Execution Timing in Equity Options*

Dmitriy Muravyev
Boston College

Neil D. Pearson
University of Illinois at Urbana-Champaign

[Preliminary Draft: Please do not cite or circulate without the author's permission]

Abstract

Conventional measures of trading costs rely on the quote midpoint as an estimate of the true security value. However, investors employ a more precise estimate, which takes advantage of public information besides best quotes. Investors buy when their public information midpoint is close to the ask price and therefore is above the quote midpoint. As a result, conventional measures have a substantial upward bias, which is particularly large in the options market. The effective and average quoted spreads overestimate actual trading costs by 42% and 87% respectively; or by several billion dollars annually. The timing bias varies across stocks and grew dramatically larger over time. Trades of non-round size pay smaller spreads. Trades cause only smaller part of the observed price impact, while expected changes in the quote midpoint is the dominant component. Conventional measures can be adjusted for the timing bias. Our results indicate that the adjusted measures should be used to make inferences about liquidity and informed trading.

Keywords: Bid-ask spreads, price impact, algorithmic trading, public information, liquidity, equity options.

---

* We gratefully thank Nanex and Eric Hunsader for providing the trade and quote data for the options and their underlying stocks. Dmitriy Muravyev acknowledges financial support from the Irwin Fellowship at the University of Illinois.
E-mail addresses: muravyev@bc.edu (D. Muravyev), pearson2@illinois.edu (N. D. Pearson)

Introduction

If the quote midpoint is replaced with a more precise measure of fair value based on a wider set of public information, will it make a difference? We show that it will because conventional measures based on the quote midpoint significantly overestimate trading costs and price impact. Although the quote and public information midpoints are equal on average, they differ systematically at the time of trades. Investors buy when the public information midpoint is close to the ask price and therefore is above the quote midpoint, as Figure 1 illustrates. Thus, the conventional bid-ask spreads are larger than actual trading costs. The data reveal that investors and their execution algorithms rely on the public information midpoint, academics should follow suit.

Measures of trading costs and price impact play a fundamental role in financial economics. An analysis of permanent and temporary price impacts sheds light on information flows within and between markets[1]. The link between liquidity and asset returns receives increasing attention, with the bid-ask spreads being the most popular measure of liquidity[2]. Realistic estimates of trading costs are required to access an economic magnitude of trading strategies which test market efficiency.

Traditional measures of trading costs universally rely on the quote midpoint as an estimate of the true value.[3]This assumption has historical roots as outsiders often had access only to the best bid and ask prices. However, currently, investors can easily access a considerable amount of public information besides best prices, with limit order book and prices of related securities being the most obvious examples.

The public information allows for a more precise estimate of the true security value[4] than the quote midpoint, which we call the *public information midpoint* (or simply the public midpoint). More specifically, the public midpoint is defined as the best

[1] Hasbrouck (1991) is a good example.
[2] Amihud, Mendelson and Pedersen (2006) provide an extensive survey.
[3] For example, the effective bid-ask spread is defined as a doubled difference between the trade price and the quote midpoint at the time of the trade. Bessembinder and Venkataraman (2009) and numerous other papers have this definition.
[4] We define fair value as a market consensus about a price of a given security. Most of the time the fair value is between the best bid and ask prices, otherwise investors will trade against them.

forecast of the future quote midpoint based on the current quote midpoint and other public variables, as summarized by Equation (1)[5].

$$\hat{P}_{t+T} = E\left(P_{t+T} \mid P_t, X_{0,t}\right) \quad (1)$$

We argue that the quote midpoint should be replaced with the public midpoint in all measures of trading costs and price impact. For example, Equations (2) and (3) show how the effective bid-ask spread becomes the public effective spread (or simply the public spread) after the adjustment. Although theoretical literature acknowledges the importance of accounting for public information in estimating the true security value, empirical literature almost universally ignores[6] this recommendation.

$$\text{Costs} = 2(\text{TradeP}_\text{t} - \text{P}_\text{t}) \rightarrow \text{Costs} = 2(\text{TradeP}_\text{t} - \hat{\text{P}}_{\text{t+T}}) \quad (2)$$

$$\text{Impact} = P_{t+\Delta t} - \text{P}_\text{t} \rightarrow \text{Impact} = \hat{\text{P}}_{\text{t}+\Delta t+\text{T}} - \hat{\text{P}}_{\text{t+T}} \quad (3)$$

The adjusted measures are important because, besides being more accurate, they reflect how investors actually execute trades[7]. The main way to minimize trading costs is the strategy called *the execution timing*. Investors time their purchases to the moments when the public midpoint is close or above the ask price. Thus, the public midpoint is systematically above the quote midpoint at the time of buyer initiated trades. As a result, conventional measures based on the quote midpoint overestimate trading costs creating *the execution timing bias.* If investors were executing trades at random, there would be little difference between the conventional and adjusted measures. However, investors time their trades.

The execution timing implies that conventional measures of price impact also have an upward bias. These measures usually don't account for the expected changes in the quote midpoint. If the public midpoint is much higher than the quote midpoint, the latter will increase converging to the former. At the same time, execution algorithms are likely to buy because the public spread is relatively small. As a result, Figure 1 shows that only a portion of the subsequent increase in the quote midpoint should be attributed to the causal impact of trades, the remainder is simply regression toward the mean.

[5] Equation (1) is linear, but non-linear models can potentially improve the forecast.
[6] Hasbrouck (1991) is one of few exceptions. He controls for past price changes and signed volumes to distinguish the permanent and transitory price impacts. However, other public variables are not included.
[7] Unlike post-trade measures, pre-trade measures not only estimate costs but also tell when to trade.

Importantly, the investor's ability to buy before the price increase is based on effective processing of public rather than private information.

Empirical part of the paper demonstrates that the execution timing bias is large in the options market. The effective and average quoted spreads overestimate actual trading costs by 42% and 87% respectively!

The options market is a perfect laboratory to study the timing bias for two reasons. First, the underlying stock price is obviously the most important public information for option prices. Second, the Black-Scholes-Merton model (BSM) provides a common way to transform the underlying price into the implied option price[8]. Thus, the model selection is much easier for options than for stocks. In addition to this simple model, we also employ a regression model in the spirit of Equation (1). It predicts changes in the option midpoint based on the BSM implied price and information about option limit order book and short-term price dynamics.

We apply these two methodologies to the sample of options on 39 stocks over the three-year period from April 2003 through October 2006, and find several implications of the execution timing. First, as mentioned above, both trading costs and price impact have much smaller magnitudes than was previously believed. In absolute terms, the execution costs are overestimated by several billion dollars per year[9]. Second, there is a substantial variation in the timing bias across stocks. Thus, conventional and adjusted measures of the bid-ask spread can rank stocks differently. Third, the magnitude of the execution timing bias increased threefold in less than four years. The execution timing will be even more important in the future. This finding demonstrates the profound effect of the growth in algorithmic trading on the options market. Finally, the execution bias varies with trade size. Trades of round size, divisible by 10, have substantially worse execution and smaller price impact than non-round trades[10], while the effective spreads are the same for two groups. The round-volume results shed light on the execution timing model used by a representative execution algorithm.

---

[8] An additional assumption is that current option quote midpoint will eventually converge to the BSM implied price.

[9] Indeed, combining option trading volume of 4.6 billion contracts in 2011 with average execution bias of 2 dollars per contract produces a multi-billion dollar number. Each option contract is on one hundred shares. http://www.optionsclearing.com/about/press/releases/2012/01_03.jsp

[10] For example, trades of 30 contracts pay a spread of four cents, while trades of 29 or 31 contracts pay only three cents.

The paper raises an important policy question concerning the protection of retail investors. Currently, retail investors cross-subsidize institutional traders[11] by paying the entire bid-ask spread while institutional algorithms employ the execution timing and pay only half of it. Market makers effectively quote a different spread for each of the two investor types. Information is largely symmetric, and market makers can estimate the probability of a trade coming from each investor type conditional on the public midpoint. The main friction is that retail investors cannot constantly re-compute and update their quotes. This friction can be elevated by delegating this function to the broker or exchange level. One possible solution is to quote option prices in implied volatility rather than dollar terms. A less radical solution would be to encourage exchanges to introduce limit orders linked to implied volatility[12].

Overall, our results indicate that the usage of the adjusted measures is crucial for making inferences about liquidity and informed trading. The execution timing provides exciting insights about the inner workings of algorithmic trading.

Data

The paper employs tick-level option and stock data on 39 stocks including 2 ETFs. The data are provided by Nanex, a firm specializing in delivering high-quality data feeds. The sample period is April 2003 through October 2006 and includes 882 trading days. The selected stocks had the most liquid options as measured by trading volume prior to the beginning of the sample period in March 2003. The data include trades and best quotes for both stocks and options for all exchanges. A more detailed description of the data is provided by Muravyev, Pearson, and Broussard (2012)[13].

We include only options with between 5 and 700 calendar days before expiration. First and last five minutes of a trading day are excluded to avoid opening and closing rotations. Trades for which implied volatility or the public midpoint cannot be computed are also excluded.

[11] If dealers are competitive and make zero profits, profits of one investor group equal to losses for the other.

[12] This solution may pose technical challenges because quote traffic will increase.

[13] Compared to Muravyev et al. (2012), we exclude DIA because through most of the sample period it was traded only at CBOE.

The paper studies of option trades and subsequent quote changes. Tables 1 and 2 summarize main descriptive statistics. The total sample includes 20.4 million option trades. Nasdaq ETF QQQ/QQQQ has the largest number of trades (1.8 million before the ticker change and 1.9 million afterwards) while AOL has only 52 thousand trades[14]. Average trade transaction has a price of 1.7 dollars and size of 40 contracts[15]. However, the trade size distribution is highly skewed with 50th and 75th percentiles of 10 and 20 contracts respectively; and 14% of trades have the smallest possible size of one contract. There are slightly more seller initiated (54%), and call option transactions (64%).

The direction of a trade is determined by the quote rule. If a transaction price is at the National Best Bid and Offer (NBBO) quote midpoint, the quote rule is applied to the best quotes of the exchange which reported the transaction. As 84% of transactions are recorded at NBBO prices, the method is easy to apply. On average three out of six exchanges are quoting the best national at the time of trades.

Methodology

Time is money – by postponing trade execution, investors can lower trading costs. But how exactly can investors achieve this? Although there is a large theoretical literature on optimal trade execution, not much is known about inner workings of the black box of algorithmic trading[16]. We show that the execution timing is one of the most effective ways to minimize trading costs.

As Figure 1 and Equations (1)-(3) imply, the model for the public midpoint plays a key role for the execution timing. In the first step, variables representing public information and a functional form are selected for the main regression in Equation (1). Then, the regression coefficients are estimated on the sample from regular time intervals[17]. Finally, the coefficients are employed to estimate the public midpoint at the time of trades.

[14] AOL dropped from the sample after changing its ticker in October 2003.
[15] Each option contract is on one hundred underlying shares.
[16] The conventional wisdom is to split large trades to minimize temporary price impact. Anand et al. (2012) study trade execution costs for the database of institutional trades. However, they have no information on how broker's execution algorithms work.
[17] Trade timing is endogenous, that is why regular time intervals are appropriate.

