Current draft: August 11, 2015
First draft: May 14, 2012

Public Information and Stale Limit Orders: The Evidence

Christopher G. Lamoureux
and
Qin Wang[1]

Removing all potentially stale limit orders from the data used to estimate a mapping from trades to price movements has almost no effect on that operator in a wide cross section of stocks on the New York Stock Exchange in 1991 (a specialist market). Unpegged limit orders are a significant source of liquidity in this market. Since stale limit orders result from public information shocks, we infer that public information shocks have a small effect on individual stock return volatility. We quantify the boundary of this effect–in terms of the size and number of such shocks–using a calibration analysis.

Keywords: *Stale limit orders, public information*

JEL Classification Code: *G14*

[1]Department of Finance, The University of Arizona, Eller College of Management, Tucson, 85721, 520–621–7488, lamoureu@lamfin.eller.arizona.edu; and College of Business, University of Michigan, Dearborn, MI, 48128, 313–583–6487, qinw@umich.edu. The current version of this paper can be downloaded from http://finance.eller.arizona.edu/lam/rsch.html. While retaining the responsibility for all errors, we would like to thank Scott Cederburg, Eric Kelley, Juhani Linnainmaa, and Marc Lipson for their input.

1. Introduction

Isolating the reasons that speculative prices move is a fundamental yet elusive goal in finance research. In the spirit of French and Roll (1986) and Koudijs (2014), we use an institutional feature of a specific market to detect public information's footprint in stock price dynamics. In large part the literature either finds or assumes that public information can and does move prices without the occurrence of transactions, whereas private information and price pressures move prices via trading activity. However, in a traditional specialist market[1] such as the New York Stock Exchange prior to decimalization, public information shocks result in transactions as limit orders that are rendered stale by the shock are picked off by nimble floor traders.

This means that executed stale limit orders contain information about the size of public information shocks. In the data we can identify limit orders that lose money after specific intervals. However we cannot discern whether the reason for this loss is: 1) that the order was rendered stale by the release of public information; or 2) that the order was subject to adverse selection (i.e., the counterparty had private information at the time of the trade); or 3) that the order was simply unlucky (for example a public or private information shock moved the price in the opposite direction of the order within the measurement interval). Accordingly, we classify all limit order trades that lose money after an interval as *potentially stale*. Our empirical design relies on a cross-section of stocks which should integrate over the "luck" dimension of our alternative hypothesis. Furthermore because specialist quotes reflect public information by definition, only the limit order book bears the risk of being picked off after a public information shock. By contrast both the specialist and the limit order book bear adverse selection risk from privately informed traders. We can put an upper bound on the importance of public information shocks in stock return volatility by evaluating the effect that removing all potentially stale limit order trades from the data has on an operator that evaluates the relationship between trades and price movements. We find that public information is not an important driver of individual stock return volatility in our sample.

Prior literature provides evidence of the importance of the picking off of stale limit orders following a public information shock in price formation. Masulis and Shivakumar (2002) show that stocks on Nasdaq adjust to news of a seasoned equity offering one hour faster than stocks listed on the New York Stock Exchange (NYSE) and American Stock Stock Exchange (AMEX). They analyze these markets in the period January 1990 - December 1992, which includes the

[1] Important features of such a market include: a non-trivial tick size; a relatively clean mapping from orders to transactions (i.e., minimal order splitting and bundling); competition between a centrally-positioned specialist and the limit order book for liquidity provision; and the limit order book providing significant liquidity. Perhaps the most important trait of such a market is the fact that observed transactions are a reasonable proxy for (generally unobserved) orders. For the inability to map orders into transactions in modern financial markets, see GAO (2005) and Easley, López de Prado, and O'Hara (2012, esp. p.1466).

period of this study. They argue that the reason for the difference in the speed of adjustment to public information is that the NYSE and AMEX had (unpegged) public limit order books which the Nasdaq did not at that time. They infer that market orders trading with stale limit orders delay the price adjustment on NYSE and AMEX. Since Nasdaq did not have a public limit order book the price adjustment to public information there was instantaneous.

Berkman (1996) shows that option market limit orders on the European Options Exchange in Amsterdam are commonly picked off after the underlying stock's price changes. Here professional traders are able to exploit these limit orders as the individuals who placed them cannot revise them quickly in the face of new public information. Linnainmaa (2010) looks at Finnish data, and finds that the documented poor performance of retail traders may be attributed to their use of limit orders. In the Helsinki double-auction market, retail traders do not make poor investment decisions, so much as their use of limit orders results in bagging losses.

We calibrate the effect of stale limit orders to show that relatively small public information shocks would leave a statistically identifiable footprint in the data. To this end we estimate a statistic from the empirical microstructure literature that links price change to transaction activity, developed by Glosten and Harris (1988). Our interest in such a measure is not whether it is viable[2] –rather how sensitive it is to the removal of stale limit order transactions. That is, our formal tests do not use this measure per se, instead they compare the measure estimated on the full sample with the measure estimated on the sample from which potentially stale limit order trades have been removed. Since this is a new use for the measure, we design a calibrated bootstrap to show the effect of removing public information shocks on the measure. We show that an alternative popular measure of private information in a specialist market, PIN, would not work for our purposes (since it does not analyze the effect of transactions on price).

We next evaluate the effect of removing all potentially stale limit orders from the sample on this statistic. If this results in a statistically significant drop, we would not be able to infer that public information shocks are the reason since we do not know why the limit orders lost money. On the other hand, if this does not affect our measure significantly, then we fail to reject the (joint) null hypothesis that public information is *not* an important driver of individual stock returns; (and that the limit order book is not subject to more adverse selection risk than the specialist, and that the limit order book is not significantly unlucky). Since we have to entertain the multiple alternative hypotheses, we do not construct a contrapositive price sequence with a (single) jump in place of the limit order sequence. Our tests involve simply removing the potentially stale limit orders from the data.

[2]Lamoureux and Wang (2015) provide an empirical assessment of the efficacy of popular measures of private information in a specialist market. Collin-Dufresne and Fos (2015) provide evidence that suggests that measures that relate price reaction to transactions can not detect the presence of insiders (their analysis transcends various market structures).

Figure 1 demonstrates the intuition behind our test. This figure shows all trades on Boeing stock on January 28, 1991. Boeing announced its quarterly earnings and provided revenue guidance at 13:10. In the 20 minutes following this announcement the price dropped from \$48.75 to \$47, and the stock closed down \$2.25 on the day. The inset shows that many of the trades during the 11 minutes following the announcement involved limit orders. By contrast, if there were no limit order book, or if limit orders were contingent on public information shocks, this price move might have occurred with no intervening transactions.

In the event, 42,500 shares changed hands as the price ratcheted down during the period highlighted in the inset; the limit order book bought 19,500 shares (46% of the volume), and sold 900 shares (2% of the volume), while the specialist bought 4,000 shares (7% of the volume). All of the buy-side limit orders in these transactions were placed at least 53 minutes before the public information release. For example, consider the fifth transaction in the inset, where 3,100 shares trade at \$48.50, at 13:14. This trade combines two limit buy orders at \$48.50–one for 1,200 shares which was placed at 10:03 that morning, the other for 5,000 shares, placed at 10:52 that morning. The first order was filled and the second received 1,900 shares. The sell side of this trade batches two market orders placed at 13:14–one for 3,000 shares and the other for 100 shares. This pattern is typical of all transactions involving limit buy orders highlighted in the inset. Most of these orders were placed prior to 11:00 on the morning of the announcement. The only exception is a limit buy order for 500 shares at \$48.125, placed at 12:17 that day–still well before Boeing's 13:10 public announcement. Since these limit orders are not contingent on the public information release, and they were placed well before the announcement we consider them to be potentially stale.

Over the entire day Boeing's specialist's net trade was +25,000 shares. Marking this position to market at \$46.8125, the spread midpoint at close, the specialist nominally lost \$34,225 on the day.[3] Over the day, the limit order book's net trade was -22,500 shares. For the day the limit order book made a net profit of \$49,531.

Several features of this example are worth noting. First, the specialist is charged with maintaining an orderly market, and is rewarded for price continuity. On this date, all but three of the 389 price changes from one transaction to the next are one tick (i.e., \$0.125) or zero. However, following the public information release the specialist is able to achieve price continuity without putting his own capital at risk by matching incoming market orders to sell to stale limit orders to buy. Also (unlike in this example), in this study we do *not* attempt to

[3] We measure specialist trading profit over an interval by first netting all round-trip transactions at actual transaction prices, and then by marking the net position to market at the spread midpoint at the end of the interval. This is an incomplete measure of the specialist's actual return for several reasons. First, we do not know whether the specialist hedged any of the exposure. Second, the specialist is charged with and rewarded for maintaining an orderly market.

identify information shocks exogenously. Instead we rely on the price dynamics after limit order trades to identify those limit order trades that lose money, since this will include all limit orders that are picked off. This is important since it is possible that the limit order book prior to an anticipated public announcement, such as an earnings call, is not representative. Our empirical design integrates over all types of public information shocks. The cost of this is that we have only an upper bound on transactions involving stale limit orders. Foucault, Röell, and Sandås (2003) and Parlour and Seppi (2008) discuss the trade-off between monitoring and picking off risk.[4] However, continuous monitoring would not have protected limit orders on the NYSE of the early 1990's from public information shocks because public traders were institutionally disadvantaged relative to floor traders.

