

Regulating Dark Trading: Order Flow Segmentation and Market Quality*

Carole Comerton-Forde^a, Katya Malinova^b and Andreas Park^{c*}

^a *Department of Finance, University of Melbourne*

^b *University of Toronto*

^c *University of Toronto Mississauga and Rotman School of Management*

September 3, 2016

Abstract

We examine the impact of a rule in the Canadian equities market that requires dark orders to offer price improvement over displayed orders. We show that this rule eliminated the intermediation of retail orders in the dark and shifted retail orders onto the lit market with the lowest take fee. Intermediaries shifted liquidity supply to this venue leading to an increase in displayed liquidity. We conclude that reducing retail order segmentation enhances lit liquidity. However, retail traders receive less price improvement, retail brokers pay higher exchange fees, and institutions incur higher implementation shortfall. High frequency traders earn higher fee revenues.

JEL classification: G14

Keywords: dark trading, high frequency trading, retail investors, trade-at rule, price improvement, segmentation, internalization

* This paper is based on the research report titled “The Impact of the Dark Trading Rules” that we prepared for the Investment Industry Regulatory Organization of Canada (IIROC). We thank IIROC for providing access to the data and for financial assistance, and in particular thank Victoria Pinnington and Helen Hogarth for assistance. We also thank Qizheng (Alan) Yuan for research assistance. We thank Peter Bossaerts, Patrik Sandas, Haoxiang Zhu, participants at the 2015 University of Cambridge Workshop on Microstructure Theory and Application, the Stern Microstructure Meeting 2015, the 7th Erasmus Liquidity Conference, the Women in Microstructure Meeting 2015, the IIROC HFT Forum, the 2016 European Finance Association Meetings, and seminar participants at ESSEC, University of Auckland, Deakin University, University of Illinois at Chicago, University of Copenhagen, Warwick University, and Stockholm Business School. Comerton-Forde is an economic consultant on market structure for the Australian Securities and Investments Commission. Park is a member of the Ontario Security Commission’s Market Structure Advisory Committee. Part of this work was done while Malinova and Park visited Copenhagen Business School, and they thank CBS for its hospitality. Support from the Center for Financial Frictions (FRIC), grant no. DNRF102, Nordea Fonden, SSHRC, and from the Global Risk Institute for the project “Risk Management and Market Liquidity” is gratefully acknowledged.

Corresponding author: Carole Comerton-Forde. Email: carole.comerton-forde@unimelb.edu.au

An empirical fact, not well known outside market microstructure circles, is that marketable retail orders in US equity markets are not typically executed on stock exchanges. Instead, these orders are routed to wholesale market makers and the retail brokers making these routing decisions receive payment for this order flow.¹ This practice is commonly referred to as retail internalization and represents segmentation of retail order flow. Rosenblatt Securities estimates that retail internalization currently accounts for approximately 16 percent of consolidated US equity market volumes.²

Proponents of internalization argue that it enhances market quality by increasing the number of agents competing to execute order flow. However, opponents suggest that it harms market quality by allowing internalizers to segment or “cream-skim” the uninformed investors away from the lit market, leading to higher bid ask spreads and lower displayed depth. The impact of retail internalization was extensively analyzed in the academic literature and by the SEC in the late 1990s, but, the empirical results are mixed. The SEC’s 1997 “Report on the Practice of Preferencing” finds that internalization does not harm market quality, but the report stops short of saying that it enhances it. Easley, Kiefer, and O’Hara (1996) and Bessembinder and Kaufman (1997), among others, find that the presence of payment for order flow arrangements lowers market quality and harms uninformed traders, whereas Battalio (1997) and Battalio, Greene, and Jennings (1997) argue, respectively, that the introduction of purchasing and internalizing dealers does not have an adverse effect on transaction costs. Hansch, Naik, and Viswanathan (1999) find no relation between the level of internalization and bid-ask spreads, but show internalized trades pay lower spreads than non-internalized trades.

¹ The SEC 2010 Concept Release on Equity Market Structure states “A review of the order routing disclosures required by Rule 606 of Regulation NMS of eight broker-dealers with significant retail customer accounts reveals that nearly 100% of their customer market orders are routed to OTC market makers.” The Concept Release estimates the amount of payment for order flow is 0.1 cent per share or less.

² See Let there be light: Rosenblatt’s Monthly US Dark Liquidity Tracker, 24 May 2016. They report that these volumes are executed by five wholesale market makers: Citadel Securities, KCG Americas, G1 Execution Services, UBS Securities, Two Sigma Securities and Automated Trading Desk.

There have been significant changes in market structure since these studies were undertaken. First, trading shifted from trading floors to trading screens. Second, tick sizes were dramatically reduced from 1/8th dollar to 1 cent. Third, Reg ATS and Reg NMS substantially increased the level of competition between exchanges and other trading venues. Fourth, there has been growth in high frequency trading and dark pools. Together these changes have markedly increased the complexity of market structure suggesting a need for a re-examination of the impact of segmentation of retail order flow on market quality. Further, internalization is again at the center of the regulatory and policy debate. Regulatory authorities in the US are reportedly investigating whether wholesale market makers are delivering best execution to their clients.³

Regulators globally are also concerned about the other types of dark trading, namely trading in dark pools and hidden liquidity on exchanges.⁴ Concerns raised by regulators include the possibility that dark trading undermines price discovery, reduces lit liquidity and segments informed and uninformed order flow. However, the prevalence of dark trading indicates that it is a critically important part of modern markets, requiring that regulators determine the most appropriate regulatory regime to ensure that they balance the benefits associated with dark trading with any potential costs. Like the empirical literature on internalization, the evidence of the impact of other types of dark trading on market quality is also mixed. Degryse, de Jong and van Kervel (2015) study internalized and dark pool trading, reporting a negative association between dark trading and market-wide liquidity measures. Comerton-Forde and Putnins (2015) show that while small amounts of dark trading improve price discovery, large amounts hinder it. They also show that block dark trades have no adverse impact on price discovery. Buti, Rindi

³ Reuters on 10 May 2016 (<http://www.reuters.com/article/us-usa-stocks-probe-exclusive-idUSKCN0Y11CJ>) and Bloomberg on 17 May 2016 (<http://www.bloomberg.com/news/articles/2016-05-16/quicktake-q-a-why-payments-for-stock-trades-face-more-scrutiny>).

⁴ In the US, Rosenblatt Securities estimate that dark pools and hidden orders on exchanges account for approximately 16.37% and 5.63% of consolidated volumes, respectively.

and Werner (2011) examine dark pool trading and find higher dark pool activity in more liquid securities, and, contemporaneously, more dark trading when spreads are narrow and depth high.

We contribute to this literature and to the regulatory debate by examining the introduction of a price improvement rule that imposes restrictions on dark trading in the Canadian equity market.⁵ The minimum price improvement rule (MPIR) dramatically impacts dark trading. In the weeks following the introduction of the rule on October 15, 2012 the share of dark activity declines sharply, from 9.3% to 5.4% of dollar trading volume (excluding pre-arranged block trades, which were unaffected by the new regulation). We exploit this rule change and the institutional features of the Canadian market to expand our understanding of the impact of dark trading and the segmentation of retail order flow on market quality.

The Canadian market comprises multiple competing venues, with substantial variation in market structure. Two dark pools, market Ad and D, account for three-quarters of dark trading before the rule change. The price improvement rule impacts market Ad and market D very differently. After the introduction of the rule, market Ad experiences a significant decline in its volume share from 4.6% to 0.8%, whereas volume on market D remains unchanged at 2.5%. We use proprietary trader-level data to examine whether this difference relates to variation in the types of traders and the nature of liquidity provision in these markets. We classify traders into four categories: retail, high frequency, institutional and other. We also calculate an intermediation score for each trader, which allows us to identify traders that act as intermediaries by consistently posting limit orders on both sides of the market in Ad and D during the pre-MPIR period. We label these traders as dark pool market makers (DPMM). We show that 88% of liquidity supply on market Ad is from intermediaries (26% from high frequency and 62% from other DPMM), while all aggressive order flow is from retail traders. In contrast, on market D, less than 15% of liquidity supply is from intermediaries and the demand for liquidity is split across the four trader types.

⁵ In the US this type of rule is referred to as a trade-at rule.

We hypothesize that these differences in the levels of intermediation in markets Ad and D will influence the impact of the MPIR, because the intermediaries will withdraw from the market when the spreads earned for liquidity provision are constrained by the MPIR. We show that after the introduction of the MPIR, the levels of intermediation fell to 38% and 1% in market Ad and market D, respectively. We find that the intermediation score is significant in predicting whether or not a trader reduces their liquidity provision in the dark after the introduction of the MPIR, while the trader types are not significant after controlling for the intermediation score. Further, the fraction of intermediated volume in the dark predicts the decline in dark trading after the MPIR change. Stocks with higher levels of intermediation exhibit larger declines in the level of dark trading.

An outstanding question is what happens to the order flow that was routed to market Ad prior to the rule change? It is unlikely that retail investors are aware of the introduction of the MPIR given that their brokers typically handle routing decisions, therefore we do not expect any change in retail order flow (to brokers) as a result of the rule change. Knowing that the rule reduced incentives for liquidity supply on market Ad, we expect that retail brokers will route retail orders to the market with the lowest take fee, which is the lit market Al.⁶ Consistent with this expectation we document a significant increase in aggressive retail flow on market Al after the rule change: the fraction of retail volume increases from around 15% of Al's dollar volume to 30%. The high frequency traders that had been acting as DPMM in market Ad before the rule change anticipate the increase in retail order flow on market Al and increase liquidity provision on this lit venue.⁷ We document a statistically significant increase in displayed dollar depth on market Al after the rule change. We show that this effect is causal: By instrumenting the fraction of the aggressive volume that retail traders trade on Al with the cross-sectional level of Ad

⁶ Like in the US, Canadian brokers do not typically pass on exchange fees to their clients. We therefore expect fees to influence the routing decisions of brokers, consistent with the evidence on the US market presented in Battalio, Corwin and Jennings (2016).

⁷ It is noteworthy that most of the other DPMM did not switch to Al but instead stopped providing liquidity, suggesting that they are at a disadvantage to high frequency market makers in a lit market.

market share pre-event, we show that the change in retail trader flow to AI causes the increase in dollar-depth on market AI. Unlike market Ad where the liquidity is only supplied to retail traders, on AI the liquidity is available to all trades

The causal effect of the shift in retail order flow is one of the main insights of our analysis and has important policy implications. These insights are most relevant in the US, where the features of market Ad before the rule change strongly resemble the retail internalization market. Specifically, market Ad segments aggressive retail order flow into a venue where liquidity is supplied by fast and slow intermediaries seeking to earn the spread and trade against uninformed order flow. If the economics of market making in the U.S are similar to the economics for the intermediaries in market Ad, the results suggest that U.S. wholesale market makers may also reduce liquidity supply in the dark in the presence of a price improvement rule. The Canadian evidence suggests that this would in turn increase liquidity on lit markets. In contrast to wholesaler liquidity, lit market liquidity is available to all traders.