Ideally, all public information should be included in the main regression. However, it's hardly feasible because some public information is costly to acquire[18], and historical data are often not available. Despite this difficulty, it is often possible to identify first order variables from general considerations. For example, for the stock market the most relevant variables include price history, state of the limit order book, and market and industry components. Statistical methods of model selection can help with picking specific variables. For the options market the task is much easier as the underlying stock price is clearly the main public information. The model specification is also easier to choose for the options market as the Black-Scholes-Merton formula links option and stock prices.

After the public midpoint model is estimated, the adjusted measures of the bid-ask spread and price impact in Equation (3) can be computed. If there are several alternative models, the most precise one should be chosen.

The adjusted measures can substantially reduce but cannot fully eliminate the timing bias. Having better resources, sophisticated investors can potentially select a better public midpoint model than academics[19]. Thus, our estimates provide a lower bound for the execution timing bias. Indeed, a more precise model will find opportunities to trade at low costs that a simpler model will miss. Similar to the public versus quoted midpoint case, the public midpoint for a sophisticated model will be systematically above the one from a simple model for buyer initiated trades.

$$\left.\begin{array}{c}\hat{P}_{t+T} - P_t = E\left(P_{t+T} - P_t \mid P_t, X_{0,t}\right) \\ E\left(\hat{P}_{t+T} - \hat{P}_t \mid X_{0,t}\right) = 0\end{array}\right\} \Rightarrow \hat{P}_t - P_t \mid X_{0,t} \qquad (4)$$

Importantly, Equation (1) serves primarily as a tool to estimate the difference between public and quoted prices at the present moment rather than to forecast future price dynamics. Indeed, Equations (4) show that if the quote midpoint will converge to the public one within time T[20], then the predicted change in the quote midpoint equals to the current difference between the two midpoints. Time T in the formula should be large enough for the quote midpoint convergence. The convergence can take more than an hour in the option market mainly because of the large bid-ask spreads.

[18] Grossman and Stiglitz (1980)
[19] The model can be non-linear such as neural networks.
[20] An additional assumption is that the public midpoint follows a martingale.

The predictable behavior of the quote midpoint is consistent with the efficient market hypothesis, since the expected profitability is smaller than trading costs and risks. Indeed, the public midpoint is normally within the bid and ask prices. The execution timing is used to minimize trading costs rather than to make arbitrage profits.

To quantify the effect of the execution timing for any given transaction, Equation (5) defines *a measure of the execution timing bias* as one minus the ratio of the public to effective spreads. As the public spread is twice the difference between the transaction price and the public midpoint, the timing bias can be rewritten in terms of the difference between the public and quote midpoints normalized by a bid-ask spread.

$$\text{Timing Bias}_t = 1 - \frac{\text{Public Spread}_t}{\text{Effective Spread}_t} = \frac{\hat{P}_{t+T} - P_t}{\text{BidAskSpread}_t / 2} \quad (5)$$

As discussed above, conventional measures of price impact significantly overestimate the causal effect of trades on prices. To remind the mechanism, investors buy when the public midpoint is close or above the ask price. At the same time, the quote midpoint will increase converging to the public midpoint. Equation (6) decomposes the observed price impact[21] into its causal impact, and the expected change in the quote midpoint if there were no trade. The expected part is estimated with Equation (1) but with smaller time horizon than for the public midpoint. Price impact is traditionally estimated over time horizon of one to twenty minutes that may not be enough for the quote midpoint to converge to the public midpoint.

$$\underbrace{P_{t+\Delta t} - P_t}_{\text{Observed Price Impact}} = \underbrace{P_{t+\Delta t} - \hat{P}_{t+\Delta t}^{t+\Delta t+T}}_{\approx 0 \text{ for large } \Delta t} + \underbrace{\hat{P}_{t+\Delta t}^{t+\Delta t+T} - \hat{P}_t^{t+T}}_{\text{Public Price Impact}} + \underbrace{\hat{P}_t^{t+T} - P_t}_{\text{Timig Bias}} \quad (6)$$

The execution bias in price impact has several implications. First, price impact is commonly decomposed into asymmetric information (private information) and inventory risk components[22]. This paper introduces the third component, the expected quote changes based on the available public information. The execution timing component plays at least as important as the other two.

In addition, the timing bias in price impact has different magnitude for different subsets of trades. For example, the bias is higher for large trades because they are more

[21] Price impact is adjusted for trade direction everywhere in the paper.
[22] See Muravyev (2011) for a recent example.

likely to come from sophisticated investors. Thus, the slope of the price impact as a function of trade size is gentler than is implied by conventional measures.

Last part of the paper tries to reverse engineer the public midpoint model of a representative investor. If two groups of trades have equal asymmetric information and inventory risk, then the difference between their price impacts is entirely due to the execution timing[23]. The simplest model for the public midpoint relative to which the difference in price impacts disappears is the model used by a representative algorithm. Indeed, if investors use a factor which is omitted from the model, then the adjustment for the expected change in the quote midpoint[24] won't fully eliminate the difference in price impacts. On the other hand, if the two groups differ in their private information content, then the difference in price impacts will remain under any model.

The difference in the group price impacts provides a model free lower bound for the execution bias as both groups contain some execution timing. If the difference is significant, it helps to eliminate concerns about the timing bias being mechanically produced by a public midpoint model.

The two groups of trades we choose are round versus non-round sized trades. Round trades are trades of more than fifteen contracts, with size divisible by ten and to a lesser degree trades divisible by five. The comparison can be done separately for each round number or jointly. The non-round trades have larger timing bias because they are more likely to come from sophisticated investors employing algorithmic trading. First, for psychological reasons, unsophisticated investors are likely to choose a round number as a target position, and to acquire it in a single transaction. On the other hand, sophisticated investors are more likely to compute the target position from a model[25]. Second, execution algorithms are likely to split the target size into multiple trades to take advantage of opportunities. If the price is attractive, they will take all the available size at it. Empirically, sophisticated investors use non-round trades for both reasons, but taking all available size at attractive price is more important.

[23] For example, the proportion of noise traders is different for the two groups.

[24] The expected change in quote midpoint can be approximated by the difference between the quote midpoint at the end of the period and the current public midpoint.

[25] For example, exotic derivatives desks commonly hedge with plain vanilla options.

Inventory risk is the same for the two groups because there is little difference in trade size. Compare for example 29 and 31 with 30 contracts. Market makers may think that non-round trades are more informed and react to them more. However, private information content of trades is not observable which makes it hard to test the proposition empirically. As noted above, no model can explain the difference in price impacts if one group of trades contains more private information than the other. However, such a model exists empirically in our case.

Option Market Methodology

The public information midpoint is computed for the options market in two ways. The first method is a simple application of the BSM formula. It combines current stock price and past stock and option quote midpoints and doesn't require historical data. The second method follows the logic of Equation (1). It predicts change in the quote midpoint by a range of public variables including the BSM implied price from the first method. The regression method is more sophisticated and requires historical data to estimate model parameters. The two methods span a large spectrum of alternative methodologies.

As the most precise method should be used to estimate the execution timing bias, we rely primarily on the regression method for this purpose. However, the results in the last section of the paper indicate that many investors use a similar to the BSM approach for their trade execution.

$$\hat{\mathrm{P}}_{\mathrm{t}}^{\mathrm{BSM}}(K,T) = \text{Option Price}_{\mathrm{t}}^{\mathrm{BSM}}(S_t, IV_{t-}, K, T),\ IV_{t-} = \frac{1}{N}\sum_{i=1}^{N} IV_{t-i} \qquad (7)$$

The BSM method consists of two steps outlined in Equation (7). In the first step, average implied volatility over previous 30 minutes is computed[26]. The past implied volatility provides a mapping between option and stock prices, similar to coefficients in a regression. In the second step, current stock price[27] is transformed into the implied option price with the past implied volatility and the BSM formula from the first step[28].

---

[26] Fifteen snapshots of implied volatility with two-minute time step

[27] Even if we use a stock price with one second lag to allow for possible latency between the markets, the results change very little.

[28] We assume no dividends and the risk free rate equal to 60-day LIBOR. Time to expiration is measured using calendar time.

The method can be viewed as a non-linear regression between option and stock prices. It is estimated on the previous 30 minutes and then predicts what option price should correspond to current stock price.

The approach is close to being model free and requires only two main assumptions. First, implied volatility changes much slower than a stock price during a trading day. Indeed, after adjusting for market microstructure effects, implied volatility changes slowly and smoothly intraday[29]. The second assumption is that the implied option price is equal to the quote midpoint on average during 30 minutes before the trade transaction[30].

The second approach for computing the public midpoint is based on a linear regression (8).

$$
\begin{aligned}
P_{t+T} - P_t = \alpha_0 + \alpha_1(\hat{P}_t^{BSM} - P_t) + \alpha_2(\hat{P}_t^{BBO} - P_t) + \alpha_3 \#\mathrm{ExchBid_t} + \alpha_4 \#\mathrm{ExchAsk_t} + \\
+ \sum_{i=1}^{12} \alpha_{i+4}(\Delta_{t-i\Delta t} dS_{t-i\Delta t}) + \sum_{j=1}^{12} \alpha_{j+16} dP_{t-j\Delta t} + \varepsilon_t
\end{aligned}
\tag{8}
$$

The change in the option quote midpoint over the next hour is predicted by a battery of explanatory variables including information about limit order book and short-term price history. The battery accounts for the BSM model by including *the BSM implied bias*, the difference between the BSM implied option price and the quote midpoint. The state of the limit order book is represented by the difference between the average quote midpoint across all exchanges and the NBBO quote midpoint. We also include the number of exchanges at the best ask and bid prices. Option and delta-adjusted stock price changes are taken for 12 five-second snapshots to accommodate the most recent price dynamics.

The regression is estimated separately for each stock and six absolute delta (0.35 and 0.65 cut-offs) and time-to-expiration (60 days cut-off) bins on each day with five second time steps[31]. The average coefficients across all days[32] within each bin are then used for predictions.

[29] The popularity of the BSM model among practitioners is partially driven by its ability to decompose fast-moving option prices into the stock price component and the slow-changing residual called implied volatility.
[30] Over long periods of time quoted and public midpoints are equal on average.
[31] As price dynamics on each day is relatively independent, this methodology simplifies the computation of t-statistics, and spotting the outliers.
[32] The largest and smallest coefficient values are dropped to avoid potential outliers.

Table 5 reports average coefficients across all stocks for ten minute and one hour time horizons which are later used for price impact and trading cost estimation respectively. Changes in the option quote midpoint are highly predictable with R-squared of 10%. The BSM implied bias is the most significant variable, thus the BSM model indeed captures the first order variation in the public midpoint. The average BBO price is the second most significant variable. It is highly correlated with the implied bias but provides some additional information. Consistent with Muravyev, et al. (2012), the option market lags slightly behind the underlying stock; and option midpoint is mean-reverting because of aggressive limit orders. The role of the short-term quote swings diminishes as time horizon increases. The coefficient estimates vary little across moneyness and time to expiration.

