

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

Joel Hasbrouck*

November 13, 2012

This version: August 20, 2013

I have benefited from the comments of Ramo Gençay, Dale Rosenthal, Gideon Saar, Mao Ye, seminar/conference participants at Aalto University, BI Norwegian Business School, the Copenhagen Business School, the Emerging Markets Group (Cass Business School, City University London), the Euronext Paris Conference on High-Frequency Trading (April, 2013), l'Institute Louis Bachelier, Jump Trading, Norges Bank Investment Management, SAC Capital, the University of Illinois at Chicago, the University of Illinois at Champaign, the University of Pennsylvania Econometrics Seminar and Utpal Bhattacharya's doctoral students at the University of Indiana. All errors are my own responsibility. **DISCLAIMER:** 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.

*Department of Finance, Stern School of Business, New York University, 44 West 4th Street, New York, NY 10012 (Tel: 212-998-0310, jhasbrou@stern.nyu.edu).

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

Abstract

High-frequency changes, reversals, and oscillations induce volatility in a market's bid and offer quotes. This volatility degrades the informational content of the quotes, exacerbates execution price risk for marketable orders, and impairs the reliability of the quotes as reference marks for the pricing of dark trades. This paper examines variance on time scales as short as fifty milliseconds for the National Best Bid and Offer (NBBO) in the US equity market. On average, in a 2011 sample, NBBO variance at the fifty millisecond time scale is approximately four times larger than can be attributed to long-term fundamental price variance. The historical picture over 2001-2011 is complex. The average volatility has not increased between 2001 and 2011, but its nature has changed. In the earlier years quote volatility is due to large spikes in bids and offers; in later years, to oscillations of low amplitude. The highest quote volatilities arise during the 2004-2006 period corresponding to the phase-in of Reg NMS and the transition to electronic trading.

I. Introduction

Recent developments in market technology have called attention to the practice of high frequency trading. The term is used commonly and broadly in 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, 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 offers of US equities, a consequence of what might be called high frequency quoting.

By way of illustration, Figure 1 depicts the bid and offer for AEP Industries (a NASDAQ-listed manufacturer of packaging products) on April 29, 2011.¹ In terms of broad price moves, the day is not a particularly volatile one, and the bid and offer quotes are stable for long intervals. The placidity is broken, though, by several intervals where the bid undergoes extremely rapid changes. The average price levels, before, during and after the episodes are not dramatically different. Moreover, the episodes are largely one-sided: the bid volatility is associated with an only moderately elevated volatility in the offer quote. 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.

It is not, however, an innocuous artifact. Bids and offers in all markets represent price signals, and, to the extent that they are firm and accessible, immediate trading opportunities. From this perspective, the noise added by quote volatility impairs the informational value of the public price. Most agents furthermore experience latency in ascertaining the market center with the best price and in timing of their order delivery. Elevated short-term volatility increases the execution price risk associated with these delays. In US equity markets the bid and offer are particularly important, because they are used as benchmarks to assign prices in so-called dark trades, a category that includes roughly thirty percent of all volume.²

¹ The bid is the National Best Bid (NBB), the maximum bid across all exchanges. The offer is the National Best Offer (NBO), the minimum offer. They are often jointly referred to as the NBBO. Unless otherwise noted, or where clarity requires a distinction, “bid” and “offer” indicate the NBBO.

² Dark mechanisms do not publish visible bids and offers. They establish buyer-seller matches, either customer-to-customer (as in a crossing network) or dealer-to-customer (as in the case of an internalizing broker-dealer). The matches are priced by reference to the NBBO: generally at the NBBO midpoint in a crossing network, or at the NBB or the NBO in a dealer-to-customer trade.

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 volatility? How can we distinguish fundamental (informational) and transient (microstructure) volatility? Finally, given the current public policy debate surrounding low-latency activity, how has it changed over time?

These questions are addressed empirically in a broad sample of US equity market data using summary statistics that are essentially short-term variances of bids and offers. Such constructions, though, inevitably raise the question of what horizon constitutes the "short term" (a millisecond? a minute?). The answer obviously depends a trader's latency, which presently ranges from the sub-millisecond (for a collocated algorithm) to seconds or minutes (for a remotely situated human trader). The indeterminacy motivates empirical approaches that accommodate flexible time horizons. This analysis uses time scale variance decompositions to measure bid and offer volatility over horizons ranging from under 50 ms to about 27 minutes.

The next section establishes the economic and institutional motivation for the consideration of local bid and offer variances with sliding time scales. Section III discusses the methodological framework. The paper then turns to applications. Section IV presents an analysis of a recent sample of US equity data featuring millisecond time stamps. To extend the analysis to historical samples in which time stamps are to the second, Section V describes estimation in a Bayesian framework where millisecond time stamps are simulated. Section VI applies this approach to a historical sample of US data from 2001 to 2011. Connections to high frequency trading and volatility modeling are discussed in Section VII. A summary concludes the paper in Section VIII.

II. The economic effects of quote volatility.

High frequency quote volatility may be provisionally defined as the short-term variance of the bid or offer, the usual variance calculation applied to the bid or offer *level* over a relatively brief window of time. This section is devoted to establishing the economic relevance of such a variance in a trading context. The case is a simple one, based on the function and uses of the bid and offer, the barriers to their instantaneous availability, the role of the time-weighted price mean as a benchmark, and the interpretation of the variance about this mean as a measure of risk.

In current thinking about markets, most timing imperfections are either first-mover advantages arising from market structure or delays attributed to costly monitoring. The former are

exemplified by the dealer's option on incoming orders described in Parlour and Seppi (2003), and currently figure in some characterizations of high frequency traders (Biais, Foucault, and Moinas, 2012; Jarrow and Protter, 2012). The latter are noted by Bessembinder, Panayides, and Venkataraman (2009) and discussed by O'Hara and Ye (2009) as an important special case of inattention which, albeit rational and optimal, leads to infrequent trading, limited participation, and transient price effects.

As a group these models feature a wide range of effects bearing on agents' arrivals and their information asymmetries. An agent's market presence may be driven by monitoring decisions, periodic participation, or random arrival intensity. Asymmetries mostly relate to fundamental (cash-flow) information or lagged information from other markets. Agents in these models generally possess, however, timely and extensive market information. Once she "arrives" in a given market, an agent accurately observes the state of that market, generally including the best bid and offer, depth of the book and so on. Moreover, when she contemplates an action that changes the state of the book (such as submitting, revising or canceling an order), she knows that her action will occur before any others'.

A notable and recent exception arises in Baruch and Glosten (2013). In their model, limit order traders pursue randomized order placement strategies, in an ongoing fashion and even in the absence of any trades. Quote volatility is an obvious consequence of such behavior.

Even without randomized strategies, however, random latencies in receiving information and transmitting intentions combine to frustrate certainties about the state of the market and terms of trade. The perspective of this paper is that for some agents these random latencies generate randomness in the execution prices, and that short-term quote variances can meaningfully measure this risk. Furthermore, although all agents incur random latency, the distributions of these delays vary among participants. An agent's latency distribution can be summarized by time scale, and this in turn motivates time scale decompositions of bid and offer variances.

While random latencies might well affect strategies of all traders, the situation is clearest for someone who intends to submit a marketable order (one that seeks immediate execution) or an order to a dark pool. In either case, ignoring hidden orders, an execution will occur at the bid, the offer or at an average of the two. A trader whose order arrival time is random over an interval

faces price risk. For a marketable sell order, if the arrival time is uniformly distributed, the interval variance of the bid quantifies the price risk around the interval mean.³

The situations discussed to this point involve a single trader and single market. In a fragmented market, the number of relevant latencies may be substantially larger. In the US there are presently about 17 “lit” market centers, which publish quotes. A given lit market’s quotes are referenced by the other lit markets, dark pools (currently around 30 in number), by executing broker-dealers (approximately 200), and by data consolidators (U.S. Securities and Exchange Commission, 2010). The National Best Bid and Offer (NBBO) is in principle well-defined. The NBBO perceived by any given market center, consolidator or other agent, however, comprises information subject to random transmission delays that differ across markets and receiving agents. These delays introduce noise into the determination. Local time-averaging (smoothing) can help to mitigate the effects of this noise, and the local variance measures the magnitude of the noise.

If the execution price risk associated with quote volatility is zero-mean and diversifiable across trades, it might appear to be economically trivial. In general, however, agents do not have symmetric exposure to this risk. Market-order traders with faster technology possess a systematic advantage relative to those with slower technology. This can be viewed as an information asymmetry that leads (in the usual fashion) to a transfer of wealth from the slower to the faster participants.

³ The use of an average price in the presence of execution timing uncertainty is a common principle in transaction cost analysis. Perold’s implementation shortfall measure is usually operationally defined for a buy order as the execution price (or prices) less some hypothetical benchmark price, and for a sell order as the benchmark less the execution price, (Perold, 1988). As a benchmark price, Perold suggests the bid-offer midpoint prevailing at the time of the decision to trade. Many theoretical analyses of optimal trading strategies use this or a similar pre-trade benchmark. Practitioners, however, and many empirical analyses rely on prices averaged over some comparison period. The most common choice is the value-weighted average price (VWAP), although the time-weighted average price (TWAP) is also used. One industry compiler of comparative transaction cost data notes, “In many cases the trade data which is available for analysis does not contain time stamps. . . . When time stamps are not available, pension funds and investment managers compare their execution to the volume weighted average price of the stock on the day of the trade” (Elkins-McSherry, 2012). This passage attests to the importance of execution time uncertainty, although a day is certainly too long to capture volatility on the scale of transmission and processing delays. Average prices are also used as objectives by certain execution strategies. A substantial portion of the orders analyzed by Engle, Furstenberg and Russell (2012) target VWAP, for example.

Asymmetric exposure to quote volatility is also likely to place customers at a disadvantage relative to their dealers. The recent SEC concept release notes that virtually all retail orders are routed to OTC market-makers, who execute the orders by matching the prevailing NBBO (U.S. Securities and Exchange Commission, 2010). Stoll and Schenzler (2006) note that these market-makers have flexibility in delaying executions to obtain favorable reference prices. They describe this as a look-back option, and find support for this behavior in a 1999 sample. Dark trading venues also face this sort of problem. A customer sending a sell order to a dark pool or crossing network can submit a *buy* order to a lit market center that will briefly boost quote midpoint, thereby achieving a better price if he receives a dark execution of his sell order. This practice (a form of “spoofing”) is forbidden in the Dodd-Frank framework, but remains difficult to detect and prove in the presence of timing uncertainties. Quote volatility might also arise from “quote stuffing” (Egginton, Van Ness, and Van Ness, 2012).

The SEC’s Reg NMS ruling on trade-through protection recognized the problem of “flickering quotes”, and mandated a one-second grace period: “... pursuant to Rule 611(b)(8) trading centers would be entitled to trade at any price equal to or better than the least aggressive best bid or best offer, as applicable, displayed by the other trading center during that one-second window.” Sub-second intervals were considered, but the benefits were not believed sufficient to justify the costs (U.S. Securities and Exchange Commission, 2005). Clearly, quote volatility within the one-second window weakens the trade-through protection.⁴

⁴ The SEC has recently mandated a consolidated audit trail intended to track all events in an order’s life cycle, such as receipt, routing, execution and cancellation, (U.S. Securities and Exchange Commission, 2012). In the final rule, the Commission recognized the importance of accurately sequencing the events, and mandated time-stamps at least to the granularity of the millisecond, and, “to the extent that the order handling and execution systems of any SRO or broker-dealer utilize time stamps in increments finer than the minimum required by the NMS plan time stamps, such SRO or member must use time stamps in such finer increments when reporting data to the central repository.”

III. Methodology

III.A. Time scale decompositions

A number of texts cover time scale decompositions. The present material summarizes the intuition of the approach. (Percival and Walden, 2000) discuss the details.^{5,6} To illustrate the computations, consider a discrete sequence of prices $p = \{p_1, \dots, p_T\}$. A time scale decomposition resolves p into a set of local averages of varying lengths and corresponding residuals. The averages in this context are called *smoothed* components (or simply “smooths”) and are denoted by S ; the residuals are *roughs*, denoted by R . The smooths and roughs are indexed by j , which maps to the length of the average. It is computationally convenient to let these lengths vary dyadically (increasing as powers of two). The shortest meaningful average we might form is of length two, so the smooths at level j , $S_j, j = 1, \dots, J$, correspond to averages over subsamples of lengths 2^j for $j = 1, \dots, J$ where $2^J \leq T$. The corresponding roughs are defined by $R_j = p - S_j$, and they are mean-zero by construction.

As level j increases, it is useful to consider the incremental changes in the roughs. For example, writing $p = S_2 + R_2 = S_2 + (R_2 - R_1) + R_1$, the middle term reflects the change in the overall residual associated with moving from an averaging length of four to two. These

⁵ Time scale and multi-resolution decompositions are widely used across many fields. In addition to Percival and Walden, (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 and financial perspectives. Walker (2008) is clear and concise, but oriented more toward engineering applications.