Consistent with prior literature, such as Comerton-Forde and Putnins (2015) and Menkveld, Yueshen, and Zhu (2016), our research confirms that not all types of dark trading are equal. We show that dark pools and retail internalization have different impacts on market quality, and respond differently to regulation. Therefore, regulators and researchers should be careful not to aggregate different types of dark trading. The importance of distinguishing between different types of dark trading is illustrated by contrasting our paper with Foley and Putnins (2016) which examines the same regulatory change considered in this paper. Despite the fundamental differences in the types of dark trading examined, Foley and Putnins aggregate dark trading across all venues. This obfuscates the reasons for the decline in dark trading after the MPIR. The aggregate market-level analysis undertaken by Foley and Putnins is too coarse to identify the impact of this decline leading to spurious conclusions. Their market aggregate approach fails to identify that the drop in dark trading is driven by the withdrawal of

intermediaries, and that the causal effect for lit liquidity is not the dark liquidity but rather the increase in retail participation in the lit market.⁸

Our trader-level data also allows us to go beyond an analysis of aggregate market quality to assess who won and lost as a result of the introduction of the new rule. For retail traders we find that although the per-trade price improvement increases, in aggregate retail traders receive significantly less price improvement after the change. Although statistically significant, the economic magnitude of this decline is relatively small. Retail brokers incur higher costs due to the taker fees paid as a result of the shift from Ad to Al. In contrast, high frequency traders benefit because they receive rebates for liquidity supply on market Al. Finally, institutional and other traders face a higher execution shortfall costs.

II. The institutional setting

A. Rules governing trading in Canada

The Toronto Stock Exchange (TSX) is the primary listing venue for large companies in Canada. 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)⁹ 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,

⁸ We provide a detailed discussion of the problems associated with the approach taken by Foley and Putnins in section A1 of our Internet Appendix which is available at: <https://dl.dropboxusercontent.com/u/25325521/CanadaDarkInternetAppendix.pdf>.

⁹ IIROC is a self-regulatory organization that oversees dealers and trading activities and performs real-time market surveillance.

which mandates that orders must be 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.¹⁰

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,¹¹ 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.¹² 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.¹³

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

¹⁰ 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.

¹¹ See UMIR 6.3 and related guidance notes.

¹² 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.

¹³ 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.

Securities Commissions (IOSCO), dark orders have lower execution priority than visible orders at the same price.¹⁴ 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.¹⁵ 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 in detail 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 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 changes

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

¹⁴ See IOSCO "Principles on Dark Liquidity" <http://www.iosco.org/news/pdf/IOSCONEWS210.pdf>

¹⁵ 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.

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 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.¹⁶ The rule change is referred to as the minimum price improvement rule (MPIR).

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

We examine the impact of the MPIR. 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 the change in regulation

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.¹⁷

¹⁶ An iceberg or reserve order is an order that displays only a portion of its full size.

¹⁷ Although readers familiar with the Canadian market will be able to identify these marketplaces based on these market shares, our confidentiality agreement prohibits us from naming the marketplaces in the paper.

Only marketplaces A and D are impacted by the introduction of the MPIR.¹⁸ Below we provide a detailed explanation of the dark trading features of marketplaces A and D, including details of how these marketplaces were impacted by the introduction of the dark liquidity rules. The institutional arrangements in place in the other marketplaces are described in section A2 of the Internet Appendix.

Marketplace A operates a public limit order book, which we refer to as market A1, and a dark pool facility, which we refer to as market Ad. A1 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.

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

¹⁸ In the Internet Appendix, we provide details of the fully hidden orders allowed on marketplaces B and C. These order types already complied with MPIR before it came into effect. In the data, we observe no change in the use of these dark orders after MPIR.

match with a dark order in market 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 market Al, provided that Al was quoting the best price on the relevant side of the market, and were only routed to other markets if an execution was not found on market Al. Although we cannot separately assess the impact of the change in functionality of this order type, we note that market Al had the lowest fees for taking liquidity among the major lit markets (see Table A.I in the Internet Appendix for details of exchange make-take fees), and it was therefore arguably most attractive for liquidity taking orders before and after the rule change. We therefore expect retail brokers to prefer to route orders to market Al regardless of Al's change in routing practices.

Marketplace D is a dark pool that allows traders to interact using two types of orders. These order types are similar to those in market 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.

III. Data and sample

A. Data

The data for this study is provided by the Investment Industry Regulatory Organization of Canada (IIROC). 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 SDL order, and a unique identifier for each order.

For trades, the data additionally specifies the aggressive (liquidity-demanding) 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. The masking is applied consistently so that the same marketplace, broker and user are 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 (2015), Korajczyk and Murphy (2015) and IIROC (2014) contain further information about the data.

B. Sample

Our analysis covers the period August 27 to November 30, 2012, (i.e. seven weeks before and after the event date, October 15, 2012).¹⁹ We end the sample on November 30 to avoid

¹⁹ We eliminate 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.

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-cross-listed securities.

Our sample comprises “highly-liquid” securities, as defined by IIROC, which 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.²⁰ IIROC compiles a list of highly-liquid 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 non-ETP 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.²¹

IV. 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 researchers describe the usage of user IDs in detail in recent

²⁰ For further details see IIROCs definition on <http://www.iirocca/industry/rulebook/Pages/Highly-liquidstocks.aspx>.

²¹ We observe 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.

research reports (IIROC 2012 and IIROC 2014).²²

According to these IIROC 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 flow (e.g., retail vs. institutional) by user ID. IIROC requires this separation for DMA clients. 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 have a total of 3,642 unique trader IDs in our classification sample, although many of these are inactive.

We group traders into four categories: high frequency (HFT), retail, institutional, and other. The other category includes trader IDs that we are not able to classify as HFT, retail, or institutional. These classifications are based on observed trading characteristics. We briefly describe these classifications below, and provide a more comprehensive description in section A3 of the Internet Appendix.

HFT are identified using two measures of reaction speed: the median order-to-cancel time and trader IDs that consistently submit large numbers of orders within 500 milliseconds of daily scheduled announcements of the market-on-close order imbalance. We identify 89 HFT IDs

²² 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

accounting for 36.1% of dollar trading volume and 43.8% of trades.²³

Retail traders are identified using an order type on market Ad that is exclusively available to retail traders. The use of this order type is the choice of the broker, not the customer, and it is our understanding that brokers must explicitly seek to be connected to venue Ad to use this order type. There are 135 retail trader IDs. The relatively small number of retail IDs reflects the fact that retail brokers use a single ID for multiple retail clients. Retail traders account for 9.4% of dollar volume traded and 7.8% of trades.

Institutional traders are identified by trader IDs that use large pre-arranged trades off-exchange, and trader IDs that accumulate large inventory positions across all Canadian marketplaces. Additionally, we require that these IDs are designated as “client” accounts. We identify 558 institutional investor IDs which account for 23.6% of dollar volume traded and 21.2% of trades.

The other category comprises trader IDs that are not identified as retail, institutional or HFT. There are 2,860 of these IDs which account for 30.8% of dollar volume traded and 27.2% of trades.

V. The impact of the minimum price improvement rule on dark trading

A. What is the impact of the MPIR on dark liquidity?

We measure the impact of the introduction of the MPIR on dark liquidity 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 dark trading volume. Second, we compute the share volume of dark orders, as a fraction of the share volume of all

²³ This measure differs from much of the existing literature which uses day end inventories due to the perception that HFT firms aim to end the day flat. However, in Canadian markets, a single DMA client may use multiple trader IDs (IIROC (2012) and IIROC (2014)), and it is possible that an HFT firm is assigned multiple user IDs. As a consequence, low end-of-day inventories are neither a necessary nor a sufficient attribute of an HFT trader ID in our dataset. Notably, many of the IDs that we classify as HFT hold substantial median end-of-day inventory.

orders. We refer to this as dark order volume. We examine these measures for the market in aggregate and each marketplace separately.

A.1. The impact of the MPIR on aggregate dark liquidity

Figure 1, Panel A plots dark trading volume and dark order volume across all venues. The figure shows that there is a significant drop in both dark trading and order volume following the introduction of the MPIR. Dark trading volume falls from 9.3% to 5.4%, and dark order volume declines from 17.2% to 11.9%.

We formally analyze the impact of the MPIR on dark trading and order volume using regression analysis. We use two control variables. First, we use the U.S. volatility index VIX to control for market-wide volatility. Second, we use the cumulative return since August 24, 2012, on an ETF for the S&P GSCI commodity index (labeled CRCI). We include this control because in the second part of our sample we observe a general decline in the market-wide price level. This negative market-wide return has the capacity to change several measures; an example are spreads measured in basis points of the mid-quote – a decline in the denominator can indicate an increase in spreads even if spreads measured in dollars did not change. The GSCI is a major international index and it is highly correlated with the Canadian TSX Composite index. It therefore captures and controls for market-wide changes in stock prices without being subject to the endogeneity concerns that would arise if we use, for instance, the TSX Composite. We estimate the following equation:

$$DV_{it} = \alpha \times MPIR_t + \sum_{i=1,2} \beta_i \times controls_t + \delta_i + \epsilon_{it}, \quad (1)$$

where DV_{it} is the dependent variable that measures dark trading volume or dark order volume for stock i for day t ; $MPIR_t$ is a dummy variable for the change in regulation and is 0 before October 15, 2012, and 1 thereafter; $controls_t$ are the daily realization of the U.S. market volatility index VIX and the cumulative return on the commodity ETF GSG from August 24 to t ; and δ_i is a security fixed effect. 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)).

Table I, Panel A confirms the observations in Figure 1. Aggregate dark trading and dark order volume decline significantly after the change in the rules.

A.2. The impact of the MPIR on dark liquidity, by marketplace

The 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 trading venues. To understand the relation between the organization of trading and the impact of the MPIR, we analyze the change in dark trading by marketplace. 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).

Figure 1, Panel B shows a sharp contrast between the impact of the MPIR on Ad and D. Before the MPIR, dark trading in Ad accounts for almost 4.6% of the total dollar trading volume in Canada, whereas after the MPIR, it accounts only for 0.8%. Dark trading in dark pool D, the other main dark market, accounts for 2.5% both before and after the MPIR.

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

$$DV_{it} = \sum_{m \in \{A,B,C,D,O\}} \alpha_m \times mk_m \times MPIR_t + \sum_{i=1,2} \beta_i \times controls_t + \delta_i + \epsilon_{it}, \quad (2)$$

where DV_{it} is the dependent variable that measures dark trading volume or dark order volume; mk_m is a dummy that is 1 if the dependent variable observation is for marketplace m , where $m=O$ stands for all marketplaces other than A, B, C, D; $MPIR_t$ is a dummy variable for the change in regulation and is 0 before October 15, 2012, and 1 thereafter; $controls_t$ are the daily realization of the U.S. market volatility index VIX and the cumulative return on the commodity ETF GSG from August 24 to t ; and δ_i are fixed effects for markets and securities.