Empirical analysis of the Bid-Ask Spreads and Price Impact

Bid-Ask Spreads

The empirical section starts by a comparison of four bid-ask spread measures in Table 2[33]. The average daily quoted spread reflects trading costs for an investor who trades at random[34]. Such an investor will pay 8.4 cents for a round trip trade, which is 20% of an average option price of 1.7 dollars. There are two related ways to reduce trading costs: the quoted spread timing and the execution timing. First, investors can trade when the quoted spread is below average. It can be evaluated by computing the quoted spread at the trade times. The spread is 6.6 cents which is a 1.8 cents improvement over the average spreads. Investors can also try to achieve an improvement over the current NBBO price. However, there is little NBBO price improvement in the options market as the effective spread is 6.4 cents and almost equal to the quoted spread. Indeed, more than 90% of option trades are executed at the NBBO quotes.

---

[33] The public bid-ask spread is twice the difference between transaction price and the price implied by the BSM model based on the current stock price and the lagged implied volatility. The effective spread is the double difference between trade price and quote midpoint. The quoted spread is twice the difference between the relevant best quoted price and the quote midpoint at the moment of transaction. Finally, the average quoted spread is computed separately for each option from one-second snapshots on the transaction day and then matched to trades in a given option.

[34] End-of-the-day bid-ask spreads from OptionMetrics is a special case of the average quoted spreads with only one observation per day.

The last and most important way to improve on random execution is the execution timing. Investors rely on their own estimate of the true price which is more precise than the quoted midpoint. They buy when the public midpoint is close to the ask price and vice versa. The public bid-ask spread is only 4.5 cents, a stunning 1.9 cent improvement over the effective spread. The spread is even smaller at 4.2 cents if the public midpoint is computed with the BSM method.

The comparison between the four bid-ask spreads confirms that the execution timing is an essential element of trade execution and provides a significant improvement over the baseline case of trading at random.

Importantly, the execution timing affects not only the level of trading costs but also the relative ranking of the underlying stocks. For example, Pfizer and QLogic have the same public spreads of 4.3 cents, but very different quoted spreads of 7 and 9.8 cents. Unfortunately, our sample contains too few stocks to conduct a comprehensive cross-sectional analysis.

The execution timing bias has increased by several times over three years. Figure 3 plots how the public spread decreases from 6.5 cents to 3.5 cents while the average quoted spread is unchanged at 8 cents, and the effective spread modestly decreased from 7.5 to 6 cents. So the public spread decreases in half, while the conventional spreads change little.

The trend clearly demonstrates the real effects of algorithmic trading on options costs; yet it cannot be detected with conventional measures. Our results for the options market are broadly consistent with results for the stock market documented by Hendershott, Jones, and Menkveld (2011). Indeed, there was little algorithmic trading before the Options Linkage connected all option exchanges in January 2003[35] which triggered a serious upgrade of exchange infrastructure. This historical observation explains the small execution bias at the beginning of the period. As the options algorithmic trading took off the execution bias steadily increased through 2006.

Other time-series properties of average spreads are worth noting. Although trading costs for any particular stock are quite volatile, market average of trading costs is

[35] Hendershott et al. (2011) argue that there was little algorithmic trading even in the stock market pre-2003.

quite predictable and moves in a narrow range. In this sense, the risk of volatility in trading costs seems to be diversifiable at least during normal times. The spread time-series fluctuate around a long term trend and have a positive autocorrelation, but little volatility clustering is observed, and the volatility of day-to-day changes in the spreads are constant over the period.

The execution timing explains the main stylized fact about option bid-ask spreads, namely why dollar option spreads increase in absolute option delta. Figure 2 shows that for out-of-the-money (OTM) options the average quote spreads are below 7 cents, while the spread is 11 cents for ITM options. By contrast, the public spread is much flatter in delta. The spread increases from 4 to only 6 cents from OTM to ITM. For large trades, the relationship becomes completely flat with 5 cent spread for ITM options.

Cho and Engle (1999) argue that option market makers immediately delta hedge after each trade and thus pay the spread in the underlying stock. However, the theory falls short empirically because it predicts that the difference between OTM and ITM spreads should be less than the underlying bid-ask spread. The bid-ask spread in the stock market is one penny, while the observed difference in spreads between OTM and ITM options is at least four times larger. After accounting for the execution timing the difference becomes two cents for trades of average size, and only one cent for large trades. These magnitudes are comparable to the bid-ask spreads in the underlying stock. Thus, the execution timing can explain why the option bid ask spread increases in absolute delta.

Finally, Table 6 provides a more rigorous conditional analysis of the timing bias measured by Equation (5). The focus is on economic rather than statistical significance because the latter is granted by the large sample size. The timing bias is increasing in absolute delta because stock price movements have a larger effect on ITM options, making execution timing easier. Average timing bias is 0.38 (or 38%), and the change from OTM (delta=25) to ITM (delta=75) will increase the bias by 0.1. Option market makers are aware of this effect as absolute delta becomes insignificant after including a variable for the number of exchanges quoting best price.

The timing bias increases by 0.16 each year reflecting the increased utilization of algorithmic trading. As expected, the number of exchanges quoting the best price in the direction of a trade is a significant determinant of the bias. Each additional exchange reduces the bias by 0.20. In a special case of only one exchange quoting the best price the bias is additionally larger by 0.15. Other explanatory variables have small economic magnitude.

Price Impact

The execution timing has direct implications for price impact. Figure 1 shows that the observed price impact consists of two components: the causal impact of trades and the expected change in the quote midpoint as it converges to the public midpoint.

In the options market, the expected quote change is larger than the causal impact of trades. Table 4 compares observed and expected price impacts for one, ten and sixty minute horizons. The observed price impact is measured in a standard way as a dollar change in the quote midpoint following a trade transaction[36]. For the BSM method, the expected price impact equals to the implied bias because the quote midpoint should eventually converge to the BSM implied price. The regression method first estimates Equation (8) on regular time intervals separately for each time horizon and then predicts quote movements at the trade times. Although ten minutes may not be enough time for the public midpoint convergence, we follow the literature and use it as a baseline case.

The observed price response to trades is rapid and large. The quote midpoint moves by 1.13 and 1.34 cents in one and ten minutes respectively. Although, it's tempting to attribute the large price impact to asymmetric information, in fact, the timing bias constitutes most of it. The BSM method predicts that even without a trade the quote midpoint should move by 1.08 cents which is 81% of the observed price impact. The regression method predicts a 0.82 cent move or 61% of the observed impact. Thus, most of observed price impact corresponds to the expected changes in the quote midpoint. The public midpoint convergence is much faster after trades than during normal time. In the first minute, the regression predicts a 0.42 cent move, while the midpoint moves by 1.13 cents. Trade transactions urge market makers to update quotes and center them around the public midpoint.

[36] Five-minute horizon is standard for the equity market, but the options market is less liquid, and ten minutes is more appropriate. All price impacts and quote changes are adjusted for trade direction.

Price impact is often studied as a function of trade size to infer which trades are informed. Figure 5 this dependence with size measured in number of contracts. There are several stylized facts to note. The observed price impact exceeds one cent even for small trades. It is increasing for small trades and is almost flat (at two cents) for trades of more than thirty contracts. Trades of the smallest size of one lot have larger impact than other small trades. But the most pronounced pattern is that round-size trades have significantly lower (by half a penny) price impact than non-round ones. Finally, both the expected price impact and the implied bias are very close to the observed price impact both in shape and magnitude.

To understand how trades change market perception about the fair price, the observed price impact should be adjusted by subtracting the expected quote changes. Figure 6 plots the price impact adjusted by both the BSM implied bias and by the expected quote changes from the regression method. The regression-adjusted price impact preserves basic properties of the observed price impact but the magnitudes are much smaller. The adjusted price impact is below one cent for any trade size and is increasing. The round-sized trades continue to have smaller impact but the difference decreases from 0.5 to 0.3 cents.

The BSM adjusted price impact makes a big difference. It has many properties which are expected from price impact. It starts almost from zero as trades of one contract have impact of only 0.07 cents. The price impact monotonically increases to about 0.6 cents. More importantly, the difference between round and non-round sizes disappears.

Finally, Table 7 presents a conditional analysis of the observed and adjusted price impacts. The quote midpoint changes are highly predictable with R-squared as high as 9%. The observed price impact doesn't depend on option characteristics such as absolute delta and time to expiration if the number of NBBO exchanges is included. As expected, the state of the options limit order book is a very significant predictor of the quote midpoint. For buy trades, each additional exchange at the ask price reduces the observed price impact by 0.57 cents. If only one exchange is standing at the ask price the quote midpoint will additionally increase by 0.56 cents.

Time trend is very strong, the observed price impact increases by 0.55 cents each year. Price impact is increasing in volume, but the slope is only 0.2 cents per hundred contracts. Coefficient for the level of option price is small confirming that option price impact should be measured in dollar rather than percentage terms.

The coefficient for the expected price response from the regression method is 0.78, and it becomes 0.92 in a univariate regression. The link between observed and expected changes is non-linear. If the expected change is large, then the coefficient is exactly one. But for expected changes of more than 5 cents, the coefficient is only 0.78.

If the BSM adjusted price impact is chosen as a dependent variable striking changes are observed. Most coefficients become insignificant. R-squared falls to zero. Remarkably, the time trend and the number of NBBO exchanges, which are highly significant in all other cases, become insignificant here. These observations together with the analysis of non-round tradees in the next section indicate that option market makers use a model similar to the BSM approach to compute the public midpoint.

Trades of Non-round Size

The comparison of price impacts for trades of round and non-round size provides insights into what public midpoint model is used by a representative execution algorithm. The non-round trades pay smaller spreads and have larger observed price impact. As discussed earlier, the simplest public midpoint model, that can explain the difference in price impacts, is the one used by a representative algorithm. To estimate actual trading costs, we employed the most precise model. Contrary to this, the representative model doesn't need to be the best or even unbiased predictor of the future quote midpoint.

Round trades can be split into two subgroups: one with size divisible by ten and the other with size ending in five. As expected, the round lot effect is more pronounced for the round-ten than round-five trades.

The best way to grasp the round lot effect is a graphical description in Figure 5. In the figure, average price impact is computed within each size category, thus treating each size equally. For trades exceeding 15 lots, the observed price impact is half a penny smaller for round-ten than non-round trades. Remarkably, the difference has essentially the same magnitude for all round-ten size categories. For the round-five trades, the difference is a quarter of a penny. Remarkably, the magnitude for round-five trades is half the magnitude for round-ten trades in all our regressions.

The BSM model for the public midpoint can explain the round-lot effect. Indeed, Figure 6 graphically demonstrates how the difference in price impacts disappears after subtracting expected changes in quote midpoint from the observed price impact. For the regression model, the difference decreases to a quarter of a penny. For the BSM model, the round size effect cannot be visually detected. Thus, the BSM model, or a very similar model, is used by a representative algorithm.

The conditional analysis in Table 7 confirms the conclusions of Figure 5 and 6. We study the sample of all individual trades as well as the sample of averages for each trade size as in Figures 5 and 6. The later sample is more appropriate because it gives equal weight to each size category. Indeed, there are many more 20-lot than 30-lot trades, thus 20-lot size category receives larger weight in the sample of individual trades. The regression includes size and the square root of size to control for global dependence between price impact and size. Option price is included to control for normalization of price impact which is measured as dollar difference rather than percentage of price.

