JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 53, No. 2, Apr. 2018, pp. 613–641
COPYRIGHT 2018, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S0022109017001053

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

High-Frequency Quoting: Short-Term Volatility in Bids and Offers

Joel Hasbrouck*

Abstract

At subsecond horizons, bids and offers in U.S. equity markets are more volatile than what would be implied by long-term fundamentals. To assess costs and consequences, this paper suggests that traders' random delays (latencies) interact with quote volatility to generate execution price risk and relative latency costs. Analysis of the behavior of quote setters suggests that this volatility is more likely to arise from recurrent cycles of undercutting similar to the Edgeworth cycles found in product markets rather than mixed strategies of limit-order placement.

I. Introduction

Recent developments in market technology have called attention to the practice of high-frequency trading. The term is used in broad reference to all sorts of fast-paced market activity, not just trades, but trades have certainly received the most attention. There are good reasons for this emphasis, as trades signify the actual transfers of income streams and risk. Quotes also play a significant role in trading process, however. This paper examines short-term volatility in bids and

*Hasbrouck (corresponding author), jhasbrou@stern.nyu.edu, New York University Stern School of Business. I have benefited from the comments of an anonymous referee, Hendrik Bessembinder (the editor), Fany DeKlerck, Ramo Gençay, Dale Rosenthal, Jeffrey Russell, Gideon Saar, Mao Ye, and seminar/conference participants at Aalto University, Baruch College, BI Norwegian Business School, the Conference on Financial Econometrics (Toulouse), the Copenhagen Business School, Duke University, City University London Emerging Markets Group, the 2013 Euronext Paris Conference on High-Frequency Trading, l'Institut Louis Bachelier, Jump Trading, Norges Bank Investment Management, NYU/Stern, SAC Capital, the Society of Financial Econometrics (Toronto), the University of California at Irvine, the University of Illinois at Urbana–Champaign, the University of Illinois at Chicago, the University of Pennsylvania Econometrics Seminar, the Wharton Finance Department Seminar, and Utpal Bhattacharya's doctoral students at the University of Indiana. All errors are my own responsibility. This research was not specifically supported or funded by any organization. During the period over which this research was developed, I taught (for compensation) in the training program of a firm that engages in high-frequency trading and served as a member (uncompensated) of a CFTC advisory committee on high-frequency trading. I am grateful to Jim Ramsey for originally introducing me to time-scale decompositions.

offers of U.S. equities, a consequence of what might be called high-frequency quoting.[1]

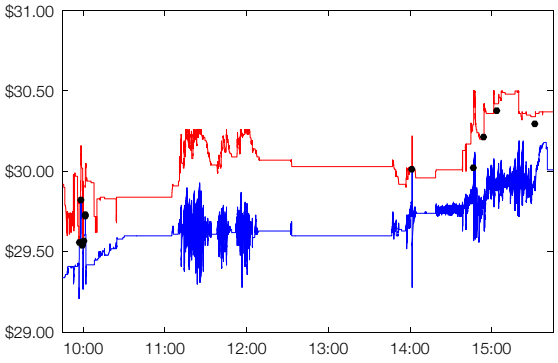
By way of illustration, Figure 1 depicts the bid (national best bid (NBB)) and offer (national best offer (NBO)) for AEP Industries (ticker symbol AEPI, a manufacturer of packaging products listed on the National Association of Securities Dealers Automated Quotation System (NASDAQ) stock exchange) on Apr. 29, 2011. In terms of broad price moves, the day is not a particularly volatile one, and the bid and offer quotes are stable for long periods. The placidity is broken, though, by several intervals where the bid undergoes extremely rapid changes. The average price levels, before, during, and after these episodes are not dramatically different. Moreover, the volatility is largely one-sided: The ask does not oscillate as strongly. Nor is the volatility associated with increased executions. These considerations suggest that the volatility is unrelated to fundamental public or private information. It appears to be an artifact of the trading process.

In the context of the paper's data sample, the AEPI episode does not represent typical behavior. Nor, however, is it a singular event. It therefore serves to motivate the paper's key questions. What is the extent of short-term quote volatility? What strategies give rise to it? What does it cost slower traders?

To characterize quote volatility, this study employs short-term variances of bid and offer levels. This approach departs from the usual reliance on variances of short-run returns. A return variance measures endpoint-to-endpoint volatility, while the level variance depends on the full path. Although equivalent for many purposes, they correspond to different types of economic risk. An agent

FIGURE 1
The Bid and Offer for AEPI

Figure 1 shows the national best bid and offer (NBBO) from the NYSE Daily TAQ data set for Apr. 29, 2011. The NBB (bottom line) is the maximum bid, taken over all market centers reporting to the Consolidated Tape Association; the NBO (top line) is the minimum offer. Dots denote trades.



[1]Jones (2013) reviews the academic work in this area. Many empirical studies of low-latency activity focus primarily on trades, and trading activity is often used to identify high-frequency traders. Recent representative studies include Angel, Harris, and Spatt (2011), (2015), Baron, Brogaard, Hagströmer and Kirilenko (2018), Brogaard, Hendershott, and Riordan (2014), Hendershott, Jones, and Menkveld (2011), Menkveld (2013), O'Hara (2015). The increasing brevity of quote lifetimes is a feature of fleeting orders (Hasbrouck and Saar (2009)).

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

who precisely knows the market price at time t and enters a marketable order that is subject to a known delay δ achieves an execution at time $t+\delta$. The variance of the difference over the interval represents execution price risk relative to the initial price. A trader who is subject to a random delay, however, faces another sort of uncertainty. If her actual execution time is uniformly distributed over $(t, t+\delta)$, she will realize an execution price somewhere on the price path over the interval. As a consequence of the arrival-time uncertainty, she can expect to execute at the average price over the interval, and the risk is measured by the mean-square deviation of the price relative to this average. This is simply the local variance of the price level.

In a sample of 2011 U.S. equity data, this article's estimates suggest that with an arrival-time uncertainty of 1.6 seconds, a typical seller would incur a risk (relative to the average price) with a standard deviation of 0.134 cents per share (or about 0.761 basis points (bps)). With additional structure, the local variances can be mapped into relative advantages realized by traders who can more tightly control their arrival-time uncertainty, presumably by use of superior technology. This calculation suggests that, relative to the 1.6-second seller, a trader who could perfectly time her order arrival would have a gain as large as 0.390 cents per share (2.221 bps).

I next consider how quote volatility might arise from the strategies and interactions of liquidity suppliers. Quote volatility is often attributed to the manipulative single-agent practices of quote stuffing (canceling and submitting orders to produce congestion and/or confusion) and spoofing (briefly exposing quotes that are not intended for execution). The present analysis cannot rule out these motives, but it also considers alternative explanations. Baruch and Glosten (2013) suggest that quote setters may simply be pursuing rational mixed (randomized) strategies. In their equilibrium, each participant repeatedly cancels and replaces their quote with a price drawn from a distribution that is stationary between informed pick-off trades.

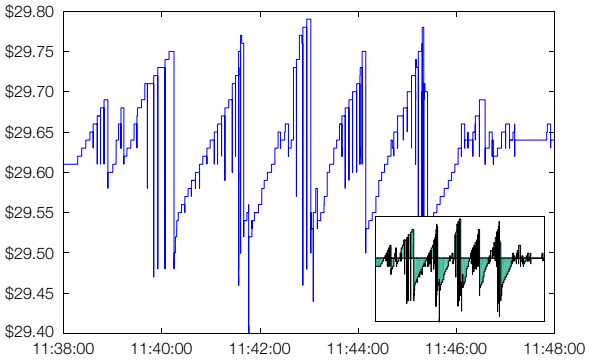
Randomized strategies suggest a quote path that resembles unstructured noise. Actual dynamics, however, often exhibit local trends and other regularities. Figure 2 presents details of the bid dynamics (as in Figure 1, for AEPI on Apr. 29, 2011) over a 10-minute window starting at 11:38AM. There are certainly intervals when the bid path could plausibly arise as a sequence of independent draws (e.g., in the vicinity of 11:47AM). The more striking feature, though, is the repeated sawtooth pattern of small successive bid improvements terminated by large drops.[2]

Baruch and Glosten (2013) follow a line of economic thought that examines randomization by sellers in product markets. These settings also sometimes exhibit Edgeworth cycles, wherein sellers incrementally undercut each other, reset to a high price, and repeat (Edgeworth (1925), Noel (2011), and Maskin and Tirole (1988)). Sellers' Edgeworth cycles have been widely studied in retail gasoline

[2] Although quotes can change as a direct consequence of incoming marketable orders, this particular the 10-minute window includes no executions. The bid changes in the 10-minute window are solely due to placement, cancellation and revision of limit orders. Three exchanges are often present at the best bid (NASDAQ, Arca, and BATS), but within this group NASDAQ is clearly dominant. The large drops reflect bid revisions grouped closely in time (but not exactly simultaneous).

FIGURE 2
AEPI National Best Bid Detail

Figure 2 shows the national best bid (NBB) from the NYSE Daily TAQ data set from Apr. 29, 2011. The inset graph depicts the area centered around the mean of the interval ($29.64).



markets; buyers' cycles (as in the AEPI bid) have been found in online advertising auctions (Zhang and Feng (2011), Zhang (2005)). Mixed strategies and Edgeworth cycles offer alternative explanations for quote volatility that are rational, competitive, and nonmanipulative.

This paper attempts a partial empirical resolution. Single-agent manipulative and mixed-strategy models generally predict an inverse relation between competition and quote volatility. This paper finds that the empirical association is generally positive, suggesting that these mechanisms are not predominant. Edgeworth cycles are characterized by skewness in price changes (negative for bids, positive for offers). This study finds that quote volatility is associated with more extreme skewness, a result supportive of the Edgeworth mechanism.

Two recent papers focus on quote volatility. The present study most closely resembles Conrad, Wahal, and Xiang (CWX) (2015). In a sample of U.S. stocks, CWX find that increased quoting activity is associated with improved efficiency and liquidity, but also that quote volatility is not as large as implied by present estimates. This difference in findings can be explained by several considerations. Most importantly, CWX examine volatility in the bid–ask midpoint, which essentially averages volatility across the two sides of the market. The present study treats the bid and ask separately, leaving undamped the one-sided volatility such as that found in the AEPI example. Relative to CWX, the present study examines somewhat shorter time horizons and also includes more small firms, for which quote volatility is higher. Finally, the volatility measures differ between the two studies. CWX consider variances of price differences; the present study examines variances of price levels. In another study, Egginton, Van Ness, and Van Ness (2016) investigate quote stuffing, which they define as intense rates of order submission and cancellation, and which they proxy by number of quote updates. In a 2010 sample, they find that standard measures of liquidity are lower during quote-stuffing episodes. The present results are broadly supportive of CWX in that they

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

are consistent with economic mechanisms that are nonmanipulative, nondisruptive, and familiar from standard product market settings.

The paper is organized as follows: Section II examines the costs that quote volatility imposes on liquidity demanders. Section III discusses measurement of quote volatility and the relation to traditional return variance ratios. The paper then turns to empirical analysis. Section IV describes the sample and data. Section V presents the results for variance ratios and discusses reconciliation with the CWX (2015) findings. Taking the perspective of liquidity demanders, Section VI examines the risk borne by market-order traders and the costs of latency. Section VII discusses the behaviors of liquidity suppliers that might give rise to quote volatility. Section VIII summarizes the findings and indicates directions for further research.

II. Economic Costs of Quote Volatility for Liquidity Demanders

A market-order seller receives a price equal to the market bid prevailing at the time that the order is actually executed. This price can be viewed as the outcome of two stochastic processes. The first process governs the evolution of the bid quote; the second describes the timing of the execution event. The execution timing is random and depends on communication and cognitive delays (latencies). In the AEPI detail of Figure 2, for example, an order that arrives between 11:38AM and 11:48AM will receive an execution price somewhere on the path. Conditional on this path, if the execution time is uniformly distributed over the interval, the seller can expect to receive the average price (which happens to be $29.64).

The timing uncertainty gives rise to two effects of potential economic significance: execution price risk and relative advantage for faster traders. The execution price risk is illustrated by the deviation between the path and the average, indicated by the shaded area in the Figure 2 inset graph. The mean-square deviation is approximately $0.06 (about 20 bps). This quantity bears further emphasis. It is a local standard deviation around a local average of the price level. It is neither an observation nor a standard deviation of a serial price difference or return.

Latencies are likely to be smaller for more sophisticated traders or those who have invested in technology. Figure 2 suggests that the local standard deviation about the local average is generally smaller for shorter subintervals. A faster trader essentially faces a narrower window of timing uncertainty and incurs lower risk. The faster trader also enjoys a relative advantage in expected outcome. Conditional on the given path, a hypothetical seller active in the interval with perfect execution timing could pick out the maximum bid: $29.79, a gain of $0.15 relative to the average. Between perfect timing and 10-minute uncertainty there may exist, of course, traders of intermediate latency, whose risk and gain lie between the two extremes.

