

Should we be afraid of the dark? Dark trading and market quality [☆]

Sean Foley ^a and Tālis J. Putniņš ^{b,c,*}

^a *University of Sydney, NSW 2006, Australia*

^b *University of Technology Sydney, PO Box 123 Broadway, NSW 2007, Australia*

^c *Stockholm School of Economics in Riga, Strelnieku Street 4a, Riga, LV 1010, Latvia*

Abstract

We exploit a unique natural experiment—recent restrictions of dark trading in Canada and Australia—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, we do not find consistent evidence that dark midpoint crossing systems significantly affect 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

[☆] The Internet Appendix that accompanies this paper is available at <http://goo.gl/NY4svm>.

We thank the TSX, ITG MatchNow, Chi-X Canada, Alpha Intraspread and the Securities Industry Research Centre of Asia-Pacific (SIRCA) for providing data. We also thank an anonymous referee, Kent Bailey, Alina Bazavan, Torstein Braaten, Carole Comerton-Forde, David Easley, Kingsley Fong, Doug Foster, Corey Garriott, Rick Harris, Amy Kwan, Ingrid Lo, Katya Malinova, Ron Masulis, Maureen O'Hara, Andreas Park, Andrew Patton, Paul Redman, Steve Satchell, Jonathan Sylvestre, Susan Thorpe, Vincent van Kervel, Michael McKenzie, Terry Walter, Mao Ye, Bart Yueshen, Haoxiang Zhu, and seminar participants at the Western Finance Association Annual Meeting, the Northern Finance Association Annual Meeting, the CAFS Behavioral 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 EFMA Annual Meeting, the Bank of Canada, the Ontario Securities Commission, the UK Financial Conduct Authority, 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. 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. Putniņš is grateful for funding from the Australian Research Council (DE150101889).

* Corresponding author. Tel.: +61 2 9514 3088. Fax: +61 2 9514 7722.

Email addresses: sean.foley@sydney.edu.au (S. Foley) and talis.putnins@uts.edu.au (T. Putniņš).

1. Introduction

While trading without pre-trade transparency has long been a feature of equity markets in the form of upstairs block trading, the recent emergence of automated dark pools for smaller sized non-transparent orders has attracted the attention of regulators worldwide. Dark pools have been very successful in attracting order flow; they account for approximately 15% of US consolidated volume, 10% in Europe, 14% in Australia, and 10% in Canada.¹ Proponents argue dark pools offer several advantages such as 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) Chairman in a recent 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 Many regulators and policymakers including the SEC, the Financial Industry Regulatory Authority (FINRA), the Committee of European Securities Regulators (CESR), and the European Commission have made proposals and conducted public consultations regarding dark trading, but have been hesitant in introducing new regulations. Their hesitance reflects the scarcity of evidence on the costs and benefits of dark pools and how the costs/benefits are distributed between 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 world) and Australia in May 2013. The rules require that dark trades 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 fell by over one-third in both countries, literally overnight. Using the regulation as our main 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 thus far hindered the empirical analysis of dark trading and market quality.

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 August-December 2012.

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

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 because probing orders can infer the direction of the dark order imbalance. The second type, ‘two-sided’ dark trading is when dark liquidity can co-exist 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 both buys and sells as long as dark liquidity exists, can profit from dark liquidity provision strategies, and can better conceal their 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, and makes prices closer to the random walk that is expected under informational efficiency. The magnitudes of these effects are economically meaningful and qualitatively similar in both Canada and Australia. Two-sided dark trading is associated with lower lit market depth in Canada, although this effect is small compared to the effects on spreads and is not present in the Australian data where we find that two-sided dark trading increases depth. We show that the reduction in Canadian lit market depth is consistent with the notion that when trading activity is split across multiple venues, so too is depth, without necessarily decreasing the total depth across all venues.

In contrast to the beneficial effects of two-sided dark trading, we do not find consistent evidence that one-sided (midpoint) dark trading has a significant effect on market quality. While it may benefit some aspects of market quality, it can be harmful to others.

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

stocks, with stronger effects in smaller stocks. The similarity of the main results in Canada and Australia provides evidence on the robustness of the effects of dark trading.

Our results have support in the theoretical literature. The positive effect of two-sided dark trading (dark limit order markets) 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 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 opacity in limit order markets not only increases liquidity but also leads to more aggressive informed trading, which improves informational efficiency. Our results suggest that strong competition 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 reasons why one-sided (midpoint) dark trading can have less favorable effects on liquidity. For example, Zhu (2014) shows that informed traders are less likely than uninformed traders to use dark midpoint markets because their tendency to cluster on the same side of the market results in 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 (one-sided) dark trading consistently harms liquidity, our results indicate that midpoint dark trading is not as beneficial to liquidity as two-sided dark trading, consistent with an opposing increase in adverse selection risks.

This paper makes two major contributions to the recent empirical studies of dark trading. It is the first to identify and exploit a unique natural experiment as part of its identification strategy and therefore provides causal evidence on the effects of dark trading. Overcoming the endogeneity problem and obtaining sufficiently detailed data have been significant challenges for empirical studies of dark trading. Second, the paper identifies a partition of dark trading that is theoretically and empirically important. The results for the two dark trading types provide a possible reason why empirical studies of dark trading in different markets find different results; namely, the composition of dark trading, which determines its impact, varies across countries and markets. Kwan, Masulis, and McInish (2015) 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, leading to increased fragmentation. Along similar lines, Buti, Consonni, Rindi, Wen, and Werner (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 dark trading is associated with improved market quality is consistent with our results. Bloomfield, O'Hara, and Saar (2015) study dark trading in a laboratory market. In their setting, the ability to hide liquidity has a substantial impact on trader behavior and order submission strategies, but a modest effect on aggregate outcomes such as liquidity. Ready (2014) analyzes the determinants of volume in two block crossing systems (Liquidnet and POSIT) and finds more dark trading in stocks with lower levels of adverse selection risks. Degryse, de Jong, and van Kervel (2015) 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, Rindi, and Werner (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 (2015) find that in Australia low levels of dark trading can improve price discovery and decrease spreads, but when dark volume exceeds 10% of total trading, informational efficiency deteriorates.

Our paper is also related to two studies of the minimum price improvement requirements for dark trades in Canada and Australia. Foley and Putniņš (2014) show that in both countries, the minimum price improvement requirements substantially decreased the level of dark trading and increased average bid-ask spreads for stock market index constituents. Their analysis suggests that the increase in spreads was larger for stocks that experienced more substitution from two-sided dark markets to one-sided midpoint dark crossing systems, consistent with our finding that two-sided dark trading benefits market quality. Comerton-Forde, Malinova, and Park (2015) highlight the heterogeneity in how the Canadian minimum price improvement rules impacted dark pools and investor groups. They show that for the most liquid cross-listed Canadian stocks there is no significant change in market quality around the regulation. We confirm this finding and show that in a broader cross-section of stocks there is a significant change in market quality.

