

Should we be afraid of the dark? Dark trading and market quality ^{*}

Sean Foley ^a and Tālis J. Putniņš ^b

^a *University of Sydney*

^b *University of Technology Sydney and Stockholm School of Economics in Riga*

10 January 2015

Abstract

We exploit a unique natural experiment – recent restrictions of dark trading in Canada – and proprietary trade-level data to analyze the effects of dark trading. Disaggregating two types of dark trading, we find that dark limit order markets are beneficial to market quality, reducing quoted, effective and realized spreads and increasing informational efficiency. In contrast, dark midpoint crossing systems do not benefit market quality. Our results support recent theory that dark limit order markets encourage aggressive competition in liquidity provision. We discuss implications for the regulation of dark trading and tick sizes.

JEL classification: G14

Keywords: dark pool, dark trading, regulation, liquidity, market efficiency, transparency

^{*} Email: sean.foley@sydney.edu.au and talis.putnins@uts.edu.au. The Internet Appendix that accompanies this paper is available at <http://goo.gl/NY4svm>.

Foley is grateful for funding from the Capital Markets CRC Limited and was a visiting researcher at the Ontario Securities Commission for part of the time spent on this study. We thank the Securities Industry Research Centre of Asia-Pacific (SIRCA), as well as TSX, ITG MatchNow, Chi-X Canada, and Alpha Intraspread for providing data. We also thank Kent Bailey, Alina Bazavan, Torstein Braaten, Carole Comerton-Forde, David Easley, Doug Foster, Corey Garriott, Amy Kwan, Ingrid Lo, Katya Malinova, Maureen O'Hara, Andreas Park, Andrew Patton, Paul Redman, Steve Satchell, Jonathan Sylvestre, Vincent van Kervel, Michael McKenzie, Terry Walter, Mao Ye, and seminar participants at the Northern Finance Association Annual Meeting, the CAFS Behavioural Finance and Capital Markets Annual Conference, the Multinational Finance Society Annual Conference, the Australasian Finance and Banking Conference, the IFMA International Conference on Finance, the Bank of Canada, the Ontario Securities Commission, and the Australian Securities and Investments Commission for helpful comments and suggestions. We also thank a number of Canadian market operators and staff at the Ontario Securities Commission and Bank of Canada for sharing their knowledge about the structure of dark trading in Canada.

1. Introduction

While trading without pre-trade transparency has long been a feature of equity markets in the form of upstairs block trades, only in recent years with the introduction of continuous dark pools for smaller sized non-transparent orders has it attracted the attention of regulators worldwide. Dark pools have been very successful in attracting order flow; they are estimated to account for approximately 15% of US consolidated volume, 10% in Europe, 14% in Australia, and 10% in Canada.¹ Their success is likely related to what proponents argue are their advantages, including the ability to avoid large orders being front run, reduced information leakage, and lower market impact costs.

The rapid growth in dark trading has caused considerable concern, especially among market regulators. For example, the US Securities Exchange Commission (SEC) proposed rules in 2009 for the “Regulation of non-public trading interest” and in 2010 issued a Concept Release on Equity Market Structure calling for comments on the issue of dark liquidity. The Financial Industry Regulatory Authority (FINRA) in 2013 proposed a new set of disclosure requirements for dark pool operators, and in 2014 the SEC Chairman in a speech said “transparency has long been a hallmark of the US securities markets, and I am concerned by the lack of it in these dark venues”.² The Committee of European Securities Regulators (CESR) has undertaken a review of dark trading with recommendations to limit the activities of broker crossing systems, and the European Commission in 2013 proposed EU-wide rules that cap the volume traded in dark pools. While these regulatory bodies have made proposals and conducted public consultations regarding dark trading, their hesitance in introducing new regulations reflects the scarcity of evidence on the costs and benefits of dark pools, and how these costs/benefits are distributed between different market participants. This study aims to address this problem by empirically analyzing the impact of dark trading on market quality.

We exploit the unique natural experiment created by the introduction of minimum price improvement rules for dark trading in Canada in October 2012; the first such regulation in the

¹ The US estimate is from Rosenblatt Securities for April 2013. The Europe estimate is for July 2013 using Thomson Reuters data as reported by the Wall Street Journal (<http://online.wsj.com/article/BT-CO-20130812-701291.html>). The Australian estimate is from the Australian Securities and Investments Commission Report 331 for the September quarter 2012 and includes some internalization. The Canadian estimate combines statistics from the Investment Industry Regulatory Organization of Canada and proprietary data obtained for this study and corresponds to the period Aug-Dec 2012.

² See Wall Street Journal, 6 June 2014 (<http://online.wsj.com/articles/sec-chairman-unveils-sweeping-proposals-to-improve-markets-1401986097>).

world.³ The rules require that dark orders of 5,000 shares or less provide one full tick of price improvement (or half a tick if the spread is constrained at one tick). When the rules came into effect, dark trading in Canada fell by over one third, literally overnight. Using the regulation as a source of exogenous variation in dark trading, and proprietary trade-level data from dark trading venues, we analyze the causal impact of dark trading on liquidity and informational efficiency. Our empirical design overcomes the endogeneity issues that have hindered analysis of dark trading and market quality.

We disaggregate dark trading into two types that theory suggests should have different effects. The first is dark trading at a single price such as the midpoint of the national best bid and offer (NBBO). We refer to this type of dark trading as ‘one sided’ because at any point in time dark liquidity can only exist on either the buy- or the sell-side, but not both. One-sided dark trading is characterized by a relatively low execution probability (particularly for traders that tend to cluster on one side of the market such as informed traders), the absence of profitable dark market making strategies due to the zero dark spread, and imperfect concealment of trading intentions due to the ability for probing orders to infer the direction of the dark order imbalance. The second type, ‘two-sided’ dark trading, occurs at different prices on both the buy- and sell-sides of the market, and more closely resembles a dark limit order market. In contrast to one-sided dark trading, traders in a two-sided dark market can instantly execute their orders as long as dark liquidity exists, can facilitate liquidity provision strategies that earn the dark bid-ask spread, and provide better concealment of trading intentions.

Our main finding is that two-sided dark trading, in moderate levels, is beneficial to liquidity and informational efficiency. It tends to lower quoted, effective and realized spreads, reduces price impact measures of illiquidity, makes prices closer to the random walk process that would be expected under informational efficiency. The magnitudes of these effects are economically meaningful. In contrast, we do not find consistent evidence that midpoint dark trading has a significant effect on market quality.

Aggregating across the two types of dark trading, our results suggest that dark trading is more likely to benefit market quality the greater the proportion of two-sided dark trading. Furthermore, changes in the composition of dark trading can impact market quality even if the aggregate level remains unchanged. An increase in two-sided dark trading relative to the level of one-sided dark trading is likely to benefit market quality. Our results are robust to a range of

³ The Australian Securities and Exchange Commission (ASIC) subsequently introduced minimum price improvement rules for dark trades in Australia, which took effect in May 2013.

alternative specifications, fixed effects, sub-period tests, a variety of control variables including matching stocks in a control market, and are similar for the largest and the smallest stocks in our sample.

Our results have support in the theoretical literature. The positive effect of two-sided dark trading (which resembles a dark limit order market) on market quality is consistent with a number of models that analyze pre-trade transparency in limit order markets. For example, Boulatov and George (2013) find that dark limit order markets encourage informed traders to supply liquidity because they can profit from doing so without revealing much of their private information. Transparency makes them reluctant to supply liquidity because other traders gain an informational advantage by observing the limit order schedules before deciding how to trade. Boulatov and George (2013) show that dark limit order markets not only increase liquidity provision by informed traders but also the aggressiveness with which they trade, which improves informational efficiency. Our results suggest that strong competition among informed traders in providing dark liquidity has positive spillover effects on the lit market, where liquidity providers are forced to narrow spreads to compete with dark liquidity.

In contrast, theory identifies mechanisms by which one-sided (midpoint) dark trading can harm liquidity. For example, Zhu (2014) shows that informed traders are less likely than uninformed traders to send orders to a dark midpoint market because their tendency to cluster on the same side of the market gives them low execution probability in the dark. The increased concentration of informed traders in the lit market increases adverse selection risks and harms liquidity. While we do not find that midpoint dark trading harms liquidity, our results indicate that midpoint dark trading does not benefit liquidity like two-sided dark trading, consistent with an opposing increase in adverse selection risks.

Our paper contributes to the recent empirical studies of dark trading by providing causal evidence from a unique natural experiment with detailed proprietary data. Overcoming the endogeneity problem and obtaining sufficiently detailed data have been significant challenges for empirical studies. Our analysis also provides a potential reason why empirical studies of dark trading in different markets sometimes find different results; namely, the composition of dark trading, which determines its impact, varies across countries and dark pools. Kwan et al. (2014) examine how the tick size influences dark trading and find that market participants use US dark pools to obtain a finer pricing grid when stock prices are constrained by the tick size. Along similar lines, Buti et al. (2014) analyze how dark venues can be used for ‘queue jumping’ ahead of displayed liquidity by trading at sub-penny increments. Although they analyze only the subset

of dark trading that is associated with queue jumping, their finding that non-midquote (fractional price improvement) dark trading is associated with improved market quality is consistent with our results. Ready (2014) analyzes the determinants of volume in two block crossing systems (Liquidnet and POSIT) and finds that dark trading activity is higher for stocks with lower levels of adverse selection risks. Degryse et al. (2014) analyze 52 Dutch stocks and conclude that fragmentation of volume across visible order books improves consolidated liquidity, but dark trading has a detrimental effect. Buti et al. (2011) use data from 11 US dark pools and conclude that dark pool activity improves spreads, depth and short-term volatility. Nimalendran and Ray (2014) examine data from one US dark pool and find informational linkages between lit and dark venues due to traders using algorithms to split orders across venues. Comerton-Forde and Putniņš (2013) find that in Australia low levels of dark trading may improve price discovery, but when dark volume exceeds 10% of total trading, informational efficiency deteriorates.

The results from our analysis have a number of policy implications, in particular given that dark trading is high on the current agenda of many regulators. Our results point to the fact that dark trading should not be treated as a homogenous group; it is important to distinguish between different types of dark trading when developing policy. The effects of aggregate dark trading depend on the composition of dark trading types within the aggregate. The larger the proportion of two-sided dark trading in the aggregate the more likely the aggregate dark trading benefits rather than harms market quality. A harmful level of dark trading in one country may not be harmful in another due to differences in the composition of dark trading types. It follows that in designing regulation it is important to consider not only the regulation's effect on the level of dark trading but also on the composition of dark trading types. For example, minimum price improvement regulation in Canada not only decreased the level of dark trading but also caused a shift from two-sided dark trading to dark trading at the midpoint of the NBBO.

Finally, our results have implications for tick size regulation, which has been on the regulatory agenda for many years and is at the center of a US SEC Pilot Program. Dark trading is more active in stocks whose spread is constrained by the tick size (Kwan et al., 2014). Our results suggest that when dark trading is used as a way of obtaining a finer price grid it can benefit market quality as long as the price grid allows dark liquidity to concurrently exist on both the buy- and sell-sides of the market. Minimum price improvement requirements can force dark trades to occur at the midpoint in the large number of stocks that are constrained by the tick size, and consequently such rules can have unintended negative effects on market quality. Our results suggest that a way of improving the effectiveness of minimum price improvement requirements is

to ensure tick sizes do not constrain the lit spread. If the price grid is sufficiently fine, dark trades can offer price improvement while maintaining a two-sided dark market.

2. Theory and hypotheses

2.1 Types of dark trading

Dark trading is a broad term that can include: (i) trading in dark pools (automated non-transparent trading venues), (ii) non-transparent order types on lit exchanges that interact with lit order flow, and (iii) internalization of order flow by brokers acting as principals. We focus on trading in dark pools and trading with non-transparent order types.⁴ Internalization of order flow is associated with issues in addition to transparency, such as cream-skimming (e.g., Easley et al., 1996).

There is a great deal of variation in how dark trading occurs in different venues. One categorization of dark trading types that is important both theoretically and empirically is whether dark trades execute (i) at a single price such as the midpoint of the national best bid and offer (NBBO), or (ii) at different prices on both the buy- and sell-sides of the market (e.g., at the NBBO or at prices that are a fraction of a spread in from the NBBO). We refer to these two types as ‘one-sided’ and ‘two-sided’ dark trading (or ‘midpoint’ and ‘fractional price improvement’ dark trading), respectively.⁵

There are three important differences between one-sided and two-sided dark trading. The first is execution probability and its impact on order routing decisions. In a market with one-sided dark trading, at any point in time, dark liquidity can be available to buyers or to sellers depending on the order imbalance at the midquote, but not to both. This feature limits execution probability, in particular for traders that tend to cluster on the same side of the market (e.g., informed traders). Traders that are not able to obtain immediate execution in the dark may route their orders to lit venues, which tends to increase imbalances in lit order flow.

