A of the table provides details on the variance of transaction price and quote changes, number of transactions per day, market capitalization at year-end 1989, and stock prices for the 274 stocks in the sample. The stocks are actively traded (the median number of transactions per day is 66), and have a relatively high level of market capitalization (the median size is $2.2 billion which corresponds roughly to the average size of NYSE-listed firms). Nonetheless, the stocks exhibit a wide range in terms of size and activity, and all variables are skewed to the right.

Panel B of Table 1 provides a breakdown of price volatility, transaction frequency, bid-ask spreads, and share volume in five intraday time intervals for 1990. On an hourly basis, it is clear that bid-ask spreads, share volume, and trading frequency exhibit the U-shaped pattern documented by previous research. For example, while the

Table 2
Summary statistics of GMM model parameter estimates

	9:30–10:00	10:00–11:30	11:30–2:00	2:00–3:30	3:30–4:00
θ					
Mean	0.0415	0.0318	0.0275	0.0274	0.0287
(Avg. Std.Er.)	(0.0057)	(0.0023)	(0.0019)	(0.0022)	(0.0038)
Std. Dev.	0.0277	0.0212	0.0190	0.0190	0.0200
Median	0.0355	0.0274	0.0234	0.0236	0.0241
ϕ					
Mean	0.0344	0.0402	0.0437	0.0450	0.0461
(Avg. Std.Er.)	(0.0053)	(0.0021)	(0.0017)	(0.0021)	(0.0036)
Std. Dev.	0.0166	0.0125	0.0109	0.0111	0.0119
Median	0.0368	0.0419	0.0450	0.0469	0.0485
ρ					
Mean	0.4073	0.3676	0.3684	0.3789	0.3847
(Avg. Std.Er.)	(0.0370)	(0.0184)	(0.0166)	(0.0203)	(0.0330)
Std. Dev.	0.0724	0.0657	0.0720	0.0763	0.0884
Median	0.4021	0.3663	0.3700	0.3838	0.3871
λ					
Mean	0.3360	0.3086	0.2893	0.2874	0.2825
(Avg. Std.Er.)	(0.0218)	(0.0108)	(0.0097)	(0.0118)	(0.0184)
Std. Dev.	0.0984	0.0971	0.0984	0.0949	0.0920
Median	0.3411	0.3105	0.2888	0.2898	0.2886

Table 2 presents summary statistics of the GMM model estimates of the parameters for the 274 NYSE-listed stocks in the 1990 sample period over five intraday trading intervals. The table presents 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 for the four main parameters of interest: θ, the asymmetric information component; ϕ, the transaction cost component; ρ, the autocorrelation coefficient of the order flow; and λ, the probability a trade takes place between the quotes.

number of transactions per hour equal approximately 17 during the opening and closing hours of trading, only 11 transactions occur during midday. Interestingly, although the volatility of transaction price changes is U-shaped, the variance of ask price changes declines throughout the day. This is especially true during the opening interval in which the standard deviation of quote changes drops from 7.8 cents to 6.9 cents.

2.3 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, standard deviation of the estimates, and median estimates of the parameter vector β for the 274 stocks in each of the five intraday trading intervals. The drift parameter (α) is essentially zero and is not reported.

The parameter estimates have economically reasonable values. The information asymmetry parameter (θ) is of special interest. From Table 2 it is clear that information 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 4.15 to 2.75 cents) and remains at this level until the final period where it increases slightly. We also performed a formal Wald test that θ is equal over the course of the day. The average value of the Wald statistic (which has a χ_4^2 distribution) over the sample of firms is 17.22, and we reject the restriction that θ does not change over the trading day for just over 68% of the sample firms. Such statistics, although informative, should be viewed with caution since the null hypothesis under consideration is not necessarily well specified. As a gauge of the reliability of these estimates, Table 2 also provides the average standard error across the 274 stocks. The standard errors are quite tight, and range from 0.19 cents to 0.57 cents at different times of the 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 market makers learning about fundamental asset values (i.e., price discovery) through the trading process, or 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. From a theoretical viewpoint, it is also unlikely that the percentage of uninformed traders may be systematically higher at certain points of the day, since this would attract informed traders who seek to disguise their trades, as in Admati and Pfleiderer (1988).