We use three post-trade lag lengths to measure a trade's profitability: five minutes, one hour, and one day, and report all of our results using all three. Our results are robust to the choice of lag length. This look-ahead to identify the trade type is standard in the microstructure literature and does not introduce any bias or endogeneity since the goal of the empirical exercise is to ascertain whether the presence of stale limit order trades in the data itself affects the empirical measure relating price change to the order flow.

Our calibration shows in the case of Boeing, whose price averages \$45 in our data, that if 8% of the daily stock return variance over the 63 trading days were due to small public information shocks of \$0.50 each, then our test would reject the null hypothesis that public information is not an important driver of the return variance. Removing the stale limit order trades induced by these public information shocks, comprising 0.6% of Boeing's transactions over the period, would result in a statistically significant drop in the measure. In the actual data, we find that after removing all potentially stale limit orders (comprising 38% of all limit order transactions–almost 4% of all transactions) from Boeing the measure rises–as it does for eight of the 19 most actively-traded stocks in our sample. This increase is never statistically significant, and it highlights the fact that the measure includes sampling error.

Our data are all transactions, quotes, and orders (enabling us to identify transactions with the limit order book) on individual NYSE stocks from November 1990 through January 1991. In this market limit orders are not contingent on public information, so when a new piece of public information hits the market nimble floor traders pick off the standing limit orders rendered stale by this event. This is the bagging or picking off risk identified by Copeland

[4]Foucault, Hombert, and Roşu (2014) model trading as a function of whether an informed trader can react to news faster than liquidity providers.

and Galai (1983).[5] In a traditional specialist market the only protection that a limit order has against this picking off risk is for the specialist to declare a market halt.

Our test's power is derived from the relative importance of the limit order book in the trading process. Bae, Jang, and Park (2003) characterize traders' choices between market and limit orders in the TORQ data. They show traders submit roughly the same number of market orders and limit orders in the data. They also find that limit orders are more attractive the larger the bid-ask spread, the larger the order size, and the higher is expected transitory volatility.[6] Harris and Hasbrouck (1996); Bae, Jang, and Park; and Masulis and Shivakumar (2002) provide evidence on the importance of the limit order book. For our average stock 18% of the trading volume involves the limit order book–roughly the same participation rate as the specialist. We classify 40% of limit order volume as potentially stale.

Harris and Hasbrouck (1996) and Handa and Schwartz (1996) look at the relative performance of alternative limit order strategies and market orders on the NYSE. Harris and Hasbrouck use the TORQ database and consider the profitability of actual limit orders. Handa and Schwartz use Institute for the Studies of Securities Markets transaction data on the Dow Jones Industrial Average 30 stocks in 1988, to consider contrapositive limit order strategies. Our results are fully consistent with Harris and Hasbrouck's (1996) finding that in the specialist market that we consider widely used limit order strategies are profitable relative to market orders. Similarly, Handa and Schwartz find that the bagging costs of placing limit orders do not overwhelm the trading benefits.[7]

The implication of these studies is that in a traditional specialist market of the 20th century, traders could use limit orders to improve execution and indeed this would seem to apply to a majority of limit orders. On the other hand, limit orders could also harm traders as they are not conditioned on public information. The effects of a centralized (unpegged) limit order book, and incumbent trades at stale quotes, are not unimportant in price formation. We are motivated by these facts to use potentially stale limit orders to isolate public information shocks in a sequence of transactions.

Our analysis complements two strands of the empirical finance literature seeking to isolate

[5]Modern electronic networks afford investors the opportunity to submit a wide array of pegged limit orders. For example, Direct Edge began allowing its members to post *route peg orders* on September 7, 2012. The customer specifies a quantity and price that is inside the best bid and offer (BBO). The order will remain hidden and execute against a suitably-sized routable order at the BBO quote. In general, *pegged* limit orders are not subject to picking off risk.

[6]A recent survey of the theory behind the choice between limit and market orders is Roşu (2012).

[7]Lo, Mackinlay, and Zhang (2002) warn against the use of first-passage times to evaluate counterfactual limit order executions (as in Handa and Schwarz 1996). Lo, Mackinlay, and Zhang demonstrate that this approach does not produce reliable estimates of such orders' actual execution times.

the effect of information on speculative price dynamics. French and Roll (1986) and Barclay, Litzenberger, and Warner (1990) compare stock return volatility in trading and non-trading hours. They obtain identification by using business days when the stock market is closed and Saturdays when the stock market is open, respectively. Since return variances are much higher during trading periods they infer that the incorporation of public information into prices (without trading) is insignificant in price formation. Koudijs (2014) attempts to measure information flows directly. He uses geographically separated markets (stocks in British companies that trade in Amsterdam), in a technologically constrained setting where he can measure the arrival of news (as boats arrive to the Netherlands from England). He finds that this information does move prices but that there is significant volatility in the absence of any (public) information. He infers that at most 25 - 50% of overall return variances can be explained by public information shocks.

By focusing on the dynamics of price and trading volume, our empirical design avoids having to use non-market-based measures of information. For example, since we only evaluate return dynamics and trading volume during trading hours we allow for the possibility that trading itself is a form of public information. Recent theoretical models such as Kondor (2012) stress the importance of higher-order expectations for trading activity. Furthermore, public information may increase disagreements in higher-order expectations. This is important because, as they discuss, French and Roll's (1986) variance ratio tests have no power against the hypothesis that trading itself is a form of public information.

Another problem that plagues analyses of price dynamics during business days when the market is closed is that such events are rare, and they are neither surprises nor exogenous. Everyone knew well in advance that the NYSE would be closed on the 16 Wednesdays in the second half of 1968, which French and Roll use to control for differences in public information releases on business and non-business days. Economic activity on these days might well be different from what it is on Wednesdays when the market is open. Finally, Andersen, Bollerslev, an Das (2001) provide evidence on the unreliability of variance-ratio tests, such as those in French and Roll (1986) and Barclay, Litzenberger, and Warner (1990). We employ a very different empirical design and specifically calibrate it to stock return variances.

The second strand of literature attempts to measure information directly. Roll (1984) exploits the special nature of orange juice futures contracts. He finds no linkage between public information releases and price volatility in this market. However, Boudoukh, Richardson, Shen, and Whitelaw (2007) and Fleming, Kirby, and Ostdiek (2006) argue that most of the information that Roll considers is not material. When they isolate what they consider to be price-relevant news–such as temperature reports around the freezing mark–they identify a strong link between measurable information shocks and price volatility. Similarly, Fleming and Remolona (1999) analyze the reactions of US Treasury securities to public information releases. They find that prices react quickly and with virtually no trading to public macroeconomic announcements.[8] These announcements are scheduled and occur at regular intervals.

The first generation (i.e., pre-textual analytics) of papers attempting to measure information in the stock market includes Mitchell and Mulherin (1994) and Berry and Howe (1994). These papers count the number of stories in specific news outlets (Dow Jones and Reuters, respectively) and evaluate whether this is related to return volatility. Both of these papers find that while trading volume is positively correlated with their measure of the amount of information, return volatility is not. Vega (2006) uses an indicator of whether a stock appeared in at least one headline or first paragraph of a story to proxy for a public information release on that day. Engle, Hansen, and Lund (2011) examine 28 individual stocks from 2001 through mid 2009, using the Dow Jones Intelligent Indexing System, which they argue gives a much broader sampling of media exposure than the sources used in earlier studies. They find that unexpected (public) information arrival accounts for 5 - 20% of these stocks' idiosyncratic variances, which in turn account for 58 to 77% of their total return variances. Roll (1988) considers the r^2 from market model regressions on days with and without measurable information releases and finds no evidence that information shocks increase (idiosyncratic) return volatility. A new generation of papers that uses machine-intensive textual analysis to measure information more accurately is reaching a different conclusion. For example, Boudoukh, Feldman, Kogan, and Richardson (2012) find a significant drop in factor model r^2 values of S&P 500 companies on days with identified relevant news, in the first decade of the 21st century.

So the evidence is mixed, with the most recent research suggesting that public information shocks do have a material effect on return volatility. However, only this paper, French and Roll (1986) and Koudijs (2014) have enough institutional restrictions to isolate the effects of public information from private information on return dynamics. The results of Boudoukh, Feldman, Kogan, and Richardson (2012), for example, are consistent with a model in which privately informed traders trade more aggressively on news days–using the increased volume engendered by the news as camouflage.[9] Of course, Roll (1984) and Fleming and Remolona (1999) consider securities which should have limited private information. Fleming and Remolona find that after Treasury securities' market prices adjust quickly, and with little or no volume to the public information, there is a period of heightened trading activity and relatively wide bid-ask spreads. This is consistent with heterogeneous interpretation of the news, in the sense of Harris and Raviv (1993). This activity, which Fleming and Remolona describe as the second stage

[8] As described by Fleming and Remolona (1999) the institutional details of the Treasury market are different from the traditional specialist market. For example, there are no limit order book or price continuity restrictions in the Treasury market.

[9] This is not meant as a criticism of this literature, since its objective is to identify the link between information and return volatility. By contrast our analysis isolates *public* information.

adjustment process will not affect our empirical design since it does not put the limit order book at an informational disadvantage. However it does militate against directly inferring that price movement on days with public information releases is the result of the public information, per se.

The next section describes our data. In Section 3 we show the importance of the limit order book as a source of liquidity in the specialist market, and the extent of stale limit orders. Section 4 contains the paper's results. Here we show the properties of the empirical measure on the data, calibrate a test to this data, and then carry out that test by removing all potentially stale limit order trades from the data. Section 5 concludes the paper.