The assumption that the path is known assigns all variation in outcomes to timing uncertainty. This simplifies the analysis, but even the idealized fast seller would face at the start of the interval the difficulty of predicting the price path. Given this limitation, the $0.15 gain is likely to overstate an actual seller's advantage.

The prediction problem may be less severe, however, than it might initially seem. Firstly, a stationary process is more predictable than a random walk, and the short-term volatility of the AEPI bid appears to be dominated by a stationary component. Secondly, although prediction is usually viewed as forecasting what has yet to occur, it is functionally equivalent to more timely knowledge of what has in fact already happened. Fast traders' advance knowledge of fundamental information in Foucault, Hombert, and Rosu (2016) and Biais, Foucault, and Moinas (2015) can be viewed in this light. Finally, for some participants, the predictability may arise from a look-back option (Stoll and Schenzler (2006)).[3]

III. Measuring Quote Volatility

Denote by p_t a price series in discrete time where t indexes the shortest time intervals in the analysis (50 milliseconds (ms), in this study). The risk and cost analysis described here focuses on local variances of deviations (residuals) around local means of price levels. The length of the relevant averaging period, however, will depend on the range of a trader's latency uncertainty. Since encountering wide variation is unlikely in this respect, it is sensible to consider a wide range of possible times. These are constructed systematically. Starting from the shortest (50 ms) interval, longer intervals are generated by powers of 2: $2^j \times 50$ ms for $j = 0, \ldots, J = 16$. The longest averaging period is therefore $2^{16} \times 50$ ms $= 3{,}276.8$ sec ≈ 55 min.

The residual price variance about an average of length $2^j \times 50$ ms is denoted σ_j^2. As the averaging periods increase with j, it will also be useful to consider incremental increases in variances, $\nu_j^2 = \sigma_j^2 - \sigma_{j-1}^2$, for $j = 1, \ldots$ and $\nu_0^2 = \sigma_0^2$. The incremental variances are of interest because they pick up variation attributed to a particular time scale. In Figure 2, for example, if σ_j^2 corresponds to the local residual variance over the 10-minute window depicted in the figure (10:38 to 10:48), then σ_{j-1}^2 would correspond to the local residual variance relative to successive 5-minute means (10:38 to 10:43 and 10:43 to 10:48), and ν_j^2 would measure the variance of the change in the 5-minute means. The time scale of the variation is $\tau_j = 2^{j-1}$. Appendix A describes the computation of σ_j^2 and ν_j^2 for the special cases of discrete and continuous-time random walks.

To maintain the intuition in the present application, this paper will continue to refer to σ_j^2 and ν_j^2 as the (local) residual variance and incremental variance, respectively. In the signal-processing and time scale–decomposition literatures, however, they are usually described as rough and detail (or wavelet) variances. Most statistical results, efficient computational algorithms, and connections to traditional time-series analysis are developed in this broader framework. Percival and Walden (2000) provide a comprehensive textbook development.[4] In finance,

[3]Most retail orders are executed by market makers who act as counterparty at a price equal to the prevailing bid or ask. Stoll and Schenzler (2006) note that if the timing of the prevailing bid or ask cannot be verified, the market maker may try to choose the price within an interval that is most favorable to them: a recent minimum if the customer is selling or a recent maximum if the customer is buying.

[4]In addition to Percival and Walden (2000), Gençay, Selçuk, and Whitcher (2002) discuss economic and financial applications in the broader context of filtering. Nason (2008) discusses time series and other applications of wavelets in statistics. Ramsey (1999), (2002) provides other useful economic

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

time-scale decompositions have been used to explore stock comovements and volatility persistence.[5]

In the present application, the local residual standard deviation, σ_j, is most directly interpreted as a measure of execution price risk subject to timing uncertainty. Appendix B shows that, under additional structure, the running sum of the incremental volatilities, $G_j = v_0 + \ldots + v_j$, can be viewed as the execution price advantage possessed by a market-order trader with perfect timing relative to a slower trader subject to random latency that is uniform over an interval of length $2^j \times 50$ ms. Intuitively, over this interval, the faster trader can pick out the best price, while the slower trader can only expect to obtain the interval average.

The σ_j and G_j measures are computed and reported in cents per share or bps, to facilitate assessment of their economic significance. To gauge their magnitude relative to the long-run risk of the security, however, and to establish a connection with other studies, it is useful to have transformations that benchmark σ_j to a random walk. This is accomplished using variance ratios.

There is a long tradition of variance ratios in empirical market microstructure (Barnea (1974), Hasbrouck and Schwartz (1988), Amihud and Mendelson (1987), and CWX (2015), among others). The typical ratio compares the variance of price differences at two different time scales. For consistency and comparability with the measures described, these time scales are set to follow the same indexing map: $\tau_j = 2^{j-1} \times 50$ ms, f or $j = 1, \ldots, J = 16$, the longest time scale in the analysis. The difference, $\Delta_j p_t$, is correspondingly defined as $\Delta_j p_t = p_t - p_{t-\tau_j}$. In this notation, the long/short variance ratio is

$$(1) \qquad V\Delta_j \quad = \quad \frac{\text{var}\left(\Delta_J p_t\right)\tau_j}{\text{var}\left(\Delta_j p_t\right)\tau_J} \quad = \quad 2^{j-J}\frac{\text{var}\left(\Delta_J p_t\right)}{\text{var}\left(\Delta_j p_t\right)}.$$

If p_t follows a random walk at all horizons, variance is linear in time, and $V\Delta_j = 1$ for all j. In typical microstructure settings, variance ratios with $j \ll J$ are generally depressed due to short-term microstructure effects.

The numerator in equation (1) is the scaled variance of a long-term price change, implicitly a measure of the long-term informational volatility. It serves to normalize the variance of the short-run variance in the denominator. The denominator is the same sort of quantity as the numerator, the variance of a price difference, simply at a shorter scale. The essence of a variance ratio, however, is normalization; this can often be accomplished even when the variances are of different types. It is particularly useful to consider a variance ratio involving the

and financial perspectives. Walker (2008) is clear and concise but oriented more toward engineering applications.

[5] Gençay, Selçuk, and Whitcher (2005) and In and Kim (2006) decompose stock comovements at daily and longer horizons. Studies that use wavelet methods to characterize volatility persistence include Dacorogna, Gençay, Muller, Olsen, and Pictet (2001), Elder and Jin (2007), Gençay, Selçuk, Gradojevic, and Whitcher (2010), Gençay et al. (2002), Høg and Lunde (2003), and Teyssière and Abry (2007). These studies generally involve absolute or squared returns at minute or longer horizons. Wavelet methods have also proven useful for jump detection and jump volatility modeling (Fan and Wang (2007)).

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

short-run residual and long-run incremental variances:

$$\text{VR}_j = \frac{\left(2^j - 1\right) v_j^2}{2^{J-1}\sigma_j^2}. \tag{2}$$

The $\left(2^j - 1\right)/2^{J-1}$ factor corresponds to the 2^{j-J} adjustment factor in the traditional variance ratio. For a random walk in continuous time, $\text{VR}_j = 1$ (Appendix A).

Although VΔ and VR are both normalized to unity for a random walk, they differ in other key respects. The denominator in VΔ is the variance of a short-term price difference, while that in VR is an interval variance of a price level. This distinction carries an important practical implication. In the case of the traditional variance ratio at the longest time scale, $\text{V}\Delta_J$ is unity by construction for all processes and all estimates, irrespective of whether they are random walks. The longest-scale residual variance ratio, VR_J, however, need not equal 1. In most of the cases encountered in the present analysis, short-term interval variance is elevated relative to what would be expected in a random walk, and $\text{VR}_J < 1$.

The numerators in $\text{V}\Delta_j$ and VR_j differ in a more subtle way. Both can be viewed as estimates of the same thing: the long-term fundamental variance of the price process. The traditional variance ratio is based on the τ_J-horizon price change, under the assumption that a long-run difference is dominated by fundamental movements. The long-run difference still impounds, however, variation from short-term effects present at the endpoints. For example, an hourly price change in a trade price may contain a large component arising from bid–ask bounce and discreteness in the starting and ending prices. The incremental variance, v_j^2, used in the numerator of VR_j is implicitly also a long-run difference, but prior to differencing, the prices are locally averaged to dampen the short-run endpoint effects. At macroeconomic time scales the distinction between $\text{var}\left(\Delta_{\tau_j} p\right)$ and v_j^2 is likely to be negligible, but at subdaily time scales (particularly for low-cap firms), the relative magnitudes of the short-run effects may well be high.[6]

IV. Sample and Data

The sample studied here is a subsample of U.S. firms, using quote data from Apr. 2011 (the first month of my institution's subscription). The subsample is constructed from all firms present on the Center for Research in Security Prices (CRSP) daily stock file and the NYSE's Daily trade and quote (TAQ) database from Jan.–Apr. 2011 with share codes of 10 or 11, with closing prices between \$2 and \$1,000, and with a primary listing on the New York, AMEX, or NASDAQ exchanges.[7] I compute the daily average dollar volume based on trading in January through March and randomly select 15 firms from each decile. For brevity, reported results are grouped into quintiles.

[6] Fan and Gençay (2010) apply this principle to unit-root tests based on time-scale decompositions. Gençay and Signori (2015) explore the use of variance ratios at different time scales to test for serial correlation.

[7] The American Stock Exchange merged with NYSE Euronext in 2008 and was renamed NYSE AMEX LLC. In May 2012, the name was changed to NYSE MKT LLC. For the sake of clarity, it is identified here simply as "AMEX."

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

The U.S. equity market is highly fragmented, but all exchanges post their quotes to the Consolidated Quotation System (CQS).[8] The consolidated quote (CQ) and national best bid and offer (NBBO) files from the New York Stock Exchange (NYSE) Daily TAQ data set used here are definitive transcripts of the consolidated activity, time-stamped to the millisecond.[9] A record in the CQ file contains the latest bid and offer originating at a particular exchange. If the bid and offer establish the NBBO, this fact is noted on the record. If the CQ record causes the NBBO to change for some other reason, a message is posted to another file (the NBBO file). Thus, the NBBO can be obtained by merging the CQ and NBBO files. It can also be constructed directly from the CQ file. Spot checks confirm that these two approaches are consistent. Transaction data are taken from the Daily TAQ consolidated trade (CT) file.

Studies involving TAQ data have traditionally used error filters to suppress quotes that appear spurious. Recent Daily TAQ data, though, appear to be much cleaner than older samples. In particular, the NBBO construction provided by the NYSE clearly defines what market participants would have perceived. Some quotes present in the CQ file are not incorporated into the NBBO because they are not firm, indicative, or otherwise deemed "not NBBO-eligible." Beyond these exclusions, however, I impose no additional filters.

Table 1 reports summary statistics. Following the U.S. Securities and Exchange Commission's adoption of Regulation NMS, U.S. exchanges have become more similar in structures and trading mechanisms. With respect to listing characteristics, though, differences persist. The NYSE has the largest proportion of

TABLE 1
Descriptive Statistics

Table 1 reports descriptive statistics of a sample of 150 firms randomly selected from CRSP with stratification based on average dollar trading volume in the first quarter of 2011, grouped in quintiles by dollar trading volume. NBB is the national best bid; NBO the national best offer. Except for the counts (first 4 rows), all table entries are cross-firm medians. *Source*: CRSP and Daily TAQ data (Apr. 2011).

		Dollar Trading Volume Quintile				
Variable	Full Sample	1 (Low)	2	3	4	5 (High)
No. of firms	149	29	30	30	30	30
NYSE	47	0	5	7	16	19
AMEX	6	2	2	0	1	1
NASDAQ	96	27	23	23	13	10
Avg. daily CT records (trades)	1,346	33	431	1,126	3,478	16,987
Avg. daily CQ records (quotes)	24,053	1,067	7,706	24,026	53,080	181,457
Avg. daily NBBO records	7,203	354	3,029	7,543	16,026	46,050
Avg. daily NBB changes	1,265	121	511	1,351	2,415	4,124
Avg. daily NBO changes	1,179	106	460	1,361	2,421	4,214
Avg. price (bid–ask midpoint)	$15.77	$4.76	$5.46	$17.86	$27.76	$51.60
Market capitalization of equity (in $millions)	690	41	202	747	1,502	8,739

[8] At the same time that an exchange sends a quote update to the consolidated system, it can also transmit the update directly to its own subscribers. Eliminating the delay associated with consolidation and retransmission saves approximately 5 milliseconds. Thus, while the CQS is a widely used single source of market data, it is not the fastest. Moreover, bids and offers with sizes under 100 shares are not reported.

[9] The "daily" reference in the Daily TAQ data set refers to the release frequency. Each morning the NYSE posts files that cover the previous day's trading. The Monthly TAQ data set, more commonly used by academics, is released monthly and contains time stamps in seconds.

high-volume stocks, NYSE AMEX has the smallest, and NASDAQ falls in the middle. In some instances, a stock that is present in CRSP and TAQ's master file is absent from a day's quote file. Stocks missing more than half of the days in the sample are dropped.

