Dark Pool Exclusivity Matters*

Leslie Boni†, David C. Brown‡ and J. Chris Leach§

This Draft: May 2012

Abstract

Recent dark pool proliferation has magnified regulatory and academic concern about their fairness and market quality implications. Some dark pools, in the hopes of creating an environment more amenable to buy-side institutional investors, craft their rules to discourage – or even exclude – brokers, high frequency traders and order-flow-information traders. We examine the role participation constraints play in large trade execution and find that more exclusive dark pools experience less serial correlation in returns, less front-running in volume and volatility, and more trade clustering across days. Exclusivity influences execution quality. Not all dark pools are created equal.

JEL Classification: G12, G14, G18

Keywords: Dark pools, exclusivity, market microstructure, institutional trading, block trades

*We would like to thank various practitioners in the industry for helpful motivating discussions, conference presentations, in-house working papers, newsletters, and trade publications on the topic, including executives, managing directors, portfolio managers, traders, analysts, customer support, staff, and developers at Liquidnet, Pipeline, ITG, NASDAQ, NYSE, BATS Global Markets, the SEC, Goldman Sachs, Credit Suisse First Boston, Bank of America, Morgan Stanley, Rosenblatt Securities, Tabb Group, UNX, Capital Group, and Thornburg Investment Management. We would like to thank Liquidnet for providing their dark pool data and Scott Harrison in particular for giving us insight into the operations of Liquidnet. We would also like to thank Paul Koch and seminar participants at the University of Kansas for helpful comments.

†Anderson School of Management, University of New Mexico, e-mail: lboni01@unm.edu

‡Leeds School of Business, University of Colorado at Boulder, e-mail: david.c.brown@colorado.edu

§Leeds School of Business, University of Colorado at Boulder, e-mail: chris.leach@colorado.edu

1 Introduction

A fundamental issue of interest to regulators and academics is the trade-off between fairness and market quality. For example, the U.S Securities and Exchange Commission (SEC) phased-in the Regulation NMS Order Protection Rule in 2006 to increase the probability that investors' displayed limit orders receive fairness in terms of price and time priority.[1] Among the rules' substantial, and perhaps unintended, consequences for U.S equity markets have been the drastic reduction of average trade size, increased high-frequency trading activity, extreme intradaily volume and volatility spikes, and the proliferation of "dark pool" alternative trading systems (ATSs). These ATSs are commonly referred to as "dark" because they do not disseminate their quotes to the public.[2]

Regulators continue to wrestle with the trade-off between fairness and market quality as they contemplate new rules for dark pools. In 2009, the SEC stated that its concerns included dark pools' conveying indications of interest (IOIs) only to an exclusive group of market participants, and that this "could lead to a two-tiered market in which the public does not have fair access to information about the best available prices and sizes for a stock that is available to some market participants." However the SEC argues for exempting "certain narrowly targeted IOIs related to large orders":

> These size discovery mechanisms are offered by dark pools that specialize in large trades. In particular, the proposal would exclude IOIs for $200,000 or more that are communicated only to those who are reasonably believed to represent current contra-side trading interest of equally large size. The ability to have a method for connecting investors desiring to trade shares in large blocks can enable those investors to trade efficiently in sizes much larger than the average size of trades in the public markets.[3]

The SEC's proposed display exemption suggests that dark pool participation could be based on minimum trade size. While motivated by the notion that large trades will more likely originate from contra-side large firms, such an approach does not directly consider the

[1] Securities and Exchange Commission (2005).
[2] As most ATSs do not publish their quotes, we will use "ATSs" and "dark pools" interchangeably. We recognize that some electronic communications networks are not dark and yet are classified as ATSs.
[3] Securities and Exchange Commission (2009).

attributes and trading behavior of said contra-side large firms. An alternative approach is to address directly criteria restricting the likely trading population. In practice, the degree to which trading venues restrict the trading population varies widely. Exchanges commonly exclude no one, allowing all market participants to place orders (displayed or undisplayed). In contrast, many dark pools seek to restrict the eligible trading population. Access can depend on whether a given dark pool admits institutional investors, some or all broker-dealers, high-frequency traders, and specific execution algorithms. In the extreme, a few dark pools design their rules and monitor trading in an attempt to limit access to buy-side (natural contra-side) institutional investors. An interesting question, previously unaddressed in the academic literature, is whether the benefits and market impacts of trading large blocks in dark pools vary systematically with venue exclusivity.

In order to examine the role of exclusivity in dark pool trading, we analyze execution data originating in a dark pool designed to match buy-side institutional investors with other buy-siders (Liquidnet Classic) and provide a contrast to the same execution data originating in other allegedly less exclusive dark and non-dark trading venues. Our results suggest that dark pool exclusivity matters for large trade execution quality. By measuring return correlation, volume, volatility and clustering of large trades, we document that large trades at more exclusive dark pools enjoy higher execution quality.

We document lower magnitude return correlation around large trades at the dark pool specifically designed to encourage buy-side exclusivity, a result consistent with a positive relationship between execution quality and dark pool exclusivity. Similarly, we document a lack of significant volume or volatility increase prior to large trades at that exclusivity-seeking dark pool. Allegedly less exclusive dark pools and order-displaying venues exhibit significant increases in volume and volatility above those at the exclusivity-seeking dark pool, possibly due to the leakage of order flow information and related front-running. Finally, we document greater inter-daily large trade clustering at the exclusivity-seeking venue, consistent with trader's willingness to trade boldly and sequentially, presumably due to a perception of higher execution quality.

The differences we observe are not due to lower trade difficulty at the exclusivity-seeking venue. Comparing trades with similar ones at other dark pools, executions are more difficult, at least if difficulty is calibrated by the large trade's volume share (trade size divided by average daily volume) or the security's liquidity as proxied by average daily volume.

Our results contribute novel empirical evidence to the policy discussion regarding the costs and benefits of darkness and exclusivity in trading venues. We also document relevant empirical regularities for those seeking to model the role different types of dark pools may play in multi-venue trading contexts. The remainder of the paper proceeds as follows. Section 2 reviews relevant background literature. Section 3 introduces our hypotheses. Our data and descriptive statistics are discussed in section 4. Section 5 details our empirical design and results and section 6 concludes.

2 Background and Related Literature

Dark pools have recently experienced increased attention from regulators, media outlets and academics. The previously quoted SEC statements indicate a clear policy interest in dark pools' contributions to market quality and fairness. The Tabb Group estimates that, as of the end of 2010, there are 52 dark pools in the U.S. and 36 in Europe, accounting for 12.5% of U.S. and 10% of European volume.[4] Researchers have recently turned their attention to developing theories and documenting empirical regularities about the use of dark pools, determinants of their market shares, and their impact on market quality.

The theoretical literature investigates dark pool usage in a broader market context. Degryse, Van Achter & Wuyts (2009) suggests that the introduction of a dark pool alongside a dealer market can increase overall transparency without necessarily increasing welfare. Buti, Rindi & Werner (2011) considers security liquidity and differential welfare implica-

[4]The Economist, Off-Exchange Share Trading: Shining Light on Dark Pools, http://www.economist.com/blogs/2011/08/exchange-share-trading.

tions across trader types when a dark pool is introduced alongside a limit order book. Kratz & Schöneborn (2009) investigate a dark pool's impact on trading strategies and costs, arguing that adverse selection and trading restrictions are material. Ye (2011) considers an informed trader who splits orders between an exchange and a dark pool, and finds the dark pool option reduces price discovery. Zhu (2011), in a similar context, argues that a dark pool's contribution to price discovery is less for short-lived information.

Empirical research on dark pools draws equally diverse conclusions. Weaver (2011) shows that increased internalization of volume has resulted in wider spreads and high return volatility, and concludes that increased internal order crossing correlates with degraded market quality. O'Hara & Ye (2011) studies increased market fragmentation, to which dark pools contribute, and finds no harm to market quality. Buti, Rindi & Werner (2010) also finds that dark pools do not decrease market quality. Nimalendran & Ray (2011) provides evidence suggesting that informed traders use dark pools and that their information spills over to other trading venues.