For the sample of individual trades, the round-ten and round-five trades have 0.5 cent and 0.27 cent lower observed price impacts respectively. Adjusting the price impacts by the regression model reduces the round-lot differences in half to 0.23 and 0.11 cents. However, the BSM model explains almost all of the round lot effect, reducing the differences to 0.12 and 0.04 cents.

For the sample of size categories[37], the round-ten and round-five dummies are -0.2 and -0.41 and are comparable to the sample of individual trades. However, for the adjusted price impacts the round lot difference in price impacts becomes very small. The regression adjusted price impacts differ only by 0.13 cents for round-ten trades. Both round-ten and round-five differences in price impacts completely disappear if the BSM model is used for the adjustment. Thus, the BSM model can completely explain the round lot effect.

This result indicates that the BSM model, or a very similar one, is used by a representative algorithm and option market makers. Several results from the previous sections point in the same direction. The BSM adjusted price impact starts from zero and increases smoothly in size as shown by Figure 6. The number of NBBO exchanges and

[37] We include only sizes of less than 100 lots, because there are too few trades for larger size categories.

the time trend become insignificant in the regression for the BSM adjusted price impact but are highly significant in all other regressions in Table 7.

In untabulated results we find that non-round trades are smarter mostly because a non-round amount is placed at the best quotes rather than because investors want to trade a non-round amount. We added an interaction between two dummy variables: non-round trade size and only one exchange quoting best price in the regression of the observed impact. The coefficient for the interaction is 0.38 cents, and the coefficient for the round-ten trades decreases from -0.28 to -0.16 cents.

Interestingly, out of trades larger than fifteen contracts, 60% have size divisible by ten which is six times larger compared to the uniform probability case. It is hard to find a rational explanation for why there are so many round-ten trades.

Conclusion.

There are growing concerns that commonly used market microstructure measures are inadequate in the age of electronic trading. We show that these concerns are warranted. Moreover, we explain the exact mechanism which causes the problem and how conventional measures can be adjusted for it.

The problem arises because conventional measures rely on the quote midpoint as an estimate of the true security value. This assumption is violated in practice because investors aggregate available public information into a more precise estimate. The measurement error in the quote midpoint manifests itself at the trade moments. Execution algorithms time purchases to the moments when their public information midpoint is close to or above the ask price, and thus is systematically above the quote midpoint.

Because of the execution timing, actual trading costs and price impact in the options market are about half of what is implied by conventional measures. Not only levels are off, some securities and types of trades are more exposed to the execution timing than others. These differences are not captured by conventional measures. For example, trades of non-round size have larger price impact than non-round ones. The conventional explanation would be that non-round trades contain more private information. However, the difference disappears after adjusting for the execution timing, which relies only on public information.

Therefore, the adjusted measures are essential for making inferences about trading costs and price impact. The execution timing provides a glimpse in the secretive world of algorithmic trading and its deep impact on modern securities markets.

References

Anand, A., P. Irvine, A. Puckett and K. Venkataraman "Persistence in Trading Cost: An Analysis of Institutional Equity Trades," *Review of Financial Studies*, forthcoming.

Amihud, Yakov, Haim Mendelson, and Lasse Heje Pedersen. "Liquidity and Asset Prices." *Foundations and Trends in Finance* 4.1(2005): pp. 269-364

Bessembinder, H., and K. Venkataraman. "Bid-ask Spreads: Measuring Trade Execution Costs in Financial Markets", 2009.

Cho, Young-Hye, and Robert F. Engle. "Modeling the Impacts of Market Activity on Bid-Ask Spreads in the Option Market." *NBER Working Paper Series* No. 7331 (1999).

Goyenko, Ruslan Y., Craig W. Holden, and Charles A. Trzcinka. "Do Liquidity Measures Measure Liquidity?" *Journal of Financial Economics* 92, no. 2 (May 2009): 153–181.

Hasbrouck, Joel. "Measuring the Information Content of Stock Trades." *The Journal of Finance* 46, no. 1 (March 1, 1991): 179–207.

Hendershott, Terrence, Charles Jones, and Albert Menkveld. 2011. "Does Algorithmic Trading Improve Liquidity?" *The Journal of Finance* 66 (1) (February 1): 1–33.

Muravyev, D.; N. D. Pearson and J. P. Broussard. 2012. Is there Price Discovery in Equity Options? University of Illinois.

Vijh, Anand M. "Liquidity of the CBOE Equity Options." *The Journal of Finance* 45, no. 4 (1990): 1157–1179.

Figure 1 Examples of the execution timing. Panel A shows a stylized example there the public information midpoint is constant. Investors buy when the public midpoint is close ask and thus is above the quote midpoint. Panel B presents a more realistic example. The public midpoint decreases but the bid price remains unchanged for a while allowing investors to execute their sales. Panel C demonstrates how the observed price impact consists of the actual impact of trades and the expected changes in the quote midpoint. Time and price are set on the horizontal and vertical axes accordingly. Solid arrows denote the moment and direction of trades.

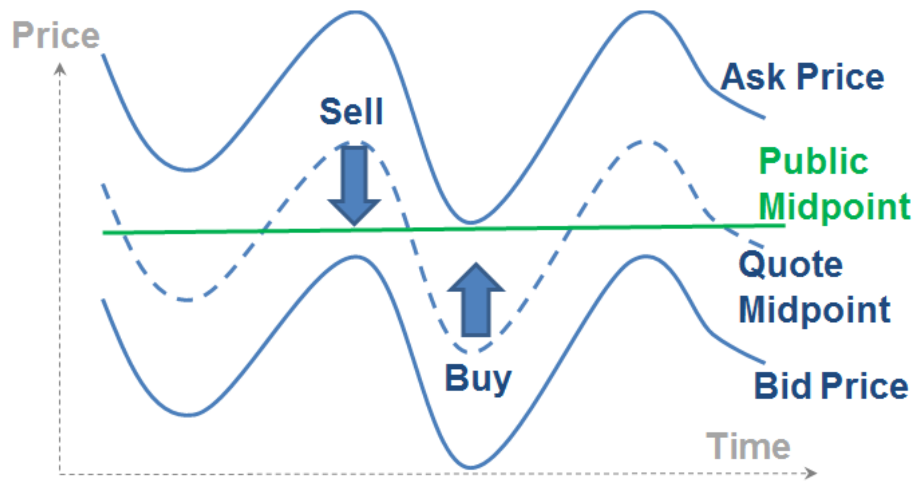

Panel A.

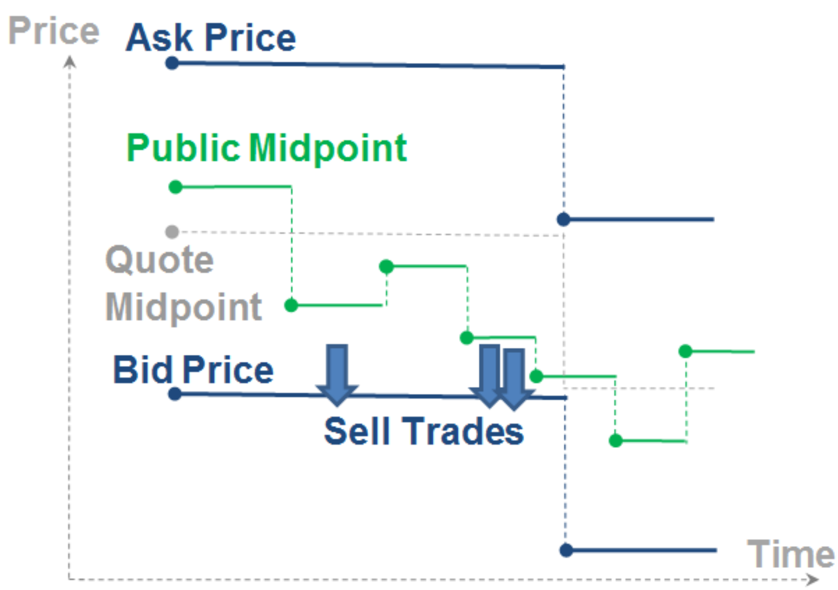

Panel B.

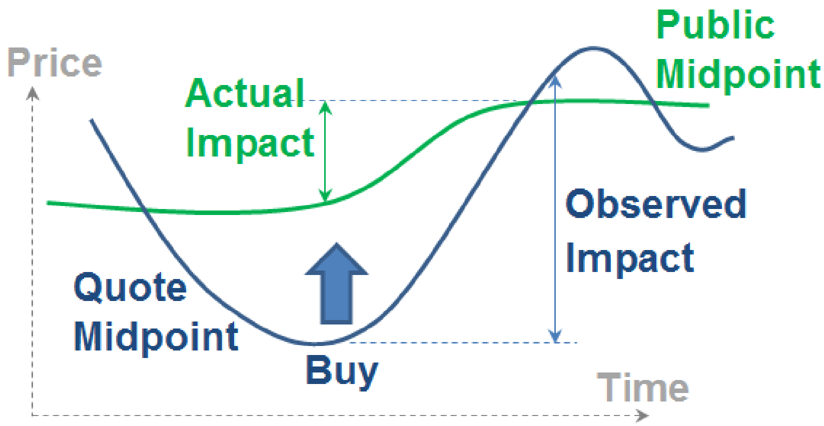
Price
Actual
Impact
Public
Midpoint
Observed
Impact
Quote
Midpoint
Buy
Time

Panel C.

Figure 2 The public bid-ask spread is flatter in option delta than other spreads. The graph plots non-parametric estimates for five types of the bid-ask spreads as a function of absolute option delta. The spreads include the public spread (red line), the public spread for trades larger than nine lots (dashed magenta), quoted (blue), effective (dashed blue), and average quoted spread (dash-dot black). The public bid-ask spread is the double difference between the transaction price and the public midpoint. The public midpoint is based on a one hour forecast of option price from regression (8) which includes the option price implied from the stock market, average quote midpoint across all exchanges, number of exchanges quoting the best bid and ask prices, and lagged changes in option and stock prices. The effective spread is the double difference between the trade price and the quote midpoint. The average quoted spread is computed from one day of one-second snapshots corresponding to each trade transaction. The spreads vary between 3 and 12 cents. The lines are estimated with a kernel regression based on the sample of 20 million trades for options on 39 stocks from April 2003 to October 2006. Option deltas are computed from the Black-Scholes-Merton formula.