Market-event counts (trades, quotes, and so forth) display some interesting patterns. There are large numbers of quote records, since one is generated when any market center changes its best bid, best offer, or size at the bid or offer. If the action establishes the bid and offer as the NBBO, this fact is noted on the quote record. But if the action causes some other change in the aggregate prices or sizes at the NBBO, an NBBO record is generated. Since many quote records do not induce such a change, there are substantially fewer NBBO records. Finally, many records reflect updates of only one side of the market.

Volatility and spreads tend to be elevated at the start and end of trading sessions (9:30 to 16:00). To remove the effect of these deterministic effects, I confine the variance estimates to the 9:45 to 15:45 subperiod. Estimates are computed separately for the bid, ask, and bid–ask midpoint (BAM). To minimize the effect of outliers, most samples are winsorized at $\pm 5\%$ prior to computation of statistics. This is conservative: In most cases the winsorization changes estimates in directions that are prejudicial to the paper's findings.

V. Variance Ratio Estimates

Variance ratios provide a familiar point of departure. Table 2 summarizes sample averages for price-change and residual variance ratios for time scales $\tau_j = 2^{j-1} \times 50\,\text{ms}$ for $j = 1, \ldots, 16$. Table 3 reports results for a subset of time scales (50 ms, 800 ms, and 27.3 min) but with standard errors and for subsamples constructed as quintiles in average dollar trading volume.

TABLE 2
Quote Variance Ratios

Table 2 reports estimates for 150 U.S. firms during Apr. 2011. The time scale is $\tau_j = 2^{j-1} \times 50\,\text{ms}$, $j = 0, \ldots, J = 16$, the longest time scale in the analysis. The standard price-difference variance ratio is $\text{V}\Delta_j = 2^{j-J}\text{var}(\Delta_{\tau_j} p_t)/\text{var}(\Delta_{\tau_J} p_t)$, where Δ_τ is the differencing operator, $\Delta_\tau p_t = p_t - p_{t-\tau}$, and p_t is the bid, ask, or BAM. The residual variance ratio is $\text{VR}_j = (2^j - 1)\nu_j^2/(2^{J-1})\sigma_J^2$, where σ_j^2 is the residual variance of the price level around an average of length $2^j \times 50\,\text{ms}$, and $\nu_j^2 = \sigma_j^2 - \sigma_{j-1}^2$ is the incremental variance at time scale J. For a random walk, all variance ratios should be unity for all j (see Appendix A). The variance ratios are estimated daily. Table entries are means across days and firms; the sample is winsorized at ±5%.

Level, j	Time Scale	$\text{V}\Delta_j^{\text{Bid}}$	$\text{V}\Delta_j^{\text{Ask}}$	$\text{V}\Delta_j^{\text{BAM}}$	VR_j^{Bid}	VR_j^{Ask}	VR_j^{BAM}
1	50 ms	0.557	0.561	0.775	0.581	0.578	0.787
2	100 ms	0.569	0.573	0.767	0.588	0.585	0.783
3	200 ms	0.584	0.590	0.766	0.599	0.597	0.780
4	400 ms	0.602	0.609	0.771	0.614	0.613	0.784
5	800 ms	0.621	0.628	0.778	0.630	0.629	0.791
6	1,600 ms	0.639	0.647	0.785	0.648	0.647	0.798
7	3.2 sec	0.658	0.666	0.791	0.665	0.665	0.804
8	6.4 sec	0.676	0.685	0.793	0.682	0.684	0.807
9	12.8 sec	0.695	0.702	0.797	0.700	0.703	0.810
10	25.6 sec	0.716	0.723	0.801	0.719	0.722	0.813
11	51.2 sec	0.740	0.745	0.809	0.742	0.744	0.820
12	102.4 sec	0.767	0.771	0.818	0.769	0.767	0.831
13	3.4 min	0.805	0.805	0.840	0.797	0.794	0.843
14	6.8 min	0.845	0.843	0.868	0.825	0.820	0.859
15	13.7 min	0.898	0.897	0.910	0.849	0.845	0.872
16	27.3 min	1.000	1.000	1.000	0.883	0.881	0.899

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

TABLE 3

Quote Variance Ratios in Volume Subsamples

Table 3 reports estimates for 150 U.S. firms during Apr. 2011 and in subsamples stratified on average dollar volume in Mar. 2011. The time scale is $\tau_j = 2^{j-1} \times 50$ ms, $j = 0, \ldots, J = 16$, the longest time scale in the analysis. The standard price-difference variance ratio is $V\Delta_j = 2^{j-J}\text{var}(\Delta_{\tau_j} p_t)/\text{var}(\Delta_{\tau_1} p_t)$, where Δ_τ is the differencing operator, $\Delta_\tau p_t = p_t - p_{t-\tau}$ and p_t is the bid, ask, or BAM. The residual variance ratio is $VR_j = (2^j - 1)v_j^2/(2^{J-1})\sigma_J^2$, where σ_j^2 is the residual variance of the price level around an average of length $2^j \times 50$ ms and $v_j^2 = \sigma_J^2 - \sigma_{j-1}^2$ is the incremental variance at time scale J. For a random walk, all variance ratios should be unity for all j (see Appendix A). The variance ratios are estimated daily. Table entries are means across days and firms; the sample is winsorized at ±5%. Standard errors are given in parentheses.

Sample	Level, j	Time Scale	$V\Delta_j^{Bid}$	$V\Delta_j^{Ask}$	$V\Delta_j^{BAM}$	VR_j^{Bid}	VR_j^{Ask}	VR_j^{BAM}
Full sample	1	50 ms	0.557	0.561	0.775	0.581	0.578	0.787
			(0.006)	(0.006)	(0.008)	(0.007)	(0.007)	(0.010)
	5	800 ms	0.621	0.628	0.778	0.630	0.629	0.791
			(0.006)	(0.006)	(0.008)	(0.008)	(0.008)	(0.009)
	16	27.3 min	1.000	1.000	1.000	0.883	0.881	0.899
						(0.003)	(0.003)	(0.003)
1 (low)	1	50 ms	0.458	0.484	0.572	0.472	0.470	0.569
			(0.015)	(0.015)	(0.017)	(0.019)	(0.019)	(0.021)
	5	800 ms	0.495	0.521	0.572	0.497	0.494	0.567
			(0.015)	(0.015)	(0.016)	(0.019)	(0.018)	(0.020)
	16	27.3 min	1.000	1.000	1.000	0.817	0.806	0.844
						(0.009)	(0.009)	(0.008)
2	1	50 ms	0.430	0.438	0.644	0.453	0.463	0.675
			(0.012)	(0.012)	(0.017)	(0.014)	(0.015)	(0.020)
	5	800 ms	0.480	0.496	0.647	0.493	0.509	0.675
			(0.012)	(0.013)	(0.016)	(0.015)	(0.015)	(0.019)
	16	27.3 min	1.000	1.000	1.000	0.874	0.874	0.903
						(0.008)	(0.008)	(0.008)
3	1	50 ms	0.519	0.504	0.761	0.542	0.526	0.779
			(0.013)	(0.013)	(0.018)	(0.015)	(0.015)	(0.021)
	5	800 ms	0.577	0.568	0.747	0.588	0.574	0.766
			(0.013)	(0.013)	(0.016)	(0.016)	(0.016)	(0.020)
	16	27.3 min	1.000	1.000	1.000	0.898	0.899	0.913
						(0.007)	(0.007)	(0.007)
4	1	50 ms	0.636	0.631	0.907	0.678	0.670	0.937
			(0.013)	(0.014)	(0.019)	(0.017)	(0.017)	(0.023)
	5	800 ms	0.720	0.719	0.930	0.743	0.739	0.958
			(0.015)	(0.015)	(0.019)	(0.018)	(0.018)	(0.023)
	16	27.3 min	1.000	1.000	1.000	0.924	0.924	0.931
						(0.007)	(0.007)	(0.007)
5 (high)	1	50 ms	0.738	0.742	0.988	0.752	0.753	0.972
			(0.013)	(0.013)	(0.018)	(0.016)	(0.016)	(0.022)
	5	800 ms	0.825	0.830	0.993	0.823	0.824	0.989
			(0.014)	(0.014)	(0.018)	(0.017)	(0.017)	(0.021)
	16	27.3 min	1.000	1.000	1.000	0.899	0.898	0.902
						(0.007)	(0.007)	(0.007)

The first 3 columns of Table 2 report the standard price-change variance ratios for the bid, offer, and bid-ask midpoint (BAM). In each case, the pattern across time scales is familiar. The ratios at short intervals are low and rise as the interval increases, converging (by construction) to unity for the longest horizon. At the 50-ms time scale, the mean ratio is 0.557 for bids, implying that the variance of 50-ms price changes is about 80% higher than it would be for a random walk with variance calibrated to the 27.3-minute horizon ($1/0.557 = 1.80$). The estimates of $V\Delta$ for bids and asks are very similar, implying that bids and asks experience similar elevation of short-run volatility. Compared to bids or offers, though, variance ratios based on the midpoint are closer to unity. The difference is particularly pronounced at the shorter time scales: At 50 ms, $V\Delta^{BAM} = 0.775$.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

The AEPI plots discussed exhibit a pattern of bid and offer volatility that is largely one-sided. The sharp differences in $V\Delta$ estimates for the bid/ask versus the midpoint suggest that this is a general characteristic.

The rightmost 3 columns of Table 2 report estimates of the residual variance ratios, the VRs. To gauge the approximate magnitude, at time scales below 1 second, $VR^{Bid} \approx VR^{Ask} \approx 0.6$, implying level variances 67% higher than a random walk calibrated to 27-minute returns. For the midpoint ratios, however, $VR^{BAM} \approx 0.8$, implying a more modest 25% elevation. As noted in the earlier discussion, $V\Delta$ and VR are capturing different aspects of the processes. This is most apparent at the longest horizon. While $V\Delta_J = 1$ by construction, the mean estimate for $V\Delta_J$ is about 0.88 (on the bid and ask sides). Intuitively, over an interval of $2 \times 27.3 = 54.6$ minutes, the variance of the bid (or offer) process about the mean level is 13.6% higher than would be expected from a random walk.

Table 3 reports variance ratio estimates for a subset of the time scales, but across volume subsamples and with standard errors. The patterns across time, and in respect to bid or ask versus the midpoint, are similar to those found in the broader sample. As might be expected, the implied departures from random-walk behavior are larger for the low-activity subsamples (quintiles 1 and 2).

The variance ratio estimates here are somewhat more extreme than those reported by CWX (2015). In their Table 2, for example, CWX report average variance ratios of approximately 0.8 for the smallest quintile; the present study finds averages for least-active quintile (at 50 and 800 ms, Table 3) to be about 0.4. These differences can largely be explained by several considerations. Most important are the sample differences. CWX impose an equity market capitalization minimum of $100 million and use market-capitalization quintiles defined on NYSE breakpoints. The present study samples over all NMS stocks and constructs quintiles on dollar trading volume. Thus, relative to CWX, the present sample includes a greater proportion of small and less-actively traded issues, for which quote volatility tends to be more extreme. A second consideration involves time scales. The variance ratios in CWX (Table 2) compare 15-second returns with 15-minute returns; the variance ratios estimated here compare shorter time scales (down to 50 ms) with 27-minute returns. Microstructure effects of most sorts generally become more important at shorter time scales. Thirdly, CWX compute variance ratios using bid–ask midpoints, which generate estimates closer to unity than those obtained for the bid or the ask analyzed separately.

Bid, ask, and BAM volatilities are economically important in different trading situations. Traders who use market orders and those whose orders are price-matched to the NBBO are exposed to bid volatility if they are selling and to ask volatility if they are buying. Traders who use dark pools with midpoint pricing, however, are exposed to volatility in the midpoint.

VI. Risk and Relative Advantages Associated with Latency Uncertainty

Table 4 summarizes the sample means for residual volatilities and implied cumulative latency gains in units of cents per share and bps. Table 5 reports similar statistics for three representative time scales, but with standard errors and also in

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

TABLE 4
Residual Volatilities and Relative Gains

Table 4 provides a sample of 150 U.S. firms during Apr. 2011. The time scale is $\tau_j = 2^{j-1} \times 50$ ms, $j = 0, \ldots, J = 16$, the longest time scale in the analysis. The residual volatility, σ_j, is the standard deviation of the price level (bid, ask, or BAM) around a local average of length $2^j \times 50$ ms. Cumulative gain $G_j = \sum_{i=1}^{j} \nu_i$, where $\nu_j = \sqrt{\sigma_j^2 - \sigma_{j-1}^2}$ is the expected gain of a hypothetical fast trader (who could perfectly time his order arrival) relative to a slower trader (who could only time his order arrival to an interval of length $2^j \times 50$ ms) (Appendix B). Estimates are formed daily for each firm and for the bid, ask, and BAM. Values for bid and ask include variation within the 50-ms window; BAM values do not include this variation. Table entries are means across firms and days. Panel A reports estimates in units of cents per share; Panel B reports estimates in basis points.

