Regulating Dark Trading: Order Flow Segmentation and Market Quality

Carole Comerton-Forde[a], Katya Malinova[b] and Andreas Park[b]

[a] University of Melbourne
[b] University of Toronto and Copenhagen Business School

March 15, 2015

-- PLEASE DO NOT CIRCULATE --

This paper is based on the research report title "The Impact of the Dark Trading Rules" that we prepared for the Investment Industry Regulatory Organization of Canada (IIROC). We thank IIROC and in particular Victoria Pinnington and Helen Hogarth. We also thank Qizheng (Alan) Yuan for research assistance. Comerton-Forde is an economic consultant on market structure for the Australian Securities and Investments Commission.

Corresponding author: Andreas Park. Email: andreas.park@utoronto.ca

Abstract

Using proprietary trader-level data we examine the impact of the introduction of a rule in the Canadian equities market that required dark orders to be executed at prices better than national best bid and offer. The rule lead to a 42% reduction in dark trading volumes, but did not impact aggregate market quality. The impact of the rule was not uniform across trading venues. We show that the impact of the rule depends critically on the nature of liquidity provision in the affected venues. The rule substantially reduced two-sided liquidity provision in the dark. We also provide evidence that the segregation of retail order flow may harm lit market liquidity.

Keywords: Dark trading, trade-at rule, regulation, high-frequency trading, retail investors, minimum price improvement rule, internalization.

JEL Classification: G14.

Dark trading allows traders to hide their trading interest from the market before a trade is executed. Over the last decade, as markets became more electronic and as advanced data analysis tools have made it easier to detect trading intentions, traders have increasingly sought to execute their orders without displaying them publicly pre-trade. The desire to trade in the dark has been facilitated by both specialized venues, known as "dark pools", and "lit" markets that accept fully hidden or dark orders.

The emergence of new dark trading mechanisms and the substantial increase in dark trading volumes[1] has captured worldwide regulatory attention. For example, in the U.S., the Securities and Exchange Commission requested public comments (SEC (2010)) on "the effect of undisplayed liquidity on order execution quality, the effect of undisplayed liquidity on public price discovery, and fair access to sources of undisplayed liquidity", and SEC Chairman Mary Jo White reiterated the need to continue to examine dark venues in her recent speech on June 5, 2014.

Regulators in Canada had similar concerns, and after a public consultation, introduced new dark trading rules on October 15, 2012. The most substantive change was the introduction of a requirement for trades in the dark to offer minimum price improvement of at least one trading increment (1 cent for our sample) better than the national best bid and offer (NBBO), or half an increment if the bid-ask spread is one trading increment. This rule is colloquially referred to as the price improvement rule. In the U.S. this type of rule is referred to as a "trade-at" rule.[2]

We use proprietary trader-level data provided by the Investment Industry Regulatory Organization of Canada (IIROC) to examine the impact of the introduction of the price improvement rule. The dataset contains information on all trades, quotes, orders, and cancellations for all marketplaces in Canada. Each marketplace is separately, but anonymously identified, and we label them marketplace A to H. These marketplaces exhibit substantial variation in market structure. The data also contains detailed information about the (masked) individual trader accounts. This enables us to classify traders into four groups, based on their trading characteristics: *retail* (based on usage of a trading tool available only to retail investors), *high-frequency* (by reaction time to a regular, scheduled public announcement), *buy-side institutions* (by the accumulation of large positions), and *other*. These data therefore allow us to examine the impact of the rule change on the heterogeneous marketplaces and trader types.

The price improvement rule had a dramatic impact on trading in Canada. In the weeks following October 15, 2012 the share of dark activity declined sharply, from 9.3% to 5.4% of dollar trading volume (excluding pre-arranged block trades). Before the change in regulations, about three quarters of all dark dollar volume was executed in two dark pools. After the change one of these dark pools, which we refer to as market Ad, experienced a significant decline in volume from 4.6% to 0.8%, whereas volume on the other dark pool, which

[1] For instance, in the U.S., dark trading has grown from 17% of the U.S. trading volume in July 2008 to 37% in June 2014 (Rosenblatt Securities).
[2] As part of the planned "Tick Size Pilot" program, the SEC will implement a Trade-At Prohibition for securities in the Test Group Three, subject to exemptions.

we refer to as market D, remained unchanged at 2.5%. We attribute the differential impact of the rule to differences in the nature of liquidity provision and types of traders in the two dark pools.[3]

To assess the nature of liquidity provision in each pool, we develop an imbalance score. For each trader posting limit orders, we measure whether they are providing one-sided or two-sided d. One-sided liquidity likely captures natural liquidity providers, while two-sided liquidity captures intermediaries or informal market markets. We observe that liquidity on market Ad, which saw a decline in trading volume, was provided primarily by traders posting two-sided liquidity. In contrast, liquidity provision in market D, which was largely unaffected by the rule, was supplied mostly by traders who posted one-sided orders.

We also find that the imbalance score is significant in predicting whether or not a trader reduces their liquidity provision in the dark after the change in regulation, whereas the trader types (HFT, retail, buy-side, or other) are not significant after controlling for the imbalance score. We attribute the differential impact of the price improvement rule on the two dark markets to pre-rule change differences in the nature of liquidity provision.

Markets Ad and D also exhibit substantial variation in the types of participating traders. Market Ad offered its liquidity providers an option to interact with marketable orders exclusively from retail investors (the order routing choice for retail orders remained with the broker). Prior to the rule change, 99.9% of aggressive volume in market Ad stemmed from retail investors, and aggressive retail flow executed in Ad accounted for 27.6% of the total retail trading volume in the Canadian market. This order flow was typically matched against informal market makers, including HFT. After the rule change, these marketable retail orders were routed by brokers to the lit market that (at the time) offered the lowest taker fee. We refer to this market as Al. Volume of marketable retail orders as a fraction of market Al's total volume increased from 15% to 30%. We observe a simultaneous significant improvement in liquidity on this market, as measured by a 15% increase in posted dollar-depth at the best prices. Cross-sectional analysis further reveals that the trading volume in Ad before October 15, 2012 (which necessarily involved a retail marketable order) predicts the change in depth on the lit market Al after that date.

Combining the segregation of retail order flow with the fact that most of the liquidity was provided by two-sided liquidity providers, the operation of market Ad prior to October 15, 2012 arguably resembled liquidity provision to retail brokers by wholesalers the U.S. The decline in dark trading in Canada was primarily driven by the decline in dark trading in market Ad, and therefore our results shed light on the potential impact of the introduction of a trade-at rule in the US market. If the economics of market making in the U.S. are similar to the economics for liquidity providers in market Ad, the results suggest that U.S wholesale market makers may also reduce liquidity supply in the presence of a trade-at rule.

We also examine the impact of the rule on intra-day returns and the implementation shortfall by trader group.. We compute intra-day return per trader group, as the intra-day profits from buying and selling, marking the end-of-day inventory position to the closing price, expressed as a fraction of the group's total daily dollar volume.

[3] We are not able to determine why the two dark markets exhibited ex ante differences in liquidity provision or why traders that posted one-sided orders concentrated on a single market. A possible explanation may stem from coordination and network externalities considerations.

We find that intra-day returns for retail traders weakly declined after the change in regulations, and that the decline is more pronounced if we account for trading venues' maker-taker fees. Since most brokers do not pass exchange trading fees to retail customers on a trade-by-trade basis (but instead charge flat commissions), the change in regulations may have led to an increase in costs for retail brokers.

The implementation shortfall measure is particularly useful for large orders that are split into smaller orders, and it is thus most applicable to assess trading costs for buy-side institutional traders. This measure compares the actual cost of a (large) order with the hypothetical cost that would have obtained had the order been filled at the time when the first small order cleared. We do not find any evidence that the change affected the buy-side's implementation shortfall.

To assess the impact of the decline in dark trading on the probability of execution of lit orders, we compute the ratio of passive, lit market volume to all non-dark order volume.[4] While we observe that fill rates for passive orders submitted by HFTs increase in several lit markets, the fill rates for buy-side institutions decline. In contrast, buy-side traders' fill rates in the dark markets increase. We attribute changes in fill rates for buy-side institutions to changes in HFT behavior. The introduction of the price improvement rule made it impossible to "earn the spread" by trading on both sides in a single dark venue, therefore reducing incentives to make a market in dark pools. Since a number of major HFT firms self-identify as liquidity providers,[5] we would expect HFTs to reduce their liquidity providing activities in the dark following the introduction of the price improvement rule. Our findings on HFT participation across different venues confirm this intuition. A decrease in liquidity provision by HFTs in the dark and an increase in liquidity provision by HFTs on lit venues arguably makes it easier for slower traders to fill their passive orders in the dark and more competitive to fill such orders on lit venues, consistent with our findings on buy-side fill rates.

I. The Institutional Setting

A. Core rules governing trading in Canada

The Toronto Stock Exchange (TSX) is the primary listing venue for large companies in Canada.[6] Like other major markets around the world, trading in TSX-listed stocks is fragmented across multiple exchanges and Alternative Trading Systems (ATS). Securities trading and the activities of market participants in Canada are regulated by the Investment Industry Regulatory Organization of Canada (IIROC) and are governed by the Universal Market Integrity Rules (UMIR).

Most of the core elements of the UMIR are similar to those governing trading in the U.S. equities markets. Brokers and marketplaces are required to respect the order protection rule, which mandates that orders must be

[4] The data does not explicitly identify non-marketable orders.
[5] Two well-known HFT firms, Virtu and Getco, held a session at the 2013 TMX Trading Conference titled "Meet the Liquidity Providers".
[6] Small and mid-cap companies are typically listed on the TSX Venture exchange.

routed to the marketplace with the best-priced orders available on lit markets. Brokers are also subject to obligations regarding best execution for client orders. [7]

In the context of our study, there are three critical differences between trading rules in the U.S and Canada. First, the order protection rule in Canada applies to the whole-of-book rather than the top-of-book as is the case in the US. Second, Canada also imposes a strict version of an order *exposure* rule,[8] with few exceptions. This rule requires that client orders below a certain size be immediately sent to a marketplace that publicly displays prices. This rule severely limits the practice of broker internalization, which occurs when a broker trades against their customer's order instead of sending the order to a public marketplace, and the practice of selling retail orders to market makers.[9] Third, unlike the US, Canadian marketplaces are allowed to offer broker-preferencing on the market's order book. This practice allows incoming orders to a marketplace to match with other orders from the same broker-dealer ahead of similarly priced orders from other broker-dealers, without regard to time priority. To take advantage of broker-preferencing, brokers must elect to publicly display broker IDs when submitting their orders. [10]

Dark trading in Canada is subject to restrictions that are similar to rules in other jurisdictions. First, consistent with the principles set out by the International Organization of Securities Commissions (IOSCO), dark orders have lower execution priority than visible orders at the same price.[11] All trades in Canada, including dark trades, are subject to full and immediate post-trade transparency.

Second, the order exposure rule dictates that passive client orders that are below a certain size can only be posted as dark if the client explicitly directs the broker to so do.[12] It is our understanding that during our sample period most brokers did not offer (passive) dark trading as an option to their retail customers; the order exposure rule does not prohibit sending clients' marketable orders to dark venues. The change in dark trading regulations on October 15, 2012, which we describe below, introduced a price improvement rule, which required that dark orders provide meaningful price improvement over the NBBO to marketable orders that were subject to the order exposure rule.

Finally, trades may be pre-arranged off-exchange, before entering orders on a public marketplace, but these trades must still be executed on a public marketplace, respecting all the applicable rules. Pre-arranged trades

[7] National Instrument 23-101 formulates the order-protection rule; UMIR 5.1 outlines the framework for best execution practices. The order-protection rule differs slightly from its U.S. counterpart, but we believe that the differences are immaterial for our analysis.
[8] See UMIR 6.3 and related guidance notes.
[9] Battalio, Corwin, and Jennings (2014) report that U.S. brokers systematically sell all of their retail marketable orders to market makers (wholesalers). It is our understanding that Canadian broker-dealers did not follow this practice during our sample period, although some entered or considered entering into such arrangements with U.S. wholesalers later. In late 2014, IIROC published a guidance note clarifying that U.S. wholesalers do not satisfy the definition of a regulated public market, effectively banning the practice of selling Canadian retail order flow to the U.S. See also *http://www.osc.gov.on.ca/en/NewsEvents_nr_20141215_concerns-routing-retail-equity-orders.htm* .
[10] Broker-preferencing is subject to several restrictions, e.g., UMIR 5.3 (Client Priority) restricts entering non-client orders at the same or better prices as client orders.
[11] See IOSCO "Principles on Dark Liquidity" http://www.iosco.org/news/pdf/IOSCONEWS210.pdf
[12] The order exposure rule applies to orders that are received by the participant (e.g., the broker). It is the obligation of the participant to ensure compliance with the rule when the received order is at or below 50 standard trading units (for securities in our sample, 5,000 shares); there is also an exemption for orders of more than $100,000 in value.

thus typically involve orders that are large enough so that they were not subject to the order exposure rule or to the new price improvement rule. We omit such trades from our analysis.

B. Regulation Change

On October 15, 2012, IIROC implemented two changes to its rules and regulations.

First, IIROC amended its rules on dark liquidity, and, in particular, introduced an additional rule regarding the entry and exposure of orders. This new rule, UMIR 6.6, titled "Provision of Price Improvement by a Dark Order", requires that marketable orders that are at or below 50 standard trading units or $100,000 in value and that trade against a non-transparent order must be provided with a price improvement upon the national best bid and offer prices by at least one trading increment, or by half an increment if the bid-ask spread is one trading increment. For securities that are priced above $1, the trading increment is 1 cent and a trading unit is 100 shares. The rule thus mandates that dark orders offer a price that is 1 cent better (1/2 cent for 1 cent bid-ask spreads) than the best price posted across the visible marketplaces. IIROC further clarified that this rule does not apply to the hidden portion of so-called iceberg orders.[13] In what follows we will refer to this new minimum price improvement rule as *MPIR*. In the U.S., rules that mandate price improvement are referred to as *trade-at* rules.

Second, IIROC repealed a set of short sell restrictions for non-crosslisted securities. This rule change did not affect cross-listed securities because these were already exempt from the repealed restrictions.

This paper examines the impact of the change in dark liquidity regulations. We therefore consider only cross-listed securities to ensure that our analysis is not confounded by changes in the short selling rules.

C. Marketplaces and their trading rules before and after to the change in dark liquidity rules.

The data in our sample contains observations for eight marketplaces. These marketplaces are separately, but anonymously identified in our data, and we label them as marketplaces A to H. During our sample period (from August 27 to November 30, 2012), marketplaces A, B, C, and D account for 20.5%, 56.3%, 16.4%, and 3.3% of the dollar volume traded, respectively. Marketplaces E to H jointly account for less than 3.5% market share. We therefore exclude marketplaces E to H from most of our analysis.

Below we provide a detailed explanation of the dark trading features of marketplaces A to D, including details of how these marketplaces were impacted by the introduction of the dark liquidity rules.

Marketplace A operates a public limit order book, which we refer to as market Al, and a dark pool facility, which we refer to as market Ad. Al allows lit and partially hidden (iceberg) limit orders. Broker preferencing is allowed provided the broker chooses to publicly display its broker ID when submitting the order. In the dark pool Ad, traders interact using two types of orders: dark orders and seek dark liquidity (SDL) orders. Dark orders are limit orders that remain in the dark pool facility until they are executed or cancelled. SDL orders are liquidity taking: an SDL order that is not filled immediately by a resting dark order cannot remain in Ad. Critically for our analysis, dark limit orders are available to all market participants, whereas SDL orders are available exclusively to retail investors.

[13] An iceberg or reserve order is an order that displays only a portion of its full size.

Dark orders that are posted in Ad must be priced relative to the national best bid and offer (NBBO), and traders are required to offer price improvement over the NBBO. Prior to the implementation of the dark liquidity rules on October 15, 2012, traders had a choice between offering price improvement of 10% or 50% of the prevailing NBBO. After October 15, 2012, the price improvement was exogenously set at 50% of the spread. Dark orders that offer a 10% improvement are matched continuously against incoming SDL orders. Dark orders that offer 50% improvement may choose to interact (i) only with incoming SDL orders, (ii) only with other dark orders, whether resting or incoming, or (iii) with both SDL and dark orders.

On the same date the dark liquidity rules were altered, marketplace A also amended the way in which SDL orders operated. Prior to October 15, 2012, an SDL order that did not find a match with a dark order in marketplace Ad would be routed to other marketplaces according to the broker's instructions. After October 15, 2012, SDL orders were automatically routed to the public limit order book for marketplace Al, provided that Al was quoting the best price on the relevant side of the market, and were only routed to other marketplaces if an execution was not found on marketplace Al. Although we cannot separately assess the impact of the change in functionality of this order type, we note that marketplace Al had the lowest liquidity taking fees among the major lit marketplaces (see Table I), and it was therefore arguably most attractive for liquidity taking orders. We would thus expect that retail brokers would prefer to route orders to market Al regardless.