In contrast to Foley and Putniņš (2014) and Comerton-Forde, Malinova, and Park (2015) who focus on the effects of minimum price improvement rules, we focus on understanding the causal effects of different types of dark trading. The regulation in our paper primarily serves as one of the instrumental variables that facilitates identification of the effects of dark trading. Understanding the effects of dark trading has broader use in policymaking and market design than understanding the effects of a particular regulation. For example, in addition to shedding light on

the effects of the regulation, our findings also: (i) can be used to predict the effects of other regulation or market structure changes that affect either of the two types of dark trading, (ii) explain why dark trading can have different effects in different countries or markets as a result of variation in the composition of dark trading types, (iii) validate recent theory and identify mechanisms that should be incorporated in theoretical models, and (iv) provide guidance on whether trends in dark trading through time and in the cross-section of stocks should be a concern to policymakers.

Our results have a number of policy implications, in particular given that dark trading is high on the current agenda of many regulators. Our results imply 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 the more likely aggregate dark trading benefits 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. 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, the minimum price improvement rules in Canada and Australia decreased the amount of dark trading but also caused substitution from two-sided dark trading to one-sided midpoint dark trading. Consequently, market-wide liquidity decreased after the regulation.

Finally, our results have implications for tick size regulation, which is at the center of a US SEC Pilot Program. The tick size has a number of effects on trader behavior and market quality; for recent evidence see O'Hara, Saar, and Zhong (2014), Buti, Consonni, Rindi, Wen, and Werner (2015). Our findings suggest that an additional channel by which the tick size affects market quality is its effect on the structure of dark trading. Dark trading is more active in stocks when the spread is constrained by the tick size (Kwan, Masulis, and McNish, 2015). 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 force beneficial two-sided dark markets out of existence. Our results suggest that minimum price improvement requirements can be improved by ensuring tick sizes do not constrain the lit spread so that two-sided dark markets are not forced out of existence.

2. Theory and hypotheses

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.³ Internalization of order flow is associated with issues in addition to transparency, such as cream-skimming (e.g., Easley, Kiefer, and O'Hara, 1996).

There is a great deal of variation in how dark trading occurs in different venues. One partition of dark trading that is important both theoretically and empirically is whether dark trades execute (i) at a single price such as the midpoint of the NBBO, or (ii) at different prices on both the buy- and sell-sides of the market (e.g., 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.⁴

There are three important differences between these two types of dark trading. The first is execution probability and its impact on order routing decisions and order flow spillovers between venues. In a market with one-sided dark trading, at any point in time, dark liquidity (unexecuted 'resting' dark orders) can be available to buyers *or* to sellers depending on the order imbalance at the midquote, but not to both. If there is a sell imbalance at the midquote, then only incoming buy orders can execute immediately against the resting dark orders, and incoming sell orders would either be placed in the queue at the midquote or re-routed to a two-sided market for immediate execution, depending on routing preferences. The opposite is true if there is a buy imbalance. In contrast, in a two-sided dark market, as long as there are resting dark orders on both sides, both buyers and sellers can instantly execute their trades against the resting dark liquidity. Therefore, (assuming resting dark orders are present) an impatient buy or sell order (one that is willing to cross the spread) routed to a midpoint dark market at a random time has a probability around 50% of executing immediately, whereas if routed to a two-sided dark market, the probability is 100%.

³ In contrast to the US, Europe and Australia, fair access rules in Canada and requirements for price improvement on internalized orders have hampered the development of automated 'internalizers'.

⁴ 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 Centre Point, 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.

Impatient traders that are not able to obtain immediate execution in a one-sided dark venue will often route their orders to lit venues, thereby increasing imbalances in lit order flow. Hendershott and Mendelson (2000) show that the increased order imbalances in the lit market can harm lit market liquidity.⁵ This is because the one-sided midpoint venue executes some of the balanced order flow (which is a source of profits for lit liquidity providers), but none of the order imbalance (which imposes costs such as inventory holding costs and adverse selection costs). Lit liquidity providers implicitly use the profits from executing balanced order flow to subsidize the costs of unbalanced order flow and therefore a decrease in the balanced flow requires a higher bid-ask spread for a liquidity provider to break even.

The impact of the order flow spillovers from one-sided dark markets is even more important if we account for the tendency for some traders to cluster on one side of the market, i.e., correlation in order flow direction. For example, informed traders buy when they have good news and sell when they have bad news. If their information is correlated (or in the extreme case, identical) they will tend to buy at the same time and sell at the same time as other informed traders. Therefore, an impatient informed buy order routed to a one-sided dark market is more likely to occur when there is a buy imbalance at the midquote than a sell imbalance, and thus has *less* than a 50% probability of immediate execution. In contrast, the probability is around 50% for an uninformed order that is not correlated with other traders' order flow. The same holds for more patient orders that are willing to wait in the queue at the midquote if they are unable to immediately execute: informed orders will tend to be placed in longer queues than uninformed orders due to their correlation with other informed order flow. Zhu (2014) shows this mechanism more formally; in his model, informed traders face lower execution probabilities in a one-sided dark market than uninformed traders. The consequence of the difference in execution probabilities is that the order imbalance that spills over into the lit market (after failing to immediately execute in the one-sided dark market, or due to choosing not to route to the one-sided market) has a higher concentration of informed orders. Therefore, one-sided dark trading can harm lit market liquidity through an increase in adverse selection risks (Zhu, 2014).

⁵ 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 also has other effects such as attracting new uninformed order flow, which counteract the increase in inventory holding costs and adverse selection risks.

Importantly, the tendency for dark trading to increase order imbalances in the lit market and thereby harm liquidity occurs for one-sided dark trading but not necessarily two-sided dark trading. This leads to our first hypothesis:

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

The second important difference between one-sided and two-sided dark trading is the extent to which they conceal the trading intentions of patient (liquidity-supplying) orders. One-sided dark trading reveals more information about trading intentions than two-sided dark trading. For example, in a one-sided dark market, market participants know the direction of the order imbalance at the midquote if they have a resting order at the midquote that has not yet executed (the imbalance must be in the direction of that order) or if they have recently had an order immediately filled at the midquote (the imbalance must be in the opposite direction to that order). It is also possible for traders to infer the direction of the midquote imbalance by submitting probing orders at the midquote and observing whether they immediately execute or not. The direction of the imbalance is informative about the trading intentions of others and reveals 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 (assuming liquidity is posted on both sides) it is not possible to infer whether there is an overall buy or sell imbalance by observing when one's orders execute or using probing trades. Informed traders, who are concerned about revealing their trading intentions, are therefore likely to be more willing to provide liquidity in a two-sided dark market than in a one-sided market.

Third, dark liquidity provision can be profitable in a two-sided dark market because a liquidity provider can earn a non-zero spread in the dark. In contrast, there is zero spread in a one-sided dark market by construction and therefore there is little incentive for market participants to act as dark liquidity providers without an alternative reason for wanting to trade.

A number of theoretical studies predict that less pre-trade transparency in a limit order market encourages more aggressive competition in liquidity provision, in particular among informed traders. 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 dark 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.⁶ This leads to our second hypothesis:

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

Two-sided dark trading can also benefit liquidity by increasing the number of liquidity providers when liquidity fragments across multiple venues (e.g., Biais, Martimort, and Rochet, 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, Bisiere, and Spatt, 2010; Buti, Consonni, Rindi, Wen, and Werner, 2014). Most of the mechanisms by which dark trading can have a positive impact on liquidity occur when dark trading is two-sided and thus provide added support for Hypothesis 2.

