Journal of Financial Markets 4 (2001) 359–384

Journal of FINANCIAL MARKETS

www.elsevier.com/locate/econbase

Volatility and market structure☆

Kenneth A. Kavajecz[a,*], Elizabeth R. Odders-White[b]

[a] Finance Department, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6367, USA

[b] Finance Department, University of Wisconsin-Madison, 975 University Avenue, Madison, WI 53706, USA

Received 1 November 2000; accepted 17 January 2001

Abstract

This study examines volatility within three related intra-day series – transaction returns, quote midpoint returns, and limit order book midpoint returns – for a set of NYSE-listed stocks. We document statistically significant GARCH effects both overall and surrounding earnings announcements in all three series for the majority of stocks in the sample. We then compare the extent of volatility clustering among the series. In addition, the relation between volatility and market structure is examined via a set of cross-sectional regressions, and relations among the series over time are studied in a vector autoregressive framework. © 2001 Elsevier Science B.V. All rights reserved.

JEL classificaion: G10; C30

Keywords: Market microstructure; Limit order book; Volatility; GARCH

1. Introduction

The fact that the variance of stock returns changes over time is well documented in the finance literature. Much of the existing research describes

☆We gratefully acknowledge helpful comments from Avanidhar Subrahmanyam (the editor) and two anonymous referees as well as Tim Bollerslev, Matthew Clayton, Simon Gervais, and Mark Ready. All remaining errors are our own.

*Corresponding author. Tel.: +1-215-898-7543; fax: +1-215-898-6200.

E-mail addresses: kavajecz@wharton.upenn.edu (K.A. Kavajecz), ewhite@bus.wisc.edu (E.R. Odders-White).

PII: S 1 3 8 6 - 4 1 8 1 (0 1) 0 0 0 1 3 - 1

the time-variation in volatility using some form of ARCH or GARCH model (Engle, 1982; Bollerslev, 1986).[1] For example, French, Schwert, and Stambaugh (1987) find GARCH effects in monthly S&P returns and Cheung and Ng (1992) and Engle and Mustafa (1992) find ARCH effects in daily returns for individual stocks. While the existence of conditional heteroskedasticity in asset returns is well documented, this phenomenon is not well understood.

In addition to documenting volatility clustering in financial series, the work in this area has focused on determining its cause. One commonly offered explanation is a similar clustering of news or information arrival. There is some evidence in support of this hypothesis (e.g. Engle et al., 1990); however, the underlying cause of the news clustering is not known. The possibility that changes in macroeconomic conditions affect volatility has also been investigated. For example, Campbell (1987) and Glosten et al. (1993) find a relation between volatility and interest rate levels. A third explanation, which has received less attention, relates conditional heteroskedasticity to micro-level factors. Lamoureux and Lastrapes (1990) find that the ARCH coefficient (α) is no longer significant when trading volume is incorporated into a GARCH model. Bollerslev and Domowitz (1991) examine the relation between price volatility and trading mechanisms in a futures market setting. They find that returns for transactions executed via the Globex futures trading system exhibit GARCH effects, while returns for transactions accomplished through an open outcry mechanism do not.

Despite this extensive literature, little conclusive evidence has emerged regarding the underlying cause of volatility clustering. Consequently, we adopt a slightly different approach to increase our understanding of this phenomenon. Rather than directly studying a particular cause, we compare the degree of volatility clustering across a set of financial series that provide a natural controlled experiment. Specifically, we examine three distinct, yet related return series – a transaction series, a quote midpoint series, and a limit order book midpoint series – for each individual stock in our sample. As a result, we are able to examine the way in which the three series relate to one another and to isolate differences in their relation to the market structure. This comparison is unique because it analyzes the properties of conditional variance series for the same asset over the same time period. The only differences among the three series are the participants (liquidity providers versus liquidity demanders) and the private information they each possess.

Our analysis documents conditional heteroskedasticity in all three return series overall, as well as surrounding earnings announcements. In other words,

[1] See Bollerslev et al. (1992) for a detailed survey of the application of ARCH models to stock market data. Stochastic volatility models are an alternative to ARCH-type models. See, for example, Andersen (1996).

we find that both past shocks to the return series and lagged values of the variance affect the magnitude of the current variance. The extent to which current variance is impacted by lagged shocks to the return series and lagged variance is known as the asset's sensitivity and persistence levels respectively. In comparing the variances of the three series, we find that the sensitivity and persistence levels for the limit order book midpoint returns equal or exceed those in quote midpoint returns, which in turn surpass those in transaction returns. Moreover, we find that the transaction series is least affected by earnings announcements, while the limit order book series is most affected. Specifically, the sensitivity to shocks increases for the limit order book and decreases for the quote series, while the persistence parameters decrease significantly for both series during the announcement period.

We also show that the level, sensitivity, and persistence of the conditional volatility series are related to the overall structure of the market. For each series, we estimate a set of cross sectional regressions relating volatility to liquidity measures. We find that the variance of each series is negatively related to trading volume and the average quoted spread and depth are positively related to the variance of quoted depth and to the stock's average trade size relative to quoted depth. In addition, a high relative trade size is associated with decreased persistence levels for transactions and the limit order book, but increased sensitivity levels for quotes. Higher trading volume is linked to a higher sensitivity to incoming information for the transaction variance series and a lower persistence for the limit order book series.

Despite the prevalence of volatility clustering, a substantial minority of the stocks in our sample have constant conditional variance for at least one of the series. We find that stocks with constant conditional variances for the quote and limit order book series tend to be either low volume stocks or high volume stocks with high prices, while stocks with constant conditional transaction variances are evenly disbursed over the trading volume/price groups.

Finally, we examine the time-series relations among the three volatility series for each stock using a vector autoregressive model. The results suggest that, while there is a strong relation between volatility clustering in quote and transaction returns, the relation among limit order book returns and the other series is substantially weaker.

Overall, this paper provides the first description of the conditional variance series for limit order books and offers insight into the behavior of market participants. By comparing the three series, we highlight important differences between liquidity demanders and liquidity providers and how those differences are manifested in price volatility. Specifically, our results, which document varying degrees of conditional heteroskedasticity across the three series, are consistent with liquidity demanders and liquidity suppliers reacting to different stimuli and having different aptitudes for processing incoming information. In addition, we show that, although microstructure effects are related to volatility clustering, they are unable to explain it fully. Consequently, other considerations, such as the speed of information arrival, are likely to have a substantial impact on clustering as well.

The remainder of the paper is organized as follows. Section 2 describes the data, the procedure used to estimate the limit order books, and the model employed in the analysis. Section 3 discusses volatility clustering in the transaction, quote, and limit order book return series both unconditionally and surrounding information events. Section 4 examines relations between volatility and the NYSE market structure, as well as the characteristics of stocks with constant conditional variance series. Section 5 uses a vector autoregression to relate volatility clustering among the three return series through time. Section 6 concludes.

2. Data and estimation techniques

Because our analysis includes transaction returns, quote midpoint returns, and limit order book midpoint returns, we need not only transaction and quote data, but also limit order data. Consequently, we use the somewhat dated TORQ (Trades, Orders, Reports and Quotes) data set. The TORQ data contain quote, transaction, and order information for 144 NYSE stocks from November 1990 through January 1991.[2]

Using the TORQ data, we construct estimates of the time series of limit order books for each stock using the technique described in Kavajecz (1999). The principle behind the limit order book estimation is that at any instant in time, the limit order book should reflect those orders remaining after netting all prior execution and cancellation records. The first step in the estimation is to search the entire data set for records that have order arrival dates prior to November 1, 1990. We use this set of limit orders to form an estimate of the initial limit order book (or "prebook") at the start of the TORQ data set. After the prebook is constructed, current records are processed. To estimate the limit order book for a given date and time, all records with a date and time stamp prior to the chosen date and time are selected, new orders are added to the prebook and execution and cancellation records are matched to existing orders on the limit order book, where matched orders are eliminated. The remaining or residual set of orders becomes our estimate of the limit order book for the chosen date and time.