⁶ Studies that apply time scale decompositions in the economic analysis of stock prices loosely fall into two groups. The first set explores time scale aspects of stock comovements. A stock’s beta is a summary statistic that reflects short-term linkages (like index membership or trading-clientele effects) and long-term linkages (like earnings or national prosperity). Wavelet analyses can characterize the strength and direction of these horizon-related effects (Gençay, Selçuk, and Whitcher, 2005; In and Kim, 2006). Most of these studies use wavelet transforms of stock prices at daily or longer horizons. A second group of studies uses wavelet methods to characterize volatility persistence (Dacorogna, Gençay, Muller, Olsen, and Pictet, 2001; Gençay *et al.*, 2002; Høgg and Lunde, 2003; Elder and Jin, 2007; Teyssière and Abry, 2007; Gençay, Selçuk, Gradojevic, and Whitcher, 2010). 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). Beyond studies where the focus is primarily economic or econometric lie many more analyses where wavelet transforms are employed for ad hoc stock price forecasting (Atsalakis and Valavanis, 2009; Hsieh, Hsiao, and Yeh, 2011, for example). An early draft of Hasbrouck and Saar (2013) used wavelet analyses of message count data to locate periods of intense message traffic on NASDAQ’s Inet system.

incremental components are the *details*, $D_1 \equiv R_1$ and $D_j = R_j - R_{j-1}$ for $j > 1$. The time scale decomposition of the price series may then be written as $p = S_j + D_{j-1} + \dots + D_1$ (also known as a multi-resolution analysis). Since the elements of the smooth S_j average over 2^j points, the elements corresponding rough R_j capture variation at shorter intervals, that is, over intervals of length $2^{j-1}, 2^{j-2}, \dots, 1$. In keeping with the incremental approach, a detail component D_j is said to capture variation at a single time scale defined by $\tau_j \equiv 2^{j-1}$.

Given a suitable stochastic framework, there exists a time scale decomposition of variance corresponding to that of the series itself. To achieve this, we assume that the first differences of p constitute a stationary stochastic process. Note that while this condition is stated in terms of the first differences, the calculations are still performed (as indicated above) on price levels. This condition suffices to define the rough and detail variances $\sigma_j^2 \equiv \text{Var}(R_j)$ and $\nu_j^2 = \text{Var}(D_j)$.

Time scale decomposition was historically based on Fourier analysis, the decomposition of a series as a sum of sine and cosine basis functions. Trigonometric functions cycle over the full support of the signal. Modern signal processing approaches use alternative bases that are localized and therefore better suited to picking up phenomena like the AEPI movements that are concentrated in time. The present analysis uses a wavelet basis (the Haar). In this framework, the detail variances ν_j^2 are conventionally called *wavelet variances*. It is important to emphasize, though, that they can be defined (as above) and even computed (albeit suboptimally) without wavelet transforms. Denoting it as the wavelet variance simply places it in an extensive and established literature.

From an economic perspective, rough and wavelet (detail) variances correspond to the timing uncertainty risks discussed in the preceding section. If p is a millisecond-stamped bid quote, for example, a trader submitting a marketable sell order with arrival time uncertainty of $2^6 = 64$ ms faces a price uncertainty of $\sigma_{j=6}^2$. Were she to cut her arrival uncertainty in half, this risk would be reduced by $\nu_{j=6}^2$. For ease of economic interpretation, therefore, estimates of the σ_j^2 and ν_j^2 are this paper's key summary statistics. For purposes of gauging magnitude or historical trend, however, it is useful to have measures of high-frequency volatility that are normalized relative to long-term volatility. This is achieved by variance ratios.

III.B. Variance ratios

There is a long tradition of variance ratios in empirical market microstructure (Barnea, 1974; Amihud and Mendelson, 1987; Hasbrouck and Schwartz, 1988, among others).⁷ A typical ratio compares the variance of a one-period price difference to an n -period price difference:

$$V_n = \frac{n \times \text{Var}(p_t - p_{t-1})}{\text{Var}(p_t - p_{t-n})} \quad (1)$$

If p_t follows a random-walk, $V_n = 1$ for all n . In actual sample data, variance ratios are generally elevated due to short-term microstructure effects. The motivation for using a long-term variance in the denominator of V_n is the desire for an estimate of “fundamental” volatility, on the assumption that a long-term price change is dominated by informational components.

The end-points of the n -period price change are subject, however, to microstructure effects just as strong (in absolute terms) as those of the short-term price change. That is, both p_t and p_{t-n} are subject to bid-ask bounce, discreteness effects, and so on. A long-term wavelet variance, however, is in principle purged of the short-term variation, and so may serve as a better estimate of fundamental long-term variance. Fan and Gençay (2010) apply this principle to unit root tests based on time scale decompositions. Gençay and Signori (2012) explore the use of variance ratios at different time scales to test for serial correlation. The variance ratios used here are special cases or minor modifications of theirs.

Like the conventional variance ratio, a wavelet variance ratio is benchmarked to a random-walk. A standard result (summarized in appendix) establishes that if the price evolves in continuous time with variance per unit time σ_u^2 (not necessarily a Gaussian diffusion) and that the prices are initially averaged over successive intervals of length M_0 units of time, then $\sigma_j^2 = 2^{j-1} M_0 \sigma_u^2 / 3$ and $v_j^2 = 2^{j-2} M_0 \sigma_u^2 / 3$. With this result, the wavelet variance ratio used here is

$$V_{j,J} = 2^{J-j} v_j^2 / v_J^2, \quad (2)$$

where J denotes the largest index (longest time scale) considered in the study and j (with $0 \leq j \leq J$) corresponds to a shorter time scale. Like the conventional variance ratio, a random-walk implies $V_{j,J} = 1$. The rough variance ratio is similarly defined as

⁷ Return variance ratios are also used more broadly in economics and finance to characterize deviations from random-walk behavior over longer horizons (Lo and MacKinlay, 1988).

$$VR_{j,J} = 2^{J-j-1} \sigma_j^2 / \nu_j^2 \quad (3)$$

Note that while the variance in the numerator is a rough variance, the denominator is a wavelet variance. For a random walk, $VR_{j,J} = 1$, and any excess indicates elevated short-term volatility. (Whereas $V_{j,J} = 1$ by construction, however, $VR_{j,J} = \sigma_j^2 / \nu_j^2$, which need not equal one.)

III.C. Estimation

Percival and Walden (for detail and rough variances) and Gençay and Signori (for variance ratios) discuss asymptotic results. For example, one could compute confidence bounds for a data sample consisting of one stock and one day. The present goal, however, is summarization of the broad aspects of high-frequency volatility across firms and time. The inferences therefore involve averages across firms and time computed from correlated observations. As an illustration, consider a wavelet standard deviation estimate $\nu_{i,d,j}$ computed for firm i on day d at resolution level j . In constructing an overall sample average across firms and days, observations on the same firm are likely to be correlated across days. For this reason, reported standard errors are clustered by firm. This corresponds to an error components model of the form: $\nu_{i,d,j} = \nu_j + u_i + e_{i,d}$, where u_i is a firm-specific error and the $e_{i,d}$ are assumed uncorrelated across firms and days. Since stock volatilities are commonly assumed to possess cross-firm commonalities, there is an argument in favor of clustering by day as well. Spot-calculations suggested, though, that standard errors clustered by firm and day were only slightly (less than five percent) higher than those clustered by firm alone.

IV. A cross-sectional analysis

From a trading perspective, stocks differ most significantly in their general level of activity (volume measured by number of trades, shares or values). The first analysis aims to measure the general level of high frequency quote volatility and to relate the measures to trading activity in the cross-section for a recent sample of firms.

IV.A. Data and sample construction.

The analyses are performed for a subsample of US firms using quote data from April, 2011 (the first month of my institution's subscription.) The subsample is constructed from all firms present on the CRSP and TAQ databases from January through April of 2011 with share codes of 10 or 11, with closing prices between two and one thousand dollars, and with a primary

listing on the New York, Amex or NASDAQ exchanges.⁸ 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.

The U.S. equity market is highly fragmented, but all exchanges post their quotes to the Consolidated Quote System (CQS).⁹ The CQ and NBBO files from the NYSE's daily TAQ dataset used here are definitive transcripts of the consolidated activity, time-stamped to the millisecond.¹⁰ A record in the consolidated quote (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 (with a somewhat more involved computation) directly from the CQ file. Spot checks confirm that these two approaches are consistent.

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 for the estimates discussed in this section. Error filters are used, however, in the subsequent historical analysis, and will be discussed in greater detail at that point.

Table 1 reports summary statistics. Post-Reg NMS US exchanges have become more similar in structures and trading mechanisms. With respect to listing characteristics, though, differences persist. The NYSE "classic" has the largest proportion of high-volume stocks, NYSE

⁸ 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".

⁹ At the same time that an exchange sends a quote update to the consolidated system, it can also transmit the update on its own subscriber line. For subscribers this can reduce the delay associated with consolidation and retransmission (which is on the order of about five 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.

¹⁰ The "daily" reference in the Daily TAQ dataset refers to the release frequency. Each morning the NYSE posts files that cover the previous day's trading. The Monthly TAQ dataset, more commonly used by academics is released with a monthly frequency and contains time stamps in seconds.

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 don't induce such a change, there are substantially fewer NBBO records. Finally, many actions might change one of sizes or one side of the quote. Thus, the numbers of NBB and NBO changes are smaller yet.

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. The estimates are computed using the maximal overlap Haar transform.¹¹ Variance estimates are computed separately for the bid and offer, and then averaged for convenience in presentation. Reported standard errors are clustered on firms.

To facilitate economic interpretation, I report the time scale variances in three ways: mils (\$0.001) per share, basis points relative to average price, and as a short/long-term variance ratio. The mils per share scaling is useful because many trading fees (such as commissions and clearing fees) are assessed on a per share basis. Access fees, the charges levied by exchanges on taker (aggressor) sides of executions are also assessed per share. US SEC Regulation NMS caps access

¹¹ A maximal overlap transform mitigates alignment problems. In the example discussed in Section III and depicted in Figure 2, the components are always aligned on dyadic boundaries. A maximal overlap transform essentially averages over all possible alignments. The computation of a one-second variance, for example, would involve not only periods aligned exactly on the one-second boundaries (1, 2, 3, ...), but also one-second periods aligned on half-second boundaries (1.5, 2.5, 3.5, ...). I assume no overlap across days, and discard boundary values affected by wrap-around.

The computations were performed in *Matlab* using the WMTSA package (Cornish, 2006). These routines conform closely to Percival and Walden. Although *Matlab* has its own wavelet toolbox, the data structures and other conventions differ significantly from those of Percival and Walden. I also find the *Mathematica* wavelet functions to be consistent with Percival and Walden.

Bid and offer series are piece-wise constant, a property shared by Haar basis functions. In applications such as audio and video processing, the signals are better described as locally linear, and other wavelet families (such as the Daubechies) are generally preferred to the Haar.

fees at 3 mils (\$0.003) per share, and in practice most exchanges are close to this level. Practitioners regard access fees as significant to the determination of order routing decisions, and this magnitude therefore serves an approximate threshold of economic importance. Basis point scaling is meaningful because most analyses involving investment returns or comparison across firms assume that share normalizations are arbitrary. Variance ratios provide a summary measure of short-term variance inflation relative to what would be expected from a random-walk calibrated to long-term variance.

IV.B. Results

Table 2 summarizes the averages for all time scales of wavelet and rough variances under all three normalizations. For example, a trader facing arrival time uncertainty of 50 milliseconds is exposed to a price risk standard deviation of 0.40 mils per share (from column (1)), or 0.32 bp (from column (2)). The entry in column (3), 3.99, implies that the price risk is roughly four times what would be implied by a random-walk calibrated to longest time scale in the analysis (27.3 minutes). At 400 ms, the rough volatility risk crosses the one mil threshold (1.06, column (2)). At 1,600 ms, it is on the order of one basis point. The variance ratios (columns (3) and (6)) increase monotonically in moving to shorter time scales.

Column (7) of Table 2 reports the wavelet correlations between bids and offers. The wavelet covariance between two processes is defined analogously to the wavelet variance: the covariance between the bid and offer is denoted $v_{bid,offer,j}^2$. The wavelet correlation, denoted $\rho_{bid,offer,j} = v_{bid,offer,j}^2 / \sqrt{v_{bid,j}^2 v_{offer,j}^2}$, is used to assess the extent to which the bid and offer co-move at different time scales. If the bid and offer always moved in lock step, this correlation would be unity at every time scale. At longer time scales this correlation is indeed quite high, but at shorter time scales it is only moderately positive.

Table 3 reports results for a subset of the measures and time scales, but provides standard errors and detail across dollar volume subsamples. Panels A and B report estimates of rough variances in mils per share and basis points, respectively. Stocks in the two lowest dollar volume quintiles have sharply higher short-term volatility. In comparing the two normalizations, it is apparent that volatility in mils per share (Panel A) at the shorter scales is somewhat more stable across dollar volume quintiles than volatility in basis points (Panel B). In moving from lowest to highest quintiles, short time scale volatilities in mils per share approximately double; while volatilities in basis points decline by a factor of four. This decline appears to be mostly caused by

the increase in share prices across the quintiles (cf. Table 1). Put another way, it appears that quote volatility is best characterized as a “mils per share” phenomenon, perhaps due to the tick size effects or the use of per-share cost schedules in assessing trading fees.