Equation (2) allows us to simultaneously estimate the effect of the MPIR for all affected

marketplaces and to test whether the sharp decline in dark volume in Ad reported in Figure 1, Panel B is indeed larger than on other markets. Table I, Panel B confirms that Ad experiences a significant drop in dark trading volume of 3.81%, while D exhibits no change. Both Ad and D see a significant drop in dark order volume 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 dark pool Ad is larger. All other marketplaces are unaffected by the change in dark liquidity rules.

B. What explains the differential effect of the minimum price improvement rule?

The MPIR fundamentally altered the economics of providing liquidity in the dark. As a result of the mandated midpoint pricing, liquidity providers are not able to earn the bid-ask spread by posting orders on both sides of the market in dark pools (prior to the MPIR, liquidity providers were able to earn up to 80% of the NBBO spread in Ad and up to 60% in D, on round trip transactions). We therefore hypothesize that the nature of liquidity provision, prior to the rule change, will influence the impact of the rule change. Where liquidity is supplied by natural liquidity providers, we expect trading to be unaffected. However, where liquidity supply is from traders posting liquidity on both sides of the market, seeking to earn the spread, we expect liquidity supply to decline.

B.1. Who supplies liquidity in the dark?

We therefore begin by examining differences in the types of traders that posted liquidity in dark pool Ad and D before the introduction of the MPIR. Using our trader-level data we find that in dark pool Ad liquidity is supplied by HFT (26.8%), others (69.6%) and institutions (3.5%). In dark pool D, liquidity is supplied by HFT (10.6%), institutions (35.8%) and others (28.6%) (these numbers exclude the periodic matches). The relatively high level of HFT supply in dark pool Ad and institutional supply in dark pool D provides a preliminary indication that there is more intermediation in dark pool Ad than in D. However, the very high level of liquidity

supply from the other traders in dark pool Ad suggests that further analysis of this group of traders and the nature of liquidity provided by each trader type is required.

B.2 Intermediation in the dark

To assess whether a trader intermediates by posting liquidity on both sides of the market we compute an intermediation score for each trader that posts limit orders in dark pools Ad or D. We compute this intermediation score separately for each dark pool (i.e., a liquidity provider that posts liquidity in both dark pools will be assigned two intermediation scores). For each dark pool the score is based on the volume of passive dark orders that a trader submits to that dark pool. We exclude orders that are designated to trade only if they are marketable (and to be cancelled or re-routed as a limit order to a lit trading venue otherwise).

$$\text{Intermediation score} = \frac{|\text{buy dark order volume} - \text{sell dark order volume}|}{\text{total dark order volume}} \quad (3)$$

We compute the intermediation score per day, per stock, per trader ID, for all dates during our sample before the introduction of the MPIR. We then determine the median score per trader ID. Most traders have an intermediation score that is equal to or close to 1. In the data, there is a visual break at 0.4 (the next highest scores are above 0.7). We therefore classify a trader as an intermediary for a given dark pool if the trader's intermediation score for that dark pool is below 0.4. We refer to these traders as dark pool market makers (DPMM).²⁴

B.3 Difference in intermediation in dark pool Ad and D

Table II reports who trades with whom in dark pool Ad and D. We consider trades where HFT or other DPMM provide liquidity, and report the fraction of trading volume that they supply to other traders that are not DPMM.

²⁴ Section A3.E of the Internet Appendix provides a summary of the extent to which liquidity in dark pool Ad and D is provided by DPMMs.

Table II, Panel A shows that HFT (other) DPMM providing liquidity to retail traders account for 25.8% (61.6%) of trading activity in dark pool Ad before the rule change. Therefore, most trades are intermediated and the dominant liquidity providers in this venue are ‘slow’ rather than ‘fast’ intermediaries. There are three possible explanations for this. First, unlike in lit venues where time priority is critical to the success of a market making strategy, in dark pool Ad time priority plays little role. Instead, liquidity providers take turns to trade against incoming order flow according to a set off priority rules based on first price (for priced-orders), then broker, size, and finally a “round-robin” mechanism. Second, limit order book monitoring is less important because dark orders are priced relative to the NBBO and therefore traders do not need to re-price their orders with the same frequency as they do in lit venues. Third, liquidity providers can choose to have certainty that the counter-party to trades in dark pool Ad will be retail (by using SDL orders), and hence the probability of trading with an informed trader is relatively low.

Table II, Panel A also shows that the liquidity supply by the DPMM declines substantially after the rule change to 0.7% for HFT DPMM and 36.6% for other DPMM. We note that the percentage declines mask the extent of total liquidity reduction, as some DPMM exited the dark markets entirely and total volume declined substantially (see Table A.II in the Internet Appendix: the dollar amount of liquidity provided by other traders including DPMM declines by 85%).

We observe in Table II, Panel B that the nature of liquidity provision in dark pool D is fundamentally different. The DPMM offer very little liquidity before the rule change, and almost none after the rule change. Liquidity in this venue is supplied primarily by natural liquidity providers who are relatively unaffected by the MPIR.

To formally establish the importance of intermediation, we first ask whether a trader’s intermediation score predicts the trader’s change in liquidity provision after the MPIR. We perform this analysis using all 895 traders that supply liquidity at any point in dark pools Ad or

D by computing the difference in liquidity provided before and after the introduction of the MPIR, for all liquidity providing IDs. We then estimate the following regression

$$\Delta\%liq_i = \alpha + \beta \times imb_i + \gamma_1 \times HFT_i + \gamma_2 \times retail_i + \gamma_3 \times institutional_i + \varepsilon_i, \quad (4)$$

where $\Delta\%liq_i$ is the difference in the percentage of aggregate liquidity provided by trader i before and after the MPIR, imb_i is the imbalance score for trader i , and HFT_i , $retail_i$, and $institutional_i$ are dummies for the respective trader groups.

Table III displays the results for our estimation of equation (4). We find that the trader ID's intermediation score has explanatory power with regard to the change in liquidity supply by this trader, and that after controlling for the intermediation score, the trader type has no additional explanatory power. This confirms the descriptive results presented in Table II that liquidity in dark pools Ad and D was provided by different traders: dark pool Ad relied primarily on intermediated liquidity provision, whereas dark pool D served primarily as a matching venue for natural traders.

Second, we ask whether the extent to which a security is traded by DPMM predicts the change in the dollar volume executed in dark pools Ad and D, in the cross-section of securities. Specifically, we compute the dollar volume traded in dark pools Ad and D respectively as a fraction of the stock's total daily passive volume in lit and dark markets combined. To assess the extent of intermediation in the dark, per day per stock, we compute the passive dollar volume traded in dark pools Ad and D by traders classified as DPMM, as a fraction of the total stock's daily passive dollar volume for that market. We then compute the average of these fractions before the MPIR per security and market and estimate the following regression for the cross-section of the 92 securities for dark pools Ad and D:

$$\Delta marketshare_i = \alpha + \beta_1 \times DPMMVolume \times Ad + \beta_2 \times DPMMVolume \times D + \varepsilon_i, \quad (5)$$

where $\Delta marketshare_i$ is the marketshare of dollar volume for the security and market that relates to each observation; $DPMMVolume$ is the percent of dollar volume that is provided by DPMM for security i ; and Ad and D are dummies that are 1 when the observation is for markets

Ad and D respectively and 0 otherwise. We omit the constant α from the table, and we also run a combined specification where we do not interact *DPMMVolume* with the respective market-dummies.

The estimation results are in Table IV. We find, consistent with our observations across trader IDs, that the degree of liquidity provision by DPMM is a significant predictor of the changes in the market shares of trading.

Our findings illustrate that the impact of the MPIR on dark trading depends critically on the level of liquidity supplied by intermediaries, rather than matches between natural traders. Activity declines in the venue with liquidity supplied by intermediaries, but not in the venue with natural liquidity provision. Although our analysis cannot explain why the two main Canadian dark pools had different liquidity providers, there are several institutional reasons for why dark pool Ad may have been more attractive to intermediaries than dark pool D.

First, dark pool Ad segments liquidity by offering liquidity providers an option to trade only against retail orders. Since retail order flow is arguably less informed and more balanced (in terms of buys and sells), it is likely that a liquidity provider would be able to earn revenue on a round-trip transaction in Ad. Additionally, liquidity providers were able to earn up to 80% of the bid-ask spread on market Ad but only up to 60% on market D. Second, at the time of our study, in contrast to all other venues, according to industry insiders, the servers for D were not located Toronto but in New Jersey, making D less attractive for high frequency traders. Third, the matching rules in dark pool Ad prioritized liquidity providers from the same broker-dealer as the incoming retail marketable order, making Ad attractive to broker-dealers who sought to internalize their retail order flow in a market-making fashion.²⁵

It is less clear why dark pool Ad saw little institutional participation, in particular since the average trade sizes in the dark pools were comparable and because these traders should also

²⁵ With few exceptions, in Canada, orders from retail clients must be sent to public markets, and U.S.-style off-exchange internalization is generally not permitted.

have been aware that it provided an opportunity to trade with relatively uninformed retail investors. Several market participants from brokerages indicated to us that it would be difficult to monitor dark limit orders that have been posted across several dark markets and that they would post orders to only one venue, making dark pool D their choice of destination for institutional limit orders for legacy reasons (market D has existed since 2007, whereas dark pool Ad opened in 2011). As we understand it, many brokerages that handle institutional order flow in Canada do not make routing decisions on a case-by-case basis; instead, the basic parameters and frameworks are determined by their “best execution” committees, giving rise to the so-called “routing tables” that brokers follow when handling clients’ orders. The behavior of traders who use limit orders to build or unload a client’s position thus tends to be sticky.

VI. Volume decline in dark pool Ad and the impact of retail segmentation

The institutional features of dark pool Ad together with our findings in Section V illustrate that trading in Ad pre-MPIR resembles the practice of internalization in the U.S., where most retail marketable orders are executed off-exchange, often through a market maker (“wholesaler”), e.g. Citadel, who purchases these orders.

Trading in Ad had three features that are similar to trading under the wholesaler model: (i) (almost) all marketable orders are from retail traders; (ii) before the MPIR, marketable orders in Ad received nominal improvement over the NBBO; and (iii) most liquidity on market Ad was provided by DPMM. Furthermore, although brokers are not permitted to receive payments for order flow in Canada, trading fees for marketable orders (“taker fees”) on market Ad were an order of magnitude lower than on lit markets (4 cents per 100 shares on Ad; the second lowest fee was 28 cents per 100 shares on market Al). Therefore, understanding the impact of the decline in trading in dark pool Ad will thus yield insights into the impact of internalization of retail order flow in the US and other markets.

A. What happened to retail orders after the introduction of the MPIR?

Most retail investors were likely unaware of the introduction of the MPIR, and therefore we do not expect substantial changes to the retail order flow that was received by the brokers. With the withdrawal of liquidity from dark pool Ad after the MPIR, retail orders that would have been executed in Ad prior to the rule change had to be routed elsewhere after the change. We first ask which markets receive this order flow.