By sequentially updating the limit order book estimates, we create a sequence of "snapshots" of the limit order books for each stock. Using the sequence of best buyside and sellside limit prices, quoted bid and ask prices and

[2] We eliminate one stock (Symbol, GFB) from the sample because the stock was delisted after one month and therefore provides insufficient data to perform the analysis.

transaction prices, we construct three return series for each stock. Specifically, returns on the midpoint of the best buyside and sellside limit prices, on the midpoint of the quoted bid and ask, and on transaction prices are computed by selecting observations from each series at fifteen-minute intervals.[3] Given the focus of the study, a fixed time interval is appropriate because if the observation intervals were allowed to vary, it would be difficult to differentiate changes in volatility from the effects of the irregular spacing.[4]

After the three return series are constructed for each stock, we estimate the conditional variance of each series using an autoregressive model with GARCH residuals. We select a GARCH(1,1) specification because it is commonly used in the existing literature for measuring volatility clustering. Moreover, this specification has the desirable features of interpretability and good fit for high frequency data.[5] While the discrete nature of stock prices is a potential concern under this specification, the use of returns and the focus on the comparison across series alleviates this concern to a large extent, since we have no reason to suspect differential effects of discreteness on one series relative to another. Our modeling choice should be viewed in the context of the paper's overall goal, that is not to perfectly model time-varying variance, but to *compare* the three series in a way that is both illuminating and consistent with the existing literature.

We jointly estimate the following system of equations using maximum likelihood:

$$r_t = c + \phi r_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0,\ \sigma_t^2), \tag{1}$$

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2, \tag{2}$$

where r_t represents the return series (for transactions, quotes, or the limit order book) and ε_t is a normally distributed disturbance, with variance described by Eq. (2). The return series are modeled as first order autoregressive processes to account for the strong autocorrelation in returns. Jointly estimating the vector $(c,\ \phi,\ \omega,\ \alpha,\ \beta)$ is complicated by the lagged variance term in Eq. (2). Fortunately, the current variance can be expressed as an infinite moving average process by recursively substituting for the lagged variance term. The result is an expression for the current variance involving only lagged error terms within an infinite geometric sequence. After rewriting Eq. (2) in this way, it is straightforward to estimate Eqs. (1) and (2) using maximum likelihood.

[3] An observation from each series is gathered at the time of the opening quote and then at every hr : 00, hr : 15, hr : 30, and hr : 45 thereafter, for a total of 27 observations per series per day.

[4] An alternative way of investigating this issue is to consider event time rather than calendar time. For example, Hasbrouck (1999) considers the impact of event time through the extent of time deformation.

[5] See Bollerslev et al. (1994).

Campbell et al. (1997) provide a simple interpretation of the estimated variance equation. Rearranging Eq. (2) slightly yields

$$\sigma_t^2 = \omega + (\alpha + \beta)\sigma_{t-1}^2 + \alpha(\varepsilon_{t-1}^2 - \sigma_{t-1}^2). \quad (3)$$

Here, we see that $\alpha + \beta$ measures the persistence of the information shocks in the variance and α measures the sensitivity of the variance to the most recent shock. A related way of interpreting these coefficients is that the sensitivity coefficient (α) describes the overall flow of information shocks into the market, while the persistence coefficient ($\alpha + \beta$) describes the speed with which information shocks are incorporated into prices. In particular, a high (low) sensitivity level suggests that information shocks tend to have a large (small) immediate impact. A high persistence level in the variance suggests that information is incorporated into prices gradually over time while a low persistence level in the variance suggests that information is quickly impounded into prices with a consensus on its implication. Given these interpretations, the sensitivity and persistence coefficients provide a convenient way to compare the impact of information shocks on the conditional variances of the related series.

3. Volatility clustering among transaction, quoted and limit order book prices

3.1. Unconditional parameter estimates

Table 1 contains estimation results for the model described by Eqs. (1) and (2). The median coefficient estimates for the transaction, quoted bid-ask spread midpoint, and limit order book spread midpoint returns are presented both in aggregate and by trading volume group.[6] In addition, Table 1 presents the results of Wilcoxon rank–sum tests of the hypothesis that the coefficient estimates do not differ across series or volume groups.

We observe that volatility clustering is common for the transaction series, with 46 percent (57 percent) of the stocks in the sample displaying statistically significant α (β) coefficients. The median estimates of α and β for the transaction series are 0.020 and 0.106, respectively. As a comparison, Engle and Mustafa (1992) report GARCH(1,1) results on a small subset of individual stocks and the S&P 500 index using daily data. They find average α and β coefficients of 0.044 and 0.960, respectively. The α coefficients are largely the same; however, our β coefficient estimates are substantially smaller than those found by Engle and Mustafa. This inconsistency between β estimates generated using high-frequency data and β estimates from lower-frequency data is common in these types of models. Another factor that may contribute to the

[6] Mean coefficient estimates (not reported) are qualitatively similar.

Table 1
Volatility clustering within limit order book, quoted spread and transaction returns

The table below contains median coefficient estimates from the following GARCH model:

$$r_t = c + \phi r_{t-1} + \varepsilon_t, \quad \sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2, \quad \varepsilon_t \sim N(0, \sigma_t^2),$$

where r_t is the return series and ε_t is a normally distributed disturbance, with mean zero and variance described above. The model is estimated for return series associated with the limit order book (LOB) midpoint, the midpoint of the quoted (Quotes) bid-ask spread and transaction prices (Trans). The "% Significant" is the percentage of the 143 total stocks for which the given coefficient is significant at the 5% level. Intercept (c) estimates, not reported, are insignificant for 99%, 95%, and 98% of the stocks for the LOB, quote, and transaction series, respectively. $P(\bullet)$ is the probability from Wilcoxon rank–sum tests that the limit order book (L), quotes (Q) and transactions (T) [high (H), medium (M) and low (L)] have the same median coefficients for the given sample (series).