Table 3 Panel C reports selected rough variance ratios across dollar volume quintiles. Figure 2 graphs the average wavelet variance ratios. From this graph, for the highest volume quintile, the excess variance seems to be about 30% at the shortest time scales. For the lowest volume quintile, however, the excess is, at ten or above, substantially higher. The wavelet bid-offer correlations are reported in Table 3 Panel D, and graphed in Figure 3. These also exhibit marked variation across dollar volume. For the highest quintile, they are close to unity at a time scale of 25.6 seconds; for the lowest, the correlation at 27.4 minutes is a modest 0.52. This indicates a pronounced de-coupling of the bid and offer for smaller firms.

Hansen and Lunde (2006) note that to the extent that volatility is fundamental, we would expect bid and offer variation to be perfectly correlated, that is, that a public information revelation would shift both prices by the same amount. Against this presumption, the short-term correlation estimates are striking. At time scales of 200 ms or lower, the correlation is below 0.7 for all activity quintiles. For the shortest time scales and lower activity quintiles, the correlation is only slightly positive. This suggests that substantial high frequency quote volatility is of a distinctly transient nature.

V. Analysis with truncated time stamps.

The analysis in the preceding section relies on a recent one-month sample of daily TAQ data. For addressing policy issues related to low-latency activity, it would be useful to conduct a historical analysis, spanning the period over which low-latency technology was deployed. Extending the analysis backwards, however, is not straightforward. Millisecond time-stamps are only available in the daily TAQ data from 2006 onwards. Monthly TAQ data (the standard source used in academic research) is available back to 1993 (and the precursor ISSM data go back to the mid-1980s). These data are substantially less expensive than the daily TAQ, and they have a simpler logical structure.

The time stamps on the Monthly TAQ and ISSM datasets are reported only to the second, a limitation that might seem to render these data useless for characterizing sub-second variation. This is unduly pessimistic. It is the purpose of this section to propose, implement and validate an approach for estimating sub-second characteristics of the bid and offer series using the second-

stamped data. This is possible because the data generation and reporting process is richer than it initially seems.

The usual sampling situation in discrete time series analysis involves either aggregation over periodic intervals (such as quarterly GDP) or point-in-time periodic sampling (such as the end-of-day S&P index). In both cases there is one observation per interval, and in neither case do the data support resolution of components shorter than one interval. In the present situation, however, quote updates occur in continuous time and are disseminated continuously. The one second time-stamps arise as a truncation (or equivalently, a rounding) of the continuous event times. The Monthly TAQ data include all quote records, and it is not uncommon for a second to contain ten or even a hundred quote records.

Assume that quote updates arrive as a Poisson process of constant intensity. If the interval $(0, t)$ contains n updates, then the update times have the same distribution as the order statistics in a sample of n independent random variables uniformly distributed on the interval $(0, t)$, (Ross, 1996, Theorem 2.3.1). Within a one-second interval containing n updates, therefore, we can simulate continuous arrival times by drawing n realizations from the standard uniform distribution, sorting, and assigning them to quotes (in order) as the fractional portions of the arrival times. These simulated time-stamps are essentially random draws from true distribution. This result does not require knowledge of the underlying Poisson arrival intensity.

We make the additional assumption that the quote update times are independent of the updated bid and offer prices. (That is, the “marks” associated with the arrival times are independent of the times.) Then all estimates based on the simulated time stamp series constitute draws from their corresponding posterior distributions. This procedure can be formalized in a Bayesian Markov-Chain Monte Carlo (MCMC) framework. To refine the estimates, we would normally make repeated simulations (“sweeps”) over the sample, but due to computational considerations and programming complexity, I make only one draw for each CQ record.

It is readily granted that few of the assumptions underlying this model are completely satisfied in practice. For a time-homogeneous Poisson process, interevent durations are independent. In fact, inter-event times in market data frequently exhibit pronounced serial dependence, and this feature is a staple of the autoregressive conditional duration and stochastic duration literature (Engle and Russell, 1998; Hautsch, 2004). In NASDAQ data, Hasbrouck and Saar (2013) show that event times exhibit intra-second deterministic patterns. Subordinated

stochastic process models of security prices suggest that transactions (not wall-clock time) are effectively the “clock” of the process (Shephard, 2005).

We can assess the reliability of the randomization approach, however, by a simple test. The time-stamps of the data analyzed in the last section are stripped of their millisecond remainders. New millisecond remainders are simulated, the random-time-stamped data are analyzed, and we examine the correlations between the two sets (original and randomized) of estimates. Let $v_{bid,i,j,d}^2$ denote the bid wavelet variance estimate for firm i on day d at level j based on the original time stamps, and let $\tilde{v}_{bid,i,j,d}^2$ denote the estimate based on the simulated time stamps. (Results are similar for estimates on the offer side.) Table 4, Panel A reports estimates across firms and days of $Corr(v_{bid,i,j,d}^2, \tilde{v}_{bid,i,j,d}^2)$. The agreement between original and randomized estimates is very high for all time scales and in all subsamples. Even at the time scale of less than fifty ms, the mean correlation is 0.952. At time scales above one second, the agreement is nearly perfect.

Given the questionable validity of some of the assumptions, and the fact that only one draw is made for each second’s activity, this agreement might seem surprising. It becomes more reasonable, however, when one considers the extent of averaging underlying the construction of both original and randomized estimates. There is explicit averaging in that each wavelet variance estimate is formed over a sample of roughly six hours. As long as the order is maintained, a small shift in a data point has little impact over the overall estimate.¹²

In parallel fashion let $v_{bid,offer,i,j,d}^2$ and $\tilde{v}_{bid,offer,i,j,d}^2$ denote the bid-offer covariance estimates based (respectively) on original and simulated millisecond time stamps. Table 4, Panel B reports estimates across firms and days of $Corr(v_{bid,offer,i,j,d}^2, \tilde{v}_{bid,offer,i,j,d}^2)$. The agreement is somewhat weaker than in the case of the variances. The correlation of under-50 ms components is 0.775 (in the full sample), although this climbs to 0.979 at a time scale of 200 ms. The reason for the relatively poorer performance of the randomized covariance estimates is simply that the wavelet covariance between two series is sensitive to alignment. For a given CQ record, the bid and offer quotes are paired, but in a typical record sequence the NBB and NBO are not changed in the same record. When a bid change is shifted even by a small amount relative to the offer, the inferred pattern of co-movement is distorted.

¹² Also, inherent in the wavelet transformation is an (undesirable) averaging across time scales known as leakage, wherein the variance at one time scale affects to a small degree the estimate at neighboring time scale (Percival and Walden, p. 303).

Across dollar volume quintiles, the correlations generally improve for all time scales. This is true for both wavelet variances and covariances, but is more evident in the latter. This is a likely consequence of the greater incidence, in the higher quintiles, of multiple quote records within the same second. Specifically, for a set of n draws from the uniform distribution, the distribution of any order statistic tightens as n increases. (For example, the distribution of the first order statistic in a sample of five hundred is tighter than the distribution of the first order statistic in a sample of one.) Essentially, an event time can be located more precisely within the second if the second contains more events. This observation will have bearing on the analysis of historical samples with varying numbers of events.

In working with Monthly TAQ data, (Holden and Jacobsen, 2013, HJ) suggest assigning sub-second time stamps by evenly-spaced interpolation. If there is one quote record in the second, it is assigned a millisecond remainder of 0.500 seconds; if two records, 0.333 and 0.667 seconds, and so on. HJ show that interpolation yields good estimates of effective spreads. It is not, however, equivalent to the present approach. Consider a sample in which each one-second interval contains one quote record. Even spacing places each quote at its half-second point. As a result, the separation between each quote is one second. For example, a sequence of second time stamps such as 10:00:01, 10:00:02, 10:00:03 ... maps to 10:00:01.500, 10:00:02.500, 10:00:03.500, and so on. The interpolated time stamps are still separated by one second, and therefore the sample has no information regarding sub-second components. In contrast, a randomized procedure would sweep the space of all possibilities, including 10:00:01.999, 10:00:02.000, ..., which provides for attribution of one-millisecond components. Of course, as the number of events in a given one-second interval increases, the two approaches converge: the distribution of the k th order statistic in a sample of n uniform observations collapses around its expectation, $k/(n+1)$ as n increases.¹³

¹³ For one class of time-weighted statistics in this setting, interpolated time stamps lead to unbiased estimates. Consider a unit interval where the initial price, p_0 , is known, and there are n subsequent price updates $p_i, i=1, \dots, n$ at occurring at times $0 < t_1 < \dots < t_n < 1$. The time-weighted average of any price function $f(p)$ is $Avg^{TW} = \sum_{i=0}^n f(p_i)(t_{i+1} - t_i)$, where $t_0 \equiv 0$ and $t_{n+1} \equiv 1$. Assuming a time-homogeneous Poisson arrival process, the t_i are distributed (as above) as uniform order statistics. This implies $E t_i = i/(n+1)$, the linear interpolated values. If the marks (the p_i) are distributed independently of the t_i , $E[Avg^{TW}] = (n+1)^{-1} \sum_{i=0}^n f(p_i)$. This result

VI. Historical evidence

This section describes the construction and analysis of variance estimates for a sample of US stocks from 2001 to 2011. In each year, I construct variance estimates for a single representative month (April) for a subsample of firms.

The period covers significant changes in market structure and technology. Decimalization had been mandated, but was not completely implemented by April, 2001. Reg NMS was proposed, adopted, and implemented.¹⁴ Dark trading grew over the period. Market information and access systems were improved, and latency emerged as a key concern of participants. The period also includes many events related to the financial crisis, which are relatively exogenous to equity market structure.

The regulatory and technological shifts over the period caused changes in the fundamental nature of bid and offer quotations. Markets in 2001 were still dominated by what would later be called “slow” procedures. Quotes were often set manually. Opportunities for automated execution against these quotes were few (cf. the NYSE’s odd-lot system, and NASDAQ’s Small Order Execution System). Trade-through protection was limited and weakly enforced. Quotes for 100 shares or less were not protected. With the advent of Reg NMS, the bids and offers became much more accessible (for automated execution). These considerations are important in interpreting the results that follow.

VI.A. Data

The data for this phase of the analysis are drawn from CRSP and *Monthly* TAQ datasets. The sample selection procedure in each year is essentially identical to that described for the 2011 cross-sectional sample. In each year, from all firms present on CRSP and TAQ in April, with share codes in (10 and 11), and with primary listings on the NYSE, Amex and NASDAQ exchanges, I draw fifteen firms from each dollar trading volume decile.¹⁵ Quote data are drawn from TAQ.

applies to time-weighted means of prices and spreads (assuming simultaneous updates of bids and offers). It also applies to wavelet transforms and other linear convolutions. It does not apply to variances (or wavelet variances), however, which are nonlinear functions of arrival times.

¹⁴ Reg NMS was proposed in February, 2004, and adopted in June 2005 with an effective date of August 2005. It was implemented in stages, mostly over 2006.

¹⁵ As of April, 2001, NASDAQ had not fully implemented decimalization. For this year, I do not sample from stocks that traded in sixteenths.

Table 5 reports summary statistics. The oft-remarked increase in the intensity of trading activity is clearly visible in the trends for median number of trade and quote records. From 2001 to 2011, the average annual compound growth rate is about 23% percent for trades, and about 32% for quotes. As described in the last section, all of a firm's quote records in a given second are assigned random, but order preserving, millisecond remainders. The NBBO is constructed from these quote records. This yields a NBBO series with (simulated) millisecond time stamps. The 2011 numbers differ slightly from those reported in Table 1. These differences are a consequence of the randomization, and (as will be indicated) the use of error filters.

Prior to the construction of the NBBO the bid and offer are filtered for extreme values. The following quotes (bids or offers) are eliminated: those with zero size and/or zero price; those quotes priced at 20% or lower of the smallest closing price reported on CRSP in the month; those priced at 500% or higher of highest closing price. Quotes that crossed the market are only eliminated if the crossing is a dollar or more, or more than 10 percent of the midpoint price. Other filters use the previously prevailing bid and offer midpoint as a benchmark. For stocks priced at ten dollars or less, the bid and offer has to be within forty percent of the benchmark; for stocks between ten and one hundred dollars, the cutoff is twenty percent; for stocks between one hundred and 250, ten percent; above 250, five percent.¹⁶ These filters do not eliminate all suspicious bids and offers, a point to which the discussion will subsequently return.

VI.B. Results

In analyzing 2001-2011, it is best to begin with the wavelet variance ratios. By construction they are normalized with respect to long-term variance, and over this period there are large swings in market-wide long-term volatility (evident from a cursory examination of the VIX). These would be expected to affect the short term variances as well. Table 6 Panel A reports the mean wavelet variance ratios for the shorter time scales. As in the 2011 sample, there is substantial variance inflation relative to the random-walk in all years. Perhaps surprisingly, though, the excess variance is high in all years, including the early years of the decade. The estimates are higher in 2001 than in 2011. The pattern does not suggest an increasing trend.

¹⁶ The error filters are applied uniformly for the Monthly TAQ data in all years 2001-2011. For 2011 this causes a small apparent discrepancy in the counts for NBB and NBO changes, between Tables 1 and 5. The inputs to Table 5 are filtered, and hence have slightly fewer NBB and NBO changes relative to the unfiltered inputs to Table 1.