⁶ 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 iceberg orders, in which only a portion of the limit order is displayed and the remainder is hidden. They find that iceberg orders encourage traders to compete in liquidity provision and therefore increase gains from trade.

Because theory predicts that different types of dark trading can have different effects on market quality, the impact of aggregate dark trading (aggregating across different types) 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.

3. Institutional setting and the natural experiment

Our main analysis focusses on dark trading in Canada and therefore this section describes the Canadian institutional setting. We describe the differences in the Australian institutional setting when we turn to the analysis of the Australian data.

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.⁷ 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.⁸ 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

⁷ A sixth alternative trading system (Chi-X 2) was added in April 2013, after our sample period.

⁸ 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.

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 October 15, 2012, all dark orders in Canada were required to provide some price improvement, resulting in dark executions within the 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 facilitate one-sided dark trading, whereas venues that accommodate ‘fractional’ price improvement (e.g., 10% and 20%) facilitate two-sided dark trading. The use of order routing algorithms is common during our sample period and results in virtual consolidation of the multiple trading venues. A routing algorithm might seek to execute an impatient order in a particular venue (e.g., a dark midpoint venue) and if it fails to obtain execution, then re-route the order to a different venue. More details on the order types and order execution priority rules are provided in the Internet Appendix.⁹

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.¹⁰ 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%.¹¹ Brokers are also able to internalize orders off-market.

⁹ The Internet Appendix is available at <http://goo.gl/NY4svm>

¹⁰ 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.

¹¹ This statistic is taken from the IIROC Marketplace Statistics Report available at www.iiroc.ca.

However, the order exposure rule together with the fair access regulations have hampered the development of automated ‘internalizers’ such as those in the US, Australia and elsewhere.¹²

Although dark trading in Canada shares many similarities with the US, there are four key differences. First, fair access regulations require dark venues in Canada to allow access to all brokers; in contrast, many dark pools in the US have exclusive access (e.g., Boni, Brown, and Leach, 2013). Second, internalized orders in Canada must provide price improvement over the best lit NBBO price of at least one cent. Third, odd-lots are not able to be traded in the dark during our sample period, significantly increasing the cost of ‘pinging’ the dark with a small probing order. Fourth, most Canadian marketplaces have broker preferencing, which allows passive (lit or dark) orders to break time priority to execute with an incoming active order from the same broker.

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 April 13, 2012, IIROC notice 12-0130 announced changes to the UMIR, which became effective on October 15, 2012. These changes imposed 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 \$100,000.¹³ 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 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.¹⁴

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. Fig. 1 shows 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

¹² See the Universal Market Integrity Rules (UMIR) Section 6.3.

¹³ 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.

¹⁴ 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.

one third. The reduction in dark trading occurred very quickly and distinctly around the change in regulation.

< Fig. 1 here >

Fig. 2 shows that prior to the introduction of the minimum price improvement requirements approximately 60% of dark trades were executed in two-sided markets (at fractional price increments of 10% or 20% of the NBBO spread), while the remaining 40% executed at the midquote. Under the new regulation, two-sided dark trading effectively disappeared and almost all dark trading now takes place at the midpoint of the NBBO (99.8% of all dark trades). Although large dark orders are allowed to execute at the NBBO after giving priority to lit orders, such dark orders are rare and account for a negligible fraction of dark trades.

< Fig. 2 here >

4. Data and metrics

Our analysis of dark trading in Canada uses 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 (August 15, 2012 – December 15, 2012).¹⁵ The four-month period is chosen as a compromise considering the following tradeoff. If the window is too narrow, the analysis will lack statistical power and will not adequately capture changes in market participants' trading behavior.¹⁶ 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.

¹⁵ We remove 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 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 242 stocks. Stocks outside of the TSX Composite Index tend to have very low levels of dark trading and thus would not contribute much to our analysis of the effects of dark trading. Our results are also robust to using only the stocks that are cross-listed in the US.

¹⁶ As illustrated in Fig. 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.

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 TSX¹⁷ directly from the trading venues.¹⁸ 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.¹⁹ 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 (August 15 – October 15, 2012):

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

The superscript C indexes Canadian stocks, while the superscript U indexes US-listed stocks.²⁰

For robustness, we also examine a second matched group. Exploiting the fact that most of the stocks in the TSX Composite Index are cross-listed in the US, we take the US side of the cross-listed stocks as a natural control group. The advantage of this approach is that, by construction, it provides a perfect control for stock characteristics. The weakness is that trading in the US side of the cross-listed stocks might not be entirely exogenous with respect to the minimum price improvement regulation due to cross-market arbitrage and the possibility of order flow migration to the US. The results using the two different groups of matched stocks are quite similar (the latter are reported as part of our robustness tests).

¹⁷ Copyright 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.

¹⁸ The proprietary data consists only of information that was publically reported to the consolidated tape.

¹⁹ 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.

²⁰ 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.

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, Amihud’s (2002) illiquidity metric, and depth. 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. Depth is measured as the log of the time-weighted average quoted dollar depth available to be bought and sold at the NBBO on all lit venues.

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.²¹ 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, Subrahmanyam, and van Dijk (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 zero (indicating high levels of efficiency) to 100 (indicating low levels of efficiency).

5. Empirical analysis

The core of our analysis uses the introduction of restrictions on dark trading as a natural experiment and source of exogenous variation to identify the causal effects of dark trading. Our main analysis, reported in the next three Subsections (5.1-5.3), focuses on the effects of dark trading in Canada, while the fourth Subsection (5.4) reports Australian results and compares them to those of Canada.

²¹ Autocorrelations are used in Hendershott and Jones (2005) and Anderson, Eom, Hahn, and Park (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 Hou and Moskowitz (2005).

Our main analysis starts with simple univariate pre/post comparisons of mean market characteristics. These show that spreads are wider and informational efficiency is lower after the regulation, consistent with the notion that dark trading can be beneficial to market quality. Next, we turn to one-stage least squares (OLS) panel regressions of market quality metrics on dark trading and control variables. Although much of the variation in dark trading is around the exogenous introduction of restrictions on dark trading, the OLS regressions do not address the potential endogeneity of dark trading and are thus reported only in the Internet Appendix for completeness (the results from the OLS concur with our main conclusions). Our main conclusions are drawn from instrumental variables regressions, exploiting the minimum price improvement regulation as the main instrumental variable. We first examine the effects of dark trading in aggregate, before partitioning dark trading into two types. We control for trends in market quality in a matched sample of US stocks, effectively giving a difference-in-differences estimate. All of these analyses, from the univariate pre/post comparisons to the instrumental variables regressions and an extensive set of robustness tests (spread across the paper and the Internet Appendix) point to the same conclusions—dark trading in our sample benefits market quality, and the benefits are driven by two-sided dark trading.²²

5.1. Descriptive statistics

Table 2 reports descriptive statistics on trading activity in Canada before and after the minimum price improvement regulation came into effect. Consistent with Fig. 1, the level of dark 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 median lit trading. Average total daily traded dollar volume remains unchanged at approximately \$6.2 billion per day.

< Table 2 here >

²² The Internet Appendix is available at <http://goo.gl/NY4svm>.