Table V provides regression estimates for the execution of retail marketable and non-marketable orders by venue before and after the MPIR. Prior to the MPIR, 27.6% (11.4%) of marketable retail orders were executed against dark (lit) orders on marketplace A. Following the rule change, executions against dark orders fell by 18.5% and against lit orders rose by 13.6%, after controlling for other factors. Executions against lit orders on marketplaces B and C also increased by 3.4% and 1%, respectively. In contrast, there is no significant change in the execution of non-marketable orders on any marketplace before and after the MPIR.²⁶

Our findings suggest that the retail orders that would have been executed in dark pool Ad are primarily executed as lit marketable orders on market Al (the lit limit order book of the same marketplace) after the introduction of the MPIR. This outcome is not surprising. Given that routing decisions typically lie with the broker, the marketplaces' taker fees for liquidity-demanding orders should arguably play a role. During our sample period, market Al charged the lowest taker fees among the three main lit markets (Al, B, and C). We therefore expect that retail marketable orders that do not find a match in dark pool Ad will be routed to market Al (conditional on abiding by the order protection rule, the workings of which would explain why some of marketable orders also go to marketplaces B and C).

²⁶ Table A.III in the Internet Appendix provides comprehensive summary statistics on the execution of dark and lit orders on all venues by trader type.

B. Does the switch of retail trading from dark to lit market affect lit market quality?

Providing liquidity to retail order flow is arguably attractive 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). We expect that an increase in the fraction of retail orders that reach the lit markets will improve market quality. However, in untabulated regressions, we did not observe a change in the market-wide bid-ask spread,²⁷ possibly because the change in the market-wide composition of the order flow is not large enough to affect the aggregate measures.

Therefore, we next study whether the change in the distribution of retail flow between lit and dark venues led to changes in market quality on individual venues. Figure 2, Panel A shows that market A1's share of retail volume increases from around 15% of A1's dollar volume to 30%. The changes in the retail shares for markets B and C (unreported) are substantially smaller. We therefore expect that changes in market quality, if any, should be most pronounced on market A1. Figure 2, Panel A also illustrates the strong co-movement between the quoted depth on market A1 and the retail share of market A1's aggressive trading volume, in aggregate.

Figure 2, Panel B plots the natural logarithm of the dollar-depth for the three markets A1, B and C, confirming an increase in depth on market A1 and little or no change for markets B and C.

We confirm these visual observations in a formal regression analysis. Table VI presents the results of our estimation of equation (2), where we use measures of market quality as the dependent variable DV_{it} . Market quality is measured using time-weighted quoted spreads measured in cents and basis points and depths measured in the natural logarithm of shares and dollars. Table VI illustrates that there is a significant increase in depth on market A1 following the introduction of the MPIR, by about 17%. It further shows that there is no evidence for a

²⁷ Analysis of market-wide bid ask spreads is presented in two research reports written by Comerton-Forde et al (2015) and Devani et al (2015) which are available on the IIROC website at: http://www.iroc.ca/Documents/2015/0BBCAC3E-670E-4B2A-A297-E7CC5F1C0B11_en.pdf and http://www.iroc.ca/Documents/2015/d215afed-a01e-453d-8f24-bd8ed2b948bf_en.pdf, respectively.

change in time-weighted quoted spread except for a minor decrease on market B (significant only at the 10% level) measured in cents.

There are two possible explanations for the absence of significant changes to the bid-ask spread. First, the order protection rule requires that marketable orders are routed to the venue that is posting the best price, and we therefore do not expect spreads on individual venues to substantially differ from each other. Second, 50 of the 92 stocks in our sample are constrained by the minimum tick size more than 80% of the time. This constraint means it is extremely difficult for spreads to tighten further after the MPIR. However, in a separate analysis we split the sample into 80%-tick-constrained and other securities, and we find that (a) quoted spreads in cents actually decrease for the tick-constrained securities (but not for the others) and (b) the fraction of the day that these securities are tick-constrained increases. The decline in spreads is, however, economically small (less than 0.05 cents for market A1). We also find that market A1 is at the NBBO more frequently. The corresponding regression tables are in Table A.IV of the Internet Appendix. This provides weak evidence of an improvement in spreads for tick-constrained securities. We find no change in spreads measured in basis points, and we emphasize that it is important to control for secular changes in the price level of securities: in November 2012, stock prices in Canada declined across the board, which in and by itself would cause bid-ask spreads measured in basis points to increase.

Trades of retail investors are commonly thought to arrive at the market in a more random fashion than those of institutional traders who typically split their large parent orders into smaller child orders. Given that more retail orders arrive at the lit markets after MPIR, it is, therefore, likely that prices become more random. A common measure of randomness is the absolute value of the return autocorrelation, where returns are typically based on 1-minute returns for mid-quotes of the bid-ask spread. We find support for the hypothesis of more randomness in prices: there is a small, but statistically significant decline in the return auto-correlation, both for market A1 (where most of the retail flow trades) and the NBBO. These results are reported in Table A.V of the Internet Appendix.

Overall, our results provide evidence of an improvement in liquidity on market A1, and we attribute this improvement to the increased retail share of market A1's aggressive volume.

C. The impact of retail trading on depth

The critical question is whether the association between a market's level of retail trading and the market's liquidity is causal. To establish causality, we use the dollar trading volume in dark pool Ad *before* the introduction of the MPIR as an instrument for the fraction of their dollar volume that retail traders execute on market A1 *after* the MPIR.

Canada has full post-trade transparency, and therefore market participants knew – after the fact – how much trading occurred in dark pool Ad. Given its institutional arrangements, prior to the MPIR, market participants knew that almost all marketable orders in dark pool Ad were from retail traders. Therefore, the trading activity in dark pool Ad prior to this time provided the market with an accurate estimate of the amount of retail flow that did *not* reach the lit markets.

Following the introduction of the MPIR, volume in Ad declines almost to zero. This decline was publicly known. 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. As we argue earlier in this section, the migration of retail marketable orders from dark pool Ad to the lit limit order book A1 was predictable because market A1 had the lowest taker fee among the three main lit markets (A1, B, and C). Conditional on market A1 quoting the best prices, it was thus arguably the preferred destination for retail brokerages that typically absorb the exchange fees and charge their clients flat commissions, which do not depend on the execution venue. Furthermore, as we discuss in Section II.C, marketplace A also amended the way in which retail marketable orders sent to Ad were routed after the MPIR, making it more likely they retail marketable orders that do not find a match in Ad would be sent to A1.

The sharp decline in dark pool Ad's trading volume provided traders with a unique opportunity to estimate the "extra" volume of retail marketable orders that hit market A1 after the introduction of the MPIR. Trading in dark pool Ad pre-MPIR does not affect the depth in A1

post-MPIR directly, and depth in AI post-MPIR does not affect trading in dark pool Ad pre-MPIR. The level of trading in dark pool Ad pre-MPIR is therefore a valid instrument for the level of retail trading in AI post-MPIR.

To determine the causal impact of retail trading on market depth, we perform an instrumental variable regression, using the average market share pre-MPIR interacted with an MPIR-dummy that is 1 after the rule change and 0 before as an instrumental variable for retail activity. The first stage of the implied 2-stage estimation procedure estimates the following relation

$$\%retail\ on\ Al_{it} = \alpha + \beta_1 \overline{marketshare\ Ad}_i \times MPIR_t + \sum_{i=1,2} \beta_i \times controls_t + \delta_i + \varepsilon_t, \quad (6)$$

where $\%retail\ on\ Al_{it}$ is the fraction of their dollar volume that retail trade on market AI for security i on day t ; $\overline{marketshare\ Ad}_i$ is the average pre-MPIR daily market share of dollar volume for security i ; and $MPIR_t$, $controls_t$, and δ_i are as defined before. The first stage takes the instrument as an exogenous shock to retail activity on AI after the MPIR and the regression then estimates the effect of this shock. In the second stage of the IV regression we estimate:

$$\ln(ddepth_{it}) = \alpha + \beta_1 \%retail\ on\ Al_{it} + \sum_{i=1,2} \beta_i \times controls_t + \delta_i + \varepsilon_t, \quad (7)$$

where $\ln(ddepth_{it})$ is the natural logarithm of the time weighted quoted dollar-depth on market AI and where we instrument our main variable of interest $\%retail\ on\ Al_{it}$, by $\overline{marketshare\ Ad}_i \times MPIR_t$. We estimate the first and second stage of the IV regression jointly. The estimate β_1 reflects the change in depth due to the shock in the instrumented variable $\%retail\ on\ Al_{it}$.

Table VII shows our results for the first and second stage of the IV regression. We ran a number of different specifications, with and without fixed effects, and the results are consistent across specifications. The first stage shows that our instrument indeed has a significant impact on the usage of AI by retail traders. All specifications tests that we examine show that the

instrument is statistically valid; in the table we include the F-test and the Kleibergen and Paap (2006) Wald statistic of under-identification. The second stage regression then reveals that retail activity on AI has a positive and significant impact on the level of the quoted depth on AI. Therefore we conclude that the increase in retail activity on AI increased the quoted depth on AI.

VII. The impact of the minimum price improvement rule: winners and losers

We discuss the impact of the MPIR on four groups of market participants, retail and institutional traders and HFT and other DPMM. Given that each of these groups has different motivations and strategies for trading, we employ different measures to assess each group's costs and benefits. By evaluating these costs and benefits we can identify the winners and losers of the MPIR. We also discuss the impact on trading venues.

A. Retail traders and their brokers

The impact of the MPIR on retail traders is evaluated using three measures: effective bid-ask spreads, price improvement and maker-taker fees.

A.1 Effective bid-ask spreads

We conjecture that most retail orders are single orders that are not split across time.²⁸ For such orders, the bid-ask spread is the most relevant measure to assess trading costs for marketable orders. We compute effective spreads, as is common in the literature, across all markets. For a buy at time t , the effective spread is defined as twice the difference between the price paid by the marketable buy order and the midpoint of the prevailing bid-ask spread; symmetrically for the sales.

²⁸ With few exceptions, Canadian regulations require that brokers route orders from retail clients to public marketplaces without delay.

Our primary measure of effective spread is measured in basis point of the midpoint excluding the exchange (taker) fees. Since most retail traders do not pay exchange fees directly (but instead pay a flat commission to their brokers), only changes to the bid-ask spread excluding taker fees have an immediate impact on their costs. Changes to the spread that includes the taker fees may, however, affect retail traders in the longer run, if brokers adjust their commissions to reflect the changes in the exchange fees incurred.²⁹ Mean effective spreads without fees remain unchanged at 1.5 cents (approximately 8 bps) following the MPIR.

We formally examine the impact of the MPIR on effective spreads using the regression specified in equation (1) where DV_{it} is the dependent variable that measures the effective spread for retail traders' marketable orders for stock i for day t . The results for the estimation of equation (1) with effective spreads as the dependent variable are in Table VIII. The estimates confirm the absence of statistically significant changes in effective spreads paid by retail traders following the rule change; hence there is no evidence that trading costs for retail traders are impacted by the MPIR.