2. The data

As Hasbrouck (1992) explains the TORQ data cover 144 NYSE stocks in the three-month period, November 1990 through January 1991.[10] We discard seven stocks for insufficient data, so our sample comprises 137 stocks over 63 trading days. Hasbrouck, Sofianos, and Sosebee (1993) provide a detailed description of the NYSE in the late 1980s and early 1990s, which we consider a traditional specialist market. Our analysis requires identifying trades involving the limit order book, which means that a trades-and-quotes database such as TAQ is insufficient for our purposes. We use the Lee and Ready (1991) algorithm to classify trades as buyer- or seller-initiated. As such we include all transactions on the NYSE, including those involving non-electronic orders. Importantly, given their role in protecting limit orders, there are no trading halts in the TORQ data.

There are several advantages to the TORQ data for our study even though there now exist broader-based sources of order-level information, such as the NASDAQ ITCH database. Decimalization and market fragmentation have given rise to entirely new strategies involving limit orders, which are now often used to scope out the presence of latent liquidity (as shown in Hasbrouck and Saar 2009, 2010). Furthermore, as we note above in footnote 3, modern markets allow traders to peg limit orders so that the institutional structure that allows us to isolate public information no longer exists. Nevertheless, we analyze and reconsider how inference was and is made in microstructure studies, which in turn influences how we think about issues dealing with current data sets and related policy questions.

[10] Lamoureux and Wang (2015) include a recent survey of the many published papers that use the TORQ database.

3. Limit Order Results

We report our results along two dimensions. First we isolate the 19 individual stocks in the database with more than 4,500 transactions in the 63-day period. Second we report portfolio aggregates for all 137 stocks in our sample, and three size-sorted portfolios. Sample summary statistics are provided in Tables I and II. Table I shows that the 19 individual stocks include the largest on the NYSE from that era. However, the number-of-trades hurdle also admits two smaller stocks, AMD and WBN, both of whose market capitalization at the beginning of the period is less than $320 million. WBN is an outlier in this group in terms of the role that the limit order book plays in providing liquidity. Excluding this stock, the mean limit order book participation across the 18 stocks is 9.3%, and this rate for all 18 falls between 5 and 17%. In stark contrast, over 47% of the trading volume in WBN over the period involves the limit order book. Adding WBN to the sample increases the standard deviation of limit order book participation from 2.8% to 9.1%. AT&T is the most actively traded stock in the sample, with 39,216 transactions–an average of 622.5 per day. Nevertheless, AT&T is not an outlier as there are four stocks with more than 33,600 transactions over the 63 days. There are no stocks whose number of transactions falls between 18,500 and 33,600. The mean number of transactions for these 19 stocks is 15,314 and the median is 10,836.

Table II provides summary statistics for the size-ranked portfolios and overall sample. The large and medium terciles contain 46 stocks and the small tercile contains 45 stocks. Panel A shows that the relative importance of the limit order book as a source of liquidity is decreasing in size. The (equally-weighted) average limit order book participation rate in the largest tercile is 9.6%–close to that of the 18 most actively-traded stocks excluding WBN. This average rate increases to 23.4% for the small tercile. Overall, for the average stock, the limit order book is involved in 18% of trading volume, which is slightly less than the specialist (Lamoureux and Wang 2015 show that for the average stock in this sample, the specialist is involved in 23% of trading activity).

We measure each transaction's profitability by multiplying its size by the difference between the quote midpoint after a lag and the trade price. We use three alternative post-trade intervals: five minutes, one hour, and one (trading) day. We compute the profitability of each trade and sum these to measure profitability at the stock level. These are equally-weighted to obtain the reported sample statistics. The medians for the 19 active firms in Table I (entire sample in Table II, Panel A) are: 27.2 (17.5), 34.0 (21.7), and 37.4% (29.7%), for the three lag lengths respectively. In general, limit orders experience higher loss rates for larger firms. Also, the loss rate tends to increase in the lag length used to benchmark the transaction price. Since these loss rates are under 50%, limit order trades tend to be profitable. This would follow if traders posting limit orders are risk-averse liquidity providers as in Grossman and Miller (1988), and/or if limit orders are used by informed traders, as suggested by Collin-Dufresne and Fos (2015); Kaniel and Liu (2006); Anand, Chakravarty, and Martell (2005); and Cornell and Sirri (1992).[11] Table II, Panel B shows that limit order profitability is positively skewed, most profitable, on average, for larger firms, and declining in the time lag used to measure profitability. As with participation rates, these profitability measures are of the same magnitude as the specialist profits reported in Lamoureux and Wang (2015).[12]

4. The empirical measure g

4.1 All transactions

Glosten and Harris (1988) consider the following projection of the transaction price change from trade $t-1$ to t, ΔP_t, on the trade direction, X_t (X_t is 1 for a market buy order and -1 for a market sell order), signed volume, X_tV_t, and the changes in these two variables from the preceding trade:

$$\Delta P_t = b_0 \Delta X_t + b_1 \Delta X_t V_t + b_2 X_t + b_3 X_t V_t + \epsilon_t \tag{1}$$

Their measure of the proportion of the bid-ask spread due to adverse selection is:

$$g = \frac{2(b_2 + b_3\overline{V})}{2(b_0 + b_1\overline{V}) + 2(b_2 + b_3\overline{V})}, \tag{2}$$

where $\overline{V}$ is the mean transaction volume for the stock in the sample. An advantage of this statistic for our purposes is its simplicity. The denominator reflects the notion that the price change from one transaction to the next may be due to bid-ask bounce and price pressure–both of which may be a function of the trade size.

While Tables I and II report the total number of trades in the sample, many of these trades cannot be used in estimating g. We have to be able to classify a trade and the preceding trade as either buyer-initiated or seller-initiated. These additional requirements mean that we have to discard between 20 - 30% of all trades to compute g. For example, of the 18,442 (6,088) trades in Boeing (Colgate-Palmolive), we use 12,755 (4,748) in computing g.

The intuition behind g is that if a price change is induced by a transaction, controlling for the bid-ask bounce, then it is the result of private information, as the providers of liquidity learn from the order flow. Since this specification is built atop the foundations of the seminal microstructure models of Kyle (1985) and Glosten and Milgrom (1985), it assumes that standing

[11] In a study of algorithmic (high frequency) traders on the Deutsche Boerse in January, 2008, Hendershott and Riordan (2011) show that such traders supply liquidity when it is expensive and demand liquidity when it is cheap. Perhaps the profitable liquidity traders in 1991 are the forebears of modern algorithmic traders.

[12] For example, Lamoureux and Wang (2015) report that the average (median) specialist profit for the same set of stocks over the same period, is \$19,193 (\$5,600), with the one-hour lag used to measure profitability.

bids and offers reflect the prevailing *public* information set. In the spirit of Glosten and Harris (1988), following a public information shock, stale limit orders can give rise to a unidirectional sequence of trades that mimics the effect of a private information shock.[13]

Table III provides properties of g for the 19 most actively-traded stocks in the TORQ database. Since g is a non-linear function of OLS parameter estimates it has a small-sample bias. We use the bootstrap to gauge the nature of the small sample bias, and to construct the sampling distribution of g for each of these 19 stocks. The bootstrap is conducted as follows. First the sample OLS estimators are obtained and used as the basis of the simulation for each stock. Second, we sample from the model residuals, the ϵ_t from Equation (1), with replacement, to construct a pseudo-sample of the same size as the original sample. Third, we then estimate g from this pseudo-sample. This is repeated 100,000 times (the process requires over 10 days of CPU time on a 2.93 GHz Intel Xeon processor). Table III shows that the median small sample bias is 0.3% of the median estimator, and the bias is positive for 16 of the 19 stocks.

Table IV reports summary statistics for the cross-section of g for the entire sample and the three size-based portfolios. The table highlights that g (being a moment estimator) is not constrained to lie between zero and one (its theoretical support). There are two small stocks whose g estimator is less than 0, and these are truncated to 0. One medium stock and one small stock have estimates of g that exceed one, these are truncated to unity.

4.2 The power of the test in the return volatility metric

In addition to using standard inference to test our compound null hypothesis, we assess the power of this test in the metric of return volatility for each of the nineteen stocks examined in depth. Before bringing our test to the transactions data, in this section we demonstrate the intuition behind it, which is suggested by Figure 1. Consider that a negative public information release moves Boeing's price down by 20 ticks or \$2.50. In a stylized example of unpegged limit orders this might engender 20 transactions–each at a price that is one tick below the preceding transaction price. Whereas in the absence of unpegged limit orders the price adjustment would occur without any transactions. In this section we ascertain how many events of different sizes are needed to move g by a statistically significant amount, and then measure the effect of these events on the stock's daily return variance. Calibrated Monte Carlo is used to answer this question. For a given information size (20 ticks in the preceding example), we insert from 1 to 30 such shocks into the transaction series for a stock. These shocks are equally-spaced through

[13] As noted in the introduction, for our purposes g is simply a summary statistic. Glosten and Harris' (1988) assumptions about the determinants of the bid-ask spread may be violated in the data, so that g may not correctly measure the proportion of the bid-ask spread due to adverse selection. Our interest in g is simply to understand how it is affected by the ratcheting of the transactions price through stale limit orders–which we calibrate below, and how our estimate of g is affected by removing potentially stale limit order trades from the sample.

the sample. Each shock is drawn with a 50% probability of being positive and negative. The volume of each transaction is drawn from a normal distribution with mean equal to the sample average transaction volume for the stock and standard deviation of 15% of average volume. For each simulated draw, we compute g and the daily return variance (which assumes that the shocks' effects on the price are permanent). We take 10,000 draws for each number/size of shock pair.