Coef.		Sample	Median parameter estimates			Tests of equality across series for each sample			
				Series					
			LOB	Quotes	Trans	$P(\mathrm{L}=\mathrm{Q})$	$P(\mathrm{L}=\mathrm{T})$	$P(\mathrm{Q}=\mathrm{T})$	$P(\mathrm{L}=\mathrm{Q}=\mathrm{T})$
ϕ									
		Full sample	−0.09511	−0.02185	−0.07445	**0.0001**	0.0703	**0.0001**	**0.0001**
	Median	High	−0.13342	−0.01386	−0.11830	**0.0001**	0.6575	**0.0001**	**0.0001**
	Parameter	Medium	−0.10695	−0.03295	−0.07948	**0.0001**	**0.0443**	**0.0006**	**0.0001**
	Estimates	Low	−0.04161	−0.01321	−0.03742	**0.0179**	0.4819	0.1040	0.0531
		% significant	66	25	53				
	Tests of	$P(\mathrm{H}=\mathrm{M})$	0.3933	0.0977	**0.0467**				
	equality across	$P(\mathrm{H}=\mathrm{L})$	**0.0001**	0.9708	**0.0001**				
	volume groups	$P(\mathrm{M}=\mathrm{L})$	**0.0001**	**0.0490**	**0.0007**				
	for each series	$P(\mathrm{H}=\mathrm{M}=\mathrm{L})$	**0.0001**	0.1097	**0.0001**				

Table 1 (*continued*)

Coef.		Sample	Median parameter estimates: Series: LOB	Quotes	Trans	Tests of equality across series for each sample: $P(\mathrm{L}=\mathrm{Q})$	$P(\mathrm{L}=\mathrm{T})$	$P(\mathrm{Q}=\mathrm{T})$	$P(\mathrm{L}=\mathrm{Q}=\mathrm{T})$
ω									
		Full sample	4.6E-6	4.1E-6	16.8E-6	0.7051	**0.0001**	**0.0001**	**0.0001**
	Median	High	4.7E-6	4.0E-6	14.8E-6	0.2775	**0.0001**	**0.0001**	**0.0001**
	Parameter	Medium	1.7E-6	3.2E-6	23.9E-6	0.4097	**0.0001**	**0.0003**	**0.0001**
	Estimates	Low	9.4E-6	6.3E-6	16.2E-6	0.7110	0.4819	0.2163	0.4840
		% significant	92	93	85				
	Tests of	$P(\mathrm{H}=\mathrm{M})$	**0.0406**	0.4684	0.5358				
	equality across	$P(\mathrm{H}=\mathrm{L})$	0.1054	0.1238	0.6741				
	volume groups	$P(\mathrm{M}=\mathrm{L})$	**0.0108**	0.0822	0.9377				
	for each series	$P(\mathrm{H}=\mathrm{M}=\mathrm{L})$	**0.0132**	0.1423	0.8257				
α									
		Full sample	0.05013	0.05190	0.01998	0.8787	**0.0006**	**0.0010**	**0.0005**
	Median	High	0.07380	0.05951	0.08717	0.2051	0.6485	0.2407	0.3533
	Parameter	Medium	0.04973	0.06929	0.01273	0.1660	**0.0001**	**0.0001**	**0.0001**
	Estimates	Low	0.02806	0.03052	0.01900	0.7823	0.3329	0.1957	0.4052
		% significant	80	78	57				
	Tests of	$P(\mathrm{H}=\mathrm{M})$	**0.0176**	0.5514	**0.0005**				
	equality across	$P(\mathrm{H}=\mathrm{L})$	**0.0001**	**0.0206**	**0.0022**				
	volume groups	$P(\mathrm{M}=\mathrm{L})$	**0.0115**	**0.0047**	0.3647				
	for each series	$P(\mathrm{H}=\mathrm{M}=\mathrm{L})$	**0.0001**	**0.0105**	**0.0006**				

β									
		Full sample	0.77791	0.63239	0.10560	**0.0015**	**0.0001**	**0.0041**	**0.0001**
	Median	High	0.72301	0.63293	0.43555	0.1704	**0.0001**	**0.0026**	**0.0001**
	Parameter	Medium	0.89796	0.70372	0.00000	**0.0004**	**0.0001**	0.0616	**0.0001**
	Estimates	Low	0.61488	0.20743	0.00001	0.4267	0.1774	0.5245	0.3824
		% significant	78	69	46				
	Tests of	$P(\mathrm{H} = \mathrm{M})$	**0.0003**	0.3985	0.4202				
	equality across	$P(\mathrm{H} = \mathrm{L})$	0.7291	0.2878	0.5122				
	volume groups	$P(\mathrm{M} = \mathrm{L})$	**0.0058**	0.2014	0.7206				
	for each series	$P(\mathrm{H} = \mathrm{M} = \mathrm{L})$	**0.0011**	0.3252	0.6836				
$\alpha + \beta$									
		Full sample	0.87992	0.73105	0.14793	**0.0011**	**0.0001**	**0.0003**	**0.0001**
	Median	High	0.80641	0.72355	0.58588	0.0992	**0.0001**	**0.0374**	**0.0004**
	Parameter	Medium	0.95621	0.79949	0.02404	**0.0003**	**0.0001**	**0.0007**	**0.0001**
	Estimates	Low	0.73438	0.25526	0.09056	0.4914	0.1229	0.2504	0.2478
		$P(\mathrm{H} = \mathrm{M})$	**0.0001**	0.2003	0.5751				
	Tests of	$P(\mathrm{H} = \mathrm{L})$	0.6632	0.5626	0.8728				
	equality across	$P(\mathrm{M} = \mathrm{L})$	**0.0148**	0.0665	0.7546				
	volume groups	$P(\mathrm{H} = \mathrm{M} = \mathrm{L})$	**0.0014**	0.1670	0.8618				
	for each series								

difference is that the coefficients are estimated over different time periods. In particular, the TORQ sample covers the initiation of the Gulf War, which touched-off some large realizations of the market index, as well as some heightened volatility.[7]

Like the transaction series, the quote and limit order book series display substantial evidence of volatility clustering, with statistically significant GARCH coefficients for between 69 and 80 percent of the stocks in the sample. Despite the fact that the GARCH effects are generally significant in each of the three series, substantial differences among the series exist. First, the sensitivity of the variance to a shock (α) is higher for the limit order book and quote series than for transactions. These differences are significant at the 0.01 percent level. The statistical significance of the sensitivity parameter within each series follows a similar pattern. The α coefficient differs significantly from zero at the five-percent level for 80 percent of the limit order book midpoint series, but only 57 percent of the transaction series.

The persistence of volatility shocks ($\alpha + \beta$) also differs across series. Shocks to the variance of the limit order book are more persistent than shocks to the variance of the quotes, which are, in turn, more persistent than shocks to transaction volatility. This difference is again significant at the 0.01 percent level. This implies that the effects of volatility shocks linger on the limit order book longer than in either quotes or transactions.[8]

While these series are certainly related since transactions may alter quotes or the limit order book and the limit order book or quotes may induce transactions, we should emphasize that the results presented here are not mechanistically driven. The presence of the specialist and floor crowd, as well as the existence of price and depth improvement, alter the interaction among these series such that they represent distinct economic phenomena. We offer some possible explanations for the differences in coefficient estimates across series below.

One explanation for the small sensitivity and persistence estimates of the transaction series relative to the quote and limit order book series is that transactions, unlike the others, incorporate the actions of liquidity demanders in addition to liquidity suppliers. This interplay between demanders and suppliers of liquidity could potentially manifest itself in two ways. On one hand, if a substantial portion of trading volume is liquidity based rather than

[7] As a robustness check, we re-estimated the model having excluded important Gulf War event dates. The dates excluded are: November 29, 1990, UN authorizes the use of force; January 15, 1991, UN withdrawal deadline; January 17, 1991, US military planes attack Iraq; January 18, 1991, First Scud Missiles hit Israel. The results were both qualitatively and quantitatively similar to those found in Table 1.

[8] While the persistence levels are less than one for the majority of the stocks in the sample, the persistence levels for the limit order book (quote and transaction) series approach one for several (a few) stocks. This implies that the model is non-stationary in these cases.

information based, then transactions would exhibit both lower sensitivity and persistence as there is little or no information to be incorporated into prices. On the other hand, if liquidity demanders and suppliers quickly arrive at a consensus regarding the information content of shocks through the trading process, then the transaction series would exhibit a lower persistence level. In this case, however, the sensitivity level for transactions should be at least as large as that of the other series, which is inconsistent with our findings. Consequently, our results offer support for the first explanation.