The transaction cost element (ϕ) is approximately 3.4 cents in the first half-hour and increases steadily over the day to 4.6 cents in the final half-hour interval, a rise of about 30%. Similar to the estimates of θ, the standard errors around these values tend to be relatively small, ranging from 0.17 cents to 0.53 cents at various times during the day. Taking these results as given, the increase in ϕ over the day is consistent with inventory control models of market making. In our model, ϕ represents the economic costs of market making, and the increase in this parameter over the day may reflect the increasing risks associated with carrying inventory overnight. Previous studies of inventory control by market makers [see, e.g., Hasbrouck (1988), Hasbrouck and Sofianos (1993), and Madhavan and Smidt (1991)] find relatively weak evidence of intraday inventory effects, although Madhavan and Smidt (1993) suggest that these effects may be apparent at lower frequencies. If inventory effects are manifested toward the end of the day, the conclusions of these studies may be worth reinvestigating.

The autocorrelation of order flow (ρ) also lies in a fairly narrow range. From Table 2, the mean value of the autocorrelation lies between 0.367 and 0.407. While the pattern is mildly U-shaped, the average standard errors and lack of supporting theory suggest there is no systematic pattern in the parameter over the day. Indeed, the average and median Wald statistic equals 9.37 and 7.19, respectively, with only 38.38% rejections across the sample of firms. Nevertheless, as expected, the mean estimate is positive and the size of the implied autocorrelation suggests that ignoring this effect may have important economic consequences.

The probability that a trade occurs within the quotes, λ, declines monotonically over the day. The mean estimates drop from 34% in the opening interval to 28% at the close. With average standard errors of around 1 to 2%, the drop is arguably large. 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. From a theoretical viewpoint, the steady decline in λ over the day is consistent with increased incentives by liquidity providers to place limit orders when spreads are wide as well as a higher probability of a cross in intervals of high activity.

3. Economic Implications of the Model

3.1 Assessing the cost of trading

3.1.1 The bid-ask spread. Recent allegations of collusion among NASDAQ dealers in setting bid-ask quotes [see, e.g., Christie and Schultz (1994) and Christie, Harris, and Schultz (1994)] have focused renewed attention on the accurate measurement of transaction costs. We show in this section that the bid-ask spread is a misleading measure of the true costs of trading, and use the model to provide more accurate estimates.

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

$$[2,\ 2]\, V_{\hat{\theta},\hat{\phi}} \begin{bmatrix} 2 \\ 2 \end{bmatrix}, \tag{6}$$

where $\hat{\theta}$ and $\hat{\phi}$ denote the sample estimates of the parameters θ

Table 3
Summary statistics of estimated trading costs

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

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.

and ϕ, and $V_{\hat{\theta},\hat{\phi}}$ is the estimated covariance matrix of the GMM estimators.[7]

Table 3 presents the mean coefficient estimate, mean standard error, standard deviation of the coefficient estimate, and median estimate of the bid-ask spread measure for 274 stocks in each of the five time intervals. Note that the spread is a function of the underlying parameters of the structural econometric model described in Section 2.1. Thus its patterns can be interpreted economically within the model's framework.

The implied spread exhibits the familiar U-shaped pattern over the course of the day (see, e.g., Table 1). In particular, the mean spread in the initial period is 15.2 cents, drops to 14.3 cents during the day, and rises to 15.0 cents at the end of the day. However, the difference between the spread at the beginning (end) of the day and the spread at

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

the middle of the day is small because our sample consists of actively traded stocks for whom the U-shaped pattern is less pronounced. In terms of the reliability of these results, note that the average standard error varies from 0.24 cents to 0.66 cents over the day. Approximately one-half (i.e., 51.46%) of the firms display significant patterns at the 5% level, with average and median Wald statistics of 14.25 and 9.77.

Nevertheless, standard errors aside, it is important to point out that our model provides an explanation for this U-shaped pattern. In particular, the intraday patterns of the asymmetric information cost, θ, and the dealer cost, ϕ, govern the spread's behavior. At the start of the day, information asymmetries are large, so that the spread is wide. Over the day, asymmetries are resolved through price discovery, causing spreads to narrow. Toward the end of the day, the asymmetric information component is small but high transaction costs (possibly reflecting the risks of carrying inventory overnight) cause the spread to widen again.

On the positive side, these structural estimates of the spread do not rely on quotations data but are inferred from the autocovariance and other moment conditions of transaction price changes. Thus one advantage of our structural framework is that, given the transactions data, the U-shaped pattern is implied by the model. On the negative side, however, because we do not use quotations data directly in estimation, the model's spread need not equal the actual (sample) quoted spread. Indeed, our estimates of the bid-ask spread are lower than the actual quoted spread. Part of the difference may reflect the increased probability of a midquote transaction when the spread is wide.[8] This systematic tendency could explain the deviation between the spreads implied by our model and the actual sample spreads. This is a topic for further research, and is discussed more in Section 4.