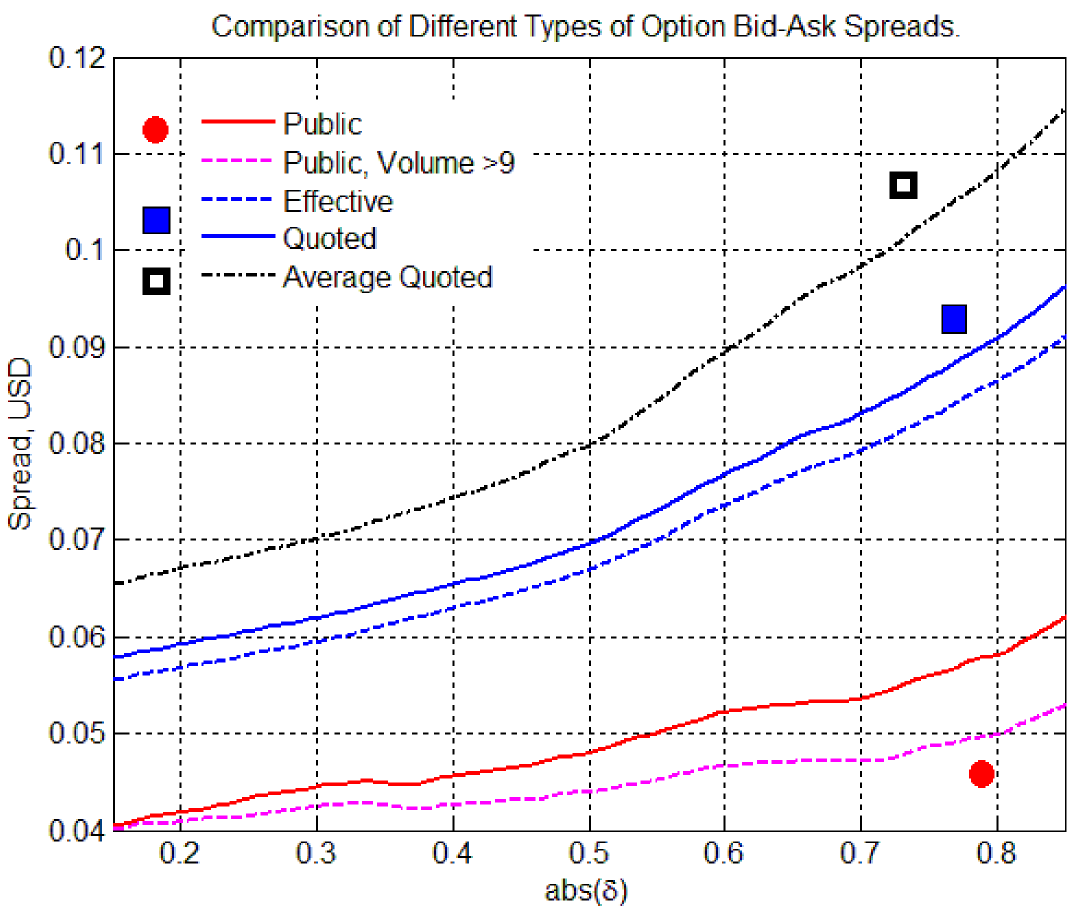

Figure 3 Execution timing is becoming increasingly important over time. The graph plots the evolution of the public (red), the effective (blue), and the average quoted (black) bid-ask spreads over the sample period. In the beginning, the spreads are comparable; however at the end, the public spread decreases in half, while other spreads change little. The decrease in the public spread coincides with the algorithmic trading boom in the options market. The public bid-ask spread is the double difference between the transaction price and the public midpoint. The public midpoint is based on a one hour forecast of option price from regression (8) which includes the option price implied from the stock market, average quote midpoint across all exchanges, number of exchanges quoting the best bid and ask prices, and lagged changes in option and stock prices. The effective spread is the double difference between the trade price and the quote midpoint. The average quoted spread is computed from one day of one-second snapshots for an option corresponding to each trade transaction. Each point is an average across all option trades on a given day. The spreads vary between 3 and 10 cents. The sample period is April 2003 through October 2006

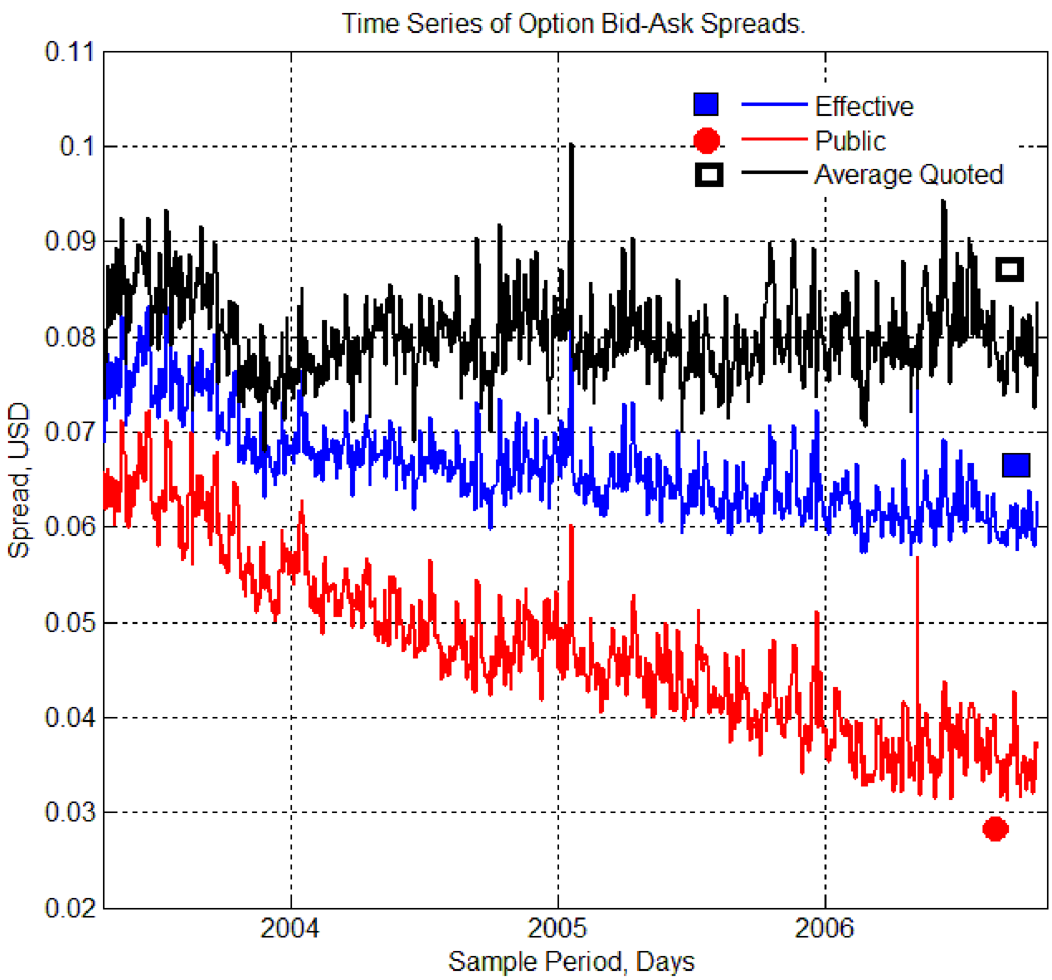

Figure 4 Bid-ask spreads for trades of different size. The graph plots the public (red) and effective (blue) bid-ask spreads as a function of trade size measured in lots. The effective spread is the double difference between the trade price and the quote midpoint. The public bid-ask spread is the double difference between the transaction price and the public midpoint. The public midpoint is based on a one hour forecast of option price from regression (8) which includes the option price implied from the stock market, average quote midpoint across all exchanges, number of exchanges quoting the best bid and ask prices, and lagged changes in option and stock prices. Note that large trades, non-round trades (with size not divisible by 10), and one lot trades achieve better than average execution. Each data point is an average across all option trades of a given trade size. The distribution of trade size is highly skewed (roughly exponential). Average trade size is 42, and its 50th and 95th percentiles are 10 and 114 contracts respectively. The effective spread is stable around 6.5 cents. The confidence bounds are computed separately for each trade size.

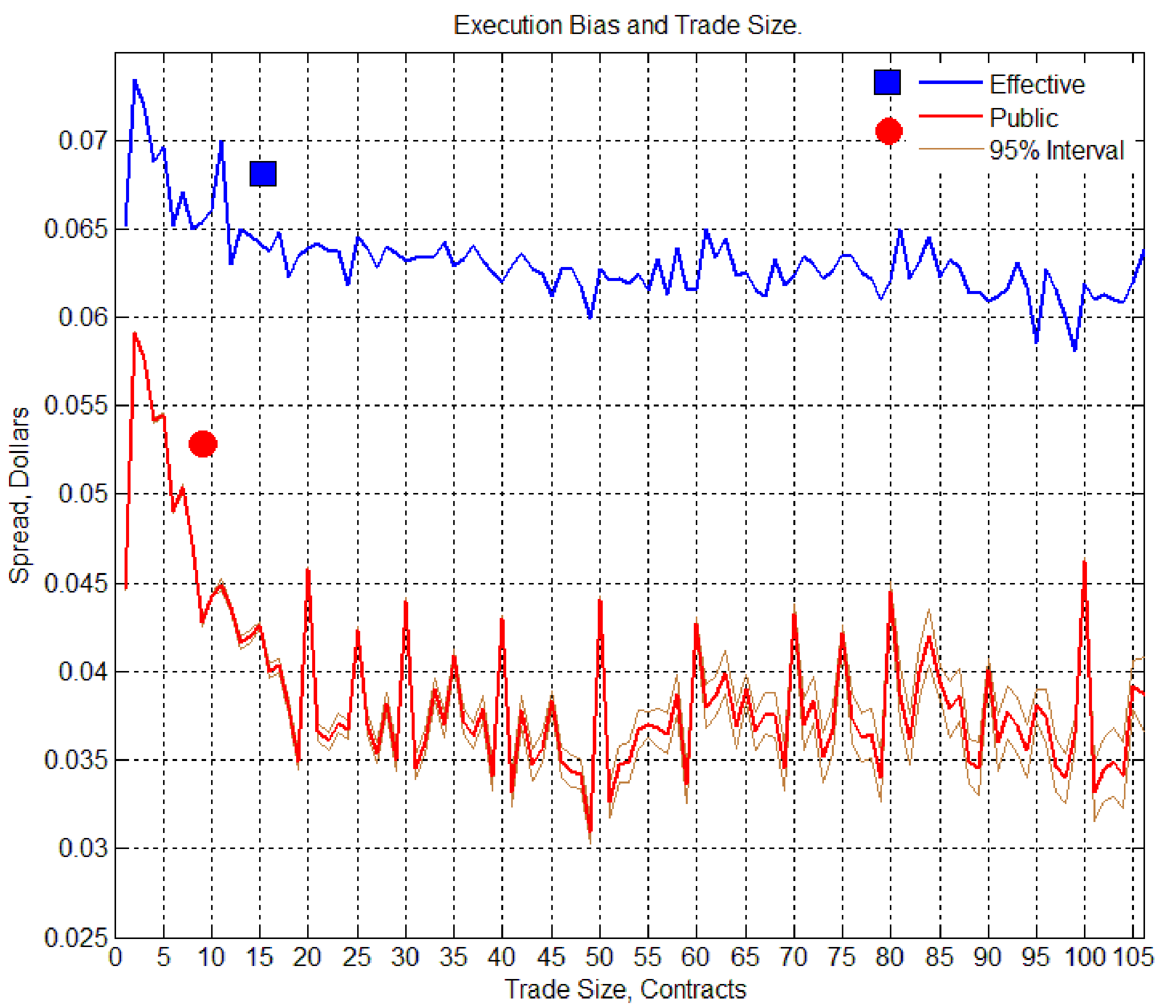

Figure 5 Observed price impact and expected changes in the quote midpoint. Most of the observed response of the quote midpoint to a trade is attributed to the convergence of the quote midpoint to the public midpoint and not to a causal effect of the trade. Change in the quote midpoint in 10 minutes after a trade (blue) is compared with the BSM implied bias (black), and expected changes in price if there were no trade (red) as a function of trade size. The implied bias is the difference between the option price implied by the BSM model from the current stock price and the lagged implied volatility and the option quote midpoint immediately before the trade. Expected quote changes (red) are computed based on the coefficients estimated from regression (8) which includes the implied bias, average quote midpoint across all exchanges, number of exchanges quoting the best bid and ask prices, and lagged changes in option and stock prices. The regression is based on regularly spaced 10-minute time steps. Each point is computed as a simple average across all option trades of a given trade size. The distribution of trade size is highly skewed (roughly exponential). Mean trade size is 42, and its 50th and 95th percentiles are 10 and 114 contracts respectively. Trade size is reported in contracts each on 100 underlying shares, price impact is cents.