Panel A. Residual Volatilities and Gains in Cents per Share

Level, j	Time Scale	σ_j^{Bid}	σ_j^{Ask}	σ_j^{BAM}	G_j^{Bid}	G_j^{Ask}	G_j^{BAM}
0	<50 ms	0.026	0.027		0.027	0.028	
1	50 ms	0.036	0.037	0.021	0.077	0.080	0.043
2	100 ms	0.050	0.052	0.036	0.122	0.126	0.081
3	200 ms	0.070	0.072	0.055	0.184	0.190	0.135
4	400 ms	0.097	0.100	0.080	0.270	0.279	0.210
5	800 ms	0.134	0.138	0.114	0.390	0.402	0.316
6	1,600 ms	0.186	0.191	0.161	0.557	0.573	0.465
7	3.2 sec	0.259	0.266	0.228	0.789	0.811	0.676
8	6.4 sec	0.360	0.370	0.322	1.112	1.143	0.976
9	12.8 sec	0.500	0.514	0.456	1.560	1.604	1.397
10	25.6 sec	0.694	0.713	0.644	2.181	2.242	1.991
11	51.2 sec	0.960	0.987	0.908	3.037	3.122	2.825
12	102.4 sec	1.328	1.365	1.274	4.215	4.336	3.994
13	3.4 min	1.832	1.883	1.776	5.845	6.009	5.614
14	6.8 min	2.521	2.597	2.474	8.093	8.319	7.847
15	13.7 min	3.477	3.576	3.436	11.154	11.463	10.940
16	27.3 min	4.755	4.887	4.727	15.273	15.675	15.051

Panel B. Residual Volatilities and Gains in Basis Points

Level, j	Time Scale	σ_j^{Bid}	σ_j^{Ask}	σ_j^{BAM}	G_j^{Bid}	G_j^{Ask}	G_j^{BAM}
0	<50 ms	0.148	0.149		0.157	0.161	
1	50 ms	0.205	0.207	0.116	0.439	0.445	0.239
2	100 ms	0.284	0.287	0.200	0.693	0.703	0.452
3	200 ms	0.395	0.399	0.305	1.046	1.060	0.753
4	400 ms	0.548	0.553	0.446	1.536	1.555	1.176
5	800 ms	0.761	0.767	0.638	2.221	2.244	1.771
6	1,600 ms	1.055	1.064	0.904	3.165	3.194	2.604
7	3.2 sec	1.461	1.470	1.269	4.464	4.501	3.761
8	6.4 sec	2.020	2.033	1.775	6.243	6.296	5.361
9	12.8 sec	2.778	2.801	2.477	8.672	8.773	7.589
10	25.6 sec	3.827	3.862	3.449	12.004	12.148	10.686
11	51.2 sec	5.253	5.323	4.804	16.597	16.826	14.993
12	102.4 sec	7.176	7.308	6.682	22.837	23.230	20.994
13	3.4 min	9.794	10.008	9.250	31.334	31.965	29.260
14	6.8 min	13.393	13.723	12.824	42.933	43.875	40.591
15	13.7 min	18.315	18.752	17.696	58.750	60.092	56.199
16	27.3 min	24.908	25.387	24.276	79.773	81.436	77.141

volume-based subsamples. Note that the time scales in these tables follow the same indexing scheme used earlier for the variance ratios: $\tau_j = 2^{j-1} \times 50$ ms is the time scale of the differencing; σ_j measures variation around a local mean of length $2 \times \tau_j$.

Panel A of Table 4, row for level $j = 5$, for example, suggests that a hypothetical market-order seller facing arrival-time uncertainty of $2^5 \times 50$ ms $= 1.6$ sec is exposed to bid volatility of 0.134 cents per share or (from the corresponding row in Panel B) 0.761 bps. The volatility increases at longer time scales. Arrival-time uncertainty of 54.6 minutes is associated with risk of almost 5 cents (or 25 bps).

How should the economic significance of these magnitudes be assessed? Many trading fees (e.g., commissions and clearing fees) are assessed on a per share basis. Access fees, the charges levied by exchanges on taker

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

(aggressor) sides of executions are also assessed per share. U.S. SEC Regulation NMS caps access fees at 0.3 cents per share, and in practice, most exchanges are close to this level. Exchanges also pay liquidity rebates that transfer a portion of the access fee to the maker (passive) side of the trade, generally about 0.2 cents per share. Variations in these fees are regarded as important determinants of order-routing decisions. As an alternative benchmark, a recent survey estimates average institutional commissions for U.S. equities at about 6 bps (ITG Inc. (2014)). The volatilities of latency risk are of comparable magnitude.

TABLE 5

Residual Volatilities and Gains in Average Dollar Volume Subsamples

The sample in Table 5 is 150 U.S. firms during Apr. 2011, in a full sample and subsamples constructed on average dollar volume. The time scale is $\tau_j = 2^{j-1} \times 50$ ms, $j = 0, \ldots, J = 16$, the longest time scale in the analysis. The residual volatility, σ_j, is the standard deviation of the price level (bid, ask, or BAM) around a local average of length, $2^j \times 50$ ms. Cumulative gain $G_j = \sum_{i=1}^{j} \nu_i$, where $\nu_j = \sqrt{\sigma_j^2 - \sigma_{j-1}^2}$ is the expected gain of a hypothetical fast trader (who could perfectly time his order arrival) relative to a slower trader (who could only time his order arrival to an interval of length $2^j \times 50$ ms) (Appendix B). Estimates are formed daily for each firm and for the bid, ask, and BAM. Values for bid and ask include variation within the 50-ms window; BAM values do not include this variation. Standard errors are given in parentheses.

Panel A. Residual Volatilities and Gains in Cents Per Share

Sample	Level, j	Time Scale	σ_j^{Bid}	σ_j^{Ask}	σ_j^{BAM}	G_j^{Bid}	G_j^{Ask}	G_j^{BAM}
Full sample	1	50 ms	0.036	0.037	0.021	0.077	0.080	0.043
			(<0.001)	(<0.001)	(<0.001)	(0.001)	(0.001)	(<0.001)
	5	800 ms	0.134	0.138	0.114	0.390	0.402	0.316
			(0.001)	(0.002)	(0.001)	(0.005)	(0.006)	(0.004)
	16	27.3 min	4.755	4.887	4.727	15.273	15.675	15.051
			(0.075)	(0.076)	(0.075)	(0.253)	(0.256)	(0.248)
1 (low)	1	50 ms	0.017	0.020	0.010	0.035	0.043	0.021
			(<0.001)	(<0.001)	(<0.001)	(0.001)	(0.002)	(0.001)
	5	800 ms	0.063	0.075	0.055	0.179	0.215	0.152
			(0.002)	(0.003)	(0.002)	(0.007)	(0.010)	(0.006)
	16	27.3 min	1.744	2.271	1.780	5.504	7.121	5.641
			(0.075)	(0.108)	(0.078)	(0.244)	(0.344)	(0.244)
2	1	50 ms	0.024	0.025	0.013	0.052	0.054	0.027
			(<0.001)	(0.001)	(<0.001)	(0.002)	(0.002)	(0.001)
	5	800 ms	0.089	0.091	0.072	0.259	0.267	0.200
			(0.003)	(0.003)	(0.002)	(0.010)	(0.011)	(0.007)
	16	27.3 min	2.734	2.786	2.674	8.763	8.939	8.425
			(0.114)	(0.119)	(0.114)	(0.364)	(0.384)	(0.357)
3	1	50 ms	0.039	0.039	0.021	0.082	0.083	0.043
			(0.001)	(0.001)	(<0.001)	(0.002)	(0.002)	(0.001)
	5	800 ms	0.141	0.144	0.117	0.410	0.416	0.321
			(0.004)	(0.004)	(0.003)	(0.011)	(0.012)	(0.008)
	16	27.3 min	4.626	4.646	4.586	14.768	14.829	14.419
			(0.138)	(0.138)	(0.137)	(0.440)	(0.441)	(0.428)
4	1	50 ms	0.048	0.049	0.027	0.103	0.105	0.056
			(0.001)	(0.001)	(<0.001)	(0.002)	(0.002)	(0.001)
	5	800 ms	0.178	0.181	0.151	0.523	0.529	0.419
			(0.004)	(0.004)	(0.003)	(0.014)	(0.014)	(0.010)
	16	27.3 min	6.690	6.724	6.672	21.720	21.810	21.389
			(0.180)	(0.182)	(0.180)	(0.627)	(0.633)	(0.615)
5 (high)	1	50 ms	0.053	0.053	0.031	0.114	0.114	0.065
			(0.001)	(0.001)	(<0.001)	(0.002)	(0.002)	(0.001)
	5	800 ms	0.197	0.198	0.173	0.581	0.582	0.485
			(0.004)	(0.004)	(0.003)	(0.012)	(0.013)	(0.009)
	16	27.3 min	7.926	7.955	7.923	25.608	25.677	25.380
			(0.165)	(0.167)	(0.165)	(0.586)	(0.593)	(0.577)

(continued on next page)

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

TABLE 5 (continued)
Residual Volatilities and Gains in Average Dollar Volume Subsamples

Panel B. Residual Volatilities and Gains in Basis Points

Sample	Level, j	Time Scale	σ_j^{Bid}	σ_j^{Ask}	σ_j^{BAM}	G_j^{Bid}	G_j^{Ask}	G_j^{BAM}
Full sample	1	50 ms	0.205	0.207	0.116	0.439	0.445	0.239
			(0.002)	(0.002)	(0.001)	(0.005)	(0.006)	(0.003)
	5	800 ms	0.761	0.767	0.638	2.221	2.244	1.771
			(0.008)	(0.008)	(0.006)	(0.027)	(0.027)	(0.020)
	16	27.3 min	24.908	25.387	24.276	79.773	81.436	77.141
			(0.268)	(0.271)	(0.252)	(0.897)	(0.917)	(0.832)
1 (low)	1	50 ms	0.280	0.287	0.153	0.610	0.635	0.334
			(0.007)	(0.007)	(0.003)	(0.019)	(0.022)	(0.011)
	5	800 ms	1.036	1.071	0.849	3.061	3.173	2.409
			(0.028)	(0.025)	(0.020)	(0.092)	(0.091)	(0.065)
	16	27.3 min	28.708	30.943	26.789	93.030	100.632	87.101
			(0.797)	(0.790)	(0.697)	(2.725)	(2.758)	(2.355)
2	1	50 ms	0.277	0.277	0.152	0.593	0.594	0.312
			(0.005)	(0.005)	(0.002)	(0.012)	(0.012)	(0.006)
	5	800 ms	1.024	1.016	0.832	2.984	2.974	2.307
			(0.020)	(0.019)	(0.015)	(0.061)	(0.060)	(0.046)
	16	27.3 min	30.760	30.929	29.866	100.097	100.670	95.913
			(0.640)	(0.651)	(0.625)	(2.174)	(2.227)	(2.119)
3	1	50 ms	0.195	0.197	0.109	0.411	0.416	0.219
			(0.003)	(0.003)	(0.002)	(0.008)	(0.008)	(0.004)
	5	800 ms	0.723	0.726	0.608	2.093	2.104	1.672
			(0.014)	(0.014)	(0.012)	(0.042)	(0.043)	(0.034)
	16	27.3 min	23.826	23.854	23.557	76.094	76.205	74.367
			(0.493)	(0.493)	(0.482)	(1.593)	(1.601)	(1.549)
4	1	50 ms	0.159	0.160	0.093	0.335	0.337	0.186
			(0.002)	(0.002)	(0.001)	(0.005)	(0.005)	(0.002)
	5	800 ms	0.593	0.594	0.510	1.710	1.715	1.398
			(0.009)	(0.008)	(0.008)	(0.026)	(0.025)	(0.022)
	16	27.3 min	22.972	22.955	22.840	72.329	72.296	71.341
			(0.447)	(0.449)	(0.441)	(1.391)	(1.399)	(1.370)
5 (high)	1	50 ms	0.115	0.115	0.071	0.243	0.243	0.142
			(0.002)	(0.002)	(0.001)	(0.004)	(0.004)	(0.002)
	5	800 ms	0.436	0.436	0.392	1.255	1.255	1.071
			(0.007)	(0.007)	(0.006)	(0.023)	(0.022)	(0.019)
	16	27.3 min	18.345	18.367	18.325	57.317	57.375	56.982
			(0.410)	(0.409)	(0.406)	(1.260)	(1.259)	(1.244)

The timing model discussed in Section II and Appendix B suggests a mapping from the time-scale volatilities to relative latency advantages. This model implies that the incremental volatility, v_j, can be interpreted as the advantage, in terms of higher selling price or lower buying price, realized by a market-order trader who can time her arrivals to a $2^{j-1} \times 50$-ms window relative to a slower trader whose arrival-time window is twice as long $\left(2^j \times 50\,\text{ms}\right)$. The accumulated volatility, $G_j = v_0 + v_1 + \ldots + v_j$, measures the relative advantage held by a trader who can perfectly time her order arrivals relative to a slower trader who faces arrival-time uncertainty of $2 \times \tau_j$. For example, from the $j = 5$ row of Panel A in Table 4, a seller with perfect timing gains about 0.390 cents per share relative to a trader with arrival-time uncertainty of 1.6 seconds. The corresponding value (in Panel B) is 2.221 bps, about one-third of the average institutional commission. Estimates for ask volatilities are close to the bid values. Due to the averaging remarked on in the discussion of the variance ratios, estimated risks and gains for the midpoint are lower than those for the bid and ask values.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

Table 5 reports estimates in volume quintile subsamples. Taking the $j=5$ estimates as representative, in moving from the low- to high-volume quintiles, the cumulative gains measured in cents per share tend to increase, while those measure in bps tend to decrease. Low-volume stocks tend to have high relative volatility and latency costs, with the latter approximately 3 bps. Some of the cumulative gain estimates at the longest time scale are quite large: 93.030 bps (bid) and 100.632 bps (ask) for the lowest-volume quintile, or 25.608 cents per share (bid) and 25.677 cents per share (ask) for the highest-volume quintile. These represent gains, however, for a fast trader who knows the price path over a relatively long horizon (27.3 minutes).