For further evidence of the relation between the model's structural parameters and the 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 $\theta/(\phi + \theta)$ and asymptotic

[8] We do indeed find that spreads are wide when midquote trades occur. Moreover, using NYSE reported statistics, about 30% of NYSE transactions occur at the midquote, but in transactions where the spread is larger than $\frac{1}{8}$, the figure increases to 70%. These stylized facts are consistent with our results.

variance

$$\left[\frac{\theta+\phi-1}{(\theta+\phi)^2}, \frac{-\theta}{(\theta+\phi)^2}\right] V_{\hat{\theta},\hat{\phi}} \left[\begin{array}{c} \frac{\theta+\phi-1}{(\theta+\phi)^2} \\ \frac{-\theta}{(\theta+\phi)^2} \end{array}\right]. \tag{7}$$

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% in the initial period, and falls steadily to 36% in the third period and remains at about this level for the rest of the day.[9] Note that the average standard errors range between 1% and 3%. Thus these results over the day suggest a fairly large drop both economically and statistically.

3.1.2 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 of 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. Under our assumptions, the conditional probability of executing at the midquote given the trade is buyer initiated is λ, so the probabilities associated with each of the four possible price paths are $(1-\lambda)^2$, $(1-\lambda)\lambda$, $(1-\lambda)\lambda$, and λ^2, respectively. It follows that the effective spread, denoted by s_t^E, takes the form $s^E = (1-\lambda)(2\phi+\theta)$.

The effective spread, s^E, can be consistently estimated via $(1-\hat{\lambda})(2\hat{\phi}+\hat{\theta})$, with corresponding asymptotic variance given by

$$\left[-(2\phi+\theta), 2(1-\lambda), 1-\lambda\right] V_{\hat{\lambda},\hat{\theta},\hat{\phi}} \left[\begin{array}{c} -(2\phi+\theta) \\ 2(1-\lambda) \\ 1-\lambda \end{array}\right]. \tag{8}$$

The solution for the effective spread shows that the bid-ask spread overstates the true cost of trading for two reasons. First, a transaction on the NYSE may execute between the quoted bid and ask prices [see,

[9] By contrast, Huang and Stoll (1996) find that the majority of the bid-ask spread is due to dealer costs. The differences in our results may partly be explained by the sample stocks, since they focus on the Major Market Index stocks which are heavily traded.

e.g., Blume and Goldstein (1992) and Lee and Ready (1991)], so that the higher the probability of midquote execution λ, the lower the cost of trading. Second, the bid-ask spread overstates the cost of a round-trip transaction because prices tend to rise following a purchase and fall after a sell.

To see how our estimates of s^E compare to the implied bid-ask spread we compute the ratio of the effective spread to the implied spread. The ratio r^E depends both on the probability of executing within the quotes and the relative magnitudes of the information and cost components of the spread. The statistic r^E can be estimated by

$$\hat{r}^E = \frac{(1-\hat{\lambda})(2\hat{\phi}+\hat{\theta})}{2(\hat{\phi}+\hat{\theta})}, \tag{9}$$

with the following asymptotic variance,

$$\left[-\frac{(2\phi+\theta)}{2(\theta+\phi)}, \frac{(1-\lambda)(4\phi+3\theta)}{2(\theta+\phi)^2}, \frac{(1-\lambda)(3\phi+2\theta)}{2(\theta+\phi)^2}\right] \times V_{\hat{\lambda},\hat{\theta},\hat{\phi}} \begin{bmatrix} -\frac{(2\phi+\theta)}{2(\theta+\phi)} \\ \frac{(1-\lambda)(4\phi+3\theta)}{2(\theta+\phi)^2} \\ \frac{(1-\lambda)(3\phi+2\theta)}{2(\theta+\phi)^2} \end{bmatrix}. \tag{10}$$

Table 3 shows that the estimates of the effective spread, s_E, are substantially smaller than the implied spread, s. In particular, the ratio of the effective spread to the implied spread (r^E) is, on average, 50% at the beginning of the day and increases monotonically to 60% at the end of the day. Moreover, while this pattern is less reliable as measured by the average standard error estimates of 2 to 4.7%, there is a consistent theoretical explanation for the monotonicity result. Specifically, unlike the implied spread, the effective spread does not exhibit a U-shaped pattern. Indeed, the effective spread increases monotonically over the day from 7.3 cents in the opening interval to 8.6 cents at the close, with average standard errors ranging from 0.19 to 0.50 cents at various times.

Thus there is some 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.

3.2 The determinants of price volatility

The model can be used to decompose transaction price volatility into its components. Using Equation (4), the variance of stock price changes is

$$\mathrm{var}[\Delta p_t] = \sigma_\epsilon^2 + 2\sigma_\xi^2 + (1-\lambda)[(\theta+\phi)^2 + (\theta\rho+\phi)^2 - 2(\theta+\phi)(\theta\rho+\phi)\rho]. \tag{11}$$

Volatility reflects the variance in public news shocks uncorrelated with trading activity. In Equation (11), the portion of volatility arising only from news shocks is measured by σ_ϵ^2. In addition, volatility reflects various microstructure-induced noise.