Table V provides an example of the output from one of these simulations. This table reports summary statistics for Boeing with zero through 30 shocks of 10 ticks each. Note that g increases monotonically as the number of shocks ranges from 0 to 7. After seven shocks non-monotonicities arise because we do not randomize over the placement of the shocks (they are inserted into the trade sequence so that they occur at equal intervals through the 63 days). The table also shows the standard deviation of g across the Monte Carlo draws. Referring to Table III, the 95% critical value for Boeing–adjusted for small sample bias–is 10.98%. Table V shows that if there are six or more 10-tick shocks the $E(g)$ is above this value. So suppose that there were six public information shocks of magnitude \$1.25 (or larger), during the three months in Boeing's stock. Our estimated g for Boeing would be 11.06, and the 5% critical value (adjusted for small sample bias) would be 9.184. Then when we remove the stale limit order transactions induced by these shocks, our estimate of g falls to 9.136.[14] Since this is below the 5% critical value we would detect a statistically significant drop in Boeing's g. In this instance a minimum of 60 transactions–less than half of one-percent of Boeing's 12,815 transactions, and less than one-eighth of the number of potentially stale limit orders in the data–is sufficient to trigger a statistical rejection of the null hypothesis that public information is not a significant source of return variation.

Table VI shows that the daily standard deviation of Boeing's stock is 2.344%. In addition to $E(g)$, Table V also shows the effect of incrementally adding stylized public information shocks on Boeing's daily return variance. This is strictly monotonic in the number of such shocks. It shows that adding the six blocks of ten trades increases the daily return variance by 13.38% (i.e., the standard deviation increases from 2.344% to 2.496%). The analysis underlying Table V is repeated for Boeing using shocks comprising four ticks (i.e., \$0.50) and 20 ticks (\$2.50), and for all 19 heavily-traded stocks, for all three information sizes.

The results from this exercise are reported in Table VI. This table is designed to provide an envelope to give a sense of the nature of public information that our test would detect. In

[14]While at first blush it might appear that the assumption of sequential limit orders is either unrealistic or biased toward rejection, recall that in estimating g, both the lead and lag trades must be classified and valid (so for example, we do not use the close-to-open in estimating g). So if in fact there were cases where 10 limit orders fell between non-limit order trades, we would remove more than 10 trades from the sequence. So assessing the power of the test by assuming that the limit orders occur consecutively is the most conservative option.

general fewer shocks of larger magnitude are needed to trigger a significant change. However such perturbations have a larger effect on the daily return variance than a larger number of smaller magnitude shocks. Consider IBM as an example. There are 23,997 transactions used to compute IBM's g in the TORQ data. Recall from Table I that IBM is the highest priced stock of the 19, with a minimum price in the sample of $104\frac{7}{8}$. The standard deviation of IBM's daily returns is 1.38%. When we estimate g on IBM data we obtain 17.24%. From Table III, we see that for IBM, g's bootstrapped 95% confidence level is 18.69%, and that IBM's g has a small sample bias of 0.35%. So IBM's bias-adjusted 95% critical value is 19.04%. Table VI shows that eight 20-tick shocks are required to increase the estimated g above this value. Tying this in to the original test and empirical design means that had we started with 24,157 transactions (including these eight 20-tick runs), we would have obtained a base g estimate of 19.25%, (the mean from the simulated samples in this case). Next we would remove the 160 potentially stale limit order transactions. We would estimate g on this sample and the estimate would be 17.24%. Our bootstrap analysis would then indicate that the potentially stale limit orders significantly biased the original g upward. While this would be potentially important for researchers taking g at face value, for our purposes it is uninteresting in light of our composite alternative hypothesis: we cannot be sure whether the deleted trades were simply cases where the limit order book was the victim of *private* information. The 160 transactions in this case increase IBM's daily stock return standard deviation by 15% (from 1.38% to 1.59%).

Continuing with IBM, if we restrict attention to 4-tick public information shocks, then 37 are needed to move g enough to reject the null hypothesis of no change (an increase in the number of transactions of 0.6%). In this case, the daily return standard deviation is only 3% higher than in the original data. On average across the 19 stocks the scale of four-tick public information shocks required to trigger a significant effect is 24% of the daily return variance (or a 7.3% increase in the daily standard deviation). Similarly our tests would not be able to reject the null hypothesis that public information does not materially contribute to return variances if such shocks were of magnitude 10 ticks, and they comprise less than 23% of the mean (19% of the median) stock's return standard deviations. Finally, if public information shocks were of magnitude 20 ticks, they would have to comprise more than 49% of the return standard deviation for the mean (38% for the median) stock to trigger a rejection. Table VI also shows that AMD is an outlier in this exercise, which is because its stock price is so low (recall from Table I that it ranges from $3\frac{3}{4}$ to $7\frac{3}{8}$), and the minimum tick size is still $\frac{1}{8}$.[15]

Table VI also shows that this test lacks power for those stocks with low return variance. For example, public information shocks could account for 80% of the daily return variance of

[15] A consequence of AMD's low price is that we cannot assign 50% probabilities to positive and negative shocks. In order to maintain a positive stock price for AMD, in the calibrated simulations, we set the probability of a positive shock to 66, 85, and 98% for 4-, 10-, and 20-tick shocks, respectively.

AMD–regardless of the size/frequency of such shocks–and we would not be able to reject the null hypothesis. Table VI also shows that public information shocks that move the price by \$2.50 (i.e., comprise 20 ticks) also deplete the test's power. For all of the individual stocks we would not detect importance of public information in this case unless it accounts for more than 38% of return volatility, and for the average stock, returns are not volatile enough to ever reject the null in this case.

Because Easley O'Hara and Paperman's (1996) PIN is widely used as a measure of private information in traditional specialist markets, we explore the use of PIN instead of g in our empirical design in the appendix. PIN uses only the number of buy and sell transactions on each date in the sample, so all stocks have 126 observations. As a result of the reduced amount of information used, PIN has no statistical power in our empirical design. In fact, for stocks with large numbers of transactions in the context of our 63-day sample, no augmented order imbalance is sufficient to move PIN beyond confidence levels that reflect that estimate's sampling error.

4.2 The effects of removing potentially stale limit orders

Having calibrated the test in the preceding section, we now consider in detail the effects of removing the potentially stale limit order trades from the data on g. Consider Boeing, which has 12,755 transactions in the sample used to estimate g. Boeing's average transaction volume is 2,711 shares, and Table III shows that its sample g estimate is 9.1%. The small sample bias in Boeing's g statistic is 0.01%. The table also shows that the sample standard deviation is 1.1%, and that the 5% (bootstrapped) critical value for g–adjusted for small sample bias– is 7.26%.

Table III also shows that after we remove all potentially stale limit orders identified at the 5-minute lag from the time series of Boeing's transactions data, and re-estimate g, we obtain an estimate of 9.8%. The test for whether public information shocks are a material part of Boeing's trading dynamics is whether this value is less than the bootstrapped 5% critical value-adjusted for small sample bias, shown in the column headed *Critical Value*, (which as noted, is) 7.26% in this instance. Not only is it not, g estimated after removing all potentially stale limit orders is actually larger than the initial value of g estimated on the entire sample. The qualitative inference is that public information releases during trading hours lack the size and frequency to leave a statistically detectable footprint in the trade sequence. Table III also shows that this result for Boeing is qualitatively identical to the cases where one hour and one day are used as the lag lengths to identify potentially stale limit order transactions. In all three cases removing potentially stale limit orders results in a higher estimate of g than is obtained on the full sample of all transactions.

We can use the analysis from the preceding section to infer the maximum scale of public information in terms of the return standard deviation. For Boeing, it must be that small public information shocks account for less than 7.2% of the return variance; medium-sized public information shocks account for less than 13.4% of the return variance; and large public information shocks account for less than 37.1% of the return variance. This is a description of an envelope: a combination of small and medium public information shocks accounting for 10% of return variance would also trigger a rejection.

Table III also shows that Boeing is typical of the 19 most active stocks. For this set of stocks the average (median) value of g estimated from all transactions is 11.23% (7.24%). When all limit orders that are potentially stale after five minutes are removed from the data, the average and median values are respectively 11.23% and 7.57%. Furthermore, as Table III shows, the average g *increases* after removing all potentially stale limit order transactions at the one-hour and one-day lags. At the individual stock level, the table shows that removing all potentially stale limit order trades at the 5-minute lag results in a drop in g for 11 of the 19 stocks. However, in no case is the effect of removing potentially stale limit order trades statistically significant. That is, for none of the individual stocks is it the case that removing all potentially stale limit order trades has an effect on g that is larger than random sampling error in that statistic estimated from the entire sample.

As with Boeing, the effects of removing all potentially stale limit order trades identified at the one-hour and one-day horizons on the mean and median are qualitatively identical to the inference when the five-minute lag is used to classify the profitability of limit order transactions. Again, using all three lag lengths, it is never the case that g estimated after the removal of all potentially stale limit order trades is less than the 5% critical value constructed for that stock's g.