In interpreting the latency gain estimates, it should be borne in mind that, while the primitives of the quote process, σ_j and ν_j, are estimated in a fairly robust statistical framework, the model mapping these estimates to costs is stylized and involves a number of additional assumptions about the structure and predictability of the bid and offer processes.

VII. Strategies of Liquidity Suppliers

In surveying the possible economic causes of quote volatility, it should first be noted that any standard static microstructure model can generate quote volatility if I introduce corresponding time variation in the underlying parameters. It is unlikely, though, that such determinants as clearing costs, interest rates, dealer risk aversion, or informed trading probabilities or fundamental volatility oscillate at subsecond horizons.

More plausible explanations arise from dynamic strategies. When there is one agent bidding or offering, or if the distance to the next best bid or offer is large, the agent has great latitude to maintain or vary the quote. Some practices that might arise in this case, such as quote stuffing and some forms of spoofing, are often cited as major concerns associated with high-frequency trading. Single-agent strategies need not, however, have manipulative intent. A single seller, for example, might experiment with different offers in an attempt to discern demand elasticity (Leach and Madhavan (1992), (1993)). In fact, Leach and Madhavan suggest that such experimentation will only occur in a market with a single monopolistic dealer. (Multiple dealers give rise to a free-rider problem.) Alternatively, a single seller who experiments by monotonically decreasing the offering price might simply be running an impromptu clock auction. The asking price starts high and declines at a constant rate. The auction ends when the offer is hit ("mine") or when the seller's reservation price is reached.[10] Whether motivated by manipulation or experimentation, however, these strategies are feasible only when the bid or offer is set by one dominant agent.

[10]Periodic single-price call auctions have been suggested to alleviate the perceived inequities of high-frequency trading (Schwartz and Wu (2013), Budish, Cramton, and Shim (2015)). These proposals envision consolidated auctions coordinated by an exchange. As long as potential buyers have access, however, a seller can run a clock auction unilaterally, with no formal coordination or consolidation. It should also be noted that though the auction might fail because the seller's reservation price is not met, the intent to execute is sincere, and the strategy is not manipulative.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

With competing liquidity suppliers, I consider two sorts of equilibrium outcomes: mixed strategy and Edgeworth cycles. Baruch and Glosten (2013) discuss the former. In their equilibrium, identical sellers quickly and repeatedly draw a new offer price from a common distribution. The best offer at any given time is the minimum (i.e., the first-order statistic of the extant sample of draws). Until a trade occurs, the mixing distribution of limit prices is time invariant, and the sequence of best offers is independent and identically distributed (IID). Baruch and Glosten show that as the number of sellers increases, the offer prices converge to the schedule predicted by the competitive (i.e., nonstrategic) equilibrium. They note that

> At the individual level, quotes are entered and cancelled quickly, yet at the aggregate level, and with sufficient competition, the National best bid and offer need not change very much. On the other hand if there are very few competing liquidity suppliers the best bid and offer can be far more volatile (p. 28).

This suggests that quote volatility should be inversely related to the number of liquidity suppliers.

To test this hypothesis, I construct time series of quote volatility and proxies for the number of competing liquidity suppliers over successive 10-minute windows. I then examine the dynamic correlations (averaged across stocks) and also estimate linear models that relate these measures.

The TAQ data do not report the number of competing quote setters. The data do identify exchanges, however. Exchanges partially aggregate individual activity, and so exchange competition may plausibly proxy for competition among individuals. This study therefore uses Herfindahl–Hirschman indices (HHIs) computed from reporting exchanges as an inverse proxy for competition among agents. These HHIs are constructed for four measures of quote provision: time at best, time alone at best, time-weighted size at best, and the number of quote improvements.

Exchange competition is usually viewed as reflecting production and pricing of exchange services, but these do not generally vary intraday. Instead, the present use of exchange-based measures to proxy for competition among agents relies on presumptions that agents are distributed across exchanges, that an exchange's quoting presence is primarily determined by the agents active on that exchange, and that short-run variation in this presence reflects the variation in number and competitiveness of these agents. These assumptions are difficult to verify, however, and they might be undermined by institutional considerations. The exchange-based HHIs would be poor proxies if the orders establishing the NBBO were routed primarily to one exchange. In that case, however, one would not expect to see much variation in exchange HHIs. Since time priority is generally observed within (but not across) exchanges, there is an incentive to direct an order to the exchange with the shortest queue. This would promote distribution of orders across exchanges. Order routing is also determined, on the other hand, by payment for order-flow practices and by maker/taker rebates and fees. The latter can also induce errors of attribution because the best bids and offers used here to compute the HHIs are not measured net of these fees and rebates. For these

reasons, the adequacy of exchange-based HHIs as proxies for competitiveness at the level of individual agents remains an open question. Accordingly, the present results should be viewed as suggestive and tentative rather than definitive.

The time at best HHI is defined on the ask side as follows: For a given firm and 10-minute time interval, let m_j^{AtBest} denote the number of milliseconds in the interval that exchange j's offer is at (i.e., either sets or matches) the national best offer. Then,

$$\text{HHI}^{\text{AtBest}} = \sum_j \left(\frac{m_j^{\text{AtBest}}}{\sum_k m_k^{\text{AtBest}}} \right)^2 . \tag{3}$$

Time alone at best, $\text{HHI}^{\text{Alone}}$, is defined similarly but is based on the time that exchange j's offer is alone at (i.e., sets) the national best offer. HHI^{Size} is computed using exchange j's time at best ask weighted by j's size (the number of shares j posts at the ask). The quote-improvement variant, HHI^{Imp}, is computed using the number of quote improvements (reductions in the national best offer) initiated by exchange j. All four measures are computed separately and analogously on the bid side.

As an alternative to the mixed-strategy models, I consider Edgeworth cycles. Empirical identification of Edgeworth cycles is typically based on features of the asymmetrical sawtooth pattern of the price path (as exhibited by AEPI in Figure 2). On the suppliers' side of the market, successive undercutting generates many successive small price drops, while the reset consists of a single large price increase. In widely studied retail gasoline markets, some studies focus on timing. Castanias and Johnson (1993) examine the runs in positive and negative price changes. Wills-Johnson and Bloch (2010) propose a sophisticated spectral measure. Other approaches are based on the skewness of the price changes (positive for sellers, negative for buyers). Commonly used measures include the median price change (Lewis (2009), Doyle, Muehlegger, and Samphantharak (2010), and Lewis and Noel (2011)) and the difference between the mean positive and mean negative price changes (Eckert (2002)).

The present study uses estimates of Pearson's moment skewness coefficient, defined for the random price changes as $\gamma = E(\Delta p - E\Delta p)^3 / [E(\Delta p - E\Delta p)^2]^{3/2}$ and, as an alternative measure, the normalized mean-less-median (the nonparametric skew), denoted here by $\text{NMLM} = \left(E\Delta p - \Delta p^{\text{median}}\right)/\sigma_{\Delta p}$. For a hypothetical offer-side cycle consisting of n 1-tick steps down followed by 1 n-tick step up, $\gamma = (n-1)/\sqrt{n}$, and $\text{NMLM} = 1/\sqrt{n}$. A cycle on the bid side (n 1-tick steps up and 1 n-tick step down) of the market would generate negative skewness of the same magnitude. To minimize noise in estimates, intervals with fewer than 30 price changes are dropped.

Skewness may also arise from other considerations. CWX (2015) examine quote volatility associated with liquidity drawdowns (i.e., sweeps of either the bid or offer side of the market within the same millisecond). To the extent that they can arise from executions deep into the books of the participating markets, they would be expected to induce sudden drops in the bid or increases in the ask. Drawdowns are defined here in a manner similar to CWX: trades occurring within the same millisecond on one side of the market.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

The four HHIs, skewness measures, and the number of drawdowns are computed for each firm (indexed by i), each side $s \in \{\text{bid}, \text{ask}\}$, over 10-minute time windows (indexed by t). High-frequency quoting is measured by the residual volatility over the window at a time scale of 12.8 seconds, denoted $\sigma_{j=8}$. For each firm, I estimate the Pearson correlations between the variables. Table 6 reports the cross-sectional (across-firm) average of these correlations. Results for Spearman correlations (not reported) are similar.

The HHIs are generally at least moderately positively correlated, suggesting that they are picking up similar sorts of variation. They are also, on both the bid and ask sides, negatively correlated with the residual volatility, $\sigma_{j=8}$. This suggests that high-frequency quoting is associated with increased competition, which runs counter to the prediction of the mixed-strategy mechanism. To model the association in a more formal multivariate framework, I specify a general linear model of the form

TABLE 6

Correlations of Measures Estimated over 10-Minute Windows

Table 6 reports cross-firm means. Pearson correlations are estimated over time for each firm; $\sigma_{j=8}$ is the residual volatility of the price level (bid or ask) around a local mean of length $2^8 \times 50\,\text{ms} \approx 12.8\,\text{sec}$. The HHIs are Herfindahl–Hirschman indices constructed across exchanges within each 10-minute window: $\text{HHI}^{\text{AtBest}}$ measures exchanges' contributions to time at the best quote; $\text{HHI}^{\text{Alone}}$, time spent alone at a quote; HHI^{Size}, time-weighted size at the best quote; HHI^{Imp}, the number of times the best quote was improved. γ is the usual (Pearson) estimate of skewness of price changes. NMLM is the normalized mean-less-median price change; n^{DD} is the number of liquidity drawdowns over the interval. Values in parentheses are p-values against the null of no correlation.

Variable		$\sigma_{j=8}$	$\text{HHI}^{\text{AtBest}}$	$\text{HHI}^{\text{Alone}}$	HHI^{Size}	HHI^{Imp}	γ	NMLM	n^{DD}
Bid	$\sigma_{j=8}$	1.000	−0.045 (<0.01)	−0.290 (<0.01)	−0.109 (<0.01)	−0.220 (<0.01)	−0.074 (<0.01)	−0.044 (<0.01)	0.286 (<0.01)
	$\text{HHI}^{\text{AtBest}}$	−0.045 (<0.01)	1.000	0.314 (<0.01)	0.704 (<0.01)	0.192 (<0.01)	0.058 (<0.01)	0.043 (<0.01)	−0.092 (<0.01)
	$\text{HHI}^{\text{Alone}}$	−0.290 (<0.01)	0.314 (<0.01)	1.000	0.337 (<0.01)	0.346 (<0.01)	0.022 (<0.01)	0.033 (<0.01)	−0.124 (<0.01)
	HHI^{Size}	−0.109 (<0.01)	0.704 (<0.01)	0.337 (<0.01)	1.000	0.202 (<0.01)	0.061 (<0.01)	0.052 (<0.01)	−0.075 (<0.01)
	HHI^{Imp}	−0.220 (<0.01)	0.192 (<0.01)	0.346 (<0.01)	0.202 (<0.01)	1.000	0.158 (<0.01)	0.106 (<0.01)	−0.141 (<0.01)
	γ	−0.074 (<0.01)	0.058 (<0.01)	0.022 (<0.01)	0.061 (<0.01)	0.158 (<0.01)	1.000	0.555 (<0.01)	−0.055 (<0.01)
	NMLM	−0.044 (<0.01)	0.043 (<0.01)	0.033 (<0.01)	0.052 (<0.01)	0.106 (<0.01)	0.555 (<0.01)	1.000	−0.087 (<0.01)
	n^{DD}	0.286 (<0.01)	−0.092 (<0.01)	−0.124 (<0.01)	−0.075 (<0.01)	−0.141 (<0.01)	−0.055 (<0.01)	−0.087 (<0.01)	1.000
Ask	$\sigma_{j=8}$	1.000	−0.050 (<0.01)	−0.285 (<0.01)	−0.112 (<0.01)	−0.215 (<0.01)	0.082 (<0.01)	0.042 (<0.01)	0.316 (<0.01)
	$\text{HHI}^{\text{AtBest}}$	−0.050 (<0.01)	1.000	0.322 (<0.01)	0.703 (<0.01)	0.198 (<0.01)	−0.061 (<0.01)	−0.049 (<0.01)	−0.094 (<0.01)
	$\text{HHI}^{\text{Alone}}$	−0.285 (<0.01)	0.322 (<0.01)	1.000	0.342 (<0.01)	0.346 (<0.01)	−0.018 (<0.01)	−0.027 (<0.01)	−0.129 (<0.01)
	HHI^{Size}	−0.112 (<0.01)	0.703 (<0.01)	0.342 (<0.01)	1.000	0.204 (<0.01)	−0.065 (<0.01)	−0.064 (<0.01)	−0.079 (<0.01)
	HHI^{Imp}	−0.215 (<0.01)	0.198 (<0.01)	0.346 (<0.01)	0.204 (<0.01)	1.000	−0.157 (<0.01)	−0.107 (<0.01)	−0.148 (<0.01)
	γ	0.082 (<0.01)	−0.061 (<0.01)	−0.018 (<0.01)	−0.065 (<0.01)	−0.157 (<0.01)	1.000	0.557 (<0.01)	0.056 (<0.01)
	NMLM	0.042 (<0.01)	−0.049 (<0.01)	−0.027 (<0.01)	−0.064 (<0.01)	−0.107 (<0.01)	0.557 (<0.01)	1.000	0.076 (<0.01)
	n^{DD}	0.316 (<0.01)	−0.094 (<0.01)	−0.129 (<0.01)	−0.079 (<0.01)	−0.148 (<0.01)	0.056 (<0.01)	0.076 (<0.01)	1.000