To get a sense of the variation in Canadian dark trading, Fig. 3 Panel A presents the pooled sample histogram of stock-day level dark trading as a fraction of total stock-day dollar volume, $Dark_{it}$. Approximately 9% 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. Panel B shows the distribution of changes in the dark trading around the regulation (the distribution of the average $Dark_{it}$ before the regulation minus the average after the regulation). The distribution shows that for a typical stock, the level of dark trading decreases by around 3%, but there is considerable heterogeneity in the impact of the regulation—approximately 10% of stocks experience an *increase* in the level of dark trading after the regulation.

< Fig. 3 here >

Fig. 4 reports the distributions of changes in dark trading, by dark trading type. Consistent with Fig. 2, Panel A shows that almost all Canadian stocks experience a decrease in two-sided dark trading (the exception is the 3% of stocks that have no change), with the typical decrease being around 3%. Panel B shows that in contrast, there is greater heterogeneity in the impact of the regulation on one-sided dark trading. Many stocks experience an *increase* in the amount of one-sided (midpoint) dark trading, with a typical increase being 1-2%.

< Fig. 4 here >

Table 3 reports descriptive statistics on the stock-day market quality metrics and control variables for Canada. 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 and depth decreases a statistically significant amount after the regulation. The four informational inefficiency metrics also show a deterioration in market quality after the regulation. 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, Masulis, and McInish, 2015). Buti, Rindi, and Werner (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 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 ($Dark_{it}$, the fraction of dollar volume in stock i on day t that is traded in the dark) on the instrumental variables and a set of control variables. The main instrumental variable is a dummy variable for the minimum price improvement rule ($D_t^{Post} = 1$ after the rule change and 0 before). This instrument alone is sufficient for identification. However, as an additional instrument, we also include the lagged level of dark trading ($Dark_{i,t-1}$).²³ It is not uncommon to include lagged endogenous variables as instruments in a microstructure setting (e.g., Sarkar and Schwartz, 2009). The first-stage regression is:

$$Dark_{it} = \alpha_i + \beta_1 D_t^{Post} + \beta_2 Dark_{i,t-1} + \sum_{j=1}^5 \gamma_j Control_{j,it} + \varepsilon_{it}, \quad (2)$$

where α_i is a set of stock fixed effects, and $Control_{j,it}$ comprises the following five control variables. $Time_t$ takes the value zero on the first day in the sample and increments by one 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 inclusion of $US Mean_t$ in the second stage removes

²³ In robustness tests (reported later in this section) we show that a simple model using just D_t^{Post} as an instrumental variable produces similar results to our baseline model.

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 (more on this below). The other control variables are $\$Volume_{it}$ (the natural log of traded dollar volume), $Volatility_{it}$ (the stock-day's high-low price range divided by 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).²⁴ In our main results we estimate the first-stage Eq. (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.

Table 4 reports the results from the first stage of the 2SLS for Canada. 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 quoted spread. 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%, holding other factors fixed, and the decline is highly statistically significant. The F-statistics in our first-stage regression with and without stock fixed effects (826–1,312) are above both this level and the critical values specified by Stock and Yogo (2005), allowing us to reject the null of a weak instrument.²⁵ There is a strong negative relation between $Dark_{it}$ and $\$Volume_{it}$ and the relation becomes stronger with the addition of stock fixed effects. This result is contrary to the finding of Buti, Rindi, and Werner (2011) that, in the US, dark pool activity is positively related to liquidity in the cross-section.²⁶ A possible explanation is that in our sample, there is considerably more variation in total dollar volume (the denominator of $Dark_{it}$) than in dark dollar volume (the numerator of $Dark_{it}$) so that when $\$Volume_{it}$ is high, $Dark_{it}$ is low and vice versa. To ensure our results are not driven predominantly by variation in $\$Volume_{it}$, we conduct two robustness tests. The first is excluding $\$Volume_{it}$ from the first-stage regressions (keeping it as a control variable in the second stage). The second is using dark dollar volume that is not scaled by total dollar volume. The results (reported in the Internet Appendix) indicate that our findings are robust to these alternative specifications.

< Table 4 here >

²⁴ Our results are robust to excluding the variable $Constrained_{it}$ from the first and second stages.

²⁵ The large F-statistics are a result of the highly significant reduction in dark trading around the regulation and the strong persistence in dark trading from one day to the next, as well as the relatively large sample size.

²⁶ Buti, Rindi, and Werner (2011) measure dark pool activity in a similar manner to this paper, i.e., dark pool volume as a percentage of total volume.

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

$$y_{it} = \alpha_i + \beta \widehat{Dark}_{it} + \sum_{j=1}^5 \gamma_j Control_{j,it} + \varepsilon_{it} , \quad (3)$$

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 produces a difference-in-differences estimate, but with two distinct advantages over a ‘standard’ form of such a model.²⁷ 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 differ between the US and Canada. Second, by summarizing the control market (US) with a single 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 in Canada. 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. The effect of dark trading on depth is not statistically different from zero. A small increase in dark trading by 5% of total dollar volume is expected to decrease average quoted spreads by approximately 0.20 bps ($0.05 \times (-3.98)$), decrease effective spreads by 0.69 bps, and decrease realized spreads by 0.70 bps. These decreases for just a small change in dark trading are economically meaningful compared to the means of quoted, effective and realized spreads: 12.86 bps, 10.75 bps, and 2.81 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 0.25 bps or 0.02 standard deviations ($0.064 \times (-$

²⁷ More specifically, ignoring the control variables for simplicity, a standard difference-in-differences model, $y_{i,t} = \mu D_t^{CANADA} + \theta D_t^{POST} + \delta D_t^{CANADA} D_t^{POST} + \varepsilon_{it}$ 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.

3.98)/12.8), decrease effective spreads by 0.88 bps or 0.08 standard deviations, decrease realized spreads by 0.89 bps or 0.32 standard deviations, and decrease price impacts (Amihud's $ILLIQ_{it}$) by 0.02 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 Fig. 2 illustrates, prior to the introduction of the minimum price improvement regulation, the majority of dark trading in Canada (approximately 60%) was two-sided. 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 spreads become wider through the course of our sample period, not counting the effects of the regulation and holding other variables including US trends constant. The coefficients on $US\ Mean_t$ are all positive and statistically significant (with the exception of the variance ratio, effective and 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 2% and 30% suggesting that there are many factors beyond the variables included in our model that influence liquidity.

The evidence in Table 5 indicates that dark trading in our sample has overall benefits for liquidity. The minimum price improvement regulation decreased the level of dark trading and consequently increased average spreads (this is seen in Table 3). However, the regulation also

increased the amount of price improvement received by the active side of dark trades. In fact, the total dollar amount of price improvement received by the active side of all dark trades in our sample increases from approximately \$60,000 per day before the regulation to \$73,000 per day after the regulation, despite the decline in the amount of dark trading. With wider spreads and worse liquidity in the lit market but more price improvement for dark trades, it is natural to ask whether market execution quality overall improves or deteriorates after the regulation. To answer this question, we compare the increase in the dollar amount of price improvement for dark trades to the dollar value of the deterioration in lit spreads. The spread costs for all lit trades (dollar value of effective half-spreads for the active side of all lit trades) increase from an average of \$2,009,000 per day before the regulation to \$2,381,000 after. The \$13,000 per day in additional price improvement for the active side of dark trades is negligible compared to the \$372,000 per day increase in lit market execution costs. So while the active side of dark trades gets better execution on average after the regulation, that improvement is not nearly enough to compensate for the overall deterioration in lit liquidity.