A potential explanation for the difference in the persistence estimates between the quote and limit order book series is contained in Kavajecz and Odders-White (2001). They demonstrate that, despite relying heavily on the limit order book to set their prices, specialists do not simply reflect the interest on the limit order book in their quotes. They conclude that specialists actively smooth price changes, which is consistent with our results that show that specialists dampen the degree of volatility clustering found in the limit order book in posting their bid and ask prices. Summary statistics on the overall variance estimates of these series also support this interpretation. The average variance of the limit order book midpoint across all stocks in the sample is 0.00029, while the average quote midpoint variance is 0.00027 and the average transaction variance is 0.00028. These variance estimates are statistically different from each other at the 0.1 percent level.

There are also differences in the degree of volatility clustering across trading volume groups. While the ω and β coefficient estimates are similar across groups, the sensitivity to volatility shocks (α) is higher for actively traded stocks than for the other groups. This is consistent with more information being disseminated about actively-traded stocks than infrequently traded stocks, which is likely the result of greater analyst following and news coverage for high-volume stocks. These differences are the most pronounced for the limit order book, with the quoted and transaction return series displaying much less variability across groups, which is again consistent with quoted prices masking some of the volatility found on the limit order book.

It is also worth noting that some of the differences across series documented in the full sample do not hold uniformly across trading volume groups. In particular, the differences among the alpha estimates for the limit order book, quote, and transaction series are significant only for medium volume stocks. Furthermore, differences in persistence parameters ($\alpha + \beta$) across series are not statistically significant for low volume stocks. Similarities among the three series for low volume stocks likely stem from the fact that these stocks have thin limit order books (see Kavajecz, 1999).

In summary, the results document that volatility clustering exists not only in transactions, but also in quote returns and limit order book returns for a substantial fraction of stocks in our sample. Furthermore, the average coefficient estimates reveal that the series differ in their sensitivities to, and in the persistence of, shocks to the variance, with the sensitivity and persistence levels of limit order book midpoint returns generally exceeding those of both quote midpoint returns and transaction returns.

3.2. Volatility clustering surrounding information events

The results presented in Table 1 offer insight into the unconditional impact of information shocks on the three series. In this section, we extend the comparison by focusing exclusively on the differences in the effects of earnings announcements on each series, since these differences can shed additional light on the information transmission mechanisms at work on the NYSE. While a number of studies have examined the effects of earnings announcements on various aspects of the trading environment, this analysis is novel in that it examines the impact of earnings announcements on characteristics of volatility among the three related series.[9]

We examine the 82 earnings announcements that occurred during our sample period. Announcement dates and times were obtained from Dow Jones News Retrieval Service. The fifteen-minute interval during which the announcement occurred is labeled time 0 and the announcement period consists of the five days surrounding the announcement (observations -68 to $+68$).

We begin by comparing the volatility of each series during the announcement and non-announcement periods. Since corporate announcements are often associated with an increase in uncertainty, we would expect to see higher levels of volatility during the announcement period. As shown in Panel A of Table 2, the variances of all three series do increase on average during the announcement period, but none of these differences is significant.

Next, we investigate differences in GARCH parameters during announcement and non-announcement periods. Our approach is similar to that of Engle and Mustafa (1992), who study the effect of the 1987 market crash on sensitivity and persistence parameters. Specifically, we estimate an unrestricted GARCH model that allows the parameters to differ between announcement and non-announcement periods in order to isolate differences in the information revelation processes across the three series. In the unrestricted model, the GARCH variance equation (Eq. (2) above), is modified as follows:

$$\begin{aligned}\sigma_t^2 = {} & \omega_{\text{Ann}} D_{\text{A}} + \omega_{\text{Non}} D_{\text{Non}} + \alpha_{\text{Ann}} D_{\text{Ann}} \varepsilon_{t-1}^2 + \alpha_{\text{Non}} D_{\text{Non}} \varepsilon_{t-1}^2 \\ & + \beta_{\text{Ann}} D_{\text{Ann}} \sigma_{t-1}^2 + \beta_{\text{Non}} D_{\text{Non}} \sigma_{t-1}^2,\end{aligned} \tag{5}$$

[9] For example, Lee et al. (1993) and Kavajecz (1999) examine the effect of earnings announcements on quoted spreads and depths.

where D_{Ann}, the announcement period dummy variable, assumes a value of one during the announcement period (observations -68 to $+68$) and zero otherwise and $D_{\mathrm{Non}} = 1 - D_{\mathrm{Ann}}$.

The estimation results are presented in Panel B of Table 2. Median coefficient estimates for announcement and non-announcement periods, as well as the median *difference* between the two coefficient estimates, are reported. First, although differences in ω across the two periods are statistically significant for all three series for most stocks in the sample, the magnitude of the differences tends to be very small. Second, announcement period sensitivity (α) estimates tend to exceed non-announcement period estimates for the limit order book return series, but the reverse holds for the quote return series. Both differences are significant for roughly forty percent of the stocks in the sample. We interpret this as evidence that more information is impounded into the limit order book during the announcement period than during the non-announcement period but that the specialist's quotes are less reactive to changes in the limit order book during the announcement period. In addition, estimates of both β and the persistence parameter $(\alpha + \beta)$ are smaller during the announcement period for the quote and the limit order book return series. Given our interpretation that a low persistence parameter implies that shocks are resolved and prices arrive at their new equilibrium values quickly, this suggests that information is incorporated into these series more quickly around corporate events. Conversely, the persistence parameter increases for transactions during the announcement period, implying that the increased transaction volatility lingers longer during the announcement period.[10] The differences in the persistence parameters across the three series suggest that liquidity suppliers, who primarily affect the quote and limit order book spread, are better able to quickly process and interpret incoming information during the announcement relative to liquidity demanders, who primarily affect transaction prices.

4. Relating volatility to market structure

4.1. Cross-sectional regressions

Having compared the GARCH parameters across the three series, we now examine the relations between volatility and market structure by estimating a set of cross-sectional regressions. Our analysis involves nine complimentary

[10] In general, the results for the non-announcement period in Table 2 are consistent with the unconditional results presented in Table 1. However, one exception is the β estimate in the transaction equation that is caused by the heavy weighting of high trading volume stocks in our earnings announcement sample.

Table 2
Volatility clustering surrounding earnings announcements

Panel A of the table below contains a comparison of announcement and non-announcement period volatility. Panel B contains median coefficient estimates, along with median differences across periods, from the following model:

$$r_t = c + \phi r_{t-1} + \varepsilon_t, \quad \sigma_t^2 = \omega_{\text{Ann}} D_{\text{A}} + \omega_{\text{Non}} D_{\text{Non}} + \alpha_{\text{Ann}} D_{\text{Ann}} \varepsilon_{t-1}^2 + \alpha_{\text{Non}} D_{\text{Non}} \varepsilon_{t-1}^2$$
$$+ \beta_{\text{Ann}} D_{\text{Ann}} \sigma_{t-1}^2 + \beta_{\text{Non}} D_{\text{Non}} \sigma_{t-1}^2, \quad \varepsilon_t \sim N(0, \sigma_t^2),$$

where r_t is the return series and ε_t is a normally distributed disturbance, with mean zero and variance described above over the 82 earnings announcements in our sample. Although the sample includes 82 earnings announcements, the GARCH analysis uses only the subset of stocks with non-constant variance for the given series, resulting in smaller sample sizes. "% Significant" denotes the percentage of stocks with statistically significant rejections of H_0: Ann estimate = Non-estimate (or, in the case of ϕ, H_0: Estimate = 0) using Wald tests. Intercept (c) estimates, not reported, are insignificant for 92%, 93%, and 96% of the stocks for the LOB, quote, and transaction series, respectively.