Given the recent media attention devoted to low-latency activity and the undeniable growth in quote volume, the absence of a strong trend in quote volatility seems surprising. There are several possible explanations. In the first place, “flickering quotes” drew comment well before the start of the sample, in an era when quotes were dominated by human market makers (Harris, 1999; U.S. Commodities Futures Trading Commission Technology Advisor Committee, 2001). Also an artifact of this era is the specialist practice of “gapping” the quotes to indicate larger quantities at worse prices (Jennings and Thirumalai, 2007). In short, the quotes may have in reality been less unwavering than popular memory holds. The apparent discrepancy between quote volatility and quote volume can be explained by appealing to the increase in market fragmentation and consequent growth in matching quotes.

Exploring this finding further, bid-offer plots for firm-days in each year that correspond to extreme realizations of the variances exhibit an interesting pattern. In later years, these outlier plots tend to resemble the initial AEPI example, with rapid oscillations of relatively low amplitude. In the earlier years, they are more likely to feature small number of prominent spikes associated with a sharply lower bid or elevated offer that persists for a minute or less.

As an example, Figure 4 (Panel A) depicts the NBBO for PRK (Park National Corporation, Amex-listed) on April 6, 2001. At around 10:00 there is a downward spike in the NBB. Shortly after noon there is a sharp drop in the NBB of roughly three dollars and a sharp rise in the NBO of about one dollar. To better document this behavior, Table 7 details the CQ records in the vicinity of the noon episode. There are multiple exchanges active in the market, but Amex (A) is the apparent price leader. At 12:02:22, A establishes the NBB at 86.74. At 12:03:11, A bids 83.63, exposing the previous *T* (NASDAQ) bid of 86.68 as the new NBB. At 12:03:16, *T* backs off, leaving A’s 83.63 as best. Within half a minute, however, the NBB is back at 86.50. The lower bid is not marketed by any special mode flag. It is not a penny (“stub”) bid. The size of the bid at two (hundred shares) is typical for the market on that day. A similar sequence of events sends the NBO up a dollar for about one second.

These quotes are not so far off the mark as to be clearly erroneous. We must nevertheless question whether they were “real”? Did they reliably indicate the consensus market values at those instances? Were they accessible for execution? Were they truly the best in the market? There were no trades between 11:38 and 12:13, but if a market order had been entered, would it

in fact have been executed at the NBBO?¹⁷ These are meaningful questions because they bear directly on market quality. Ultimately, though, the record is unlikely to provide clear answers. The US equity market in 2001 reflected a blend of human and automated mechanisms, practices and conventions that defies detailed description even at a distance of only twelve years.

Discerning whether or not quote volatility increased over the period, therefore, requires that we sharpen the question. The quote volatility in the initial AEPI example is of high frequency, but low amplitude. This is visually distinct from the spikes of high frequency and high amplitude found in PRK. The latter is sometimes called “pop” noise, in reference to its sound in audio signals (Walker, 2008). As in the de-noising of audio signals, the goal is to remove the pops from the signals of lower amplitude. The wavelet literature has developed many denoising approaches (see Percival and Walden, Gençay et al, and Walker). When the stochastic properties of the noise and signal processes are known, optimal methods can often be established. In the present case, though, I adopt a simpler method.

As indicated in Section III, wavelet transforms facilitate the direct computation of smooth and rough components. This process, known as multiresolution analysis, isolates components at different time scales. As an example, Panel B of Figure 4 plots the rough component of the PRK bid at a time scale of 51.2 seconds. It is zero mean by construction, and the spikes are cleanly resolved. On the principle that high frequency quoting (as in the AEPI example) should not be substantially larger than the bid-offer spread in magnitude, I set acceptance bands at $\pm \text{Min}(1.5 \times (\text{average spread}), \$0.25)$. The minimum of \$0.25 is set to accommodate stocks with very tight spreads. For PRK, the bands are approximately $\pm \$0.33$, and they are indicated in the figure by horizontal black lines. Values lying outside of the band are set to the band limits. This clips the high-amplitude peaks, while leaving the low-amplitude components, some of which are highly oscillatory, untouched. The signal (bid or offer) is reconstituted using the clipped rough, and analysis proceeds on this denoised signal. I recompute all estimates for all firms using the denoised bids and offers.

Table 6 Panel B reports the wavelet variance ratios for the denoised quotes. The results are striking. In the early years, the variance ratios computed from the denoised quotes are much

¹⁷ The Amex (like the NYSE) had specialists in 2001. Specialists generally had affirmative price continuity obligations that would have discouraged (though not expressly forbidden) trades occurring at prices substantially different from those prevailing immediately before and immediately after. A broker-dealer, however, would not have been subject to this restriction.

lower than those computed from the raw data. In later years, however, the reduction associated with the denoising is small. For the 200 ms variance ratio, for example, the 2001 drop is from 5.35 (for the raw quotes) to 1.56 (for the denoised quotes), but the 2011 value only drops from 3.75 to 3.57.

These results are consistent with the view that the overall level of quote volatility did not change very much over the decade. The nature of the volatility has apparently, however, evolved. In the early years, the volatility was of relatively high amplitude but non-oscillatory. It is removed by the pop-denoising procedure. The procedure does not attenuate the low-amplitude highly oscillatory components, however, which drive quote volatility in the later years. The difference between the raw and denoised ratios generally declines throughout the decade, but the convergence appears to be strongest during the Reg NMS transition period.

The denoising procedure accentuates low-amplitude oscillatory volatility. Since one might expect that this would be tied more closely to low-latency technology, it is sensible to ask whether the denoised volatilities have increased. Table 8 therefore presents rough volatilities for the denoised quotes in mils per share (Panel A) and basis points (Panel B), and as a variance ratio (Panel C) for a representative subset of time scales. Figure 5 plots these quantities at the 800 ms time scale. The table and figure suggest that neither the mils per share volatility (Panel A) nor the basis point volatility (Panel B) evinces an upward trend. The variance ratio (Panel C) appears to climb from 2001 to 2004, but thereafter drifts distinctly downwards. In summary, the climb that might be expected from ongoing enhancements to trading technology or the growth in quote traffic is conspicuously absent.

VII. Discussion

From an economic perspective, high frequency quote volatility is connected most closely to other high frequency and low latency phenomena in modern markets. From a statistical perspective, it is connected to volatility modeling. I discuss both in turn.

VII.A. High frequency quoting and high frequency trading

Most definitions of algorithmic and high frequency trading encompass many aspects of market behavior (not just executions), and would be presumed to cover quoting as well.¹⁸ Executions and quotations are nevertheless very different events. It is therefore useful to consider their relation in the high frequency context.

Quote volatility is not necessarily associated with high frequency executions. One can envision regimes where relatively stable quotes are hit intensively when fundamental valuations change, and periods (such as Figure 1) where frenetic quoting occurs in the absence of executions. Nevertheless, the same technology that makes high frequency executions possible also facilitates the rapid submission, cancellation and repricing of the nonmarketable orders that define the bid and offer. One might expect this commonality of technology to link the two activities in practice.

Executions are generally emphasized over quotes when identifying agents as high frequency traders. For example, (Kirilenko, Kyle, Samadi, and Tuzun, 2010) select on high volume and low inventory. The low inventory criterion excludes institutional investors who might use algorithmic techniques to accumulate or liquidate a large position. The NASDAQ HFT dataset uses similar criteria (Brogaard, 2012; Brogaard, Hendershott, and Riordan, 2012). Once high frequency traders are identified, their executions and the attributes of these executions lead to direct measures of HF activity in panel samples.

In some situations, however, identifications based on additional, non-trade information are possible. Menkveld (2013) identifies one Chi-X participant on the basis of size and prominence. The Automated Trading Program on the German XETRA system allows and provides incentives for designating an order as algorithmic (Hendershott and Riordan, 2013). Other studies analyze indirect measures of low-latency activity. Hendershott, Jones, and Menkveld (2011) use NYSE message traffic. Hasbrouck and Saar (2013) suggest strategic runs (order chains) of cancel and replace messages linked at intervals of 100 ms or lower.

¹⁸ A CFTC draft definition reads: “High frequency trading is a form of automated trading that employs: (a) algorithms for decision making, order initiation, generation, routing, or execution, for each individual transaction without human direction; (b) low-latency technology that is designed to minimize response times, including proximity and co-location services; (c) high speed connections to markets for order entry; and (d) high message rates (orders, quotes or cancellations)” (U.S. Commodities Futures Trading Commission, 2011).

Most of these studies find a positive association between low-latency activity and market quality. Low-latency activity, for example, tends to be negatively correlated with ask posted and effective spreads, which are inverse measures of market quality. Most also find a zero or negative association between low-latency activity and volatility, although the constructed volatility measures usually span intervals that are long relative to those of the present paper. With respect to algorithmic or high frequency activity, Hendershott and Riordan (2012) find an insignificantly negative association with the absolute value of the prior 15-minute return; Hasbrouck and Saar (2013) find a negative association with the high-low difference of the quote midpoint over 10-minute intervals.

The time-scaled variance estimates used here clearly aim at a richer characterization of volatility than the high/low or absolute return proxies used in the studies above. The present study does not, on the other hand, attempt to correlate the variance measures with intraday proxies for high frequency trading. One would further suspect, of course, that the ultimate strategic purpose of high frequency quoting is to facilitate a trade or to affect the price of a trade. The mechanics of this are certainly deserving of further research.

The discussion in Section II associates short-term quote volatility with price uncertainty for those who submit marketable orders, use dark mechanisms that price by reference, or face monitoring difficulties. From this perspective, quote volatility is an inverse measure of market quality. Although the present study finds evidence of economically significant and elevated quote volatility, it does not establish a simple connection to technological trends associated with low latency activity.

VII.B. High frequency quoting and volatility modeling

Security prices at all horizons are a mix of integrated and stationary components. The former are usually identified with persistent fundamental information innovations; the latter, with transient microstructure effects. The former are important to long-term hedging and investment; the latter, to trading and market-making. The dichotomy is sometimes reflected in different statistical tools and models.

Between the two approaches, the greatest common concerns arise in the analysis of realized volatility (Andersen, Bollerslev, Diebold, and Ebens, 2001; Andersen, Bollerslev, Diebold, and Labys, 2003b, a). RVs are calculated from short-term price changes. They are useful as estimates of fundamental integrated volatility (IV), and typically serve as inputs to

longer-term forecasting models. RVs constructed directly from trade, bid and offer prices are typically noisy, however, due to the presence of microstructure components. Local averaging moderates these effects [see Hansen and Lunde (2006) and accompanying comments]. Other approaches are discussed in (Zhang, Mykland, and Aït-Sahalia, 2005; Zhang, 2006; Aït-Sahalia, Mykland, and Zhang, 2011). There is a methodological connection here, in that long-term wavelet variances are computed from short-term averages, much like the pre-averaged inputs to realized volatility.

The present study draws on several themes in the RV literature. The volatility ratio plots in Figure 2 serve a purpose similar to the volatility signature plots introduced by Fang (1996), and used by Andersen, Bollerslev, Diebold, and Ebens (2002) and Hansen and Lunde (2006), among others Hansen and Lunde also articulate the connection between bid-offer comovement and fundamental volatility: since the bid and offer have economic fundamentals in common, divergent movements must be short-term, transient, and unconnected to fundamentals.

One strand in the RV literature emphasizes analysis of multiple time-scales. Zhang *et al.* (2005) posit a framework consisting of a Brownian motion with time-varying parameters, $dX_t = \mu_t dt + \sigma_t dz$, and a discretely-sampled noisy observation process, $Y_t = X_t + \varepsilon_t$. The Y_t are viewed as transaction prices, and ε_t constitute i.i.d. microstructure noise. The objective is estimation of the integrated volatility $\int \sigma_t^2 dt$ over a sample. They propose a two-scale variance estimator in which a long-scale estimate is corrected for bias with an adjustment based on properties of the noise estimated at a short scale. While the present analysis also features multiple time scales, there are major differences in the perspective. In the present situation, execution price risk is caused by volatility in the observed process (the quote, not the underlying latent value, X_t); the quote process is right-continuous (and continuously observable); the noise is not necessarily i.i.d. (cf. the AEPI episodes in Figure 1); and, the noise is possibly correlated with the X_t increments.

The paper also departs from the RV literature in other respects. The millisecond time scales employed in this paper are several orders of magnitude shorter than those typically encountered. Most RV studies also focus on relatively liquid assets (index securities, Dow-Jones stocks, etc.). The low-activity securities included in the present paper's samples are important because, due to their larger spreads and fewer participants, they are likely to exhibit relatively strong, persistent and distinctive microstructure-related components.

VIII. Conclusion

High frequency volatility in the bid and offer quotes induces price risk for agents who experience delay in communicating with the market. The risk may be quantified as the price variance over the interval of delay, relative to the average price over the interval. This volatility degrades the informational value of the quotes. Furthermore, because the bid and offer are often used as reference prices for dealer trades against customers, the volatility increases the value of a dealer's look-back option and exacerbates monitoring problems for customers, exchanges, and regulators.