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

$$\text{HHI}_{its} = X_{its}\beta + e_{its}, \tag{4}$$

where HHI is a placeholder for one of the four HHI measures, $i = 1, \ldots, 150$ indexes firms, t indexes 10-minute date/time intervals, and $s \in \{\text{bid}, \text{ask}\}$. (The bid and ask sides of the market are treated as separate observations.) On the right-hand side, X_{its} is a vector of explanatory variables, β is the associated coefficient vector, and e_{its} is the regression error.

The explanatory variables includes control fixed effects (dummies for firm and date/time) and a proxy for high-frequency quoting. To enhance comparability across firms, the proxy is constructed as an indicator variable. This indicator, denoted HFQ_{its} is equal to 1 if the $\sigma_{j=8}$ (for firm i, date/time t, and side s) lies at or above the 90th percentile of the distribution (for firm i and side s) over all date/time intervals, and 0 otherwise.

Panel A of Table 7 reports the estimated least squares means associated with HFQ_{its}. The values in the top row, for example, imply that the mean $\text{HHI}^{\text{AtBest}}$ in the full sample is 0.363 when $\text{HFQ}_{its} = 0$ and 0.327 when $\text{HFQ}_{its} = 1$. For the full sample, the presence of HFQ is associated with a decrease in all HHI variants. In moving from low- to high-volume subsamples, mean HHIs decline in both HFQ states, confirming that quote-setting for high-volume stocks is more competitive. Panel B of Table 7 reports estimates of effects (the coefficients of the HFQ indicator variable) with standard errors. At customary confidence levels, these are significant. Thus, for all subsamples and all HHI measures, the HFQ indicator is associated with lower mean HHIs. This suggests that more competition is associated with increased quote volatility, not less.

I now turn to the skewness measures, intended to proxy for cyclic price behavior. From Table 6, the residual quote volatility is negatively correlated with skewness on the bid side and positively correlated with skewness on the offer side. This is the case for both skewness measures (γ and NMLM) and for Pearson and Spearman correlations. To clarify the association of the skewness measures with HFQ, I estimate a linear model of the form

$$\text{Skewness}_{its} = X_{its}\beta + e_{its} \tag{5}$$

for firm i in 10-minute time window t on side $s \in \{\text{bid}, \text{ask}\}$, where Skewness is a placeholder for γ or NMLM. The specification includes the indicator HFQ_{its}, described previously, and also a high-drawdown indicator, HDD_{its}. Analogously to HFQ_{its}, the high drawdown indicator $\text{HDD}_{its} = 1$ if the number of drawdowns (for firm i, date/time t, and side s) lies at or above the 90th percentile of the distribution (for firm i and side s) over all t. Because skewness is asymmetric across market sides, all variables and fixed effects are interacted with side (bid or ask).

Panel A of Table 8 reports estimated means predicted by the model. The first 2 columns report the effect of $\text{HFQ}_{it,s=bid}$. The first two values imply, for example, that the indicator is associated with more negative values of the product moment skewness in bid changes (from -0.162 to -0.266), suggestive of higher intensity (frequency or magnitude) of bidders' cycles. On the ask side of the market (top row right, full sample), the indicator's effect is in the direction of increased positive skewness, which is consistent with sellers' cycles. On both bid and ask sides,

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

TABLE 7
High-Frequency Quoting and Competition

Table 7 reports estimates of least squares means of Herfindahl–Hirschman indices (HHIs) implied by the linear fixed-effects panel model $(\cdot)_{its} = X_{its}\beta + e_{its}$, where $i = 1, \ldots, 150$ indexes firms, $t = 1, \ldots, 20 \times 36$ indexes 10-minute date/time intervals between 9:45 and 15:45 and trading days in Apr. 2011 (20 days $\times$36 intervals/day), and side $s \in$ {bid,ask}. (The bid and offer sides of the market are treated as separate observations.) HHIs are Herfindahl–Hirschman indices constructed across exchanges within each 10-minute window: HHI^{AtBest} measures exchanges' contributions to time at the best quote; HHI^{Alone}, time spent alone at a quote; HHI^{Size}, time-weighted size at the best quote; HHI^{Imp}, the number of times the best quote was improved (an increase in the best bid or a decrease in the best offer). The specification includes firm and date/time fixed effects. Panel A reports the implied means associated with an HFQ indicator variable set to unity if the quote volatility in the 10-minute window (as measured by $\sigma^2_{j=8}$) is at or above the 90th distribution percentile for the firm. Panel B reports estimated effects (i.e., coefficient estimates) of the HFQ indicator. The table reports estimates for the full sample and for quintile subsamples based on dollar trading volume. Values in parentheses are standard errors.

Panel A. Estimated Least Squares Means

	Dependent Variable							
	HHI^{AtBest}		HHI^{Alone}		HHI^{Size}		HHI^{Imp}	
Sample	0	1	0	1	0	1	0	1
Full sample	0.363 (<0.001)	0.327 (0.001)	0.443 (<0.001)	0.393 (0.001)	0.722 (<0.001)	0.586 (0.001)	0.616 (0.001)	0.503 (0.001)
1 (low)	0.726 (0.001)	0.613 (0.004)	0.789 (0.001)	0.673 (0.003)	0.962 (0.001)	0.822 (0.002)	0.934 (0.003)	0.788 (0.004)
2	0.384 (0.001)	0.340 (0.002)	0.482 (0.001)	0.419 (0.003)	0.832 (0.001)	0.647 (0.003)	0.734 (0.002)	0.530 (0.004)
3	0.295 (0.001)	0.271 (0.002)	0.381 (0.001)	0.342 (0.002)	0.713 (0.001)	0.567 (0.003)	0.595 (0.001)	0.470 (0.004)
4	0.236 (<0.001)	0.228 (0.001)	0.315 (0.001)	0.296 (0.002)	0.590 (0.001)	0.503 (0.004)	0.451 (0.001)	0.402 (0.003)
5 (high)	0.175 (<0.001)	0.173 (0.001)	0.249 (<0.001)	0.237 (0.001)	0.510 (0.001)	0.428 (0.004)	0.381 (0.001)	0.351 (0.003)

Panel B. Estimated HFQ Effects

Sample	HHI^{AtBest}	HHI^{Alone}	HHI^{Size}	HHI^{Imp}
Full sample	−0.037 (0.001)	−0.051 (0.001)	−0.136 (0.001)	−0.113 (0.002)
1 (low)	−0.114 (0.004)	−0.116 (0.004)	−0.139 (0.002)	−0.145 (0.005)
2	−0.044 (0.003)	−0.063 (0.003)	−0.185 (0.003)	−0.204 (0.004)
3	−0.024 (0.002)	−0.039 (0.002)	−0.146 (0.004)	−0.125 (0.004)
4	−0.008 (0.001)	−0.020 (0.002)	−0.087 (0.004)	−0.049 (0.003)
5 (high)	−0.002 (0.001)	−0.012 (0.001)	−0.083 (0.004)	−0.030 (0.003)

the means and mean shifts associated with the drawdown indicator HDD_{its} are in the same direction as those of HFQ_{its} and of comparable magnitude.

The results across volume subsamples are generally but not uniformly consistent with this pattern. First, note that the skewness measures are not monotonic, with the larger magnitudes (negative for bids, positive for asks) occurring in the middle quintiles. When $HFQ_{its} = 1$, skewness is larger, with the exception of the bid side of the highest-volume quintile. The alternative skewness measure, NMLM, behaves in a similar fashion.

The estimated means reported in Panel A of Table 8 provide the most intuitive summary of the HFQ_{its} and HDD_{its} effects, but they do not measure the marginal impacts. If high-frequency quoting and drawdowns are correlated, their influences may be confounded. Panel B of Table 8 reports estimated effects

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

(coefficients of HFQ_{its} and HDD_{its} in the linear model). For both skewness measures and in most subsamples, the effects of HFQ_{its} are generally statistically significant and larger than the effects of HDD_{its}. (The notable exception is the highest-volume quintile.)

In summary, the analysis of HHIs suggests a positive association between high-frequency quoting and the degree of competition, which is not supportive of manipulative single-agent and/or mixed multiple-agent strategies. High-frequency quoting does appear to be associated with more extreme skewness, a characteristic of Edgeworth cycles.

The differences in support for the mixed strategies versus Edgeworth cycles might be considered somewhat surprising because both mechanisms are driven by the same behavior, namely competitive undercutting. Why might one arise but not the other? One difference is that models of Edgeworth cycles assume a discrete-price grid and the mixed-strategy models assume a continuous price space. Thus,

TABLE 8

High-Frequency Quoting and Asymmetry in Quote Changes

Table 8 reports estimates of the relation between high-frequency quoting and skewness in quote changes. The specification is a linear fixed-effects panel model $(\cdot)_{its} = X_{its}\beta + e_{its}$, where $i = 1,\ldots,150$ indexes firms, $t = 1,\ldots,20 \times 36$ indexes 10-minute date/time intervals between 9:45 and 15:45 and trading days in Apr. 2011 (20 days $\times$36 intervals/day), and side $s \in \{\text{bid}, \text{ask}\}$. The dependent variable is either the standard product moment estimate of skewness, γ_{its}, or $NMLM_{its}$, the normalized mean less the median. The specification includes fixed effects for firm and time and interaction effects for high-frequency quoting and sweeps: HFQ_{its} is an HFQ indicator variable set to unity if the quote volatility in the 10-minute window (as measured by $\sigma^2_{i=8}$) is at or above the 90th distribution percentile for the firm; HDD_{its} is similarly defined for the number of liquidity drawdowns over the interval. Both variables are interacted with side *s*. Panel A reports the least squares means associated with the indicated effects; Panel B reports estimated effects (i.e., coefficient estimates) of the HFQ and HDD indicators. The table reports estimates for the full sample and for quintile subsamples based on dollar trading volume. Values in parentheses are standard errors.