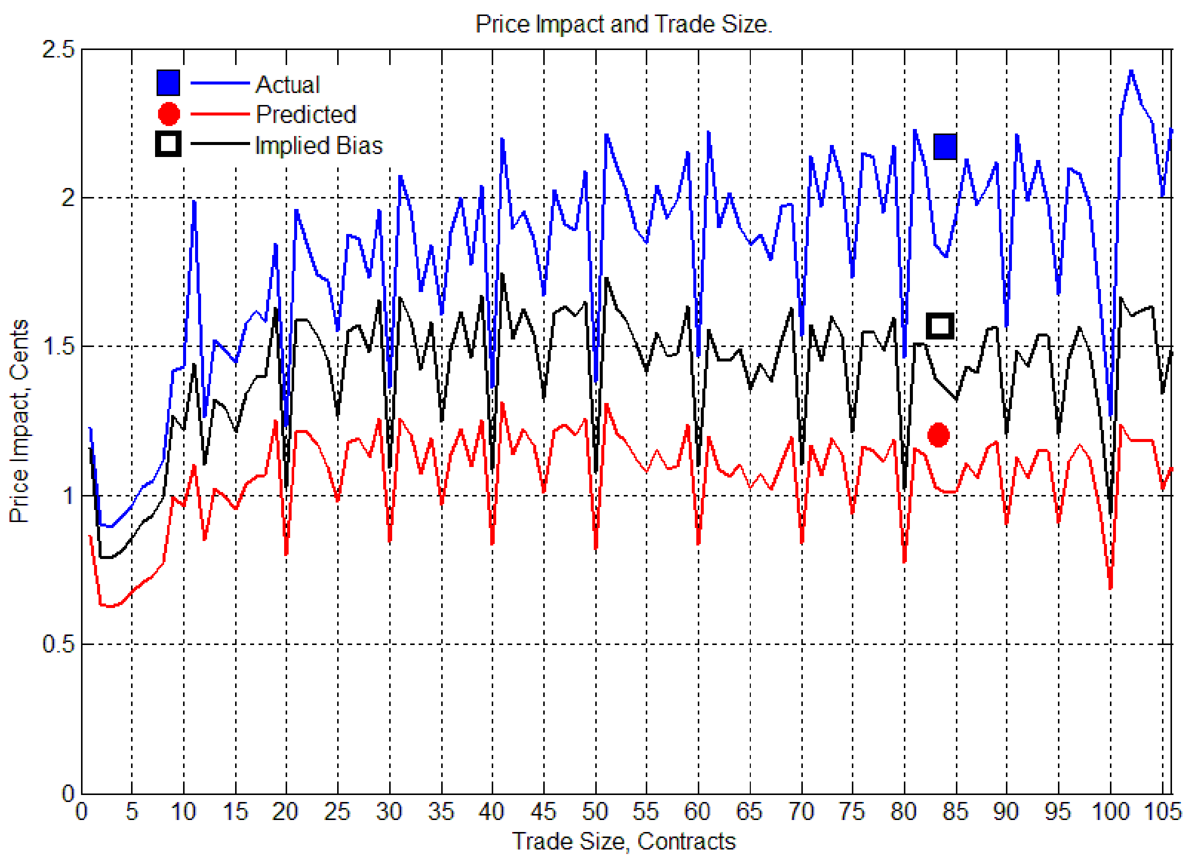

Figure 6 Price impact adjusted for the expected quote changes. 10-minute quote midpoint change is adjusted for predicted changes in the option price by subtracting the BSM implied bias (blue) or by subtracting a prediction from a regression model. The BSM implied bias is the difference between the option price implied by the BSM model from the current stock price and the lagged implied volatility and the option quote midpoint immediately before the trade. Expected quote changes (in red) are computed from regression (8) of the option midpoint 10-minute changes on the implied bias, average quote midpoint across all exchanges, number of exchanges quoting best bid and ask, as well as lagged changes in option and stock quote midpoints. The regression is based on regularly spaced 10-minute time steps. Each data point is computed as a simple average across all option trades of a given trade size. The distribution of trade size is highly skewed (roughly exponential). The mean trade size is 10, and its 50th and 95th percentiles are 42 and 114 contracts respectively. Trade size is reported in contracts each on 100 underlying shares, price impact is cents. The smoothing for the implied bias adjustment (in black) is done via a kernel regression.

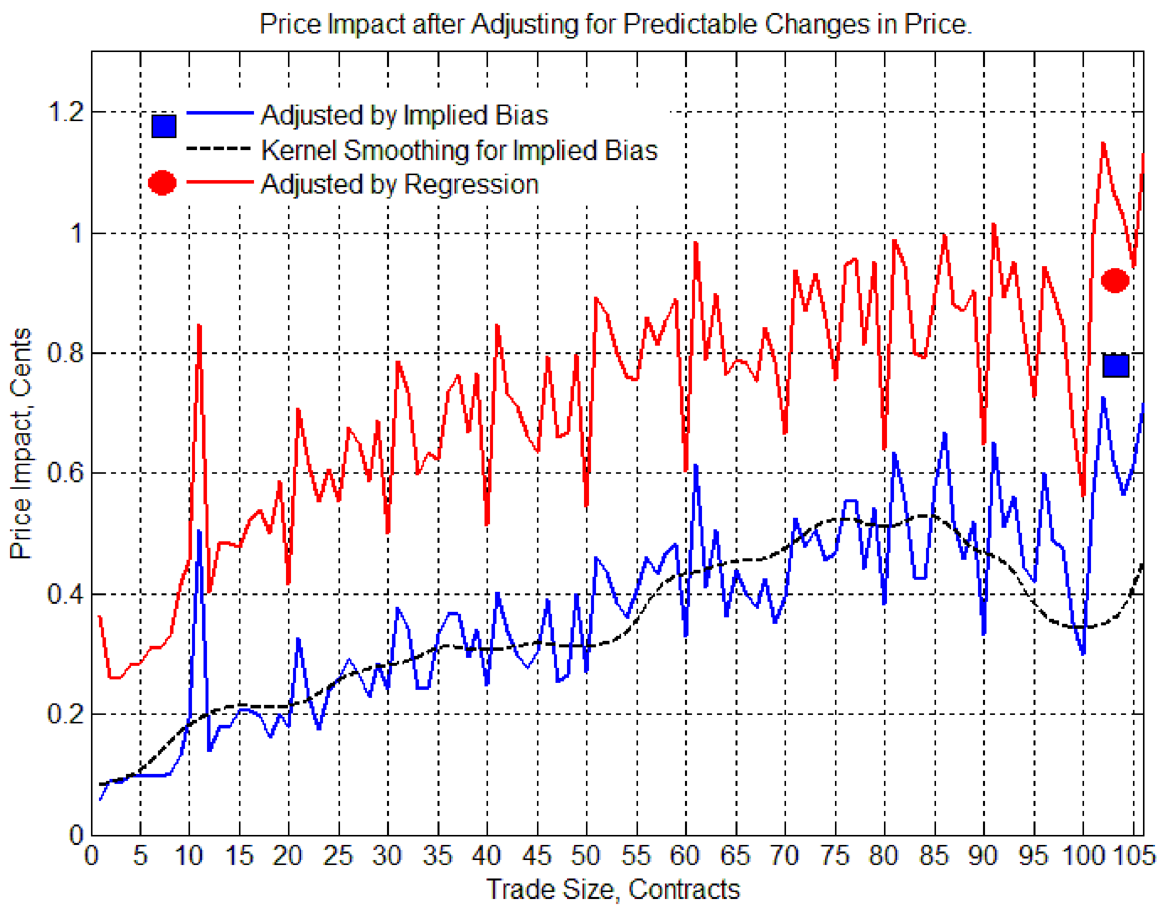

Table 1. Summary statistics. The variables include the execution timing bias for the BSM and the regression methods; the observed price impact and the expected changes in the quote midpoint for the BSM and regression methods for 10 minutes after transaction; absolute option delta; square root of time to expiration measured in calendar days; dummy for call options; time trend; option price and the bid-ask spread; dummy for buyer initiated trades; functions of trade size in lots. Round trades have size divisible by ten and larger than 15. Time trend is in calendar years and normalized to zero at the beginning of the sample period. #ExchAtNBBO is number of exchanges quoting best price in the direction of the trade. Absolute delta is computed from the Black-Scholes-Merton model. Dummy variables are denoted D(x). Mean, standard deviation, 25^{th}, 50^{th}, and 75^{th} percentiles are reported.

Variable	Mean	Std	p25	p50	p75
Timing Bias BSM, %	0.38	0.78	-0.09	0.29	0.82
Timing Bias, %	0.32	0.69	-0.06	0.26	0.70
ΔP_t^{t+T} , 10minutes, Cents	1.31	5.80	0.00	0.00	5.00
$\Delta \hat{P}_t^{BSM}$, Cents	1.10	2.39	-0.29	0.91	2.32
$\Delta \hat{P}_t^{t+T}$, 10minutes	0.84	1.79	-0.15	0.65	1.67
Abs(Delta)	0.45	0.20	0.30	0.44	0.59
Sqrt(T-t)	8.57	5.10	5.00	6.71	11.09
D(Call)	0.64	0.48	0.00	1.00	1.00
TimeTrend	1.72	1.07	0.77	1.69	2.69
OptionPrice, $	1.70	1.94	0.60	1.10	2.10
BidAsk, Cents	6.42	4.13	5.00	5.00	10.00
D(Buy)	0.46	0.50	0.00	0.00	1.00
Sqrt(Size)	4.14	5.03	1.73	3.16	4.47
D(Size=1)	0.14	0.35	0.00	0.00	0.00
D(Size > 15)	0.31	0.46	0.00	0.00	1.00
Size ends in 10 & is >15	0.18	0.39	0.00	0.00	1.00
#ExchAtNBBO	2.97	1.85	1.00	3.00	5.00
D(#Exch = 1)	0.34	0.47	0.00	0.00	1.00

Table 2. Summary statistics by stock. Each column reports an average across all option trades for a given underlying stock: stock and option prices in dollars; option trade size in number of contracts; dummy variable for buyer initiated trades; dummy variable for call options; number of exchanges quoting the best ask price for buy trades and vice versa for sell trades; the dummy variable which equals to one if the transaction price equals to the best quoted price. An average and standard deviation across 39 stocks are reported at the bottom.