This study is a preliminary analysis of short-term quote volatility in the US equity market. Estimates of sub-second high frequency variance for the National Best Bid and Offer (NBBO) are well in excess of what would be expected relative to random-walk volatility estimated over longer intervals. The excess volatility is more pronounced for stocks that have lower average activity. The sub-second volatility is comparable in magnitude to access fees and other transaction expenses. Furthermore, the correlations between bids and offers at these time scales are positive, but low. That the bid and offer are not moving together also suggests that the volatility is not fundamental.

The paper proposes a simulation approach to measuring millisecond-level volatility in US equity data (like the Monthly TAQ) that possess all quote records, but are time-stamped only to the second. In data time-stamped to the millisecond I compare two sets of estimates: one set based on the original time-stamps; the other based on simulated time stamps. I find high correlations between the two estimates, establishing the reliability of the simulation procedure.

With these results, the paper turns to a longer US historical sample, 2001-2011, with one-second time-stamps. Despite the current public scrutiny of high frequency trading, the rapid growth in the number of quote records, and the presumption that low-latency technology is a new and recent phenomenon, the excess short-term quote volatility found in the 2011 data also appears in earlier years. The nature of quote volatility has changed over the decade, however. In the earlier years, the volatility apparently arises from spikes in with bids and offers that are neither clearly erroneous nor reliably valid. In later years the volatility is more attributable to oscillatory low-amplitude changes: rapid movements not substantially larger than the spread. The highest excess volatilities, though are found in 2004-2006, a period that corresponds to the discussion and implementation of Reg NMS. These mid-decade volatilities, therefore, might constitute a transitional or learning period as markets scrambled to adapt to a new era.

Appendix: Deviations about averages of random walks

Consider a price series that evolves as $p_t = p_{t-1} + u_t$ where u_t is a white-noise process with unit variance. Without loss of generality, we initialize $p_0 = 0$ and consider the mean-squared deviations about the mean over the first n observations.

$$MSD(n) = \frac{1}{n} \sum_{i=1}^n p_i^2 - \left(\frac{1}{n} \sum_{i=1}^n p_i \right)^2 = \frac{1}{n} \sum_{i=1}^n \left(\sum_{j=1}^i u_j \right)^2 - \left(\frac{1}{n} \sum_{i=1}^n \left(\sum_{j=1}^i u_j \right) \right)^2$$

Taking expectations (noting that $E u_i u_j = 1$ for $i = j$, and zero otherwise) and simplifying the sums gives

$$\sigma_n^2 \equiv E(MSD(n)) = \frac{n+1}{2} - \frac{(n+1)(2n+1)}{6n} = \frac{n^2-1}{6n}$$

For the sequence of averaging periods $n_j = n_0 2^j$ for $j = 0, 1, 2, \dots$, the corresponding sequence of variances is

$$\sigma_j^2 = \frac{4^j n_0^2 - 1}{3n_0 2^{j+1}}$$

In moving from $j-1$ to j the incremental change in variance (also known as the wavelet variance) is

$$v_j^2 = \sigma_j^2 - \sigma_{j-1}^2 = \frac{4^j n_0^2 + 2}{3n_0 2^{j+2}}$$

We now reinterpret these results in a slightly expanded framework. Suppose that the original time subscript t indexes periods of Δ time units (“milliseconds”) and that the variance per unit time of the u_t process is σ_u^2 . Let M denote the averaging period measured in units of time, and correspondingly, $M_j = M_0 2^j$ for $j = 0, 1, \dots$. Then the rough and wavelet variances become

$$\sigma_j^2 = \frac{(4^j M_0^2 - \Delta^2) \sigma_u^2}{3M_0 2^{j+1}} \quad \text{and} \quad v_j^2 = \frac{(4^j M_0^2 + 2\Delta^2) \sigma_u^2}{3M_0 2^{j+2}}.$$

In the continuous time limit, as $\Delta \rightarrow 0$, that $\sigma_j^2 = 2^{j-1} M_0 \sigma_u^2 / 3$ and $v_j^2 = 2^{j-2} M_0 \sigma_u^2 / 3$. These results suffice to define and characterize the variances considered in the paper.

References

- Aït-Sahalia, Y., Mykland, P.A., Zhang, L., 2011. Ultra high frequency volatility estimation with dependent microstructure noise. *Journal of Econometrics* 160, 160-175.
- Amihud, Y., Mendelson, H., 1987. Trading mechanisms and stock returns: An empirical investigation. *Journal of Finance* 42, 533-553.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2002. Great realizations. *Risk* 13, 105-108.
- Andersen, T.G., Bollerslev, T.P., Diebold, F.X., Ebens, H., 2001. The distribution of realized stock return volatility. *Journal of Financial Economics* 61, 43-76.
- Andersen, T.G., Bollerslev, T.P., Diebold, F.X., Labys, P., 2003a. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* 98, 501-501.
- Andersen, T.G., Bollerslev, T.P., Diebold, F.X., Labys, P., 2003b. Modeling and forecasting realized volatility. *Econometrica* 71, 579-625.
- Atsalakis, G.S., Valavanis, K.P., 2009. Surveying stock market forecasting techniques - Part II: Soft computing methods. *Expert Systems with Applications* 36, 5932-5941.
- Barnea, A., 1974. Performance evaluation of New York Stock Exchange specialists. *Journal of Financial and Quantitative Analysis* 9, 511-535.
- Baruch, S., Glosten, L.R., 2013. Fleeting orders and the competitive equilibrium. Working paper, Columbia University.
- Bessembinder, H., Panayides, M.A., Venkataraman, K., 2009. Hidden liquidity: an analysis of order exposure strategies in electronic stock markets. *Journal of Financial Economics*.
- Biais, B., Foucault, T., Moinas, S., 2012. Equilibrium high-frequency trading. Working paper, University of Toulouse.
- Brogaard, J., 2012. Essays on High-Frequency Trading. PhD Dissertation. Northwestern University.
- Brogaard, J., Hendershott, T.J., Riordan, R., 2012. High frequency trading and price discovery. Working paper, University of Washington.
- Cornish, C., 2006. The WMTSA wavelet toolkit for MATLAB Department of Atmospheric Sciences, University of Washington, Seattle.
- Dacorogna, M.M., Gencay, R., Muller, U.A., Olsen, R.B., Pictet, O.B., 2001. High-Frequency Finance. Academic Press, New York.
- Egginton, J., Van Ness, B.F., Van Ness, R.A., 2012. Quote stuffing. Working paper, University of Mississippi.
- Elder, J., Jin, H.J., 2007. Long memory in commodity futures volatility: A wavelet perspective. *Journal of Futures Markets* 27, 411-437.

- Elkins-McSherry, 2012. Methodology.
<<https://www.elkinsmcsberry.com/EM/methodology.html>>.
- Engle, R.F., Ferstenberg, R., Russell, J., 2012. Measuring and modeling execution cost and risk. *Journal of Portfolio Management* 38, 14-28.
- Engle, R.F., Russell, J.R., 1998. Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica* 66, 1127-1162.
- Fan, J., Wang, Y., 2007. Multi-scale jump and volatility analysis for high-frequency financial data. *Journal of the American Statistical Association* 102, 1349-1362.
- Fan, Y., Gençay, R., 2010. Unit root tests with wavelets. *Econometric Theory* 26, 1305-1331.
- Fang, Y., 1996. Volatility Modeling and Estimation of High-Frequency Data with Gaussian Noise. PhD Dissertation. MIT Sloan School.
- Gençay, R., Selçuk, F., Gradojevic, N., Whitcher, B., 2010. Asymmetry of information flow between volatilities across time scales. *Quantitative Finance* 10, 895-915.
- Gençay, R., Selçuk, F., Whitcher, B., 2002. *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. Academic Press (Elsevier), San Diego.
- Gençay, R., Selçuk, F., Whitcher, B., 2005. Multiscale systematic risk. *Journal of International Money and Finance* 24, 55-70.
- Gençay, R., Signori, D., 2012. Multi-scale tests for serial correlation. Working paper, Simon Fraser University.
- Hansen, P.R., Lunde, A., 2006. Realized variance and market microstructure noise. *Journal of Business & Economic Statistics* 24, 127-161.
- Harris, L.E., 1999. Trading in pennies: a survey of the issues. Working paper, Marshall School, University of Southern California.
- Hasbrouck, J., Saar, G., 2013. Low-Latency Trading. *Journal of Financial Markets*, forthcoming.
- Hasbrouck, J., Schwartz, R.A., 1988. Liquidity and execution costs in equity markets. *Journal of Portfolio Management* 14, 10-16.
- Hautsch, N., 2004. *Modelling Irregularly Spaced Financial Data: Theory and Practice of Dynamic Duration Models*. Springer.
- Hendershott, T.J., Jones, C.M., Menkveld, A.J., 2011. Does algorithmic trading improve liquidity? *Journal of Finance* 66, 1-33.
- Hendershott, T.J., Riordan, R., 2012. High Frequency Trading and Price Discovery. SSRN eLibrary.
- Hendershott, T.J., Riordan, R., 2013. Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis*, forthcoming, *forthcoming*.
- Høg, E., Lunde, A., 2003. Wavelet estimation of integrated volatility. Working paper, Aarhus School of Business, Department of Information Science.

- Holden, C.W., Jacobsen, S.E., 2013. Liquidity measurement problems in fast, competitive markets: expensive and cheap solutions. *Journal of Finance*, forthcoming.
- Hsieh, T.-J., Hsiao, H.-F., Yeh, W.-C., 2011. Forecasting stock markets using wavelet transforms and recurrent neural networks: An integrated system based on artificial bee colony algorithm. *Applied Soft Computing* 11, 2510-2525.
- In, F., Kim, S., 2006. The hedge ratio and the empirical relationship between the stock and futures markets: a new approach using wavelet analysis. *The Journal of Business* 79, 799-820.
- Jarrow, R.A., Protter, P., 2012. A dysfunctional role of high-frequency trading in electronic markets. *International Journal of Theoretical and Applied Finance* 15.
- Jennings, R.H., Thirumalai, R.S., 2007. Advertising for liquidity on the New York Stock Exchange. Working paper, Kelley School, University of Indiana.
- Kirilenko, A.A., Kyle, A.S., Samadi, M., Tuzun, T., 2010. The flash crash: the impact of high frequency trading on an electronic market. Working paper, CFTC and University of Maryland.
- Lo, A.W., MacKinlay, A.C., 1988. Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies* 1, 41-66, Record Number: 2140.
- Menkveld, A.J., 2013. High-frequency trading and the new market-makers. *Journal of Financial Markets*, forthcoming.
- Nason, G.P., 2008. *Wavelet Methods in Statistics with R*. Springer Science+Business Media, LLC, New York.
- O'Hara, M., Ye, M., 2009. Is market fragmentation harming market quality? Working paper, Cornell University.
- Parlour, C.A., Seppi, D.J., 2003. Liquidity-based competition for order flow. *Review of financial studies* 16, 301-303.
- Percival, D.B., Walden, A.T., 2000. *Wavelet Methods for Time Series Analysis*. Cambridge University Press, Cambridge.
- Perold, A., 1988. The implementation shortfall: paper vs. reality. *Journal of Portfolio Management* 14, 4-9, Record Number: 6070.
- Ramsey, J.B., 1999. 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, 2593-2606.
- Ramsey, J.B., 2002. Wavelets in economics and finance: past and future. *Studies in Nonlinear Dynamics and Econometrics* 6.
- Ross, S.M., 1996. *Stochastic Processes*. John Wiley and Sons, Inc., New York.
- Shephard, N., 2005. General introduction. In: Shephard N (ed.) *Stochastic Volatility*. Oxford University Press, Oxford.

- Stoll, H.R., Schenzler, C., 2006. Trades outside the quotes: Reporting delay, trading option, or trade size? *Journal of Financial Economics* 79, 615-653.
- Teyssière, G., Abry, P., 2007. Wavelet analysis of nonlinear long-range dependent processes: applications to financial time series. In: Teyssière G & Kirman AP (eds.) *Long Memory in Economics*. Springer, Berlin.
- U.S. Commodities Futures Trading Commission, 2011. Presentation (Working Group 1, Additional Document). In: Technology Advisor Committee
- U.S. Commodities Futures Trading Commission Technology Advisor Committee, 2001. Market Access Committee Interim Report.
- U.S. Securities and Exchange Commission, 2005. Regulation NMS (Final Rule Release No. 34-51808; June 9, 2005). <<http://www.sec.gov/rules/final/34-51808.pdf>>.
- U.S. Securities and Exchange Commission, 2010. Concept release on equity market structure.
- U.S. Securities and Exchange Commission, 2012. Consolidated Audit Trail (Rule 613, Final Rule).
- Walker, J.S., 2008. *A Primer on Wavelets and Their Scientific Applications*. Chapman and Hall/CRC (Taylor and Francis Group), Boca Raton.
- Zhang, L., 2006. Efficient Estimation of Stochastic Volatility Using Noisy Observations: A Multi-Scale Approach. *Bernoulli* 12, 1019-1043.
- Zhang, L., Mykland, P.A., Aït-Sahalia, Y., 2005. A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data. *Journal of the American Statistical Association* 100, 1394-1411.

Table 1. Sample Summary Statistics

Source: CRSP and Daily TAQ data, April 2011. The sample is 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 four rows), all table entries are cross-firm medians.