Second, one-sided dark trading can reveal more information about trading intentions than two-sided dark trading. In a one-sided dark market, market participants can readily infer the

⁴ In contrast to the US, Europe and Australia, fair access rules in Canada have hampered the development of automated ‘internalizers’.

⁵ Dark trading at a single price such as the midpoint of the NBBO or the volume-weighted average price (VWAP) is not uncommon. Examples include dark pools such as ITG Posit, Liquidnet, ASX CentrePoint, ITG MatchNow (after Oct. 2012), Instinet VWAP Cross, Turquoise Midpoint, and midpoint dark order types offered by exchanges such as the TSX, Chi-X, Nasdaq, BATS, and DirectEdge. Similarly, two-sided dark trading is not uncommon in both dark pools and as non-transparent order types on lit exchanges. Examples include ITG MatchNow (before Oct. 2012), Alpha Intraspread (before Oct. 2012), Instinet CBX, Turquoise Integrated, Credit Suisse CrossFinder, Goldman Sachs Sigma X, Deutsche Bank Super X, Citi Match, and UBS PIN.

direction of the order imbalance at the midquote by submitting probing orders at the midquote and observing whether they immediately execute or not. The direction of the imbalance may reveal some private information. In contrast, a two-sided dark market can have dark liquidity posted on both the buy- and sell-sides, with no information about the quantities available on either side. In such a situation it is not possible to use probing trades to infer information about an imbalance between buyers and sellers.

Third, dark liquidity provision can be profitable in a two-sided dark market because a liquidity provider can earn the non-zero spread in the dark. In contrast, there is no spread in a one-sided dark market and therefore there is little incentive for market participants to act as dark liquidity providers without an alternative reason for wanting to trade.

2.2 Impact of dark trading

Theoretical predictions about the impact of one-sided dark trading differ from those for two-sided dark trading. We discuss models of these two types of dark trading in turn.

Zhu (2014) models the choice between a lit market and a one-sided dark midpoint crossing system as a problem of execution probability. Informed traders tend to cluster on the same side of a market, buying when they have good news and selling when they have bad news. Thus, their execution probability in the dark is lower than that of uninformed traders. Consequently, informed traders prefer to trade on lit markets, which increases lit market adverse selection risk, quoted spreads, price impacts, and improves price discovery.⁶

An earlier model by Hendershott and Mendelson (2000) uses a setting that is similar to Zhu (2014).⁷ Their model is cast as a dealer market that competes with a crossing network; however, the market structure is similar to a two-sided lit exchange that competes with a dark pool that executes trades at the midpoint of the lit market's quotes. Hendershott and Mendelson (2000) show that the introduction of a competing crossing network has many different effects, and depending on which effects dominate, the crossing network can harm or improve liquidity. First, the midpoint crossing system attracts new uninformed order flow, which tends to increase

⁶ Ye (2011) obtains opposite predictions with respect to price discovery using a Kyle (1985) framework. In the model, a strategic, monopolistic informed trader knows that his trades in the lit market have price impact and therefore decrease the profits on his dark trades. Consequently, the informed trader reduces the aggressiveness of his trading in the lit market, which impedes price discovery. Because uninformed traders in Ye (2011) are exogenous, the model does not make predictions about the impact of dark trading on liquidity.

⁷ In Hendershott and Mendelson (2000), if informed traders fail to execute their orders in the dark pool, dealers are willing to execute the orders at the quotes that they originally posted. In contrast, Zhu (2014) assumes that orders that are routed to the lit market after failing to execute in the dark will face a different, updated set of quotes. This modeling choice leads to different empirical predictions.

the liquidity of the market. Second, by executing a fraction of the balanced order flow but none of the order imbalance, the crossing system can increase order imbalances in the lit market, leading to higher inventory holding risks for liquidity suppliers and wider spreads. Third, the tendency for traders to use the lit market as a ‘market of last resort’ (i.e., route orders to the lit market if they fail to execute in the dark) tends to increase adverse selection risks in the lit market because the order imbalance that spills over into the lit market is often due to informed traders.

A common theme in the models discussed above is that one-sided dark trading impacts market quality by changing the composition of order flow received by the lit market. In Zhu (2014) the dark pool increases adverse selection in the lit market because informed traders disproportionately choose to stay in the lit market. In Hendershott and Mendelson (2000), traders that use the lit market as a ‘market of last resort’ increase adverse selection and inventory holding risks. Importantly, both of these mechanisms are more likely to arise from midpoint dark trading than two-sided dark trading. This leads to our first hypothesis:

Hypothesis 1: One-sided (midpoint) dark trading harms market liquidity.

Turning to two-sided dark trading, a number of theoretical studies predict that less pre-trade transparency in a limit order market encourages more aggressive competition in liquidity provision. For example, Boulatov and George (2013) show that liquidity and price discovery are better in a dark limit order market than in a transparent one. With pre-trade transparency, informed traders that submit limit orders earn a profit from providing liquidity, but in doing so they give away some of their private information to other traders that observe the limit orders before submitting market orders. Without pre-trade transparency, informed traders can profit from providing liquidity without giving away their trading intentions because the limit orders are not displayed. Therefore, a larger proportion of informed traders compete to provide liquidity, which makes the market more liquid. Informed traders that supply liquidity do so more aggressively than those that demand liquidity, which increases informational efficiency.

Although Boulatov and George (2013) do not model side-by-side trading in a dark and lit venue, the mechanisms they identify are relevant for understanding the impact of dark pools. Informed traders can compete in liquidity provision more aggressively in two-sided dark pools because they can do so without revealing their trading intentions. Thus, informed liquidity providers in the dark can undercut the spreads offered in the lit market. The increased aggressiveness of informed trading can help price discovery. It can also improve liquidity in the

lit market by forcing lit liquidity providers to compete with the dark liquidity. Importantly, this mechanism only applies to two-sided dark trading because (i) dark liquidity provision can be profitable because of the positive bid-ask spread in the dark, and (ii) informed traders do not reveal their information by posting dark liquidity in a two-sided dark market.⁸

Dark trading can also benefit liquidity by increasing the number of liquidity providers when liquidity fragments across multiple venues (e.g., Biais et al., 2000). Dark trading allows traders to bypass time priority in a lit limit order book. Foucault and Menkveld (2008) show that such ‘queue jumping’ can encourage competition in liquidity provision and compete away profits on inframarginal limit orders, thereby increasing liquidity. Finally, dark trading can benefit liquidity by allowing liquidity providers to compete on a finer pricing grid (e.g., Biais et al., 2010; Buti et al., 2014).

Most of the mechanisms by which dark trading can have a positive impact on liquidity occur when dark trading is two-sided. This leads to our second hypothesis:

Hypothesis 2: Two-sided dark trading improves market liquidity.

The models discussed above suggest that different types of dark trading can have different effects. Therefore, the impact of aggregate dark trading (aggregating across different types of dark trading) depends on the composition of dark trading types within the aggregate. This leads to a third hypothesis:

Hypothesis 3: The effect of aggregate dark trading depends on the composition of dark trading types within the aggregate. Dark trading is more likely to have a positive effect on liquidity the higher the level of two-sided dark trading relative to one-sided dark trading.

⁸ Models of liquidity supply/demand decisions in different settings also find that reduced pre-trade transparency can benefit liquidity. For example, Rindi (2008) models the effects of pre-trade transparency of trader identities. Informed traders are effective liquidity suppliers, as they face little or no adverse selection costs. When information acquisition is endogenous and costly, transparency reduces the number of informed traders, which harms liquidity. Buti and Rindi (2013) model the use of reserve orders, in which only a portion of the limit order is displayed and the remainder is hidden. They find that reserve orders encourage traders to compete in liquidity provision and therefore increase gains from trade.

3. Institutional setting and the natural experiment

3.1 Trading venues and order types

Canada, like many other countries, has experienced rapid fragmentation of its trading landscape during the last decade. In addition to the main listing exchange, the Toronto Stock Exchange (TSX), at the time of our sample there are five Alternative Trading Systems on which trading occurs with pre-trade transparency ('lit' venues): Alpha, Chi-X, Pure Trading, TMX-Select, and Omega.⁹ TSX is still the dominant market, executing approximately 61% of Canadian dollar volume during the sample period, followed by Alpha (15%) and Chi-X (13%). Additionally, there are four continuous auction venues in which orders can be submitted without pre-trade transparency: ITG's MatchNow, Alpha Intraspread, Chi-X and TSX.¹⁰ MatchNow and Alpha Intraspread are dark pools because only dark orders can be submitted to these venues and therefore dark orders execute exclusively against other dark orders. They account for approximately 3.0% and 2.5% of Canadian dollar volume, respectively, during our sample period. In contrast, Chi-X and TSX allow dark orders in addition to lit orders and the two types of orders interact and can execute against one another. Following the introduction of the minimum price improvement rules in 2012, Alpha Intraspread (which had been a stand-alone continuous dark pool) was merged with the Alpha lit exchange. Subsequently, Intraspread orders are able to interact with both lit and dark liquidity, similar to the TSX and Chi-X. Table 1 provides a summary of the market shares, order types and other characteristics of Canadian trading venues.

< Table 1 here >

Prior to 15 October 2012, all dark orders were required to provide some price improvement, resulting in dark executions within the national best bid and offer (NBBO) spread. The required amount of price improvement, however, was not specified legislatively. MatchNow and Intraspread both offered two types of price improvement: midpoint (i.e., 50% improvement over the NBBO) and 20% (on MatchNow) or 10% (on Intraspread) fractional improvement over the NBBO. Price improvement of 10%, for example, means that if a stock has a national best bid of \$10.05 and a national best offer of \$10.06, a passive dark buy order could be placed at a price of \$10.051 (an improvement of 10% of the NBBO spread) and a passive dark sell order could be placed at a price of \$10.059 (also an improvement of 10% of the NBBO spread). Midpoint orders

⁹ A sixth alternative trading system (Chi-X 2) was added in April 2013, after our sample period.

¹⁰ MatchNow was launched in July 2007. Intraspread was launched in May 2011. Chi-X introduced dark midpoint orders in February 2008. TSX introduced dark orders between April and May 2011.

facilitate one-sided dark trading, whereas venues that accommodate ‘fractional’ price improvement (e.g., 10% and 20%) facilitate two-sided dark trading. More details on the order types and order execution priority rules are provided in the Internet Appendix.

In addition to the continuous dark pools, systems to negotiate block trades without pre-trade transparency have existed for decades. The two that operate during our sample are Liquidnet and Instinet.¹¹ While these systems also have limited or no pre-trade transparency, they differ from dark pools that have captured significant market share in recent years in that they are generally only used by large institutional traders, are non-continuous, and only offer services for block trades. The combined market share of Liquidnet and Instinet in Canadian equities during the third quarter of 2012 was only 0.2%.¹² Brokers are also able to internalize orders off-market. However, the order exposure rule (requiring internalized trades are given one full tick of price improvement) together with the fair access regulations (prohibiting venues from providing exclusive access to certain types of clients) have hampered the development of automated ‘internalizers’ such as those that exist in the US, Australia and elsewhere.¹³

3.2 Regulation of dark trading

The Investment Industry Regulatory Organization of Canada (IIROC) sets and enforces the Universal Market Integrity Rules (UMIR), which govern trading on debt and equity marketplaces in Canada. On 13 April 2012 IIROC notice 12-0130 announced changes to the UMIR, which became effective on 15 October 2012. These changes impose a minimum price improvement by dark orders of one full tick relative to the prevailing NBBO, except when the spread is already constrained to one tick, in which case dark orders are allowed at the midpoint of the NBBO (half a tick price improvement). This new requirement provides an exemption for dark orders larger than either 50 standard trading units (STU), which is usually 5,000 shares, or \$100,000.¹⁴ Such large dark orders are able to execute at the NBBO, without providing *any* price improvement, as long as they give priority to lit orders at the same price on the same trading

¹¹ These venues provide ‘trade blotter’ services that facilitate the execution of ‘upstairs’ trades. Typically, clients enter their desire to trade large blocks into the system. The system then identifies whether any potential counterparties exist, and if so, allows the counterparties to negotiate the trade anonymously.

¹² This statistic is taken from the IIROC “Marketplace Statistics Report” available at www.iroc.ca.

¹³ See the Universal Market Integrity Rules (UMIR) section 6.3.

¹⁴ A standard trading unit is 100 shares for stocks priced above \$1.00, 1,000 shares for stocks priced between \$0.10 and \$1.00, and 10,000 shares for stocks priced below \$0.10.

venue. Prior to the change in regulation, dark orders were required to provide a “better price” than the prevailing NBBO but with no minimum increment of price improvement.¹⁵

The minimum price improvement requirements caused a significant decline in dark volume, and a change in the mix of one-sided versus two-sided dark trading. Figure 1 documents the significant decrease in dark volume as a result of the change in regulation. The level of dark trading fell from approximately 8.5% of dollar volume during the two months preceding the regulatory change to approximately 5.3% in the two months after the change – a decrease of over one third. The reduction in dark trading occurred very quickly and distinctly around the change in regulation.