Ticker	# Obs.	Stock Price	Option Price	Trade Size	D(Buy)	D(Call)	#Exch. at NBBO	D(Trade at NBBO)
AIG	318,758	61.8	2.5	29.7	0.45	0.62	2.9	0.88
AMAT	413,626	18.7	1.1	27.6	0.45	0.71	2.8	0.81
AMGN	550,588	66.9	2.6	18.2	0.45	0.70	2.8	0.83
AMR	294,827	15.2	1.7	31.7	0.47	0.61	2.6	0.87
AMZN	671,045	42	2.3	21.7	0.50	0.52	2.6	0.83
AOL	52,235	15.1	1.1	35.6	0.43	0.75	2.4	0.67
BMY	220,111	24.9	1.1	31.9	0.42	0.68	3.0	0.84
BRCM	552,115	36.7	2.2	20.4	0.48	0.66	2.7	0.88
C	475,215	46.6	1.5	35.8	0.42	0.67	3.0	0.86
COF	181,717	67.9	3.1	18.6	0.43	0.56	2.2	0.84
CPN	114,616	4.6	0.9	44.8	0.43	0.70	2.4	0.75
CSCO	775,050	20.3	1.1	34.8	0.45	0.74	3.3	0.85
DELL	505,896	32.7	1.4	36.5	0.46	0.66	3.1	0.89
EBAY	1,245,534	63.7	2.9	17.1	0.49	0.64	2.8	0.87
EMC	250,585	12.6	0.9	26.7	0.44	0.78	3.0	0.83
F	203,197	10.6	0.9	37.6	0.46	0.66	3.1	0.84
GE	613,642	32.5	1.2	28.2	0.41	0.73	3.2	0.82
GM	619,137	30.1	2.1	29.3	0.47	0.52	2.8	0.88
HD	371,755	37.2	1.6	19.4	0.42	0.70	3.0	0.86
IBM	739,533	86.1	2.5	19.3	0.46	0.65	2.7	0.84
INTC	1,227,606	24.3	1.3	31.6	0.45	0.71	3.4	0.83
JPM	380,140	37.4	1.5	31.8	0.42	0.65	3.0	0.82
KLAC	313,293	48.3	2.3	18.4	0.49	0.53	2.4	0.82
MMM	273,757	89.3	2.6	14.7	0.45	0.61	2.7	0.84
MO	572,192	64.3	2.4	37.7	0.44	0.67	2.8	0.88
MSFT	969,972	26.3	1.1	41.4	0.45	0.74	3.2	0.83
MWD	147,682	52.8	2.1	31.5	0.45	0.66	2.6	0.82
NXTL	192,156	22.4	1.6	26.5	0.45	0.73	2.4	0.78
ORCL	306,421	13	0.8	39.9	0.45	0.75	3.2	0.84
PFE	710,795	28.3	1.2	34.7	0.42	0.70	3.1	0.85
QCOM	749,730	45.3	2.4	22.6	0.48	0.69	2.9	0.89
QLGC	243,209	38.2	2.1	14.9	0.47	0.63	2.4	0.87
QQQ	1,962,807	34.2	1.2	50.9	0.50	0.54	3.0	0.71

QQQQ	1,840,357	39.1	1.0	54.8	0.52	0.49	3.3	0.91
SBC	107,770	24.3	1.1	27.7	0.38	0.71	2.6	0.78
SMH	386,808	35	1.6	60.7	0.47	0.60	2.9	0.88
TYC	224,817	27.4	1.5	34.3	0.43	0.69	2.9	0.81
XLNX	169,972	29.9	1.6	21.1	0.47	0.61	2.5	0.84
XOM	536,609	55.4	1.9	27.1	0.44	0.71	3.3	0.89
Average	20,485,275	37.47	1.69	30.45	0.45	0.66	2.85	0.84
Std.Dev.		20.24	0.63	10.68	0.03	0.07	0.30	0.05

Table 3 Option bid-ask spreads by stock. The public bid-ask spread is the double difference between the transaction price and the regression public midpoint. The public midpoint is based on a one hour forecast of option price from regression (8) which includes the option price implied from the stock market, average quote midpoint across all exchanges, number of exchanges quoting the best bid and ask prices, and lagged changes in option and stock prices. The simple public bid-ask spread is the double difference between the transaction price and the price implied by the BSM model based on the current stock price and the lagged implied volatility. The effective spread is the double difference between the trade price and the quote midpoint. The average quoted spread is computed from one day of one-second snapshots corresponding to each trade transaction. Spreads are reported in cents. An average and standard deviation across 39 stocks are reported at the bottom.

Ticker	Bid-Ask Spread, Cents				
	Public	BSM Public	Effec-tive	Quoted	Average Quoted
AIG	5.5	5.1	7.8	8.1	10.6
AMAT	3.9	3.6	5.3	5.6	7.0
AMGN	4.8	4.4	7.4	7.8	10.2
AMR	4.5	4.2	6.7	7.0	9.3
AMZN	4.3	3.8	6.7	7.0	9.3
AOL	5.1	5.1	6.0	6.0	7.0
BMY	4.6	4.3	5.9	6.1	7.5
BRCM	4.3	3.8	6.9	7.3	9.7
C	4.9	4.5	6.3	6.5	8.0
COF	5.8	5.4	8.5	9.0	12.1
CPN	4.8	4.6	5.9	6.1	7.4
CSCO	3.8	3.5	5.2	5.3	6.4
DELL	4.0	3.6	5.8	6.0	7.5
EBAY	4.9	4.3	7.4	7.8	10.1
EMC	3.9	3.8	5.3	5.4	6.8
F	4.4	4.2	5.6	5.7	6.9
GE	4.4	4.2	5.6	5.7	6.8
GM	5.0	4.8	7.5	7.8	10.1
HD	4.7	4.4	6.3	6.5	8.1
IBM	5.2	4.8	7.3	7.6	9.7
INTC	3.8	3.5	5.2	5.4	6.4
JPM	5.0	4.7	6.5	6.8	8.4
KLAC	4.1	3.8	6.9	7.3	9.8
MMM	6.0	5.6	8.4	8.8	11.5
MO	5.8	5.4	7.9	8.2	10.4
MSFT	3.8	3.5	5.2	5.4	6.4
MWD	5.3	5.0	7.3	7.6	9.9

NXTL	5.0	4.7	6.4	6.5	8.2
ORCL	3.7	3.5	5.1	5.3	6.4
PFE	4.3	4.0	5.7	5.8	7.0
QCOM	4.2	3.9	6.8	7.0	9.2
QLGC	4.3	4.0	6.9	7.3	9.8
QQQ	3.9	3.7	5.2	5.4	6.3
QQQQ	2.6	2.3	4.7	5.0	6.2
SBC	4.9	4.7	6.0	6.1	7.6
SMH	3.7	3.5	5.7	6.0	7.9
TYC	4.9	4.7	6.4	6.7	8.4
XLNX	3.9	3.7	6.1	6.4	8.8
XOM	4.7	4.2	6.7	6.9	8.7
Average	4.5	4.2	6.4	6.6	8.4
Std.Dev.	0.7	0.7	1.0	1.0	1.6

Table 4 Price impacts and expected changes in the quote midpoint. The observed price impact is computed as the change in the option quote midpoint in one, ten and sixty minutes following a trade. The BSM method expects that the quote midpoint will increase to the BSM options implied price. For the regression method, the expected quote midpoint changes are computed based on the coefficients estimated from regression (8) which includes the implied bias, average quote midpoint across all exchanges, number of exchanges quoting the best bid and ask prices, and lagged changes in option and stock prices. The regression coefficients are estimated separately for one, ten and sixty minute horizons. An average and standard deviation across 39 stocks are reported at the bottom. All variables are in cents.

Ticker	Observed Price Impact, Cents			$\Delta\hat{P}_t^{BSM}$	Expected Quote Change, Cents		
	1 minute	10 minutes	1 hour		1 minute	10 minutes	1 hour
AIG	1.51	1.80	1.90	1.40	0.64	1.08	1.17
AMAT	0.87	1.01	1.07	0.89	0.36	0.66	0.75
AMGN	1.52	1.73	1.76	1.49	0.81	1.21	1.31
AMR	1.45	1.79	2.07	1.24	0.42	0.87	1.13
AMZN	1.42	1.64	1.82	1.43	0.76	1.16	1.21
AOL	0.48	0.65	0.73	0.45	0.13	0.31	0.44
BMY	0.92	1.12	1.22	0.77	0.24	0.53	0.65
BRCM	1.64	1.77	1.86	1.58	0.90	1.31	1.36
C	0.96	1.13	1.20	0.91	0.34	0.66	0.73
COF	1.82	2.26	2.53	1.58	0.76	1.24	1.39
CPN	0.81	1.00	1.13	0.64	0.11	0.32	0.55
CSCO	0.74	0.93	1.02	0.84	0.26	0.57	0.71
DELL	1.06	1.23	1.35	1.10	0.45	0.81	0.91
EBAY	1.59	1.76	1.87	1.57	0.88	1.25	1.29
EMC	0.74	0.93	1.06	0.76	0.21	0.51	0.69
F	0.78	0.98	1.13	0.69	0.15	0.42	0.57
GE	0.67	0.84	0.93	0.69	0.22	0.46	0.60
GM	1.61	1.95	2.09	1.36	0.47	1.00	1.27
HD	1.03	1.19	1.27	0.94	0.39	0.72	0.82
IBM	1.34	1.53	1.59	1.30	0.63	1.00	1.07
INTC	0.75	0.95	1.03	0.89	0.31	0.62	0.74
JPM	0.94	1.18	1.27	0.90	0.33	0.66	0.79
KLAC	1.61	1.75	1.75	1.57	0.97	1.31	1.41
MMM	1.64	1.87	1.92	1.43	0.72	1.13	1.22
MO	1.48	1.78	1.94	1.29	0.52	0.93	1.09
MSFT	0.77	0.95	1.06	0.87	0.25	0.57	0.74
MWD	1.41	1.64	1.80	1.15	0.55	0.92	1.01
NXTL	0.96	1.11	1.19	0.88	0.36	0.66	0.71
ORCL	0.75	0.94	1.02	0.81	0.23	0.53	0.70

PFE	0.83	1.02	1.15	0.82	0.26	0.56	0.68
QCOM	1.48	1.59	1.67	1.47	0.82	1.21	1.32
QLGC	1.52	1.72	1.80	1.48	0.74	1.18	1.31
QQQ	0.56	0.79	0.83	0.74	0.29	0.58	0.65
QQQQ	1.01	1.23	1.29	1.17	0.48	0.91	1.01
SBC	0.68	0.85	0.92	0.64	0.20	0.43	0.56
SMH	1.21	1.41	1.46	1.12	0.54	0.94	1.04
TYC	0.99	1.20	1.31	0.89	0.32	0.64	0.80
XLNX	1.35	1.50	1.55	1.26	0.66	1.04	1.13
XOM	1.33	1.54	1.55	1.23	0.51	0.92	1.01
Average	1.13	1.34	1.44	1.08	0.47	0.82	0.94
Std.Dev.	0.37	0.40	0.42	0.32	0.24	0.30	0.28

Table 5 Expected changes in the option quote midpoint. A regression of option midpoint changes for ten minutes and one hour on the explanatory variables as well as lagged changes in the option and delta-adjusted stock quote midpoints.

$$P_{t+T} - P_t = \alpha_0 + \alpha_1(\hat{P}_t^{BSM} - P_t) + \alpha_2(\hat{P}_t^{BBO} - P_t) + \alpha_3 \#\text{ExchBid} + \alpha_4 \#\text{ExchAsk} + \sum_{i=1}^{12} \alpha_{i+4}(\Delta_{t-i\Delta t} dS_{t-i\Delta t}) + \sum_{i=1}^{12} \alpha_{i+16} dP_{t-i\Delta t} + \varepsilon_t \quad (8)$$

The explanatory variables include the BSM implied bias (the difference between the price predicted by the BSM model and the quote midpoint), average quote midpoint across all exchanges minus the current quote midpoint, and number of exchanges at the best ask and best bid. The regression is estimated separately for each stock and six absolute delta (0.35 and 0.65 cut-offs) and time-to-expiration (60 days cut-off) bins within each day, average coefficients are reported. The lagged quote changes are based on twelve regularly spaced five-second time periods (only the first two and the sum of all twelve coefficients are reported). All quote changes are measured in cents.