Marketplace B is a lit market that operates as a public limit order book. Broker preferencing is allowed provided the broker chooses to publicly display its broker ID when submitting the order. Traders may post lit, partially hidden (iceberg), and fully hidden orders. Fully hidden orders may be posted as "mid-point" orders, which are pegged to execute at the floating midpoint of the NBBO, or they may be posted as undisplayed limit orders. Therefore, marketplace B already complied with the dark liquidity rules, before they were introduced. As a result, marketplace B is not directly impacted by the rule change.

Marketplace C is a lit market that operates as a public limit order book. Like marketplace B, it allows lit, iceberg, and fully hidden orders, which may be pegged to the midpoint. Like marketplace B, marketplace C also complied with the new dark liquidity rules before they were introduced. Marketplace C does not offer broker-preferencing.

Marketplace D is a dark pool that allows traders to interact using two types of orders. These order types are similar to those in marketplace Ad, but with no restrictions on the type of traders that can use these orders. First, traders may submit passive dark orders that remain in the dark pool until they are executed or cancelled. Second, traders may submit aggressive, liquidity taking orders that are either executed immediately against a passive dark order or cancelled. Dark passive orders are priced relative to the NBBO and offer price improvement on the NBBO. Prior to October 15, 2012, traders had a choice to offer price improvement of either 20% or 50% of the NBBO. After October 15, 2012, Market D mandated a 50% price improvement so that all trades occurred at the midpoint of the NBBO. All dark orders continuously trade against the incoming IOC orders. Dark orders that offer 50% price improvement may additionally interact with each other, according to a periodic matching mechanism.

Marketplaces E and F operate as public limit order books, and **marketplaces G and H** are dark pools. During our sample period, marketplace G is an institutional-only venue, marketplace H offers periodic matching with 1-second random NBBO prices.

II. Data and Sample

Data. The data for this study is provided by the Investment Industry Regulatory Organization of Canada (IIROC).[14] The dataset contains detailed records on all trades, orders, order cancellations, order amendments, and updates to marketplaces' best bid and offer quotes from IIROC's real-time surveillance system, for all trading on all regulated Canadian marketplaces. Each order-related record includes, in particular:

- The marketplace where the order was sent (masked).
- Size, price, and the direction (buy or sell) of an order.
- Broker ID (masked), user ID (masked), and account type (e.g., specialist, client, options-trader, or inventory).
- Other characteristics, including the duration of an order (for instance, good-till-cancel or immediate-or-cancel), whether an order was transparent or non-transparent, whether the order was a seek-dark-liquidity order, and a unique identifier for each order.

For trades, the data additionally specifies the aggressive and passive (liquidity-providing) side of a trade. The data also identifies intentional broker-crosses—these trades are usually arranged off-exchange but they must be executed on a public marketplace. The information for marketplaces, brokers and users is masked in the sense that IIROC provides a scrambled identifier. The masking is applied consistently so that the same marketplace, broker and user is always assigned the same identifier.

Marketplaces' time-stamps are generally reported with millisecond precision, although marketplace B reported only at hundredth-of-a-second precision until October 15, 2012. Brogaard, Hendershott, and Riordan (2014), Korajczyk and Murphy (2014) and IIROC (2014) contain further information of the data.

Sample. We base our analysis on the period from August 27 to November 30, 2012, (i.e. seven weeks before and after the event date, October 15, 2012). We end the sample on November 30 to avoid confounding effects that may stem from a connection speed update implemented by the primary market, the TSX, on December 1, 2012. We restrict attention to cross-listed securities because on the event date, October 15, 2012, IIROC changed the rules regarding short-selling for non-crosslisted securities.

We base our analysis on "highly-liquid" securities, as defined by IIROC, that are cross-listed in U.S. markets. Loosely, a security qualifies as highly-liquid for a given day if over a 60-day period it traded more than 100 times per trading day and had an average trading value of at least \$1M.[15] IIROC compiles a list of highly-liquid

[14] IIROC is a self-regulatory organization that oversees dealers and trading activities and performs real-time market surveillance.
[15] For further details see IIROCs definition on http://www.iirocca/industry/rulebook/Pages/Highly-liquidstocks.aspx.

securities daily; we include a security in our sample if that security is on the list of highly liquid securities at the end of each month in our sample period. We determine the security's cross-listing status from the monthly TSX e-Review publication. We identify 334 securities that are in the list of frequently traded securities throughout our sample period; 92 of these securities are highly-liquid and cross-listed with a U.S. market throughout our sample period.

Outliers. We eliminated four days from our sample: October 29 and 30, when U.S. markets were closed because of Hurricane Sandy, and November 22 and 23, U.S. Thanksgiving and Black Friday. We further observed an extraordinary number of order submissions (80,000+) by a single trader on a single venue on two days for a single, very large order size in a single, relatively low-volume security. These days were not marked by high order or trading activity levels for this security, and the trader displayed no noteworthy characteristics other than on these two days. We thus eliminated the observations for this security on these two days from our sample.

III. Impact of the Minimum Price Improvement Rule on Market Quality

A. What was the impact of the price improvement rule on dark trading volume?

We measure the impact of the introduction of the minimum price improvement rule (MPIR) on dark trading in two ways. First, we compute the dollar trading volume that involves a dark order on the passive side of the trade, as a fraction of the total dollar trading volume. We refer to this as *PassiveDarkVolume*. Second, we compute the share of volume of dark orders, as a fraction of the share volume of all orders. We refer to this as *DarkOrderVolume.*

A.1. What was the impact on dark trading in aggregate?

Figure 1 plots *PassiveDarkVolume* for our sample securities across all venues. The figure shows that there is a significant drop in dark trading volume following the introduction of the dark liquidity rules. Table II reports summary statistics on *PassiveDarkVolume* and *DarkOrderVolume*, and it illustrates a decline in both measures following the introduction of MPIR. *PassiveDarkVolume* declines from 9.3% to 5.4%, and *DarkOrderVolume declines from 17.2% to 11.9%.*

We formally analyze the impact of the MPIR on dark trading by estimating the following linear panel specification:

$$DV_{it} = \alpha \times MPIR_t + \beta \times VIX_t + \delta_i + \epsilon_{it}, \qquad (1)$$

where DV_{it} is the dependent variable that measures dark trading activity (*PassiveDarkVolume or DarkOrderVolume)*; $MPIR_t$ is a dummy variable that stands for the change in regulation and it is 0 before October 15, 2012, and 1 thereafter; VIX_t is the daily realization of the U.S. market volatility index VIX, and δ_i is a security fixed effect.

Volatility is known to affect trading variables; since our securities are cross-listed with U.S. markets, we include the U.S. volatility index VIX as a control. To avoid biases in standard errors stemming from observations that are correlated across time by security or across securities by time or both, we employ standard errors that are double-clustered by both security and date (see Cameron, Gelback and Miller (2011) and Thompson (2011)).

Panel A in Table III confirms the observations in Figure 1 and Table II that *PassiveDarkVolume and DarkOrderVolume* have both declined significantly after the change in the dark liquidity rules.

A.2. What was the impact of the minimum price improvement rule on dark trading by marketplace?

MPIR was binding for the organization of trading in dark pools Ad and D, which had to adjust their trading rules to accommodate the change in regulation. In contrast the rule did not directly affect dark orders on lit marketplaces. To understand the relation between the organization of trading and the impact of the dark liquidity rules, we now analyze the change in dark trading by marketplace.

In response to the change in regulation, dark pools Ad and D adjusted their trading rules to ensure that all orders comply with the MPIR. Ex ante, the impact of the increased price improvement on dark trading is not obvious. On the one hand, passive orders in dark pool facilities Ad and D became more expensive because these orders lost the option to offer a price improvement of less than 50% of the NBBO. On the other hand, the larger price improvement made dark pools more attractive for marketable orders, potentially increasing the probability of execution of dark orders in marketplaces Ad and D.

Dark orders in lit markets were only marginally impacted. These orders were either pegged to trade at the midpoint of the NBBO, which already provided the required price improvement, or they were priced, fully hidden limit orders, which had lower priority than visible orders at the same price and thus already provided the required price improvement upon execution (relative to the next level in the book) within a marketplace. It is our understanding that the only impact on dark orders on lit markets was the effective introduction of visible order priority *across* marketplaces. Prior to October 15, 2012, dark limit orders could trade at a given price level, after all visible liquidity at the same price level on the *same* marketplace was exhausted – even when other marketplaces displayed orders at this price level. After October 15, 2012, to comply with MPIR, dark limit orders could trade only after *all* visible liquidity at the relevant price level on *all* marketplaces has been exhausted.

To establish the impact of the change on dark trading by venue, we compute the two measures of dark volume for each of the four major marketplaces (A-D) and the total for the remaining venues (E-H). We note that all dark trading on marketplace A occurs in its dark trading facility Ad.

Table II provides summary statistics on dark trading, split by trading venue. Before MPIR, (passive) dark trading on market A accounted for almost 4.6% of all Canadian dollar trading volume whereas after MPIR, it accounted for 0.8%. Market D, the other main dark market, accounted for 2.5% before and after MPIR. Markets B and C accounted for 1.4% and 0.8% respectively before and after MPIR. Both market A and market D experience a drop in order volume, from 10.4% to 6.4% of all market-wide order volume for market A and from

4.7% to 3.3% for market D.

We formally analyze the impact of the MPIR on dark trading by estimating the following linear security-market panel specification

$$DV_{it} = \sum\nolimits_{m\in\{A,B,C,D,O\}} \alpha_m \times mk_m \times MPIR_t + \beta \times VIX_t + \delta_i + \epsilon_{it}, \qquad (2)$$

where DV_{it} is the dependent variable that measures dark trading activity (*PassiveDarkVolume and DarkOrderVolume)*; mk_m is a dummy that is 1 if the dependent variable observation is for market *m,* where m=O stands for all marketplaces other than A, B, C, D; $MPIR_t$ is a dummy variable that stands for the change in regulation and it is 0 before October 15, 2012, and 1 thereafter; VIX_t is the daily realization of the U.S. market volatility index VIX, and δ_i are fixed effects for market and securities.

Equation (2) allows us to simultaneously estimate the effect of MPIR for all affected marketplaces and to test whether the sharp decline in dark volume on market A reported in Panel A of Table II is indeed larger than on other markets. Table III, Panel B displays the results from our estimation of equation (2). We confirm that indeed market A experiences a significant drop in volume: *PassiveDarkVolume on market A declines* by 3.81%. Both market A and market D see a significant drop in *DarkOrderVolume*, of 4.04% and 1.33%, respectively. A formal test for equality of coefficients α_A and α_D is rejected at all conventional levels suggesting that the drop in order volume for market A is larger.

All dark trading on market A occurs in its dark pool Ad, and dark trading on market A is thus most similar to dark trading on market D. In Figure 2, we plot the level of dark trading dollar volume (in logs) for markets A and D. The figure confirms the regression observations in Table III: market A sees a significant decline in dollar trading volume whereas market D value is essentially unaffected.

In summary, we observe that dark trading on lit marketplaces B and C is unaffected by the change in dark liquidity regulations. The two dark markets, Ad and D, experience significant changes in order volume, but only market Ad experiences a drop in trading volume. We develop possible explanations for this difference in Section VI.

B. What was the impact of the minimum price improvement on market quality?

We next analyze whether the decline in dark volume is associated with changes in the volatility, price efficiency, price discovery, and liquidity.

B.1 Volatility

We measure the volatility of prices using two measures. First, we compute the realized volatility of intra-day returns, measured as the sum of squared 1-minute returns. Second, we compute the trading-range measure, defined as the difference of the maximum and minimum prices over 1-minute intervals, scaled by the average price. In this report, we base our analysis on the transaction prices, and we compute one-minute returns based on the last price per 1-minute interval; all measures are expressed in basis points.

To formally test for changes in volatility we estimate a linear specification security-date panel

$$DV_{it} = \alpha \times MPIR_t + \beta \times VIX_t + \delta_i + \epsilon_{it}, \qquad (3)$$

where DV_{it} is the daily realization of the respective volatility measure. The variable of interest is the estimate for α, as this number signifies the impact of the change.

Table IV reports summary statistics and Table V presents the estimation results. We find that there is no change in volatility as measured by the range measure, and we see only a marginally significant (at the 10% level) increase in realized volatility.

B.2 Price Efficiency

We measure price efficiency by the absolute value of the autocorrelation of 1-minute returns. A lower autocorrelation correspond to higher price efficiency. In this report, we base our return computation on the last trade-price per 1-minute trading interval. To test for changes in price discovery, we estimate equation (3), with the return auto-correlation as the dependent variable. The third column in Table IV presents our estimation results and it shows that there is no measurable change in price efficiency.

B.3 Price Discovery

In this report, we measure price discovery using the 1-, 10-, and 30-second and 1-, 5- and 10-minute price impact of trades, where the price impact is measured as the signed change in the quoted midpoint x-seconds/minutes after the trade

$$price\ impact_{it} = q_{it} \times (m_{i,t+x} - m_t)/m_t, \qquad (4)$$

where q_{it} is an indicator variable that is 1 if the trade at time *t* is a buy and -1 if it is a sale; m_t is the prevailing NBBO midprice at time *t*, and $m_{i,t+x}$ is the prevailing NBBO midprice x units of time in the future. A larger price impact of trades is associated with a faster price discovery process – but higher adverse selection costs.

To estimate the effect of MPIR on the price discovery process, we estimate (3) using the volume-weighted daily average of price impact as the dependent variable. Table IV reports summary statistics and columns 5-10 Table V reports the estimation results for the change in the market-wide price impact at different time horizons. We observe no changes in the shorter horizon price impacts (at 1 to 30 seconds) and we observe a decline in longer-horizon (1 to 10 minutes) price impacts, suggesting a reduction in adverse selection following the introduction of MPIR.

B.3 Liquidity

We measure market-wide liquidity with the effective spread, which is defined for a trade at time *t* in security *i* as

$$espread_{it} = q_{it} \times (p_{it} - m_t)/m_t, \qquad (5)$$

where q_{it} is an indicator variable that is 1 if the trade at time t is a buy and -1 if it is a sale; m_t is the prevailing midpoint of the Canadian national best bid and offer (NBBO) prices at time t, and p_{it} is trade price. A larger effective spread is associated with lower liquidity and higher trading costs.

To estimate the effect of MPIR on market-wide liquidity, we estimate (2) using the volume-weighted average of the effective spread per day per security as the dependent variable. Table IV reports summary statistics and column 4 in Table V reports the estimation results for the change in the effective spread. We do not find any changes in the effective spread.

IV. Impact of the Minimum Price Improvement Rule by Trader Type

In Section III we have discussed the impact of the change in dark liquidity regulations on trading volume and market quality in aggregate. In this section, we analyze the impact of the change on trading volume and trading costs and returns for different groups of traders (e.g., retail, institutional, HFT).

A. Trader classification

All traders access the marketplaces via brokers. We base our classification on the analysis of order submission and trading behavior by trader IDs, where we define a trader ID as the combination of broker ID plus user ID, plus the account type (client, specialist, inventory, option market maker, and non-client). User ID is the most granular identification that is available to regulators in Canada; IIROC staff describes the usage of user IDs in detail in a recent research reports (IIROC 2012 and IIROC 2014).[16]

According to IIROC staff reports, a user ID is assigned by a marketplace, and it may identify a single trader, a business stream (for example, all orders that originate through a broker's online discount brokerage system), or a client that accesses trading venues directly (through a direct market access (DMA) relationship). It is our understanding that the brokers separate different types of order flows (e.g., retail vs. institutional) by user ID. For DMA clients, IIROC requires this. However, according to IIROC (2012), a DMA client may be assigned more than one user ID, for instance, to trade through multiple brokers or on different marketplaces, and they may choose to use multiple user IDs for business or administrative purposes.

For the classification of traders we expand our sample of 92 frequently-traded cross-listed securities to additionally include the 151 frequently traded securities that are part of the S&P/TSX Composite index, Canada's main market index. We classify traders based on trading characteristics that we collect for the eight weeks that precede our sample period (July 4 to August 24).

We group traders into four categories: HFT, retail, institutional, and other. The "other" category includes trader IDs that were not able to classify as HFT, retail, or institutional. Table VI provides summary statistics on these trader ID groups. We have a total of 3,642 unique trader IDs in our classification sample, many of these are,

[16] See http://www.iiroc.ca/documents/2012/c03dbb44-9032-4c6b-946e-6f2bd6cf4e23_en.pdf and http://www.iiroc.ca/Documents/2014/169edd4f-15e6-4330-8cb5-2c31e8f2bf82_en.pdf

however, inactive.

Retail. In Market A, seek dark liquidity (SDL) orders are exclusively available to retail investors. The use of SDL orders is the choice of the broker, not the customer, and it is our understanding that brokers have to explicitly seek to be connected to venue Ad to use this order type. We extract all trader IDs that use SDL orders from the complete database (which spans January 1, 2012 to June 30, 2013), and we classify these traders as retail. In our sample, we observe 135 such IDs. We know with certainty that these trader IDs are used to trade order flow from retail investors, but there may be other trader IDs that are assigned for order flow from retail investors that are not captured by our classification.