< Figure 1 here >

Figure 2 shows that prior to the introduction of the minimum price improvement requirements approximately 60% of orders were executed at fractional price increments (10% and 20% of the NBBO spread), which are not allowed under the new rules, and the remaining 40% executed at the midquote. Under the new regulation, fractional price improvement orders disappeared and almost all dark trading now takes place at the midpoint of the NBBO (99.8% of all dark trades). Although after the rule change large dark orders are allowed to execute at the NBBO (at the ‘touch’) after giving priority to lit orders, such dark orders are rare and account for a negligible fraction of dark trades.

< Figure 2 here >

4. Data and metrics

We analyze the constituents of the TSX Composite Index, which comprises approximately 250 of the most actively traded Canadian listed securities, for a period of two months before and two months after the introduction of the minimum price improvement rules (15 August 2012 – 15 December 2012).¹⁶ The four month period is chosen as a compromise

¹⁵ The UMIR defined “better price” simply as a lower price than the best ask price in the case of a purchase and higher price than the best bid price in the case of a sale.

¹⁶ We remove two trading days in which the US markets were closed (during US Thanksgiving and Hurricane Sandy), so that the sample is consistent across all analyses including those that use US data. We obtain data on shares outstanding, stock splits, index constituents and cross-listed securities from the monthly TSX e-Review publications. We restrict our sample to stocks that are included in the TSX Composite Index at both the start and end of our sample period to avoid effects arising from index inclusion/deletion. This results in 246 Canadian stocks. To avoid problems associated with differing

considering the following tradeoff. If the window is too narrow, the analysis will lack statistical power and may not adequately capture changes in market participants' trading behavior.¹⁷ But, if the window is too wide, the analysis around the regulation is likely to be influenced by confounding factors that are unrelated to dark trading. Although we control for confounding factors using a matched sample of US stocks, the longer the window, the less precise the matches and controls.

We combine tick-by-tick data on lit and dark trades from a number of sources. We obtain proprietary data on all dark trades executed on MatchNow, Intraspread, Chi-X, Alpha and TSX¹⁸ directly from the trading venues.¹⁹ The data on dark trades includes the stock ticker, date, time, price and volume. We are unable to obtain data on dark block trades negotiated on Liquidnet/Instinet.²⁰ We also obtain data on all lit trades and the best quotes for all Canadian lit marketplaces (Alpha, Omega, TSX, TMX Select, Pure and Chi-X) from the *Thomson Reuters Tick History* database. Lit trades contain information on the stock ticker, date, time, price and volume, and the quotes comprise the best bid and best ask quote at every point in time for every venue. Timestamps on trades and quotes are recorded to the millisecond. We consolidate the best bid and ask quotes across all lit Canadian venues at every point in time to obtain the NBBO.

To control for changes in market characteristics that are driven by factors other than dark trading, we obtain similar trade-level data (from the *Thomson Reuters Tick History* database) for a matched sample of US stocks, consolidating trades and quotes from all US exchanges. Each Canadian stock is matched to a US stock listed on either NASDAQ or NYSE. Matched stocks are chosen in a manner similar to Huang and Stoll (1996) as those that minimize the sum of squared relative differences in market capitalization and trading volume, X_j , during the two months prior to the price improvement rules (15 August – 15 October 2012):

$$MatchingScore_{CU} = \sum_{j=1}^2 \left(\frac{X_j^C - X_j^U}{(X_j^C + X_j^U)/2} \right)^2. \quad (1)$$

Standard Trading Units and tick sizes we omit stocks with a price less than \$1. This criterion removes five stocks, leaving a final sample of 241 stocks.

¹⁷ As illustrated in Figure 1, the regulation impacted the amount of dark trading effectively overnight, with no evidence of a gradual adjustment process. Therefore the four month window is likely to be sufficiently long to capture changes in market participants' behavior.

¹⁸ Copyright of TSX Inc., all rights reserved. Not to be reproduced or redistributed. TSX Inc. disclaims all representations and warranties with respect to this information, and shall not be liable to any person for any use of this information.

¹⁹ This proprietary data consists only of information that was publically reported to the consolidated tape.

²⁰ IIROC Marketplace Statistics indicate that in the third quarter of 2012 Liquidnet and Instinet combined accounted for only 0.2% of total Canadian dollar volume.

The superscript C indexes each Canadian stock, and the superscript U indexes stocks listed in the US.²¹

All liquidity and informational efficiency metrics are calculated for each stock-day using intraday trade and quote data. Details are in Appendix A. We measure liquidity using quoted, effective and realized spreads, as well as Amihud's (2002) illiquidity metric. Quoted spreads are time-weighted and measure the cost of immediately executing a small round trip trade at the best lit quotes. Effective spreads reflect the cost of a transaction, accounting for the fact that trades can execute at prices within the best lit quotes. Realized spreads reflect the proportion of the transaction cost that is earned by the liquidity provider after removing the adverse selection cost. Amihud's (2002) illiquidity metric is a measure of price impact scaled by traded dollar volume, and therefore captures depth and resiliency.

Following the empirical literature, we use four high-frequency measures of the informational efficiency of prices: absolute autocorrelations of midquote returns, midquote variance ratios, high-frequency standard deviations, and measures of short-term return predictability using lagged market returns.²² The informational efficiency metrics, to varying degrees, measure inefficiency with respect to transitory deviations in price (possibly caused by order imbalances and imperfect liquidity), as well as inefficiency around permanent changes in prices (possibly caused by delay in impounding new information and under/over reactions to news). Thus, the informational inefficiency measures are likely to be impacted by liquidity, but also capture an informational component that is orthogonal to liquidity. In support of this conjecture, Rösch et al. (2013) provide evidence that informational efficiency metrics measured at intraday horizons are highly correlated with low-frequency measures of informational efficiency, and are different from liquidity measures. All four informational inefficiency measures are scaled so that they range from 0 (indicating high levels of efficiency) to 100 (indicating low levels of efficiency).

5. Empirical analysis

5.1 Descriptive statistics

Table 2 reports descriptive statistics on trading activity before and after the minimum price improvement regulation came into effect. Consistent with Figure 1, the level of dark

²¹ The median differences between the Canadian and matched US stocks' market capitalization and average traded dollar volume are less than 15%, suggesting the matching is relatively precise.

²² Autocorrelations are used in Hendershott and Jones (2005) and Anderson et al. (2013), variance ratios are popularized by Lo and MacKinlay (1988), high frequency volatility is used by O'Hara and Ye (2011), and return predictability using lagged market returns follows from Hou and Moskowitz (2005).

trading is considerably lower after the minimum price improvement rules come into effect. The mean (median) percentage of daily dollar volume executed in the dark falls from 9.01% (7.58%) to 5.93% (4.05%) after the regulation. Lit and dark trades tend to have a similar size (mean of approximately \$6,200 and median of approximately \$4,600) and their size does not change noticeably after introducing the minimum price improvement rules. While the total amount of dark trading is reduced from an average of \$492 million per day to \$321 million per day, this is somewhat offset by a small increase in lit trading. Average total daily traded dollar volume remains unchanged at approximately \$6.2 billion per day.

< Table 2 here >

To get a sense of the variation in dark trading, Figure 3 presents the pooled sample histogram of stock-day level dark trading as a fraction of total stock-day dollar volume, $Dark_{it}$. Approximately 5% of stock-days have no dark trading at all. Around 40% of stock-days have between 1% and 5% of their dollar volume executed in the dark. There are very few stock days with greater than 20% dark trading, and only 28 stock-days have dark trading in excess of 50% of total dollar volume.

< Figure 3 here >

Table 3 reports descriptive statistics on the stock-day market quality metrics and control variables. Before the regulation, quoted spreads have a mean and median of 12.69 bps and 9.66 bps; effective spreads are slightly lower with a mean of 10.44 bps due to some trades executing within the spread; and realized spreads are even smaller with a mean of 2.29 bps due to the fact that trades tend to have positive price impact on average. All three spreads increase after the regulation, by between 0.56 and 1.21 bps on average, and the differences are statistically distinguishable from zero, using standard errors clustered by stock and date. Similarly, $ILLIQ_{it}$ increases a statistically significant amount after the regulation, as do each of the four informational inefficiency metrics. The variable $Constrained_{it}$ indicates that quoted spreads tend to be constrained to the minimum of one tick approximately 59% of the time for an average stock. The median company has a market value of approximately \$2.3 billion.

< Table 3 here >

5.2 Instrumental variables regressions for aggregate dark trading

One of the main challenges in empirically studying the impact of dark trading on market quality is the endogeneity of dark trading with respect to market conditions. For example, dark trading tends to increase when spreads are constrained to the minimum tick size because dark trades are allowed to occur within the spread at sub-penny price increments (Kwan et al., 2014). Buti et al. (2011) find that dark pool activity is higher when limit order depth is high, spreads are narrow and tick sizes are large. They argue that the conditional nature of the decision to execute in the dark results in an endogeneity issue between market quality and dark trading.

To overcome the endogeneity issue we use the introduction of the minimum price improvement rule as an instrumental variable (IV) for dark trading in a two-stage least squares (2SLS) framework, controlling for confounding effects with a set of matched US stocks. The first stage is a regression of the level of dark trading on the instrument (a dummy variable for the minimum price improvement rule), and a set of control variables:

$$Dark_{it} = \alpha_i + \beta D_t^{PostPeriod} + \sum_{j=1}^6 \gamma_j Control_{j,it} + \varepsilon_{it} , \quad (2)$$

where $Dark_{it}$ is the fraction of dollar volume in stock i on day t that is traded in the dark, α_i is a set of stock fixed effects, $D_t^{PostPeriod} = 1$ after the rule change and 0 before. $Control_{j,it}$ comprises the following six control variables. $Time_t$ takes the value 0 on the first day in the sample and increments by 1 every subsequent day. It removes general time-series trends in dark trading and in market quality. $US Mean_t$ is the daily mean of the corresponding market quality metric (the second stage dependent variable) for the matched US stocks. Consequently, we estimate a different first-stage model for each market quality metric. The reason for including $US Mean_t$ is that in the second stage it removes variation in market quality that is common to the US and Canada and is driven by factors other than dark trading, thereby giving a difference-in-differences estimate. The other control variables are $\$Volume_{it}$ (the natural logarithm of traded dollar volume), $Volatility_{it}$ (the stock-day's high-low price range divided by the time-weighted midquote), $Price_{it}$ (the time-weighted midquote), and $Constrained_{it}$ (the percentage of the trading day for which the stock's NBBO spread is constrained at one tick).²³ In our main results we estimate the first stage equation (2) on the pooled sample of stock-days. Robustness tests indicate that estimating the first stage separately for each stock produces similar results in the second stage, as does including/omitting stock fixed effects in the first stage.

²³ We estimate the first stage with and without the variable $Constrained_{it}$ because in the second stage we omit the variable $Constrained_{it}$ when the dependent variable is a spread.

Table 4 reports the results from the first stage of the 2SLS. Although the first stage is different for each of the market quality metrics (due to $US Mean_t$), the results are fairly similar across the market quality metrics. Therefore, Table 4 reports results from only two market quality metrics (in columns (1) and (2) $US Mean_t$ is the US daily mean quoted spread, whilst in (3) and (4) it is the US daily mean of Amihud’s $ILLIQ_{it}$). The estimates indicate that the minimum price improvement regulation is associated with a decline in the average level of dark trading by approximately 3.1 percentage points, holding other factors fixed, and the decline is highly statistically significant. Problems associated with “weak” instruments arise when first-stage F-statistics for the instruments are close to one (Bound et al., 1995, p. 446). The F-statistics in our first stage regression with and without stock fixed effects are well above this level (219-271), suggesting the instrument is strong.