	Dollar trading volume quintile					
	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-offer midpoint)	\$15.77	\$4.76	\$5.46	\$17.86	\$27.76	\$51.60
Market capitalization of equity, \$Million	\$690	\$41	\$202	\$747	\$1,502	\$8,739

Table 2. Time scale volatility estimates for US equities in 2011

Estimates of time scale variances and related measures for 150 US firms during April, 2011. The wavelet variances, v_j^2 , are estimates of the price variance at the time scale $\tau_j = 50 \times 2^{j-1}$. The rough variances, σ_j^2 , measure cumulative variation at all time scales $\leq \tau_j$. For presentation, I report the square-roots of the variances, in mills (\$0.001) per share and basis points (0.01%). The wavelet variance ratio is $V_{j,J} = 2^{J-j} v_j^2 / v_J^2$ where $J = 16$ is the longest time-scale in the analysis, and the rough variance ratio is similarly defined as $VR_{j,J} = 2^{J-j} \sigma_j^2 / v_J^2$. For a random-walk, both ratios would be unity at all horizons. The bid-offer correlation is the wavelet correlation (correlation between detail components) at the indicated time scale. All entries are cross-firm means. The National Best Bid and Offer are computed from TAQ data; the bid and offer are separately transformed using the Haar basis; the reported variance estimates are averages of the bid and offer variances. The data are time stamped to the millisecond. Prior to transformation, I take the average of the bid or offer over non-overlapping 50 millisecond intervals. Entries for $j = 0$ are variances within the 50 ms intervals.

Level, j	Time scale	Rough variances			Wavelet variances			Correlation $\rho_{bid,offer,j}$
		(1) σ_j (mils per share)	(2) σ_j (basis pts)	(3) Ratio, $VR_{j,J}$	(4) v_j (mils per share)	(5) v_j (basis pts)	(6) Ratio, $V_{j,J}$	
	< 50 ms	0.29	0.17	4.22				
1	50 ms	0.40	0.23	3.99	0.28	0.16	3.76	0.32
2	100 ms	0.56	0.32	3.79	0.38	0.22	3.59	0.36
3	200 ms	0.77	0.44	3.53	0.53	0.30	3.27	0.41
4	400 ms	1.06	0.61	3.21	0.73	0.42	2.89	0.44
5	800 ms	1.47	0.84	2.90	1.02	0.58	2.60	0.48
6	1,600 ms	2.04	1.17	2.64	1.41	0.81	2.38	0.52
7	3.2 sec	2.84	1.61	2.40	1.97	1.11	2.16	0.55
8	6.4 sec	3.94	2.22	2.12	2.74	1.52	1.84	0.60
9	12.8 sec	5.48	3.04	1.88	3.80	2.08	1.65	0.65
10	25.6 sec	7.61	4.17	1.69	5.27	2.83	1.51	0.70
11	51.2 sec	10.57	5.70	1.54	7.31	3.88	1.39	0.75
12	102.4 sec	14.65	7.80	1.42	10.12	5.30	1.29	0.79
13	3.4 min	20.29	10.67	1.32	13.99	7.25	1.21	0.83
14	6.8 min	28.11	14.61	1.23	19.38	9.93	1.15	0.86
15	13.7 min	38.85	19.98	1.16	26.66	13.53	1.08	0.89
16 (=J)	27.3 min	53.24	27.16	1.08	36.00	18.17	1.00	0.90

Table 3. Time scale volatility estimates for US equities in 2011, across dollar trading volume quintiles.

Estimates of time scale variances and related measures for 150 US firms during April, 2011, for quintiles constructed on dollar trading volume. The wavelet variances, v_j^2 , are estimates of the price variance at the time scale $\tau_j = 50 \times 2^{j-1}$. The rough volatilities, σ_j^2 , measure cumulative variation at all time scales $\leq \tau_j$. For presentation, I report the rough volatilities (square-roots of the rough variances) in mils (\$.001) per share (Panel A) and basis points (0.01%, Panel B). The rough wavelet variance ratio is $VR_{j,J} = 2^{J-j} \sigma_j^2 / v_j^2$ where $J = 16$ is the longest time-scale in the analysis (Panel C). For a random-walk $VR_{j,J}$ would be unity at all horizons. The bid-offer correlation (Panel D) is the wavelet correlation (correlation between bid and offer components) at the indicated time scale. Table entries are cross-firm means with standard errors in parentheses. The National Best Bid and Offer are computed from TAQ data; the bid and offer are separately transformed using the Haar basis; the reported variance estimates are averages of the bid and offer variances. The data are time stamped to the millisecond. Prior to transformation, I take the average of the bid or offer over non-overlapping 50 millisecond intervals. Entries for $j = 0$ are variances within the 50 ms intervals. Transforms are performed through level $J = 16$; for brevity only a subset of time-scales are reported.

Panel A. Rough volatility, σ_j , mils per share

Level, j	Time scale	Full sample	Dollar trading volume quintiles				
			1 (low)	2	3	4	5 (high)
0	< 50 ms	0.29 (0.02)	0.16 (0.02)	0.20 (0.04)	0.30 (0.04)	0.37 (0.05)	0.40 (0.05)
1	50 ms	0.40 (0.03)	0.23 (0.03)	0.27 (0.05)	0.40 (0.05)	0.52 (0.07)	0.57 (0.07)
3	200 ms	0.77 (0.05)	0.43 (0.05)	0.51 (0.10)	0.77 (0.10)	1.00 (0.13)	1.11 (0.13)
5	800 ms	1.47 (0.10)	0.83 (0.10)	0.96 (0.18)	1.45 (0.18)	1.91 (0.25)	2.14 (0.26)
7	3.2 sec	2.84 (0.20)	1.58 (0.19)	1.81 (0.34)	2.76 (0.34)	3.71 (0.49)	4.19 (0.51)
10	25.6 sec	7.61 (0.55)	3.90 (0.47)	4.54 (0.83)	7.06 (0.86)	10.13 (1.36)	12.02 (1.51)
14	6.8 min	28.11 (2.17)	12.77 (1.68)	15.35 (2.74)	24.77 (3.00)	38.36 (5.32)	47.61 (6.23)
16	27.3 min	53.24 (4.19)	22.57 (3.07)	28.16 (4.82)	46.88 (5.76)	74.70 (10.52)	90.49 (11.87)

Table 3. Time scale volatility estimates for US equities in 2011 across dollar trading volume quintiles (continued).

Panel B. Rough volatility, σ_j , basis points

Level, j	Time scale	Full sample	Dollar trading volume quintiles				
			1 (low)	2	3	4	5 (high)
0	< 50 ms	0.17 (0.01)	0.29 (0.02)	0.22 (0.02)	0.15 (0.01)	0.12 (0.01)	0.08 (0.01)
1	50 ms	0.23 (0.01)	0.39 (0.03)	0.30 (0.02)	0.20 (0.02)	0.16 (0.01)	0.11 (0.01)
3	200 ms	0.44 (0.02)	0.75 (0.06)	0.57 (0.05)	0.38 (0.03)	0.31 (0.02)	0.22 (0.02)
5	800 ms	0.84 (0.04)	1.44 (0.11)	1.08 (0.09)	0.74 (0.06)	0.59 (0.03)	0.43 (0.03)
7	3.2 sec	1.61 (0.08)	2.74 (0.20)	2.05 (0.17)	1.41 (0.12)	1.15 (0.06)	0.85 (0.06)
10	25.6 sec	4.17 (0.19)	6.70 (0.48)	5.20 (0.43)	3.65 (0.31)	3.13 (0.18)	2.43 (0.18)
14	6.8 min	14.61 (0.62)	21.37 (1.54)	17.89 (1.55)	12.97 (1.11)	11.92 (0.75)	9.64 (0.75)
16	27.3 min	27.16 (1.13)	37.68 (2.82)	33.26 (2.93)	24.42 (2.02)	23.20 (1.48)	18.38 (1.46)

Panel C. Rough variance ratio, $VR_{j,j}$

Level, j	Time scale	Full sample	Dollar trading volume quintiles				
			1 (low)	2	3	4	5 (high)
0	< 50 ms	4.22 (1.29)	12.86 (7.01)	3.45 (0.43)	2.62 (0.22)	1.76 (0.16)	1.37 (0.05)
1	50 ms	3.99 (1.26)	12.14 (6.85)	3.23 (0.39)	2.44 (0.18)	1.69 (0.15)	1.35 (0.05)
3	200 ms	3.53 (1.07)	10.51 (5.80)	2.83 (0.30)	2.20 (0.15)	1.57 (0.13)	1.30 (0.05)
5	800 ms	2.90 (0.66)	7.90 (3.56)	2.50 (0.23)	2.02 (0.13)	1.43 (0.11)	1.21 (0.04)
7	3.2 sec	2.40 (0.40)	5.92 (2.09)	2.17 (0.17)	1.82 (0.10)	1.32 (0.10)	1.15 (0.04)
10	25.6 sec	1.69 (0.13)	3.07 (0.64)	1.70 (0.12)	1.49 (0.06)	1.19 (0.07)	1.17 (0.04)
14	6.8 min	1.23 (0.03)	1.59 (0.11)	1.24 (0.06)	1.17 (0.03)	1.04 (0.03)	1.16 (0.02)
16	27.3 min	1.08 (0.01)	1.19 (0.03)	1.08 (0.02)	1.06 (0.01)	1.01 (0.01)	1.06 (0.01)

**Table 3. Time scale volatility estimates for US equities in 2011
across dollar trading volume quintiles (continued).**

Panel D. Wavelet bid-offer correlations.

Level, j	Time scale	Full sample	Dollar trading volume quintiles				
			1 (low)	2	3	4	5 (high)
1	50 ms	0.32 (0.02)	0.05 (0.01)	0.23 (0.03)	0.31 (0.03)	0.41 (0.03)	0.56 (0.04)
3	200 ms	0.41 (0.02)	0.11 (0.01)	0.33 (0.03)	0.42 (0.03)	0.49 (0.03)	0.65 (0.03)
5	800 ms	0.48 (0.02)	0.16 (0.02)	0.40 (0.03)	0.51 (0.03)	0.56 (0.03)	0.72 (0.03)
7	3.2 sec	0.55 (0.02)	0.19 (0.02)	0.47 (0.03)	0.59 (0.03)	0.66 (0.03)	0.82 (0.02)
10	25.6 sec	0.70 (0.02)	0.27 (0.03)	0.61 (0.03)	0.75 (0.02)	0.85 (0.02)	0.95 (0.01)
14	6.8 min	0.86 (0.02)	0.44 (0.05)	0.88 (0.02)	0.97 (<0.01)	0.99 (<0.01)	1.00 (<0.01)
16	27.3 min	0.90 (0.02)	0.52 (0.05)	0.96 (0.01)	0.99 (<0.01)	1.00 (<0.01)	1.00 (<0.01)

Table 5. Summary statistics for US equities, 2001-2011

From the CRSP file, for each year, 2001-2011 and all stocks present in January through April of that year with share codes equal to 10 or 11, I draw 150 firms in a random sample stratified by dollar trading volume in January through March. NBB is the National Best Bid; NBO, the National Best Offer; CT, Consolidated Trade; CQ, Consolidated Quote. Trade and quote counts are from the Monthly TAQ database (one-second time stamps). Except for the number of firms, table entries are cross-firm medians.

	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
No. firms	137	122	141	148	144	150	150	147	145	149	149
NYSE	106	46	51	44	48	44	55	53	56	54	47
Amex	16	4	10	12	8	15	14	6	5	14	6
NASDAQ	15	72	80	92	88	91	81	88	84	81	96
Avg. daily CT records (trades)	167	228	231	399	448	605	970	1,217	1,993	1,141	1,346
Avg. daily CQ records (quotes)	1,525	1,053	1,470	3,917	6,004	7,307	12,521	16,791	41,571	23,530	24,053
Avg. daily NBB changes	128	162	210	514	611	761	772	1,183	1,787	1,468	1,225
Avg. daily NBO changes	127	163	226	545	729	751	789	1,142	1,789	1,461	1,146
Avg. price (bid-offer midpoint)	\$20.57	\$20.98	\$14.41	\$16.53	\$16.10	\$21.14	\$15.81	\$14.12	\$11.25	\$16.79	\$15.77
Market capitalization of equity, \$Million	\$976	\$410	\$205	\$352	\$348	\$411	\$480	\$411	\$382	\$490	\$690

Table 6. Wavelet variance ratios for US firms, 2001-2011

In each year 2001-2011, 150 US firms are randomly selected from CRSP (stratified by average daily dollar trading volume during the first quarter of the year). Quote records for April are taken from the NYSE Monthly TAQ database. Within each second, quotes are randomly assigned order-preserving millisecond fractional portions. Wavelet variances, v_j^2 , are estimates of the price variance at the time scale. The wavelet variance ratio is $V_{j,J} = 2^{J-j} v_j^2 / v_J^2$ where $J=16$ is the longest time-scale in the analysis. For random-walk, the ratio would be unity at all horizons. All entries are cross-firm means. The National Best Bid and Offer are computed from TAQ data; the bid and offer are separately transformed using the Haar basis; variance estimates are formed as the average of the bid and offer variances. Estimates in Panel A are constructed from bids and offers that were filtered for errors, but not otherwise adjusted. Estimates in Panel B are constructed from denoised bids and offers (with short-term peaks clipped).