A.2 Price improvement in dark pools

One of the arguments in favor of U.S. wholesalers and internalization, and, by association, in favor of dark pool Ad, is that retail investors receive price improvement. In contrast to maker-taker fees and, critically, fee rebates which are often absorbed by the brokerage and only passed on to retail clients via flat commissions, the price improvement in dark pool Ad benefits the retail client directly. We compute the price improvement that retail traders receive in dark pool Ad, as the total dollar amount (summed up over all trades in Ad) as well as per trade in Ad.

²⁹ Regression results for the alternative spread specifications are provided in Table A.IV of the Internet Appendix. These results are qualitatively consistent except when fees are included and measured in bps.

After the MPIR, each trade in the dark receives price improvement of at least $\frac{1}{2}$ cent (for 1 cent spreads) and possibly more for larger spreads. Therefore, per share traded, we expect price improvement to increase. However, since trading volume in dark pool Ad declines, it is unclear whether total price improvement increases or decreases.

Table VIII reports a mean dollar price improvement prior to the MPIR of \$102 per stock per day and \$0.63 per trade. The results from our estimation of equation (1), using the total per stock per day amount in dollars and the per transaction amount (in dollars) as the dependent variables DV_{it} shows that consistent with our predictions, the price improvement per trade increases, by about \$1.82. However, the increase is not sufficient to offset the decline in dark pool Ad's trading volume, and the total amount of price improvement declines significantly, by about \$33 per day per stock. Therefore, while retail traders that receive price improvement are better off, retail traders in aggregate are worse off. Considering that the relative loss is distributed over many retail traders, the economic impact of the MPIR per retail trader is small.

A.3 Exchange maker-taker fees

We estimate the change in the amount of exchange fees paid by retail trader IDs. The level of taker fees charged by the trading venues to execute liquidity demanding (marketable) orders is a contentious issue commonly raised by the Canadian retail brokerages. The off-exchange internalization of retail order flow and payment-for-retail-order-flow, which is common practice in the U.S., is illegal in Canada. Retail brokers necessarily incur taker fees when executing clients' marketable orders and the brokerages argue that it is extremely difficult to pass these costs through to the end-client. One of the attractions of market Ad is the comparatively low taker fee. The sharp decline in liquidity in this dark pool compels retail brokers to seek execution for their clients' orders at more expensive venues, and we expect that they face higher exchange fees as a consequence.

To estimate the change in maker-taker fees for retail brokerages, we compute the fees for all markets. The fees are expressed as the net fee per-dollar traded (in bps), and they are

computed as the difference of maker rebates received by brokerages for executed passive limit orders and the total amount of taker fees; a negative amount therefore corresponds to a cost.

Table VIII contains the results from our estimation of equation (1), using the above maker-taker fee measure as the dependent variable. Prior to the MPIR, retail traders paid on average 0.3 bps in maker-taker fees per day per stock. Consistent with our expectations we find that exchange maker-taker fees incurred by retail brokers increase by 0.4 bps.³⁰

In summary, we observe that there is evidence that retail traders and their brokers are worse off after the introduction of the MPIR.

B. High frequency and other dark pool market makers

To assess the costs and benefits of the MPIR for high frequency and other DPMM we focus on two measures: intraday returns to trading and the maker-taker fees paid.

B.1 Returns to trading

We compute returns to trading on day t , following Comerton-Forde, Hendershott, Jones, Moulton and Seasholes (2010), as follows:

$$ret\ to\ trading_t = \min\{buy\ volume_t, sell\ volume_t\} \times (vwap\ sell_t - vwap\ buy_t), \quad (8)$$

where $vwap\ buy_t$ and $vwap\ sell_t$ are volume-weighted average prices for buying and selling, respectively, for this trader. The expression captures the trading revenue on all round trip transactions. Table IX, Panel A and C present summary statistics for returns to trading by HFT and other DPMM, respectively. We compute the returns to trading for all venues as well as for markets Ad and Al jointly and for market Al separately. Table IX, Panel B and D provide results from our estimation of equation (1), using the returns to trading as the dependent variable. We

³⁰ Table A.VIII in the Internet Appendix provides robustness using maker taker fees for retail brokers, and for fees relating to market Al and Ad together and market Al alone. These produce qualitatively similar results.

find no evidence of changes in these returns for either type of DPMM on any venue. This suggests that the returns on the trading strategies of these traders are not impacted by the MPIR.³¹

B.2 Exchange maker-taker fees

Analogous to retail traders, we compute the maker-taker fees that HFT and other DPMM pay (or receive). In the data, we notice significant changes in behavior for non-HFT DPMM. Although they are still involved in the liquidity provision, in terms of dollar-volume, other DPMM reduce their liquidity provision by about 90% as a group, and several IDs stop trading in Ad entirely. At the same time, as a group, on average, they double the dollar-volume that they trade on AI. HFT DPMM leave market Ad almost entirely. These behavioral changes are important to understand our subsequent results. Specifically, market Ad paid no maker rebates, whereas AI paid a positive rebate. Thus trading against the same orders on AI as opposed to Ad increases a trader's fee revenue.

Table IX, Panels A and C show that in the period before the rule change HFT DPMM earned net rebates of 1.1 bps across all venues compared to just above 0 bps for other DPMM. HFT DPMM are therefore better at capturing rebates across all venues than other DPMM.

Table IX, Panel B and D shows that both HFT and other DPMM earn more rebates after the MPIR. For HFTs on Ad and AI combined, the increase is 0.24 bps, for other DPMM it is 0.56 bps. While the increase is larger for other DPMM, we note that HFT DPMM earn six times the total fees and they likely pursue a different set of strategies that are not entirely captured by the change in liquidity provision from dark to lit venues (see Tables A.VII and A.VIII in the Internet Appendix).

³¹ An alternative explanation is that when HFTs and DPMMs trade more in the lit market, their returns to trading become more volatile. In untabulated regressions, we estimate the volatility in returns and find no change for HFTs. However, measuring this effect is quite difficult to measure for DPMMs because some DPMMs leave the market after the rule change and are therefore excluded from the analysis.

In summary, we observe no evidence that returns to trading for HFT and other DPMM change after the MPIR. However, these traders, particularly the HFT DPMM earn significantly more from maker-taker fees post-MPIR.

C. Institutional traders

We assess the costs and benefits for institutions using two measures: the probability that their limit orders execute and the implementation shortfall for all executed orders.

C.1 Probability of execution for lit limit orders

We proxy the execution probability of lit orders using the ratio of passive trading volume (in shares) submitted by institutional traders in lit markets to the total order volume submitted to lit markets (in shares).³²

$$pr(execution) = \text{passive buy-side volume} / \text{buy-side order volume}. \quad (9)$$

We similarly compute the fill rates for dark passive orders. To establish the impact on fill rates following the introduction of the MPIR, we perform a panel regression analysis using equation (2), where DV_{it} is the dependent variable that measures the fill rate for lit or dark orders for institutional traders, split by marketplace.

Panel A in Table X reports the market-wide average fill rates. We find no evidence for a change for lit orders and we find evidence that fill rates for dark orders improve. Panel B in Table X reports the average fill rates before the MPIR for lit and dark orders in each market. Fill rates in the dark are substantially higher in dark pool Ad (13.3% of all order volume gets filled) compared to dark pool D (3.4%); we note, however, that institutions send more orders in total to market D, which is reflected by the fact that the market-wide dark fill rate mimics that of market D (see also Table A.III for trading volume by venue). Table X also reports the regression results.

³² We cannot directly infer from the data which orders are marketable at the time of their submission.

We observe that fill rates for lit institutional traders decline in market C (by about 1 percentage point) – contrary to the intuition that a reduction in dark trading should increase the probability of execution for lit orders. One possible explanation for this result is an increase in competition on lit markets for liquidity provision to retail order flow that would have been executed in dark pool Ad prior to the implementation of the MPIR. We note, however, that institutions trade only a small amount on market C: pre-MPIR they trade only 5.2% of their volume on market C (see Table A.III in the Internet Appendix). Analyzing fill rates for dark orders, we find an increase in the probability of execution for institutional passive orders in dark pool D (by 3.5%), and no changes in the other venues.

C.2 Implementation shortfall

Institutional traders typically build or unwind positions gradually, by splitting their large “parent” orders into multiple “child” orders. The gradual execution of the parent order over time may temporarily affect the stock price (e.g., a large buy order, if detected by other traders, may cause the price to rise), and the bid-ask spread measure of trading costs that we employed for retail traders may therefore underestimate institutional trading costs.

Instead, institutional trading costs are typically assessed using implementation shortfall. This measure compares the actual (ex-post) prices paid when establishing (or received when unwinding) a position with the hypothetical prices that would be obtained if the trader has filled the entire parent order either at the price prevailing at the time when the trader sent the first child orders or at the volume-weighted average price for the “duration” of the order. Unfortunately, the IIROC dataset (as is often the case for trader-level data) does not contain information about “parent” orders. We therefore develop a proxy for parent orders, by aggregating a series of orders for a given stock from an individual trader ID into a “package.” A package is defined as a series of trades are in the same direction (only buys or only sells) on consecutive days (where the Monday following a Friday counts as “consecutive”). We exclude trades from trader IDs that we classified as HFT, retail or a DPMM. We examine all trader IDs that we previously classified as institutional or that we were not able to classify. Finally, we exclude pre-arranged block trades.

One limitation of this approach is that a single trader ID may handle orders from multiple traders and may be routing buy orders on behalf of one client and sell orders on behalf of a different client. Our sample of packages does not capture these trades. Excluding packages that traded both before and after the introduction of the MPIR on October 15, 2012, we identify 23,527 packages.

We analyze trading costs for these packages using the approach of Anand, Irvine, Puckett and Venkataraman (2012). We compute, for each package, the length (in days), the number of transactions, the VIX at the beginning of the package, the CRCI value at the beginning of the package, the size of the package relative to total volume in the stock during the duration of the package (also split by buying and selling volume), the 20-day closing-price return standard-deviation at the beginning of the package, the lagged volume imbalance (buy-volume minus sell-volume over total volume) relative to the starting date of the package. We include dummies if the package is a purchase and if the trader is in the Other category (i.e., is not classified as an institutional trader in Section IV). Table XI, Panel A reports that there is little change in the package characteristics before and after the introduction of the MPIR. The average package duration is 2.5 days and includes approximately 100 trades. The average package represents approximately 1% of the average daily volume. Sell (buy) packages account for approximately 1.2% (1.5%) of average sell (buy) volume.