< Table 4 here >

The second stage regressions estimate the impact of dark trading on a number of liquidity and informational efficiency measures:

$$y_{it} = \alpha_i + \beta \widehat{Dark}_{it} + \sum_{j=1}^6 \gamma_j Control_{j,it} + \varepsilon_{it} , \quad (3)$$

where y_{it} is a market quality metric for stock i on day t , α_i is a set of stock fixed effects, and \widehat{Dark}_{it} is the fitted level of dark trading. $Control_{j,it}$ includes the same control variables as in the first stage, including the US daily mean of the corresponding market quality metric, $US Mean_t$. With the inclusion of $US Mean_t$, our second-stage model is very similar to a difference-in-differences model, but with two distinct advantages.²⁴ First, by having a free coefficient on the variable $US Mean_t$, it does not impose a one-to-one correspondence between changes in US market quality and changes in Canadian market quality. Instead, the degree of co-movement in market quality in the two markets is estimated from the data. Therefore, the model is better able to account for the fact that the scale of the market quality variables may differ between the US and Canada. Second, by summarizing the control market (US) with a single

²⁴ More specifically, ignoring the control variables for simplicity, a standard difference-in-differences model, $y_{i,t} = \mu D_i^{CANADA} + \theta D_t^{POST} + \delta D_i^{CANADA} D_t^{POST} + \varepsilon_{it}$ gives an estimate of $\hat{\delta} = (\overline{y_{CANADA,POST}} - \overline{y_{CANADA,PRE}}) - (\overline{y_{US,POST}} - \overline{y_{US,PRE}})$. By comparison, our second stage model produces an estimator that is essentially equivalent to $\hat{\beta} = (\overline{y_{CANADA,POST}} - \overline{y_{CANADA,PRE}}) - \hat{\gamma}(\overline{y_{US,POST}} - \overline{y_{US,PRE}})$, where $\hat{\gamma}$ is the coefficient of $US Mean_t$ and measures the extent to which the market quality characteristic $y_{i,t}$ tends to co-move in US and Canadian stocks.

time-series (rather than a collection of control stocks), the model avoids inflating the number of observations and thus provides more conservative standard errors.

Table 5 reports second stage estimates of the impact of dark trading on liquidity. The results indicate that dark trading has a negative and statistically significant effect on all of the spread measures as well as Amihud's illiquidity metric, suggesting that aggregate dark trading in Canada *benefits* liquidity. A small increase in dark trading by 5% of total dollar volume is expected to decrease average quoted spreads by approximately 0.29 bps ($0.05 \times (-5.71)$), decrease effective spreads by 1.12 bps, and decrease realized spreads by 1.31 bps. These decreases for just a small change in dark trading are economically meaningful compared to the means of quoted, effective and realized spreads: 12.98 bps, 10.88 bps, and 2.91 bps, respectively.

An alternative way to interpret the magnitude of the effects is in terms of pooled standard deviations. A one standard deviation increase in dark trading (6.4% of total dollar volume) is expected to decrease quoted spreads by 0.37 bps or 0.03 standard deviations ($(0.064 \times (-5.71)) / 11.8$), decrease effective spreads by 1.42 bps or 0.13 standard deviations, decrease realized spreads by 1.67 bps or 0.13 standard deviations, and decrease price impacts (Amihud's $ILLIQ_{it}$) by 0.12 standard deviations. Therefore, while there is variation in the magnitudes across the different liquidity measures, the results suggest that aggregate dark trading in our sample has economically meaningful benefits to liquidity.

< Table 5 here >

The findings in Table 5 support our third hypothesis about the level and composition of dark trading. As Figure 2 illustrates, prior to the introduction of the minimum price improvement regulation, the majority of dark trading in Canada (approximately 60%) was dark trading with fractional price improvement, or two-sided dark trading. The literature suggests that two-sided dark trading benefits liquidity (Hypothesis 2) and therefore a reduction in the aggregate level of Canadian dark trading is expected to decrease liquidity. The minimum price improvement regulation not only changed the aggregate level of dark trading but also the composition, significantly increasing the level of one-sided (midpoint) dark trading and decreasing two-sided dark trading. Because one-sided (two-sided) dark trading is expected to harm (benefit) liquidity, Hypothesis 3 suggests that this change in the composition is expected to reinforce the effect of the decrease in the aggregate level. Therefore, the results support the notion that the composition of

dark trading types within the aggregate affects market quality, and that dark trading can benefit liquidity when a large proportion of it is two-sided, resembling a dark limit order book.

Coefficients on the time trend suggest that liquidity tends to improve through time, not counting the effects of the regulation (quoted spreads, effective spreads and $ILLIQ_{it}$ have a negative time trend during our sample period, after controlling for the regulation). The coefficients on $US Mean_t$ are all positive and statistically significant (with the exception of realized spreads), indicating that liquidity in Canadian stocks tends to co-move with liquidity in US stocks. Most coefficients on the control variables are consistent with our expectations – liquidity tends to be higher for days with greater volume and lower volatility. The adjusted R^2 of the regressions, which do not include the variation explained by the stock fixed effects range between 3% and 24%, suggesting that there are many factors beyond the variables included in our model that influence liquidity.

< Table 6 here >

Turning to the informational efficiency proxies, Table 6 reports second stage regression estimates of the impact of dark trading. The results suggest that dark trading has a negative effect on all informational inefficiency metrics: absolute autocorrelations, variance ratios, high-frequency volatility (although not statistically significant), and delay in reflecting market-wide information. These results suggest that, similar to its effects on liquidity, aggregate dark trading in Canada *benefits* informational efficiency. Because the units of the informational efficiency proxies do not have a natural interpretation we examine the magnitude of the effects in terms of standard deviations. A one standard deviation increase in dark trading (6.4% of total dollar volume) is expected to decrease absolute midquote return autocorrelations by 0.43 standard deviations ($0.064 \times (-85) / 12.7$), decrease variance ratios by 0.63 standard deviations, and decrease the delay in impounding market-wide information by 0.10 standard deviations, after controlling for other market characteristics and stock fixed effects. The impact on autocorrelations and variance ratios, measured in standard deviations, is stronger than the impacts on the liquidity measures. While there is variation in the magnitudes across the different informational efficiency measures, the results suggest that aggregate dark trading in our sample has economically meaningful benefits for informational efficiency.

The beneficial impact of aggregate dark trading on informational efficiency is consistent with the close relationship between liquidity and informational efficiency (e.g., Chordia et al.,

2008). Similar to the liquidity result, the positive effect of aggregate dark trading on informational efficiency is likely to be driven by two-sided dark trading. The model analyzed by Boulatov and George (2013) suggests that the ability to submit dark limit orders not only increases liquidity provision by informed traders but also the aggressiveness with which they trade, which in turn improves informational efficiency.

We test the robustness of our results to a variety of alternative specifications of the IV regressions and different subsamples. To concisely summarize the results of these tests Table 7 reports the t-statistics for the coefficient on the key independent variable, \widehat{Dark}_{it} , in the second-stage regressions. The rows of Table 7 correspond to different dependent variables and the columns correspond to different specifications and subsamples. Specification (1) is the base case specification (the specification reported in Tables 5 and 6, corresponding to equation (3)), which includes all of the control variables and stock fixed effects. Specification (2) is identical to specification (1) except that both the first and second stage has been run using only cross-listed securities. Specification (3) is identical to specification (1) except that the first stage IV regression is estimated on each stock separately allowing for heterogeneity across stocks in the way in which dark trading is affected by the minimum price improvement rules. Specification (4) is identical to specification (1) except that it omits two weeks either side of the introduction of the minimum price improvement rules to allow for transitory effects and adjustment in trading behavior. Specifications (5) and (6) are estimated on the largest 121 stocks and smallest 120 stocks, respectively, using the same variables as in specification (1). Specification (7) is identical to specification (1) except that the two stages are estimated simultaneously using maximum likelihood.

< Table 7 here >

The results from the different specifications and subsamples are consistent with our base specification. Dark trading is associated with improved liquidity and informational efficiency (decreased illiquidity and informational inefficiency) across all proxies and specifications with few exceptions. Our results are robust to allowing for heterogeneity in the impact of the instrumental variable, allowing for transitory effects around the rule change, and are similar for the largest and the smallest stocks in our sample. The latter result is consistent with Comerton-Forde and Putniņš (2013) who find that the effects of dark trading on price discovery are similar in both large and small stocks.

Finally, we examine whether the effects of dark trading are different in stocks that are constrained by the tick size (trade at a spread of one tick), compared to those that are not. For stocks that *are* constrained by the tick size, dark trades are more likely to occur at the midquote than in stocks that are *not* constrained (e.g., after the regulation, dark trades in tick constrained stocks can only take place at the midquote). Under the hypothesis that two-sided dark trading is beneficial to liquidity, we would thus expect that dark trading in stocks that are not constrained by the tick size is more beneficial to liquidity than dark trading in stocks that are constrained. The last two columns of Table 7 report estimates for the 50% of stocks that are most/least frequently constrained by the tick size (using the average value of $Constrained_{it}$ in the two months prior to the minimum price improvement regulation). Consistent with our hypothesis, the beneficial effects of dark trading are more pronounced (higher statistical significance and consistency) for stocks that are least often constrained by the tick size. The magnitudes of the effects, although not reported, lead to a similar conclusion.

5.3 Instrumental variables regressions for different types of dark trading

To provide a more formal analysis of whether different types of dark trading have different effects on market quality, we disaggregate dark trading into two-sided dark trading with fractional price improvement ($Fractional_{it}$) and one-sided midpoint dark trading ($Midpoint_{it}$). Both types are measured as a percentage of total dollar volume. If we were to simply follow the same standard approach as in equations (2-3), the model would not be identified because we would have *one* instrument and *two* endogenous variables ($Fractional_{it}$ and $Midpoint_{it}$).

Our solution to the issue of identification is to exploit the cross-sectional heterogeneity in the way in which the regulation impacts the two different types of dark trading, $Fractional_{it}$ and $Midpoint_{it}$. The intuition behind this approach is as follows. For different stocks, the regulation has different effects on $Fractional_{it}$ and $Midpoint_{it}$. For example, some stocks have a large reduction in $Fractional_{it}$ around the regulation, others have no change in $Fractional_{it}$. Similarly, some stocks have a large reduction in $Midpoint_{it}$, while others have no change. As long as the stocks that have a large reduction in $Fractional_{it}$ are not exactly the same stocks that have a large reduction in $Midpoint_{it}$ (i.e., the cross-sectional correlation of the regulation's impact on $Fractional_{it}$ and $Midpoint_{it}$ is less than perfect), then the stocks that have a large reduction in $Fractional_{it}$ will allow the effect of $Fractional_{it}$ on market quality to be identified. Similarly, the stocks that have a large reduction in $Midpoint_{it}$ will allow the effects of $Midpoint_{it}$ on market quality to be identified.

More formally, for each stock separately, we estimate first stage models for $Fractional_{it}$ and first stage models for $Midpoint_{it}$. This is equivalent to:

$$Fractional_{it} = \sum_{k=1}^{241} D_k \left(\alpha_k + \beta_k D_t^{PostPeriod} + \sum_{j=1}^6 \gamma_{jk} Control_{j,it} \right) + \varepsilon_{it} \quad (4)$$

$$Midpoint_{it} = \sum_{k=1}^{241} D_k \left(\mu_k + \delta_k D_t^{PostPeriod} + \sum_{j=1}^6 \theta_{jk} Control_{j,it} \right) + \vartheta_{it} \quad (5)$$

where D_k is a dummy variable for each of the $k = 1, 2, \dots, 241$ Canadian stocks in our sample. The model effectively has 241 instruments to identify two endogenous variables and is thus over-identified.²⁵

In the second-stage regressions we include the fitted values of fractional and midpoint dark trading together to estimate the independent impact of each of these volume types:

$$y_{it} = \alpha_i + \beta_1 \widehat{Fractional}_{it} + \beta_2 \widehat{Midpoint}_{it} + \sum_{j=1}^6 \gamma_j Control_{j,it} + \varepsilon_{it}. \quad (6)$$

Table 8 reports the second stage estimates. Consistent with our hypotheses, fractional and midpoint dark trading have different effects on market quality. Fractional dark trading is associated with strong improvements in all of the liquidity and informational efficiency metrics. All of the improvements are statistically significant except for the effect on $Delay_{it}$. In contrast, the effect of midpoint dark trading is somewhat mixed for the different market quality measures and is statistically indistinguishable from zero for half of the market quality metrics. For example, the results suggest that midpoint dark trading is associated with higher absolute autocorrelations and variance ratios (both statistically significant). The results also suggest that midpoint dark trading reduces effective spreads; however, this is a somewhat mechanical effect because midpoint dark trades by definition occur at zero effective spreads. Therefore, the results support the hypothesis that different types of dark trading have different effects on the market. Two-sided dark trading has clear benefits for market quality. In contrast, the effects of midpoint dark trading are weaker with no conclusive evidence of positive or negative effects.

< Table 8 here >

Theory provides some interpretations of the results. Models that analyze the effects of pre-trade transparency in limit order markets suggest the positive effects of two-sided dark trading stem from increased willingness among informed traders to supply liquidity (e.g., Boulatov and George, 2013; Rindi, 2008). Less pre-trade transparency encourages informed

²⁵ The F-statistic for the joint hypothesis that the 241 instruments are not significant in the first stage is 6.8 in regression equation (4) and 9.7 in regression equation (5), indicating the instruments are sufficiently strong.

traders to provide liquidity because they can profit from doing so without revealing as much of their private information. More aggressive competition among informed traders also helps facilitate price discovery. Models of fragmentation suggest that the positive effects of two-sided dark trading stem from increasing the number of liquidity providers (e.g., Biais et al., 2000), encouraging liquidity provision through ‘queue jumping’ (e.g., Foucault and Menkveld, 2008), and allowing liquidity providers to compete on a finer pricing grid (e.g., Biais et al., 2010; Buti et al., 2014). Finally, models of competition between midpoint dark crossing systems and lit markets show that midpoint dark trading can increase adverse selection and/or inventory holding risks, thereby offsetting any positive effects. For example, in Zhu (2014) informed traders are more likely than uninformed traders to send their order to the lit market rather than to the dark, and in Hendershott and Mendelson (2000) traders will sometimes send their orders to the lit market if they fail to execute in the dark, using the lit market as a ‘market of last resort’.