Rather than relying on the bootstrap at the individual stock level, the inference in Table IV uses the cross-sectional distribution of g–overall and in the three size-based portfolios–to test for the importance of public information in price formation. For the entire sample, the average g is 23.25%, with a standard error of 1.6%. When we eliminate all potentially stale limit order transactions identified with a five-minute lag from all stocks' trade sequences, the mean g is 22.66. The $t-$statistic of the test of the null hypothesis that this mean equals the mean g using all trades is -0.37. Similarly, when we eliminate all potentially stale limit order transactions identified with a one-hour lag from all stocks' trade sequences, the mean g is 22.91%. The $t-$statistic of the test of the null hypothesis that this mean equals the mean g obtained using all trades is -0.21. Finally, when we eliminate all potentially stale limit order transactions identified with a one-day lag from all stocks' trade sequences, the mean g increases to 23.55%. We infer that the ratcheting of public information shocks through stale limit orders does not significantly bias g downward.

We also fail to reject the joint null hypothesis that includes the notion that public information shocks are too small (in number and magnitude) to leave a statistically significant footprint in realized returns for the three size-based portfolios. When we use either a five-minute or one-hour lag to measure the profitability of limit order trades, the mean g decreases for all three size-sorted portfolios. None of these decreases are statistically significant. Furthermore, as with the entire sample, the average g increases for each size tercile after removing all trades involving potentially stale limit orders classified using a one-day lag.

The final information contained in Table IV is the number of stocks in each size tercile where the estimated g is significantly lower after all potentially stale limit order trades are removed from the sample, for each of the three classification lags. When five minutes and one hour lags are used to classify the limit orders, we reject the null hypothesis at the 5% level for 4.4% of our sample firms. Four of the six rejections are firms in the smallest tercile, and the remaining two rejections are medium-sized firms. Thus confirming the earlier tests that there is no *statistical* evidence that public information releases during trading hours contribute significantly to individual stock return variances. This is true for actively traded individual stocks as well as at the aggregate level.

In light of the statistical rejection, we can refer back to Table VI, to evaluate what this means in terms of return variances. For our median active stock, it must be that small public information shocks account for less than 13.1% of the return variance; medium-sized public information shocks account for less than 42.4% of the return variance; and large public information shocks account for less than 89.4% of the return variance.

5. Conclusion

In this paper we demonstrate that (hypothetical) transactions involving stale limit orders resulting from public information shocks that account for as little as 5 to 6% of a stock's daily return variance, leave a statistically visible footprint in transactions data. We look for this footprint in the data, and can not find it. We therefore infer that public information shocks are not an important source of stock return volatility, and we calibrate what this means quantitatively. Similarly since this footprint is not in the data, medium and large public information shocks would have to comprise less than 19% and 37% of total return variance.

Our results indicate that *public* information shocks account for less than one-fifth of the return variance of individual stocks. Furthermore we add trading activity itself to the set of public information that French and Roll (1986) argue does not drive volatility. As with French and Roll (1986) and Koudijs (2014) our ability to isolate public information is specific to our particular setting. So for example, our empirical approach is not applicable to US equity markets post-decimalization, which are fragmented, no longer have specialists, and in which we cannot classify orders by observing transactions. Furthermore, the use of limit orders itself is likely mechanism-specific. The generality of our results depends on the extent to which information is exogenous to the trading process. Koudijs suggests that the role of information in trading is largely exogenous. Indeed our results are consistent with those of French and Roll, Koudijs, and Engle Hansen, and Lund (2011), even though we evaluate different markets and use very different empirical designs.

Finally, our results inform empirical analyses of private information in the microstructure literature. We demonstrate that stale limit order transactions have the potential to significantly bias the Glosten and Harris (1988) measure. We find that this is not a concern in the data however. We also evaluate the use of Easley, Kiefer, O'Hara and Paperman's (1996) PIN as an alternative to the Glosten and Harris measure in our empirical design. Our results show that PIN's clean design and straightforward implementation come with a cost: information is discarded that makes the measure virtually impervious to stale limit orders. Because PIN uses only the number of buy and sell trades on each date, we find that for active stocks it has no power to identify public information shocks from stale limit orders using the 63 days in our sample.

Appendix. Using PIN instead of the Glosten and Harris measure

PIN is a very popular measure of asymmetric information in traditional specialist markets, and has been used in a wide variety of applications. Following Easley, Kiefer, O'Hara, and Paperman (1996) as refined by Easley, Hvidkjaer, and O'Hara (2002), the likelihood function characterizes the sequence of daily buy and sell transactions as coming from one process if there was no information shock for the day and another process (with one-sided trading) if there was an information shock for the day. In their model, at the beginning of each day there is either an information shock with probability α or not, with probability $1-\alpha$. If there is an information shock, the probability that it is good (bad) news is δ $(1-\delta)$. On each day liquidity sell orders are drawn from one Poisson distribution, with parameter ϵ_s; liquidity buy orders are drawn from another Poisson distribution, with parameter ϵ_b; and if an information event occurred, informed orders are drawn from a third Poisson distribution with parameter μ. The likelihood function thus has five parameters, and is a function of the number of buy and sell orders on each day in the sample. Under the model, PIN, the probability of informed trading, is available as a function of the model parameters: $\frac{\alpha\cdot\mu}{\alpha\cdot\mu+\epsilon_b+\epsilon_s}$.

Unlike g, PIN does not rely on the time-series of trades and quotes, it is a function of daily order imbalances. As such, it relies heavily on a clean mapping from orders to transactions, which is generally a feature of a traditional specialist market. In light of its popularity and the fact that it uses a different type of data, we consider how it works in our test design. Easley, O'Hara, and Saar (2001, p. 34) note that PIN would be biased if informed traders use limit orders. They cite the relevant theory (e.g., Glosten 1994) that shows that in such models informed traders would eschew limit orders, and specifically "assume that informed traders use only market orders."

Conceptually then, we could use PIN instead of g to examine whether public information shocks are an important driver of return variances. However the mapping from the nature of stale limit order trades to the test shows that the use of PIN instead of g would render our test powerless. Consider the case of Phillip Morris, one of the most active stocks in the sample, with an average of more than 590 trades per day. PIN estimated on all data is 9.23%, with a bias-adjusted 95%ile of 16.9%. As in Section 3.B above, we ask how many one-sided transactions sequences of varying magnitudes have to be added to the original data to cause PIN to exceed this critical value. In this case, this is an impossible task. For example if we add 120 buys or 120 sells (randomized as to which and when) to each of 30 days, PIN increases to values between 14.3 and 16.1%. The maximum likelihood parameter estimates in the base case (trade-augmented case with PIN of 16.1%) are: α: 0.365 (0.610); δ: 0.227 (0.349); μ: 128.7 (149.2); ϵ_b: 207.2 (208.6); ϵ_s: 254.8 (264.6). If we add 120 sell orders to each of 30 days, PIN is 13.1%. In this case the MLE parameter estimates are: α: 0.560; δ: 0.568; μ: 132.3; ϵ_b: 211.5;

ϵ_s: 278.6. This example clarifies why PIN has no power for our test. The only data used to estimate the model are the numbers of buyer-initiated and seller-initiated transactions on each of 63 days. As we introduce an order imbalance the model parameters adapt. For example, we see how ϵ_s increases as we move from the base case to the randomized sign case to the added sells-only case.

Turning to Boeing, the PIN on the full sample is 11.02%, and the bias-adjusted critical value is 13.94%. If we add 50 buy orders on each of 30 days, PIN becomes 11.96%. If we add 50 same-directional trades of random sign to 30 days, PIN is 14.64%. If we use 40 trades instead of 50 in the previous case, PIN becomes 13.93%. So for Boeing, over 1,200 stale limit order transactions are needed to move PIN sufficiently. Recall that when we use g for this test, 80 additional transactions are sufficient to move g enough to reject the null. In contrast to PIN, g is estimated with 12,755 values of the price change and signed volume, so the sequential trades resulting from a public information shock ratcheting transaction prices through stale limit order quotes convey more information than just the consequent number of one-sided transactions used in estimating PIN.

References

Anand, Amber, Sugato Chakravarty, and Terrence Martell, 2005, Empirical evidence on the evolution of liquidity: Choice of market versus limit orders by informed and uninformed traders, *Journal of Financial Markets* 8, 289–309.

Andersen, Torben G., Tim Bollerslev, and Ashish Das, 2001, Variance ratio statistics and high-frequency data: Testing for changes in intraday volatility patterns, *Journal of Finance* 56, 305–327.

Bae, Ke-Hong, Hasang Jang, and Kyung Suh Park, 2003, Traders' choice between limit and market orders: Evidence from NYSE stocks, *Journal of Financial Markets* 6, 517–538.

Barclay, Michael J., Robert H. Litzenberger, and Jerold B. Warner, 1990, Private information, trading volume, and stock-return variances, *Review of Financial Studies* 3, 233-253.

Berkman, H., 1996, Large option trades, market makers and limit orders, *Review of Financial Studies* 9, 977-1002.

Berry, Thomas D. and Keith M. Howe, 1994, Public information arrival, *Journal of Finance* 39, 1331–1346.

Boudoukh, Jacob, Ronen Feldman, Shimon Kogan, and Matthew Richardson, 2012, Which news moves stock prices? A textual analysis, Working Paper, NYU.