< Table 6 here >

Turning to the informational efficiency proxies, Table 6 reports second-stage regression estimates of the impact of dark trading in Canada. The results suggest that dark trading reduces the level of all informational inefficiency metrics: absolute autocorrelations, variance ratios, high-frequency volatility, and delay in reflecting market-wide information (although the effect on delay is not statistically significant). 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 0.07 standard deviations ($0.064 \times (-0.06)/0.05$), decrease the variance ratio by 0.07 standard deviations, and decrease high-frequency volatility by 0.03 standard deviations, after controlling for other market characteristics and stock fixed effects. 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, Roll,

and Subrahmanyam, 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 instrumental variables 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 (as reported in Tables 5 and 6, corresponding to Eq. (3)), which includes all of the control variables and stock fixed effects. Specification (2) is a simpler instrumental variables model, which uses only the regulation dummy variable as an instrument (omitting lagged dark trading). Other than the reduced instrument set, the control variables, sample and model structure is identical to specification (1). Specification (3) is identical to the base case (1) except that the sample is constrained to cross-listed securities only.²⁸ Specification (4) is identical to the base case (1) except that the first-stage 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 (5) is identical to the base case (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. In specification (6), rather than matching each Canadian stock to a different US stock, we use the US side of cross-listed stocks as the control group. Specifications (7) and (8) are estimated on the largest 121 stocks and smallest 120 stocks, respectively, using the same variables as in specification (1). Specification (9) is identical to specification (1) except that the two stages are estimated simultaneously using maximum likelihood. Specifications (10) and (11) are estimated on the 68 stocks that are constrained by the tick size at least 90% of the time pre-regulation, and the 174 stocks that are constrained less than 90% of the time, respectively. Additional robustness tests reported in the Internet Appendix show that our results are robust to alternative measurement frequencies for the

²⁸ Here, we include stocks that are cross-listed on any other market outside of Canada. In the Internet Appendix we show that our results are also robust to using only stocks that are cross-listed on a US exchange. The robustness of the results to using the subsample of cross-listed stocks indicates that our results are not driven by the repeal of the short sale price restrictions (which occurred at the same time as the minimum price improvement requirements) because cross-listed stocks experienced no change with respect to these rules.

informational efficiency metrics, using simple one-stage OLS regressions, excluding $\$Volume_{it}$ from the first stage, and using dark dollar volume that is not scaled by total dollar volume.

< Table 7 here >

The results from the different specifications and subsamples are largely consistent with our base case estimates. 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 using the regulation as the only instrument, allowing for heterogeneity in the impact of the instrumental variables, allowing for transitory effects around the rule change, and constraining the sample to only cross-listed stocks. The effect of dark trading is qualitatively similar in the largest and the smallest stocks in our sample, consistent with Comerton-Forde and Putniņš (2015) who find that the effects of dark trading on price discovery are similar in both large and small stocks. The magnitudes and statistical significance of the effects tend to be greater in 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 stocks that are more/less frequently constrained by the tick size. Consistent with our hypothesis, the beneficial effects of dark trading are more pronounced (more statistically significant effects and greater magnitudes) for stocks that are least often constrained by the tick size.

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 ($TwoSided_{it}$) and one-sided midpoint dark trading ($OneSided_{it}$). Both types are measured as a fraction of total dollar volume. Following a similar approach to the previous section, the first-stage regressions are:

$$TwoSided_{it} = \alpha_i + \beta_1 D_t^{Post} + \beta_2 TwoSided_{i,t-1} + \sum_{j=1}^5 \gamma_j Control_{j,it} + \varepsilon_{it}, \quad (4)$$

$$OneSided_{it} = \mu_i + \delta_1 D_t^{Post} + \delta_2 OneSided_{i,t-1} + \sum_{j=1}^5 \theta_j Control_{j,it} + \varepsilon_{it}. \quad (5)$$

The model above is exactly identified.²⁹ In robustness tests, we estimate an over-identified model with an expanded instrument set including the lagged stock price, and find similar results. We also estimate an instrumental variables model that allows for heterogeneity in the effects of the regulation and find similar results.

In the second-stage regressions we include the fitted values of one-sided and two-sided dark trading together to estimate their independent impact on market quality:

$$y_{it} = \alpha_i + \beta_1 \widehat{TwoSided}_{it} + \beta_2 \widehat{OneSided}_{it} + \sum_{j=1}^5 \gamma_j Control_{j,it} + \varepsilon_{it}. \quad (6)$$

Table 8 reports the second-stage estimates for Canada. Consistent with our hypotheses, one-sided and two-sided dark trading have different effects on market quality. Two-sided dark trading is associated with strong improvements in almost all of the liquidity and informational efficiency metrics.³⁰ An exception is the marginally statistically significant (at the 10% level) negative relation between the amount of two-sided dark trading and depth in the lit markets. This effect, however, is small in magnitude compared with the effects on spreads. For example, the average level of two-sided dark trading prior to the regulation is around 5% of total dollar volume. If this level of two-sided dark trading were exogenously doubled holding other factors fixed, i.e., a 200% increase in two-sided dark trading, depth would be expected to decrease by an estimated 1.25%.³¹ For comparison, the expected changes in spreads for the same change in two-sided dark trading are -0.45 bps for quoted spreads, -0.83 bps for effective spreads, and -0.86 bps for realized spreads, representing decreases of 4% to 38% of the pre-regulation means. Thus, the effect of two-sided dark trading on depth is much smaller than its effect on spreads.

The weak negative relation between two-sided dark trading and lit depth could result from posted liquidity being split between lit and two-sided dark venues and thus does not necessarily imply a decline in overall liquidity. Because we are unable to observe depth in dark markets, we are unable to test how two-sided dark trading affects aggregate depth (summing

²⁹ The F-statistic for the joint hypothesis that the instruments are not significant in the first stage is 8,402 in regression Eq. (4) and 1,287 in regression Eq. (5), which is greater than the critical values specified in Stock and Yogo (2005), indicating the instruments are strong.

³⁰ For conciseness, from here on we do not report the results of the variance ratio as it is related to autocorrelation and produces qualitatively similar results. The Internet Appendix reports additional results using the variance ratio.

³¹ Doubling the level of two-sided dark trading is an increase of 0.05, implying a change in the natural log depth of $0.05 \times -0.25 = -0.00125$, which is a decrease in depth of approximately 1.25%. In dollar terms, the mean depth before the regulation is around \$49,020 ($e^{10.8}$, converting back from log terms), which would be expected to decrease to around \$48,441 if two-sided dark trading was to double.

across lit and dark markets). However, we are able to test whether depth relative to the amount of trading activity that remains in the lit markets is affected by two-sided dark trading. We compute the ratio of lit dollar depth to lit traded dollar volume and use this relative depth measure as the dependent variable. We find that dark trading (aggregate dark trading, two-sided dark trading, and one-sided dark trading) is not significantly related to relative depth, consistent with the notion that as trading activity fragments so does depth, without reducing the aggregate depth across all venues.³²

In contrast to the largely beneficial effects of two-sided dark trading, the effect of one-sided dark trading is somewhat mixed for the different market quality measures and is statistically indistinguishable from zero for most of the market quality metrics. For example, the results suggest that one-sided dark trading is associated with higher absolute autocorrelations (marginally statistically significant), but lower high-frequency volatility and effective spreads. The latter 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 one-sided dark trading are weaker with no conclusive evidence of positive or negative effects.