Buy-Side Institutional: We conjecture that buy-side institutions will be involved in large pre-arranged trades and accumulate large inventory positions. We therefore use these two criteria to identify buy-side institutions.

First, we extract all trader IDs that involve a client account and that are involved in a so-called "intentional cross". An intentional cross is a trade, usually a large one, that is pre-arranged off-exchange by a brokerage, for instance to match two client orders or take an inventory from a client via a liability desk.

Second, we search for trader IDs that accumulate large inventory positions across all Canadian marketplaces. We determine each trader's maximum cumulative position for the classification period in *non-crosslisted* stocks, assigning a zero inventory at the beginning of the period. We focus on non-crosslisted securities to reduce the possibility that a seemingly large inventory position of an entity is offset by an equally large position elsewhere. Since a trader may buy in one jurisdiction and sell in another, for instance, to exploit an arbitrage opportunity, it is imaginable that a Canada-only position is off-set by a U.S.-based position. We acknowledge that this classification is imperfect, for instance, because a trader ID may trade on behalf of multiple retail clients who jointly accumulate a large position or because a DMA client may use different trader IDs for buying and selling securities. To mitigate these imperfections, we set a high bar for the required cumulative position.

We classify the trader ID as a buy-side institution if its maximum cumulative position during the classification period exceeds $10,000,000 in absolute value. We classify 558 trader IDs as assigned to buy-side institution order flow.

High Frequency: The critical component of high frequency trading is that trading is automated and that traders have the ability to react quickly to market conditions. Definitions used by various regulators or policy institutions (e.g., BAFin in Germany, the European Commission, or the S.E.C.) often include as a requirement that HFTs use many orders, in particular in relation to their trades. In our opinion, using orders or order-to-trade ratios biases the classification against traders or strategies that use only marketable orders.

We focus on reaction speeds as the main metric to identify HFTs, and we use reaction times that are faster than human reaction times (the average duration of a single blink of a human eye is 100-400 milliseconds, according to the Harvard Database of Useful Biological Numbers[17]). We further require that trader IDs exhibit fast reaction times for long stretches of time, across many trades, and in many securities. We use two criteria to quantify a trader ID’s reaction speed.

[17] http://bionumbers.hms.harvard.edu//bionumber.aspx?&id=100706&ver=1

Our first criterion is the trader ID's median order-to-cancel time. The order-to-cancel time is the time from the submission to cancellation of the same order; for the purpose of this classification, we exclude immediate-or-cancel (IOC) orders, because their order-to-cancel time is determined by the processing speed of the marketplace.

Our second criterion is the number of trade and order messages that a trader ID submits during a short interval after a daily scheduled public information release. We focus on the first 500 milliseconds after 3:40 p.m., which is when the TSX first publishes the imbalance between the buy and sell orders in its market-on-close facility.

The closing price for TSX-listed securities is determined in a multi-stage process. Before 3:40 p.m., traders may submit market buy and sell orders tagged as market-on-close orders. These orders will trade at the 4:00 p.m. closing price. At 3:40 p.m., the TSX publishes the imbalances of buy and sell orders, and traders then have the opportunity to submit priced limit orders to trade at the market-on-close to off-set the market order imbalance. The market-on-close imbalance is indicative of the closing price and may help predict behavior over the last 20 minutes of trading.

In aggregate, there is a significant spike in trades immediately after the publication of the market-on-close imbalance, though this spike may not be visible or pronounced on a stock-by-stock basis. Figure 3 plots the by-minute number of trades, aggregated over all securities in our sample over all days in the classification period. The dataset that is provided to us by IIROC does not contain information on the market-on-close announcement. Thus, we are not able to determine the time between the publication of the market-on-close imbalance and a trader's action at the millisecond level. For this reason, we classify trader IDs as HFTs based on their actions during a relatively long interval of 500 milliseconds after the announcement.

For each trader ID, stock and day we compute the median order-to-cancel speed, and for each trader ID we compute the total number of orders and aggressive trades during the 500 milliseconds after 3:40 p.m. A trader ID is classified as HFT

1. if the median of the traders ID's median stock-day order-to-cancel speeds is below 250 milliseconds, or
2. if the trader ID submits more than 1,000 orders or is involved in more than 500 aggressive transactions in the first 500 milliseconds after the market-on-close publication across all securities in our classification sample during our classification period.

We classify a total of 89 trader IDs as HFT. Our HFT group displays the characteristics typically expected of high frequency traders. As a group, they account for 82% of the orders and 83% of order cancellations during our classification period, they account for 53% of all passive trades, and 34% of all aggressive trades. On an average day, the average HFT trades in 136 securities. None of the HFT trader IDs fall into either the retail or buy-side categories. HFTs are also among the top users of immediate-or-cancel orders, the use of which indicates that the trader is speed sensitive.

Inventories of HFT Trader IDs. A common perception is that high-frequency trading firms aim to hold no or only very small overnight inventories. We observe that most trader IDs that we classify as HFT hold substantial median end-of-day inventories, even in non-crosslisted securities. Furthermore, several of the fastest trader IDs that we classify as HFT trade more than 85% passive, have order-to-trade ratios in the 99^{th} percentile, and yet hold median inventories of 70% or more of their daily trading volume.

This observation highlights the importance of understanding the usage of trader and user IDs in different jurisdictions and in different datasets. In Canadian markets, a single DMA client may use multiple trader IDs (IIROC (2012) and IIROC (2014)), and it is thus possible that an HFT firm is assigned multiple user IDs. Furthermore, a single user ID may be used for trading activity of multiple entities, for instance, for all the brokerage's retail order flow (which is balanced, on average). As a consequence, low end-of-day inventories are neither a necessary nor a sufficient attribute of an HFT trader ID in our dataset.

B. Intra-day returns by trader group

We compute the benefits from trading based on the intra-day returns from buying and selling, and we evaluate end-of-day positions at the closing price. We measure returns in terms of basis points of the total value traded, and we compute returns by trader group. We compute the return per stock i per day t for each of a group of trader IDs (HFT, retail, buy-side, or other) as follows

$$iret_{it} = ((sellval_{it} - buyval_{it}) + (buyvol_{it} - sellvol_{it}) \times close_{it}) / value_{it}, \quad (6)$$

where $sellval_{it}$ and $buyval_{it}$ are the dollar amounts sold and bought in aggregate by the respective group of trader IDs in security i on day t, and $sellvol_{it}$ and $buyvol_{it}$ are the number of shares sold and bought by the group. The (unrealized) profit from intra-day trading is $sellval_{it} - buyval_{it}$; a positive value for a group indicates that trader IDs within the group bought low and sold high in the aggregate. The end-of-day inventory position is $buyvol_{it} - sellvol_{it}$, and we evaluate this position at the security's closing price.

The intraday return measure accounts for favorable and unfavorable price movements after the trade, and it implicitly includes transaction costs. If, for instance, the prevailing bid and ask prices remain constant throughout the day, then the intra-day return of a trader who buys the security at the ask price equals the effective spread.

We acknowledge that the intraday return measure has shortcomings. First, it is inventory-based and we do not know traders' true end-of-day inventories, for reasons that are related to the usage of trader IDs by brokerages as discussed in subsection A and because securities in our sample are cross-listed with U.S. markets. We believe, however, that aggregating volume across all trader IDs within each groups helps mitigate this issue. Second, the measure benchmarks a trader ID's intraday performance to the closing price, and it thus assumes that the end-of-the-day price is the true efficient price for the security. If a large uninformed institution is building or liquidating a position over many days, then they may have a transitory price impact that will affect the end-of-the-day price only temporarily. Finally, we measure an unrealized return, and we acknowledge that traders may not be able to close their positions at the end-of-the-day prices to realize these returns.

Table VII provides summary statistics for intra-day returns, split by trader type and by market. We observe that HFTs have the highest intra-day returns, and that retail traders have the lowest intra-day returns.

To determine the changes in returns following the introduction of the dark liquidity rules, we perform a panel regression analysis (by trader group). We estimate the following regression for the intra-day return measure for each trader group:

$$DV_{it} = \sum\nolimits_{m \in \{A,B,C,D,O\}} \alpha_m \times mk_m \times MPIR_t + \beta \times VIX_t + \delta_i + \epsilon_{it}, \qquad (7)$$

where DV_{it} is the dependent variable that measures the intra-day return for the specific trader group, split by marketplace; mk_m is a dummy that is 1 if the dependent variable observation is for market m, where m=O stands for all marketplaces other than A, B, C, D; $MPIR_t$ is a dummy variable that stands for the change in regulation and it is 0 before October 15, 2012, and 1 thereafter; VIX_t is the daily realization of the U.S. market volatility index VIX, and δ_i is a market and security fixed effect. Results for the estimation of equation (7) are in Part 1 of Table VIII.

We observe that retail traders experience a (statistically weak) decline in their intra-day returns on market A. This decline is consistent with a decline in trading volume in Ad, where retail orders received price improvement prior to the change in regulations. We explore the change in the distribution of retail order flow between dark and lit trading facilities of marketplace A before and after the change in Subsection E. The decline in intra-day returns is more pronounced and statistically significant after accounting for the marketplaces' trading fees (maker-taker fees), which is consistent with marketable retail orders incurring higher fees in the limit order book Al relative to what they would incur in dark facility Ad. Since most retail brokers only pass maker-taker fees to clients through flat commissions, our results suggest that retail brokerages may incur higher trading fees after the change in dark liquidity regulations.

We further observe an increase in profits for the group of buy-side trader IDs on market C.

C. Implementation Shortfall

A measure that is related to intra-day returns is the implementation shortfall, which is useful in measuring the cost of a large "parent" order that is split into multiple "child" orders. This measure compares the realized cost of establishing or unwinding a position with the hypothetical cost that would be obtained if the trader has filled the entire position at the time when the trader starts trading the child orders. Unfortunately, we have no information about "parent" orders, and our measure of implementation shortfall is thus imprecise. As an approximation of the price prevailing at the time that a trader ID started to build or unwind a position, we use the first price, labeled $first_{it}$, at which a trader trades on any given day. We then compute the "raw" shortfall per trader j

$$raw\ shortfall_{it}^{j} = \left(sellvol_{it}^{j} - buyvol_{it}^{j}\right) \times first_{it}^{j} + \left(buyval_{it}^{j} - sellval_{it}^{j}\right), \qquad (8)$$

where $\left(sellvol_{it}^{j} - buyvol_{it}^{j}\right) \times first_{it}^{j}$ is the trader j's hypothetical cost of establishing a position at the first trading price, and $buyval_{it}^{j} - sellval_{it}^{j}$ is the trader's realized costs. The smaller the shortfall, the lower the realized trading costs. We then sum the raw shortfalls for all traders within a group, and we scale the sum by the group's daily dollar trading volume:

$$shortfall_{it} = \sum\nolimits_{j} raw\ shortfall_{it}^{j} \Big/ \sum\nolimits_{j} val_{it}^{j}. \qquad (9)$$

We believe that the shortfall measure is most relevant for buy-side traders. Arguably, the objective of HFTs is to generate intra-day profits and thus the hypothetical situation of filling all their orders at their first trade price is moot. Furthermore, it is our understanding that retail orders are generally not split into smaller orders (unless a retail investor splits his or her order by sending multiple order to the brokerage). We include the statistics for the non-buy-side groups for completeness only.

Table VII presents the summary statistics for the shortfall measure. We observe that buy-side traders face a positive shortfall and that the magnitude of the shortfall is comparable across markets. Part 2 in Table VIII estimates the effect of the change in the dark liquidity rules, and we find no evidence that these rules affected buy-side institution's implementation shortfall.

D. Fill Rates

In 2009, the CSA and IIROC published a joint consultation paper (CSA (2009)), where they discussed different views on dark trading. They explain, in particular, the view of those who are opposed to dark orders that are pegged to the national best bid and offer prices. According to CSA (2009), the opponents of pegged dark orders believe that these "orders "free-ride" on the contribution of those that have posted visible limit orders and [that] the execution of a primary pegged order ahead of the order establishing the best bid or offer is unfair" because "the investor will not achieve the benefit of posting the limit order (e.g. execution or rebate credit)."

In this section, we study the relation between dark order submission and the probability of execution for lit orders. Table III illustrates that the introduction of the MPIR led to a reduction in submitted dark orders, in particular, in markets Ad and D where these orders were pegged to the NBBO. We proxy the execution probability of lit orders by the ratio of passive trading volume to all lit order volume (expressed in shares):[18]

$$pr(execution) = passive\ volume / lit\ order\ volume. \quad (10)$$

Similarly, we compute the ratio of passive trading volume to all lit order volume to assess the fill rates for dark orders.

Table VII provides summary statistics for fill rates (in percent). HFTs have low fill rates, consistent with the rapid submission and cancellation of orders that is commonly attributed to HFT strategies. Retail and buy-side trader IDs obtain higher fill rates.

Table IX displays results from a regression analysis where we estimate equation (7) with $pr(execution)$ defined in (10) and an analogous measure for the probability of dark order execution as the dependent variables. We observe that HFT fill rates improve substantially on marketplaces A and C. Fill rates for buy-side traders, however, decline in markets A and C and fill rates for retail trader IDs decline in market A – contrary to the intuition that a reduction in dark trading would increase the probability of execution for lit orders. One possible explanation stems from the increased competition on lit markets for liquidity provision to retail order flow that would have been executed in the dark (on venue Ad) prior to the implementation of MPIR. We explore changes in traders' order routing decisions below, in Subsection E. and Section VI.

[18] We cannot directly infer from the data which orders are marketable at the time of their submission.

Analyzing fill rates for dark orders, we find an increase in the probability of execution for buy-side dark passive orders in market D. We explore whether this change is related to changes in order flow segmentation across different trading venues below.

E. Where do different groups of traders trade before and after MPIR?

To study order flow segmentation across marketplaces, we compute trading dollar volume by trader type. We split a trader ID's total trading volume into four categories: *dark and aggressive*, which is the volume where the trader ID is on the active side and their (marketable) order is marked dark, *lit and aggressive*, which is the remainder of the trader ID's marketable order volume, *dark and passive*, which is the volume where the trader ID is on the passive side and their order is marked dark, and *lit and passive*, which is the remainder of the trader ID's passive trading volume.[19] We aggregate the dollar trading volume by trader group (retail, institutional, HFT, and other) by summing it across all trader IDs within each group, and we express it as a fraction (in %) of the group's total trading volume.[20]

Table X displays the distribution of these types of trading volume as well as the distribution of dark, lit, aggressive and passive dollar volume traded by trader group across different marketplaces before and after the change in dark trading regulations. The numbers per group are computed as fractions of the total dollar trading volume for the group.

To formally analyze changes in trading volume by group following MPIR, we estimate equation (7) where DV_{it} is the dependent variable that measures the trading volume of the specific type (e.g., dark and aggressive) and for the specific trader group, split by marketplace. Results for the estimation are in Table XI.

High-Frequency Traders. The distribution of trading for the group of high-frequency trader IDs is related to the market shares that we observe for the four major Canadian marketplaces: HFTs trade most of their volume, both passively and aggressively, on marketplace B, followed by marketplaces C and A. The regression analysis shows that the distribution of HFT trading volume was affected by the change in regulations. After the change, HFTs reduce their dark, passive trading in dark pools Ad and D, and they increase their aggressive dark trading in dark pool D, consistent with the idea that the MPIR increases the cost for passive trades in dark trading venues and decreases it for aggressive trades against these dark orders. We further note that HFTs reduce their lit passive trading on market B and increase it on marketplaces A and C, with a large increase on market A.

Figure 4 plots the dollar order volume (in logs) that HFTs submit to dark pools Ad and D, and it illustrates that on aggregate they send fewer orders (in dollar terms) to dark venues after the change in dark liquidity rules.

The observed changes in the distribution of HFT trading volume are consistent with our observations in Table IX on the reduction in lit order fill rates for buy-side and retail trader IDs on venues A and C, and on the

---

[19] After the change in dark liquidity rules, SDL orders that did not find a match in Ad were automatically routed to Al. This re-routed order flow is classified as "lit and aggressive" in our categorization.

[20] Computing the average volume per trader ID is meaningless, since the usage of trader IDs may be heterogeneous across brokerages and clients: a single trader ID may be used to trade on behalf of multiple clients and at the same time a single client may use more than one trader ID.

increase in buy-side fill rates for dark orders in market D. Since lit markets operate according to time priority, HFTs have an advantage there when providing liquidity. If the increase in the fractions of HFT passive volumes on markets A and C stems from HFTs sending a larger fraction of their passive orders to lit venues, then the increased competition from HFT liquidity providers may reduce fill rates for slow liquidity providers on these venues. Similarly, if HFTs choose to reduce the share of their passive order flow that is sent to dark venues, other market participants may see an increase in their fill rates. We develop possible explanations for the observed reduction of the fraction of HFTs passive volume that is sent to dark venues in Section VI.