Panel A. Estimated Least Squares Means

		Bid				Ask			
		HFQ		HDD		HFQ		HDD	
Dependent Variable	Sample	0	1	0	1	0	1	0	1
γ	Full	−0.162	−0.266	−0.206	−0.222	0.173	0.282	0.219	0.236
		(0.007)	(0.008)	(0.005)	(0.009)	(0.006)	(0.008)	(0.005)	(0.009)
	1 (low)	−0.061	−0.252	−0.125	−0.189	0.174	0.317	0.139	0.352
		(0.055)	(0.053)	(0.027)	(0.095)	(0.054)	(0.051)	(0.026)	(0.091)
	2	−0.221	−0.419	−0.294	−0.345	0.203	0.414	0.305	0.312
		(0.016)	(0.020)	(0.010)	(0.028)	(0.015)	(0.020)	(0.010)	(0.027)
	3	−0.251	−0.400	−0.319	−0.332	0.267	0.408	0.317	0.359
		(0.012)	(0.018)	(0.009)	(0.022)	(0.011)	(0.018)	(0.009)	(0.020)
	4	−0.169	−0.207	−0.174	−0.202	0.188	0.236	0.194	0.230
		(0.008)	(0.014)	(0.007)	(0.014)	(0.008)	(0.014)	(0.007)	(0.015)
	5 (high)	−0.095	−0.057	−0.061	−0.091	0.086	0.104	0.084	0.106
		(0.005)	(0.011)	(0.006)	(0.010)	(0.005)	(0.011)	(0.006)	(0.010)
NMLM	Full	−0.196	−0.230	−0.140	−0.286	0.184	0.204	0.136	0.252
		(0.007)	(0.008)	(0.005)	(0.009)	(0.006)	(0.008)	(0.005)	(0.009)
	1 (low)	−0.167	−0.231	−0.125	−0.273	0.140	0.172	0.105	0.206
		(0.038)	(0.037)	(0.019)	(0.066)	(0.037)	(0.035)	(0.018)	(0.063)
	2	−0.183	−0.205	−0.124	−0.264	0.183	0.199	0.135	0.248
		(0.011)	(0.014)	(0.007)	(0.019)	(0.011)	(0.014)	(0.007)	(0.019)
	3	−0.216	−0.285	−0.194	−0.306	0.189	0.248	0.187	0.251
		(0.009)	(0.013)	(0.007)	(0.016)	(0.008)	(0.013)	(0.007)	(0.015)
	4	−0.217	−0.224	−0.139	−0.302	0.215	0.211	0.149	0.277
		(0.009)	(0.016)	(0.008)	(0.016)	(0.009)	(0.016)	(0.008)	(0.017)
	5 (high)	−0.182	−0.177	−0.102	−0.257	0.170	0.150	0.082	0.238
		(0.010)	(0.019)	(0.010)	(0.017)	(0.010)	(0.019)	(0.010)	(0.017)

(continued on next page)

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

TABLE 8 (continued)
High-Frequency Quoting and Asymmetry in Quote Changes

Panel B. Estimated Effects

		Bid		Offer	
Dependent Variable	Sample	HFQ	HDD	HFQ	HDD
γ	Full	−0.104	−0.016	0.109	0.017
		(0.007)	(0.009)	(0.007)	(0.009)
	1 (low)	−0.191	−0.064	0.143	0.212
		(0.027)	(0.093)	(0.027)	(0.088)
	2	−0.198	−0.051	0.211	0.007
		(0.019)	(0.029)	(0.019)	(0.029)
	3	−0.150	−0.013	0.141	0.042
		(0.018)	(0.023)	(0.018)	(0.021)
	4	−0.038	−0.028	0.047	0.035
		(0.015)	(0.015)	(0.015)	(0.016)
	5 (high)	0.038	−0.031	0.018	0.021
		(0.012)	(0.010)	(0.012)	(0.010)
NMLM	Full	−0.034	−0.147	0.020	0.116
		(0.007)	(0.009)	(0.007)	(0.009)
	1 (low)	−0.064	−0.148	0.033	0.101
		(0.019)	(0.064)	(0.019)	(0.061)
	2	−0.023	−0.139	0.016	0.113
		(0.013)	(0.020)	(0.013)	(0.020)
	3	−0.069	−0.112	0.059	0.064
		(0.013)	(0.017)	(0.013)	(0.016)
	4	−0.007	−0.162	−0.004	0.128
		(0.017)	(0.017)	(0.017)	(0.018)
	5 (high)	0.005	−0.154	−0.020	0.156
		(0.021)	(0.018)	(0.021)	(0.019)

a stronger presence of Edgeworth cycles may be a consequence of the relatively large tick size. Alternatively, a setting in which the participants are converging toward a mixed-strategy equilibrium might still exhibit undercutting cycles during a learning phase, as agents estimate the number of competitors, their reaction speeds, and reservation prices. In an active security market, these conjectures are obviously very dynamic. Finally, the positive association between competition and high-frequency quoting might also arise from variation in reaction speed. In the Baruch and Glosten (2013) model, all participants cancel and replace in unison at a frequency that is determined by the limits of technology. If there is variation in latency of liquidity providers, or if there are limits imposed on individual traders' message intensity, the observed frequency of quote revision may be positively related to the number of traders.[11,12]

If quote volatility is in fact associated with Edgeworth cycles, what might this suggest about welfare? In retail gasoline markets, Noel (2011) notes, "While the gasoline cycles continue to generate public concern with claims of collusion often raised, the current evidence favors Edgeworth price cycles being the result of stronger competition and the source of lower retail gasoline prices." Two other

[11] In the Baruch and Glosten (2013) model, with prices drawn from a continuous distribution, ties in limit-order prices are a zero-probability event, which in principle renders time priority irrelevant. The product market settings in which Edgeworth cycles are usually studied do not have time priority, but one might conjecture that incorporating time priority would only increase the incentives to improve on, rather than simply match, an opponent's quoted price.

[12] I am grateful to Larry Glosten for discussions on these points.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

interesting parallels emerge from the studies of retail gasoline markets. Lewis (2009) and Lewis and Noel (2011) find evidence that retail prices adjust more quickly to wholesale prices in markets in which price cycles are present. This would correspond to greater informational efficiency in a financial market. Along these lines, CWX (2015) find that increased quoting frequency is positively associated with a number of proxies for efficiency and liquidity. Noel (2012) examines timing strategies that might be used by consumers to minimize their average expenditures. Although the approaches (and, of course, the time scales) are quite different, this is implicitly the problem faced by market-order traders in the model of Section II.

As a final note, the analysis here is directed at a broad classification of quote volatility. It does not rule out occurrences of quote stuffing or spoofing. It simply offers evidence consistent with other motives.

VIII. Conclusions

This paper presents new economic and statistical perspectives on high-frequency quoting (i.e., short-term volatility in bid and offer quotes). From the viewpoint of liquidity demanders, quote volatility implies execution price risk, as an incoming order essentially attempts to hit a moving target. As a representative estimate, a market-order buyer of a U.S. stock who can only time her order arrival to some point within a 1.6-second interval incurs an execution-price volatility of about 0.138 cents per share relative to the interval mean. Traders with lower timing uncertainty are in a better position to pick out advantageous times for execution. Under additional assumptions, there is a correspondence between quote volatility and this relative gain. As a representative value, a seller who can precisely time her order arrivals has an expected advantage of 0.390 cents per share relative to a seller subject to a 1.6-second window of uncertainty. These risks and advantages are comparable to other costs associated with trading, such as commissions and access fees.

The paper also examines quote-setters' strategies that might cause quote volatility. Although quote stuffing and spoofing are often mentioned, analogies to product markets suggest alternative explanations. In Baruch and Glosten (2013), quote-setters follow mixed strategies, repeatedly drawing their limit prices from a mixing distribution that does not change between updates to public information. This model suggests that quote volatility should be inversely related to the number of liquidity suppliers. This paper's analysis uses various HHIs of exchange quotation shares as inverse proxies for the latter. In panel estimations, these HHIs are found to be negatively related to quote volatility. This implies a positive relation between competition and quote volatility, which is not supportive of the mixed-strategy view.

Product markets are also known to exhibit Edgeworth cycles. Commonly found in retail gasoline markets, these price dynamics involve incremental undercutting punctuated by large price changes in the opposite direction. They are characterized by asymmetry in price changes (on the bid side, for example, many small increases followed by a large drop). This asymmetry is measured by skewness (negative for bids, positive for offers). This paper finds that quote volatility is

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

associated with intensified skewness, which is consistent with a higher incidence of Edgeworth cycles.

The association of volatile quotes with manipulative strategies has caused regulators to consider the imposition of rules mandating a minimum quote life (see U.K. Government Office for Science (2012)). This paper's findings suggest alternative explanations for quote volatility that might not warrant such extreme measures. If quote stability is viewed as a valuable social good, it can be encouraged by other means. To this end, the Toronto Stock Exchange has recently introduced a long-life qualifier for orders that commit to a minimum resting time of 1 second in return for priority benefits (TMX Equities Markets (2015)).

The findings on quote volatility presented here are broad characterizations that may, on closer examination, exhibit differences stemming from diversity in the strategies of liquidity makers and takers. The dynamics of quote volatility are also of interest. In the motivating example of AEPI, the volatility appears to be episodic and concentrated. The statistical and economic implications of these features are worthwhile goals of further research.

Appendix A. Time-Scale Decompositions of Random Walks

This appendix discusses the direct computation of the residual and incremental variances, σ_j^2 and ν_j^2, and related variance ratios for a random walk. Consider a discrete-time random walk given by $p_t = p_{t-1} + u_t$, where u_t is a disturbance with $Eu_t = 0$, $Eu_t^2 = \sigma_u^2$, and $Eu_t u_{t-k} = 0$ for $k = \pm 1, \pm 2, \ldots$ Without loss of generality, I initialize $p_0 = 0$ and consider the mean-squared deviations around an n-period average for $n \geq 2$:

$$\text{(A-1)} \qquad \text{MSD}(n) = \frac{1}{n}\sum_{i=1}^{n} p_i^2 - \left(\frac{1}{n}\sum_{i=1}^{n} p_i\right)^2 = \frac{1}{n}\sum_{i=1}^{n}\left(\sum_{j=1}^{i} u_i\right)^2 - \left(\frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{i} u_i\right)^2.$$

Taking expectations and simplifying the sums gives

$$\text{(A-2)} \qquad E(\text{MSD}(n)) = \left[\frac{n+1}{2} - \frac{(n+1)(2n+1)}{6n}\right]\sigma_u^2 = \left[\frac{n^2-1}{6n}\right]\sigma_u^2.$$

The sequence of residual (rough) variances is obtained as $\sigma_j^2 = E(\text{MSD}(n_j))$ for the sequence of averaging periods $n_j = 2^j$ for $j = 1, \ldots$. This simplifies to

$$\text{(A-3)} \qquad \sigma_j^2 = \frac{\sigma_u^2}{3}\left(\frac{4^j - 1}{2^{j+1}}\right) \quad \text{for} \quad j = 1, \ldots.$$

The incremental variance is defined as $\nu_1^2 = \sigma_1^2$ for $j = 1$, and for $j > 1$,

$$\text{(A-4)} \qquad \nu_j^2 = \sigma_j^2 - \sigma_{j-1}^2 = \frac{\sigma_u^2}{6}(2^{j-1} + 2^{-j}).$$

Within a wavelet framework, σ_j^2 and ν_j^2 denote the rough and detail or wavelet variances computed under the Haar transformation (Percival and Walden (2000), chap. 8). Equation (A-4) is equivalent to Percival and Walden, exercise 8.3, p. 337 (using, in their equation, $\tau_j = 2^{j-1}$).

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

We now rework this in units of natural time. The random walk steps are realized over intervals of length Δ, M denotes the length of the averaging interval, and σ_v^2 denotes the random-walk variance per unit time. The number of steps in the interval is $n = M/\Delta$, and the variance over each step is $\sigma_u^2 = \Delta\sigma_v^2$. Substituting into equation (A-2) gives the expected mean-square deviation as

$$\text{(A-5)} \qquad \text{EMSD}\left(M, \sigma_v^2, \Delta\right) = \frac{\sigma_v^2\left(M^2 - \Delta^2\right)}{6M}.$$

In the limit, as $\Delta \to 0$, $\text{EMSD} \to \sigma_v^2 M/6$.

Next consider a sequence of averaging intervals, $2^j M_0$ for $j = 0, \ldots$. The corresponding sequence of residual variances is $\sigma_j = 2^{j-1} M_0 \sigma_v^2/3$. As before, the incremental variances are given by $v_{j=0}^2 = \sigma_0^2$ and $v_j^2 = \sigma_j^2 - \sigma_{j-1}^2$ for $j \geq 1$. For any choice of j and fixed J with $j \leq J$, one may define a normalized (long incremental)/(short residual), $\text{VR}_j = 2^{1+j-J} v_j^2/\sigma_j^2$. For a random walk in continuous time, the normalization yields a value of unity. The numerator uses as a measure of long-term variation the incremental variance, v_J^2, for the longest time scale. By construction, this quantity smooths out short-term effects.

It is sometimes computationally convenient to base the analysis on prices averaged over a short interval. The bids, offers, and midpoints in this study, for example, are averaged over intervals of width $M_0 = 50$ ms. If variation *within* this interval is not included, the normalization of the residual variance ratio changes slightly, and $\text{VR}_j = (2^{j-1}/2^{J-1}) \times (v_J^2/\sigma_j)$. The VRs reported in Tables 2 and 3 are of this form.

Appendix B. Quote Volatility and the Gains to Speed

The discussion in Section II established that, given the path of a bid or offer, order arrival–time uncertainty gives rise to execution price risk. This appendix uses the same framework to model a fast trader's ability to pick off the best price (lowest offer or highest bid). The mathematical device that facilitates this is an orthogonal decomposition based on time scale.