Our finding that different types of dark trading have different effects on the market suggests that some of the variation in the results of empirical studies that do not disaggregate dark trading may stem from differences across countries in the composition of dark trading types. For example, Buti et al. (2011) find that dark trading is associated with narrower spreads in the US, whereas Degryse et al. (2014) conclude that dark fragmentation harms liquidity. Our sample also differs from other empirical studies in that it does not include trades from systematic dark ‘internalizers’, which account for a large proportion of dark volume in many other countries.

Two-sided dark trading constitutes approximately two thirds of all dark trading in Canada before the minimum price improvement regulation. The new rules not only decreased the level of dark trading but also changed its composition, replacing much of the two-sided dark trading with midpoint dark trading. Our finding that two-sided dark trading benefits market quality suggests that these changes should lead to a deterioration of market quality. The descriptive statistics support this conjecture, suggesting spreads are wider and informational *inefficiency* metrics are higher after the regulation. The substitution from two-sided to midpoint dark trading following the minimum price improvement regulation occurred largely because many Canadian stocks are constrained by the tick size, i.e., the spread is often one tick, forcing dark trades to execute at the midpoint in order to comply with the new rules. Our results suggest that a way of improving the effectiveness of minimum price improvement requirements for dark trades is to ensure tick sizes do not constrain the lit spread. That way, dark trades can offer price improvement while maintaining a two-sided dark market.

We again subject our analysis to a range of robustness tests, which we summarize in Table 9. The results are qualitatively similar when using only the subset of cross-listed stocks, omitting two weeks either side of the minimum price improvement regulation, and running the analysis separately on large stocks, small stocks, and the most/least tick size constrained stocks. In particular, the positive effect of dark trading with fractional price improvement on liquidity and informational efficiency remains strong, whereas the effect of midpoint dark trading is weak and somewhat mixed.

< Table 9 here >

6. Conclusions

We use a unique natural experiment, the introduction of minimum price improvement regulation in Canada, to examine the effects of dark trading. We disaggregate dark trading into two types: dark trading at the midpoint of the lit NBBO ('one-sided' dark trading) and dark trading at prices that are either side of the midpoint ('two-sided' dark trading). This partition is important both theoretically and empirically. One-sided and two-sided dark trading differ in execution probability, the feasibility of dark market making strategies, and the amount of information that can be inferred from resting dark orders about trading intentions.

We find that two-sided dark trading is beneficial to both liquidity and informational efficiency. It tends to lower quoted, effective and realized spreads, reduces price impact measures of illiquidity, makes prices closer to the random walk that would be expected under informational efficiency. The magnitudes of the effects are economically meaningful.

In contrast, we do not find consistent evidence that midpoint dark trading has a significant effect on market quality. Aggregating across the two types of dark trading, our results suggest that aggregate dark trading is more likely to benefit market quality the greater the proportion of two-sided dark trading. Furthermore, changes in the composition of dark trading can impact market quality even if the aggregate level remains unchanged. An increase in two-sided dark trading relative to one-sided dark trading is likely to benefit market quality. Our results are robust to a range of alternative specifications, fixed effects, sub-period tests, as well as controlling for time trends and confounding factors using a sample of matched US stocks. The effects of dark trading are similar for the largest and the smallest stocks in our sample.

Our findings have two caveats. First, the levels of dark trading in Canada are lower than in some other markets, in particular the US. It is possible that the effects of dark trading on market quality are non-linear in the level of dark trading, and that the positive effects of two-

sided dark trading dissipate once they reach some ‘tipping’ point after which the marginal effect of two-sided dark trading on market quality becomes negative. Second, our analysis considers trading in dark pools and dark order types on lit markets, but not systematic internalization by brokers in off-market ‘internalizers’. Internalization is associated with a range of different issues and therefore should be analyzed as a separate type of dark trading.

Our findings are consistent with theoretical studies. For example, Boulatov and George (2013) find that less pre-trade transparency in limit order markets encourages informed traders to act as liquidity suppliers because they can profit from liquidity provision without revealing much of their private information. Our results suggest that strong competition among informed traders in providing dark liquidity has positive spillover effects on the lit market, where liquidity providers are forced to narrow spreads to compete with dark liquidity. Consistent with our results, Boulatov and George (2013) also show that aggressive liquidity provision by informed traders in the dark improves price discovery. Our results are also consistent with the notion that fragmentation across lit and two-sided dark trading venues can benefit liquidity by increasing the number of liquidity providers (e.g., Bias et al., 2000), encouraging liquidity provision through ‘queue jumping’ (e.g., Foucault and Menkveld, 2008), and allowing liquidity providers to compete on a finer pricing grid (e.g., Biais et al., 2010; Buti et al., 2014). Lastly, Zhu (2014) and Hendershott and Mendelson (2000) conjecture that midpoint dark trading can increase inventory and adverse selection risks. While we do not find that midpoint dark trading harms liquidity, our results indicate that midpoint dark trading does not benefit liquidity like two-sided dark trading, consistent with an opposing increase in inventory and adverse selection risks.

This paper has a number of policy implications, in particular given the current regulatory interest in dark trading. At the broadest level, the results point to the fact that dark trading should not be treated as a homogenous group; it is important to distinguish between different types of dark trading when developing policy. The effects of aggregate dark trading depend on the composition of dark trading types within the aggregate. Our results suggest that the larger the proportion of two-sided dark trading in the aggregate the more likely the aggregate dark trading benefits rather than harms market quality. A harmful level of aggregate dark trading in one country may not be harmful in another due to differences in the composition of dark trading types. It follows that in designing regulation it is important to consider not only the effect on the level of dark trading but also on the composition of dark trading types. For example, minimum price improvement regulation in Canada not only decreased the level of dark trading but also caused a shift from two-sided dark trading to dark trading at the midpoint of the NBBO.

Finally, our results have implications for tick size regulation. Dark trading is more active in stocks for which their spread is constrained by the tick size (Kwan et al., 2014). Our results suggest that the use of dark trading as a way of obtaining a finer price grid may benefit market quality as long as the price grid in the dark allows dark liquidity to concurrently exist on both the buy- and sell-sides of the market. Minimum price improvement rules can force dark trades to occur at the midpoint in stocks that are constrained by the tick size and consequently such rules can have unintended negative effects on market quality. Our results suggest a way of improving the effectiveness of minimum price improvement requirements: ensure tick sizes do not constrain the lit spread. If the price grid is sufficiently fine, dark trades can offer price improvement while maintaining a two-sided dark market.

Appendix A: Liquidity and informational efficiency metrics

When calculating market quality metrics we use the regular market hours of 9:30am – 4:00pm, less the first and last 15 minutes to exclude the impacts of the opening auction and market on close facility. However we include the first and last 15 minutes as well as the opening and closing auctions in the summations of daily volume.

A.1 Liquidity measures

We measure liquidity using quoted, effective and realized spreads, as well as Amihud’s (2002) illiquidity metric. All liquidity metrics are calculated for each stock-day. We measure the quoted NBBO spread in basis points relative to the prevailing midquote, $m = (Ask + Bid)/2$,

$$QuotedSpread = [(Ask - Bid)/m] , \quad (A1)$$

and take the time-weighted average between 9:45am and 3:45pm for each stock-day. For a trade that occurs at time τ we measure its effective spread and five-minute realized spread (both in basis points relative to the prevailing midquote) as

$$EffectiveSpread = 2q[(p_\tau - m_\tau)/m_\tau] \quad (A2)$$

$$RealizedSpread = 2q[(p_\tau - m_{\tau+5})/m_\tau] , \quad (A3)$$

where p_τ is the transaction price, m_τ is the midpoint of the NBBO prevailing at the time of the trade, $m_{\tau+5}$ is the midpoint of the NBBO five minutes after the trade, and q indicates the direction of the trade (+1 for buyer-initiated trades and -1 for seller initiated trades). Buyer- and seller-initiated trades are identified by comparing the prevailing NBBO to the transaction price using the Lee and Ready (1991) algorithm. For each stock-day we take the volume-weighted average effective and realized spread across all lit trades.

Amihud’s (2002) illiquidity metric is a measure of price impact scaled by traded dollar volume. For each stock-day we compute the average ratio of hourly absolute midquote returns to hourly dollar volume:

$$ILLIQ_{it} = \log \left[1 + \frac{10^5}{H} \sum_{h=1}^H \frac{|r_{it,h}|}{\$Volume_{it,h}} \right] , \quad (A4)$$

where $r_{it,h}$ and $\$Volume_{it,h}$ are the midquote return and traded dollar volume, respectively, for stock i during hour h of day t .²⁶ To reduce the influence of outliers we winsorize the liquidity metrics at the 1% level for each stock and each date.

²⁶ As indicated in equation (4) the $ILLIQ_{it}$ metric is log transformed to reduce the impact of outliers, consistent with Karolyi et al. (2012). If there is no volume traded in a given hour, the denominator in the $ILLIQ_{it}$ metric is zero. Rather than generating a missing observation we replace such instances with the stock’s 99th percentile value of valid $|r_{it,h}|/\$Volume_{it,h}$ observations.

A.2 Informational efficiency measures

We use four high-frequency measures of the informational efficiency of prices: autocorrelations, variance ratios, high-frequency standard deviations, and measures of short-term return predictability using lagged market returns.

Positive or negative midquote return autocorrelations indicate that quotes deviate from a stochastic random walk and exhibit short-term return predictability. Such predictability is mainly driven by partial price adjustment to information, including under- and over-reaction (see Anderson et al., 2013), which is inconsistent with an informationally efficient market. We calculate the absolute value of first-order midquote return autocorrelations for each stock-day, at three intraday frequencies, $k \in \{10\text{sec.}, 30\text{sec.}, 60\text{sec.}\}$, similar to Hendershott and Jones (2005):

$$\text{Autocorrelation}_k = |\text{Corr}(r_{k,\tau}, r_{k,\tau-1})|, \quad (\text{A5})$$

where $r_{k,\tau}$ is the τ^{th} midquote return of length k in a given stock-day. Taking the absolute value of the autocorrelation yields a measure of informational efficiency that measures both the under- and over-reaction of returns to information, with larger values indicating greater inefficiency. We combine the absolute autocorrelations of the three frequencies by calculating their first principle component. The combined measure, $\text{Autocorrelation}_{it}$, is scaled so that it ranges from 0 (highly efficient) to 100 (highly inefficient).

For stock prices that follow a random walk, the variance of returns is a linear function of the return measurement frequency, i.e., $\sigma_{k\text{-PeriodReturn}}^2$ is k times larger than $\sigma_{1\text{-PeriodReturn}}^2$. The variance ratio makes use of this property to measure inefficiency as a price series' deviation from the characteristics that would be expected under a random walk (e.g., Lo and MacKinlay, 1988). We construct three variance ratios each stock-day, utilizing different intra-day frequencies:

$$\text{VarianceRatio}_{kl} = \left| \frac{\sigma_{kl}^2}{k\sigma_l^2} - 1 \right|, \quad (\text{A6})$$

where σ_l^2 and σ_{kl}^2 are the variances of l -second and kl -second midquote returns for a given stock-day. We use the (l,kl) combinations: (1sec., 10sec.), (10sec., 60sec.), (1min., 5min.). We combine the three variance ratios by calculating their first principle component. The combined variance ratio, $\text{VarianceRatio}_{it}$, is scaled so that it ranges from 0 (indicating high levels of efficiency) to 100 (indicating low levels of efficiency).

For each stock-day we also combine the intra-day midquote standard deviations calculated at 10, 30 and 60 second frequencies by taking their first principal component. This produces a single measure of high-frequency volatility, HFVolatility_{it} , which is a proxy for noise and temporary deviations of prices from their equilibrium values due to trading frictions. In

the regressions we control for volatility of the fundamental value using a lower frequency measure of realized variance.