A lack of detailed dark pool data limits many studies of market quality, and encourages the treatment of dark pools as homogeneous. Neither O'Hara & Ye (2011) nor Buti, Rindi & Werner (2010) appear to have venue-specific data. Ye (2010) combines data from eight dark pools to examine execution rates. Nimalendran & Ray (2011), using a single firm's data, documents differences in the information content of trade on the firm's crossing network and its upstairs desk.[5]

Treating dark pool volume as a single homogeneous sample can be problematic if there are important differences in dark pool structures. Butler (2007) classifies 24 U.S. equity dark pools in 2007 according to three characteristics: pricing, order types, and counterparties. Sixteen unique classifications for only 24 dark pools demonstrate their heterogeneity. Domowitz, Finkelshteyn & Yegerman (2009) suggests that differences in dark pools lead to large differences in execution quality. Part of the variation in performance may be due to

[5]Because the firm is not identified, we do not know what restrictions, if any, are placed the trading population for either of the two trading mechanisms.

the trading clienteles permitted to use the venue. Large brokerages provide internal dark pools that Weaver (2011) suggests constitute more than 75% of dark pool order flow. Credit Suisse Crossfinder, UBS PIN, and Goldman Sachs Sigma X are examples. To examine dark pool market share, Ready (2010) analyzes aggregate quarterly data from "the three oldest dark pools that cater primarily to the buy-side": Liquidnet, ITG POSIT, and Pipeline.

Buy-side institutional traders have long had the option of coordinating block trade execution. Zhu (2011) highlights that brokers' upstairs desks compete with off-exchange dark pools, where both allow for coordination between providers and consumers of undisplayed liquidity. Admati & Pfleiderer (1988) discusses potential information leakage and free-riding when institutional investors execute blocks with brokers. Conrad, Johnson & Wahal (2003) finds institutional traders' execution costs are lower for dark pool executions relative to traditional broker executions. Bessembinder & Venkataraman (2004) use data from the Paris Bourse and find upstairs markets provide lower execution costs for large trades and provide opportunities for institutional traders to "tap into pools of 'hidden' or 'unexpressed' liquidity."

3 Hypothesis Development

A trading venue's rules and policies have the potential to influence the composition of its trading population. We consider venues with rules and policies intended to attract natural contra-side traders seeking to avoid potential losses from others' learning about trade intentions. Venues that somehow successfully limit the flow of such information, or the exploitation thereof, may be able to increase execution quality and facilitate less disruptive transfer of large volumes. Ultimately, such venues may, as a consequence, attract repeat business.[6]

[6]That certain types of traders might want to concentrate in a particular venue to provide for credible protected clustering is not a new insight. For example Admati & Pfleiderer (1988) remark that "It is intuitive that, to the extent that liquidity traders have discretion over when they trade, they prefer to trade when the market is 'thick' - that is, when their trading has little effect on prices. This creates strong incentives for liquidity traders to trade together and for trading to be concentrated."

Structuring a venue to attract natural contra-side traders involves limitations on the protections offered. Traders, particularly those who execute large orders and have more to trade, might prefer to keep all information that they have traded private. SEC rules require, however, that trades be publicly disclosed in a timely manner.[7] Therefore venue rules to maintain total privacy of post-trade information are not an option, even though they might be highly attractive to natural contra-side traders. Dark pool venues have more options for designing rules to protect information pre-trade. Through an SEC exemption, dark pools avoid public order display and thereby limit, but possibly do not eliminate, the dissemination of pre-trade order information (Mittal 2008). Beyond this starting point, dark pools seeking to offer additional protection can design trading rules, policies and enforcement aimed at attracting natural contra-side traders. As a complement to the protection afforded by permitted darkness, venues can seek a trading population disinclined to exploit order flow information.

Buy-side institutional investors are often thought of as one such population. These traders are caricatured as having little tolerance or incentive to engage in gaming by trading opposite their intended net position. That is, if they were to encounter order flow information related to a counterparty's desire to trade more, they would be reluctant to switch sides to front-run. Such a reluctance provides the counterparties with some comfort that an inferred desire to continue to accumulate (decumulate) will not exacerbate that task. If, by using trading rules, policies and enforcement, one could create a dark pool environment that would attract a preponderance of such disinclined traders, then that pool might justifiably claim to have created a safer trading environment for natural contra-side traders.[8]

If a dark pool's design features succeed in reducing the exploitation of order flow information, we should observe systematic differences in market quality measures relevant to

[7] Dark pools that report through ADF and TRF must report trades within 30 seconds.
http://www.finra.org/web/groups/industry/@ip/@reg/@notice/documents/notices/p121343.pdf

[8] For the motivation behind our empirical investigation, all we need is that an allegedly exclusive venue's espoused rules, policies and enforcement could lead to a relatively higher proportion of trades taking place between those disinclined to exploit other traders' order flow information *or* that the venue's rules, policies and enforcement inhibit the exploitation of order flow information (in the venue and otherwise). We are not aware of any venue, dark or otherwise, claiming to have successfully restricted the trading population or behavior to completely eliminate the exploitation of order flow information.

the targeted dark pool customers.[9] We consider return correlations, volume, volatility and clustering related to large trades. For such measures, we contrast a dark pool specifically targeting buy-side traders with the aggregation of other dark pools, and with order-displaying (non-dark) venues.

More specifically, if a dark pool's rules and policies have the desired effect, one would expect less magnitude in any correlation between returns in intervals before and after a large transaction: (i) negative correlation would suggest related temporary price pressure in other markets (since the dark pool transaction is typically between the bid and ask prevailing elsewhere); (ii) positive correlation would suggest the possibility that the dark pool transaction supplemented or reacted to price pressure that did not dissipate after the dark pool transaction printed. High magnitudes of such correlations would most likely be considered undesirable by buy-side institutional traders seeking to complete their trading objectives without having other market intermediaries exacerbate the task.[10] Other things being equal (including the shade of darkness),

Hypothesis 1. *More exclusive trading venues should exhibit smaller magnitudes of serial correlation in returns around large trades.*

- **Implication 1.1:** Buy-side only dark pools should exhibit lower magnitude serial correlation in returns than other dark pools.
- **Implication 1.2:** Buy-side only dark pools should exhibit lower magnitude serial correlation in returns than order-displaying venues.

Irrespective of the shade of darkness, more exclusive dark pools that successfully mitigate

[9]Butler (2007) states that "Experience has shown that not all ATSs are the same. Because of its specific order types, constituents, and mechanics, any given ATS may be more or less prone to various forms of market impact or information leakage. This is commonly lumped together as 'gaming' and often comes in the form of predatory traders who seek out orders in ATSs to 'game' them for maximum advantage." Mittal (2008) states that "If there is one thing we can emphasize, it is that all dark pools are different. Yet there is massive push by broker dealers selling dark pool aggregators and algorithms to ignore that fact and push the focus on the fill rate. ... A dark pool's quality directly reflects that of the players in it. Information leakage is less likely to occur where constituents are less likely to benefit, therefore institutions with 'natural' liquidity sit at the top of the quality pyramid. ... So, if you can, it is worth finding out about the types and concentration of constituents in each dark pool."

[10]A contributing factor, particularly undesirable to the buy-side traders, would be the use of other venues to exploit buy-siders' order flow information related to a large dark pool execution, e.g. see Mittal (2008).

order-flow free-riding, other things being equal, should be more prone to relative quiet in trading prior to the execution of a large trade. Undesirable order-flow information leakage and front-running can contribute to increases in aggregate market volume and price volatility prior to a large trade.[11] If a dark pool's design eliminates this contributing factor, other things being equal,

Hypothesis 2. *More exclusive trading venues should exhibit less volume and volatility increase prior to a large trade execution.*

- **Implication 2.1:** Buy-side only dark pools should exhibit less volume and volatility increase prior to a large trade than other dark pools.
- **Implication 2.2:** Buy-side only dark pools should exhibit less volume and volatility increase prior to a large trade than order-displaying venues.

Institutional investors often take multiple days to trade into or out of a large position in a given security. Satisfactory dark pool trades for such institutions can lead to repeat dark pool business in the same security over the span of the trading program.[12] If, as our previous discussion of return correlation, volume and volatility suggests, a more exclusive dark pool provides better large trade execution quality, then it is reasonable to conjecture that, other things being equal,

Hypothesis 3. *More exclusive trading venues should exhibit higher follow-on volume in large trades.*

- **Implication 3.1:** Buy-side only dark pools should exhibit higher follow-on volume than other dark pools.

[11] Although our study of return correlations involves post-trade data, our current investigation does so to focus on changes from the pre-trade period. More generally, we focus on pre-trade rather than post-trade regularities to avoid introducing post-print problems due to the dark pools' obligation to publicly disseminate trade prints expeditiously. Unfortunately, after other markets see those prints, they will most likely respond thereby contaminating clean post-trade-print implications for volume and volatility.