Boudoukh, Jacob, Matthew Richardson, YuQing (Jeff) Shen, and Robert F. Whitelaw, 2007, Do asset prices reflect fundamentals? Freshly squeezed evidence from the OJ market. *Journal of Financial Economics* 83, 397–412.

Collin-Dufresne, Pierre and Vyacheslav Fos, 2015, Do prices reveal the presence of informed trading? *Journal of Finance* 70, 1555–1582.

Copeland, Tom and Dan Galai, 1983, Information effects on the bid-ask spreads, *Journal of Finance* 38, 1457–1469.

Cornell, Bradford and Erik Sirri, 1992, The reaction of investors and stock prices to insider trading, *Journal of Finance* 47, 1031–1059.

Easley, David, Soeren Hvidkjaer, and Maureen O'Hara, 2002, Is information risk a determinant of asset returns? *Journal of Finance* 57, 2185–2221.

Easley, David, Nicholas Kiefer, Maureen O'Hara, and Joseph Paperman, 1996, Liquidity, information, and infrequently traded stocks, *Journal of Finance* 51, 1405–1436.

Easley, David, Marcos M. López de Prado, and Maureen O'Hara, 2012, Flow toxicity and liquidity in a high-frequency world, *Review of Financial Studies* 25, 1457–1493.

Easley, David, Maureen O'Hara, and Gideon Saar, 2001, How stock splits affect trading: A microstructure approach, *Journal of Financial and Quantitative Analysis* 36, 25–51.

Engle, Robert F., Martin Hansen, and Asger Lund, 2011, And now the rest of the news: Volatility and firm specific news arrival, Working Paper, Aarhus School of Business and Social Sciences.

Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 2006, Information, trading, and volatility: Evidence from weather-sensitive markets, *Journal of Finance* 61, 2899–2930.

Fleming, Michael J. and Eli M. Remolona, 1999, Price formation and liquidity in the U.S. Treasury market: The response to public information, *Journal of Finance* 54, 1901–1915.

Foucault, Thierry, Johan Hombert, Ioanid Roşu, 2014, News trading and speed, Working Paper, HEC Paris.

Foucault, Thierry, Alisa Röell, and Patrik Sandås, 2003, Market making with costly monitoring: An analysis of the SOES controversy. *Review of Financial Studies* 16, 345–384.

French, Kenneth R. and Richard Roll, 1986, Stock return variances: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 5–26.

Glosten, Lawrence R., 1994, Is the electronic open limit order book inevitable? *Journal of Finance* 49, 1127–1161.

Glosten, Lawrence R. and Lawrence Harris, 1988, Estimating the components of the bid/ask spread, *Journal of Financial Economics* 14, 123–142.

Glosten, Lawrence R. and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.

Government Accountability Office, 2005, Decimal pricing has contributed to lower trading costs and a more challenging trading environment, *GAO Report 2005*, available online at: http://www.gao.gov/new.items/d05535.pdf.

Grossman, Sanford J. and Merton H. Miller, 1988, Liquidity and market structure, *Journal of Finance* 43, 617–633.

Handa, Puneet and Robert A. Schwartz, 1996, Limit order trading, *Journal of Finance* 51, 1835–1861.

Harris, Lawrence and Joel Hasbrouck, 1996, Market vs. limit orders: The superDOT evidence on order submission strategy, *Journal of Financial and Quantitative Analysis* 31, 213–231.

Harris, Milton and Artur Raviv, 1993, Differences of opinion make a horse race, *Review of Financial Studies* 6, 473–506.

Hasbrouck, Joel, 1992, Using the TORQ database, Working paper, New York University.

Hasbrouck, Joel and Gideon Saar, 2009, Technology and liquidity provision: The blurring of traditional definitions, *Journal of Financial Markets* 12, 143–172.

Hasbrouck, Joel and Gideon Saar, 2010, Low-latency trading, Working paper, New York University.

Hasbrouck, Joel, George Sofianos, and Deborah Sosebee, 1993, New York Stock Exchange systems and trading procedures, NYSE Working Paper #93-01.

Hendershott, Terrence and Ryan Riordan, 2011, Algorithmic trading and information, Working Paper, University of California, Berkeley.

Kaniel, Ron and Hong Liu, 2006, So what orders do informed traders use? *Journal of Business* 79, 1867–1913.

Kondor, Péter, 2012, The more we know about the fundamental, the less we agree on the price, *Review of Economic Studies* 79, 1175–1207.

Koudijs, Peter, 2014, The boats that did not sail: Asset price volatility in a natural experiment, *Journal of Finance* forthcoming.

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

Lamoureux, Christopher G. and Qin Wang, 2015, Measuring private information in a specialist market, *Journal of Empirical Finance* 30, 92–119.

Lee, Charles M.C. and Mark J. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733–746.

Linnainmaa, Juhani, 2010, Do limit orders alter inferences about investor performance and behavior? *Journal of Finance* 65, 1473–1506.

Lo, Andrew W., A. Craig MacKinlay, and June Zhang, 2002, Econometric models of limit-order executions, *Journal of Financial Economics* 65, 31–71.

Masulis, Ronald W. and Lakshmanan Shivakumar, 2002, Does market structure affect the immediacy of stock price responses to news? *Journal of Financial and Quantitative Analysis* 37, 617–647.

Mitchell, Mark L. and J. Harold Mulherin, 1994, The impact of public information on the stock market, *Journal of Finance* 49, 923–950.

Parlour, Christine A. and Duane J. Seppi, 2008, Limit order markets: A survey, *Handbook of Financial Intermediation and Banking*, Anjan V. Thakor and Arnoud W.A. Boot, editors, Elsevier.

Roll, Richard, 1984, Orange juice and weather, *American Economic Review* 74, 861–880.

Roll, Richard, 1988, R^2, *Journal of Finance* 43, 541–566.

Roşu, Ioanid, 2012, Order choice and information in limit order markets, *Market Microstructure: Confronting Many Viewpoints*, Frédéric Abergel, Jean-Philippe Bouchaud, Thierry Foucault, Charles-Albert Lehalle, and Matthieu Rosenbaum, editors, John Wiley & Sons.

Vega, Clara, 2006, Stock price reaction to public and private information, *Journal of Financial Economics* 82, 103-133.

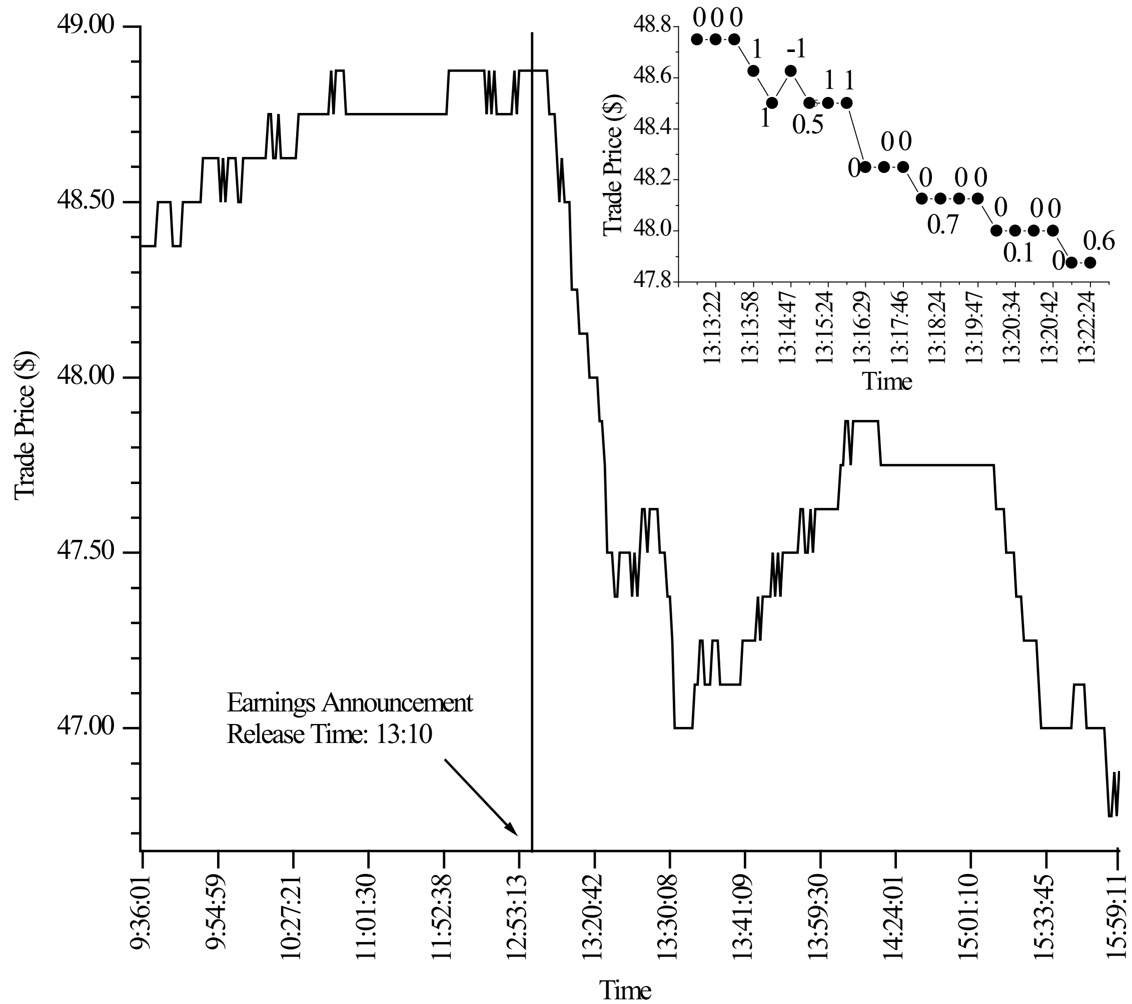

Figure 1. Price ratcheting through stale limit orders after a public information shock

This graph shows all trades in Boeing stock from 9:36:01 to 15:59:11 on January 28, 1991. At 13:10 Boeing posted a four-fold increase in earnings for the fourth quarter, but the stock price fell by $2.25 on the day. The inset highlights the 22 trades (each trade denoted by a dot), between 13:13 and 13:22. The net participation rate by the limit order book in each trade is indicated. For example if the trade was for 2,000 shares and the limit order book purchased 1,500 and sold 500, we record 0.5. There were 389 transactions in Boeing stock on this day, with total trading volume of 1,158,100 shares. Total volume during the latter 19 transactions (i.e., the downward ratchet) in the inset was 42,500 shares, of which the limit order book bought 19,500 (46%) and sold 900 (2%).