Panel A. Wavelet variance ratios, $V_{j,J=16}$, computed from raw bids and offers

Level, j	Time scale	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
1	50 ms	5.29	7.36	5.96	10.31	6.56	8.57	6.96	6.07	4.53	7.09	4.71
2	100 ms	5.52	6.75	5.20	9.71	6.38	8.07	6.27	5.39	4.12	6.27	4.33
3	200 ms	5.35	6.44	5.05	9.06	6.10	7.34	5.33	4.65	3.68	5.41	3.75
4	400 ms	4.65	5.35	4.92	8.18	5.64	6.30	4.25	3.84	3.21	4.54	3.07
5	800 ms	3.16	4.12	3.86	5.59	4.93	5.10	3.41	3.11	2.76	3.71	2.56
6	1,600 ms	2.13	2.56	3.19	4.11	4.06	4.05	2.89	2.59	2.42	3.04	2.23
7	3.2 sec	2.00	2.25	2.91	3.39	3.42	3.37	2.56	2.28	2.16	2.53	2.01
8	6.4 sec	1.95	2.12	2.61	2.91	2.88	2.92	2.35	2.08	1.94	2.16	1.82

Panel B. Wavelet variance ratios, $V_{j,J=16}$, computed from denoised bids and offers

Level, j	Time scale	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
1	50 ms	1.61	2.38	3.07	7.03	5.95	8.24	6.56	5.84	4.22	6.81	4.47
2	100 ms	1.58	2.33	3.02	6.84	5.76	7.76	5.89	5.18	3.85	6.01	4.08
3	200 ms	1.56	2.28	2.96	6.50	5.49	7.04	4.99	4.46	3.43	5.19	3.57
4	400 ms	1.56	2.24	2.88	5.92	5.05	6.02	3.96	3.68	2.99	4.37	3.01
5	800 ms	1.57	2.20	2.77	5.01	4.37	4.82	3.13	2.98	2.58	3.58	2.52
6	1,600 ms	1.64	2.21	2.68	4.00	3.52	3.79	2.63	2.51	2.28	2.94	2.20
7	3.2 sec	1.82	2.31	2.60	3.45	2.96	3.16	2.33	2.22	2.05	2.46	1.99
8	6.4 sec	2.12	2.54	2.58	3.20	2.60	2.75	2.15	2.04	1.86	2.11	1.82

Table 7. Consolidated quote record for PRK, April 6, 2001.

The table contains the consecutive records from the monthly TAQ consolidated quote file for the Park National Corporation. The first five columns are directly from the CQ file. The national best bid and offer (NBBO), and the exchange(s) at the NBBO are inferred. The NBBO columns contain entries only when there is a change. The size is units of 100 shares. ("4x2" means that 400 shares are bid for and 200 shares are offered.) The exchange codes are "A" (the American Stock Exchange, the primary listing exchange [presently named "NYSE MKT LLC"]); "M," Midwest; "C," Cincinnati; "T," NASDAQ.

Time	Bid	Offer	Size	Ex	Mode	NBB	NBO	Time	Bid	Offer	Size	Ex	Mode	NBB	NBO
12:01:33	86.73	86.90	4x2	A	12	86.73	A 86.90 A	12:03:37	83.75	86.96	1x1	T	12		
12:01:34	86.63	87.00	1x1	M	12			12:03:38	86.50	86.90	2x2	A	12	86.50	A
12:01:35	86.35	87.28	1x1	C	12			12:03:39	86.40	87.00	1x1	M	12		
12:01:35	86.67	86.96	1x1	T	12			12:03:40	86.50	86.90	2x2	A	12		
12:01:35	86.67	86.96	1x1	T	12			12:03:40	86.12	87.28	1x1	C	12		
12:02:22	86.74	86.90	3x2	A	12	86.74	A	12:03:45	86.44	86.96	1x1	T	12		
12:02:23	86.64	87.00	1x1	M	12			12:03:45	86.44	86.96	1x1	T	12		
12:02:25	86.68	86.96	1x1	T	12			12:03:46	86.50	88.00	2x9	A	12		86.96 T
12:02:25	86.68	86.96	1x1	T	12			12:03:48	86.40	88.10	1x1	M	12		
12:03:11	83.63	86.90	1x2	A	12	86.68	T	12:03:49	86.12	88.38	1x1	C	12		
12:03:13	83.53	87.00	1x1	M	12			12:03:51	86.44	88.06	1x1	T	12		88.00 A
12:03:15	83.25	87.28	1x1	C	12			12:03:51	86.44	88.06	1x1	T	12		
12:03:15	83.60	86.90	2x2	A	12			12:03:52	86.50	86.90	2x2	A	12		86.90 A
12:03:16	83.57	86.96	1x1	T	12	83.60	A	12:03:54	86.40	87.00	1x1	M	12		
12:03:16	83.57	86.96	1x1	T	12			12:03:55	86.12	87.28	1x1	C	12		
12:03:16	83.50	87.00	1x1	M	12			12:03:58	83.00	86.90	2x2	A	12	86.44	T
12:03:21	83.54	86.96	1x1	T	12			12:03:58	86.44	86.96	1x1	T	12		
12:03:21	83.54	86.96	1x1	T	12			12:03:58	86.44	86.96	1x1	T	12		
12:03:27	83.81	86.90	1x2	A	12	83.81	A	12:04:00	82.90	87.00	1x1	M	12		
12:03:29	83.71	87.00	1x1	M	12			12:04:01	82.62	87.28	1x1	C	12		
12:03:30	83.43	87.28	1x1	C	12			12:04:01	82.94	86.96	1x1	T	12	83.00	A
12:03:30	83.81	86.90	1x2	A	12			12:04:01	82.94	86.96	1x1	T	12		
12:03:32	83.75	86.96	1x1	T	12			12:04:06	86.50	86.90	2x2	A	12	86.50	A

Table 8. Time scale volatility estimates for US equities, 2001-2011

In each year 2001-2011, 150 US firms are randomly selected from CRSP (stratified by average daily dollar trading volume during the first quarter of the year). Quote records for April are taken from the NYSE Monthly TAQ database. Within each second, quotes are randomly assigned order-preserving millisecond fractional portions. The rough volatilities, σ_j , measure cumulative variation at all time scales $\leq \tau_j = 50 \times 2^{j-1}$ ms. These are reported in mils (\$0.001) per share (Panel A) and basis points (0.01%, Panel B). The rough wavelet variance ratio (Panel C) is $VR_{j,J} = 2^{J-j} \sigma_j^2 / v_j^2$ where $J=16$ is the longest time-scale in the analysis. (For a random-walk $VR_{j,J}$ would be unity at all horizons.) Table entries are cross-firm means; standard errors clustered within firm are given in parentheses. The National Best Bid and Offer are computed from TAQ data; the bid and offer are separately transformed using the Haar basis; the reported estimates are averages across bid and offer sides. Transforms are performed through level $J = 16$; for brevity only a subset of time-scales are reported. All estimates are constructed from denoised bids and offers (with short-term peaks clipped).

Panel A. Rough volatility, σ_j mils per share

Level, j	Time scale	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
1	50 ms	0.39 (0.02)	0.35 (0.02)	0.27 (0.02)	0.39 (0.02)	0.36 (0.02)	0.41 (0.02)	0.29 (0.02)	0.38 (0.04)	0.44 (0.03)	0.46 (0.06)	0.31 (0.02)
3	200 ms	1.02 (0.06)	0.90 (0.04)	0.72 (0.04)	1.01 (0.05)	0.93 (0.04)	1.06 (0.06)	0.73 (0.06)	0.94 (0.10)	1.09 (0.06)	1.14 (0.15)	0.76 (0.05)
5	800 ms	2.12 (0.13)	1.85 (0.09)	1.47 (0.09)	2.02 (0.10)	1.85 (0.08)	2.08 (0.12)	1.41 (0.11)	1.80 (0.17)	2.13 (0.13)	2.20 (0.29)	1.49 (0.10)
7	3.2 sec	4.33 (0.26)	3.68 (0.17)	2.86 (0.17)	3.77 (0.20)	3.44 (0.15)	3.83 (0.20)	2.63 (0.21)	3.36 (0.25)	4.05 (0.25)	4.02 (0.52)	2.84 (0.19)
10	25.6 sec	12.77 (0.78)	10.19 (0.54)	7.52 (0.44)	9.64 (0.55)	8.78 (0.39)	9.73 (0.52)	6.92 (0.56)	8.72 (0.53)	10.66 (0.70)	9.96 (1.17)	7.60 (0.54)
14	6.8 min	46.19 (2.72)	35.17 (1.67)	25.32 (1.42)	32.53 (1.96)	30.36 (1.47)	33.85 (1.92)	24.86 (2.00)	31.28 (1.86)	38.82 (2.61)	34.44 (3.47)	28.11 (2.17)
16	27.3 min	85.30 (5.31)	64.14 (3.11)	45.16 (2.57)	58.27 (3.62)	54.15 (2.68)	60.97 (3.59)	45.57 (3.52)	58.42 (3.56)	72.26 (4.93)	63.12 (6.05)	53.23 (4.18)

Table 8. Time scale volatility estimates for US equities, 2001-2011 (continued)Panel B. Rough volatility, σ_j , basis points

Level, j	Time scale	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
1	50 ms	0.18 (0.01)	0.20 (0.01)	0.20 (0.01)	0.28 (0.02)	0.28 (0.02)	0.21 (0.01)	0.18 (0.01)	0.28 (0.02)	0.36 (0.02)	0.24 (0.02)	0.18 (0.01)
3	200 ms	0.48 (0.02)	0.52 (0.03)	0.52 (0.03)	0.73 (0.05)	0.73 (0.06)	0.52 (0.03)	0.45 (0.03)	0.68 (0.05)	0.88 (0.05)	0.58 (0.04)	0.44 (0.02)
5	800 ms	1.01 (0.05)	1.07 (0.05)	1.07 (0.05)	1.46 (0.10)	1.45 (0.12)	1.02 (0.05)	0.84 (0.06)	1.29 (0.09)	1.70 (0.09)	1.09 (0.06)	0.86 (0.04)
7	3.2 sec	2.06 (0.09)	2.13 (0.11)	2.08 (0.10)	2.65 (0.16)	2.61 (0.20)	1.86 (0.08)	1.53 (0.09)	2.37 (0.14)	3.14 (0.15)	1.94 (0.10)	1.62 (0.08)
10	25.6 sec	6.01 (0.26)	5.81 (0.29)	5.42 (0.24)	6.43 (0.34)	6.34 (0.43)	4.60 (0.18)	3.88 (0.20)	6.00 (0.30)	7.94 (0.35)	4.69 (0.21)	4.19 (0.20)
14	6.8 min	21.68 (0.96)	20.23 (1.04)	18.05 (0.77)	20.63 (0.98)	20.50 (1.28)	15.27 (0.52)	13.19 (0.59)	20.15 (0.76)	27.66 (1.12)	15.82 (0.61)	14.64 (0.63)
16	27.3 min	39.55 (1.83)	36.64 (1.87)	31.82 (1.34)	35.95 (1.62)	35.79 (2.22)	27.19 (0.92)	23.84 (1.05)	36.67 (1.21)	50.49 (1.93)	28.75 (1.06)	27.18 (1.13)

Panel C. Rough variance ratio, $VR_{j,J}$

Level, j	Time scale	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011
1	50 ms	1.61 (0.07)	2.38 (0.17)	3.07 (0.17)	7.03 (1.42)	5.95 (0.83)	8.24 (1.94)	6.56 (1.37)	5.84 (1.40)	4.22 (0.44)	6.81 (1.05)	4.47 (1.43)
3	200 ms	1.57 (0.07)	2.31 (0.16)	3.00 (0.16)	6.67 (1.33)	5.63 (0.78)	7.42 (1.68)	5.47 (1.06)	4.86 (1.18)	3.66 (0.36)	5.66 (0.86)	3.85 (1.17)
5	800 ms	1.57 (0.07)	2.23 (0.15)	2.85 (0.15)	5.62 (1.02)	4.83 (0.63)	5.72 (1.16)	3.87 (0.61)	3.59 (0.73)	2.93 (0.25)	4.25 (0.59)	2.94 (0.69)
7	3.2 sec	1.72 (0.11)	2.27 (0.19)	2.68 (0.14)	4.12 (0.57)	3.56 (0.38)	3.94 (0.62)	2.78 (0.33)	2.63 (0.34)	2.32 (0.16)	3.02 (0.32)	2.28 (0.36)
10	25.6 sec	2.38 (0.37)	2.74 (0.62)	2.59 (0.30)	3.16 (0.58)	2.38 (0.16)	2.53 (0.25)	2.03 (0.16)	1.89 (0.13)	1.74 (0.08)	1.93 (0.12)	1.70 (0.13)
14	6.8 min	1.38 (0.04)	1.41 (0.06)	1.49 (0.04)	1.58 (0.07)	1.47 (0.04)	1.52 (0.06)	1.37 (0.04)	1.29 (0.03)	1.30 (0.03)	1.30 (0.03)	1.23 (0.03)
16	27.3 min	1.12 (0.01)	1.13 (0.02)	1.16 (0.01)	1.18 (0.02)	1.15 (0.01)	1.17 (0.02)	1.12 (0.01)	1.09 (0.01)	1.10 (0.01)	1.10 (0.01)	1.08 (0.01)

Figure 1. The bid and offer for AEPI, April 29, 2011

National best bid and offer (NBBO) from the NYSE Daily TAQ dataset. The National best bid (bottom line, in blue) is the maximum bid, taken over all market centers reporting to the Consolidated Tape Association; the National best offer (top line, red) is the minimum offer.