[12] Others have argued that trade execution satisfaction can lead to clustering in transactions. For example, see Ye (2011). Similarly, traders in a multiday program that are dissatisfied with a dark pool transaction will likely be reluctant to rely on that dark pool going forward.

- **Implication 3.2:** Buy-side only dark pools should exhibit higher follow-on volume than other order-displaying venues.

To investigate these hypotheses and implications, we compare data originating in a dark pool targeting buy-side institutional investors to data for executions originating in other dark and non-dark venues.

4 Data and Descriptive Statistics

4.1 Data Collection

In order to investigate the possibility that more exclusive venues foster superior large trade execution, we employ transactions data from Liquidnet Classic, a dark pool having rules, policies and enforcement intended to restrict participation to buy-side institutional investors.[13] Our sample includes 62 days of transactions between January 3, 2011 and March 31, 2011. TAQ data allow us to identify large transactions across dark and order-displaying venues, and CRSP and COMPUSTAT data identify the related average daily volumes, market capitalizations and exchange listings.

Following O'Hara & Ye (2011) and Boehmer (2005), we restrict attention to NYSE- and Nasdaq-listed common stock of US companies that are not closed-end funds, REITs, ETFs or carrying dual class common stock. We further require that stocks in our sample have a minimum daily volume of at least 1,000 shares throughout the sample period, and that the closing price as of the end of 2010 is at least $5. The resulting sample includes 1,694 firms. Table 1 details the sample selection criteria.

Our analysis focuses on trade executions of at least 50,000 shares. While 50,000 shares is arbitrary, we have applied this restriction due to the substantial size of our datasets and the SEC's (and our) interest in large trade executions.[14] We eliminate from consideration large

[13] Appendix A provides additional institutional details.

[14] While SEC comments suggest a $200,000 minimum transaction size, our sample transactions are intended to exceed $250,000 = 50,000($5) with some buffer to accommodate the possibility of prices lower than $5 during the sample period.

trades occurring outside of normal trading hours or under the opening or closing cross. We classify trades using the TAQ exchange code. Dark pool trades, including those at Liquidnet Classic, are reported as “D”, a code used for the Trade Reporting Facilit, and for the NASD Alternative Display Facility. Those facilities include transactions from over-the-counter, some non-exchange ECNs, and dark pools.

For our universe of dark pool trades, we classify any trade of at least 50,000 shares with an execution code of “D” as a dark pool trade. We use Liquidnet disclosure to separate this universe of dark pool trades into Liquidnet Classic and other dark pool trades. While not a perfect separation between exclusive and non-exclusive venue trades, Liquidnet disclosure allows us to compare trades known to have taken place on a venue designed for exclusivity with trades from the universe of other dark pools.[15] 86.0% of the time TAQ data matches Liquidnet-disclosed large transactions with the same second time stamp. Another 13.5% of Liquidnet-disclosed trades match with a one second difference. The remaining 5 observations in our final sample are matched within 5 minutes of the Liquidnet-disclosed time stamp.[16]

Using time intervals surrounding each execution, we calculate pre-trade and post-trade return, volume and trading range. To be included in our reported statistics, we require evidence of continuing trading interest in the intervals before and after a large trade.[17] Our presented statistics are from a sample winsorized at the 1% and 99% levels, although our results are qualitatively similar without winsorization.

The statistics we report are for a subsample of large trades that are likely clean of a possible lack of independence due to clustered trades in the same security. Inclusion in this clean subsample requires that no other large trade occur in that security within the five minutes preceding or following the trade.[18] Table 2 presents descriptive statistics for the

[15]While Liquidnet Classic is structured to cater to buy-side institutional traders, it is not the only dark pool claiming to offer some level of exclusivity and protection. Any improved large trade execution for other exclusive venues would bias against our finding significant differences for the Liquidnet Classic subsample.

[16]We excluded 37 trades from the sample due to their apparent duplication or symbol ambiguity.

[17]Specifically, we require that transactions take place in 15 out of 20 30-second intervals before and after the large trade.

[18]Analogous results holds for a more inclusive sample without the cluster-avoidance. Full sample results are available from the authors by request.

resulting 34,898 non-overlapping large trades.

4.2 Matched Firms

Rather than viewing market quality changes around a large trade as detached from the trading environment, it is appropriate to control for that environment. For example, if we observe a volume run-up in a given security prior to a large trade, it could be just a reflection of market-wide activity. To avoid potential for misattribution, we control for contemporaneous market conditions by using a matched firm approach. Following Davies & Kim (2008), we match firms based on listing exchange, market capitalization and price.[19] We require matched firms to be listed on the same exchange and we minimize the deviation given by

$$D_{ij} = abs\left(\frac{MktCap_i}{MktCap_j} - 1\right) + abs\left(\frac{Price_i}{Price_j} - 1\right) \tag{1}$$

where i indexes potential matches and j indexes the security in our large trade sample.

To ascertain robustness with respect to our matching procedure, we employ three unique matches for each target firm. The first and second follow the Davies & Kim (2008) methodology, where the second match is just the second-best minimizer of the quantity defined in Equation (1). The third match is a broad market match, proxied by an S&P 500 Index ETF.

4.3 Descriptive Statistics

Table 2 details the characteristics of our sample. Trades are segmented into terciles based on average daily volume (ADV) and trade difficulty. ADV is determined by using trading volumes from December 2010. We define trade difficulty as the trade size divided by the ADV of the security. For any given security, all large trades of that security are in the same ADV tercile, but are not necessarily in the same trade difficulty tercile (as the trade size varies but ADV remains constant).

[19] For our sample, we use pre-sample levels prevailing in December 2010.

The descriptive statistics in Table 2 indicate systematic differences between large trades executed at Liquidnet Classic and those executed at the other venues. Only 5% of large trades are executed on an order-displaying venue (column labelled "Open Access Trades"), and these trades are significantly smaller on average (83,689 shares) than those executed in dark pool venues (113,414 shares for Liquidnet Classic and 114,427 for other dark pools).[20] Liquidnet Classic executes 8% of all large trades and the remaining 87% are executed at other dark pools. Our proxy for trade difficulty is significantly higher for dark pool executions at Liquidnet Classic (6.8% versus 3.8% for other dark pools).[21] Panel B shows Liquidnet Classic's market share of large trades is greater for more difficult trades (14% in the highest difficulty tercile compared to 7% in the middle tercile). This suggests that Liquidnet Classic is executing more difficult trades on average. Panel C shows Liquidnet Classic executes a relatively higher percentage of the large trades in low volume stocks as compared to high volume stocks (14% market share versus 3% market share). Overall, the data suggest than any superior performance at Liquidnet Classic (our proxy for a dark pool venue structured to provide exclusivity) is not due to that venue's attracting trades that are easier to execute.

Liquidnet Classic's intra-day trading pattern presents distinctly from that of other dark venues taken as a whole and from order-displaying venues (also taken as a whole). Figure 1 shows the proportion of large trades transacted in each 30-minute period of the trading day for each class of trading venue. The traditional "U"-shaped pattern is observed for both order-displaying venues and other dark pools; however, Liquidnet Classic's trades exhibit a decreasing pattern throughout the day, suggesting a fundamental difference in trading behavior. Liquidnet Classic's volume pattern is consistent with anecdotal evidence that institutional traders split their orders, allocating a portion to the exclusive dark pools and the rest to trading strategies that have higher execution probabilities. As the day progresses,

[20] Significance, based on Wilcoxon rank-sum tests, is at the 1% level. We cannot conclude there is size difference between Liquidnet Classic trades and trades at other dark pools.

[21] Liquidnet Classic trades are significantly more difficult than other dark pool trades, at the 1% level, for the full sample, all trade difficulty terciles, and all ADV terciles. As compared to trades at order-displaying venues, Liquidnet Classic trades are not consistently more difficult within terciles, but are more difficult overall.

trades not executed in the exclusive dark pools may be shifted to the other strategies. Such shifting as the day progresses is consistent with institutional traders' preferring dark pool execution, but managing the risk of non-execution as suggested in Ye (2010).