Table I
19 Active Stocks: Summary Statistics and Limit Order Usage

This table reports summary information for the 19 stocks on the TORQ database with more than 4,500 transactions over the 63-day period, November 1, 1990 – January 31, 1991. The Limit Order Book (LOB) participation rate is the percentage of total trading volume over the 63 days that involves the limit order book. The % of Limit Orders (LOs) that are potentially stale is measured by examining the profitability of a limit order trade after a lag, by comparing the transaction price to the mid-point of the bid-ask spread. If the trade loses money by this metric it is classified as potentially stale.

	Ticker	Market Cap ($ Billion)	Total # Trans over the 3 months	Min Price	Max Price	LOB participation Rate (Vol) (%)	% of LOs that are potentially stale at: 5 mins	1 hour	1 day
Advanced Micro Devices	AMD	0.3	5,976	3.75	7.375	13.2	20.4	25.8	33.0
Boeing	BA	15.5	18,442	42	49.625	10.2	37.0	37.8	38.5
Colgate-Palmolive	CL	4.4	6,088	65.625	75.5	8.2	27.2	36.0	37.0
CPC International	CPC	5.8	5,607	73.5	83.25	7.6	29.8	34.0	33.2
Community Psych. Ctrs	CMY	1.1	5,088	24.375	31.75	8.2	19.6	34.5	43.1
Dresser Industries	DI	2.4	8,367	16.625	24	10.3	20.0	26.1	32.3
Federal Express	FDX	1.8	4,547	29.5	43.25	8.6	29.3	31.2	40.7
Fed'l Nat. Mtge Assn.	FNM	6.7	18,468	27.75	39.75	5.2	27.2	34.3	37.6
FPL Group (Florida Power)	FPL	3.8	11,305	27	29.875	16.9	17.5	21.8	31.2
General Electric	GE	46.2	39,112	51.25	64	8.5	33.1	35.7	38.2
Glaxo Holdings	GLX	3.0	14,085	28.875	35.625	7.4	20.3	35.3	38.4
Hanson PLC	HAN	2.8	5,458	17	19.75	10.0	15.8	23.1	32.9
Int'l Bus Machines	IBM	60.4	33,654	104.875	127.125	5.7	40.5	40.0	41.2
Philip Morris Co.	MO	43.6	37,646	46.375	55.875	8.5	27.2	32.2	37.4
Potomac Electric Power	POM	2.0	5,528	19.25	21.125	12.0	11.4	20.9	22.6
Schlumberger Ltd.	SLB	13.2	10,836	50.5	59.625	6.0	29.8	35.1	38.0
A T & T	T	37.0	39,216	29.125	35.25	9.5	23.2	29.6	36.7
Exxon Corp.	XON	61.1	16,609	48.25	52.875	11.2	21.1	31.3	32.5
Waban Inc	WBN	0.2	4,940	6.125	12.25	47.1	35.8	37.4	38.2
	Average	16.4	15,314.3	37.5	45.7	11.3	25.6	31.7	35.9
	–Per Day		243.1						
	Median	4.4	10,836.0	29.1	39.75	8.6	27.2	34.0	37.4
	Std Dev	21.4	12,621.8	24.8	28.5	9.1	7.7	5.7	4.6

Table II
Aggregate and Size-ranked Portfolios: Summary Statistics, Limit Order Usage and Profitability

This table provides aggregate properties for the 137 stocks in the TORQ database used in our study. The data are from the New York Stock Exchange, over the period November 1, 1990 – January 31, 1991. The Limit Order Book (LOB) participation rate is the percentage of total trading volume over the 63 days that involves the limit order book. The % of Limit Orders (LOs) that are potentially stale is measured by examining the profitability of a limit order trade after a lag, by comparing the transaction price to the mid-point of the bid-ask spread. If the trade loses money by this metric it is classified as potentially stale.

Panel A. Summary Statistics

Portfolio (Number of stocks)	Statistic ($ millions)	Market Cap	Total # Trans over the 3 months	Min Price	Max Price	LOB participation Rate (Vol) (%)	% of LO's that are potentially stale at:		
							5 mins	1 hour	1 day
All:	Mean:	2,843.6	3210.89	17.6042	22.29	17.6	18.4	21.4	23.5
(137)	–per day:		50.97						
	Std Dev:	9,949.7	6775.12	18.031	21.10	13.7	11.1	10.7	11.4
	Median:	271.3	1171	13.875	18.25	10.9	17.5	21.7	29.7
Large:	Mean:	8,132.3	7689.39	29.3641	36.25	9.6	23.0	28.0	29.4
(46)	–per day:		122.05						
	Std Dev:	16,002.9	10275.64	17.8744	20.56	4.6	8.3	6.5	5.2
	Median:	2,025.1	3320	24.375	30.62	7.6	20.7	27.3	33.7
Medium:	Mean:	298.9	1439.87	16.1658	21.06	19.5	20.2	23.5	26.2
(46)	–per day:		22.86						
	Std Dev:	124.2	1241.46	11.4209	13.21	16.0	10.8	11.1	10.5
	Median:	271.1	1043	14.3125	19.44	13.3	18.9	22.5	30.5
Small:	Mean:	38.7	443.24	7.0533	9.29	23.4	15.2	17.1	18.9
(45)	–per day:		7.04						
	Std Dev:	31.1	483.26	16.8569	19.68	14.6	11.5	8.9	13.2
	Median:	28.1	318	3.125	5	18.6	9.0	10.6	19.5

Table II (Continued)
Aggregate and Size-ranked Portfolios: Summary Statistics, Limit Order Usage and Profitability

This table provides aggregate properties for the 137 stocks in the TORQ database used in our study. The data are from the New York Stock Exchange, over the period November 1, 1990 – January 31, 1991. The profitability of each transaction is measured using the mid-point of the bid-ask spread after the designated lag, relative to the transaction price, and multiplied by the transaction size.

Panel B. Profitability of the Limit Order Book

Portfolio	Mean	Std Dev	Minimum	25%ile	Median	75%ile	Maximum
5-minute lag used to measure trade profitability							
All	\$22,154	\$52,259	-\$167,806	\$2,406	\$7,481	\$17,519	\$308,425
Large	46,159	81,917	-167,806	8,138	16,138	42,469	308,425
Medium	14,565	26,660	-12,488	1,588	5,619	13,312	149,502
Small	6,258	9,392	-1,044	1,169	3,366	8,081	58,112
1-hour lag used to measure trade profitability							
All	\$20,395	\$60,555	-\$93,381	\$1,106	\$6,050	\$14,350	\$455,700
Large	48,890	98,530	-98,381	6,256	14,350	46,119	455,700
Medium	7,817	17,940	-56,831	306	3,822	11,850	61,306
Small	5,096	7,422	-4,131	819	2,784	7,450	40,038
1-day lag used to measure trade profitability							
All	\$13,035	\$95,524	-\$243,119	\$0	\$3,081	\$9,488	\$816,850
Large	34,628	161,854	-243,119	-4,312	4,462	19,481	816,850
Medium	1,297	34,943	-188,869	-6,531	2,878	13,425	66,588
Small	3,650	6,699	-13,569	706	2,212	5,125	28,512

Table III
19 Active Stocks: Effect of Removing Potentially Stale Limit Orders on Measure of Private Information

This table reports properties of the Glosten and Harris (1998) (g) statistic for the 19 stocks on the TORQ database with more than 4,500 transactions over the 63-day period, November 1, 1990 – January 31, 1991. Potentially stale limit orders are identified by examining the profitability of a limit order trade after a lag, by comparing the transaction price to the mid-point of the bid-ask spread. If the trade loses money by this metric it is classified as potentially stale. The statistical test of whether public information is an important component of price movement is whether g–estimated after removing the potentially stale limit orders–is less than the tabulated critical value.