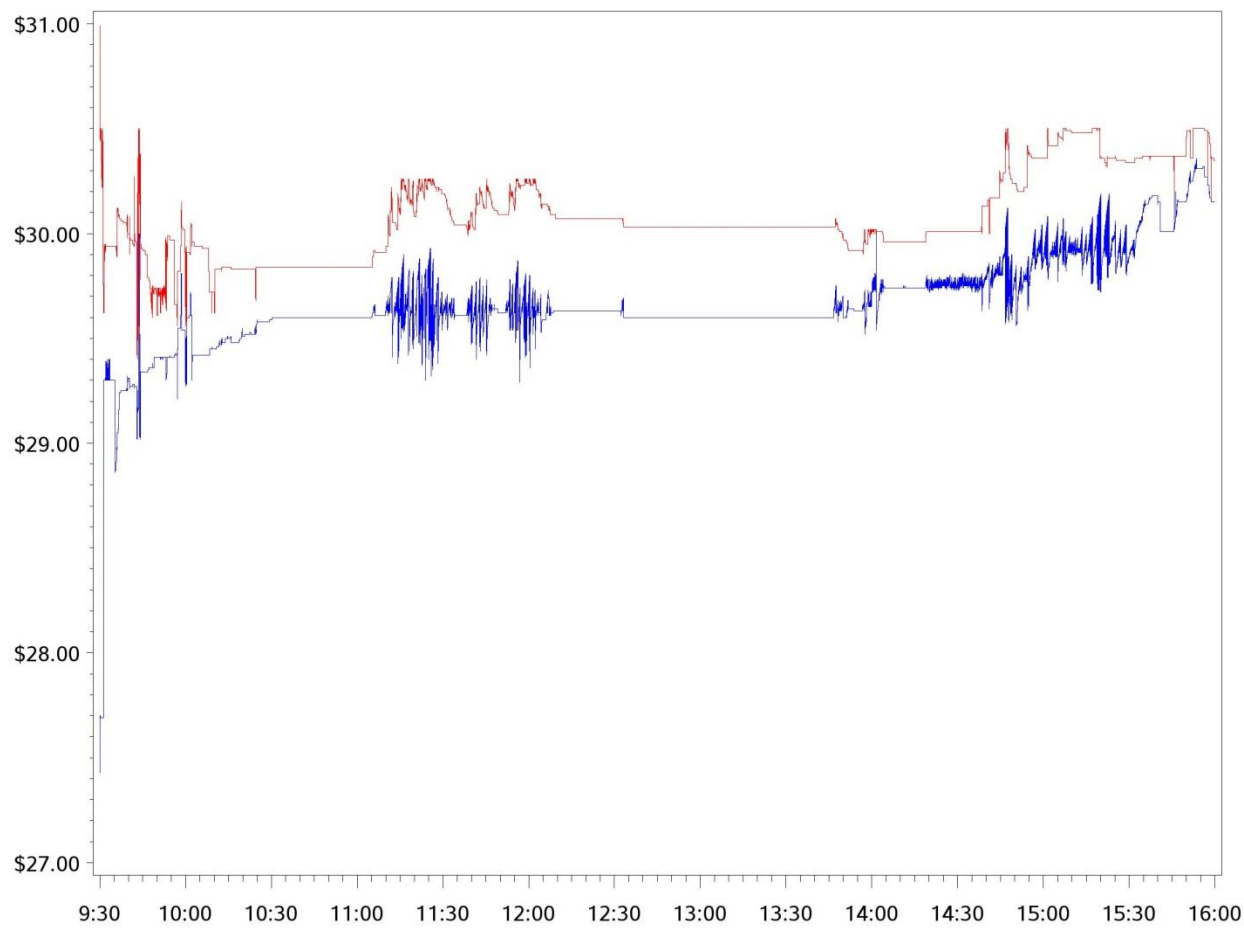


Figure 2. Variance ratios for US equities in 2011

Estimates of time scale variance ratios for 150 US firms during April, 2011. The wavelet variances, v_j^2 , are estimates of the price variance at the time scale $\tau_j = 50 \times 2^{j-1}$. The wavelet variance ratio is $V_{j,J} = 2^{J-j} v_j^2 / v_J^2$ where $J = 16$ is the longest time-scale in the analysis. For random-walk, both ratios would be unity at all horizons. Plotted points are means (across firms) of estimated variance ratios in quintiles constructed on dollar trading volume. The National Best Bid and Offer are computed from TAQ data; the bid and offer are separately transformed using the Haar basis; the reported variance estimates are averages of the bid and offer variances. The data are time stamped to the millisecond. Prior to transformation, I take the average of the bid or offer over non-overlapping 50 millisecond intervals.

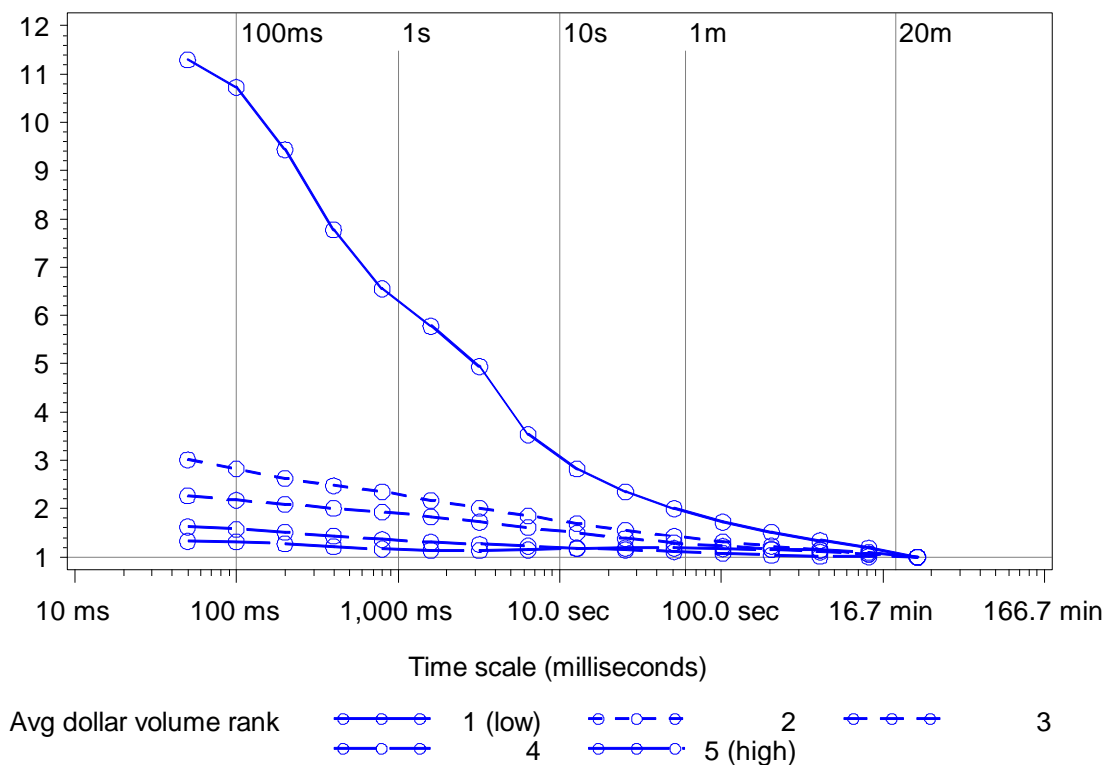


Figure 3. Wavelet correlations between the National Best Bid and National Best Offer

Estimates of bid-offer wavelet correlations for 150 US firms during April, 2011. The wavelet correlation between the bid and offer at level j (and time scale $\tau_j = 50 \times 2^{j-1}$) is defined as $\rho_{bid,offer,j} = v_{bid,offer,j}^2 / \sqrt{v_{bid,j}^2 v_{offer,j}^2}$ where $v_{bid,j}^2$, $v_{offer,j}^2$ and $v_{bid,offer,j}^2$ denote the bid wavelet variance, the offer wavelet variance, and the bid-offer wavelet covariance. Plotted points are means (across firms) of estimated correlations in quintiles constructed on dollar trading volume. The National Best Bid and Offer are computed from TAQ data; the bid and offer are separately transformed using the Haar basis; the reported variance estimates are averages of the bid and offer variances. The data are time stamped to the millisecond. Prior to transformation, I take the average of the bid or offer over non-overlapping 50 millisecond intervals. The sample is 150 randomly chosen U.S. stocks, over April, 2011.





Figure 4. Bid and offer for PRK (Park National Corporation) on April 6, 2001.

Panel A. National best bid and offer (NBBO) from the NYSE Daily TAQ dataset. The National best bid (bottom line, in blue) is the maximum bid, taken over all market centers reporting to the Consolidated Tape Association; the National best offer (top line, red) is the minimum offer.



Panel B. Rough component of the National Best bid, constructed from a Haar wavelet transform and comprising components at time scales of 51.2 seconds and lower. The bands demarcate $\pm\$0.33$, approximately 150% of the average bid-ask spread for the day.


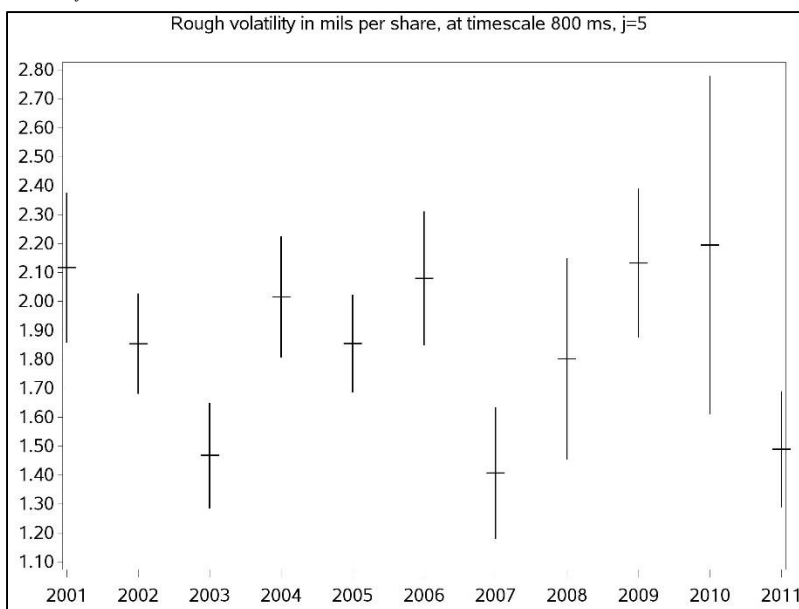


Figure 5. Time scale volatility estimates for U.S. Equities, 2001-2011

Quote volatility at a time scale of 800 ms., in mils per share (Panel A), basis points (Panel B), and as a variance ratio (Panel C). In each year, the horizontal tick marks the mean and the vertical line demarcates the mean \pm twice the standard error (clustered on firm).

Panel A. Rough volatility, $\sigma_{j=5}$ mils per share



Panel B Rough volatility, $\sigma_{j=5}$ basis points per share

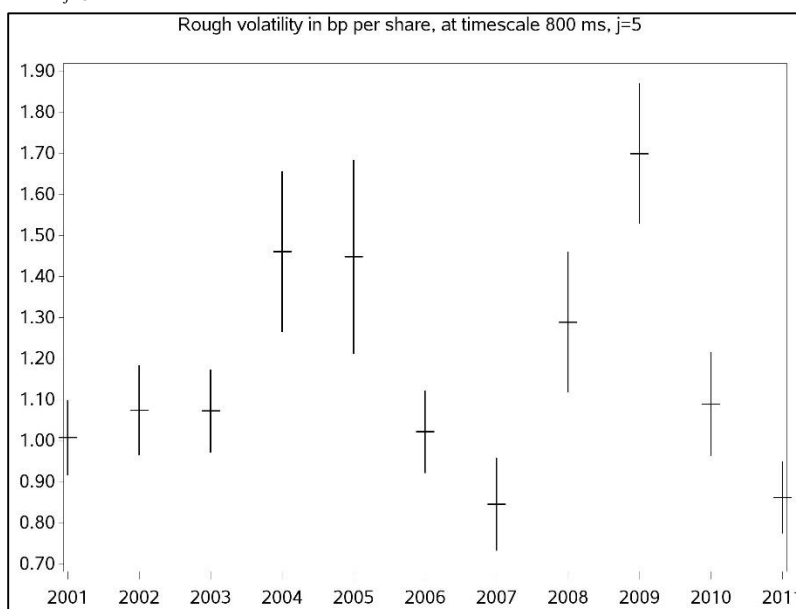


Figure 6. Time scale volatility estimates for U.S. Equities, 2001-2011 (continued)Panel C. Rough variance ratio, $VR_{j=5,J}$ 