Panel A: Announcement and non-announcement period variances

	Series		
	Limit order book	Quotes	Transactions
Ann variance – Non-variance			
mean difference	0.00012	0.00001	0.00017
P-value (H_0: $\sigma_{\text{A}}^2 = \sigma_{\text{N}}^2$)	0.4867	0.8283	0.4214

Panel B: Unrestricted GARCH(1,1) model estimates

Coefficient\Sample size	50	72	73
ϕ	−0.09855	−0.02877	−0.08263
% significant	56	22	66
ω			
Ann ω	4.3E-5	1.3E-5	1.3E-5
Non-ω	8.3E-6	5.3E-6	2.3E-5
Ann ω - Non-ω	1.3E-5	3.0E-6	−6.0E-7
% significant	80	71	59
α			
Ann α	0.06734	0.03658	0.02090
Non-α	0.07231	0.07238	0.02458
Ann α - Non-α	0.00182	−0.01053	0.00000
% significant	45	37	44
β			
Ann β	0.08618	0.08535	0.42132
Non-β	0.76943	0.61770	0.47417
Ann β - Non-β	−0.48114	−0.21180	0.00000
% significant	73	70	48

Table 2 (*continued*)

Coefficient\Sample size	50	72	73
$\alpha + \beta$			
Ann $(\alpha + \beta)$	0.28834	0.20252	0.57328
Non-$(\alpha + \beta)$	0.88338	0.73877	0.61540
Ann $(\alpha + \beta)$ - non-$(\alpha + \beta)$	−0.37060	−0.08773	0.01138
% significant	68	58	51
Joint test for differences in ω, α, and β			
% significant	96	89	79

regressions – three for each series – each estimated using one observation per stock (143 observations total). Relying on the existing literature for guidance, we select six liquidity measures as potential determinants of cross-sectional differences in average conditional variance, as well as in sensitivity (α) and persistence ($\alpha + \beta$) levels.

The first explanatory variable we include is the bid-ask spread for each stock. The following example illustrates the reasons for its inclusion. Suppose that the minimum price variation is \$1/8, as it was during our sample period, and compare two stocks, one with a bid-ask spread of a \$1/8 and the other with a bid-ask spread of \$1/2. There are two conflicting hypotheses relating spreads and the volatility series we consider. One hypothesis predicts that a large bid-ask spread will restrict the information flow to the market by discouraging informed traders from trading. In this case, the transaction conditional variance and the quoted spread will be negatively related. Furthermore, given that the flow of information to the market is dampened, we would expect to find a similar negative relation between the spread and the transaction sensitivity parameter (α). The impact of a wide spread on the quotes and limit order book series is unclear under this hypothesis. Quote and limit order book volatility as well as sensitivity might be decreased if a large bid-ask spread deters all market activity. Alternatively, a large bid-ask spread might induce liquidity demanders to switch from submitting market orders to limit orders. The increase in limit order book activity may increase the volatility and sensitivity of the limit order book and possibly also the quotes.

Under a second hypothesis, stocks with larger spreads will be likely to exhibit more transaction conditional variance since there is more room within which prices can vary, given that most transactions occur at or within the quoted spread. In other words, the large potential bid-ask bounce could increase the amount of volatility and spreads would be positively related to the degree of transaction volatility clustering. This hypothesis says nothing about the impact of the spread on the quote or limit order book series, however.

Next, we include both the mean and the variance of quoted depth. Many researchers have demonstrated the importance of considering depth in understanding trading behavior.[11] The basis for including the mean and variance of the quoted depth in the cross-sectional regressions is illustrated in the following simple example. Suppose that an exogenous process drives the submission of market buy and sell orders and that the sizes of these orders are distributed over some support [0, *S*]. Compare two stocks, the first with a constant depth of 5S, and the second with a constant depth of $\frac{1}{2}$S. Conditional on available information, transaction prices in a low depth stock are much more likely to move than transaction prices in a high depth stock because orders are more likely to exhaust the depth at the current quoted prices in the first stock than in the second stock. The same logic can be applied to the quote and limit order book series. The lower the quoted depth, the more likely it is that trades will exhaust the liquidity displayed in the quotes, whether that liquidity is originating from the specialist, floor or limit order book. Of course, in all cases, liquidity suppliers have the option of immediately replacing their executed orders, so the relation between the average quoted depth and the variance of these series is not entirely mechanistically driven. Furthermore, variability in the quoted depths is likely to induce more volatility clustering because the chance of exhausting the depth at the current price increases due to an increase in periodic low depth quotes. Consequently, we would expect the average level of depth to be negatively related and the variability of quoted depth to be positively related to the variance of all three series. The sensitivity of the transaction series will be negatively related to the average depth if a high depth primarily implies that prices move less in general and positively related to average depth if high depth levels allow more information to be impounded into prices.

Using similar logic, a measure of average trade size relative to the quoted depth is included in the regressions. Transactions in stocks with large average relative trade sizes may eliminate the depth at a particular price (in either quotes or the limit order book) more easily than transactions in stocks with smaller average trade sizes. Consequently, relative trade size should be positively related to the variance of all three series. Holthausen et al. (1987), Chan and Lakonishok (1995) and Keim and Madhavan (1996) find that the price impact of a transaction is increasing in trade size, which offers support for this hypothesis. Furthermore, trade size may also proxy for the type of market participant that is engaged in trading the asset. For instance, an asset with a large relative trade size may have large liquidity based traders, like mutual funds, while an asset with a small average trade size might have retail customers as the bulk of its participants. This suggests that relative trade size will be negatively associated with sensitivity and persistence levels. In contrast,

[11] See, for example, Petersen and Umlauf (1993), Knez and Ready (1996) and Kavajecz (1999).

a large relative trade size may be associated with informed trading, which implies that it will be positively related to sensitivity and persistence.

The relationship between volume and volatility has received substantial attention in the literature and, as a result, average daily trading volume is included as an explanatory variable in our analysis. For example, Blume et al. (1994) study relationships among volume, information, and price changes. Most recently, work by Gervais et al. (2000) and Mingelgrin (2000) demonstrates that trading volume contains information beyond that already contained in prices both at weekly and intra-daily frequencies. To the extent that higher trading volume is the result of an increased flow of information to the market, we would expect volume to be positively associated with both the variance and the sensitivity to incoming innovations (α) for all three series. Conversely, Easley et al. (1996) argue that information-based trading is highest for infrequently traded stocks because the arrival rate of informed traders exceeds the arrival rate of uninformed traders, which predicts that trading volume will be negatively related to overall variance.

We also expect volatility clustering to be related to the price level for each stock. Existing work suggests that price level could be either positively or negatively related to volatility. On one hand, the minimum tick size for low-priced stocks is large relative to the price level and, consequently, percentage bid-ask spreads tend to be higher. Glosten and Milgrom (1985) demonstrate that higher transaction costs deter investors from trading, which in turn reduces the amount of information incorporated into prices. This would imply that price level and potential volatility clustering are positively related. On the other hand, Christie (1982) proposes the so-called "leverage effect," which states that as a stock's price decreases, the firm's leverage increases, resulting in an increase in the volatility of equity. This implies a negative relationship between volatility and stock price level.