Microstructure noise, in turn, arises from several sources including price discreteness (measured by the term $2\sigma_\xi^2$), and terms involving the asymmetric information and cost components and their interaction. The portion of price volatility attributable to asymmetric information, denoted by A, is measured by the terms in Equation (11) involving only the parameter θ, that is, $(1-\lambda)(1-\rho^2)\theta^2$. Similarly, the portion of volatility arising from transaction costs alone is $B = 2(1-\lambda)(1-\rho)\phi^2$. Finally, the variance of prices in Equation (11) also includes an interaction term, $C = 2\phi\theta(1-\lambda)(1-\rho^2)$. This decomposition shows that, other things being equal, asymmetric information and transaction costs increase volatility. However, the magnitude of their effects is inversely related to the autocorrelation of order flow. Intuitively the bid-ask bounce is a source of transaction price volatility, and this effect is strongest when order flow is uncorrelated. A similar result carries through for the effect of changes in the market-maker's economic rents (i.e., his execution and liquidity costs) measured by ϕ. Interestingly there is an interaction between the asymmetric information and cost components; these two elements reinforce each other.

Hasbrouck (1991a, 1991b) uses a vector autoregression of quotes and returns to infer the proportion of price volatility attributable to trading. We build on this idea to distinguish the relative importance of public information and various market frictions on price volatility. The fraction of variance attributable to trading frictions is π, where

$$\pi = \frac{2\sigma_\xi^2 + (1-\lambda)[(\theta+\phi)^2 + (\theta\rho+\phi)^2 - 2(\theta+\phi)(\theta\rho+\phi)\rho]}{\sigma_\epsilon^2 + 2\sigma_\xi^2 + (1-\lambda)[(\theta+\phi)^2 + (\theta\rho+\phi)^2 - 2(\theta+\phi)(\theta\rho+\phi)\rho]}. \tag{12}$$

To estimate π we need to identify two additional model parameters governing the volatility of prices, σ_ϵ^2 and σ_ξ^2. To do this, we need two additional moment conditions. One condition is suggested by Equation (11) which places restrictions on the variance of transaction price changes. In addition, the serial covariance between successive pricing errors in Equation (4) is $-\sigma_\xi^2$, providing a second orthogonality restriction. Then we can estimate the additional parameters using the moment conditions discussed in Section 1.2 together with the following moments

$$E\begin{pmatrix} (u_t-\alpha)^2 - (\sigma_\epsilon^2 + 2\sigma_\xi^2) \\ (u_t-\alpha)(u_{t-1}-\alpha) + \sigma_\xi^2 \end{pmatrix} = 0, \tag{13}$$

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.[10]

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 average proportional contribution to total price volatility, the mean coefficient estimate of this component, its

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

Table 4
The components of the volatility of transaction price movements

	9:30–10:00	10:00–11:30	11:30–2:00	2:00–3:30	3:30–4:00
Volatility	0.007	0.0068	0.0065	0.0066	0.0070
Public Information					
Prop. of Volat.	0.4626	0.4064	0.3679	0.3565	0.3526
Mean	0.00363	0.00302	0.00264	0.00259	0.00267
Std. Dev.	0.00260	0.00227	0.00203	0.00204	0.00221
Median	0.00300	0.00238	0.00203	0.00202	0.00211
Discreteness					
Prop. of Volat.	0.0161	0.0272	0.0346	0.0380	0.0357
Mean	0.00004	0.00015	0.00002	0.00022	0.00024
Std. Dev.	0.00065	0.00040	0.00039	0.00048	0.00086
Median	0.00013	0.00016	0.00020	0.00022	0.00023
Asymmetric Information					
Prop. of Volat.	0.1345	0.0916	0.0740	0.0726	0.0762
Mean	0.00122	0.00077	0.00060	0.00059	0.00065
Std. Dev.	0.00168	0.00103	0.00081	0.00089	0.00100
Median	0.00068	0.00042	0.00032	0.00033	0.00034
Transaction Cost					
Prop. of Volat.	0.2158	0.2928	0.3416	0.3488	0.3484
Mean	0.00122	0.00165	0.00193	0.00200	0.00211
Std. Dev.	0.00092	0.00109	0.00112	0.00111	0.00114
Median	0.00104	0.00148	0.00178	0.00188	0.00204
Interaction					
Prop. of Volat.	0.1711	0.1820	0.1819	0.1841	0.1872
Mean	0.00112	0.00120	0.00119	0.00122	0.00130
Std. Dev.	0.00081	0.00053	0.00049	0.00056	0.00068
Median	0.00119	0.00122	0.00117	0.00123	0.00130