We benchmark the total trading costs per package to (i) the price of the first trade of the package and (ii) the volume-weighted average price (VWAP) for the days when the package traded, where we report results relating to the VWAP in Table A.XI in the Internet Appendix. We compute the “raw” shortfall per trader j as:

$$shortfall_{it}^j = \frac{(buyvol_{it}^j - sellvol_{it}^j) - (sellvol_{it}^j - buyvol_{it}^j) \times p_{it}^{j,benchm}}{vol_{it}^j}, \quad (10)$$

where $(sellvol_{it}^j - buyvol_{it}^j) \times p_{it}^{j,benchm}$ is the trader j 's hypothetical cost of establishing a position at the benchmark price $p_{it}^{j,benchm}$, which is either the price of the first trade in the

package or the VWAP price, and $buy\$vol_{it}^j - sell\vol_{it}^j is the trader's realized costs. A larger implementation shortfall corresponds to the higher trading costs. We scale these costs by the dollar volume of the packages to compute the shortfall as a fraction of the total dollar volume traded. We estimate the following regression equation:

$$DV_{it} = \alpha \times MPIR_t + \sum_{i=1..10} \beta_i \times controls_{it} + \delta_i + \epsilon_{it}, \quad (11)$$

where DV_{it} is the dependent variable that measures the shortfall for stock i on day t where t is the first day of the package; $MPIR_t$ is a dummy variable for the change in regulation and is 0 before October 15, 2012, and 1 thereafter; $controls_{it}$ are the above-mentioned control variables, and δ_i is a security fixed effect.

Table XI, Panel B, reports that the mean implementation shortfall before the MPIR, measured using the first trade price (VWAP) is -3.1 (-0.2) bps. The regression results reported in Table XII show there is a significant increase in the shortfall (at the 5% or 10% level) measured using the first trade price. Table A.XI in the Internet Appendix shows that these results are robust when implementation shortfall is measured using VWAP. These results illustrate that, controlling for size and length of the package, institutions incur higher costs to fill their trades after the MPIR, of around 19 bps.

D. Marketplaces

None of the four main marketplaces changed their maker-taker fees as a result of the introduction of the MPIR. However, marketplace revenues can be affected by the change in two ways: first, through shifts in aggregate volume and, second, through changes in the mix of dark and lit volume. For each share traded, marketplaces earn the net fee, that is, the taker fee minus the maker rebate. The three main marketplaces A, B, and C all charged a higher net fee for dark transactions (4, 10, and 9 cents per 100 shares) compared to lit transactions (3, 1-4, and 4). As a result, a reduction in the share of dark trading leads to lower fee income. Indeed, in untabulated regressions we estimate equation (1) by venue, using the per-share fee as DV_{it} , and observe that

markets A and C have smaller fee revenue per share traded, and we find no change for market B. Estimating equation (1) by marketplace, using the total fee as DV_{it} , we find no significant change for any of the marketplace.

VIII. Conclusions

Regulatory concerns about both retail internalization and dark trading has heightened the need for detailed and rigorous analysis of their impact on market quality. We use the introduction of the MPIR in Canada to examine these issues. Demonstrating the critical need for empirical evidence, the Canadian regulators facilitated our research by providing proprietary trader-level data which allows us to document and explain how these rules impacted not only the market in aggregate, but also how they impacted different marketplaces and different trader types.

The MPIR had an immediate impact on the Canadian market. The rule change made intermediation in the dark unprofitable, and as a result intermediated dark volume all but disappeared. Aggressive retail orders that had previously been intermediated by both HFT and other DPMM are routed to a single lit market after the MPIR. We believe that the choice of this market was predictable because it charged the lowest fees for marketable orders. We find that the influx of retail volume on this market led to a substantial improvement in liquidity, demonstrated by larger displayed depth and more frequent trading at the NBBO on market Ad and tighter spreads in tick-constrained stocks. In contrast to the liquidity supplied in market Ad which is only available to retail traders, the lit liquidity on market Al is available to all traders.

These results provide insights into the impact of wholesalers in the U.S market. It suggests that the practice of internalization of retail order flow in the U.S. may harm market quality. The potentially negative impact of internalization was debated in the U.S. prior to the introduction of Regulation NMS. In a comment letter to the Securities and Exchange Commission in 2004, Citadel, a large U.S. broker argued that “[...] the potential long-term

impact of internalization is so corrosive to our national market system that the Commission should take every possible step to curtail this business practice. [...]” They suggested that “[...] the Commission ultimately should require all market participants to route their order flow to any one of the regulated security exchanges or alternative trading systems.”

In the Canadian context, the improvement in lit liquidity due to the MPIR did not unambiguously benefit all market participants. As traders adjusted their choice of marketplace, they became subject to different exchange fees. High frequency market makers are the beneficiary of the change, capturing larger exchange rebates as their liquidity provision shifted from market Ad to lit markets. In contrast, retail brokers incurred higher exchange fees as they are required to pay take fees on lit venues. A simple back-of-the-envelope calculation illustrates the impact of the MPIR on retail brokers: before the MPIR, the average trade in market Ad was for approximately 350 shares. In Ad, this trade incurred take fees for the retail broker of $350 \times \$0.004 = \0.14 . In Al, where most of the marketable retail orders trade after the MPIR, the taker fee is \$0.98. Given the typical flat commission of \$7.99 per trade charged by retail brokers, the difference of \$0.84 certainly affects the brokers’ margins. In the longer term we expect this will lead to increases in retail commissions or changes in the services delivered to retail clients by brokers, but we are not aware of any such changes at this time.

The changes in the distribution of exchange fees and revenues among the different market participants highlights that dark trading is just one piece of the market structure puzzle and that the debate on dark trading is intricately connected to the debate on exchange fees. The improvements in posted liquidity that arise due to higher maker rebates for liquidity provision come at the expense of higher fees levied on marketable orders. This ultimately changes incentives for the demand and supply of liquidity, and the profitability of different types of traders and trading strategies.


This unintended consequence of the MPIR set in motion new initiatives aimed at accommodating retail brokers’ complaints about their increasing exchange fees. The first notable change is the May 2013 introduction of a new marketplace, CX2, a lit limit order market

with “inverted” maker-taker fees (the operator is Chi-X, now owned by NASDAQ). Anecdotal evidence suggests that this marketplace captured the retail flow that had previously been segmented in dark pool Ad. In 2014, the Ontario Securities Commission announced that it would consider a pilot study on the abolition of maker-taker fees. In August 2015, the TSX reconfigured its Alpha Exchange as an inverted maker-taker fee venue. Additionally, TSX introduced a speed bump for all orders except for “post-only” so that this marketplace would be unattractive to all traders that want to take liquidity at multiple venues. The express goal of the reconfiguration is to capture retail order flow. At least one major retail brokerage confirmed to us that they are now using Alpha Exchange as the preferred marketplace for their retail order flow. Finally, in February 2016 Aequitas Innovations’ Neo Exchange introduced an inverted fee structure for its Neo Book, a market that operates a speed bump for fast traders. The impact of these changes on order routing decisions and order flow segmentation should be the subject of further research.

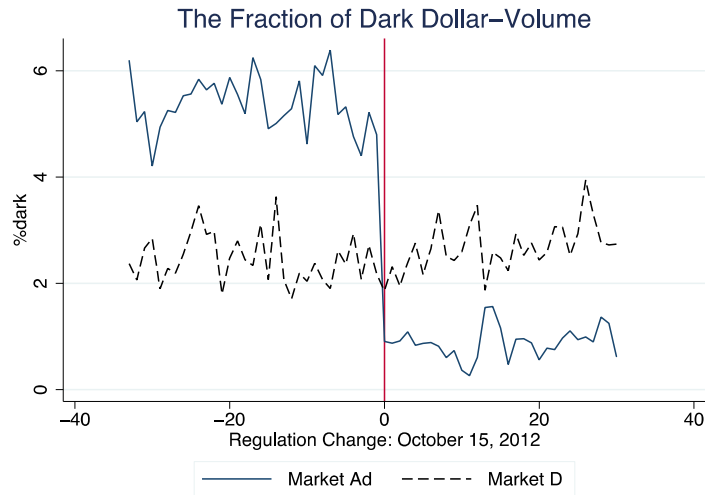
References

- Anand, Amber, Paul Irvine, Andy Puckett, and Kumar Venkataraman, 2012, Performance of Institutional Trading Desks: An Analysis of Persistence in Trading Costs, *Review of Financial Studies* 25(2), 557-598
- Battalio, Robert, 1997, Third market broker-dealers: Cost competitors or cream skimmers?, *Journal of Finance* 52, 341–352.
- Battalio, Robert , Jason Greene, and Robert Jennings, 1997, Do competing specialists and preferencing dealers affect market quality?, *Review of Financial Stlitudies* 10, 969–993.
- Battalio, Robert H. and Corwin, Shane A. and Jennings, Robert H., 2016, Can Brokers Have It All? On the Relation between Make Take Fees & Limit Order Execution Quality, *Journal of Finance* (forthcoming), doi: 10.1111/jofi.12422
- Bessembinder, Hendrik, and Herbert M. Kaufman, 1997, A cross-exchange comparison of execution costs and information flow for NYSE-listed stocks, *Journal of Financial Economics* 46, 293 – 319.
- Brogaard, Jonathan, Hendershott, Terrence, Riordan, Ryan, 2015, Price Discovery Without Trading: Evidence from Limit Orders. Working Paper, <http://ssrn.com/abstract=2655927>.
- Buti, Sabrina, Rindi, Barbara And Ingrid Werner, 2011, Diving into dark pools, 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* 29 (2):238–249.
- Comerton-Forde, Carole, Terrence Hendershott, Charles M. Jones, Pamela C. Moulton, and Mark S. Seasholes, 2010. Time Variation in Liquidity: The Role of Market-Maker Inventories and Revenues, *Journal of Finance* 65 (1), 295-331.
- Comerton-Forde, Carole and Putnins, Talis J., 2015, Dark Trading and Price Discovery?, *Journal of Financial Economics* 118(1), 70-92.
- Degryse, Hans, Frank de Jong, Vincent van Kervel, 2015, The impact of dark trading and visible fragmentation on market quality. *Review of Finance* 19, 1587–1622.
- Devani, Baiju, Lisa Anderson, and Yifan Zhang, 2015, Impact of the Dark Rule Amendments. Trading review and analysis, Investment Industry Regulatory Organization of Canada.
- Easley, David, Nicholas M. Kiefer, and Maureen O’Hara, 1996, Cream-skimming or profit-sharing? The curious role of purchased order flow, *Journal of Finance* 51, 811–833.

- Foley, Sean and Talis Putnins, 2016, Should we be afraid of the dark? Dark trading and market quality, *Journal of Financial Economics*, (forthcoming).
- Hansch, Oliver, Narayan Y. Naik, and S. Viswanathan, 1999, Preferencing, internalization, best execution, and dealer profits, *Journal of Finance* 54, 1799-1828.
- IIROC, 2014, Identifying Trading Groups – Methodology and Results, IIROC Notice 14-0210.
- Korajczyk, Robert and Dermott Murphy, 2015, High Frequency Market Making to Large Institutional Trades, Discussion Paper.
- Kleibergen, F., Paap, R., 2006, Generalized reduced rank tests using the singular value decomposition, *Journal of Econometrics* 133(1), 97-126.
- Menkveld, Albert J. and Yueshen, Bart Z. and Zhu, Haoxiang, Shades of Darkness: A Pecking Order of Trading Venues, *Journal of Financial Economics* (forthcoming).
- Securities and Exchange Commission, 1997, Report on the Practice of Preferencing. Available at: <https://www.sec.gov/news/studies/prefrep.htm>
- Securities and Exchange Commission, 2010, Concept Release on Equity Market Structure. Available at: <https://www.sec.gov/rules/concept/2010/34-61358.pdf>
- Thompson, Samuel B., 2011, Simple formulas for standard errors that cluster by both firm and time, *Journal of Financial Economics* 99, 1–10.
- Zhu, Haoxiang, 2014. Do dark pools harm price discovery? *Review of Financial Studies* 27, 47-789.