We investigate these hypotheses by estimating nine cross-sectional regression models of the following form:

$$\begin{aligned} y = b_0 &+ b_1\ AvgQuotedSpread + b_2\ AvgQuotedDepth \\ &+ b_3\ VarQuotedDepth + b_4\ RelTradeSize \\ &+ b_5\ AvgDailyTradingVolume + b_6\ Price + \varepsilon, \end{aligned} \tag{4}$$

where, for each return series, $y=$ the average conditional variance ($\mu(\sigma_t^2)$), the sensitivity (α), and the persistence ($\alpha + \beta$) estimated from the autoregressive model with GARCH(1,1) residuals (see Table 1).

Table 3 contains the results for the cross-sectional regression analysis. In general, the coefficient estimates are consistent with the information-based hypotheses discussed above. Specifically, the sign and significance of the bid-ask spread coefficient in the transaction equation suggests that wide spreads primarily serve to restrict the amount of information that is incorporated into

Table 3
Cross-sectional regression analysis of the conditional variance series

The table below contains the results of three measures of volatility clustering regressed upon microstructure variables. (See Eq. (4) in the text.) The table displays results for the average conditional variance of transactions, the GARCH sensitivity parameter (α) and the GARCH persistence parameter ($\alpha+\beta$). The explanatory variables are average quoted spread, average quoted depth, the variance of the quoted depth, relative trade size (average trade size divided by average quoted depth), average daily trading volume and the price level (measured at the 1990 year end). *P*-values are shown in parentheses and estimates that are significantly different from zero at the 5% level are in bold.

	Coefficient estimates (*t*-statistics)							
Dependent variable	Constant	Average quoted spread	Average quoted depth	Quoted depth variance	Relative trade size	Average trade volume	Price	R^2
Panel A: Transactions								
Conditional	**0.00089**	**−0.00498**	**−7.5E-5**	**9.5E-6**	**0.11433**	**−6.9E-10**	6.2E-6	0.49
variance	**(0.0001)**	**(0.0001)**	**(0.0001)**	**(0.0001)**	**(0.0001)**	**(0.0214)**	(0.2072)	
α	**0.06445**	−0.15126	−0.00126	3.8E-5	−0.93873	**1.6E-7**	0.00098	0.28
	(0.0076)	(0.1871)	(0.5518)	(0.8280)	(0.7429)	**(0.0001)**	(0.1446)	
$\alpha+\beta$	**0.61224**	−0.06693	−0.01300	0.00070	**−36.795**	2.0E-7	0.00013	0.11
	(0.0001)	(0.8900)	(0.1121)	(0.3411)	**(0.0028)**	(0.2365)	(0.9630)	
Panel B: Quote midpoint								
Conditional	**0.00061**	**−0.00373**	**−4.8E-5**	**6.0E-6**	**0.12706**	−3.3E-10	8.5E-11	0.33
variance	**(0.0025)**	**(0.0001)**	**(0.0030)**	**(0.0001)**	**(0.0001)**	(0.3242)	(1.000)	
α	**0.03956**	−0.06199	**0.00379**	**−0.00025**	**4.2900**	−4.1E-8	0.00048	0.12
	(0.0054)	(0.3570)	**(0.0010)**	**(0.0144)**	**(0.0117)**	(0.0824)	(0.2243)	
$\alpha+\beta$	**0.62854**	0.35377	0.00713	−0.00077	−13.495	−5.8E-8	−0.00396	0.05
	(0.0001)	(0.4616)	(0.3783)	(0.2920)	(0.2627)	(0.7286)	(0.1584)	
Panel C: Limit order book midpoint								
Conditional	**0.00056**	−0.00081	**−2.5E-5**	1.9E-6	0.02923	−9.8E-11	−5.0E-6	0.12
variance	**(0.0003)**	(0.2643)	**(0.0393)**	(0.0915)	(0.1078)	(0.6979)	(0.2378)	
α	**0.08174**	−0.06087	0.00166	−0.00013	−3.2011	−4.0E-8	0.00054	0.04
	(0.0001)	(0.5291)	(0.3089)	(0.3795)	(0.1874)	(0.2418)	(0.3363)	
$\alpha+\beta$	**0.67701**	0.53240	0.00776	−0.00078	**−29.121**	**−4.4E-7**	0.00361	0.14
	(0.0001)	(0.2371)	(0.3060)	(0.2562)	**(0.0105)**	**(0.0053)**	(0.1691)	

prices rather than permit an increase in volatility through a large bid-ask bounce. Similarly, the negative spread coefficient in the quoted variance equation implies that the specialist tends to update quotes less actively when the spread is wide. The quoted depth variables confirm the hypothesized relations, with the mean (variance) of depth negatively (positively) related to each of the conditional variance series.

Relative trade size is associated with higher conditional variance and lower persistence, which suggests that a large relative trade size is indicative of large liquidity traders. The fact that relative trade size is positively related to the sensitivity of the quote series, however, implies that relative trade size also proxies, at least to some extent, for informed trading.

Trading volume is negatively related to transaction variance, consistent with the argument made by Easley et al. (1996) that informed trading is highest for infrequently traded stocks. As in Table 1, however, we find that higher trading volume is associated with a higher sensitivity for the transaction series. The results initially seem somewhat at odds but may suggest that when information shocks *do* occur in frequently traded stocks, they induce a larger increase in transaction volatility than shocks in infrequently traded stocks. The negative relation between trading volume and persistence levels on the limit order book suggests that shocks filter through the limit order books of actively traded stocks more quickly.

Overall, we find substantial evidence that volatility clustering is related to microstructure characteristics in our sample. In general, the quoted spread appears to affect the average conditional variance of the series, while the depth, relative trade size, and trading volume variables are determinants of sensitivity and persistence levels, as well. Furthermore, the impact of market structure on characteristics of volatility is stronger for the transaction and quote series than for the limit order book series, which is consistent with limit order traders' actions being less restricted by the trading environment.

4.2. Analysis of constant conditional variance series

Despite the strong presence of GARCH effects in the sample as a whole, many stocks have constant conditional variance for at least one of the three series. Specifically, series for which the sensitivity and persistence levels do not differ significantly from zero are deemed to have constant conditional variance. In our sample, 38 percent of the 143 stocks have constant transaction variance series, 17 percent have constant quote variance series, and 17 percent have constant limit order book variance series.

This section compares the characteristics of stocks with constant and time-varying conditional variance in an attempt to understand the factors that might contribute to the *non-existence* of volatility clustering. Contingency tables for the transaction, quote, and limit order book conditional variance series are contained in Panels A–C of Table 4, respectively. Each panel contains both the actual and expected distribution of stocks with constant variance for partitions of the data based on trading volume level (high, medium, or low) and price level (high, medium, or low).

The results in Panel A demonstrate that no systematic relationships exist between the frequency of stocks with constant transaction variance and trading volume or price. In contrast, stocks with constant quote or limit order book

Table 4
Analysis of constant conditional variance series

This table reports both the actual and expected frequency of constant transaction, quote and limit order book conditional variance series for the sample partitioned by (a) trading volume, (b) price level, and (c) price level after controlling for trading volume. The reported values represent the number of stocks that display constant conditional variance for each group. *P*-values are from exact Chi-square tests of association between constant variance and the given characteristic.