As an illustration, I consider an offer price that is stochastic over a trading session of 8 discrete periods, $t = 1, 2, \ldots, 8$ ("seconds"). The offer is $\mu + R(t)$, where μ is the session mean and $R(t)$ is a zero-mean residual. This residual is constructed as the sum of three components: $R(t) = D_1(t) + D_2(t) + D_3(t)$, where each $D_j(t)$ has a binomial structure identified with a particular time scale. More precisely, the (relatively) long-term component is given by

$$\text{(B-1)} \qquad D_3(t) = a_3\left[1 \quad 1 \quad 1 \quad 1 \quad -1 \quad -1 \quad -1 \quad -1\right],$$

where a_3 is a random variable that takes on values $\pm v_3$ with equal probability. The medium-term component is

$$\text{(B-2)} \qquad D_2(t) = \begin{bmatrix} a_{2,1} \\ a_{2,2} \end{bmatrix}' \begin{bmatrix} 1 & 1 & -1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & -1 & -1 \end{bmatrix},$$

where $a_{2,1}$ and $a_{2,2}$ are IID with equiprobable values $\pm v_2$. The short-term component is

$$\text{(B-3)} \qquad D_1(t) = \begin{bmatrix} a_{1,1} \\ a_{1,2} \\ a_{1,3} \\ a_{1,4} \end{bmatrix}' \begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix}.$$

The $a_{1,\bullet}$ are IID with equiprobable values $\pm v_1$.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

Each $D_j(t)$ component is constant over a time scale of 2^{j-1} periods and then flips sign. The 7 row vectors on the right sides of equations (B-1)–(B-3) are Haar functions. They are zero-mean and constitute an orthogonal basis for $R(t)$. (One degree of freedom is lost because the residual must be zero-mean.) The 7 a_{jk} binomial random variables generate $2^7 = 128$ possible paths for the residual.

There are four types of market-order buyers, indexed by $i \in \{1,2,3,4\}$. The index i corresponds to a trader's random latency. The type-1 trader is the fastest and is subject to no latency uncertainty. For example, if a type-1 agent submits in period 3, the order executes in period 3. More generally, a type-i trader experiences a delay that is uniformly distributed over $\{0,1,\ldots,2^{j-1}\}$. For example, if a type-2 trader submits in period 3, the order will execute with equal probability in period 3 or 4. The slowest trader (type 4) is subject to random latency that is uniform over $\{0,1,\ldots,7\}$. This trader has no discretion and must submit in period 1.

Although somewhat stronger than necessary, it simplifies things to assume that all traders know the full path of the offer (tantamount to knowing the realizations of a_{ij}). Their abilities to exploit this knowledge, however, depend on their latency uncertainties. The type-4 (slowest) buyer, submitting in period 1, expects to pay μ, the session mean. A type-3 trader can optimize over the path of $D_3(t)$. If $a_{3,1} = -\nu_3$, she knows that the average offer price over periods $\{1,2,3,4\}$ is $\mu - \nu_3$; if $a_{3,1} = +\nu_3$, the average offer is $\mu - \nu_3$ over periods $\{5,6,7,8\}$. She will, accordingly, submit in period 1 or period 5. Similarly, a type-2 (faster) trader, by optimizing over the path of $D_3(t) + D_2(t)$, can realize a purchase price $\mu - \nu_2 - \nu_3$. The fastest trader can buy at $\mu - \nu_1 - \nu_2 - \nu_3$. In this framework, then, the sum $G_j = \sum_{i=1}^{j} \nu_j$ represents the gain (advantage) of the fastest trader relative to traders of type $j+1$ (i.e., those subject to latency uncertainty of duration 2^j).

From the construction of the offer process, $\text{var}\left(D_j\right) = \nu_j^2$, and thus the ν_j^2 correspond to the incremental variances discussed in Section II. In that framework, the interpretation of σ_j^2 and ν_j^2 as time scale–related risk is valid for a fairly general set of price processes. The mapping from the ν_j to the speed gain G_j, however, is contingent on the equiprobable two-value realizations of the a_{ij}: This simplification easily locates the minimum or maximum of the price path. If the a_{ij} were, for example, Gaussian, this determination would be considerably more difficult.

As developed here, G_j is largely hypothetical, an amount that a fast trader would gain versus a slower one, ceteris paribus. In equilibrium, however, these relative advantages can easily become real trading costs. If all buyers are of the slowest type (type 4), sellers will presumably set their prices so that the average offer breaks even. Now suppose that part of the buyer population is suddenly equipped with technology that makes them type 3. Sellers will lose, on average, to these buyers (because they can selectively pick off the better prices). To recoup their losses, they will raise their average offer, shifting onto the slower traders their losses to faster traders. This transfer is familiar from standard models of asymmetric information. There is no necessary implication that the faster traders' gains are rents, particularly if entry is available to all who purchase the necessary technology.

References

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

Angel, J. J.; L. E. Harris; and C. S. Spatt. "Equity Trading in the 21st Century." *Quarterly Journal of Finance*, 1 (2011), 1–53.

Angel, J. J.; L. E. Harris; and C. S. Spatt. "Equity Trading in the 21st Century: An Update." *Quarterly Journal of Finance*, 5 (2015), 1–39.

Barnea, A. "Performance Evaluation of New York Stock Exchange Specialists." *Journal of Financial and Quantitative Analysis*, 9 (1974), 511–535.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

Baron, M.; J. Brogaard; B. Hagströmer; and A. A. Kirilenko. "Risk and Return in High-Frequency Trading." *Journal of Financial and Quantitative Analysis*, forthcoming (2018).
Baruch, S., and L. R. Glosten. "Flickering Quotes." Working Paper, Columbia University (2013).
Biais, B.; T. Foucault; and S. Moinas. "Equilibrium Fast Trading." *Journal of Financial Economics*, 116 (2015), 292–313.
Brogaard, J.; T. Hendershott; and R. Riordan. "High-Frequency Trading and Price Discovery." *Review of Financial Studies*, 27 (2014), 2267–2306.
Budish, E. B.; P. Cramton; and J. J. Shim. "The High-Frequency Arms Race: Frequent Batch Auctions as a Market Design Response." *Quarterly Journal of Economics*, 130 (2015), 1547–1621.
Castanias, R., and H. Johnson. "Gas Wars: Retail Gasoline Price Fluctuations." *Review of Economics and Statistics*, 75 (1993), 171–174.
Conrad, J.; S. Wahal; and J. Xiang. "High-Frequency Quoting, Trading, and the Efficiency of Prices." *Journal of Financial Economics*, 116 (2015), 271–291.
Dacorogna, M. M.; R. Gençay; U. A. Muller; R. B. Olsen; and O. B. Pictet. *High-Frequency Finance*. New York, NY: Academic Press (2001).
Doyle, J.; E. Muehlegger; and K. Samphantharak. "Edgeworth Cycles Revisited." *Energy Economics*, 32 (2010), 651–660.
Eckert, A. "Retail Price Cycles and Response Asymmetry." *Canadian Journal of Economics/Revue Canadienne d'Économique*, 35 (2002), 52–77.
Edgeworth, F. Y. "The Pure Theory of Monopoly." In *Papers Relating to Political Economy*, Vol. I, F. Y. Edgeworth, ed. London, UK: MacMillan (1925), 111–142.
Egginton, J. F.; B. F. Van Ness; and R. A. Van Ness. "Quote Stuffing." *Financial Management*, 45 (2016), 583–608.
Elder, J., and H. J. Jin. "Long Memory in Commodity Futures Volatility: A Wavelet Perspective." *Journal of Futures Markets*, 27 (2007), 411–437.
Fan, J., and Y. Wang. "Multi-Scale Jump and Volatility Analysis for High-Frequency Financial Data." *Journal of the American Statistical Association*, 102 (2007), 1349–1362.
Fan, Y., and R. Gençay. "Unit Root Tests with Wavelets." *Econometric Theory*, 26 (2010), 1305–1331.
Foucault, T.; J. Hombert; and I. Rosu. "News Trading and Speed." *Journal of Finance*, 71 (2016), 335–382.
Gençay, R.; F. Selçuk; N. Gradojevic; and B. Whitcher. "Asymmetry of Information Flow between Volatilities across Time Scales." *Quantitative Finance*, 10 (2010), 895–915.
Gençay, R.; F. Selçuk; and B. Whitcher. *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. San Diego, CA: Academic Press (2002).
Gençay, R.; F. Selçuk; and B. Whitcher. "Multiscale Systematic Risk." *Journal of International Money and Finance*, 24 (2005), 55–70.
Gençay, R., and D. Signori. "Multi-Scale Tests for Serial Correlation." *Journal of Econometrics*, 184 (2015), 62–80.
Hasbrouck, J., and G. Saar. "Technology and Liquidity Provision: The Blurring of Traditional Definitions." *Journal of Financial Markets*, 12 (2009), 143–172.
Hasbrouck, J., and R. A. Schwartz. "Liquidity and Execution Costs in Equity Markets." *Journal of Portfolio Management*, 14 (1988), 10–16.
Hendershott, T. J.; C. M. Jones; and A. J. Menkveld. "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance*, 66 (2011), 1–33.
Høg, E., and A. Lunde. "Wavelet Estimation of Integrated Volatility." Working Paper, Aarhus School of Business (2003).
In, F., and S. Kim. "The Hedge Ratio and the Empirical Relationship between the Stock and Futures Markets: A New Approach Using Wavelet Analysis." *Journal of Business*, 79 (2006), 799–820.
ITG Inc. *Global Cost Review*, Q1/2014 (2014).
Jones, C. M. "What Do We Know About High-Frequency Trading?" Working Paper, Columbia Business School (2013).
Leach, J. C., and A. N. Madhavan. "Intertemporal Price Discovery by Market Makers: Active versus Passive Learning." *Journal of Financial Intermediation*, 2 (1992), 207–235.
Leach, J. C., and A. N. Madhavan. "Price Experimentation and Security Market Structure." *Review of Financial Studies*, 6 (1993), 375–404.
Lewis, M. S. "Temporary Wholesale Gasoline Price Spikes Have Long-Lasting Retail Effects: The Aftermath of Hurricane Rita." *Journal of Law and Economics*, 52 (2009), 581–605.
Lewis, M. S., and M. Noel. "The Speed of Gasoline Price Response in Markets With and Without Edgeworth Cycles." *Review of Economics and Statistics*, 93 (2011), 672–682.
Maskin, E., and J. Tirole. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica*, 56 (1988), 971–599.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press

Menkveld, A. J. "High-Frequency Trading and the New Market-Makers." *Journal of Financial Markets*, 16 (2013), 712–740.
Nason, G. P. *Wavelet Methods in Statistics with R*. New York, NY: Springer Science + Business Media, LLC (2008).
Noel, M. D. "Edgeworth Price Cycles." *The New Palgrave Dictionary of Economics*. S. N. Durlauf and L. E. Blume, eds. Basingstoke, UK: Palgrave Macmillan (2011).
Noel, M. D. "Edgeworth Price Cycles and Intertemporal Price Discrimination." *Energy Economics*, 34 (2012), 942–954.
O'Hara, M. "High Frequency Market Microstructure." *Journal of Financial Economics*, 116 (2015), 257–270.
Percival, D. B., and A. T. Walden. *Wavelet Methods for Time Series Analysis*. Cambridge, UK: Cambridge University Press (2000).
Ramsey, J. B. "The Contribution of Wavelets to the Analysis of Economic and Financial Data." *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 357 (1999), 2593–2606.
Ramsey, J. B. "Wavelets in Economics and Finance: Past and Future." *Studies in Nonlinear Dynamics and Econometrics*, 6 (2002), Article 1.
Schwartz, R. A., and L. R. Wu. "Equity Trading in the Fast Lane: The Staccato Alternative." *Journal of Portfolio Management*, 39 (2013), 3–6.
Stoll, H. R., and C. Schenzler. "Trades Outside the Quotes: Reporting Delay, Trading Option, or Trade Size?" *Journal of Financial Economics*, 79 (2006), 615–653.
Teyssière, G., and P. Abry. "Wavelet Analysis of Nonlinear Long-Range Dependent Processes: Applications to Financial Time Series." In *Long Memory in Economics*, G. Teyssière and A. P. Kirman, eds. Berlin, Germany: Springer (2007).
TMX Equities Markets. *Order Types and Functionality Guide* (2015).
U.K. Government Office for Science. *Economic Impact Assessments on Mifid II Policy Measures Related to Computer Trading in Financial Markets* ("Foresight Report"). London, UK: Government Office for Science (2012).
Walker, J. S. *A Primer on Wavelets and Their Scientific Applications*. Boca Raton, FL: Chapman and Hall/CRC (Taylor and Francis Group) (2008).
Wills-Johnson, N., and H. Bloch. "The Shape and Frequency of Edgeworth Price Cycles in an Australian Retail Gasoline Market." Working Paper, Curtin University of Technology (2010).
Zhang, X. M. "Finding Edgeworth Cycles in Online Advertising Auctions." Working Paper, Massachusetts Institute of Technology (2005).
Zhang, X. M., and J. Feng. "Cyclical Bid Adjustments in Search-Engine Advertising." *Management Science*, 57 (2011), 1703–1719.

https://doi.org/10.1017/S0022109017001053 Published online by Cambridge University Press