5 Results

5.1 Tests for Serial Correlation in Returns

In order to test Hypothesis 1, we analyze the relationship between abnormal returns immediately before and after large trades. Lack of correlation between abnormal returns is consistent with high execution quality for buy-side institutional traders, while positive or negative correlations can be viewed as harmful to their trading programs. Abnormal returns are calculated as the difference between the return of the security and the return of the matched security over the same period. Figures 2, 3 and 4 display scatter plots of the pre-trade and post-trade abnormal one-minute returns for trades at Liquidnet Classic, order-displaying venues, and other dark pools, respectively. Visual inspection, and the superimposed univariate linear regression lines, suggest slight negative correlation at other dark pools and order-displaying venues (treated as a whole), but slight positive correlation at Liquidnet Classic (our proxy for a more exclusive dark pool venue). Kolmogorov-Smirnov tests reject the null hypothesis that the OLS residuals in the graphically-inspired linear regressions are normally distributed. Rather than proceed with some form of robust regression, we adopt an approach that combines observed match-adjusted returns from before and after a large trade and proceed nonparametrically. Specifically, we construct the following ratio of returns that facilitates consideration of positive or negative correlation:

$$ReturnRatio_i = \frac{(PostReturn_{i,test} - PostReturn_{i,control})}{(PreReturn_{i,test} - PreReturn_{i,control})} \tag{2}$$

where i indexes trades, $test$ indicates the return of the security with the large trade and *control* indicates the return for the matched security. Positive values of the return ratio reflect positive serial return correlation and come from either consecutive positive or consecutive negative returns. Negative return ratios result from either a positive return followed by a negative return, or a negative return followed by a positive return, and therefore reflect negative serial correlation in returns.

A ratio of returns conveniently captures these categories of theoretically interesting co-variations by combining two returns into a single quantity amenable to univariate statistical tests. However, the ratio's disadvantage is that it introduces potentially large nonlinear transformations when pre-trade returns are close to zero.[22] In response, we consider non-parametric statistics emphasizing signs and ranks rather than magnitudes of the return ratio.[23] We conduct traditional sign tests for non-zero median return ratios. To compare return ratios from potentially different populations, we employ Wilcoxon rank-sum tests. In both cases, rejection of the null hypothesis suggests that return correlations are significantly different from zero (sign test) or each other (rank-sum test).[24]

Regarding the sign test to detect serial correlation, the test statistic of interest is the M-sign, the number of positive return ratios less the number of negative return ratios divided by 2. Table 3 displays M-signs for each type of venue. Panel A shows that large trades at other dark pools and order-displaying venues reflect significantly negative median return ratios. Despite the low power associated with the sign test, these statistics (-782.5 and -54.5 in the total sample) are significant at the 1% level. In contrast, the M-sign for Liquidnet Classic trades (-17 in the total sample) is not significantly different from zero. These results suggest that there is a preponderance of negative return ratios, consistent with large trades'

[22] Any trade with an abnormal return of zero in the pre-trade period has an undefined (infinite) return ratio and is omitted from the sample.

[23] In an alternative, but related, approach we could have applied a sign transformation to each return (yielding -1, 0 or +1) and then performed statistical analysis on the transformed data. Both approaches are inherently sign-based and lead to similar qualitative inferences.

[24] If the expected return ratio is zero, the sign test's test statistic can be easily interpreted. We recognize, however, that sampling error or non-zero expected returns can lead to systematic biases in the return ratio, and that such biases can also affect the interpretation of Wilcoxon rank-sum test statistics. Accordingly, we test the robustness of our results by bootstrapping the distributions of the test statistics using the sampled pre-trade and post-trade returns. Inferences are qualitatively unchanged using the bootstrapped distributions.

being in the middle of a price reversal, in the other dark pools and order-displaying venues. For the Liquidnet Classic sample, the absence of a preponderance of negative or positive co-movements is consistent with either the absence of serial correlation or the nearly perfect balancing of positive and negative co-movements. In the volatility tests that follow, we provide additional evidence that the Liquidnet Classic return ratio's insignificance (difference from zero) is most likely due to the absence of serial correlation.

Analysis of large trades by trade difficulty and ADV terciles supports the results from our full-sample tests. Panels B and C of Table 3 show that other dark pool trades consistently exhibit return ratio distributions having significant negative correlation. Large trades at order-displaying venues are less consistent, evidencing significant negative correlation for lower difficulty trades and trades in higher volume stocks. As in the full sample, examining trades at Liquidnet Classic by tercile shows no evidence of either a preponderance of positive or negative return ratios.

We next test for significant differences between large trade return ratios at Liquidnet Classic and those at other dark pools and order-displaying venues. Between the columns of Table 3, we display the test results for the venues being compared.[25] Large trade return ratios at Liquidnet Classic are more consistently positive than those at other dark pools; we reject equality at the 2% significance level. A similar result, though slightly weaker (p=0.063), holds when comparing Liquidnet Classic to order-displaying venues. Tercile-level analysis supports these results in direction, but statistical significance suffers in the smaller samples. Overall, the segment results provide additional confidence that no single tercile is driving the general result that return serial correlations around large trades at Liquidnet Classic exhibit less net negativity than other dark pools and order-displaying venues.

To establish robustness in our results, we replicate the analysis using our two alternative matches and using an alternative method for calculating returns.[26] All results using these

[25]The z-score shown compares the right-most column to the left-most, e.g. the value of -2.335 in Panel A is consistent with other dark pool trades' return ratios' being centered significantly below Liquidnet Classic's. The value of 1.860 suggests the distribution of Liquidnet Classic trades' return ratios is centered significantly above that of order-displaying venues'.

[26]The results presented herein are based on returns calculated using prices determined by the last trade prior to the start

alternative methods are qualitatively similar to those reported herein.

Liquidnet Classic trade return ratios suggest significantly less negative correlation between pre-trade and post-trade returns, demonstrating one way exclusivity differences coincide with differences in trade experiences. From the combination of tests, the data appear to support Hypothesis 1 that more exclusive dark pools show significantly less negative correlation in returns surrounding large trades.

5.2 Pre-Trade Volume and Volatility Increases

We hypothesize that pre-trade volume run-ups and volatility increases prior to large trades are significantly lower at Liquidnet Classic relative to other dark pools and order-displaying venues. Relative quite prior to large trade executions is consistent with a lower level of order-flow information leakage and related front-running. We test for volume increases by measuring the amount of volume traded in each minute prior to large trades and then comparing the increases in volume traded per minute between trading venues. We proxy for volatility by measuring the trade price range (highest reported trade price minus lowest reported trade price, excluding the large trade itself) in each minute prior to large trades and then testing for differences between the increases in range prior to large trades at different venues.

We use abnormal, rather than absolute, volume measures to account for differences in the timing and trade difficulty of large trades executed at venues with varying degrees of alleged exclusivity. We calculate the abnormal volume increase by first normalizing volume using the prior period's volume in the same security. We then subtract the normalized measure from the matched firm to capture any abnormal volume increase. Specifically, the abnormal volume increase measure is given by:

of a time interval. For example, if the last trade before 10:00 AM is a trade made at \$10.01, then the price as of 10:00 AM is recorded as \$10.01. The alternative method uses the first trade price after the start of a time interval. Continuing the example, if the next trade occurs at 10:00:05 AM at a price of \$10.03, then the reported 10:00 AM price would be \$10.03.

$$VolumeIncrease_i(r,s,t) = \frac{TotalVolume^i_{(s,t)}}{TotalVolume^i_{(r,s)}} - \frac{TotalVolume^{Match(i)}_{(s,t)}}{TotalVolume^{Match(i)}_{(r,s)}} + \epsilon_i \tag{3}$$

where t is the trade time, $TotalVolume^i_{(r,s)}$ is the volume in stock i from r minutes prior to the trade to s minutes prior to the trade, and $TotalVolume^{Match(i)}_{(s,t)}$ is the volume for the stock matched to i from s minutes prior to the trade to the time of the trade. This method has the benefit of naturally controlling for volume level by testing for abnormal volume increases on a percentage rather than an absolute basis. We use non-parametric methods to test for abnormal volume increases and for differences in the increases between trading venues.[27] To estimate this model, the volume in the minute prior to the large trade is normalized by the volume in the prior minute (from two minutes before the trade to one minute before the trade). Any trade with no volume in either pre-trade period is omitted from the sign and Wilcoxon rank-sum tests.

Table 4 shows that volume increases significantly prior to large trades at other dark pools and order-displaying venues. Panel A shows that there are significant (at the 1% level) volume increases prior to large trades for both classes of non-exclusive venues, while also showing that we are unable to reject the null hypothesis that there is no volume increase prior to large trades at Liquidnet Classic. The Wilcoxon rank-sum tests provide formal evidence that the differences in volume increase are significant. The volume increases prior to large trades are significantly greater at other dark pools (p-value 0.006) and order-displaying venues (p-value 0.014) as compared to Liquidnet Classic. As in the tests of negative correlation, the results from studying the sample in terciles (shown in Panels B and C) are qualitatively similar, but suffer from a lack of power. Using alternative matched stocks does not qualitatively change our results. We conclude that large trades at Liquidnet Classic experience less pre-trade volume increase. This is consistent with either lower order-flow information leakage, or a trading population less inclined to front-run using leaked information

[27]Using a standard differences-in-differences approach is intuitively appealing, but the residuals from such a regression analysis are non-normal, making inferences from the resulting coefficients unreliable.

at a dark pool designed and marketed as providing exclusivity.