< Table 8 here >

We again subject our analysis to a range of robustness tests, which we summarize in Table 9. The first of the robustness tests is an alternative instrumental variables model that does not rely on lagged dark trading and instead makes greater use of the cross-sectional heterogeneity in the effects of the regulation. The intuition behind this alternative is as follows. For different stocks, the regulation has different effects on *OneSided_{it}* and *TwoSided_{it}*. For example, some stocks have a large reduction in *TwoSided_{it}* around the regulation, while others experience a small change or no change. Similarly, some stocks have a large reduction in *OneSided_{it}*, while others have a small change or no change. As long as the stocks that have a large reduction in *TwoSided_{it}* are not exactly the same stocks that have a large reduction in *OneSided_{it}* (i.e., the cross-sectional correlation of the regulation's impact on *TwoSided_{it}* and *OneSided_{it}* is less than perfect), then the stocks that have a large reduction in *TwoSided_{it}* will allow the effect of

³² These results are reported in the Internet Appendix.

$TwoSided_{it}$ on market quality to be identified. Similarly, the stocks that have a large reduction in $OneSided_{it}$ will allow the effects of $OneSided_{it}$ on market quality to be identified.

We examine whether there is sufficient cross-sectional heterogeneity in the effects of the regulation to reliably identify the independent effects of one-sided and two-sided dark trading. The details are in the Internet Appendix.³³ We conclude that the variation necessary for identification of the effects of one-sided and two-sided dark trading is present. However, identification of the effects of one-sided dark trading is driven by a relatively small number of stocks and therefore is likely to be imprecise relative to the paper's other results. Consequently, the mixed results about the effects of one-sided dark trading should be taken with some caution due to the lower statistical power compared to tests of the effects of two-sided dark trading.

For each stock separately, we estimate first-stage models for $TwoSided_{it}$ and first-stage models for $OneSided_{it}$. This is equivalent to:

$$TwoSided_{it} = \sum_{k=1}^{242} D_k \left(\alpha_k + \beta_k D_t^{Post} + \sum_{j=1}^5 \gamma_{jk} Control_{j,it} \right) + \varepsilon_{it} , \quad (7)$$

$$OneSided_{it} = \sum_{k=1}^{242} D_k \left(\mu_k + \delta_k D_t^{Post} + \sum_{j=1}^5 \theta_{jk} Control_{j,it} \right) + \vartheta_{it} , \quad (8)$$

where D_k is a dummy variable for each of the $k = 1, 2, \dots, 242$ Canadian stocks in our sample. The model has 242 instruments to identify two endogenous variables and is thus over-identified.

Results from this alternative instrumental variables model are reported in Table 9 column (2). The results are qualitatively similar to the baseline model reported in column (1). The other robustness tests in Table 9 indicate that the results are robust to using only the subset of cross-listed stocks, omitting two weeks either side of the minimum price improvement regulation, using the US side of cross-listed stocks as a control group, and running the analysis separately on the most/least tick size constrained stocks. In particular, the positive effect of dark trading with two-sided price improvement on liquidity and informational efficiency remains strong, whereas the effect of one-sided dark trading is weak and somewhat mixed. Columns (6) and (7) show that the effects of one-sided and two-sided dark trading are qualitatively similar in large and small stocks, but of greater magnitude in small stocks. Additional robustness tests reported in the Internet

³³ Fig. 4 shows that almost all stocks experience a decrease or no change in two-sided dark trading around the minimum price improvement requirements, whereas the changes in one-sided dark trading are distributed fairly symmetrically around a positive mean. These distributions suggest that it might be difficult to identify the effects of one-sided dark trading because that requires a sufficient number of stocks that have a small change (or no change) in two-sided dark trading yet have sufficient variation in one-sided dark trading changes. In the Internet Appendix we show that 40 of our 242 stocks have little or no change in two-sided dark trading (an absolute change less than 2%) and that for such stocks there is at least as much variation in one-sided dark trading changes as there is for the 202 stocks that have a large change in two-sided dark trading.

Appendix show that our results are robust to alternative measurement frequencies for the informational efficiency metrics, using simple one-stage OLS regressions, excluding $\$Volume_{it}$ from the first stage, and using dark dollar volume that is not scaled by total dollar volume.

The Internet Appendix compares our findings with those of Comerton-Forde, Malinova, and Park (2015) and describes additional analysis that reconciles the differences in findings. We find that while our results hold consistently across most subsets of stocks, the effects of dark trading are largely statistically insignificant in the most liquid Canadian cross-listed stocks and, for some measures, different to the effects of dark trading in other groups of stocks. The group of the most liquid cross-listed stocks is a small fraction of all Canadian stocks. Furthermore, the tendency for dark trading to have stronger effects in smaller and less liquid stocks suggests that our findings are likely to generalize to stocks that are not part of the TSX Composite Index.

< Table 9 here >

5.4. Comparison with dark trading in Australia

Subsequent to the Canadian introduction of minimum price improvement rules in October 2012, Australia introduced similar regulation on May 26, 2013. We repeat our analysis for the Australian market to assess the extent to which our main results generalize to different market structures. We use data for the ASX 200 Index constituents (200 of the most actively traded Australian stocks) for a period of two months before and two months after the rules come into effect, consistent with our analysis of the Canadian data.³⁴

There are some noteworthy differences in the structure of the Australian equity markets and the nature of dark trading in Australia compared to Canada.³⁵ First, Australian dark pools are used by brokers for internalization of client order flow as well as facilitating client-to-client trades, similar to the US, but contrary to Canada. In Canada the order exposure rule, fair access regulations, requirements for price improvement on internalized orders, and broker preferencing on most trading venues have hampered off-market internalization. Second, before the minimum price improvement requirements, dark trades in Australia were permitted at or within the best bid and ask on the lit market, whereas in Canada some, albeit minimal price improvement (e.g., 10%

³⁴ We obtain data from the AusEquities and Thomson Reuters Tick History databases maintained by the Securities Industry Research Centre of Asia Pacific. Our data allow us to precisely identify all dark trades using flags attached to trades. Our data also provide a complete record of all trades and quotes on lit markets with millisecond timestamp precision.