Days-to-Expiration	Money-ness	Inter-cept	BSM Implied Bias	Average BBO Price	# Exch at Bid	# Exch at Ask	Stock price changes adjusted for option delta			Changes in option quote midpoint, 5 seconds			R^2
							t-1	t-2	Sum t-1, t-12	t-1	t-2	Sum t-1, t-12	
T = 10 minutes													
short-term	OTM	-0.02	0.26	0.27	0.02	-0.01	0.38	0.24	1.84	-0.17	-0.15	-1.22	0.17
long-term	OTM	-0.03	0.33	0.13	0.07	-0.07	0.39	0.25	1.90	-0.20	-0.18	-1.49	0.13
short-term	ATM	-0.08	0.40	0.31	0.05	-0.05	0.39	0.24	1.78	-0.14	-0.13	-1.15	0.12
long-term	ATM	-0.09	0.44	0.17	0.12	-0.11	0.40	0.25	1.91	-0.16	-0.15	-1.30	0.12
short-term	ITM	-0.27	0.54	0.23	0.09	-0.07	0.32	0.19	1.28	-0.12	-0.12	-1.06	0.09
long-term	ITM	-0.17	0.51	0.18	0.13	-0.11	0.32	0.19	1.34	-0.14	-0.14	-1.22	0.11
T = 1 hour													
short-term	OTM	-0.04	0.31	0.39	0.11	-0.09	0.27	0.13	0.66	-0.07	-0.07	-0.63	0.18
long-term	OTM	-0.08	0.42	0.29	0.12	-0.11	0.25	0.11	0.46	-0.09	-0.09	-0.77	0.13
short-term	ATM	-0.15	0.45	0.44	0.07	-0.06	0.24	0.10	0.33	-0.08	-0.08	-0.82	0.10
long-term	ATM	-0.16	0.53	0.32	0.10	-0.08	0.23	0.09	0.26	-0.08	-0.08	-0.74	0.08
short-term	ITM	-1.03	0.63	0.33	0.06	-0.02	0.13	0.01	-0.42	-0.05	-0.06	-0.69	0.08
long-term	ITM	-0.53	0.56	0.35	0.02	0.02	0.15	0.04	-0.27	-0.10	-0.10	-1.02	0.08

Table 6 The execution timing bias, a conditional analysis. Each column reports a regression of the timing bias on absolute delta, square root of time to expiration in days, dummy for call options, dummy for buy transactions, square root of trade size in contracts, dummies for trades of one lot, number of exchanges quoting best price in the trade direction, option price and bid-ask spread, time trend in years, and finally, a dummy for a single exchange quoting the relevant best price. The last column reports a regression for the subsample of trades with more than one exchange quoting the relevant best price. The timing bias is defined as the expected one hour change in quote midpoint from regression (8) divided by the average quoted spread for a given option. T-statistics based on robust standard errors, which are clustered by date, are reported in parentheses. Stock fixed effects are included but not reported.

Execution Timing, %	Full Sample	Full Sample	Full Sample	# Exchanges > 1
Abs(Delta)	19.367	19.492	0.412	-11.140
	(46.11)	(45.88)	(0.90)	(19.44)
Sqrt(T-t)	-0.009	-0.011	-1.133	-1.551
	(0.58)	(0.74)	(63.15)	(66.01)
D(Call)	-2.028	-2.189	1.081	0.876
	(16.01)	(18.17)	(11.44)	(7.44)
TimeTrend	6.720	6.743	16.450	20.319
	(54.83)	(55.48)	(102.42)	(93.90)
OptPrice	4.068	4.097	0.512	-2.149
	(77.74)	(77.84)	(7.38)	(38.00)
BidAsk	-3.553	-3.570	1.578	3.066
	(97.72)	(96.72)	(35.88)	(117.86)
D(Buy)		-3.564	-6.710	-5.974
		(9.05)	(20.03)	(19.84)
Sqrt(Size)		-0.160	0.047	0.002
		(23.30)	(10.49)	(0.52)
D(Size=1)		1.843	-2.696	-1.258
		(6.66)	(14.90)	(6.01)
#ExchAtNBBO			-19.552	-21.649
			(284.99)	(384.06)
D(#Exch=1)			15.500	
			(41.67)	
R^2	0.05	0.06	0.33	0.26
N, 1000s	20484	20483	20483	13576

Table 7 Conditional analysis of observed price impact. The observed price impact is measured as dollar change in the quote midpoint in ten minutes after a trade. Independent variables include the expected quote changes from regression (8), absolute delta, square root of time to expiration in days, dummy for call options, dummy for buy transactions, square root of trade size in contracts, dummies for trades of one contract, number of exchanges quoting best price in the trade direction, option price and bid-ask spread, time trend in years, and a dummy for a single exchange quoting the relevant best price. The last column uses the difference between the observed price response and the BSM implied bias as the dependent variable. Nonlinearity is examined by including a variable $(\hat{P}_t^{t+T} - P_t)I_{0<x<2}$ which equals expected quote change if it's between zero and two cents, and zero otherwise. T-statistics based on robust standard errors, which are clustered by date, are reported in parentheses. Stock fixed effects are included but not reported.

	ΔP_t^{t+T}	ΔP_t^{t+T}	ΔP_t^{t+T}	$\Delta P_t^{t+T} - \Delta \hat{P}_t^{BSM}$
$\Delta \hat{P}_t^{t+T}$, Cents		0.776	0.584	
		(66.92)	(32.60)	
Abs(Delta)	0.102	-0.127	-0.315	0.201
	(4.20)	(5.77)	(13.53)	(8.87)
Sqrt(T-t)	-0.030	-0.005	-0.005	0.008
	(33.51)	(5.68)	(5.35)	(10.93)
D(Call)	-0.018	-0.046	-0.041	-0.079
	(2.86)	(7.25)	(6.50)	(11.28)
TimeTrend	0.552	0.228	0.201	0.043
	(92.37)	(33.31)	(34.23)	(9.75)
OptPrice, $	0.158	0.079	0.082	0.010
	(38.50)	(20.01)	(19.82)	(2.65)
BidAsk, Cents	0.061	0.017	0.012	0.008
	(26.70)	(7.74)	(6.08)	(4.97)
D(Buy)	0.116	0.144	0.145	0.087
	(7.88)	(8.98)	(9.08)	(4.43)
Sqrt(Size)	0.019	0.019	0.018	0.018
	(50.26)	(50.22)	(50.43)	(48.73)
D(Size=1)	-0.160	-0.079	-0.076	-0.058
	(16.43)	(8.51)	(8.17)	(6.06)
#ExchAtNBBO	-0.571	-0.198	-0.162	-0.015
	(139.13)	(27.08)	(28.51)	(5.11)
D(#Exch=1)	0.562	0.194	0.155	-0.006
	(35.07)	(19.41)	(16.16)	(0.65)
$\Delta \hat{P}_t^{t+T} I_{0<x<2}$			0.435	
			(22.46)	
$\Delta \hat{P}_t^{t+T} I_{2<x<5}$			0.366	

			(21.90)	
$\Delta\hat{P}_t^{t+T} I_{5<x}$			0.201	
			(15.49)	
R^2	0.05	0.09	0.09	0.00
N, 1000s	20,483	20,483	20,483	20,483

Table 7 Comparison of the observed and adjusted price impacts for round and non-round trades. The observed price impact is measured as dollar change in the quote midpoint in ten minutes after a trade. For the regression method, the expected quote changes from regression (8) are subtracted from the observed price impact. For the BSM method, the difference between the BSM implied bias and the observed impact is taken. Independent variables include trade size and square root of it as well as option price. Round trades are considered in two subgroups: one with size divisible by ten and the other with size ending in five. For both subgroups, only trades with size larger than fifteen lots are included, as the round lot effect should be observed only for large enough trades. Dummy for size greater than fifteen lots is included so that round-size dummies estimate proper conditional means. The analysis is done for two sample variations: all individual trades as well as the sample of averages for each trade size as in Figures 5 and 6. Only size categories of less than hundred lots are included in the later sample. For the sample of individual trades, t-statistics are based on robust standard errors, which are clustered by date. For the sample of size categories, robust t-statistics are reported.

	Individual Trades			Size Categories		
Price Impact, Cents	ΔP_t^{t+T}	$\Delta P_t^{t+T} - \Delta \hat{P}_t^{t+T}$	$\Delta P_t^{t+T} - \Delta \hat{P}_t^{BSM}$	ΔP_t^{t+T}	$\Delta P_t^{t+T} - \Delta \hat{P}_t^{t+T}$	$\Delta P_t^{t+T} - \Delta \hat{P}_t^{BSM}$
Size ends in 5	-0.271	-0.109	-0.036	-0.199	-0.045	0.031
	(25.07)	(10.83)	(3.52)	(6.09)	(2.24)	(1.68)
Size ends in 10	-0.536	-0.234	-0.123	-0.414	-0.126	-0.006
	(57.10)	(30.09)	(16.14)	(12.45)	(6.25)	(0.31)
Size > 15	0.665	0.271	0.120	0.156	0.016	-0.041
	(59.77)	(27.52)	(12.18)	(1.78)	(0.34)	(1.21)
Size	-0.000	-0.000	-0.000	-0.017	-0.006	-0.002
	(5.89)	(6.80)	(7.07)	(2.62)	(1.69)	(0.73)
Sqrt(Size)	0.011	0.024	0.028	0.328	0.165	0.104
	(6.29)	(14.53)	(17.26)	(3.23)	(3.02)	(2.90)
OptPrice, $	0.205	0.053	0.042	1.029	0.731	0.705
	(65.74)	(23.85)	(18.61)	(2.89)	(3.73)	(5.30)
R^2	0.01	0.00	0.00	0.85	0.83	0.76
N		20,483,318			99	