Following a similar methodology, we test for volatility increases before large trade executions. Our volatility increase measure is given by:

$$VolatilityIncrease_i(r,s,t) = \frac{Range^i_{(s,t)}}{Range^i_{(r,s)}} - \frac{Range^{Match(i)}_{(s,t)}}{Range^{Match(i)}_{(r,s)}} + \epsilon_i \tag{4}$$

where t is the trade time, $Range^i_{(r,s)}$ is the trade price range in stock i from r minutes prior to the trade to s minutes prior to the trade, and $Range^{Match(i)}_{(s,t)}$ is the trade price range for the stock matched to i from s minutes prior to the trade to the time of the trade. We normalize the volatility around a large trade executions by using the trade price range from two minutes prior to one minute prior to the trade. Any trade with no trade range in either pre-trade period is omitted from tests.

Table 5 shows that volatility increases prior to large trades at the less exclusive dark pools. Panels A, B and C show that this volatility increase is significant at the aggregate level and across all ADV and trade difficulty terciles. Large trades occurring at Liquidnet Classic and order-displaying venues do not experience significant volatility increases prior to execution. Wilcoxon rank-sum tests confirm that the difference in volatility increase between other dark pool trades and Liquidnet Classic trades is significant at the 5% level. There is no significant difference in volatility increase prior to large trades at Liquidnet Classic and those at order-displaying venues. We reject the null hypothesis that there is no volatility increase prior to large trades at other dark pools, but we fail to reject the null hypothesis for large trades at either Liquidnet Classic or order-displaying venues.

Our evidence of differences in volume run-ups and volatility increases prior to large trades demonstrates that designing a dark pool to foster exclusivity matters for execution quality. We find significantly more volume run-up and volatility increase prior to large trades at other dark pools as compared to Liquidnet Classic. These results support Hypothesis 2 and are consistent with order-flow-information leakage and front-running's being less prevalent at more exclusive dark pools. These tests provide evidence suggesting that more exclusive dark pools can provide higher execution quality.

5.3 Clustering Across Days

We hypothesize that large trades at more exclusive dark pools have higher execution quality, and as a consequence will receive more repeat business via higher follow-on volume. We test this hypothesis by using a balanced panel dataset of 105,028 stock-day observations to study inter-daily trade clustering. Each observation in the panel consists of static stock data (ADV, price, exchange listing) and time-varying trade data. The daily trade data include counts of the number of large trades per type of exchange (order-displaying venues, other dark pools and Liquidnet Classic).

To study the inter-day clustering of large trades, we estimate autocorrelated negative binomial regression models. If there is no clustering in the data (i.e. the probability of the number of trades today is independent of the number and timing of trades in the past) then we expect that lagged dependent values will not improve model fit in the estimation process. We use the negative binomial distribution to model the number of daily trades at each type of trading venue.[28] For example, we model the total number of large trades and also the total number of large trades at Liquidnet Classic.

The model is estimated via maximum likelihood methods where the mean and variance of the negative binomial distribution are functions of the independent variables. The mean of the distribution is given by:

$$E(y_i|X_i) = m(X_i, \beta) \tag{5}$$

where i indexes a stock-day and we assume that the function m takes the form:

[28] The negative binomial distribution is used to model discrete data processes, usually in the context of how many events occur before something happens. The negative binomial distribution nests the Poisson distribution (which forces the mean and variance to be equal) and is completely described by its mean and variance.

$$m(X_i, \beta) = exp(X_i \times \beta). \tag{6}$$

The variance of the distribution is then assumed to take the form:

$$Var(y_i|X_i) = m(X_i, \beta) + k \times m(X_i, \beta)^2. \tag{7}$$

Finally, the parameters are estimated by parametrically maximizing the log-likelihood of the model which is given by

$$l_i(\beta, k) = k^{-1}log\left(\frac{k^{-1}}{k^{-1} + m(X_i, \beta)}\right) + y_i \times log\left(\frac{m(X_i, \beta)}{k^{-1} + m(X_i, \beta)}\right) + log\left(\Gamma(y_i + k^{-1})/\Gamma(k^{-1})\right) \tag{8}$$

where $\Gamma()$ is the gamma function defined for $r > 0$ by $\Gamma(r) = \int_0^\infty z^{r-1}exp(-z)dz$ (Wooldridge 2002).

Autocorrelation is incorporated by including eight lagged values of the dependent variable (number of trades in a given category) in the vector of explanatory variables, X_i. Exchange listing, ADV, year-end price and market cap are also included in X_i. We test for the significance of the lagged dependent variables by using likelihood ratio tests. Each variable is tested by comparing the likelihood of the observed data with and without the lagged value included. The significance of a lagged dependent variable indicates that the number of trades on a given day is not independent of the history, i.e. there is positive trade clustering.

We estimate the autoregressive negative binomial regressions for four different dependent variables: all trades, order-displaying venue trades, Liquidnet Classic trades and other dark pool trades. The results are presented in Table 6. We report only the first four lagged dependent variables in each model. Likelihood ratio tests show statistically significant improvement in model fit for each set of lagged dependent variables.[29] The Chi-Squared test

[29]The third and fourth lags are not significant when modelling the expected number of order-displaying venue trades.

statistics indicate that these lags are important additions in each model, allowing us to reject the null hypothesis there is not clustering in large trades, at dark pools or otherwise. The coefficients estimated in the model represent the elasticity of $E(y|X)$ with respect to X_j. It is hard to interpret the elasticity as trades only occur in a discrete fashion and are not continuous, but the magnitudes of the coefficients are economically significant. For example, a 1% increase in yesterday's number of Liquidnet Classic trades would increase the expectation of the number of Liquidnet Classic trades tomorrow by 0.76%.

Trade clustering appears to be more prominent at Liquidnet Classic than at other dark pools or order-displaying venues. Table 6 shows that the lagged dependent variables predicting large trades at Liquidnet Classic are highly significant and add considerably to the predictive power of the model. The same is true for most other dependent variables; however, the coefficients for Liquidnet Classic trades are the largest in every case. The three-standard-deviation confidence bands around each estimate show that, for each lag, the Liquidnet Classic model coefficients' lower bounds are above the upper bounds for the coefficients of every other model. For example, 3 standard deviations below the first lag estimate for the Liquidnet Classic model is 0.60, which is higher than 3 standard deviations above the estimates for every other model (0.17, 0.43, and 0.21 for all trades, order-displaying venue trades, and other dark pool trades, respectively). While this is not a direct statistical test of differences in the distribution, it seems clear that Liquidnet Classic experiences more clustering in large trades than order-displaying or other dark venues. We interpret this finding as support for Hypothesis 3.

6 Conclusion

We document evidence that a dark pool specifically designed to foster buy-side exclusivity exhibits statistical regularities consistent with higher execution quality for large trades. Specifically, our evidence suggests that large trades at an allegedly more exclusive dark pool exhibit patterns consistent with: (i) less serial correlation in returns; (ii) less pre-trade volume and volatility increase; and (iii) more large trade clustering across days. Such indications are consistent with less pre- and post-trade exploitation of order-flow information in a dark pool designed to foster buy-side exclusivity, and are not due to selection bias in trade difficulty. We conclude that it is important to consider, in empirical, theoretical and policy-oriented discussions, that not all dark pools are created equal. Venue design features related to exclusivity factor into dark pool performance.

References

Admati, Anat R. & Paul Pfleiderer. 1988. "A Theory of Intraday Patterns: Volume and Price Variability." *The Review of Financial Studies* 1(1):3–40.

Bessembinder, Hendrik & Kumar Venkataraman. 2004. "Does an electronic stock exchange need an upstairs market?" *Journal of Financial Economics* 73:3–36.

Boehmer, Ekkehart. 2005. "Dimensions of execution quality: recent evidence for US equity markets." *Journal of Financial Economics* 78:553–582.

Buti, Sabrina, Barbara Rindi & Ingrid M. Werner. 2010. "Diving into Dark Pools." *Working Paper, University of Toronto* .

Buti, Sabrina, Barbara Rindi & Ingrid M. Werner. 2011. "Dark Pool Trading Strategies." *Working Paper, University of Toronto* .