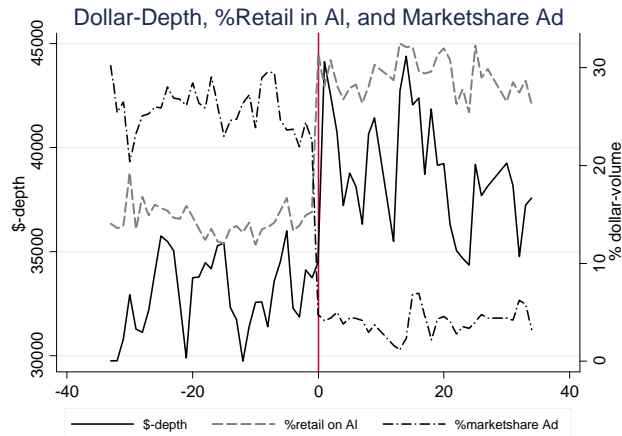


Panel A




Panel B

Figure 1: Dark trading and order volume. *Panel A:* 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 (the solid line) and the average per day per stock dark order volume as a percentage of the total order volume (the dashed line) across all Canadian marketplaces. *Panel B* plots the average per stock per day market shares in terms of dollar volume of the two main dark pools Ad and D. Both panels are for the period August 27 to November 30, 2012. The vertical line at 0 marks October 15, 2012, the date the minimum price improvement rule is introduced.



Panel A



Panel B

Figure 2: Time-weighted quoted depth and retail trading. *Panel A* plots the time-weighted dollar depth for the lit market AI, the percentage of marketable orders (by dollar volume) by retail traders that trade in dark pool Ad, as a fraction of the total dollar volume of marketable orders executed on marketplace A (in dark pool Ad and on lit market AI together), and the percentage of marketable orders (by dollar volume) by retail traders that trade in lit market AI, as a fraction of the total dollar volume of marketable orders executed on marketplace A (in dark pool Ad and on lit market AI together) Dollar volume figures are based on the aggregated traded dollar volume across all securities per day, depth is computed as the average per stock per day. *Panel B* plots the average per stock per day time-weighted quoted dollar-depth for the three main lit markets: AI, B, and C. Both panels are for the period August 27 to November 30, 2012. The vertical line at 0 marks October 15, 2012, the date the minimum price improvement rule is introduced.

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

Table I estimates the effect of the minimum price improvement rule on a market's share of dark liquidity, where dark liquidity is: (1) dark dollar trading volume as a fraction of all dollar trading volume, where a trade is classified as dark if the order on the passive was dark; (2) the fraction of order volume of all volume that is submitted as dark. *MPIR* is a dummy variable for the change in regulation and is 0 before October 15, 2012, and 1 thereafter; *VIX* is the daily realization of the U.S. market volatility index VIX, and *CRCI* the cumulative return on the commodity ETF GSG from August 24 to *t*. We estimate the effect for the entire market in Panel A and by marketplace in Panel B. 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.

	<i>% dark dollar trading volume</i>	<i>% dark order volume</i>
<i>Panel A: Markets in aggregate</i>		
MPIR	-4.19*** (0.48)	-6.73*** (0.82)
VIX	-0.09 (0.07)	-0.26 (0.18)
CRCI	-0.10 (0.07)	-0.43*** (0.15)
Observations	5,884	5,888
<i>Panel B: By marketplace</i>		
Market A x MPIR	-3.88*** (0.32)	-4.35*** (0.67)
Market B x MPIR	-0.07 (0.10)	-0.36 (0.16)
Market C x MPIR	-0.11 (0.08)	-0.15 (0.16)
Market D x MPIR	-0.07 (0.16)	-1.64*** (0.27)
other markets x MPIR	-0.06 (0.05)	-0.30 (0.11)
VIX	-0.02 (0.01)	-0.05 (0.04)
CRCI	-0.02 (0.01)	-0.09*** (0.03)
Observations	29,142	28,678

Table II: Who trades with whom in dark pools?

Table II presents aggregate statistics for dark pool interaction among the different trader types by dollar-volume traded; all numbers are in percent of the total dollar volume. There are four trader type categories: HFT, retail, institutional and other, in addition to a classification of dark pool market maker (DPMM) or not DPMM based on the intermediation score (defined in equation (3)). The described interactions are for all trades that were distinguishable into an active and a passive side. We focus on interactions where the liquidity provider is a DPMM (HFT or other) and where the liquidity demander is not a DPMM (retail, institutional, non-DPMM HFT, and other non-DPMM).

		Liquidity demanders				
		HFT Not DPMM	Retail	Institutional	Other Not DPMM	Total
Liquidity suppliers						
<i>Panel A: Market Ad</i>						
Before	HFT DPMM	0.0	25.8	0.0	0.0	25.8
	other DPMM	0.0	61.7	0.0	0.0	61.7
After	HFT DPMM	0.2	0.7	0.1	0.2	1.2
	other DPMM	0.0	36.0	0.0	0.0	36.0
<i>Panel B: Market D</i>						
Before	HFT DPMM	0.1	1.9	2.8	4.5	9.3
	other DPMM	0.1	0.8	1.8	2.8	5.5
After	HFT DPMM	0.0	0.1	0.1	0.2	0.4
	other DPMM	0.0	0.3	0.4	0.4	1.1

Table III: Regressions for changes in liquidity provision in dark pools

Table III tests whether the intermediation score (defined in equation (3)) explains changes in liquidity supply. The dependent variable is the percentage of liquidity provided by the trader before the MPIR minus the percentage provided after the MPIR. We interact the intermediation score with dummies for markets Ad and D to test whether the coefficients are equal. *HFT*, *Institutional* and *Retail* are dummies for the trader type. Other traders are the omitted trader type. 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	
Intermediation score x market Ad	-3.17*** (0.51)	-3.21*** (0.56)
Intermediation score x market D	-3.01*** (0.50)	-3.04*** (0.55)
HFT		-0.11 (0.43)
Institutional		-0.07 (0.16)
Retail		0.00 (0.86)
Constant	2.98*** (0.49)	3.04*** (0.55)
Observations	895	895

Table IV: Regression on the impact of dark market maker volume on the change in dark market share

Table IV estimates the effect of the percentage of liquidity provided by dark pool market makers (DPMM) in each of the dark pools Ad and D on the change in market share for that market. Δ marketshare is the change in the average per day market share of the venue (in % of dollar volume) from before to after the introduction of the MPIR. *DPMM Volume* is the average per-day fraction of dollar volume provided by dark pool market makers for the security before the introduction of the MPIR in the respective market. The regression is estimated for the cross-section of 92 securities, with two observations per security (one for market Ad and one for D). In the second specification, we interact DPMM Volume with dummies for the respective venues. The estimates are not statistically different. Robust standard errors are in parentheses. *indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	Δ marketshare	
DPMM Volume	-0.06***	
	(0.005)	
DPMM Volume \times market Ad		-0.06***
		(0.005)
DPMM Volume \times market D		-0.04***
		(0.007)
Observations	184	184

Table V: Regression for the usage of trading venues by retail traders

Table V estimates the effect of the MPIR on the usage of venues by retail traders. The variables used are the per security per day retail trader marketable dollar-volume against dark orders and lit orders and non-marketable dollar volume that trades against marketable dark and lit orders; all variables are measured as a percentage of all retail trader dollar-volume. All trading against dark orders on marketplace A occurs in its dark pool facility Ad; all trading against lit orders on marketplace A occurs in its lit limit order book Al. Being a dark pool, market D does not have any lit trading. For each measure, there are two columns. The first column presents the pre-MPIR average of the respective measure per market, the second column displays the estimates of the effect of the MPIR. Independent variables are dummy variables for each market interacted with the dummy for the introduction of the MPIR. *VIX* is the daily realization of the U.S. market volatility index *VIX*, and *CRCI* the cumulative return on the commodity ETF *GSG* from August 24 to *t*. All regression 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.

	marketable against dark		marketable against lit		non-marketable dark		non-marketable lit	
	pre-MPIR average	estimated change	pre-MPIR average	estimated change	pre-MPIR average	estimated change	pre-MPIR average	estimated change
Market A x MPIR dummy	27.6	-18.49*** (1.22)	11.4	13.55*** (0.96)	0	-0.00 (0.01)	13.6	-0.29 (0.20)
Market B x MPIR dummy	0	-0.15 (0.12)	20.4	3.38*** (0.62)	0	-0.01 (0.01)	20.6	0.71 (0.61)
Market C x MPIR dummy	0	-0.15 (0.12)	3.4	1.04*** (0.23)	0	-0.00 (0.01)	0	0.10 (0.13)
Market D x MPIR dummy	1.5	0.12 (0.18)	n/a		0	0.00 (0.02)	n/a	
other markets x MPIR dummy	0	-0.15 (0.12)	1.5	-0.08 (0.17)	0	-0.00 (0.01)	0	0.11 (0.13)
VIX		-0.11*** (0.03)		0.03 (0.06)		0.03 (0.06)		0.08 (0.06)
CRCI		-0.07** (0.03)		0.01 (0.04)		0.01 (0.04)		0.05* (0.03)
Observations		29,071		29,071		29,071		29,071

Table VI: The effect of the MPIR on market quality by marketplace

Table VI 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 (all lit trading on marketplace A occurs in its lit limit order book). We estimate the effect for all three marketplaces simultaneously to capture whether marketplaces 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 the MPIR. *VIX* is the daily realization of the U.S. market volatility index *VIX*, and *CRCI* the cumulative return on the commodity ETF *GSG* from August 24 to *t*. 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	0.09 (0.12)	0.61 (0.43)	0.17*** (0.03)	0.18*** (0.02)
Market B x MPIR	-0.12* (0.07)	-0.21 (0.29)	-0.03 (0.03)	-0.02 (0.03)
Market C x MPIR	-0.09 (0.12)	-0.14 (0.41)	0.05 (0.03)	0.05** (0.02)
<i>VIX</i>	0.06** (0.03)	0.26*** (0.10)	-0.01** (0.01)	-0.01 (0.00)
<i>CRCI</i>	-0.03* (0.01)	-0.09 (0.06)	0.01 (0.01)	0.01 (0.00)
Observations	17,388	17,388	17,388	17,388

Table VII: Instrumental variable regression on the impact of retail on quoted depth