³⁵ For more details on dark trading in Australia, see ASIC (2013, 2015) and Comerton-Forde and Putniņš (2015).

of the spread) was required even before the new rules came into effect. Third, while lit liquidity is less fragmented in Australia than it is in Canada, dark liquidity in Australia is spread across a larger number of trading venues than in Canada. Australia has two lit venues: the Australian Securities Exchange (ASX), which executes approximately 70.5% of consolidated dollar volume in its lit limit order book during our sample period, and Chi-X Australia, which launched in 2011 and executes around 6.3% of consolidated dollar volume in its lit limit order book. Dark trading occurs in two exchange-operated venues (ASX Centre Point, a midpoint dark pool; and dark orders on Chi-X, which interact with lit orders on Chi-X) and around 21 broker-operated dark pools. The most active broker-operated dark pools in terms of turnover are those operated by Credit Suisse, Goldman Sachs, and Citigroup (ASIC, 2015). Unlike the exchange-operated dark venues and dark orders, each broker-operated dark pool is only accessible to a small fraction of the market. This contrasts with Canadian dark pools that allow access to all brokers under the Canadian fair access regulations. Some of the broker-operated Australian dark pools are virtually connected by dark liquidity aggregators that are able to access multiple dark pools. To varying degrees, the broker-operated dark pools facilitate client-to-client trading and broker internalization of client order flow.

There are also noteworthy differences in how the minimum price improvement requirements were implemented in Australia compared to Canada.³⁶ Perhaps the most significant is that the minimum price improvement requirements for dark trades were accompanied by a reduction in the block trade size threshold. Block trades in Australia are allowed to be negotiated away from the lit markets at any price (even after the minimum price improvement requirements for dark trades) and immediately reported to either the ASX or Chi-X Australia. Prior to the rule changes, for a single trade to qualify as a block trade and be exempt from the requirements for lit or dark trades, it had to have a value of at least \$1 million. After the rule changes, trades could qualify as block trades if they exceeded either \$0.2 million, \$0.5 million, or \$1 million, depending on the liquidity category of the stock.

The minimum price improvement requirements in Australia had a similar impact on dark trading as in Canada. The amount of dark trading fell virtually overnight from 15.2% of total dollar volume before the regulation to 9.2% after. One of the reasons why the decline in dark trading in Australia was larger than in Canada is that the accompanying reductions in block size thresholds allowed large dark orders to be executed as blocks and bypass the minimum price

³⁶ See Foley and Putniņš (2014) for details of the Australian regulation and how it compares to the Canadian minimum price improvement requirements.

improvement requirements (block trading increased from around 9.7% to 12.5% of total dollar volume). The fraction of dark dollar volume executed at the midquote increased from 46% before the regulation to 81% after. Therefore, the minimum price improvement requirements in Australia also provide a useful instrument to analyze the causal effects of dark trading.

Tables 10 and 11 report estimates of the effects of dark trading in Australia, using 2SLS instrumental variables models similar to those used in our analysis of dark trading in Canada.³⁷ Aggregating across the two types of dark trading (Table 10), we find that dark trading in Australia is largely beneficial to market quality, consistent with our findings for Canada. The statistically significant results indicate that dark trading leads to narrower quoted spreads, lower price impacts as measured by Amihud's $ILLIQ_{it}$, more lit market depth, and lower levels of high-frequency volatility (controlling for low-frequency volatility). The magnitudes of the effects are economically meaningful; for example, an increase in dark trading by 5% of total dollar volume is expected to decrease average quoted spreads by approximately 1.1 bps ($0.05 \times (-21.44)$) and increase depth by approximately 4.3% ($0.05 \times (-21.44)$). A one-standard-deviation increase in dark trading is expected to decrease $ILLIQ_{it}$ by 18% of a standard deviation and decrease high-frequency volatility by 9% of a standard deviation.

< Table 10 here >

Table 11 shows that similar to the results for Canada, one-sided and two-sided dark trading in Australia differ in their effects on market quality. Two-sided dark trading tends to benefit liquidity, with statistically significant effects on quoted spreads, $ILLIQ_{it}$, and depth, and no statistically significant detrimental effects. One-sided dark trading is associated with lower $ILLIQ_{it}$, higher depth, and less high-frequency volatility (marginally statistically significant impact), but larger delay in stocks reflecting market-wide information ($Delay_{it}$).³⁸ Unlike two-sided dark trading, one-sided dark trading is not associated with narrower spreads; in fact the point estimate suggests the opposite, although the effect is not statistically distinguishable from zero. On balance, consistent with the Canadian results and our hypotheses, we find less consistent beneficial effects for one-sided dark trading than two-sided.

³⁷ We exclude the $US\ Mean_t$ control variable as trends in the Australian market are unlikely to be sufficiently similar to those in the US.

³⁸ A one-standard-deviation increase in one-sided dark trading is expected to increase $Delay_{it}$ by 8% of a standard deviation.

< Table 11 here >

The effects of dark trading in Australia are largely consistent with those in Canada, despite substantial differences in market structure and the way dark trading is used. Our results for the Australian market are also consistent with two existing studies of dark trading in Australia. Comerton-Forde and Putniņš (2015) find that while high levels of dark trading can harm price discovery, low levels below 10% of dollar volume in individual stocks can be beneficial. During the two months before the regulation, slightly over half of the stocks in our sample (54%) have a median level of dark trading below 10% indicating that for most of the market the level of dark trading was within the range where dark trading can be beneficial.³⁹

Our results are largely consistent with Foley and Putniņš (2014) with the exception of a few differences. We reconcile the differences with further analysis (reported in the Internet Appendix) in which we find the following. In Australia, one-sided dark trading does not have a statistically significant effect on spreads during a four-month window around the regulation (Table 11), yet in a longer 18-month window one-sided dark trading is associated with significantly wider spreads. We identify market structure changes that coincide with the change in the effects of one-sided dark trading. In contrast, two-sided dark trading is beneficial to liquidity in both sample periods. In the longer sample, the detrimental effects of one-sided dark trading on spreads outweigh the beneficial effects of two-sided dark trading. These findings reinforce the two main conclusions of this paper: (i) one-sided and two-sided dark trading have different effects on market quality, with two-sided dark trading being more consistently beneficial, and (ii) the effects of the aggregate level of dark trading depend critically on the mix of one-sided and two-sided dark trading in the aggregate.

6. Conclusions

We use a unique natural experiment, the introduction of the minimum price improvement regulation in Canada and Australia, 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

³⁹ Comerton-Forde and Putniņš (2015) highlight that the 10% estimated threshold for individual stocks should not be compared to market-wide averages because high levels of dark trading in a few of the largest stocks can result in a market-wide average above 10% even if most stocks have levels below 10%.

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, and makes prices closer to the random walk that would be expected under informational efficiency. The magnitudes of the effects are economically meaningful and qualitatively similar in both Canada and Australia.

In contrast, we do not find consistent evidence that one-sided 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 qualitatively similar for the largest and the smallest stocks in our Canadian sample, but of larger magnitude in small stocks. The effects of dark trading are largely insignificant in a subset of the most liquid Canadian cross-listed stocks. Considering the non-trivial differences in Canadian and Australian market structure, the similarity of the results in Canada and Australia speak to the robustness of the effects of dark trading and the ability to generalize the paper's findings.