Retail Investors. As discussed in Battalio, Jennings and Corwin (2014), retail traders often do not control the choice of execution venues for their orders, and instead these decisions are taken by retail investors' brokerages. Battalio et al argue that broker routing decisions are consistent with minimizing trading fees. During our sample period, market Ad charged the lowest fee for marketable orders (the so-called taker fee), of \$0.04 for a 100-share trade. Of the lit markets, market Al offered the lowest taker fee at \$0.28 for a 100-share order. Market B offered the highest maker rebate, of at least \$0.31 for a 100-share trade.[21]

Panel B of Table X displays the distribution of retail trading volume across marketplaces. Most of the retail marketable orders clear on market Ad or market Al (taken together as marketplace A, these two venues account for 59.4% of all aggressive retail dollar volume), and the majority of passive retail orders trade on marketplace B (60% of all passive retail dollar trading volume). Our findings are consistent with the idea that trading venues' fees influence brokers' routing choices, subject to other constraints. For instance, the order protection rule may require marketable orders to be routed to the venue that offers the best price (rather than the lowest taker fee), and broker's best execution obligations may dictate that the broker posts client limit orders on a venue that does not offer the highest rebate.

Part 2 of Table XI displays regression results on retail trader usage of orders before and after MPIR. We find that retail traders reduce trading in dark pool Ad and that they increase their usage of aggressive orders on the lit market Al. Retail trader IDs also increase their usage of aggressive orders on markets B and C, consistent with order routing that respects the order protection rule. Assuming that the reduction in dark trading volume on Ad is not driven by a sudden decline in the aggregate retail order flow, retail marketable orders that used to find execution in Ad will have to be routed elsewhere after the change. Since market Al has the lowest taker fee, we expect it to receive most of these orders. However, some of these orders will need to be routed to other markets, to obey the order protection rule.

The decline in retail traders' intra-day returns that we observed in Table VIII is consistent with the shift of retail marketable order from the dark facility Ad in marketplace A to its limit order book Al. Orders that previously received a price improvement in dark pool Ad would pay the full bid-ask spread plus the higher taker fee in market Al. Since we do not observe a change in the bid-ask spreads following the change in dark liquidity rules, we would expect the intra-day returns of retail traders to decline.

Buy-Side Institutions. Panel C in Table X shows that buy-side institutional trader IDs trade 9% of their passive volume on market A and 8% of their passive volume on market D. Their overall lit trading is concentrated on

---

[21] The exact amount of the rebate depends on the brokers' monthly volume.

marketplace B, which has the largest market share. The share of buy-side trader IDs' dark passive volume that is executed in dark pool D is six-fold the share of their dark passive volume in dark pool Ad (3.6% of the of the buy-side trader IDs total trading volume vs. 0.6%). The concentrated dark trading in dark pool D is somewhat surprising, since market Ad offers the opportunity to match with other dark orders and additionally the option to interact with aggressive order flow from retail investors, who are presumably uninformed and non-strategic. We further study differences in liquidity provision in the two dark pools in Section VI.

Part 3 of Table XI illustrates that after the introduction of MPIR, institutions trade less on market D with aggressive orders but that they trade more there with passive orders – in contrast to what the naïve intuition would suggest and in stark contrast to HFT behavior. Their liquidity provision in dark pool Ad is unaffected, but we note that they did not have significant presence there even before MPIR. Furthermore, buy-side institutions increase their aggressive order flow to market Al. Given our findings in Part 2 of Table XI on the increased fraction of retail marketable orders on market Al, this outcome is consistent with Admati and Pfleiderer (1988) who predict concentrated trading of informed and uninformed traders (which allows the informed traders to minimize their price impact). Applying this idea to our setting, we would expect that an increase in retail marketable orders (which are presumably uninformed) on market Al would allow institutions to "hide among these" to achieve lower trading costs.

V. Order Flow Segmentation and Market Liquidity

In Section IV, we discussed that the fraction of retail marketable orders that execute on the lit markets has significantly increased following the introduction of the MPIR, and that this increase was concentrated on market Al (Table XI).

Trading with retail order flow is considered to be desirable because this flow is deemed to be uninformed and non-directional (retails investors as a group are expected to have only small imbalances of buys and sells). Our findings thus suggest that posting liquidity on lit markets generally and on market Al specifically became relatively more attractive.

We did not, however, observe a change in the market-wide bid-ask spread in Table V, possibly because the change in the composition of order flow was not large enough to affect market-wide measures. We will now study whether the change of the distribution of retail flow causes meaningful changes in the composition of the order flow on individual venues. In Figure 6 we plot the retail share of a lit market's aggressive dollar volume, by trading venue. The figure illustrates a substantial increase in the share of retail volume for market Al, from around 15% of dollar trading volume to 30%. The increases for markets B and C are much smaller.

A strong change in the order flow composition on a single venue suggests that there should be changes in liquidity for that market, either in the sense that traders there post a tighter bid-ask spread or that they post more depth at the best prices. The left panel in Figure 7 plots the natural logarithm of the dollar-depth for the three markets Al, B and C. The figure indicates that there is an increase in depth on market Al and little to no changes for markets B and C, consistent with the increase in the concentration of retail marketable order flow on market Al. The right panel further confirms the increase in liquidity for market Al by zooming in on market Al. Figure

8 depicts strong co-movement between the depth on market Al and the retail share of market's Al's aggressive trading volume.

To formally test whether there is a change in liquidity on individual markets, we estimate the effect of the introduction of MPIR on markets' time-weighted quoted spread and time-weighted quoted depth in the following panel regression:

$$DV_{it} = \sum_{m \in \{Al,B,C\}} \alpha_m \times mk_m \times MPIR_t + \beta \times VIX_t + \delta_i + \epsilon_{it}, \qquad (9)$$

where DV_{it} is one of the four liquidity variables (the time-weighted per stock per day quoted spread in cents and in basis points of the prevailing midpoint, and the logarithm of share and dollar depth at the best prices.); mk_m is a dummy that is 1 if the dependent variable observation is for market *m; $MPIR_t$* is a dummy variable that stands for the change in regulation and it is 0 before October 15, 2012, and 1 thereafter; VIX_t is the daily realization of the U.S. market volatility index VIX, and δ_i is a market and security fixed effect.

Table XII presents the results of our estimation. It illustrates that there is a significant increase in depth on market Al following the introduction of the dark liquidity rules, by about 15%. It shows a marginally significant (at a 10% level) increase in the time-weighted quoted spread measured in basis points but no change for the spread measured in cents.

The absence of significant changes to the bid-ask spread is not surprising. The order protection rule requires that marketable orders are routed to the venue that is posting the best price, and we thus do not expect spreads on major individual venues to substantially differ from each other.

Overall, our results imply an improvement in liquidity on market Al that we attribute to the increased retail share of market Al's aggressive volume.

Does market Ad volume predict the change in market Al depth? Prior to October 15, 2012, Canadian market participants agreed that almost all marketable orders in market Ad stemmed from retail traders; our analysis shows that this attribution was correct. As our analysis in Section III indicates, after MPIR, volume in market Ad declined almost to zero. Furthermore, in Section IV we argued that market Al experienced an increase in its fraction of marketable retail order flow, and in this section we argued that quoted depth on market Al increased significantly, by about 16%.

We now establish that the extent of retail trading on market Al and the depth on market Al are systematically related. To address this question, we study whether trading volume on market Ad *before* the rule change has predictive power for the change in depth *after* the regulatory change.

The fraction of volume traded in market Ad prior to October 15, 2012 provided the market with an accurate estimate of the amount of retail flow that did *not* hit the lit markets prior to October 15, 2012. Assuming no sharp changes in retail volume on October 15, 2012, the volume that does not execute on market Ad must trade on the other venues.

Since the routing decision commonly lies with the broker, the marketplaces' taker fees for liquidity-demanding orders should arguably play a role. Among the lit marketplaces, market Al charges the lowest taker fees. It was thus reasonable to expect, that retail order flow that does not trade on market Ad would be preferentially routed to marketplace Al (conditional on abiding by the order protection rule).

The sharp decline in market Ad's trading volume thus provided traders with a unique opportunity to estimate the "extra" retail volume that hit market Al after October 15, 2012.

If quoting activities in market Al after October 15, 2012 are related to the change in the extent of retail trading, then we expect to find a relation between the pre-rule-change dark trading volume in Ad and the change in the depth on Al. If, however, the change in depth on Al is unrelated to changes in retail order flow, then pre-rule-change trading in market Ad should have no predictive power over changes in depth on Al.

We compute the average per-security total dollar-volume of dark trading in market Ad, which we label $totalretailvolume_i$, and the share of market Ad dollar-volume of all market A dollar-volume, which we label $\%retailvolume_i$. Further, we compute the differences in the before and after MPIR average daily time-weighted quoted depths in markets Al, B, and C. We then estimate the following relation

$$\Delta depth_i = \alpha + \beta \cdot market\ A_d\ value_i + \varepsilon_i,$$

where $\Delta depth_i$ is the difference in the average daily time-weighted quoted depth after and before the introduction of MPIR for security i, and $market\ A_d\ value_i$ is either $totalretailvolume_i$ or $\%retailvolume_i$. We also use the September 2012 log-market cap as a control for firm size in some of the regressions.

Table XVII displays the regression results. The first two columns indicate that a 1% drop in *%retailvolume* increases depth in market Ad by between \$253 and \$270. The last two columns show that a \$1000 drop in market Ad volume leads to a \$0.02-\$0.025 increase in dollar-depth in market Al. While this number may appear small, one has to keep in mind that dollar-volume is the total per day whereas depth is a time-weighted average. To put the number into perspective: on average \$1.9M of volume is traded in market Ad per security per day. Should all this volume disappear and move to market Al, then per-stock, per-day, the time-weighted quoted depth in market Al should increase by around \$3,350, or about 11%. This number is consistent with the magnitude of the effect that we estimated with our panel approach.

In untabulated regressions, we further observe that market Ad trading has no predictive power over the changes in any of the other markets, B and C and that market Ad volume before the rule change has strong predictive power for the fraction of retail trading in market Al after the event. Our findings thus suggest that an increase in retail volume in market Al led to an improvement of liquidity there.

VI. Liquidity Provision and the Minimum Price Improvement Rule

Figure 2 and Table III illustrate that the change in dark trading regulations affected dark pools Ad and D differently: trading volume in Ad significantly declined whereas trading volume in D remained unchanged. In Section IV we have established that the change in the regulations affected the segmentation of trading volume across different marketplaces, for each trader group. In this section, we focus on trading in Ad and D before and after the change and study the composition of order flow within these markets before and after the change. By design of market Ad, most of its marketable order flow stems from retail investors; the same is not true for market D. We aim to understand whether and how this difference and other potential differences in the composition of order flow between the two venues explain their differential response to the minimum price improvement rule.

A. Who demands and provides liquidity in dark markets before and after MPIR?

In this section, we study the contribution of each group (HFT, retail, buy-side institutions, and other) to liquidity demand and liquidity provision in dark pools Ad and D, by marketplace.

Table XIII reports summary statistics on the passive (aggressive) dollar trading volume for each trader group on venues Ad and D, before and after the change in dark liquidity regulations, as a fraction (in %) of the total passive (aggressive) dollar trading volume on the respective venue during the respective time period (before and after the change). Aggressive trading volume on marketplace Ad for retail traders is the volume of SDL orders that were executed against resting dark orders, aggressive volume in Ad for non-retail traders is the volume of dark orders that matched with resting dark orders on the opposite side immediately upon submission.

To assess whether the changes in volume after the introduction of MPIR are significant, we estimate the following linear regressions for each trader group for each volume type (passive and aggressive), using a security-date panel:

$$DV_{it} = \sum_{m\in\{Ad,D\}} \alpha_m \times mk_m \times MPIR_t + \beta \times VIX_t + \delta_i + \epsilon_{it}, \qquad (10)$$

where the dependent variable DV_{it} is the amount of passive (aggressive) dollar trading volume for the trader group as a fraction of total passive (aggressive) dollar trading volume for the respective marketplace (Ad and D); mk_m is a dummy that is 1 if the dependent variable observation is for market m; $MPIR_t$ is a dummy variable that stands for the change in regulation and it is 0 before October 15, 2012, and 1 thereafter; VIX_t is the daily realization of the U.S. market volatility index VIX, and δ_i is a market and security fixed effect. The estimation results are in Table XIV.

High-Frequency Traders. Table XIII illustrates that HFTs provide a sizeable fraction of liquidity in both dark pools, and that they demand a much larger fraction of liquidity in dark pool D than in dark pool Ad. The latter is not surprising because, by design, liquidity providers in Ad may elect to interact exclusively with aggressive orders that stem from retail investors. After the introduction of MPIR, the share of the trading venue's passive dollar trading volume that is provided by HFTs' declined substantially in both dark pools. The results from our regression in Table XIV confirm this observation, but a test for the equality of the estimated coefficients is rejected at the 5% level, suggesting that the decline of HFT market share in passive trades is stronger for market Ad. The share of the venue's aggressive trades that stem from HFTs increases in both markets, and the increase is larger in market D (the equality of coefficients is rejected).

We believe that the reduction in HFTs' share of dark venues' passive trading volume stems from the nature of midpoint pricing, which makes it impossible to "earn the spread" in the dark pool and thus reduces incentives for market making. Many prominent HFT firms worldwide specialize in market making, both informally and through exchange-sponsored programs (e.g., some HFTs serve as NYSE-designated market makers), and it is likely that some of the HFT trader IDs in our sample made a two-sided market in dark pools Ad and D before the introduction of MPIR. We would expect that such trader IDs reduce their volume of passive orders to both dark pools after MPIR rendered dark pool market making unattractive.

Retail Investors. Table XIII illustrates that retail investors account for almost 100% of market Ad's aggressive dollar volume before the change and for 95% after the change, but that they account for less than 12% of market D's aggressive dollar volume. Retail traders provide very little dark liquidity, which we attribute to the order exposure rule that requires that client (passive) orders below a certain size be sent to marketplaces that display prices, unless the broker is explicitly directed otherwise. Tables XIII and XIV illustrate the decline in the retail investor market share of aggressive trades on venue Ad. This decline may be mechanically related to the increased probability of a match among dark orders and thus the increased market share or dark-to-dark trades (the incoming order that participates in a match is labeled as active in the dataset). Before the introduction of MPIR some dark orders offered a 20% improvement of the bid-ask spread, and such orders could not be matched. After the introduction of MPIR, all dark orders had to be posted at the midpoint, ceteris paribus increasing the probability of a match between two dark orders (traders had the option to opt out of matching with other dark orders).

Buy-Side Institutions. Table XIII illustrates that buy-side institutions' share of aggressive volume in Ad is negligible, but that they account for 18.1% and 12.1% of market D's aggressive volume, before and after the introduction of MPIR, respectively. They account for 3.5% and 25.1% of venue Ad's passive volume and for 47% and 61% of market D's passive dollar trading volume, before and after the introduction of MPIR, respectively. The buy-side share of passive volume in market D is notably large, given that the group of buy-side trader IDs accounts for only 24% of passive dollar volume across all marketplaces (Table VI). These observations are consistent with buy-side institutions having a comparative advantage over HFTs when providing liquidity in markets where larger orders enjoy execution priority (marketplace D employs a pro-rata matching mechanism).

Tables XIII and XIV illustrate that the market share of buy-side trader IDs of a venues aggressive trading volume has declined on market D, and that their share of passive trading volume has increased in both dark pools. The increase in the buy-side share of passive volume is unlikely to be mechanically driven by the decline in HFT market share of passive volume, because we do not observe the increase in "other" traders' share of passive volume. Instead, it appears that markets Ad and D became relative more attractive to buy-side investors after the change in regulations, possibly because the probability of a midpoint match is higher when all dark orders are submitted at the midpoint.

In summary, although the segmentation of order flow differs between markets Ad and D, we do not observe significant differences in *changes* in liquidity provision on markets Ad and D, by trader group, after the introduction of the MPIR. Our findings suggest that the differential response of the two marketplaces is driven by factors other then the segmentation of order flow by trader type, and we explore an alternative explanation below.

B. Market Making on Markets Ad and D

In this section, we analyze differences in liquidity provision for markets Ad and D other than by trader group. Table XV summarizes the stylized facts.

The number of trader IDs that provide liquidity (trade on a passive side in the respective market) in market D is close to 600, which is almost twice the number of such trader IDs in market Ad. MPIR appears to have little to no effect on these total numbers.

When we consider trader IDs' market shares of liquidity provision, we observe that liquidity provision in market Ad is much more concentrated. Sorting trader IDs by their fraction of the venue's passive volume, we observe that 95% of liquidity in market Ad (measured by the executed passive dollar volume) is provided by the 14 largest trader IDs (in terms of their shares of liquidity provision). In market D, on the other hand, it takes 185 trader IDs to provide 95% of liquidity. We refer to these groups of trader IDs as "top-95% liquidity providers." The difference in the number of "top-95% liquidity providers" mainly stems from buy-side and "other" trader IDs, whereas the number of HFT trader IDs that fall into the "top-95% liquidity provider" category is similar for both markets (4 in Ad and 5 in D).

Finally, we characterize liquidity providing trader IDs by the type of liquidity (one- or two-sided) that they post. Loosely speaking traders may trade passively for two reasons. First, they may use passive orders to build or unwind a position, in which case we expect their order flow to be directional (one-sided). Second, they may act as de facto market makers and trade on both sides of the market, to benefit from capturing the bid-ask spread.