Table 4 presents summary statistics of estimates of five components of the volatility of transaction price changes, σ_p^2, for our sample of 274 NYSE-listed stocks over five intraday trading intervals. The five variance components are public information, σ_ϵ^2; discreteness, $2\sigma_\xi^2$; asymmetric information, σ_θ^2; transaction costs, σ_ϕ^2; and interaction, $\sigma_{\theta\phi}^2$. For each volatility component, the table presents the average proportional contribution of the component to total intraday price volatility, mean coefficient estimate across the stocks, the standard deviation of the estimates across the stocks, and the median coefficient estimate based on the GMM estimation described in Table 2.

standard deviation, and its median estimate over the 274 stocks, by time of day.

The public information component of volatility, σ_ϵ^2, declines by about one-third over the day. The decline is monotonic except for the last half-hour interval where there is a small increase in the variance. As the variance decreases, the fraction of variance attributable to market frictions (i.e., π) increases steadily from 54% at the open to 65% 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 nontrading 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 noise generated by the trading process 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 events (such as corporate earnings or dividend announcements) early in the day or the overnight accumulation of news. The result is also consistent with price discovery. In this interpretation the high variance of public information shocks at the open reflects investor disagreement about the interpretation of public news events whose "fundamental" volatility may actually be constant over time. As market participants learn about market clearing prices through the process of price discovery, a consensus emerges that narrows the dispersion in beliefs and hence volatility.

Table 4 provides a detailed breakdown of the contribution of market frictions to volatility for various subperiods of the day. The variance of the rounding error (σ_ξ^2) is, as expected, very small in magnitude and in many cases is not significantly different from zero. Interestingly this term tends to increase over the day, possibly because it captures a variety of microstructure noise, and the variance of these omitted noise terms is positively related to the magnitude of dealer costs (ϕ).

The impact of asymmetric information on volatility is small and actually declines over the day. In the 9:30–10:00 A.M. period, asymmetric information (measured by A) captures, across the 274 stocks, 13.5% of the price change volatility. By the end of trading, however, this component has been reduced to 7.6%. In contrast, the part of the bid-ask bounce due to market-maker trading costs, B, is large and increases steadily over the day. As a result, the effect of dealer costs on volatility also increases from 22% at the open to 35% at the close. The interaction effect, C, is also relatively important, and accounts for approximately 17 to 19% of the volatility.

To summarize, the volatility attributable to public information shocks and asymmetric information declines over the day, but this decline is offset by steady increases in the portion of volatility attributable to transaction costs and price discreteness. The net effect is that volatility is highest at the open and close and is smallest at midday.

3.3 Autocorrelation of price changes and quote revisions

3.3.1 Transaction prices. Using Equation (4), we obtain

$$\begin{aligned}\mathrm{cov}(\Delta p_t, \Delta p_{t-1}) = & -\sigma_\xi^2 + \rho(1-\lambda)[(\theta+\phi)^2 + (\theta\rho+\phi)^2] \\ & - (1-\lambda)(\theta\rho+\phi)(\theta+\phi)(1+\rho^2). \qquad (14)\end{aligned}$$

Simplifying this expression, we can show that $\mathrm{cov}(\Delta p_t, \Delta p_{t-1}) < 0$. 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 of 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. We can therefore use the parameter estimates of Table 2 to generate implied autocorrelations of transaction price changes. Of some interest, these autocorrelations are using estimates of θ, ϕ, ρ, λ, and σ_ξ^2 derived from different moment conditions. Table 5 presents, for the five intraday time intervals, the mean implied autocorrelation (from the model and estimated parameters) for the 274 stocks in the sample. Note that these implied estimates are computed on a stock-by-stock basis using Equations (4) and (6). The average implied autocorrelations are −0.0972, −0.2026, −0.2501, −0.2588, and −0.2525, respectively, over the five trading intervals. The two most important factors for explaining this autocorrelation pattern are the increase in the costs to providing liquidity over the day and the general decline in the level of the volatility of the public information flow between 9:30 A.M. and 2:00 P.M.

It is interesting to relate the implied correlations with the actual correlations present in the data. If these implied autocorrelations capture some of the characteristics of the data, the structural model may provide clues as to what drives the short-horizon time variation of returns. Including the results above, Table 5 also presents the mean sample autocorrelation of transaction price changes, the standard deviation of these estimates and the standard deviation of the difference between the actual and implied estimates across the 274 stocks.