Butler, Gabriel. 2007. "Liquidity Aggregation: What Institutional Investors Need to Know." *Journal of Trading* Spring:108–114.

Conrad, Jennifer, Kevin M. Johnson & Sunil Wahal. 2003. "Institutional Trading and Alternative Trading Systems." *Journal of Financial Economics* 70:99–134.

Davies, Ryan J. & Sang Soo Kim. 2008. "Using matched samples to test for differences in trade execution costs." *Journal of Financial Markets* 12(2):173–202.

Degryse, Hans, Mark Van Achter & Gunther Wuyts. 2009. "Dynamic Order Strategies with Competition between a Dealer Market and a Crossing Network." *Journal of Financial Economics* 91:319–338.

Domowitz, Ian, Ilya Finkelshteyn & Henry Yegerman. 2009. "Cul de Sacs and Highways: An Optical Tour of Dark Pool Trading Performance." *Journal of Trading* Winter:16–22.

Kratz, Peter & Torsten Schöneborn. 2009. "Optimal Liquidation in Dark Pools." *Working Paper, Humboldt Universitat zu Berlin* .

Mittal, Hitesh. 2008. "Are You Playing in a Toxic Dark Pool? A Guide to Preventing Information Leakage." *Journal of Trading* 3:20–33.

Nimalendran, Mahendrarajah & Sugata Ray. 2011. "Informed Trading in Dark Pools." *Working Paper, University of Florida* .

O'Hara, Maureen & Mao Ye. 2011. "Is Market Fragmentation Harming Market Quality?" *Journal of Financial Economics* 100:459–474.

Ready, Mark J. 2010. "Determinants of Volume in Dark Pools." *Working Paper, University of Wisconsin-Madison* .

Securities & Exchange Commission (2005). June 9, 2005. "Regulation NMS: Final Rules and Amendments to Joint Industry Plans." *available: http://www.sec.gov/rules/final/34-51808.pdf* .

Securities & Exchange Commission (2009). October 21, 2009. "Fact Sheet: Strengthening the Regulation of Dark Pools." *available: http://www.sec.gov/news/press/2009/2009-223-fs.htm* .

Weaver, Daniel. 2011. "Internalization and Market Quality in a Fragmented Market Structure." *Working Paper, Rutgers University* .

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data*. The MIT Press.

Ye, Mao. 2010. "Non-Execution and Market Share of Crossing Networks." *Working Paper, University of Illinois* .

Ye, Mao. 2011. "A Glimpse into the Dark: Price Formation, Transaction Costs, and Market Share in the Crossing Network." *Working Paper, University of Illinois* .

Zhu, Haoxiang. 2011. "Do Dark Pools Harm Price Discovery." *Working Paper, Stanford University* .

Appendix A: Institutional Setting

An Exclusive Dark Pool: Liquidnet Classic

As of 2011, Liquidnet offers various execution options to its clients. The majority of Liquidnet's volume is executed by clients using the "negotiated" option (also known as "Liquidnet Classic"). With this option, traders enter indications of interest (IOIs) through their Order Management System (OMS) interfaces. Suppose a trader places an IOI to buy one million shares of IBM at Liquidnet Classic. If another buy-side trader enters or has entered an IOI to sell IBM, both traders are informed that there is a potential counterparty for IBM. Each party is in "passive" mode. In order to execute shares with each other, one of the traders has to make the decision to change his status from "passive" to "active" mode. The other side then can send an invitation to trade, which enables a one-on-one negotiation of price and size. Specifically, each of the counterparties specifies his or her maximum size and negotiates price.

To preserve anonymity, negotiations occur through an interface provided by Liquidnet, which is independent of the trader's OMS. Once negotiations begin, typically they are completed within seconds. Therefore neither of the counterparties is expected to have an opportunity to front-run the other at other execution venues. In addition Liquidnet monitors market activity that occurs around the negotiation.

Traders can specify a minimum IOI size that they require from the other side in order to be displayed as a"passive" interest. A trader cannot observe that there is a potential (passive) counterparty until he places an IOI in that security.

Institutional investors have the option of participating only with other institutional investors. Brokers operate in a separate dark pool operated by Liquidnet, and members have the choice whether to dedicate a portion of their shares to that pool.

Typically OMS interfaces are designed to allow traders to easily split and modify orders across execution venues, brokers, and strategies (e.g., algorithms). For example, a trader who wants to buy one million shares of IBM that day can use his OMS to enter an IOI for 900,000 shares at Liquidnet Classic while using a broker's Volume Participation algorithm to execute the other 100,000 shares. Throughout the day, if IBM is not executed at Liquidnet Classic, the trader can continually decrease the size of the IOI at Liquidnet Classic while increasing the amount of IBM executed using other strategies.

Undisplayed Liquidity at Other Dark Pools and at Broker-Dealers' Desks

Among the dark pools that are less exclusive than just buy-side-only crossing are those owned and operated by broker-dealers as internal dark pools. Examples are Credit Suisse Crossfinder, UBS PIN, and Goldman Sachs Sigma X. O'Hara & Ye (2011) state that off-exchange volumes must now be reported through trade reporting facilities (TRFs). Weaver (2011) indicates that more than 90% of executions that are reported through the TRFs are either executed in dark pools or are internalized order flow. Furthermore, more than 75% of dark pool order flow is internalized order flow.

For many years institutional traders have had the option of executing blocks with their brokers' upstairs desks. Zhu (2011) states that although this type of broker-dealer internalization is not usually classified as dark pool trading, it is a source of undisplayed liquidity. Like the undisplayed liquidity made available by internalization dark pools, this type of undisplayed liquidity is less exclusive than buy-side-only dark pools.

Nimalendran & Ray (2011) provides evidence of how blurred the line has become between undisplayed liquidity in dark pools and undisplayed liquidity available at broker-dealers' desks. They analyze the executions of a dark pool operator that offers its clients a variety of options for executing orders. It operates a dark pool for manual negotiation for block trades. The dark pool operator also has a brokerage desk, which can execute blocks or work orders for clients.

Undisplayed Liquidity at Exchanges

Exchanges offer a variety of order types that allow traders to hide their willingness to trade. Undisclosed limit orders often maintain price priority but lose time priority to displayed limit orders at the same price. Undisclosed limit orders are not protected by the Regulation NMS Order Protection Rule. Exchanges typically allow all traders to use the undisplayed (dark) order types they offer. Zhu (2011) refers to hidden order liquidity on exchanges as the other source of undisplayed liquidity not normally classified as dark pool trading.

In order to execute against hidden liquidity, some algorithms are designed to send IOC (immediate-or-cancel) limit orders to "sweep" exchanges at or within the top-of-book quotes. Some algorithms "oversize" the order quantity (i.e., the order is larger than the size displayed at the top-of-book). As a result, executions may exceed the size quoted at the venue. Because the ratio of undisplayed to displayed size can be far greater than one, block executions can occur, even in thinly-quoted stocks.

In addition to allowing a variety of hidden limit order types, NYSE and NASDAQ operate crossing networks. Ye (2010) reports that exchange crossing network executions are not reported independently from their other exchange executions.

Table 1: Security sample selection criteria follows O'Hara & Ye (2011).

Criterion	**Nasdaq**	**NYSE**
All Securities in December 2010 CRSP File	2814	2450
No data listed in COMPUSTAT	-82	-41
<u>CRSP Filter</u>		
Non-common stock equities	-101	-353
Common stocks of non-US companies, closed-end funds, REITs, ADRs, ETFs	-138	-331
Dual Class Stock	-72	-112
<u>Volume and Price Filter</u>		
Mean daily volume $< 1{,}000$	0	0
Price $< \$5$	-390	-68
No trade of at least 50,000 shares in sample	-1249	-633
Final Sample	**782**	**912**

Table 2: Sample summary data. Trades are included in our sample if they do not take place within 5 minutes of another large trade and if they have trade activity in at least 75% of the thirty second periods both 10 minutes before and 10 minutes after the trade occurred. The first column shows the full sample while the remainder split the sample into mutually exclusive groups. Panel A displays unsegmented data. Panel B segments the data into terciles (by number of trades) based on trade difficulty. Panel C segments the data into terciles based on average daily volume.