To quantify a trader ID's one- vs. two-sided liquidity provision, we compute the imbalance score for each trader ID as follows[22]

$$imbalance_{it} = 2\times\left|\frac{buy\ order\ volume_{it}}{total\ order\ volume_{it}} - \frac{1}{2}\right|. \quad (11)$$

This imbalance score ranges from 0 to 100% where a score of 0% implies that a trader ID posts and equal number of buy and sell orders, and a score of 100 implies that a trader posts orders on one side of the market only. For each trader ID, we compute the imbalance score for trader ID i, imb_i, as the median of the trader ID's imbalance scores across all securities in our sample across all days before the introduction of MPIR.

Table XV shows that the average imbalance score for top-95% liquidity providers in market Ad is 28.6% whereas the average imbalance score in market D is 98.8%. Our findings suggest that before the introduction of MPIR most liquidity in market Ad was provided by trader IDs that posted two-sided orders, presumably to capture (a fraction of) the bid-ask spread, whereas most liquidity in market D was provided by traders who aim to build or unwind positions.

Figure 9 illustrates the difference in the imbalance scores in the two markets from a different angle, focusing on the amount of liquidity provided by trader IDs with lowest imbalance scores. We sort trader IDs by their

[22] This score ignores positions taken on other markets. The idea is, however, to assess whether a trader aims to trade in one direction.

imbalance score, and we compute the cumulative percentage of liquidity provided by trader IDs, starting from the trader ID with the lowest imbalance score. Figure 9 then plots the running average of the trader IDs' imbalance scores against the cumulative percentage of liquidity provided by these trader IDs. It illustrates, for instance, that over 80% of passive liquidity in market Ad was provided by trader IDs with imbalance scores of 10% or less (i.e., trader IDs that post an almost similar number of orders on both sides of the market), whereas less than 20% of passive liquidity in market D was provided by such trader IDs.

We believe that the difference between markets Ad and D in terms of the imbalance scores of their liquidity providers is critical to understanding the differential impact of the minimum price improvement rule on these two marketplaces.

Since, as argued before, the introduction of MPIR lowered the incentive to make a two-sided market in markets Ad and D, and we expect that traders who acted as market makers and posted orders on both sides of the market in Ad and D will reduce their liquidity provision.

To assess the impact of the imbalance score on the trader ID's liquidity provision in a dark venue, we compute, per liquidity providing trader ID, the difference in liquidity provided before and after the introduction of MPIR. We then estimate the following relation

$$\Delta\%liq_i = \alpha + \beta \times imb_i + \gamma_1 \times HFT_i + \gamma_2 \times retail_i + \gamma_3 \times buy - side_i + \varepsilon_i, \qquad (12)$$

where $\Delta\%liq_i$ is the difference in the percent of aggregate liquidity provided by trader *i* before and after MPIR, imb_i is the average imbalance score for trader *i*, and HFT_i, $retail_i$, $buy\text{-}side_i$, and $others_i$, are dummies for the respective trader groups.

Table XI displays the result for our estimation of equation (7). We find that, across the board, the imbalance score has explanatory power with regard to the change in liquidity provided whereas the trader types have no power.

We thus conclude that markets Ad and D were populated by different traders: market Ad relied much more on intermediated liquidity provision whereas market D was a market in which traders with offsetting trading intentions were matched.

References

Admati, A., Pfleiderer, P., 1988. A theory of intraday patterns: volume and price variability. Review of Financial Studies 1, 3–40.

Anand, Amber, and Daniel G. Weaver, 2004, Can order exposure be mandated?, Journal of Financial Markets 7, 405 – 426.

Baillie, R.T., Booth, G.G., Tse, Y., Zabotina, T., 2002. Price discovery and common factor models. Journal of Financial Markets 5, 309-321.

Baruch, S., 2005. Who benefits from an open limit order book?. Journal of Business 78, 1267–1306.

Battalio, Robert H. and Corwin, Shane A. and Jennings, Robert H., Can Brokers Have It All? On the Relation between Make Take Fees & Limit Order Execution Quality (March 5, 2014). Available at SSRN:http://ssrn.com/abstract=2367462

Bessembinder, H., Venkataraman, K., 2004. Does an electronic stock exchange need an upstairs market?. Journal of Financial Economics 73, 3-36.

Bessembinder, Hendrik, Marios Panayides, and Kumar Venkataraman, 2009, Hidden liquidity: An analysis of order exposure strategies in electronic stock markets, Journal of Financial Economics 94, 361 – 383.

Bloomfield, R., O'Hara, M., Saar, G., 2005. The "make or take" decision in an electronic market: Evidence on the evolution of liquidity. Journal of Financial Economics 75, 165-199.

Boehmer, E., Saar, G., Yu, L., 2005. Lifting the veil: An analysis of pre-trade transparency at the NYSE. Journal of Finance 60, 783–815.

Boni, L., Brown, D.C., Leach, J.C., 2012. Dark pool exclusivity matters, Unpublished working paper.

Booth, G.G., Lin, J., Martikainen, T., Tse, Y., 2002. Trading and pricing in upstairs and downstairs stock markets. Review of Financial Studies 15, 1111-1135.

Boulatov, A., George, T.J., 2013. Hidden and displayed liquidity in securities markets with informed liquidity providers. Review of Financial Studies 26, 2095-2137.

Brogaard, J., Hendershott, T., Riordan, R. (2014) Market Integration and High Frequency Intermediation. Discussion Paper.

Brolley, Michael, 2015. Should Dark Pools Improve Upon Visible Quotes? The Impact of Trade-at Rules . Discussion paper, University of Toronto.

Buti, S., Rindi, B. Werner, I.M., 2011. Diving into dark pools, Unpublished working paper, Fisher College of Business.

Buti, S., Rindi, B. Werner, I.M., 2013. Dark pool trading strategies, market quality and welfare, Unpublished working paper, Fisher College of Business.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2011, Robust inference with multi-way clustering, Journal of Business Economics and Statistics forthcoming.

Chowdhry, B., Nanda, V., 1991. Multimarket trading and market liquidity. Review of Financial Studies 4, 483–511.

Comerton-Forde, Carole and Putnins, Talis J., 2015, Dark Trading and Price Discovery, *?Journal of Financial Economics,* forthcoming.

CSA (2009) Dark pools, dark orders, and other developments in market structure in Canada, CONSULTATION PAPER 23-404

Degryse, H., de Jong, F., van Kervel, V., 2014. The impact of dark trading and visible fragmentation on market quality. Review of Finance, forthcoming.

De Winne, Rudy, and Catherine D'hondt, 2007, Hide-and-seek in the market: Placing and detecting hidden orders, Review of Finance 11, 663–692.

Eom, K.S., Ok, J., Park, J., 2007. Pre-trade transparency and market quality. Journal of Financial Markets 10, 319–341.

Foucault, T., Menkveld, A.J., 2008. Competition for order flow and smart order routing systems. Journal of Finance 63, 19–58.

Frey, Stephan, and Patrik Sandas, 2009, The impact of iceberg orders in limit order books, working paper University of Virginia.

Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction price in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71-100.

Goettler, R.L., Parlour, C.A., Rajan, U., 2009. Informed traders and limit order markets. Journal of Financial Economics 93, 67-87.

Grossman, S., 1992. The informational role of upstairs and downstairs markets. Journal of Business 65, 509–529.

Grossman, S., Stiglitz, J., 1980. On the impossibility of informationally efficient markets. American Economic Review 10, 393-408.

Hasbrouck, J., 1991. Measuring the information content of stock trades. Journal of Finance 46, 179-207.

Hasbrouck, J., 1995. One security, many markets: Determining the contributions to price discovery. Journal of Finance 50, 1175-1199.

Hasbrouck, J., Saar, G., 2013. Low-latency trading. Journal of Financial Markets 16, 646-679.

Hautsch, N., Huang, R., 2012. The market impact of a limit order. Journal of Economic Dynamics and Control, 36, 501–522.

Hautsch, Nikolaus and Huang, Ruihong, On the Dark Side of the Market: Identifying and Analyzing Hidden Order Placements (February 8, 2012). Available at SSRN: http://ssrn.com/abstract=2004231

Hendershott, T., Jones, C.M., 2005. Island goes dark: Transparency, fragmentation, and regulation. Review of Financial Studies 18, 743-793.

Hendershott, T., Jones, C.M., Menkveld, A.J., 2011. Does algorithmic trading improve liquidity?. Journal of Finance 66, 1-33.

Hendershott, T., Mendelson, H., 2000. Crossing networks and dealer markets: Competition and performance. Journal of Finance 55, 2071- 2115.

Ho, T.S.Y., Stoll, H.R., 1981. Optimal dealer pricing under transactions and return uncertainty. Journal of Financial Economics 9, 47-73.

Hou, K., Moskowitz, T.J., 2005. Market frictions, price delay, and the cross-section of expected returns. Review of Financial Studies 18, 981-1020.

IIROC, 2012, The HOT study, Discussion paper, Investment Industry Regulatory Organization of Canada.

IIROC 2014, Identifying Trading Groups – Methodology and Results, IIROC Notice 14-0210.

Kaniel, R., Liu, H., 2006. So what orders do informed traders use?. Journal of Business 79, 1867-1913.

Korajczyky, R., Murphy, D. (2014) High Frequency Market Making to Large Institutional Trades, Discussion Paper

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315-1335.

Kyle, A.S., 1989. Informed speculation with imperfect competition. Review of Economic Studies 56, 317-356.

Kwan, A., McInish, T.H., Masulis, R., 2015. Trading rules, competition for order flow and market fragmentation. Journal of Financial Economics 115, 330-348.

Madhavan, A., Porter, D., Weaver, D., 2005. Should securities markets be transparent?. Journal of Financial Markets 8, 266–288.

Menkveld, Albert J. and Yueshen, Bart Z. and Zhu, Haoxiang, Shades of Darkness: A Pecking Order of Trading Venues (December 23, 2014). Available at SSRN: http://ssrn.com/abstract=2510707

Nimalendran, M., Ray, S., 2014. Informational linkages between dark and lit trading venues. Journal of Financial Markets 17, 230-261.

O'Hara, M., Ye, M., 2011. Is market fragmentation harming market quality?. Journal of Financial Economics 100, 459-474.

Pardo, Angel, and Roberto Pascual, 2011, On the hidden side of liquidity, The European Journal of Finance 1, 1–19.

Ready, M.J., 2013. Determinants of volume in dark pool crossing networks. Unpublished working paper.

Rindi, B., 2008. Informed traders as liquidity providers: Anonymity, liquidity and price formation. Review of Finance 12, 497-532.

Securities and Exchange Commission (SEC), 2010. Concept Release on Equity Market Structure, Release # 34-61358, File # S7-02-10.

Seppi, D., 1990. Equilibrium block trading and asymmetric information. Journal of Finance 45, 73–94.

Thompson, Samuel B., 2011, Simple formulas for standard errors that cluster by both firm and time, Journal of Financial Economics 99, 1–10.

Tuttle, L., 2013. Alternative Trading Systems: Description of ATS Trading in National Market System Stocks, Division of Economic and Risk Analysis, Securities and Exchange Commission.

Weaver, Daniel G., The Trade-At Rule, Internalization, and Market Quality (April 17, 2014). Available at SSRN: http://ssrn.com/abstract=1846470 or http://dx.doi.org/10.2139/ssrn.1846470

Ye, M., 2012. Price manipulation, price discovery and transaction costs in the crossing network, Unpublished working paper.

Zhu, H., 2014. Do dark pools harm price discovery?. Review of Financial Studies 27, 74

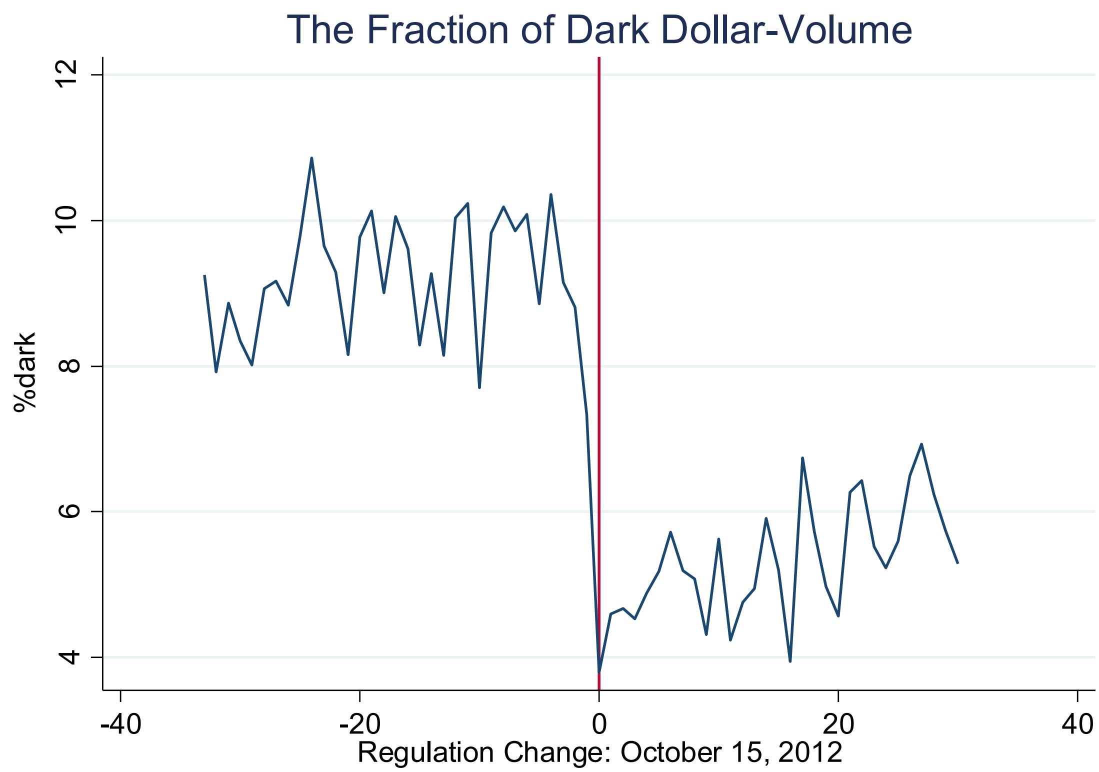

Figure 1: Dark Trading as a Percent of all Trading. The figure plots the average per stock per day dollar trading volume that involves a dark order on the passive side of the trade, as a percentage of the total dollar trading volume in Canada for the August 27 to November 30, 2012 sample period. The vertical line at 0 marks the event date, October 15, 2012.

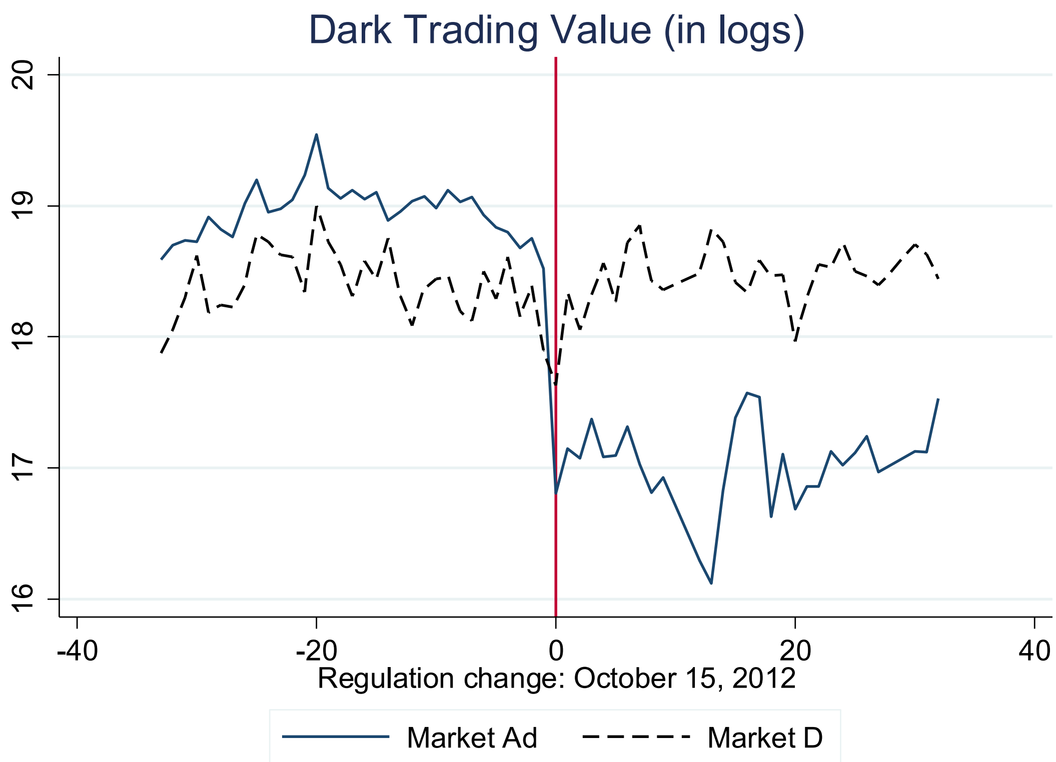

Figure 2: Dollar Volume on the two main dark markets Ad and D. The figure plots the natural logarithm of dollar trading volume that involves a dark order on the passive side of the trade for the two main dark markets, market Ad and market D for the August 27 to November 30, 2012 sample period. The vertical line at 0 marks the event date, October 15, 2012.