The final measure of informational efficiency is an intraday adaptation of the Hou and Moskowitz (2005) *Delay* metric. *Delay* measures short-term return predictability by the extent to which lagged market returns predict a stock's midquote returns. For each stock-day it we estimate a regression of intraday 1-minute midquote returns for the stock, $r_{it,\tau}$, on the TSX60 market index return, $r_{mt,\tau}$, and ten lags:

$$r_{it,\tau} = \alpha_{it} + \beta_{it}r_{mt,\tau} + \sum_{k=1}^{10} \delta_{it,k}r_{mt,\tau-k} + \varepsilon_{it,\tau} . \quad (\text{A7})$$

We save the R^2 from this unconstrained regression, $R_{Unconstrained,it}^2$, re-estimate the regression constraining the coefficients on all lagged market returns to zero (i.e., $\delta_{it,k} = 0, \forall k$), and save the R^2 from the constrained regression, $R_{Constrained,it}^2$. *Delay* is calculated from the ratio of the constrained and unconstrained regression R^2 s:

$$Delay_{it} = 100 \left(1 - \frac{R_{Constrained,it}^2}{R_{Unconstrained,it}^2} \right) . \quad (\text{A8})$$

$Delay_{it}$ takes values between 0 and 100 and describes the amount of variation in a stock's intraday returns that is explained by lagged market returns. The more explanatory power the lagged returns have, the higher is $R_{Unconstrained,it}^2$ and the closer is $Delay_{it}$ to 100, implying a delayed incorporation of market-wide information into the stock's price, and lower overall informational efficiency.

References

- Anderson, R.M., K.S. Eom, S.B. Hahn, and J.-H. Park, 2013, Autocorrelation and partial price adjustment, *Journal of Empirical Finance* 24, 78-93.
- Amihud, Y., 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.
- Biais, B., C. Bisiere, and C. Spatt, 2010, Imperfect competition in financial markets: An empirical study of Island and Nasdaq, *Management Science* 56, 2237-2250.
- Biais, B., D. Martimort, and J. Rochet, 2000, Competing mechanisms in a common value environment, *Econometrica* 68, 799-837.
- Boulatov, A., and T.J. George, 2013, Hidden and displayed liquidity in securities markets with informed liquidity providers, *Review of Financial Studies* 26, 2095-2137.
- Bound, J., D.A. Jaeger, and R.M. Baker, 1995, Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak, *Journal of the American Statistical Association* 90, 443-450.
- Buti, S., and B. Rindi, 2013, Undisclosed orders and optimal submission strategies in a limit order market, *Journal of Financial Economics* (forthcoming).
- Buti, S., B. Rindi, and I.M. Werner, 2011, Diving into dark pools, *Unpublished manuscript*.
- Buti, S., F. Consonni, B. Rindi, Y. Wen, and I.M. Werner, 2014, Sub-penny and queue jumping, *Unpublished manuscript*.
- Chordia, T., R. Roll, and A. Subrahmanyam, 2008, Liquidity and market efficiency, *Journal of Financial Economics* 87, 249-268.
- Comerton-Forde, C., and T.J. Putniņš, 2013, Dark trading and price discovery, *Unpublished manuscript*.
- Degryse, H., F. de Jong, and V. van Kervel, 2014, The impact of dark trading and visible fragmentation on market quality, *Review of Finance* (forthcoming).
- Easley, D., N.M. Kiefer, and M. O'Hara, 1996, Cream-skimming or profit-sharing? The curious role of purchased order flow, *Journal of Finance* 51, 811-833.
- Foley, S., K. Malinova, and A. Park, 2012, Dark trading on public exchanges, *Unpublished manuscript*.
- Foucault, T., and A.J. Menkveld, 2008, Competition for order flow and smart order routing systems, *Journal of Finance* 63, 19–58.
- Hendershott, T., and C.M. Jones, 2005, Island goes dark: Transparency, fragmentation, and regulation, *Review of Financial Studies* 18, 743–793.

- Hendershott, T., and H. Mendelson, 2000, Crossing networks and dealer markets: Competition and performance, *Journal of Finance* 55, 2071- 2115.
- Hou, K., and T.J. Moskowitz, 2005, Market frictions, price delay, and the cross-section of expected returns, *Review of Financial Studies* 18, 981–1020.
- Huang, R.D., and H.R. Stoll, 1996, Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE, *Journal of Financial Economics* 41, 313–357.
- Karolyi, G.A., K.-H. Lee, and M.A. van Dijk, 2012, Understanding commonality in liquidity around the world, *Journal of Financial Economics* 105, 82-112.
- Kwan, A., R. Masouliis, and T. McInish, 2014, Trading rules, competition for order flow and market fragmentation, *Journal of Financial Economics* (forthcoming).
- Kyle, A.S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315-1335.
- Kyle, A.S., 1989, Informed speculation with imperfect competition, *Review of Economic Studies* 56, 317-356.
- Lee, C.M.C., and M.J. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733–746.
- Lo, A., and C. MacKinlay, 1988, Stock market prices do not follow random walks: Evidence from a simple specification test, *Review of Financial Studies* 1, 41–66.
- Nimalendran, M., and S. Ray, 2014, Informational linkages between dark and lit trading venues, *Journal of Financial Markets* 17, 230-261.
- Ready, M., 2014, Determinants of volume in dark pool crossing networks, *Unpublished Manuscript*.
- Rindi, B., 2008, Informed traders as liquidity providers: Anonymity, liquidity and price formation, *Review of Finance* 12, 497-532.
- Rösch, D.M., A. Subrahmanyam, and M.A. van Dijk, 2013, An empirical analysis of co-movements in high- and low-frequency measures of market efficiency, Working paper.
- Ye, M., 2011, A glimpse into the dark: Price formation, transaction cost and market share of the crossing network, *Unpublished manuscript*.
- Zhu, H., 2014, Do dark pools harm price discovery?, *Review of Financial Studies* (forthcoming).

Table 1
Summary of all trading venues in Canada

This table provides an overview of all lit and dark trading venues in Canada. The types of orders allowed include lit-only, dark-only or both lit and dark. The approximate market share of total traded dollar volume (including dark and lit trades on all listed venues) is reported for the period 15 August 2012 – 15 December 2012. The market share of Liquidnet / Instinet is obtained from IIROC statistics, whereas the other market shares are calculated from our data.

Venue	Lit / Dark	Market share
TSX	Both	61.3%
Chi-X	Both	12.9%
Alpha (Lit)	Both (post 15 Oct 2012)	15.0%
MatchNow	Dark	3.0%
Alpha Intraspread	Dark	2.5%
Pure	Lit	2.4%
TMX Select	Lit	1.4%
Omega	Lit	1.2%
Liquidnet / Instinet	Dark (block)	0.2%

Table 2
Descriptive statistics on trading activity

This table reports descriptive statistics on market-wide dark and lit trading activity during the two months preceding the minimum price improvement rules (15 August 2012 – 14 October 2012) and two months after (15 October 2012 – 15 December 2012). The trading activity variables are calculated on each trading day, pooling across all stocks in our sample (TSX Composite Index constituents). The mean, median and standard deviation are calculated from the daily observations. The last two columns report the difference in means pre/post regulation, and the significance of the difference using a two-tailed t-test. Standard errors are clustered both by stock and date. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

	Pre-regulation			Post-regulation			Difference	t-statistic
	Mean	Median	Standard deviation	Mean	Median	Standard deviation		
Dark \$ volume / total \$ volume (%)	9.01	7.58	7.00	5.93	4.05	6.33	-3.07	(-10.16)***
Dark trade size (\$1,000)	6.19	4.64	8.75	6.40	4.73	9.24	0.21	(0.76)
Lit trade size (\$1,000)	6.16	4.59	7.97	6.09	4.48	7.77	-0.07	(-0.33)
Dark daily \$ volume (\$100m)	4.92	4.90	0.89	3.21	3.32	0.71	-1.71	(-9.67)***
Lit daily \$ volume (\$100m)	57.41	53.97	18.73	56.68	57.94	13.88	-0.73	(-0.20)
Total daily \$ volume (\$100m)	62.01	58.82	19.13	59.46	61.45	14.40	-2.56	(0.69)

Table 3**Descriptive statistics on liquidity, informational efficiency and control variables**

This table reports descriptive statistics on liquidity, informational efficiency and control variables during the two months preceding the minimum price improvement rules (15 August 2012 – 14 October 2012) and two months after (15 October 2012 – 15 December 2012). Quoted spreads are time-weighted based on the lit national best bid and offer (NBBO). Realized and effective spreads are volume-weighted averages for the trades in each stock-day. Realized spreads are calculated using the NBBO midquote five minutes after the trade. Quoted, effective and realized spreads are measured relative to the midquote, in basis points. $ILLIQ_{it}$ is Amihud's price impact metric calculated for each stock-day using hourly return and volume observations. $Autocorrelation_{it}$, $VarianceRatio_{it}$, and $HFVolatility_{it}$ are the first principle components of midquote return absolute autocorrelations, variance ratios, and standard deviations at different intraday frequencies. $Delay_{it}$ measures intraday midquote return predictability using lagged market returns. $\$Volume_{it}$ is the natural logarithm of traded dollar volume. $MarketCap_{it}$ is the natural logarithm of the stock's market capitalization. $Volatility_{it}$ is the stock-day's high-low price range divided by the time-weighted midquote. $Price_{it}$ is the time-weighted midquote. $Constrained_{it}$ is the percentage of the trading day for which the stock's NBBO spread is constrained at one tick. The last two columns report the difference in means pre/post regulation, and the significance of the difference using a two-tailed t-test. Standard errors are clustered both by stock and date. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

	Pre-Regulation			Post-Regulation			Difference	t-stat
	Mean	Median	Standard deviation	Mean	Median	Standard deviation		
Panel A. Liquidity variables								
$QuotedSpread_{it}$	12.69	9.66	11.03	13.26	10.11	12.37	0.56	(2.66)***
$EffectiveSpread_{it}$	10.44	7.64	9.93	11.28	8.08	11.94	0.84	(3.58)***
$RealizedSpread_{it}$	2.29	1.56	11.65	3.49	2.15	13.38	1.21	(4.46)***
$ILLIQ_{it}$	1.74	1.51	1.21	1.80	1.53	1.30	0.06	(2.03)**
Panel B. Informational efficiency variables								
$Autocorrelation_{it}$	19.51	17.08	11.62	21.92	18.91	13.46	2.40	(5.40)***
$VarianceRatio_{it}$	25.01	22.91	13.03	27.65	25.13	14.73	2.64	(4.87)***
$HFVolatility_{it}$	21.55	18.58	12.77	23.07	20.11	13.45	1.51	(2.68)***
$Delay_{it}$	86.12	91.94	15.50	87.01	92.63	14.86	0.89	(1.82)*
Panel C. Control variables								
$\$Volume_{it}$	15.76	15.63	1.40	15.73	15.63	1.46	-0.03	(-0.58)
$MarketCap_{it}$	21.81	21.53	1.19	21.79	21.52	1.19	-0.01	(-1.68)*
$Volatility_{it}$	196.29	157.99	142.77	193.03	158.91	137.15	-3.26	(-0.53)
$Price_{it}$ (\$)	27.42	22.02	29.67	27.31	22.29	29.31	-0.10	(-0.52)
$Constrained_{it}$ (%)	58.80	65.69	35.33	59.78	71.41	36.90	0.98	(1.24)

Table 4

First-stage IV regressions of the impact of minimum price improvement rules on dark trading

This table reports coefficient estimates from the first stage of the instrumental variables regressions in which the introduction of the minimum price improvement rule is used as an instrument for the level of dark trading (measured as a fraction of total dollar volume), $Dark_{it}$:

$$Dark_{it} = \alpha_i + \beta_1 D_t^{PostPeriod} + \sum_{j=1}^6 \gamma_j Control_{j,it} + \varepsilon_{it} .$$

$D_t^{PostPeriod}$ is a dummy variable that takes the value 1 after the minimum price improvement rules come into effect and 0 before. $Time_t$ is a trend variable that starts at 0 at the beginning of the sample period and increments by 1 every trading day. $US Mean_t$ is the daily mean of a market quality metric (the market quality metric that is the dependent variable in the second stage) estimated for a matched sample of US firms. Consequently, the first stage is estimated separately for each market quality metric and the Table reports results from only two metrics (in specifications (1) and (2) $US Mean_t$ is the US daily mean quoted spread, whilst in (3) and (4) it is the daily mean of Amihud's $ILLIQ_{it}$). $\$Volume_{it}$ is the natural logarithm of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by the time-weighted midquote. $Price_{it}$ is the time-weighted midquote. $Constrained_{it}$ is the percentage of the trading day for which the stock's NBBO spread is constrained at one tick. Standard errors are clustered both by stock and date, and t-statistics are reported in parentheses. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively. The F-statistic tests the null hypothesis that the instruments do not affect the level of dark trading, i.e., that the coefficient of $D_t^{PostPeriod}$ is 0.