Table VII tests whether the extent of retail trading causally affects posted depth. The estimation is performed in a two-stage instrumental variable regression. The explanatory variable of interest is *%mktble retail on AI of all mktble retail*, which is the ratio of marketable retail dollar volume traded on market AI to retail dollar volume across all markets, per stock per day. We instrument this variable in the first stage by the average daily per-security pre-MPIR market share of dark pool Ad in terms of dollar volume (*%Admarketshare*), interacted with the event dummy for MPIR. The dependent variable *quoted depth* is measured as the natural logarithm of the time-weighted quoted dollar depth at the best bid and offer prices on market AI. $\ln(MCap)$ is the log of security i 's market capitalization, VIX is the daily realization of the U.S. market volatility index VIX, and $CRCI$ the cumulative return on the commodity ETF GSG from August 24 to t . As first stage specification tests we include the Kleibergen-Papp underidentification statistic (and its p-value) and a standard F-test statistic. 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

	<i>%mktble retail on AI of all mktble retail</i>			
<i>Panel A: First Stage Regression</i>				
<i>%Admarketshare x MPIR</i>	2.14*** (0.01)	2.10*** (0.12)	2.19*** (0.12)	2.18*** (0.12)
VIX		0.204 (0.17)		0.31* (0.16)
CRCI			0.10 (0.11)	0.18 (0.12)
Fixed effects	yes	yes	yes	yes
R-squared	0.39	0.39	0.39	0.40
Observations	5881	5881	5881	5881
Kleinbergen Paap rkLM	23.7	21.8	19.78	19.67
p-val for K-P	0.0	0.0	0.0	0.0
F-stat	34.4	204.9	193.4	144.5
			<i>quoted depth</i>	
<i>Panel B: Second Stage Regression</i>				
<i>% mktble retail on AI of all mktble retail</i>	0.01*** (0.00)	0.01*** (0.00)	0.01*** (0.00)	0.01*** (0.00)
VIX		-0.01 (0.01)		-0.01** (0.01)
CRCI			-0.01 (0.00)	-0.01** (0.00)

Table VIII: Regression on change in trading costs and benefits for retail traders

Table VIII estimates the effect of the introduction of the minimum price improvement rule on trading costs for retail traders. We consider four measures: the effective spread, measured in basis points of the prevailing midpoint, that retail traders pay for their marketable orders; the total amount of price improvement that retail traders receive in market Ad relative to the NBBO in dollars; the price improvement per trade that retail traders receive in market Ad; and the maker-taker fees that retail brokers pay per dollar traded, where this measure is computed as the total maker rebates minus taker fees. *VIX* is the daily realization of the U.S. market volatility index *VIX*, and *CRCI* the cumulative return on the commodity ETF GSG from August 24 to *t*. Panel A presents the pre-MPIR averages, Panel B presents the estimated effect of the MPIR. All regression 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.

	Effective spreads in bps	Dollar amount price improvement	Price improvement per trade	Maker-taker fees per dollar traded in bps
<i>Panel A: Pre-MPIR average</i>				
	1.5	\$102.22	\$0.64	-0.3
<i>Panel B: Estimated effect</i>				
MPIR	0.04 (0.30)	-32.89*** (10.51)	1.82*** (0.12)	-0.40*** (0.09)
VIX	0.02 (0.05)	-6.35** (2.88)	0.01 (0.03)	0.00 (0.01)
CRCI	-0.03 (0.04)	-1.51 (2.43)	0.00 (0.02)	0.02* (0.01)
Observations	5,771	5,796	5,796	5,738

Table IX: Regression on change in returns to trading for HFTs and dark pool market makers

Table IX estimates the effect of the introduction of the minimum price improvement rule on trading costs and benefits for HFT Dark Pool Market Makers and other Dark Pool Market Makers. We consider the returns to trading for HFTs and the maker-taker fees per dollar traded, where this measure is computed as the total maker rebates minus taker fees. Panels A and C present the pre-MPIR averages, Panels B and D present the estimated effects of the MPIR. *MPIR* is a dummy variable for the change in regulation and is 0 before October 15, 2012, and 1 thereafter; *VIX* is the daily realization of the U.S. market volatility index VIX, and *CRCI* the cumulative return on the commodity ETF GSG from August 24 to t . All regression 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.

	returns to trading			Maker-taker fees per dollar traded in bps		
	all venues	Ad and AI	AI	all venues	Ad and AI	AI
<i>Panel A: pre-MPIR average for HFT Dark Pool Market Makers</i>						
	1.2	-26.5	6.1	1.1	0.3	0.3
<i>Panel B: estimated effect for HFT Dark Pool Market Makers</i>						
MPIR	-157.39 (140.17)	-110.31 (138.89)	-115.76 (135.75)	0.30*** (0.11)	0.24*** (0.07)	0.24** (0.07)
VIX	-88.06** (43.29)	-.56.46* (36.81)	-63.98* (37.98)	0.02* (0.01)	0.00 (0.00)	0.00 (0.00)
CRCI	-61.84 (39.27)	-49.39 (31.95)	-49.31 (33.24)	0.01*** (0.00)	0.01*** (0.00)	0.01** (0.00)
Observations	5,792	5,792	5,792	5,781	5,781	5,781
<i>Panel C: pre-MPIR average for other Dark Pool Market Makers</i>						
	-59.5	-63.3	8.5	0.0	-0.1	-0.1
<i>Panel D: estimated effect for other Dark Pool Market Makers</i>						
MPIR	-64.35 (56.89)	-53.62 (60.25)	-8.74 (25.39)	0.49** (0.25)	0.56*** (0.24)	0.56** (0.24)
VIX	-17.86 (16.53)	-16.07 (16.81)	17.28** (7.08)	-0.03 (0.03)	-0.03 (0.03)	-0.03 (0.03)
CRCI	-46.78*** (19.55)	-44.55** (19.76)	-13.04** (5.00)	-0.11*** (0.03)	-0.10*** (0.03)	-0.10*** (0.03)
Observations	5,792	5,792	5,792	4,803	4,803	4,803

Table X: Regression for institutional trader fill rates

Table X presents estimation results for the effect of the introduction of the minimum price improvement rule on buy-side traders' fill rates for *passive* orders, as defined in Section V.C (equation (10)) and in Table VIII. All dark trading on market A occurs in its dark pool facility Ad, all lit trading on market A occurs in its lit limit order book Al. Market D is a dark pool and does not have lit trading. For each measure, there are two columns. The first column presents the pre-MPIR average of the respective measure per market, the second column displays the estimates of the effect of the MPIR. *MPIR* is a dummy variable for the change in regulation and is 0 before October 15, 2012, and 1 thereafter; *VIX* is the daily realization of the U.S. market volatility index *VIX*, and *CRCI* the cumulative return on the commodity ETF GSG from August 24 to *t*. All regression 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.

	lit		dark	
	Pre-MPIR average	Estimated change	Pre-MPIR average	Estimated change
<i>Panel A: All markets together</i>				
MPIR	10.8	0.13 (0.42)	3.9	1.59** (0.71)
VIX		0.26*** (0.10)		-0.28*** (0.09)
CRCI		0.20** (0.09)		0.10 (0.11)
Observations		5,700		5,653
<i>Panel B: By market</i>				
Market A	10.6	-0.78 (0.55)	13.3	4.14 (2.54)
Market B	11.2	0.44 (0.42)	7.7	0.44 (1.42)
Market C	6.8	-0.92** (0.46)	8.1	1.84 (1.20)
Market D			3.4	3.54*** (0.93)
VIX		0.21*** (0.08)		-0.18 (0.17)
CRCI		0.16** (0.08)		0.51** (0.20)
Observations		17,214		17,894

Table XI: Package characteristics for institutional traders

Table XI, Panel A reports package characteristics for institutional and other traders (i.e., non-high frequency, non-dark market-making and non-retail traders.) The duration of the package is the number of (consecutive) days that the package is traded. The number of trades is the number of transactions that are part of a package. The absolute value of the order imbalance is the difference of buying and selling volume relative to all volume on the day before the package starts trading. The percentage of volume of the package measure is computed relative to all volume that trades on the days that the package trades. The percentage of buying and selling volume measures the package volume relative to all buyer- and seller-initiated volume. The VIX is reported as it pertained on the first day of the package; similarly the CRCI is the commodity index ETF GSG's return since the beginning of the sample. Panel B reports the implementation shortfall benchmarked against the first trade price and the volume weighted average price. The sample includes 23,527 packages and excludes pre-arranged block trades.

	before	after
<i>Panel A: Package characteristics</i>		
duration of package	2.5	2.4
number of trades in package	105.2	107.2
order imbalance in %	10.4	11.9
package volume / total volume in %	0.9	1.0
package volume/ selling volume in % for sales	1.2	1.2
package volume/ buying volume in % for buys	1.5	1.4
VIX (at day of first order in package)	15.6	17.0
CRCI	0.1	-4.1
<i>Panel B: Mean implementation shortfall</i>		
Implementation shortfall against first trade price	-3.1	-0.7
Implementation shortfall against VWAP	-0.2	0.2

Table XII: Regression on changes in implementation shortfall for institutional traders

Table XII estimates the effect of the introduction of the minimum price improvement rule on implementation shortfall for institutional and other traders. Implementation shortfall is measured as the difference of the price achieved on the package and the first trade price on the day the package began. 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.

	Implementation Shortfall			
MPIR	12.76*	20.36**	18.86**	18.86**
	(7.44)	(8.32)	(8.70)	(8.70)
duration of package	-5.15**	-9.51***	-8.54**	-8.54**
	(2.27)	(2.50)	(3.42)	(3.42)
trades in package	0.01	0.09***	0.09***	0.09***
	(0.01)	(0.03)	(0.03)	(0.03)
VIX	-1.25	-0.45	-0.44	-0.44
	(1.65)	(1.76)	(1.68)	(1.68)
package percentage of sell volume		-26.15***	-25.34***	-25.34***
		(7.68)	(9.04)	(9.04)
package percentage of buy volume		2.03	2.08	2.08
		(1.91)	(2.15)	(2.15)
return standard deviation	-1.21	-1.20	-0.93***	-0.93***
	(1.00)	(1.00)	(0.28)	(0.28)
lagged order imbalance	-0.19	-0.56***	-0.31	-0.31
	(0.14)	(0.20)	(0.20)	(0.20)
"other"-dummy	-148.79***	-163.01***	-162.80***	-162.80***
	(54.08)	(55.41)	(57.81)	(57.81)
buy-dummy	126.37***	63.80*	65.21*	65.21*
	(39.40)	(36.33)	(36.00)	(36.00)
CRCI	1.20	2.72	2.51	2.51
	(1.45)	(1.76)	(1.88)	(1.88)
package percent of volume	-1.29			
	(2.87)			
Constant	82.22**	140.59***	133.41***	133.41***
	(36.39)	(43.60)	(45.20)	(45.20)
Observations	23,527	23,527	23,527	23,527
Firm Fixed Effects	Yes	Yes		
SE clustered by firm and date			Yes	
SE clustered by firm, date and trader ID				Yes