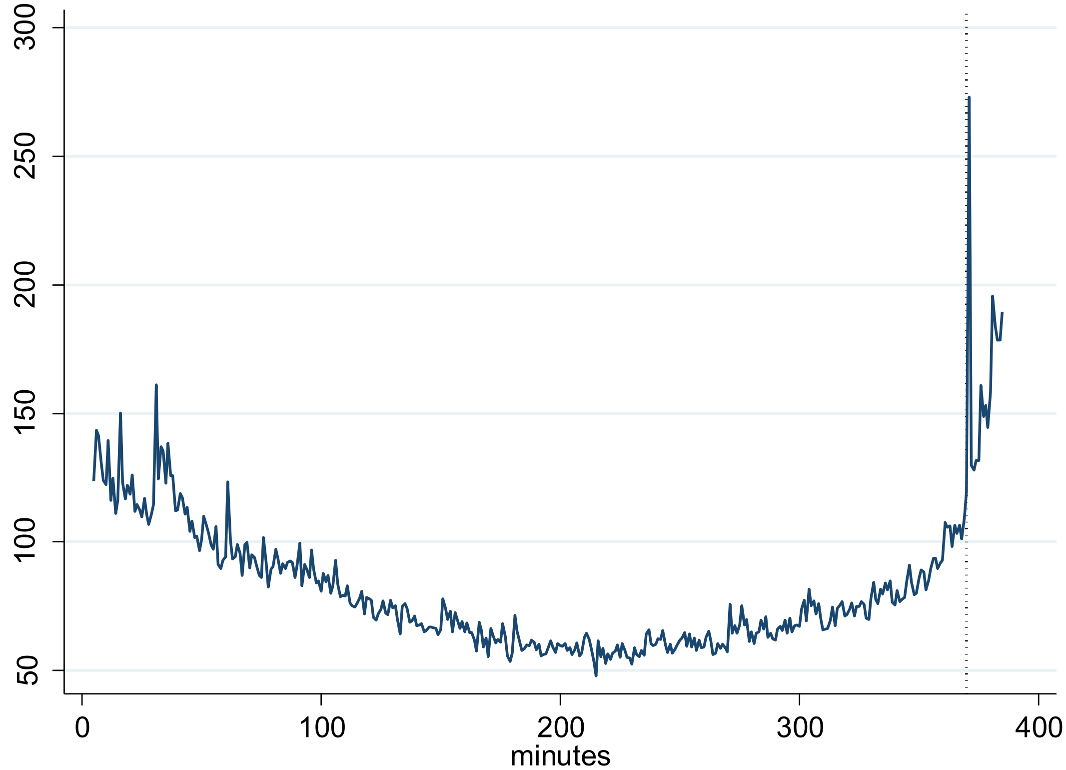

Figure 3: Intra-Day Transactions. The figure plots the aggregate by-minute number of transactions for the securities in our classification sample (92+151 securities) for the July 3 to August 24, 2012 classification period. The plot is for the time between 9:35 a.m. and 3:55 p.m. The dotted line indicates the 370th minute of the trading day, which occurs at 3:40 p.m.

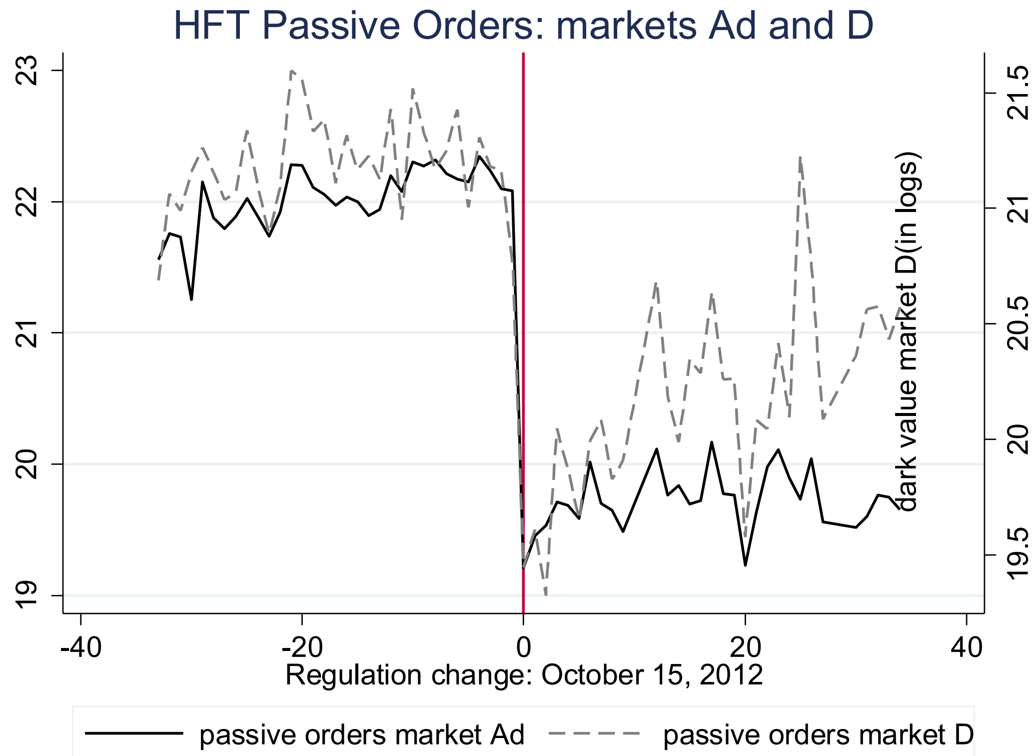

Figure 4: Liquidity Provision by HFTs on Dark Markets Ad and D. The figure plots the logarithm of the value of passive orders by high frequency traders to markets Ad and market D. The left vertical axis measures the order value for market Ad, the right vertical axis measures order value for market D. The plot is for the August 27 to November 30, 2012 sample period. The vertical line at 0 marks the event date, October 15, 2012.

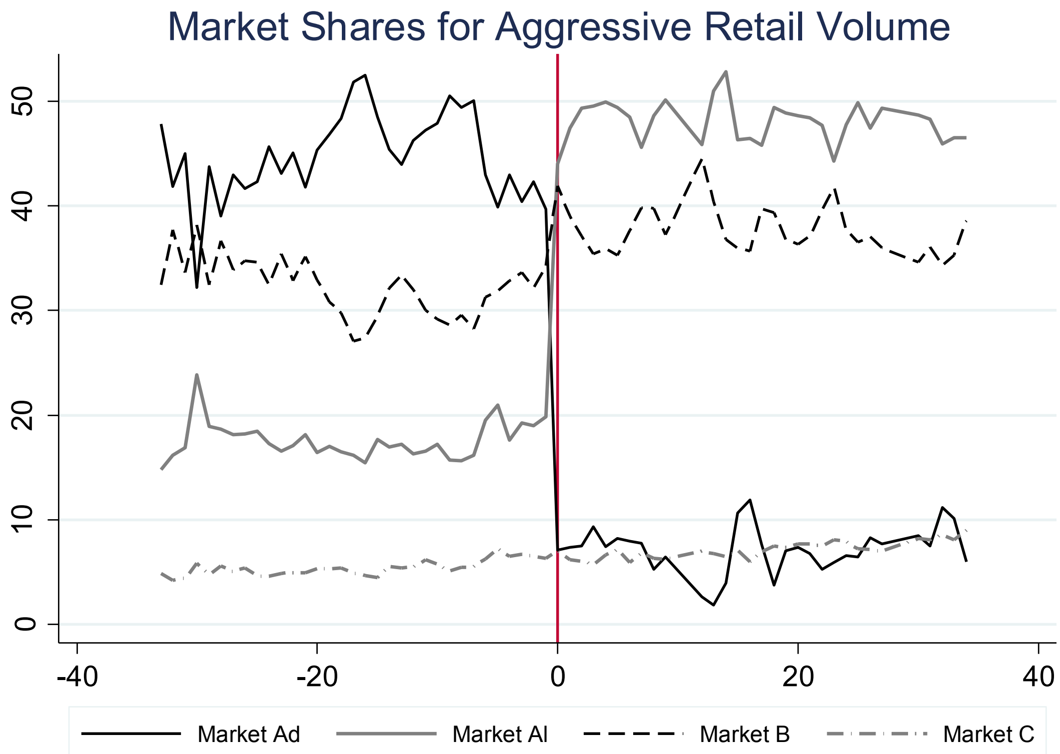

Figure 5: Execution venues for aggressive retail orders. The figure plots the respective percentage of value of aggressive retail orders for market A, split into market Ad (thin black line) and Al (thick, gray line), market B (thin dashed line), and market C (thick, dashed line). The plot is for our sample period from August 27 to November 30, 2012. The vertical line at 0 marks the event date, October 15, 2012.

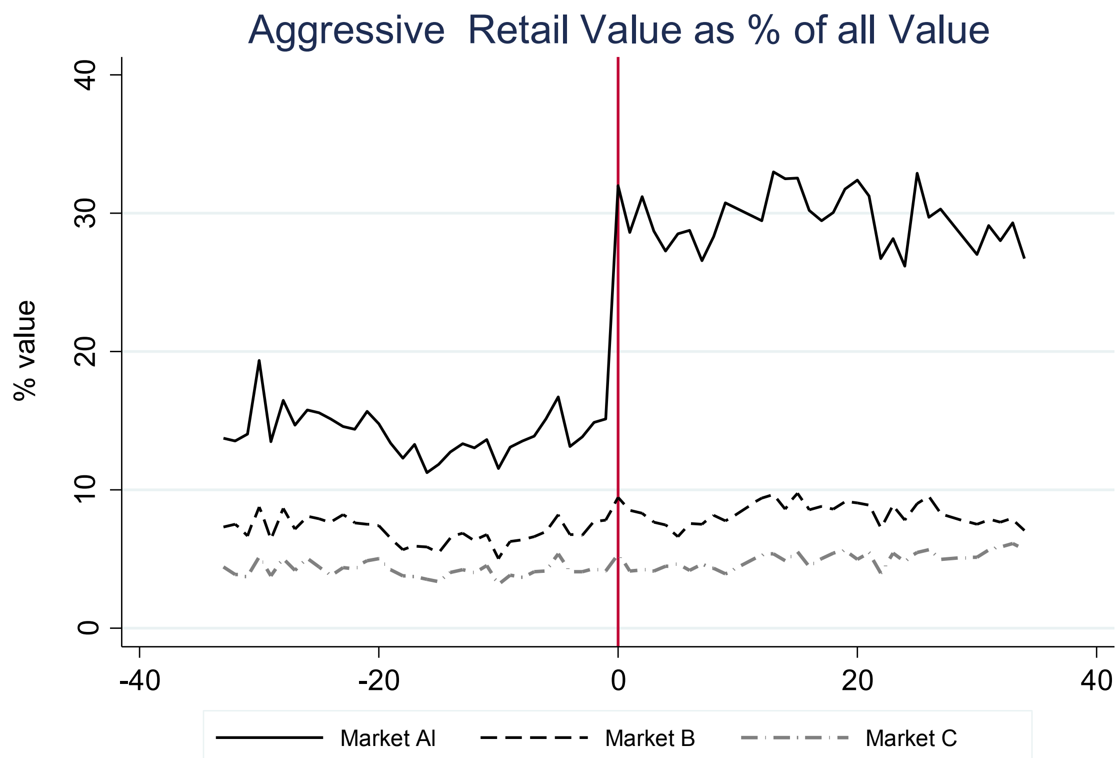

Figure 6: Retail Share of Aggressive Value by Marketplace. The figure plots the respective percentage of value of aggressive retail value in the lit market Al, (thin black line), market B (thin dashed line), and market C (thick, dashed line). The plot is for our sample period from August 27 to November 30, 2012. The vertical line at 0 marks the event date, October 15, 2012.

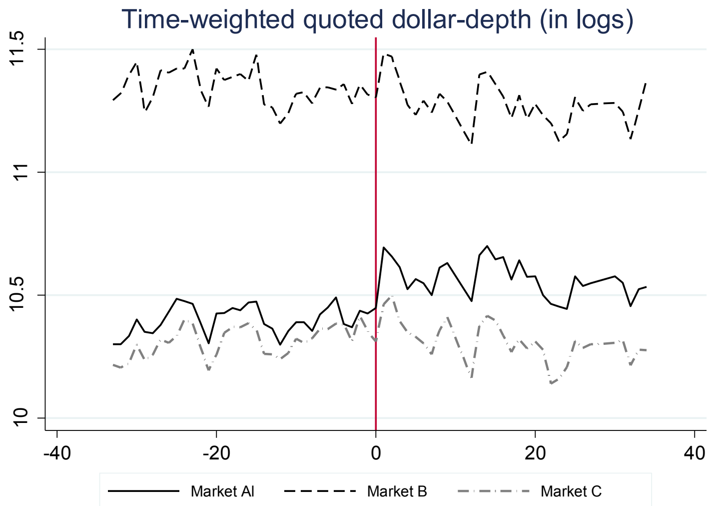

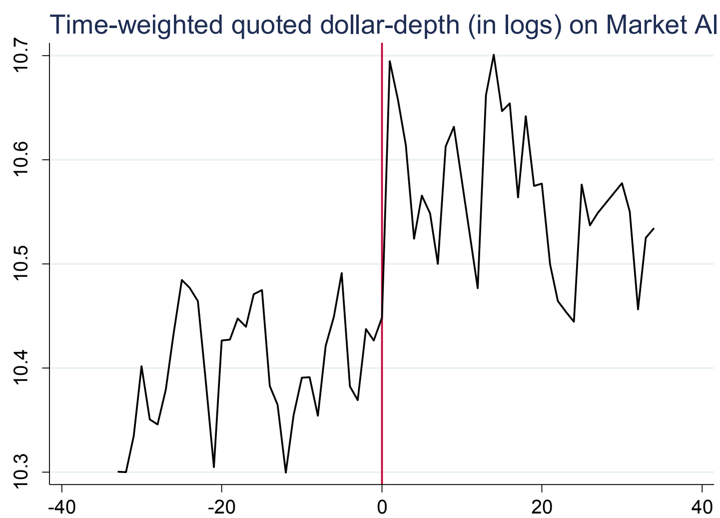

Figure 7: Time-weighted quoted depth. Panel A: The figure plots the natural logarithm of the aggregate time-weighted quoted dollar-depth for the marketplaces A, B, and C. Depth is averaged across all securities per day. Panel B plots log-dollar depth only for Market A. The plot is for the August 27 to November 30, 2012 sample period. The vertical line at 0 marks the event date, October 15, 2012.

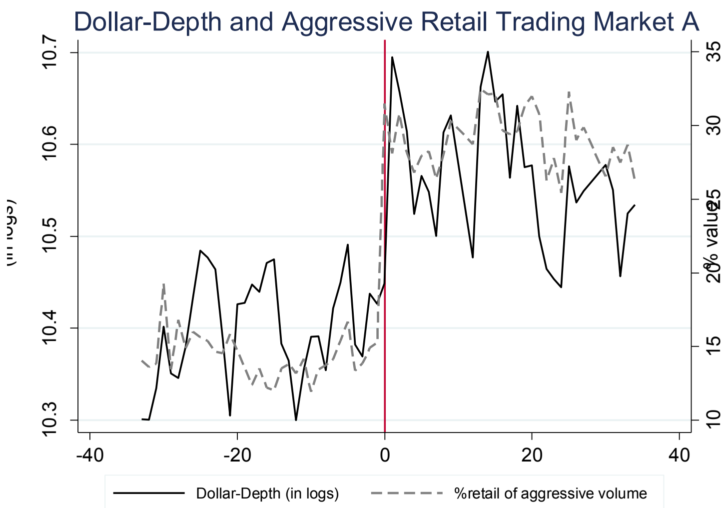

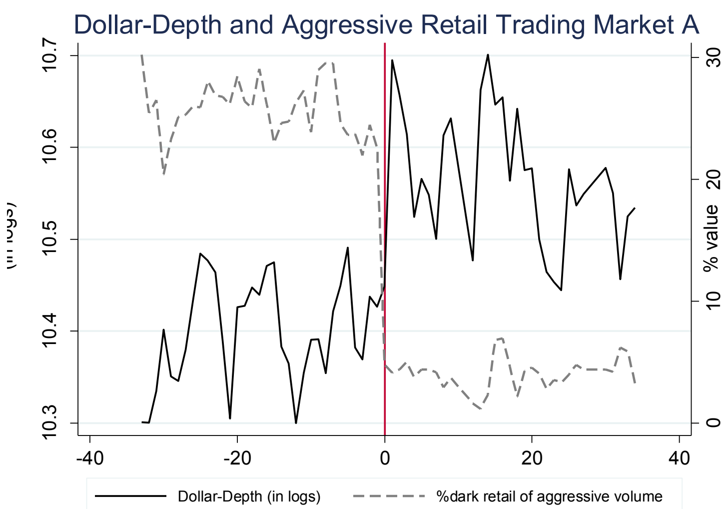

Figure 8: Time-weighted dollar depth and retail volume share. Panel A plots the time-weighted dollar depth (in logs) for market A and the percentage of aggressive value in the lit market Al that is attributed to retail traders. Panel B plots the time-weighted dollar depth (in logs) for market A and the percentage of dark aggressive value of all aggressive value (dark and lit) in market A that is attributed to retail traders. Value figures are based on the aggregated traded value across all securities per day, depth is computed as the average per stock per day. The plot is for the August 27 to November 30, 2012 sample period. The vertical line at 0 marks the event date, October 15, 2012.