Variable	(1)	(2)	(3)	(4)
<i>Intercept</i>	0.194 (21.39)***	0.416 (24.32)***	0.218 (21.92)***	0.437 (24.21)***
$D_t^{PostPeriod}$	-0.032 (-15.53)***	-0.032 (-16.35)***	-0.030 (-14.82)***	-0.030 (-15.98)***
$Time_t$	-0.000 (-0.16)	0.000 (1.50)	-0.000 (-0.15)	0.000 (0.88)
$US Mean_t$	0.002 (3.32)***	0.001 (1.36)	0.001 (0.17)	-0.007 (-1.64)
$\$Volume_{it}$	-0.007 (-18.66)***	-0.020 (-23.45)***	-0.008 (-18.99)***	-0.021 (-23.85)***
$Volatility_{it}$	-0.477 (-14.23)***	-0.017 (-0.43)	-0.469 (-14.04)***	-0.016 (-0.40)
$Price_{it}$	0.000 (2.73)***	0.000 (1.41)	0.000 (3.64)***	0.000 (1.52)
$Constrained_{it}$			0.004 (2.21)**	0.010 (2.29)**
Observations	19,699	19,699	19,699	19,699
Adjusted R ²	9%	8%	9%	8%
F-statistic	271.5	267.4	219.7	255.3
Fixed effects	None	Stock	None	Stock

Table 5
Second-stage IV regressions of the impact of dark trading on liquidity

This table reports estimates from second-stage instrumental variables regressions:

$$y_{it} = \alpha_i + \beta_1 \widehat{Dark}_{it} + \sum_{j=1}^6 \gamma_j Control_{j,it} + \varepsilon_{it}.$$

The dependent variables, y_{it} , are estimates of liquidity and transaction costs for each stock-day. Quoted spreads are time-weighted based on the lit national best bid and offer (NBBO). Realized and effective spreads are volume-weighted averages for the trades in each stock-day. Realized spreads are calculated using the NBBO midquote five minutes after the trade. Quoted, effective and realized spreads are measured relative to the midquote, in basis points. $ILLIQ_{it}$ is Amihud's price impact metric calculated for each stock-day using hourly return and volume observations. The key independent variable, \widehat{Dark}_{it} , is the fitted value of a stock-day's dark dollar volume as a fraction of the stock-day's total dollar volume (from the first stage). $Time_t$ is a trend variable that starts at 0 at the beginning of the sample period and increments by 1 every trading day. $US Mean_t$ is the daily mean of the market quality metric (the same metric as the dependent variable) in a matched sample of US firms. $\$Volume_{it}$ is the natural logarithm of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by the time-weighted midquote. $Price_{it}$ is the time-weighted midquote. $Constrained_{it}$ is the percentage of the trading day for which the stock's NBBO spread is constrained at one tick. When $Constrained_{it}$ is used in the second-stage regression, the first stage that is used also contains the $Constrained_{it}$ variable. Adjusted R²s do not include the variance explained by the fixed effects. Standard errors are clustered both by stock and date, and t-statistics are reported in parentheses. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

Variable	<i>QuotedSpread_{it}</i>	<i>EffectiveSpread_{it}</i>	<i>RealizedSpread_{it}</i>	<i>ILLIQ_{it}</i>
Intercept	25.65 (17.08)***	19.45 (10.23)***	0.96 (0.18)	9.79 (40.52)***
\widehat{Dark}_{it}	-5.71 (-2.17)**	-22.14 (-5.94)***	-26.19 (-2.63)***	-2.38 (-5.45)***
$Time_t$	-0.01 (-4.15)***	-0.00 (-1.65)*	0.00 (0.65)	-0.00 (-3.16)***
$US Mean_t$	0.29 (9.76)***	0.08 (2.33)**	0.11 (1.13)	0.08 (2.37)**
$\$Volume_{it}$	-1.30 (-19.10)***	-0.84 (-8.70)***	0.35 (1.29)	-0.51 (-44.68)***
$Volatility_{it}$	72.23 (25.52)***	75.04 (17.18)***	-180.12 (-13.19)***	8.07 (19.18)***
$Price_{it}$	-0.06 (-4.37)***	-0.06 (-4.32)***	-0.06 (-2.19)**	-0.01 (-6.33)***
$Constrained_{it}$				-0.18 (-5.19)***
Observations	19,699	19,699	19,699	19,699
Adjusted R ²	9%	3%	4%	24%
Fixed effects	Stock	Stock	Stock	Stock

Table 6

Second-stage IV regressions of the impact of dark trading on informational efficiency

This table reports estimates from second-stage instrumental variables regressions:

$$y_{it} = \alpha_i + \beta_1 \widehat{Dark}_{it} + \sum_{j=1}^6 \gamma_j Control_{j,it} + \varepsilon_{it}.$$

The dependent variables, y_{it} , are estimates of informational efficiency for each stock-day. $Autocorrelation_{it}$, $VarianceRatio_{it}$, and $HFVolatility_{it}$ are the first principle components of midquote return absolute autocorrelations, variance ratios, and standard deviations at different intraday frequencies. $Delay_{it}$ measures intraday midquote return predictability using lagged market returns. The key independent variable, \widehat{Dark}_{it} , is the fitted value of a stock-day's dark dollar volume as a fraction of the stock-day's total dollar volume (from the first stage). $Time_t$ is a trend variable that starts at 0 at the beginning of the sample period and increments by 1 every trading day. $US Mean_t$ is the daily mean of the market quality metric (the same metric as the dependent variable) in a matched sample of US firms. $\$Volume_{it}$ is the natural logarithm of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by the time-weighted midquote. $Price_{it}$ is the time-weighted midquote. $Constrained_{it}$ is the percentage of the trading day for which the stock's NBBO spread is constrained at one tick. When $Constrained_{it}$ is used in the second-stage regression, the first stage that is used also contains the $Constrained_{it}$ variable. Adjusted R²s do not include the variance explained by the fixed effects. Standard errors are clustered both by stock and date, and t-statistics are reported in parentheses. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

Variable	$Autocorrelation_{it}$	$VarianceRatio_{it}$	$HFVolatility_{it}$	$Delay_{it}$
Intercept	40.45 (7.53)***	67.25 (11.73)***	-13.06 (-5.41)***	69.68 (10.01)***
\widehat{Dark}_{it}	-85.39 (-8.02)***	-129.36 (-12.12)***	-0.93 (-0.20)	-24.58 (-1.86)*
$Time_t$	-0.02 (-3.22)***	-0.04 (-6.03)***	0.01 (4.11)***	-0.00 (-0.44)
$US Mean_t$	0.16 (3.26)***	-0.03 (-0.95)	0.39 (18.49)***	21.85 (7.69)***
$\$Volume_{it}$	-1.08 (-4.26)***	-1.69 (-6.49)***	1.95 (16.38)***	0.18 (0.59)
$Volatility_{it}$	-19.32 (-2.11)**	-63.27 (-6.72)***	468.45 (48.24)***	-71.03 (-6.44)***
$Price_{it}$	0.15 (3.17)***	0.17 (3.26)***	-0.18 (-7.67)***	-0.01 (-0.25)
$Constrained_{it}$	-3.46 (-4.13)***	-7.64 (-9.08)***	-12.95 (-32.55)***	-1.82 (-1.97)**
Observations	19,670	19,684	19,699	19,683
Adjusted R ²	2%	3%	53%	6%
Fixed effects	Stock	Stock	Stock	Stock

Table 7
Robustness tests

This table reports t-statistics for the coefficients on the key independent variable in the second-stage instrumental variables regressions, for a variety of different specifications. The independent variable to which the t-statistics correspond is the fitted value of a stock-day's dark dollar volume as a fraction of the stock-day's total dollar volume. The rows correspond to different dependent variables and the columns correspond to different specifications/samples and are used to assess the robustness of the results. Specification (1) is the base case specification (the specification reported in Tables 5 and 6). Specification (2) is identical to specification (1) except that both the first and second stage has been run only for cross-listed securities. Specification (3) is identical to specification (1) except that the first stage IV regression is estimated on each stock separately. Specification (4) is identical to specification (1) omitting two weeks either side of the introduction of the minimum price improvement rules. Specification (5) and (6) are estimated on the largest 121 stocks and smallest 120 stocks, respectively, using the same variables as in specification (1). Specification (7) is identical to specification (1) except that the two stages are estimated simultaneously using maximum likelihood estimates. Specifications (8) and (9) are estimated on the 121 stocks that are most frequently constrained by the tick size, and the 120 stocks that are least frequently constrained, respectively. Standard errors are clustered both by stock and date in all specifications. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>QuotedSpread_{it}</i>	(-2.17)**	(-0.37)	(-5.87)***	(-1.53)	(-2.89)***	(-0.85)	(-1.95)*	(2.62)***	(-2.71)***
<i>EffectiveSpread_{it}</i>	(-5.94)***	(-3.43)***	(-10.06)***	(-5.89)***	(-4.24)***	(-4.78)***	(-5.72)***	(-1.35)	(-6.13)***
<i>RealizedSpread_{it}</i>	(-2.63)***	(-1.01)	(-1.58)	(-3.29)***	(-2.26)**	(-1.83)*	(-2.54)**	(-0.75)	(-2.78)***
<i>ILLIQ_{it}</i>	(-5.45)***	(-3.98)***	(-5.90)***	(-4.32)***	(-2.22)**	(-4.66)***	(-5.26)***	(-3.32)***	(-4.26)***
<i>Autocorrelation_{it}</i>	(-8.02)***	(-6.97)***	(-0.01)	(-9.58)***	(-5.02)***	(-6.29)***	(-7.19)***	(-3.06)***	(-7.75)***
<i>VarianceRatio_{it}</i>	(-12.12)***	(-9.74)***	(-0.49)	(-13.55)***	(-6.94)***	(-9.93)***	(-9.75)***	(-6.27)***	(-10.10)***
<i>HFVolatility_{it}</i>	(-0.20)	(-0.55)	(-1.68)*	(5.39)***	(-0.84)	(-0.33)	(-0.18)	(-0.34)	(0.20)
<i>Delay_{it}</i>	(-1.86)*	(-1.95)*	(-0.23)	(-3.30)***	(-2.53)**	(-0.02)	(-1.83)*	(-1.42)	(-1.37)
Observations	19,699	14,730	19,699	15,234	10,029	9,670	19,699	9,831	9,868
First stage IV	Pooled	Pooled	By stock	Pooled	Pooled	Pooled	Pooled	Pooled	Pooled
Estimation method	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	MLE	2SLS	2SLS

Table 8
The impact of different types of dark trading on market quality

This table reports estimates from second-stage instrumental variables regressions:

$$y_{it} = \alpha_i + \beta_1 \widehat{Fractional}_{it} + \beta_2 \widehat{Midpoint}_{it} + \sum_{j=1}^6 \gamma_j Control_{j,it} + \varepsilon_{it}$$

$\widehat{Fractional}_{it}$ and $\widehat{Midpoint}_{it}$ are the fitted values of the proportions of dollar volume executed via dark orders offering fractional price improvement, and via dark orders at the midquote, respectively. The dependent variables, y_{it} , are market quality metrics estimated each stock-day. Quoted, effective and realized spreads are measured relative to the midquote, in basis points. $ILLIQ_{it}$ is Amihud's price impact metric calculated for each stock-day using hourly return and volume observations. $Autocorr_{it}$, $VarianceRatio_{it}$, and $HFVolatility_{it}$ are the first principle components of midquote return absolute autocorrelations, variance ratios, and standard deviations at different intraday frequencies. $Delay_{it}$ measures intraday midquote return predictability using lagged market returns. The set of control variables is as follows. $Time_t$ is a trend variable that starts at 0 at the beginning of the sample period and increments by 1 every trading day. $US Mean_t$ is the daily mean of the market quality metric (the same metric as the dependent variable) in a matched sample of US firms. $\$Volume_{it}$ is the natural logarithm of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by the time-weighted midquote. $Price_{it}$ is the time-weighted midquote. $Constrained_{it}$ is the percentage of the trading day for which the stock's NBBO spread is constrained at one tick. Standard errors are clustered both by stock and date, and t-statistics are reported in parentheses. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