There is a close correspondence between the actual and implied autocorrelation of transaction price changes, especially after the 9:30–10:00 A.M. period. For example, the sample autocorrelations after this period equal −0.2166, −0.2433, −0.2484, −0.2197, respectively, which matches the average implied autocorrelations pattern and magnitude described above. 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 structural model does have information for the source of the autocorrelation of transaction returns. However, the implied autocovariances are more dispersed than the actual sample

Table 5
Actual and implied moments of price changes and quote revisions

	9:30–10:00	10:00–11:30	11:30–2:00	2:00–3:30	3:30–4:00
corr(ΔP_t, ΔP_{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(ΔP_t^{ask}, ΔP_{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(ΔP_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

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(ΔP_t, ΔP_{t-1}), the autocorrelation of price changes; corr(ΔP_t^{ask}, ΔP_{t-1}^{ask}), the autocorrelation of ask price changes; and var(ΔP_t^{ask}), the variance of ask price changes.

estimates, possibly because the parameters are estimated with noise. As predicted, the actual and implied autocovariances are negative (and fairly substantially so) in all time periods.

3.3.2 Quote revisions. Our model also has implications for the autocorrelation patterns in successive quote revisions. Since the model is symmetric, the implied pattern for bid price revisions and ask price revisions is the same, and we focus here on ask price revisions. Interestingly, even though market makers in our model have rational expectations and set ex post rational or regret-free quotes, the theoretical autocorrelation between successive ask revisions can be positive or negative.[11] To show this, note that the definition of the ask price implies that

$$p_t^a - p_{t-1}^a = \theta(1-\rho)x_{t-1} + \epsilon_t + \xi_t^a - \xi_{t-1}^a, \tag{15}$$

[11] This is consistent with recent empirical studies which find ask-to-ask returns and mid-quote returns do not follow martingales [see, e.g., Handa (1991) and Hasbrouck and Ho (1987)]. Such autocorrelations may also arise from time-varying expected returns [Conrad, Kaul, and Nimalendran (1991)] or explicit models of inventory control [Madhavan and Smidt (1993), Huang and Stoll (1994), and Huang and Stoll (1996)].

where ξ_t^a is the rounding error on the ask side. A similar equation holds for the revision in bid prices. Equation (15) shows that the change in the ask price is related to the previous trade, but the greater the autocorrelation in order flow, the less the revision in beliefs. Intuitively, if order flow is highly correlated, successive transactions at the ask are more likely than a reversal from ask to bid, and the revision in beliefs reflects this fact.

Using Equation (15), we obtain

$$\operatorname{cov}(\Delta p_t^a, \Delta p_{t-1}^a) = -\sigma_{\xi^a}^2 + \theta^2(1-\lambda)\rho(1-\rho)^2, \tag{16}$$

where we assume that ξ_t^a is distributed independently with mean zero and variance $\sigma_{\xi^a}^2$. This expression is not generally equal to zero, even though [taking expectations in Equation (1)], market-makers' beliefs follow a martingale. Indeed, the autocorrelation of ask revisions is zero only in the special case where (i) there are no rounding errors arising from price discreteness (i.e., $\sigma_{\xi^a}^2 = 0$), and (ii) there is no information asymmetry (i.e., $\theta = 0$) or the autocorrelation of order flow is zero (i.e., $\rho = 0$) or the trivial case where all trades occur at the midquote. The covariance of ask price revisions does not depend on the cost parameter ϕ, because this parameter affects transaction prices, not ask prices or the midquote.

While price discreteness induces negative covariance in ask revisions, the overall covariance in Equation (16) may be positive or negative depending on the sign of the autocorrelation order flow and the relative magnitudes of the parameters. 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.[12] 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 (16) 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.

---

[12] This may also occur because of the presence of stale limit orders. In our model, quotes are rational, so this factor is omitted.

Table 5 shows that the actual serial correlation of successive ask price revisions is negative [as in Huang and Stoll (1994)], 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 0.0042 to 0.0035, and then level off at 0.0031. Similarly, the sample variances drop from 0.0061 to 0.0048, and then level off at 0.0040. To the extent that the implied variances of ask price changes are calculated using parameters estimated from transactions data, the similarity in patterns provides us with a potential explanation for the true pattern in the variance of ask price changes.[13] In particular, the fall in the variance of ask price changes is due to the combination of price discovery and less public information arriving to the market. With quote revisions, transaction costs (i.e., ϕ) are not relevant, so the increase in ϕ throughout the day does not lead to a corresponding increase in the variance.

It is important to point out that these implied estimates of variances and autocorrelations of quote revisions are derived from the structural parameters of the model, $\beta = (\theta, \phi, \rho, \lambda)$. These parameters are estimated using transactions data within the model's framework, and do not rely on direct estimation of moments of bid/ask price changes.