								Sample g after removing potentially stale limit orders:		
Ticker	Sample g Estimate	$E(g)$ (Bootstrap)	Small Sample Bias	Bootstrap standard deviation	Bootstrap 5%ile Value	Bootstrap 95%ile Value	Critical Value	@ 5 minutes	@ 1 hour	@ 1 day
AMD	1.931	1.973	-0.042	1.153	0.066	3.86	0.024	1.878	1.428	1.992
BA	9.136	9.122	0.014	1.135	7.246	10.98	7.26	9.755	9.819	9.939
CL	32.751	32.473	0.278	2.258	28.736	36.16	29.014	34.483	36.084	35.811
CPC	30.093	29.987	0.106	2.233	26.284	33.626	26.39	29.093	29.628	31.631
CMY	13.872	12.495	1.377	2.149	8.924	15.992	10.301	14.685	14.961	14.403
DI	9.896	9.881	0.015	1.445	7.484	12.239	7.499	9.116	10.183	9.851
FDX	23.962	24.074	-0.112	4.041	17.525	30.79	17.413	24.089	26.337	25.652
FNM	7.236	7.202	0.034	0.81	5.867	8.53	5.901	7.569	7.819	8.077
FPL	2.793	2.596	0.197	0.961	1.007	4.171	1.204	2.425	2.504	2.919
GE	5.270	5.264	0.006	0.565	4.332	6.191	4.338	5.349	5.743	5.477
GLX	1.914	1.654	0.26	0.673	0.548	2.778	0.808	1.444	2.090	1.730
HAN	4.36	4.35	0.01	1.05	2.62	6.06	2.63	3.973	4.531	4.371
IBM	17.24	16.89	0.35	1.1	15.08	18.69	15.43	18.551	18.924	18.771
MO	3.83	3.81	0.02	0.59	2.84	4.78	2.86	3.545	3.801	3.632
POM	3.43	2	1.43	1.04	0.09	3.8	1.52	2.176	3.078	3.173
SLB	24.86	23.93	0.93	1.33	21.72	26.11	22.65	26.240	26.917	26.548
T	0.95	0.94	0.01	0.37	0.33	1.55	0.34	0.796	0.793	0.940
XON	3.02	3.14	-0.12	0.86	1.72	4.56	1.6	2.787	2.797	3.319
WBN	16.78	16.77	0.01	1.62	14.06	19.42	14.07	15.433	17.417	14.891
Average	11.2228	10.976	0.25	1.34	8.76	13.17	9.013	11.231	11.834	11.743
Median	7.236	7.202	0.02	1.10	5.867	8.53	5.901	7.569	7.818	8.077
Std Dev	10.216	10.169	0.469	0.856	9.197	11.234	9.243	10.601	11.010	10.974
Std Err	2.344	2.333	0.108	0.196	2.110	2.577	2.120	2.432	2.526	2.518

Table IV
Aggregate and Size-ranked Portfolios: Effect of Removing Potentially Stale Limit Orders on Measure of Private Information

This table provides aggregate properties for the 137 stocks in the TORQ database used in our study. The data are from the New York Stock Exchange, over the period November 1, 1990 – January 31, 1991. Potentially Stale Limit Orders (PSLOs) are identified by examining the profitability of a limit order trade after a lag, by comparing the transaction price to the mid-point of the bid-ask spread. If the trade loses money by this metric it is classified as potentially stale. The statistical test of whether public information is an important component of price movement is whether g–estimated on a sample from which all potentially stale limit orders have been removed–is significantly lower than g estimated using all transactions.

Portfolio (N)	Mean g	Cross-sectional Standard Deviation	Standard Error	# Truncated from above	# Truncated from below	Sample g after excluding potentially potentially stale limit orders: @ 5 minutes	@ 1 hour	@ 1 day
All	23.25	18.65	1.59	2	2	22.66	22.910	23.55
(137)						(-0.37)	(-0.21)	(0.19)
Large	16.35	10.11	1.49	0	0	16.1	16.020	17.06
(46)						(-0.17)	(-0.22)	(0.48)
Medium	29.17	19.61	2.89	0	1	28.1	28.260	29.23
(46)						(-0.37)	(-0.31)	(0.02)
Small	24.09	21.97	3.28	2	1	23.65	23.670	24.23
(45)						(-0.13)	(-0.13)	(0.04)

Number of cases where g is < CV:	No PSLO @ 5 minutes	No PSLO @ 1 hour	No PSLO @ 1 day
All 137 Stocks	6	6	2
46 Large	0	0	0
46 Medium	2	2	0
45 Small	4	4	2

Table V
Sample Monte Carlo Study of the Effect of Public Information Shocks on Return Variance

This table provides an example of the effect of adding incremental public information shocks to the time series of transactions. This is for Boeing (BA), as we add 10-tick public information shocks. In Table III, g estimated on the original sample is 9.136, and the bias-adjusted bootstrapped 95% confidence level is 10.98. The question that this study investigates is how many shocks of this size are needed to move g above the 95% confidence level. Then we evaluate the effect of this number of shocks of this size on Boeing's daily stock return variance.

Number of Shocks	$E(g)$ (Across Simulations)	$\sigma(g)$ (Across Simulations)	$E(\%\Delta\sigma_r)$ (Across Simulations)	$\sigma(\%\Delta\sigma_r)$ (Across Simulations)
0	9.14	0.00	0.00	0.00
1	9.32	0.01	2.39	4.16
2	9.70	0.11	4.73	3.13
3	9.88	0.13	7.56	6.38
4	10.29	0.16	8.95	6.43
5	10.50	0.18	11.84	9.75
6	**11.06**	**0.23**	**13.38**	**11.52**
7	11.41	0.25	16.85	6.97
8	11.13	0.15	18.89	7.22
9	11.37	0.20	20.64	12.92
10	11.76	0.23	24.22	15.49
11	11.44	0.17	26.11	16.48
12	11.91	0.21	28.67	18.64
13	13.06	0.30	31.41	22.00
14	12.82	0.28	32.39	21.17
15	12.96	0.26	35.96	18.37
16	12.52	0.20	38.08	19.78
17	12.65	0.20	41.49	19.84
18	13.49	0.28	41.90	21.34
19	13.55	0.26	44.14	20.60
20	14.25	0.31	45.50	22.32
21	13.91	0.26	49.80	26.02
22	13.50	0.20	53.30	28.22
23	13.82	0.23	53.78	28.36
24	14.00	0.23	59.24	30.78
25	16.04	0.39	59.65	30.42
26	15.76	0.35	61.66	31.94
27	15.53	0.34	66.63	33.86
28	15.22	0.29	69.93	36.58
29	15.84	0.33	70.70	33.95
30	16.69	0.38	69.10	34.62

Table VI
19 Active Stocks: Monte Carlo Assessment of the Power of the Test on g to Detect Public Information

This table reports the results of a Monte Carlo exercise for the 19 stocks on the TORQ database with more than 4,500 transactions over the 63-day period, November 1, 1990 – January 31, 1991. We consider three "sizes" of public information shocks: $0.50, $1.25, and $2.50. The table shows how many of each of these shocks must be added to the original sample in order to move the Glosten and Harris (1988) statistic, g, above its 95% bootstrap critical value. It then characterizes the effect of these additions on the stock's daily stock return variance.

Ticker	# of trades used to compute g	Mean trade size (# shares)	Daily $ Change Std Dev	Daily Return Std Dev (%)	Mean Number of 4-tick shocks needed to move g above CV	% Increase in daily return variance	Mean Number of 10-tick shocks needed to move g above CV	% Increase in daily return variance	Mean Number of 20-tick shocks needed to move g above CV	% Increase in daily return variance
AMD	4,150	3,586.0	0.218	4.523	7	82.49	4	121.95	2	202.47
BA	12,755	2,711.3	1.082	2.344	20	7.24	6	13.38	4	37.1
CL	4,748	1,575.9	0.916	1.302	16	7.88	8	23.4	2	23.06
CPC	4,168	1,526.2	1.190	1.526	17	4.87	7	11.64	5	35.76
CMY	3,286	2,733.6	0.112	2.083	10	13.09	6	49.26	2	65.86
DI	5,793	2,223.7	0.517	2.524	13	20.88	4	42.44	4	182.79
FDX	2,789	2,520.0	0.873	2.496	22	13.06	11	43.17	4	64.18
FNM	11,292	4,111.7	0.756	2.284	13	9.04	7	33.3	4	71.03
FPL	8,624	1,308.9	0.264	0.93	14	80.85	7	56.02	3	334.44
GE	29,012	1,816.7	0.912	1.647	23	12.08	9	27.16	5	63.46
GLX	10,429	2,245.1	0.567	1.773	9	11.07	5	38.34	3	95.55
HAN	4,077	3,884.8	0.324	1.771	5	18.83	2	45.49	1	89.4
IBM	23,997	2,505.7	1.541	1.378	37	6.03	14	15.85	8	33.85
MO	29,253	2,287.8	0.609	1.218	21	23.77	7	49.45	4	112.6
POM	4,530	1,459.0	0.220	1.096	8	65.27	4	205.49	2	413.61
SLB	7,429	2,391.6	1.070	1.934	22	7.85	9	20.06	6	55.23
T	26,525	1,817.2	0.570	1.804	11	14.62	5	41.37	5	171.61
XON	10,870	2,631.2	0.487	0.968	11	19.02	6	63.51	3	125.66
WBN	3,492	2,095.7	0.360	3.866	10	30.48	5	104.16	2	163.7
Mean	10,906.3	2,391.2	0.663	1.972	15.21	23.6	6.63	52.92	3.63	123.23
Mean –excl AMD	112,81.6	2,324.8	0.687	1.830	15.67	20.33	6.78	49.08	3.72	118.83
Median	7,429	2,287.8	0.57	1.773	13	13.09	6	42.44	4	89.4
Std Dev	9,195.9	788.3	0.388	0.930	7.64	24.5	2.77	46.59	1.71	104.23