Variable	<i>Quoted Spread_{it}</i>	<i>Effective Spread_{it}</i>	<i>Realized Spread_{it}</i>	<i>ILLIQ_{it}</i>	<i>Autocorr - elation_{it}</i>	<i>Variance Ratio_{it}</i>	<i>HF Volatility_{it}</i>	<i>Delay_{it}</i>
Intercept	26.30 (27.05)***	15.78 (10.77)***	-7.72 (-1.81)*	9.16 (54.92)***	5.87 (1.65)*	15.70 (3.95)***	-12.45 (-7.52)***	59.95 (12.88)***
$\widehat{Fractional}_{it}$	-15.15 (-11.70)***	-22.85 (-13.43)***	-13.24 (-2.98)***	-1.36 (-7.84)***	-15.85 (-3.73)***	-28.18 (-6.45)***	-6.62 (-3.24)***	-6.04 (-1.13)
$\widehat{Midpoint}_{it}$	1.06 (0.75)	-6.19 (-3.29)***	1.51 (0.34)	-0.35 (-2.13)**	11.30 (2.72)***	17.74 (4.04)***	0.31 (0.16)	3.05 (0.70)
$Time_t$	-0.01 (-9.94)***	-0.01 (-4.45)***	0.01 (1.21)	-0.00 (-2.87)***	0.01 (1.23)	-0.00 (-0.19)	0.01 (3.10)***	0.00 (0.55)
$US Mean_t$	0.21 (6.78)***	0.08 (2.30)**	0.11 (1.09)	0.08 (2.52)**	0.22 (4.32)***	-0.01 (-0.24)	0.38 (18.26)***	22.30 (7.93)***
$\$Volume_{it}$	-1.26 (-25.73)***	-0.64 (-7.64)***	0.77 (3.22)***	-0.48 (-57.02)***	0.57 (3.23)***	0.79 (4.31)***	1.94 (21.98)***	0.64 (3.12)***
$Volatility_{it}$	73.47 (25.84)***	75.87 (17.40)***	-179.27 (-13.13)***	8.10 (19.26)***	-19.56 (-2.14)**	-62.91 (-6.69)***	469.12 (48.21)***	-70.93 (-6.43)***
$Price_{it}$	-0.07 (-4.63)***	-0.06 (-4.75)***	-0.07 (-2.41)**	-0.01 (-6.88)***	0.12 (2.55)**	0.12 (2.38)**	-0.18 (-7.70)***	-0.02 (-0.42)
$Constrained_{it}$				-0.20 (-5.80)***	-4.56 (-5.53)***	-9.29 (-11.17)***	-12.89 (-33.02)***	-2.11 (-2.35)**
Observations	19,699	19,699	19,699	19,699	19,670	19,684	19,699	19,683
Fixed effects	Stock	Stock	Stock	Stock	Stock	Stock	Stock	Stock
Adjusted R ²	9%	4%	2%	26%	8%	4%	56%	6%

Table 9

Further robustness tests separating fractional and midpoint dark trading

This table reports t-statistics for the coefficients on the key independent variable in the second-stage instrumental variables regressions, separating the dark volume traded with fractional price improvement from that traded at the midpoint. The independent variable to which the t-statistics correspond is the fitted value of a stock-day's dark dollar volume (either fractional or midpoint) as a fraction of the stock-day's total dollar volume. The rows correspond to different dependent variables and the columns correspond to different specifications/samples and are used to assess the robustness of the results. Specification (1) is the base case specification (the specification reported in Table 8). Specification (2) is identical to specification (1) except that both the first and second stage has been run only for cross-listed securities. Specification (3) is identical to specification (1) omitting two weeks either side of the introduction of the minimum price improvement rules. Specification (4) and (5) are estimated on the largest 121 stocks and smallest 120 stocks, respectively, using the same variables as in specification (1). Specifications (6) and (7) are estimated on the 121 stocks that are most frequently constrained by the tick size, and the 120 stocks that are least frequently constrained, respectively. Standard errors are clustered both by stock and date in all specifications. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Fractional dark trading t-statistics							
<i>QuotedSpread_{it}</i>	(-11.70)***	(-10.04)***	(-11.26)***	(-5.25)***	(-9.01)***	(-8.78)***	(-7.76)***
<i>EffectiveSpread_{it}</i>	(-13.43)***	(-11.20)***	(-12.20)***	(-4.89)***	(-11.81)***	(-7.82)***	(-9.02)***
<i>RealizedSpread_{it}</i>	(-2.98)***	(-2.20)**	(-3.04)***	(-2.85)***	(-2.08)**	(-2.64)***	(-1.47)
<i>ILLIQ_{it}</i>	(-7.84)***	(-8.45)***	(-6.25)***	(-3.97)***	(-6.04)***	(-6.91)***	(-1.32)
<i>Autocorrelation_{it}</i>	(-3.73)***	(-3.64)***	(-2.68)***	(-0.46)	(-4.62)***	(-1.28)	(-8.07)***
<i>VarianceRatio_{it}</i>	(-6.45)***	(-5.14)***	(-5.44)***	(-3.36)***	(-5.89)***	(-2.81)***	(-10.44)***
<i>HFVolatility_{it}</i>	(-3.24)***	(-3.31)***	(-0.83)	(-1.29)	(-2.14)**	(-2.43)**	(-4.04)***
<i>Delay_{it}</i>	(-1.13)	(-1.23)	(-1.08)	(-0.96)	(-0.44)	(-1.13)	(-0.77)
Panel B: Midpoint dark trading t-statistics							
<i>QuotedSpread_{it}</i>	(-0.75)	(1.76)*	(2.37)**	(-0.99)	(1.81)*	(-4.40)***	(1.20)
<i>EffectiveSpread_{it}</i>	(-3.29)***	(-3.11)***	(-2.02)**	(-2.50)**	(-1.23)	(-5.17)***	(-1.66)*
<i>RealizedSpread_{it}</i>	(-0.34)	(-0.15)	(-0.06)	(-0.99)	(-1.13)	(-1.21)	(1.03)
<i>ILLIQ_{it}</i>	(-2.13)**	(-2.04)**	(-1.88)*	(-1.48)	(-2.17)**	(-1.16)	(-1.98)**
<i>Autocorrelation_{it}</i>	(2.72)***	(2.11)**	(2.61)***	(3.89)***	(-1.13)	(0.96)	(2.14)**
<i>VarianceRatio_{it}</i>	(4.04)***	(3.31)***	(4.21)***	(3.37)***	(2.94)***	(1.95)*	(3.01)***
<i>HFVolatility_{it}</i>	(-0.16)	(-0.42)	(-0.76)	(-1.19)	(1.90)*	(-0.66)	(0.43)
<i>Delay_{it}</i>	(-0.70)	(-0.82)	(-0.41)	(-0.78)	(-0.03)	(1.10)	(0.50)
Observations	19,699	14,730	15,234	10,029	9,670	9,831	9,868
First stage IV	By Symbol	By Symbol	By Symbol	By Symbol	By Symbol	By Symbol	By Symbol

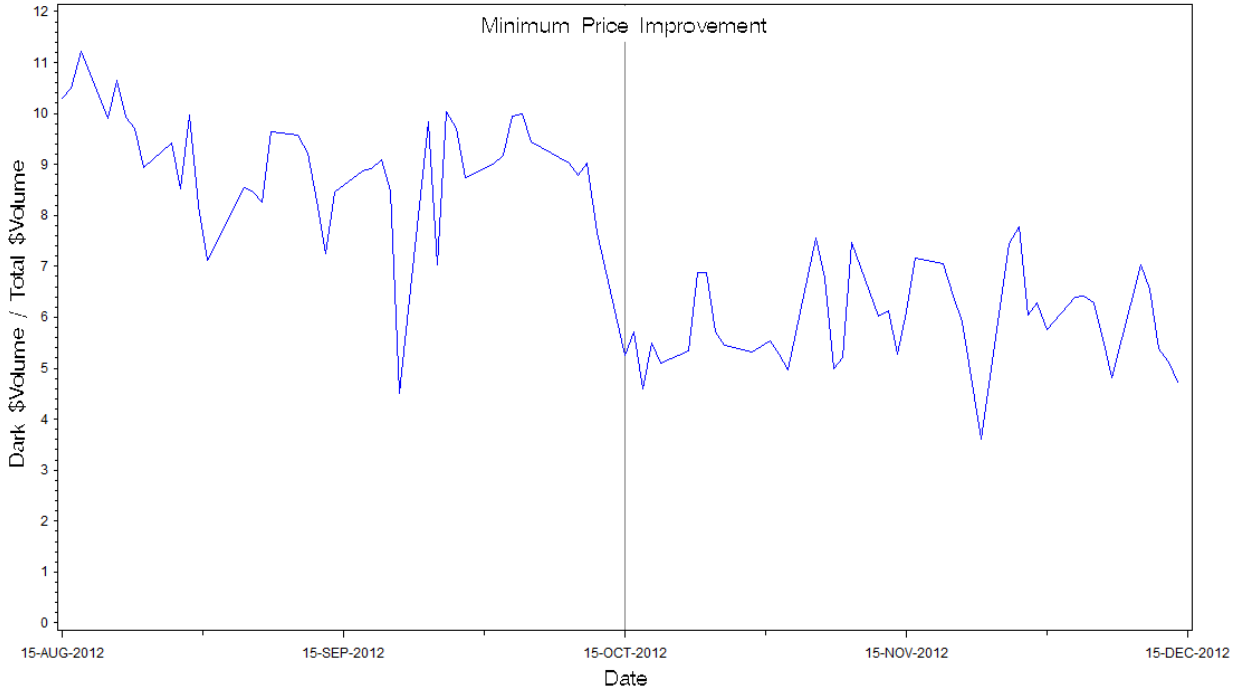


Figure 1. Dark trading in Canada as a percentage of consolidated dollar volume.

This figure shows daily dark trading in Canada as a fraction of total consolidated dollar volume, for constituents of the TSX Composite Index, from 15 August 2012 to 15 December 2012. The dark trading fraction is constructed by aggregating the dollar volume of dark trades executed and dividing it by the total dollar volume of trading on all of the main venues (TSX, Chi-X, Alpha, MatchNow, Intraspread, TMX Select, Pure Trading, and Omega). The aggregation of trading volume uses proprietary data from MatchNow, Intraspread, Chi-X, Alpha and TSX. The aggregation of dark trading volume does not include dark block trades executed on Liquidnet/Instinet. The vertical bar indicates the introduction of minimum price improvement requirements on 15 October 2012.

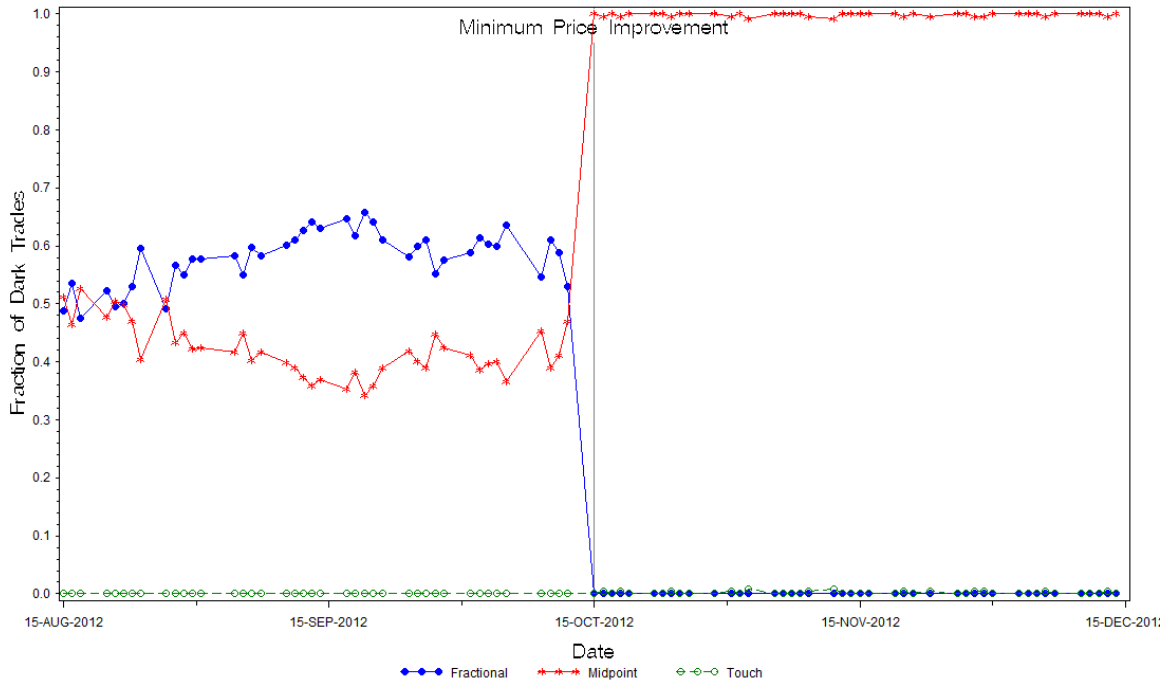


Figure 2. Price improvement provided by dark trades.

This figure shows the fraction of dark trades providing different levels of price improvement around the introduction of the minimum price improvement regulation (indicated by the vertical bar). *Fractional* price improvement consists of 10% and 20% price improvement orders executed on MatchNow and Alpha Intraspread, respectively. *Midpoint* consists of orders executed at the midpoint of the NBBO, and could be executed on any of the dark venues. *Touch* refers to orders larger than 50 Standard Trading Units or \$100,000 executed at the NBBO. This order type only became valid after the introduction of the minimum price improvement rules.




Figure 3. Distribution of dark trading across stock-days.

This histogram shows the distribution of the share of dark trading (fraction of total dollar volume), $Dark_{it}$, across stock-days. Values of $Dark_{it}$ greater than 0.5 have been aggregated into the 0.5 bucket.