4. Discussion and Conclusions

We turn now to a discussion of the limitations of the model and some possible avenues for future research.

[13] While the patterns are similar, it should be pointed out that, like the spreads, the implied variance of ask price changes is of a lower magnitude than actual estimates. We hope to explore the relation between this result and the likewise result for bid-ask spreads in future research.

4.1 Limitations and extensions

The model does not perform well in some dimensions, providing an indicator of where extensions are most needed. First, the difference between the implied bid-ask spread and actual spreads is troubling. The expected spread is underestimated by approximately one-third systematically throughout the day. One explanation for this result is that some quotes are not representative of the price level at which trades might take place. For example, midquote transactions are far more common when quoted spreads are large. We could extend the model to incorporate this feature by making λ a function of the quoted spread. This more complex model might also address the model's failure to capture a second characteristic of the data, namely the level of quote revision volatility. If bid-ask spreads are subject to a stochastic, market liquidity factor, then this factor would also lead to an increase in quote revision volatility, which is not captured by our model. Interestingly, the intraday patterns of both of these characteristics are explained by the structural model. This suggests that our explanation for intraday patterns may still be reasonable.

At the structural level, the model can be extended to incorporate variable order sizes. While this issue is perhaps less relevant for the frequently traded stocks in our sample (where most trades take place at or within the bid and ask quotes), it may be more important for inactive securities where information asymmetry is significant. One approach is to model the information and cost parameters as explicit functions of order size. For example, we could write $\theta = \theta(q_t)$, where q_t denotes (absolute) trade size. Theoretical models suggest linear functional forms (i.e., $\theta(q_t) = \theta_0 + \theta_1 q_t$), while empirical studies suggest concave relations such as $\theta(q_t) = \theta_0\sqrt{q_t}$ or more general Taylor series approximiations. In all cases, estimation simply requires substituting the appropriate $\theta(q_t)$ into Equation (4) and proceding as before, albeit with more explanatory variables. Alternatively, a general model could be approximated by estimating Equation (4) with multiple indicator variables for various order size ranges, as in Huang and Stoll (1996).

A related extension is to explicitly model the intraday evolution of the parameters as a nonlinear function of elapsed time rather than estimate them separately for subperiods of the day. As above, we could model the parameters as parametric functions of the time elapsed from the start of the appropriate trading horizon (either the open or some event day of interest), and approximate the true functional form using this expansion. We can also estimate the model over calendar intervals other than a trading day (e.g., as an exploration of day-of-the-week anomalies) or over more natural sample periods defined by economic events such as earnings or dividend announcements. Indeed there is considerable discussion in the accounting literature about the extent of asymmetric information around announcement dates. Our model provides a starting point for developing a method to perform an "event study" to answer questions of this sort. As a related issue, we could take account of firm-specific versus marketwide information by modeling ϵ_t as a function of movements in a major market index such as the S&P 500 index. Such extensions may provide better discrimination of public versus private information flows.

Finally, the model can be used to compare and contrast price formation across different assets and trading systems. A crucial attribute of a trading system is its ability to discover market clearing prices. Our model suggests a natural metric for price discovery, that is, the rate at which the information asymmetry parameter declines over the day. Although we do not perform a formal cross-sectional analysis of price discovery, we find evidence of systematic variation in the parameters across stocks. Indeed, consistent with intuition, our metric of price discovery is most important for stocks in the smallest firm size decile. Clearly other factors (such as international cross-listing, option trading, market design, etc.) may also affect the speed of price discovery. For example, given the results of Christie, Harris, and Schultz (1994) and Christie and Schultz (1994), it would be interesting to test whether there are systematic differences in information asymmetry and other spread elements between comparable Nasdaq and NYSE stocks. A more detailed cross-sectional analysis of these determinants of the model's parameters is a subject for future research.

4.2 Conclusions

This article shows that security prices change because of new public information and through information revealed in the trading process. We develop a model to explain the evolution of prices over the day that embodies these two features. The article's contribution is threefold. First, with a relatively simple structural model, many patterns in intraday bid-ask spreads, execution costs, price and quote volatility, and autocorrelations of price changes and quote revisions can be jointly explained. The model provides a unified framework that sheds light on why, over the day, (i) the variance of transaction price changes is U-shaped while the variance of ask price changes is declining, (ii) the bid-ask spread is U-shaped although information asymmetry and uncertainty over fundamentals is decreasing, and (iii) the autocorrelations of transaction price changes are large and negative, yet the autocorrelations of ask price changes are small and negative.