Panel A: All Large Trades

	All Trades	Open Access Trades	Liquidnet Trades	Other Dark Pool Trades
Full Sample				
Number of Trades	34,898	5%	8%	87%
Average Trade Size	112,859	83,689	113,414	114,427
Average Trade Difficulty	4.0%	2.2%	6.8%	3.8%

Panel B: Terciles by Trade Difficulty

	All Trades	Open Access Trades	Liquidnet Trades	Other Dark Pool Trades
Top Tercile				
Number of Trades	10,745	2%	14%	84%
Average Trade Size	157,706	158,489	131,351	162,173
Average Trade Difficulty	10.0%	12.9%	11.0%	9.8%
Middle Tercile				
Number of Trades	15,007	2%	7%	91%
Average Trade Size	100,362	84,573	93,610	101,271
Average Trade Difficulty	1.8%	1.6%	2.0%	1.8%
Bottom Tercile				
Number of Trades	9,146	12%	2%	85%
Average Trade Size	80,675	68,244	81,897	82,423
Average Trade Difficulty	0.4%	0.3%	0.5%	0.4%

Panel C: Terciles by Average Daily Volume

	All Trades	Open Access Trades	Liquidnet Trades	Other Dark Pool Trades
Top Tercile				
Number of Trades	8,753	12%	3%	85%
Average Trade Size	136,491	75,475	225,633	142,352
Average Trade Difficulty	0.6%	0.2%	1.3%	0.7%
Middle Tercile				
Number of Trades	15,043	3%	7%	91%
Average Trade Size	112,377	93,503	121,268	112,313
Average Trade Difficulty	2.2%	1.8%	2.6%	2.2%
Bottom Tercile				
Number of Trades	11,102	2%	14%	84%
Average Trade Size	94,879	103,600	92,574	95,058
Average Trade Difficulty	9.0%	12.2%	10.2%	8.7%

Table 3: Non-parametric tests of abnormal return correlations surrounding large trades. Negative M-signs indicate negative correlation between the pre-trade one-minute and post-trade one-minute abnormal returns. P-values are shown for the null hypothesis that the M-sign is zero. The numbers between columns are the p-values from Wilcoxon rank-sum tests for the equality of the medians between two samples. This data set is generated by using the best match under the Davies & Kim (2008) methodology. Prices are measured based on the last trade price prior to the measurement time.

Dependent Variable:	0-to-1 Minute Post-Trade Abnormal Return				
Independent Variable:	1-to-0 Minute Pre-Trade Abnormal Return				
Panel A: Non-Overlapping Trades Sample					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Full Sample					
M-Sign	-54.5		-17.0		-782.5
P-Value (H0: Median = 0)	0.006		0.512		0.000
Sample Size	1,689		2,785		30,424
Wilcoxon Z-Score (Right > Left)		1.860		-2.335	
P-value (Equality of Medians)		0.063		0.020	
Panel B: Terciles by Trade Difficulty					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Top Tercile					
M-Sign	1.5		-20.0		-206.5
P-Value (H0: Median = 0)	0.886		0.292		0.000
Sample Size	227		1,530		8,988
Wilcoxon Z-Score (Right > Left)		-0.579		-0.937	
P-value (Equality of Medians)		0.562		0.349	
Middle Tercile					
M-Sign	-16.0		1.0		-340.0
P-Value (H0: Median = 0)	0.078		0.974		0.000
Sample Size	343		1,034		13,630
Wilcoxon Z-Score (Right > Left)		1.448		-2.068	
P-value (Equality of Medians)		0.148		0.039	
Bottom Tercile					
M-Sign	-40.0		2.0		-236.0
P-Value (H0: Median = 0)	0.014		0.834		0.000
Sample Size	1,119		221		7,806
Wilcoxon Z-Score (Right > Left)		0.914		-0.842	
P-value (Equality of Medians)		0.361		0.400	
Panel C: Terciles by Average Daily Volume					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Top Tercile					
M-Sign	-32.0		-7.5		-284.0
P-Value (H0: Median = 0)	0.042		0.328		0.000
Sample Size	1,046		224		7,483
Wilcoxon Z-Score (Right > Left)		0.353		-0.356	
P-value (Equality of Medians)		0.724		0.722	
Middle Tercile					
M-Sign	-25.5		7.0		-279.0
P-Value (H0: Median = 0)	0.010		0.667		0.000
Sample Size	417		984		13,642
Wilcoxon Z-Score (Right > Left)		1.881		-1.812	
P-value (Equality of Medians)		0.060		0.070	
Bottom Tercile					
M-Sign	3.0		-16.5		-219.5
P-Value (H0: Median = 0)	0.720		0.395		0.000
Sample Size	226		1,577		9,299
Wilcoxon Z-Score (Right > Left)		-0.580		-1.256	
P-value (Equality of Medians)		0.562		0.209	

Table 4: Non-parametric tests of abnormal volume prior to large trades. Positive M-signs indicate an abnormal increase in volume prior to the large trade. P-values are shown for the null hypothesis that the M-sign is zero. The numbers between columns are the p-values from Wilcoxon rank-sum tests for the equality of the medians between two samples. This data set is generated by using the best match under the Davies & Kim (2008) methodology.

Dependent Variable:	1-to-0 Minute Pre-Trade Volume				
Independent Variable:	2-to-1 Minute Pre-Trade Volume				
Panel A: Non-Overlapping Trades Sample					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Full Sample					
M-Sign	77.0		25.5		986.5
P-Value (H0: Median = 0)	0.000		0.315		0.000
Sample Size	1,619		2,480		28,508
Wilcoxon Z-Score (Right > Left)		-2.456		2.752	
P-value (Equality of Medians)		0.014		0.006	
Panel B: Terciles by Trade Difficulty					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Top Tercile					
M-Sign	1.0		23.0		178.5
P-Value (H0: Median = 0)	0.941		0.207		0.000
Sample Size	180		1,275		7,755
Wilcoxon Z-Score (Right > Left)		-0.459		0.902	
P-value (Equality of Medians)		0.646		0.367	
Middle Tercile					
M-Sign	3.5		3.0		469.5
P-Value (H0: Median = 0)	0.740		0.874		0.000
Sample Size	328		988		13,024
Wilcoxon Z-Score (Right > Left)		-0.487		2.084	
P-value (Equality of Medians)		0.626		0.037	
Bottom Tercile					
M-Sign	72.5		-0.5		338.5
P-Value (H0: Median = 0)	0.000		1.000		0.000
Sample Size	1,111		217		7,729
Wilcoxon Z-Score (Right > Left)		-1.946		1.687	
P-value (Equality of Medians)		0.052		0.092	
Panel C: Terciles by Average Daily Volume					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Top Tercile					
M-Sign	68.0		-2.5		234.5
P-Value (H0: Median = 0)	0.000		0.789		0.000
Sample Size	1,038		223		7,431
Wilcoxon Z-Score (Right > Left)		-1.484		0.940	
P-value (Equality of Medians)		0.138		0.347	
Middle Tercile					
M-Sign	8.5		2.5		497.0
P-Value (H0: Median = 0)	0.427		0.897		0.000
Sample Size	405		947		13,160
Wilcoxon Z-Score (Right > Left)		-1.239		2.292	
P-value (Equality of Medians)		0.215		0.022	
Bottom Tercile					
M-Sign	0.5		25.5		255.0
P-Value (H0: Median = 0)	1.000		0.167		0.000
Sample Size	176		1,310		7,917
Wilcoxon Z-Score (Right > Left)		-0.264		1.530	
P-value (Equality of Medians)		0.792		0.126	

Table 5: Non-parametric tests of abnormal increases in volatility, proxied by trading range, prior to large trades. The abnormal volatility measure compares the trading range in the minute prior to the large trade to the trading range of the previous minute. Positive M-signs indicate an abnormal increase in volatility prior to the large trade. P-values are shown for the null hypothesis that the M-sign is zero. The numbers between columns are the p-values from Wilcoxon rank-sum tests for the equality of the medians between two samples. This data set is generated by using the best match under the Davies & Kim (2008) methodology.