Frequency		High			Medium			Low		*p*-values
Panel A: Transactions										
Actual by trading volume		15			22			17		0.3430
Actual by price level		19			17			18		0.9300
Expected		18			18			18		
Price conditional on volume	H	M	L	H	M	L	H	M	L	
Actual	4	5	6	10	5	7	6	5	6	
Expected	5	5	5	7.3	7.3	7.3	6	6	5	
p-values		0.9240			0.2430			0.9300		
Panel B: Quotes										
Actual by trading volume		7			3			14		0.0065
Actual by price level		7			6			11		0.3100
Expected		8			8			8		
Price conditional on volume	H	M	L	H	M	L	H	M	L	
Actual	5	1	1	1	1	1	2	6	6	
Expected	2.3	2.3	2.3	1	1	1	5	5	4	
p-values		0.0990			1.000			0.1800		
Panel C: Limit order book										
Actual by trading volume		3			5			16		0.0004
Actual by price level		4			5			15		0.0024
Expected		8			8			8		
Price conditional on volume	H	M	L	H	M	L	H	M	L	
Actual	3	0	0	0	1	4	5	5	6	
Expected	1	1	1	1.7	1.7	1.7	5.5	5.5	5	
p-values		0.0970			0.1100			0.8600		

variance tend to be concentrated in low trading volume and low price groups. In light of the positive correlation between trading volume and price level ($\rho = 0.46$), we also examine the role of price level after *controlling for* trading volume level. Specifically, this final partition, which consists of nine groups (high/high through low/low), is constructed by first sorting the sample into three trading volume groups and then sorting by price level *within* each volume group.[12] While low trading volume appears to be the dominant factor, there is evidence that stocks with high volume and high prices also tend to exhibit constant quote and limit order book variance more often than expected. The relations are significant at the 10% level.

Like the findings in the cross-sectional analysis, these results are consistent with the structure of the market impacting the existence of conditional volatility. The groups of stocks that have higher concentrations of constant variance for the quote and limit order book series tend to trade in a different type of market structure. Specifically, stocks with high trading volume and high price, for example stocks in the Dow Jones Industrial Average (DJIA), essentially trade in an open outcry environment since a crowd of floor traders is often present and actively trading.[13] This is consistent with Bollerslev and Domowitz (1991), who show that transaction returns for futures contracts traded in an open outcry environment do not exhibit GARCH effects. Despite the fact that stocks with low trading volume do not have a trading crowd, they still trade in a unique trading environment. This subset of stocks is known to have large limit order book spreads and low cumulative depth (see Kavajecz, 1999). Consequently, stocks with low trading volume and low price essentially trade in a call market, whereby the occasional arrival of an order prompts a trade with the specialist.

5. Relation among transaction, quote, and limit order book conditional variances

Although the static comparison of the GARCH estimates addresses relations among the volatility of the three series overall, it offers no insight into how periods of volatility clustering in each of the series are aligned in time. Specifically, we are interested in determining whether periods of high volatility for each of the series coincide – possibly with short leads or lags. In order to investigate this time series aspect of their relation, we estimate a vector

[12] Note that the conditional price level differs from the unconditional price level used in the second partition. When the unconditional price level is used in the final partition, the results are qualitatively identical, but significance levels differ slightly.

[13] Of the five (three) high trading volume, high price stocks with constant quote (limit order book) variance, four (two) were DJIA stocks.

Table 5
Vector autoregressions of limit order book, quote, and transaction conditional variance series

The table below reports the results of a set of VAR models of limit order book (LOB), quoted and transaction conditional variance series estimated for each stock from the following GARCH(1,1) model:

$$r_t = c + \phi r_{t-1} + \varepsilon_t, \quad \sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2, \quad \varepsilon_t \sim N(0, \sigma_t^2),$$

where r_t is the appropriate return series and ε_t is a normally distributed disturbance, with mean zero and variance described above. The model is estimated with four lags (15 minutes each) of each dependent variable as well as the current quoted spread and depth measure. The figures reported are the median coefficient estimates for each variable as well as the percentage of stock exhibiting significance at the 5% level (included in parentheses) for the 70 stocks in the sample displaying non-constant limit order book, quote, and transaction conditional variance series.

	Quoted		Limit order book				Quotes				Transactions			
Constant	Spread	Depth	Lag1	Lag2	Lag3	Lag4	Lag1	Lag2	Lag3	Lag4	Lag1	Lag2	Lag3	Lag4
Limit order book midpoint return series														
0.0000	0.00000	−0.0000	0.91317	−0.0037	−0.0124	−0.0086	−0.0013	−0.0237	−0.0007	−0.0055	0.0485	−0.0150	−0.0232	0.0027
(44)	(21)	(16)	(96)	(33)	(36)	(33)	(39)	(30)	(20)	(14)	(54)	(39)	(24)	(19)
Quote midpoint return series														
0.0000	0.0000	−0.0000	−0.0005	−0.0009	0.0000	0.0001	0.7210	−0.0016	0.0047	−0.0004	0.0459	−0.0153	−0.0084	0.0024
(66)	(53)	(39)	(34)	(33)	(19)	(14)	(94)	(34)	(29)	(23)	(64)	(67)	(50)	(39)
Transaction return series														
0.0000	0.0000	−0.0000	−0.0000	−0.0019	0.0008	0.0002	0.0777	−0.0463	−0.0100	0.0049	0.6724	0.0566	−0.0290	−0.0187
(77)	(37)	(26)	(27)	(19)	(11)	(10)	(56)	(40)	(20)	(11)	(91)	(56)	(53)	(34)

autoregression (VAR) using the three conditional variance series generated from the estimated GARCH(1,1) model.

Our VAR model is estimated with four lags (of 15 minutes each) of the dependent variables and includes the current bid-ask spread and depth quotes as control variables.[14] We estimate the model for the 70 stocks in the sample that display non-constant variance for all three series. Table 5 contains the median coefficient estimates across stocks and the percentage of stocks with coefficients that are statistically significant at the five-percent level (shown in parentheses).

Given the GARCH(1,1) structure used to estimate the conditional variance series, it is not surprising that the own first lags are highly significant across almost all stocks in the sample. The results demonstrate that each conditional variance series has a significant impact on the other two series, as well. For example, the first lag of transaction return volatility has a significant effect on the conditional variance of the quote midpoint returns, and vice versa, for the majority of the stocks in the sample. The positive median coefficients suggest that periods of high transaction and quote volatility often immediately follow one another (within fifteen minutes). This is consistent with the hypothesis that periods of intense trading activity are likely to reduce the amount of liquidity available in the market, thereby increasing the volatility of the quotes. The fact that recent transaction volatility also tends to imply an increase in volatility on the limit order book further supports this interpretation, since the executed transactions are likely to remove limit orders from the book.

The results in Table 5 also demonstrate that longer lags of transaction and quote return volatility tend to be negatively related to current volatility for both series, which implies that activity has a cyclical nature. Lagged quote volatility is negatively related to current limit order book volatility and the size of the quoted spread is positively related to quote volatility, although the coefficient is extremely small. This is generally consistent with work by Biais et al. (1995), which shows that wide spreads induce order placements (activity on the limit order book).

Overall, the results in Table 5 document strong relationships among all three series. While the impact of each series is multi-directional, flowing both from and to the limit order book, quotes and transactions, a stronger relationship seems to exist between quote and transaction volatility than between limit order book volatility and the other series. Furthermore, the fact that volatility on the limit order book does not filter through to quotes and transactions is consistent with research by Kavajecz and Odders-White (2001), which shows that specialists do not completely reflect changes in the limit order book when

[14] In an earlier specification, we included dummy variables for the open and close of trading, consistent with the intraday patterns found by Andersen and Bollerslev (1997); however, both dummy variables were completely uninformative.

posting quotes. One possible interpretation of these relationships, consistent with work by Corwin and Lipson (2000) and Harris and Panchapagesan (1999), is that the limit order book contains information not contained in quotes and transactions and that the limit order book is a important source of information.