Second, because the model is so simple, the structural parameters can be estimated in a setting that provides a high comfort level in terms of estimation error. The model's simplicity allows for natural extensions to incorporate variable order size, time-varying parameters, price discreteness, and more detailed descriptions of the order flow and information arrival processes. Integrating these extensions and incorporating some of the generality of the reduced-form approach may yield a more universal model of price formation.

Third, although our focus is on the evolution of intraday prices, a model of this type can be applied to many other issues of interest. A partial list of such topics includes (i) an analysis of limit order execution probability and its impact on execution costs; (ii) the extent to which market structure or firm characteristics affects the speed of price discovery (measured by the rate of decrease of the information asymmetry parameter); (iii) the dynamic relation between price volatility and order flow; and (iv) an intermarket analysis of the components of the bid-ask spreads. These, however, are topics for future research.

References

Admati, A. R., and P. Pfleiderer, 1988, "A Theory of Intraday Trading Patterns," *Review of Financial Studies*, 1, 3–40.

Amihud, Y., and H. Mendelson, 1987, "Trading Mechanisms and Stock Returns: An Emprical Investigation," *Journal of Finance*, 42, 533–553.

Bloomfield, R., 1996, "Quotes, Prices, and Estimates in a Laboratory Market," *Journal of Finance*, 51, 1791-1808.

Bloomfield, R., and M. O'Hara, 1996, "Market Transparency: Who wins and who loses?," working paper, Cornell University.

Blume, M., and M. Goldstein, 1992, "Displayed and Effective Spreads by Market," working paper, University of Pennsylvania.

Choi, J. Y., D. Salandro, and K. Shastri, 1988, "On the Estimation of Bid-Ask Spreads: Theory and Evidence," *Journal of Financial and Quantitative Analysis*, 23, 219–230.

Christie, W., J. Harris, and P. Schultz, 1994, "Why did NASDAQ Market Makers Stop Avoiding Odd-Eighth Quotes?," *Journal of Finance*, 49, 1841–1860.

Christie, W., and P. Schultz, 1994, "Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?," *Journal of Finance*, 49, 1813–1840.

Conrad, J., G. Kaul, and M. Nimalendran, 1991, "Components of Short-Horizon Individual Security Returns," *Journal of Financial Economics*, 29, 365–384.

Easley, D., and M. O'Hara, 1992, "Time and the Process of Security Price Adjustment," *Journal of Finance*, 47, 577–606.

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

Garbade, K., and W. Silber, 1979, "Structural Organization of Secondary Markets: Clearing Frequency, Dealer Activity and Liquidity Risk," *Journal of Finance*, 34, 577–593.

George, T. J., G. Kaul, and M. Nimalendran, 1991, "Estimation of the Bid-Ask Spread and Its Components: A New Approach," *Review of Financial Studies*, 4, 623–656.

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

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.

Greene, J., 1996, "The Impact of Limit Order Executions on Trading Costs in NYSE Stocks: An Empirical Examination," working paper, Indiana University.

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.

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

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

Hasbrouck, J., 1993, "Assessing the Quality of a Security Market: A New Approach to Transaction Cost Measurement," *Review of Financial Studies*, 6, 191–212.

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, 1993, "The Trades of Market-Makers: An Analysis of NYSE Specialists," *Journal of Finance*, 48, 1565–1595.

Huang, R., and H. Stoll, 1994, "Market Microstructure and Stock Return Predictions," *Review of Financial Studies*, 7, 179–213.

Huang, R., and H. Stoll, 1996, "The Components of the Bid-Ask Spread: A General Approach," *Review of Financial Studies*, 10, 995–1034.

Jain, P., and G. Joh, 1988, "The Dependence Between Hourly Prices and Trading Volume," *Journal of Financial and Quantitative Analysis*, 23, 269–283.

Keim, D. B., and A. Madhavan, 1996, "The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects," *Review of Financial Studies*, 9, 1–36.

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

Madhavan, A., 1992, "Trading Mechanisms in Securities Markets," *Journal of Finance*, 47, 607–642.

Madhavan, A., and M. Cheng, 1997, "In Search of Liquidity: Block Trades in the Upstairs and Downstairs Markets," *Review of Financial Studies*, 10, 175–204.

Madhavan, A., and S. Smidt, 1991, "A Bayesian Model of Intraday Specialist Pricing," *Journal of Financial Economics*, 30, 99–134.

Madhavan, A., and S. Smidt, 1993, "An Analysis of Daily Changes in Specialist Inventories and Quotations," *Journal of Finance*, 48, 1595–1628.

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

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

Seppi, D., 1990, "Equilibrium Block Trading and Asymmetric Information," *Journal of Finance*, 45, 73–94.

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