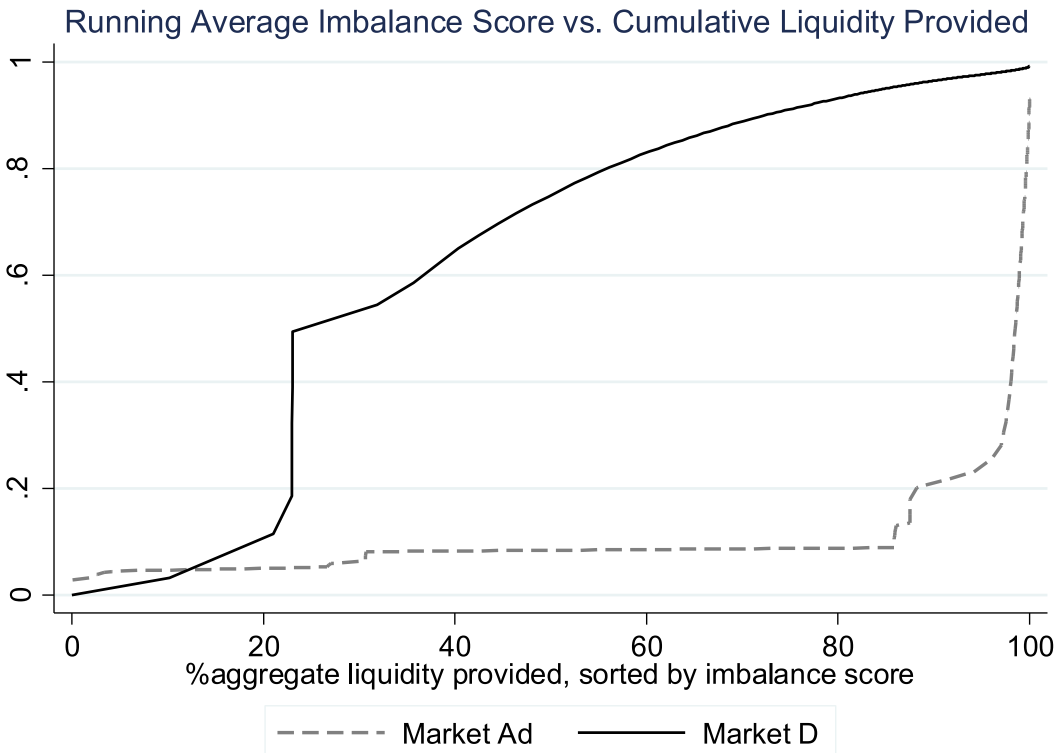

Figure 9: Difference in Characteristics of Liquidity Providers in Markets Ad vs D. The figure plots cumulative percentage liquidity provided against the running average of the imbalance score of the traders that supply the liquidity, where we compute the cumulative volume after sorting traders by their imbalance score. The imbalance score is twice the absolute value of the fraction of buy order volume of all order volume less ½. The plot illustrates that the main liquidity providers in market A submit orders in a much more balanced fashion The plot is based on aggregate dollar-volume before the introduction of the dark liquidity rules, from August 27 to October 15, 2012. The vertical line at 0 marks the event date, October 15, 2012.

Table I: Trading Fees by Marketplace

Table I reports the trading fees by marketplace, depending on the price of the security for lit and dark trades. There are several breakpoint s ($0.10, $1, $5); our sample contains nine securities that trade below $5, but all trade above $1. In the table, negative numbers signify a rebate. For market B, fees depend on the total dollar volume that a broker trades on a particular venue; those brokers with the highest volume receive the most favorable conditions (for the marginal unit traded). All number are in cents per 100 shares.

	lit		dark	
Market	taker/ aggressive	maker/ passive	taker/ aggressive	maker/ passive
Panel A: price>=$5				
Ad			4	0
Al	28	-25		
B	33-35	-32 to -31	10	0
C	29	-25	29	-20
D			10	10
Panel B: price<$5				
Ad			4	0
Al	25	-21		
B	33-35	-32 to -31	10	0
C	29	-25	29	-20
D			5	5

Table II: Summary Statistics for Dark Trading Volume and Dark Order Volume

Table II reports the average per day per security mean for the fraction of dark volume. Variable *PassiveDarkVolume* measures dollar trading volume that involves a dark order on the passive side of the trade as a percentage of the total dollar trading volume; *DarkOrderVolume* measured as the share of volume of dark orders, as a percentage of the share volume of all orders. We refer to this as *DarkOrderVolume*

	before		after	
	PassiveDarkVolume	*DarkOrderVolume*	*PassiveDarkVolume*	*DarkOrderVolume*
Market A	4.6	10.4	0.8	6.4
Market B	1.4	1.0	1.4	1.0
Market C	0.8	1.1	0.8	1.2
Market D	2.5	4.7	2.5	3.3
all others	0.0	0.0	0.0	0.0
total	9.3	17.2	5.4	11.9

Table III: Regression on changes in dark trading and dark order submissions

Table III estimates the effect of MPIR on a market's share of dark trading. We consider two measures: (1) the fraction of dark value of all value, where dark value is defined as the passive side having submitted a dark order; (2) the fraction of order volume of all volume that is submitted as dark. We estimate the effect for the entire market in Panel A and split by market in Panel B. The split in Panel B in a single regression allows us to assess whether the coefficient-estimates differ across markets. These tests show that for both specifications the coefficients for market A and D are statistically significantly different. Both specifications for Panel A include security fixed effects and for Panel B they include security and marketplace fixed-effects. Standard errors are in parentheses and are clustered by time and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	% dark trading volume	%dark order volume
Panel A: Entire Market		
event	-3.84***	-5.20***
	(0.46)	(0.84)
VIX	-0.06	-0.10
	(0.07)	(0.18)
Observations	5,884	5,888
Panel B: By market		
Market A x MPIR	-3.81***	-4.04***
	(0.33)	(0.69)
Market B x MPIR	0.01	-0.04
	(0.09)	(0.12)
Market C x MPIR	-0.03	0.17
	(0.06)	(0.12)
Market D x MPIR	0.00	-1.33***
	(0.16)	(0.28)
other markets x MPIR	0.01	0.02
	(0.02)	(0.04)
VIX	-0.01	-0.02
	(0.01)	(0.04)
Observations	29,142	28,678

Table IV: Summary Statistics Market Quality Measures

Table IV presents summary statistics for our market quality variables. For volatility we consider two measures. (1) realized volatility is measured as the sum of squared 1-minute returns measured in basis points; (2) the range measure is defined as the difference between the largest and the smallest price over 1-minute intervals, measured in basis points of the average price during that interval. We further present the effective spread and several price-impact measures; both are measured in basis points. Standard deviations are in parentheses.

	Before		After	
realized volatility	173.1	(97.3)	184.0	(105.4)
range-measure	7.1	(3.7)	7.5	(4.0)
return autocorrelation	0.1	(0.1)	0.1	(0.1)
effective spread	7.7	(8.1)	7.5	(9.6)
1-second price impact	4.0	(3.6)	4.0	(3.7)
10-second price impact	4.3	(3.9)	4.2	(4.2)
30-second price impact	4.5	(4.2)	4.4	(4.6)
1-minute price impact	4.7	(4.4)	4.6	(4.7)
5-minute price impact	4.9	(5.3)	4.7	(5.6)
10-minute price impact	5.0	(6.1)	4.8	(6.4)

Table V: Impact of the Dark Liquidity Rules on Market Quality

Table V estimates the effect of the introduction of the minimum price improvement rule on the market quality measures that Table IV introduced, namely, realized volatility, the range measure , the absolute value of the 1-minute return autocorrelation, effective spreads and a variety of price impact measures.. All specifications contain security fixed-effects. Standard errors are in parentheses and are clustered by time and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	realized volatility	range measure	return-autocorrelation	effective spread	1 sec price impact	10 sec price impact	30 sec price impact	1 min price impact	5 min price impact	10 min price impact
MPIR	9.44*	0.35	0.01	-0.18	-0.12	-0.14	-0.14	-0.23*	-0.37**	-0.37**
	(5.42)	(0.26)	(0.00)	(0.25)	(0.10)	(0.09)	(0.11)	(0.12)	(0.15)	(0.17)
VIX	1.28	0.01	0	0.06	0.04	0.03	0.03	0.04*	0.08*	0.11*
	(2.05)	(0.10)	(0.00)	(0.04)	(0.03)	(0.02)	(0.03)	(0.03)	(0.04)	(0.06)
Observations	5,884	5,884	5,884	5,884	5,884	5,884	5,884	5,884	5,884	5,884

Table VI: Attributes of classified trader IDs

Table VI summarizes key statistics for the four groups of traders defined in Section VI. The statistics are based on our main sample of 92 cross-listed securities plus the 151 securities that we use for the classification. A*verage end-of-day-inventory* is the average of the medians of the end-of-day inventories for the respective groups of traders, where inventories are computed only for non-crosslisted securities. *Average number of securities* is the average number per trader per day in the group. All other measures are based on the aggregated value or aggregated trades for all securities for the July 3 to August 24, 2012, classification period.

	HFT	Buy-side	Others	retail
Number of IDs	89	558	2860	135
average end-of-day inventory	62.8	97.1	87.6	88.7
average number of securities	136	19	14	35
% value traded	36.1	23.6	30.8	9.4
% of all aggressive value	32.0	23.1	32.3	12.7
% of all passive value	39.6	24.0	29.7	6.7
% of all trades	43.8	21.2	27.2	7.8
% of all aggressive trades	33.9	24.4	29.4	12.4
% of all passive trades	52.8	18.3	25.2	3.7
% of all order-volume	66.5	7.7	24.2	1.7
% of all orders	81.7	3.0	14.7	0.6
% of all order cancellations	83.3	2.2	14.3	0.3
% of all IOC orders	52.6	20.6	20.9	6.0
% of own volume that is passive	54.8	48.8	46.9	38.5
% of own transactions that are passive	50.8	38.2	39.6	25.0

Table VII: Summary Statistics on Trader Returns and Implementation Shortfall

Table VII presents summary statistics for the intra-day returns and implementation shortfall that traders achieve. Returns are computed as *((sell-value – buy-value)+(buy-volume – sell-volume) x closing price)/(total value)*, i.e. they are the intra-day trading profits plus the end-of-day inventory evaluated at the closing price, scaled by the days total dollar-volume. We compute returns for all trades with and without maker-taker fees. The implementation shortfall is computed similarly, except that the closing price is substituted by the price of the day's first trade by the respective trader and that the formula is multiplied with -1, so that a higher the shortfall signifies larger costs of implementing a multi-trade strategy. The fill rate for lit orders is defined as the passive trading volume (in shares) divided by the total lit order volume (in shares); similarly for the dark fill rate.

		Before				After			
		Market A	Market B	Market C	Market D	Market A	Market B	Market C	Market D
intraday return	HFT	1.5	1.4	0.6	1.3	1.4	1.2	0.2	1.4
	retail	-2.9	-2.9	-3.3	-3.8	-5.6	-4.2	-4.9	-1.4
	buy-side	-0.2	0.3	-1.4	-0.6	1.4	1.2	1.6	-2.3
	others	0.3	-1.2	-1.2	-0.8	1.0	-1.0	-1.4	-1.3
intraday return with maker/taker fees	HFT	2.2	1.7	1.0	0.7	1.8	1.9	1.2	0.6
	retail	0.4	0.8	0.0	-3.1	-1.1	-0.2	-2.7	-1.4
	buy-side	0.8	-1.2	-2.2	-2.1	0.0	-1.4	-2.3	-1.5
	others	-7.2	-4.3	-7.2	-2.1	-3.4	-2.9	-5.5	-4.5
shortfall	HFT	-0.3	-0.4	-2.6	-8.0	0.7	2.1	-1.7	-12.6
	retail	-8.8	-22.2	-2.9	-10.7	-12.4	-22.9	-3.7	-18.1
	buy-side	11.4	11.4	7.9	9.0	10.5	11.1	8.9	9.8
	others	2.9	-0.1	0.8	2.2	4.0	1.2	-1.7	7.6
Fill rate lit orders	HFT	1.6	1.4	1.8		2.0	1.4	2.0	
	retail	12.3	16.2	0.3		11.0	15.8	0.1	
	buy-side	10.6	11.2	6.8		9.4	11.1	5.5	
	others	3.8	5.1	5.0	0.0	2.6	4.4	5.9	0.0
Fill rate dark orders	HFT	1.0	6.4	3.1	0.6	0.3	6.1	1.9	0.4
	retail	0.0	29.4		0.1	0.0			0.1
	buy-side	13.3	7.7	8.1	3.4	14.8	5.9	7.8	4.8
	others	1.1	12.3	6.5	2.0	0.8	20.1	5.5	2.3

Table VIII: Regression for Trader Returns and Implementation Shortfall (Part 1)

Table VIII presents estimation results for the effect of the introduction of the minimum price improvement rule on traders' intra-day returns and implementation shortfall, as defined in Table XV. All specifications contain security and marketplace fixed-effects. Standard errors are in parentheses and are clustered by time and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	intraday return				intraday with maker/taker fees			
	HFT	Buy-Side	Others	Retail	HFT	Buy-Side	Others	Retail
Market A x MPIR	0.38	1.53	0.65	-2.42*	0.77	1.42	0.82	-3.51**
	(0.49)	(1.24)	(0.85)	(1.37)	(0.52)	(1.23)	(0.84)	(1.38)
Market B x MPIR	0.15	0.83	0.16	-1.11	0.16	0.88	0.20	-1.16
	(0.69)	(1.08)	(0.53)	(1.59)	(0.68)	(1.06)	(0.54)	(1.60)
Market C x MPIR	-0.02	2.91**	-0.21	-1.41	0.12	2.65*	0.07	-1.49
	(0.45)	(1.37)	(0.97)	(1.74)	(0.47)	(1.37)	(0.94)	(1.74)
Market D x MPIR	0.46	-1.72	-0.51	1.82	0.39	-1.65	-0.57	1.81
	(1.87)	(2.19)	(1.71)	(2.75)	(1.87)	(2.22)	(1.66)	(2.75)
VIX	-0.32*	0.08	0.00	-0.15	-0.32*	0.09	-0.02	-0.16
	(0.19)	(0.40)	(0.19)	(0.58)	(0.19)	(0.41)	(0.19)	(0.58)
Observations	22,336	22,014	22,305	21,277	22,313	21,990	22,284	21,277

Table VIII: Regression for Trader Returns and Implementation Shortfall (Part 2)

	Shortfall			
	HFT	Buy-Side	Others	Retail
Market A x MPIR	0.48	-0.80	0.93	-2.47
	(0.72)	(1.76)	(0.73)	(1.93)
Market B x MPIR	2.01**	-0.05	1.10	0.31
	(0.97)	(1.13)	(0.70)	(1.89)
Market C x MPIR	0.35	1.08	-2.66**	0.11
	(0.61)	(1.68)	(1.20)	(2.18)
Market D x MPIR	-5.14**	0.74	5.10***	-5.80**
	(2.57)	(1.96)	(1.87)	(2.84)
VIX	0.44	-0.16	0.14	-0.90
	(0.33)	(0.46)	(0.27)	(0.63)
Observations	22,864	22,526	22,830	21,779

Table IX: Regression for Fill Rates

Table IX presents estimation results for the effect of the introduction of the minimum price improvement rule on traders' fill rates for *passive* orders, as defined in Table VIII. All specifications contain security and marketplace fixed-effects. Standard errors are in parentheses and are clustered by time and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	fill rate passive lit orders				Fill rate passive dark orders			
	HFT	Buy-Side	Others	Retail	HFT	Buy-Side	Others	Retail
Market A x MPIR	0.34***	-1.40***	-1.05***	-1.35***	-0.39***	1.88	-0.48	-0.01
	(0.09)	(0.46)	(0.23)	(0.42)	(0.14)	(2.13)	(0.47)	(0.01)
Market B x MPIR	0.08	-0.22	-0.52**	-0.42	0.04	-1.51	7.26	
	(0.06)	(0.36)	(0.21)	(0.53)	(0.62)	(1.00)	(4.78)	
Market C x MPIR	0.23**	-1.56***	1.06***	-0.16	-1.01***	0.08	-1.09*	
	(0.10)	(0.41)	(0.32)	(0.12)	(0.27)	(0.85)	(0.60)	
Market D x MPIR					0.05	1.71***	0.18	-0.01
					(0.12)	(0.36)	(0.43)	(0.04)
VIX	-0.02	0.17**	-0.12***	0.00	-0.20***	-0.35*	0.11	0.01
	(0.02)	(0.08)	(0.04)	(0.10)	(0.06)	(0.21)	(0.35)	(0.01)
Observations	17,554	17,483	19,667	17,283	22,716	18,176	21,410	11,365

Table X: Summary Statistics for trader types usage of trading venues (Part 1)

Table X summarizes where different types of traders trade. The column dominator indicates the denominator used to compute the fraction of trades. For each group of traders we examine all combinations of dark & lit and aggressive & passive trades. We differentiate by aggressive trades and by dark vs passive trades. An SDL order that executes in the lit market Al is counted as an aggressive lit trade. All value figures are for the means of the aggregate volume across all securities for that type per day.