Dependent Variable:	1-to-0 Minute Pre-Trade Range				
Independent Variable:	2-to-1 Minute Pre-Trade Range				
Panel A: Non-Overlapping Trades Sample					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Full Sample					
M-Sign	8.5		-3.0		394.5
P-Value (H0: Median = 0)	0.660		0.911		0.000
Sample Size	1,506		2,083		25,120
Wilcoxon Z-Score (Right > Left)		-0.969		2.145	
P-value (Equality of Medians)		0.332		0.032	
Panel B: Terciles by Trade Difficulty					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Top Tercile					
M-Sign	6.5		-15.0		73.5
P-Value (H0: Median = 0)	0.305		0.350		0.054
Sample Size	141		1,001		6,048
Wilcoxon Z-Score (Right > Left)		-1.441		2.014	
P-value (Equality of Medians)		0.150		0.044	
Middle Tercile					
M-Sign	-4.5		-1.0		131.0
P-Value (H0: Median = 0)	0.622		0.973		0.013
Sample Size	285		871		11,593
Wilcoxon Z-Score (Right > Left)		0.140		0.943	
P-value (Equality of Medians)		0.889		0.346	
Bottom Tercile					
M-Sign	6.5		13.0		190.0
P-Value (H0: Median = 0)	0.693		0.072		0.000
Sample Size	1,080		211		7,479
Wilcoxon Z-Score (Right > Left)		0.787		-0.256	
P-value (Equality of Medians)		0.431		0.798	
Panel C: Terciles by Average Daily Volume					
	Open Access Trades		Liquidnet Trades		Other Dark Pool Trades
Top Tercile					
M-Sign	6.0		6.5		175.0
P-Value (H0: Median = 0)	0.709		0.390		0.000
Sample Size	1,017		217		7,235
Wilcoxon Z-Score (Right > Left)		0.706		-0.173	
P-value (Equality of Medians)		0.480		0.862	
Middle Tercile					
M-Sign	0.5		-13.5		123.0
P-Value (H0: Median = 0)	1.000		0.364		0.021
Sample Size	349		859		11,845
Wilcoxon Z-Score (Right > Left)		-1.192		2.490	
P-value (Equality of Medians)		0.233		0.013	
Bottom Tercile					
M-Sign	2.0		4.0		96.5
P-Value (H0: Median = 0)	0.796		0.823		0.012
Sample Size	140		1,007		6,040
Wilcoxon Z-Score (Right > Left)		-0.821		0.838	
P-value (Equality of Medians)		0.412		0.402	

Table 6: Clustering of trades across days. This table shows the coefficients of the lagged dependent variables estimated in a negative binomial regression model used to model the number of trades per day of each category. For instance, the third data column models only Liquidnet Classic trades using information from only Liquidnet Classic trades. Only the first four lagged coefficients are displayed. Additional lags and control variable coefficients are suppressed. All lagged coefficients displayed significantly add to the explanatory power of the model (as measured by increase in log likelihood) except for the third and fourth lags of the model predicting open access trades.

	All Trades	Open Access Trades	Liquidnet Trades	Other Dark Pool or ATS Trades
1st Lag Coefficient	0.16	0.32	0.76	0.19
Standard Error	0.01	0.04	0.05	0.01
99% Confidence Interval	(0.14 , 0.17)	(0.21 , 0.43)	(0.60 , 0.92)	(0.17 , 0.21)
2nd Lag Coefficient	0.06	0.06	0.45	0.08
Standard Error	0.01	0.03	0.05	0.01
99% Confidence Interval	(0.04 , 0.07)	(-0.03 , 0.16)	(0.29 , 0.60)	(0.06 , 0.09)
3rd Lag Coefficient	0.05	0.03	0.38	0.06
Standard Error	0.01	0.03	0.05	0.01
99% Confidence Interval	(0.03 , 0.06)	(-0.06 , 0.12)	(0.22 , 0.54)	(0.03 , 0.07)
4th Lag Coefficient	0.04	0.05	0.30	0.06
Standard Error	0.01	0.03	0.05	0.01
99% Confidence Interval	(0.02 , 0.05)	(-0.04 , 0.15)	(0.15 , 0.45)	(0.03 , 0.07)

Figure 1: Large Trade Execution Times (by Volume). The graphs show the percentage share of volume for each time period based on our full sample (after cuts in Table 1) with all trades in the first or last 15 seconds of the trading day removed in order to minimize inclusion of mislabaled open and close prints in the TAQ data.

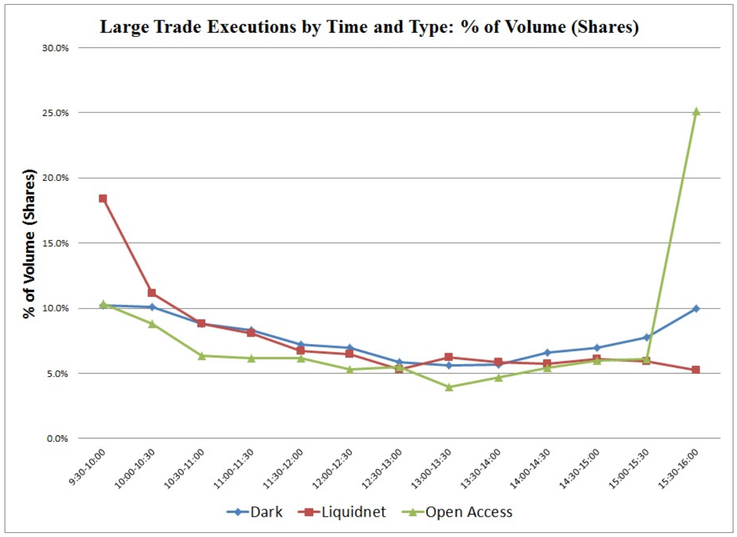

Figure 2: Liquidnet Classic Pre-Trade and Post-Trade Abnormal Returns. Pre-trade abnormal returns are on the x-axis and post-trade abnormal returns are on the y-axis. Abnormal returns are measured relative to a matched control firm's return over the same time period. The red line is a simple univariate linear regression of the post-trade abnormal returns on the pre-trade abnormal returns.

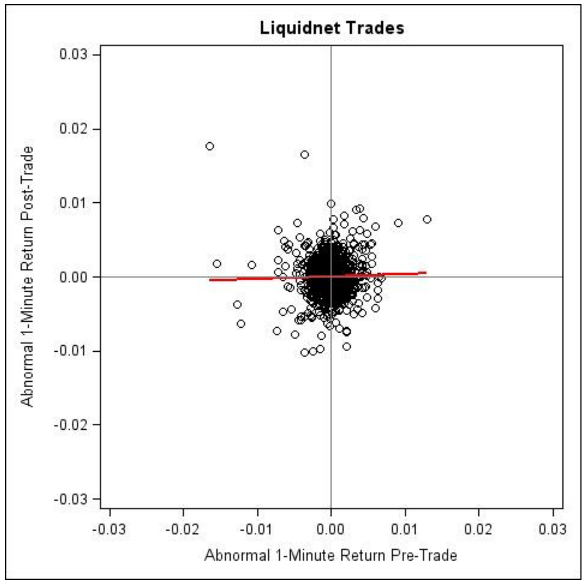

Figure 3: Open-Access Pre-Trade and Post-Trade Abnormal Returns. Open-access trades occur on an exchange whose trading membership is not restricted. Typical examples include NYSE, NASDAQ and BATS. Pre-trade abnormal returns are on the x-axis and post-trade abnormal returns are on the y-axis. Abnormal returns are measured relative to a matched control firm's return over the same time period. The red line is a simple univariate linear regression of the post-trade abnormal returns on the pre-trade abnormal returns.

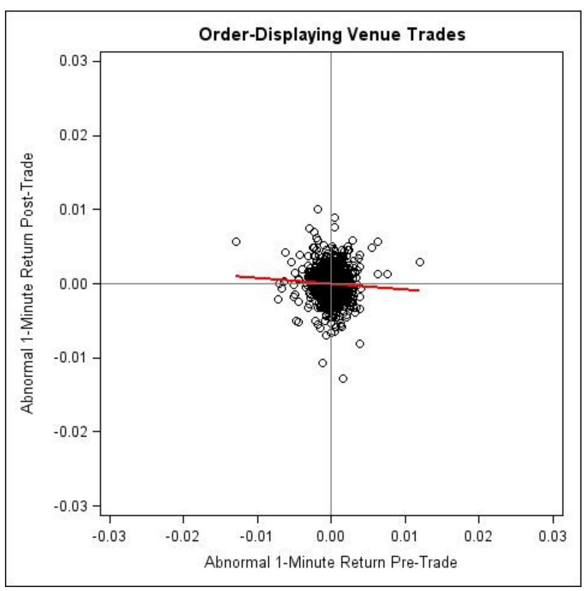

Figure 4: Dark Pool (without Liquidnet Classic) Pre-Trade and Post-Trade Abnormal Returns. Dark pool executions are considered those whose volume is reported to the trade reporting facility or the alternative display facility. Pre-trade abnormal returns are on the x-axis and post-trade abnormal returns are on the y-axis. Abnormal returns are measured relative to a matched control firm's return over the same time period. The red line is a simple univariate linear regression of the post-trade abnormal returns on the pre-trade abnormal returns.