6. Conclusion

By analyzing the volatility of three related yet distinct series – transaction returns, quote midpoint returns, and limit order book midpoint returns – we obtain insight into the impact of market structure on volatility and into the relations among the series. We document statistically significant GARCH effects in the transaction, quote, and limit order book series for the majority of the stocks in our sample. We further show that the sensitivity and persistence levels for the best prices on the limit order book equal or exceed those in quoted prices, which in turn surpass those in transaction prices. We extend the comparison by examining the impact of earnings announcements on the volatility of each series. The results demonstrate that the sensitivity to shocks (α) increases for the limit order book and decreases for the quote series surrounding earnings announcements, while the persistence parameters ($\alpha + \beta$) decrease for both series during the announcement period.

We examine the impact of market structure on volatility directly by estimating a set of cross-sectional regressions. We find that the average conditional variance is negatively related to trading volume and the quoted spread and depth, and is positively related to the variance of quoted depth and to relative trade size. Furthermore, the sensitivity of transaction variance to incoming innovations (α) is higher for actively-traded stocks and the persistence ($\alpha + \beta$) of shocks decreases as the average trade size (relative to the quoted depth) increases for both transactions and the book. We also show that low volume stocks and high priced, high volume stocks have a greater tendency to exhibit conditional variance series that are not time-varying.

Finally, we estimate a vector autoregressive model to analyze relations among the three series over time. The results suggest that, while there is a strong relation between volatility clustering in quote and transaction returns, the relation among limit order book returns and the other series is substantially weaker.

The results summarized above document a number of facts about volatility clustering in equity markets that we believe will be valuable in many different contexts. First, the findings constitute the first detailed description of limit order book volatility in the literature. Consequently, the limit order book results can be viewed as interesting stylized facts. They also provide additional information regarding the role of limit order books, which could be beneficial in discussions of the optimal structure of equity markets. Second, the findings offer new insight into the behavior of market participants, as well as confirming some existing results. In particular, by comparing the three series, we are better able to understand the interplay between liquidity providers who place limit orders, the specialist and floor members, and liquidity demanders who place market orders. For example, our results may suggest that, while some information is impounded into prices through the trading process, a significant portion of transaction volatility stems from liquidity shocks. Furthermore, we find evidence that liquidity suppliers may be better able than liquidity demanders to quickly process and interpret incoming information in the period surrounding corporate announcements.

Finally, as stated in the introduction, the underlying cause of volatility clustering in financial markets remains largely a mystery. Although this study in isolation cannot offer definitive conclusions about the source of volatility clustering in equity markets, the results constitute an important step towards improving our understanding of this phenomenon. Valuable insight can be gained by considering the relations among the series in conjunction with the differential effects of the market structure on each. For example, we find that both volatility levels and clustering are higher on the limit order book than in transactions, but that little of the clustering on the book can be linked to microstructure effects. Together, these facts suggest that other considerations, like the way in which information arrives and is incorporated into each series, have a substantial impact on clustering. We believe that additional research in this area will be fruitful and hope that this study has laid the groundwork for such work.

References

Andersen, T., 1996. Return volatility and trading volume: an information flow interpretation of stochastic volatility. Journal of Finance 51, 169–204.

Andersen, T., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. Journal of Empirical Finance 4, 115–118.

Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris bourse. Journal of Finance 50, 1655–1689.

Blume, L., Easley, D., O'Hara, M., 1994. Market statistics and technical analysis: the role of volume. Journal of Finance 49, 153–181.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307–327.

Bollerslev, T., Chou, R., Kroner, K., 1992. ARCH modeling in finance. Journal of Econometrics 52, 5–59.

Bollerslev, T., Domowitz, I., 1991. Price volatility, spread variability, and the role of alternative market mechanisms. Review of Futures Markets 10, 79–102.

Bollerslev, T., Engle, R., Nelson, D., 1994. Arch models. In: Engle, R., McFadden, D. (Eds.), Handbook of Econometrics, Vol. 4. Elsevier, Amsterdam.

Campbell, J., 1987. Stock returns and the term structure. Journal of Financial Economics 18, 373–399.
Campbell, J., Lo, A., MacKinlay, A., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton, NJ.
Chan, L., Lakonishok, J., 1995. The behavior of stock prices around institutional trades. Journal of Finance 50, 1147–1174.
Cheung, Y., Ng, L., 1992. Stock price dynamics and firm size: an empirical investigation. Journal of Finance 47, 1985–1997.
Christie, A., 1982. The stochastic behavior of common stock variances: value, leverage, and interest rate effects. Journal of Financial Economics 10, 407–432.
Corwin, S., Lipson, M., 2000. Order flow and liquidity around NYSE trading halts. Journal of Finance 55, 1771–1801.
Easley, D., Kiefer, N., O'Hara, M., Paperman, J., 1996. Liquidity, information and infrequently traded stocks. Journal of Finance 51, 1405–1436.
Engle, R., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U. K. inflation. Econometrica 50, 987–1008.
Engle, R., Ito, T., Lin, W., 1990. Meteor showers or heat waves ? Heteroskedastic intraday volatility in the foreign exchange market. Econometrica 58, 525–542.
Engle, R., Mustafa, C., 1992. Implied ARCH models from options prices. Journal of Econometrics 52, 289–311.
French, K., Schwert, G., Stambaugh, R., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–29.
Gervais, S., Kaniel, R., Mingelgrin, D., 2000. The high volume return premium. Working paper, The Wharton School, University of Pennsylvania.
Glosten, L., Jagannathan, R., Runkle, D., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. Journal of Finance 48, 1779–1802.
Glosten, L., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.
Harris, L., Panchapagesan, V., 1999. The information-content of the limit order book: evidence from NYSE specialist actions. Working paper, University of Southern California.
Hasbrouck, J., 1999. Trading fast and slow: security market events in real time. Working paper, New York University.
Holthausen, R., Leftwich, R., Mayers, D., 1987. The effect of large block transactions on security prices: a cross-sectional analysis. Journal of Financial Economics 19, 237–267.
Kavajecz, K., 1999. The specialist's quoted depth and the limit order book. Journal of Finance 52, 747–771.
Kavajecz, K., Odders-White, E., 2001. An examination of changes in specialists' posted price schedules. Review of Financial Studies, in preparation.
Keim, D., Madhavan, A., 1996. The upstairs market for large-block transactions: analysis and measurement of price effects. Review of Financial Studies 9, 1–36.
Knez, P., Ready, M., 1996. Estimating the profits from trading strategies. Review of Financial Studies 9, 1121–1163.
Lamoureux, C., Lastrapes, W., 1990. Heteroskedasticity in stock return data: volume versus GARCH effects. Journal of Finance 45, 221–229.
Lee, C., Mucklow, B., Ready, M., 1993. Spreads, depths, and the impact of earnings information: an intraday analysis. Review of Financial Studies 6, 345–374.
Mingelgrin, D., 2000. The informational content of high-frequency trading activity. Working paper, The Wharton School, University of Pennsylvania.
Petersen, M., Umlauf, S., 1993. Quote revisions and the quoted size: an empirical analysis. CRSP Working Paper, No. 312.