		Before MPIR					After MPIR				
	Denominator	Market A	Market B	Market C	Market D	other	Market A	Market B	Market C	Market D	other
Panel A: HFT											
dark & aggressive	all volume by HFT	0.0	0.0	0.3	1.1	0.0	0.0	0.0	0.5	1.6	0.0
dark & passive	all volume by HFT	1.6	0.3	0.1	0.4	0.0	0.1	0.2	0.1	0.1	0.0
lit & aggressive	all volume by HFT	7.4	25.4	7.5	n/a	1.3	7.6	23.9	7.4	n/a	1.3
lit & passive	all volume by HFT	9.8	28.1	14.3	n/a	2.5	13.4	26.1	15.2	n/a	2.5
aggressive	all aggressive volume by HFT	17.3	59.2	18.1	2.5	2.9	18.0	56.4	18.8	3.8	3.1
passive	all passive volume by HFT	20.0	49.7	25.3	0.6	4.3	23.3	45.6	26.5	0.2	4.3
dark	all dark volume by HFT	42.2	7.6	11.6	38.6	0.0	3.6	8.0	23.1	65.3	0.0
lit	all lit volume by HFT	17.9	55.6	22.6	n/a	3.9	21.5	51.3	23.3	n/a	3.9
Panel B: Retail											
dark & aggressive	all volume by retail	27.6	0.0	0.0	1.5	0.0	4.4	0.0	0.0	1.6	0.0
dark & passive	all volume by retail	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
lit & aggressive	all volume by retail	11.4	20.4	3.4	n/a	1.5	29.7	23.5	4.4	n/a	1.9
lit & passive	all volume by retail	13.6	20.6	0.0	n/a	0.0	13.0	21.4	0.0	n/a	0.0
aggressive	all aggressive volume by retail	59.4	30.9	5.1	2.2	2.4	52.0	35.8	6.8	2.5	3.0
passive	all passive volume by retail	39.7	60.2	0.0	0.1	0.0	37.8	62.1	0.0	0.1	0.0
dark	all dark volume by retail	94.9	0.0	0.0	5.1	0.0	71.1	0.0	0.0	28.9	0.0
lit	all lit volume by retail	35.2	57.8	4.8	n/a	2.2	45.5	47.7	4.7	n/a	2.1
Panel C: buy side											
dark & aggressive	all volume by buy-side	0.0	0.1	0.4	1.4	0.0	0.0	0.1	0.6	1.0	0.0
dark & passive	all volume by buy-side	0.6	1.7	0.6	3.6	0.0	0.7	1.7	0.5	5.3	0.0
lit & aggressive	all volume by buy-side	6.5	33.8	10.5	n/a	2.5	7.1	31.6	11.3	n/a	2.5
lit & passive	all volume by buy-side	3.6	31.8	1.8	n/a	1.3	3.9	30.6	1.5	n/a	1.6
aggressive	all aggressive volume by buy-side	11.8	61.4	19.8	2.5	4.5	13.1	58.4	22.0	1.9	4.6
passive	all passive volume by buy-side	9.4	74.4	5.2	8.0	3.0	10.1	70.5	4.4	11.5	3.4
dark	all dark volume by buy-side	7.1	20.8	12.0	60.2	0.0	7.5	17.9	11.1	63.5	0.0
lit	all lit volume by buy-side	11.0	71.5	13.3	n/a	4.2	12.2	69.0	14.3	n/a	4.5

Table XI: Regression for by-trade usage of trading venues (Part 1: HFTs)

Table XI estimates the effect of MPIR on the usage of venues by traders. The variables used are as defined in Table VI, except that the measures are computed per security per day.. For instance, we estimate that HFTs trade 1.6% less volume with passive, dark orders on market A. All specifications contain fixed effects for securities and marketplaces. Standard errors are in parentheses and are clustered by time and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	dark & aggressive as a fraction of all volume by HFT	dark & passive as a fraction of all volume by HFT	lit & aggressive as a fraction of all volume by HFT	lit & passive as a fraction of all volume by HFT	aggressive as a fraction of all aggressive volume by HFT	passive as a fraction of all passive volume by HFT	dark as a fraction of all dark volume by HFT	lit as a fraction of all lit volume by HFT
Panel A: HFT								
Market A x MPIR dummy	0.03***	-1.60***	0.11	3.05***	0.50**	2.24***	-31.61***	2.96***
	(0.01)	(0.19)	(0.13)	(0.37)	(0.25)	(0.38)	(2.09)	(0.38)
Market B x MPIR dummy	0.00	0.01	-0.95***	-1.69***	-1.99***	-2.98***	0.29	-3.35***
	(0.01)	(0.07)	(0.36)	(0.37)	(0.53)	(0.61)	(0.86)	(0.54)
Market C x MPIR dummy	0.20***	-0.04	0.17	0.91***	0.83*	1.51***	15.91***	0.67
	(0.04)	(0.02)	(0.19)	(0.32)	(0.43)	(0.54)	(2.21)	(0.49)
Market D x MPIR dummy	0.41***	-0.38***	0.02	-0.02	0.95***	-0.54***	16.55***	-0.00
	(0.08)	(0.09)	(0.03)	(0.03)	(0.19)	(0.12)	(2.14)	(.)
other markets x MPIR dummy	-0.00	0.00	-0.08	-0.11	-0.20	-0.22	-0.01	-0.27
	(0.01)	(0.01)	(0.07)	(0.13)	(0.16)	(0.22)	(0.02)	(0.19)
VIX	0.00	-0.00	-0.02	0.02	0.00	-0.00	0.01	0.00
	(0.01)	(0.01)	(0.02)	(0.02)	(0.00)	(0.00)	(0.02)	(.)
Observations	29,142	29,142	29,142	29,142	29,142	29,142	28,811	29,142

Table XI: Regression for by-trade usage of trading venues (Part 2: Retail)

	dark & aggressive	dark & passive	lit & aggressive	lit & passive	aggressive	passive	dark	lit
Panel B: Retail								
Market A x MPIR dummy	-18.23***	-0.00	13.51***	-0.47**	-7.00***	-1.63***	-24.74***	8.53***
	(1.21)	(0.00)	(0.95)	(0.20)	(1.10)	(0.48)	(2.00)	(0.68)
Market B x MPIR dummy	0.10**	-0.01	3.34***	0.53	5.17***	1.63***	-0.03	-8.05***
	(0.05)	(0.00)	(0.61)	(0.65)	(0.89)	(0.49)	(0.03)	(0.67)
Market C x MPIR dummy	0.10**	-0.00	1.00***	-0.08	1.54***	-0.01	-0.00	0.11
	(0.05)	(0.00)	(0.19)	(0.08)	(0.25)	(0.03)	(0.01)	(0.20)
Market D x MPIR dummy	0.37**	0.01	-0.03	-0.07	0.39*	-0.00	25.31***	-0.00
	(0.16)	(0.02)	(0.07)	(0.07)	(0.20)	(0.07)	(2.02)	(0.00)
other markets x MPIR dummy	0.10**	-0.00	-0.08	-0.07	-0.07	0.01	-0.00	-0.58***
	(0.05)	(0.00)	(0.13)	(0.07)	(0.18)	(0.01)	(0.01)	(0.14)
VIX	-0.09**	0.00	0.03	0.06	0.00	0.00	0.00	0.00
	(0.03)	(0.00)	(0.06)	(0.06)	(0.00)	(0.00)	(0.01)	(0.00)
Observations	29,071	29,071	29,071	29,071	29,049	28,840	27,475	29,062

Table XI: Regression for by-trade usage of trading venues (Part 3: Buy-Side)

	dark & aggressive	dark & passive	lit & aggressive	lit & passive	aggressive	passive	dark	lit
Panel C: Buy-Side								
Market A x MPIR dummy	0.00	0.21	0.76***	0.03	1.26***	0.82**	1.51	1.03***
	(0.01)	(0.13)	(0.16)	(0.19)	(0.29)	(0.41)	(1.24)	(0.30)
Market B x MPIR dummy	0.02	0.07	-1.79***	-1.00*	-3.38***	-1.44*	0.19	-2.22***
	(0.01)	(0.10)	(0.55)	(0.59)	(0.70)	(0.77)	(0.75)	(0.58)
Market C x MPIR dummy	0.21***	0.01	1.51***	-0.41***	2.87***	-1.16***	1.92**	1.37***
	(0.05)	(0.05)	(0.35)	(0.16)	(0.58)	(0.43)	(0.88)	(0.38)
Market D x MPIR dummy	-0.47***	1.21***	0.01	-0.05	-0.68***	2.65***	-3.55**	-0.00*
	(0.12)	(0.20)	(0.04)	(0.05)	(0.20)	(0.45)	(1.69)	(0.00)
other markets x MPIR dummy	-0.00	0.04	0.02	-0.28**	-0.06	-0.66*	-0.01	-0.17
	(0.01)	(0.02)	(0.15)	(0.14)	(0.23)	(0.36)	(0.01)	(0.24)
VIX	0.00	-0.03	-0.01	0.04	0.00	0.00	0.01	0.00**
	(0.01)	(0.02)	(0.04)	(0.04)	(0.00)	(0.00)	(0.01)	(0.00)
Observations	29,137	29,137	29,137	29,137	29,121	29,102	26,786	29,137

Table XII: The Effect of MPIR on Market Quality by Market

Table XII presents the results of an estimation of the effect of the introduction of the minimum price improvement rule on time-weighted depth and spreads for the three main lit markets A, B and C. We estimate the effect for all three markets simultaneously to capture when markets are differently affected. The dependent variables are the time-weighted quoted spread in cents and in basis points of the prevailing price, the log of share depth and the log of dollar-depth. Independent variables are dummy variables for each market interacted with the dummy for the introduction of MPIR. All specifications contain security and marketplace fixed-effects. Standard errors are in parentheses and are clustered by time and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	time-weighted quoted spread		time-weighted quoted depth	
	in cents	in BPS	in $ (logs)	in shares (logs)
Market A x MPIR	0.36	0.80*	0.16***	0.16***
	(0.24)	(0.45)	(0.02)	(0.02)
Market B x MPIR	0.06	0.09	-0.05**	-0.04
	(0.18)	(0.23)	(0.02)	(0.02)
Market C x MPIR	-1.14	0.04	0.02	0.03
	(1.02)	(0.35)	(0.02)	(0.02)
VIX	-0.01	0.29***	-0.01**	-0.01*
	(0.14)	(0.10)	(0.01)	(0.00)
Observations	17,664	17,664	17,664	17,664

Table XIII: Summary Statistics for Liquidity Supply and Demand in Dark Markets by Trader Type

Table XIII provides summary statistics for the distribution of aggressive and passive trading by trader types on the two dark markets Ad and D before and after the rule change on October 15, 2012. All figures are computed as the percentage of total value per day for that market. SDL orders that execute in the lit book of market A count as aggressive and lit orders. In market D, we exclude the periodically-matched trades and only consider those where an aggressive side can be identified.

		in percent of total value for the marketplace				In $million			
		Market Ad		Market D		Market Ad		Market D	
		before	after	before	after	before	after	before	after
HFT	aggressive	0.1	4.9	40.5	56.2	0.2	1.3	32.4	47.0
	passive	27.3	7.3	14.0	4.9	46.9	1.8	10.9	4.0
Retail	aggressive	99.9	94.8	11.6	11.7	175.7	25.7	9.2	9.7
	passive	0.0	0.0	0.2	0.2	0.0	0.0	0.1	0.2
Buy-side	aggressive	0.0	0.1	18.1	12.1	0.0	0.0	14.4	10.2
	passive	3.5	25.1	47.3	61.0	6.3	6.5	37.8	51.4
Others	aggressive	0.0	0.2	29.8	20.1	0.0	0.1	24.0	16.7
	passive	69.3	67.6	38.5	33.9	122.6	18.8	31.2	28.0

Table XIV: Regressions for Changes in Liquidity Supply and Demand in Dark Markets by Trader Type

Table XIV provides the results from a regression where we estimate the effect of the introduction of the minimum price improvement rule on the fractions of liquidity supply and demand by trader type. Variables are constructed as described in Table VIII. All specifications contain security and marketplace fixed-effects. Standard errors are in parentheses and are clustered by time and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	HFT		Retail		Buy Side		Others	
	aggressive	passive	aggressive	passive	aggressive	passive	aggressive	passive
Market Ad x MPIR	8.56***	-19.27***	-9.10***	-0.03	0.03	14.79***	0.51**	4.52
	(1.37)	(2.65)	(1.38)	(0.04)	(0.14)	(1.98)	(0.21)	(3.28)
Market D x MPIR	17.57***	-12.53***	0.33	0.09	-7.88***	12.91***	-10.02***	-0.48
	(1.55)	(2.02)	(0.79)	(0.13)	(0.82)	(2.05)	(1.06)	(2.06)
VIX	0.42	0.57	-0.14	0.03	0.04	-0.18	-0.32**	-0.43
	(0.29)	(0.39)	(0.21)	(0.04)	(0.11)	(0.42)	(0.15)	(0.50)
Observations	11,015	11,015	11,015	11,015	11,015	11,015	11,015	11,015

Table XV: Characteristics of Liquidity Provision in Dark Markets

Table XV provides information about the type of liquidity providers in the two dark markets. We sort traders by liquidity provided and then determine the number needed for 95% of liquidity, where liquidity is measured by passive-side aggregated dollar volume. The imbalance score is defined as *Imbalance Score=2 x |buy-ordervolume/total ordervolume*-0.5|, so that a score of 100% implies that the trader only submits passive buy- or only sell-order. The imbalance score presented here is the median of the trader's per-security per-day scores, the volume based figures are based on the total aggregate dollar-volume over all securities and all days.

		Market Ad		Market D	
		before	after	before	after
total number of liquidity providers		307	284	588	612
95% of liquidity provided by how many traders?		14	46	185	214
average imbalance score of top 95%		28.6		98.8	
among the top-95% how many are	HFT	4	4	5	5
	retail	0	0	2	1
	buy-side	2	14	78	79
	others	8	28	100	129
%liquidity provided by top 95% traders	HFT	25.3	4.2	13.5	4.7
	retail	0	0	0.1	0.1
	buy-side	2.9	23.4	45.2	59
	others	67.3	67.4	36.1	31.2

Table XVI: Regressions for Changes in Liquidity Supply in Dark Markets based on Imbalance Score

Table XVI tests whether the imbalance score (see Table XV) has explanatory power in terms of the changes in liquidity supply. The dependent variable is the percentage of liquidity provided by the trader before MPIR minus the percentage provided after MPIR. We interact the imbalance score with dummies for markets Ad and D to be able to test whether the coefficients are equal. Post-estimation tests show that we cannot reject the hypothesis that the coefficients coincide. Standard errors are in parentheses. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	%liq prov. Before - %liq provided after	%liq prov. Before - %liq provided after
Imbalance score x market Ad	-3.17***	-3.21***
	(0.51)	(0.56)
Imbalance score x market D	-3.01***	-3.04***
	(0.50)	(0.55)
HFT		-0.11
		(0.43)
buy-side		-0.07
		(0.16)
retail		0.00
		(0.86)
Constant	2.98***	3.04***
	(0.49)	(0.55)
Observations	895	895

Table XVII: Cross-Sectional Regression to Predict Changes in Quoted Depth

Table XVII tests whether the extent of retail trading in dark market Ad predicts the change in quoted depth on lit market Al after the introduction of MPIR. The change in depth is the different of the average per security per day time-weighted quoted depth in dollars after and before the introduction of MPIR. *%retailvolume* is the average per security per day fraction of market Ad volume of all market A volume before the introduction of MPIR; *totalretailvolume* is the average per security per day dollar-volume in market Ad before MPIR. Robust standard errors are in parentheses. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	Change in Depth			
%retailvolume	269.51***	253.37***		
	(83.14)	(82.38)		
totalretailvolume			0.002***	0.0025***
			(0.00)	(0.00)
log MarketCap		1401.42*		-1699.19**
		(732.38)		(805.16)
Constant	-374.83	31521.79*	2092.39**	39266.61**
	(1341.49)	(16031.85)	(898.12)	(17640.37)
Observations	92	92	92	92
Adj R-squared	0.092	0.115	0.259	0.279