Our findings have two caveats. First, the levels of dark trading in Canada and Australia 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 beyond some 'tipping' point. Second, our analysis of the Canadian market 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 should be analyzed as a separate type of dark trading. Although the Australian dark trading includes internalization, our data do not allow us to separate broker internalization from client-to-client 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 supply liquidity because they can profit from liquidity provision without revealing their private

information. Our results suggest that strong competition in providing dark liquidity forces lit market liquidity providers to narrow spreads to compete with dark liquidity. Consistent with our results, Boulatov and George (2013) also show that aggressive liquidity provision 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., Biais, Martimort, and Rochet, 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, Bisiere, and Spatt, 2010; Buti, Consonni, Rindi, Wen, and Werner, 2014). Lastly, Zhu (2014) and Hendershott and Mendelson (2000) show 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 consistently 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 imply 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. Our results suggest that the larger the proportion of two-sided dark trading in the aggregate the more likely the aggregate dark trading benefits market quality. A harmful level of aggregate dark trading in one country might not be harmful in another due to differences in the composition of dark trading types. In designing regulation it is important to consider the effect on the level of dark trading and also on the composition of dark trading types. The minimum price improvement rules in Canada and Australia decreased the level of dark trading but also caused substitution from two-sided dark trading to one-sided dark trading.

Finally, our results have implications for tick size regulation because the tick size can affect the structure of dark trading when minimum price improvement rules are imposed. Our results suggest that using dark trading as a way of obtaining a finer price grid 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 rules can force two-sided dark markets out of existence in stocks that are constrained by the tick size and consequently have unintended negative effects on market quality. Our results suggest that minimum price improvement requirements can be improved by ensuring tick sizes do not constrain the lit spread so as to avoid forcing two-sided dark markets out of existence.

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, Amihud's (2002) illiquidity metric, and depth. All liquidity metrics are calculated for each stock-day. We measure the quoted NBBO spread in bps 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 in Canada (in Australia, the data identify the trade initiator). 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 .⁴⁰ To reduce the influence of outliers we winsorize the liquidity metrics at the 1% level for each stock and each date.

⁴⁰ As indicated in Eq. (4) the $ILLIQ_{it}$ metric is log transformed to reduce the impact of outliers, consistent with Karolyi, Lee, and van Dijk (2012). If there is no volume traded in a given hour, the denominator in

We measure depth as the log of the time-weighted average quoted dollar depth available to be bought and sold at the NBBO on all lit venues.

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, Eom, Hahn, and Park, 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):

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

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. Throughout the paper we report results using $k = 10\text{sec.}$, and in the Internet Appendix we show that our results are robust to using the other frequencies as well as combining the various autocorrelation measures into a single one by calculating their first principal component.

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 for 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. Higher values are associated with greater inefficiency. We use the (l, kl) combinations: (1sec., 10sec.), (10sec., 60sec.), (1min., 5min.). Throughout the paper we report results using the

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.

combination of frequencies (10sec., 60sec.), and in the Internet Appendix we show that our results are robust to using the other frequency combinations as well as combining the various variance ratios into a single one by calculating their first principal component.

For each stock-day we also estimate the intra-day midquote standard deviations ($HFVolatility_{it}$) calculated using returns of 10, 30 and 60 second horizons. $HFVolatility_{it}$ 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. Throughout the paper we report results using ten-second returns, and in the Internet Appendix we show that our results are robust to using the other frequencies as well as combining the various frequency $HFVolatility_{it}$ estimates into a single one by calculating their first principal component.

The final measure of informational efficiency is an intraday adaptation of the Hou and Moskowitz (2005) $Delay_{it}$ metric. $Delay_{it}$ 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 one-minute midquote returns for the stock, $r_{it,\tau}$, on the TSX60 market index return, $r_{mt,\tau}$, and ten lags:

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

We save the R^2 from this unconstrained regression, $R^2_{Unconstrained,it}$, 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, $R^2_{Constrained,it}$. $Delay_{it}$ is calculated from the ratio of the constrained and unconstrained regression R^2 s:

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

$Delay_{it}$ takes values between zero 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 $R^2_{Unconstrained,it}$ 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., Eom, K.S., Hahn, S.B., Park, J.H., 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.
- ASIC, 2013. Australian Securities and Investments Commission Report 331, Dark liquidity and high frequency trading.
- ASIC, 2015. Australian Securities and Investments Commission Report 452, Review of high-frequency trading and dark liquidity.
- Biais, B., Bisiere, C., Spatt, C., 2010. Imperfect competition in financial markets: An empirical study of Island and Nasdaq. *Management Science* 56, 2237–2250.
- Biais, B., Martimort, D., Rochet, J., 2000. Competing mechanisms in a common value environment. *Econometrica* 68, 799–837.
- Bloomfield, R., O'Hara, M., Saar, G., 2015. Hidden liquidity: Some new light on dark trading. *Journal of Finance* 70, 2227–2273.
- Boni, L., Brown, D.C., Leach, J.C., 2013. Dark pool exclusivity matters. Unpublished working paper. University of New Mexico.
- Boulatov, A., George, T.J., 2013. Hidden and displayed liquidity in securities markets with informed liquidity providers. *Review of Financial Studies* 26, 2095–2137.
- Buti, S., Rindi, B., 2013. Undisclosed orders and optimal submission strategies in a limit order market. *Journal of Financial Economics* 109, 797–812.
- Buti, S., Rindi, B., Werner, I.M., 2011. Diving into dark pools. Unpublished working paper. University of Toronto.
- Buti, S., Consonni, F., Rindi, B., Wen, Y., Werner, I.M., 2014. Sub-penny and queue jumping. Unpublished working paper. Bocconi University.
- Buti, S., Consonni, F., Rindi, B., Wen, Y., Werner, I.M., 2015. Tick size: Theory and evidence. Unpublished working paper. Bocconi University.
- Chordia, T., Roll, R., Subrahmanyam, A., 2008. Liquidity and market efficiency. *Journal of Financial Economics* 87, 249–268.
- Comerton-Forde, C., Malinova, K., Park, A., 2015. The impact of the dark trading rules. Report prepared for IIROC.
- Comerton-Forde, C., Putniņš, T.J., 2015. Dark trading and price discovery. *Journal of Financial Economics* 118, 70–92.

- Degryse, H., de Jong, F., van Kervel, V., 2015. The impact of dark trading and visible fragmentation on market quality. *Review of Finance* 19, 1587–1622.
- Easley, D., Kiefer, N.M., O'Hara, M., 1996. Cream-skimming or profit-sharing? The curious role of purchased order flow. *Journal of Finance* 51, 811–833.
- Foley, S., Putniņš, T.J., 2014. Regulatory efforts to reduce dark trading in Canada and Australia: How have they worked? Report prepared for the CFA Institute.
- Foucault, T., Menkveld, A.J., 2008. Competition for order flow and smart order routing systems. *Journal of Finance* 63, 19–58.
- Hendershott, T., Jones, C.M., 2005. Island goes dark: Transparency, fragmentation, and regulation. *Review of Financial Studies* 18, 743–793.
- Hendershott, T., Mendelson, H., 2000. Crossing networks and dealer markets: Competition and performance. *Journal of Finance* 55, 2071–2115.
- Hou, K., Moskowitz, T.J., 2005. Market frictions, price delay, and the cross-section of expected returns. *Review of Financial Studies* 18, 981–1020.
- Huang, R., Stoll, H., 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., Lee, K.H., van Dijk, M.A., 2012. Understanding commonality in liquidity around the world. *Journal of Financial Economics* 105, 82–112.
- Kwan, A., Masulis, R., McNish, T.H., 2015. Trading rules, competition for order flow and market fragmentation. *Journal of Financial Economics* 115, 330–348.
- Lee, C.M.C., Ready, M.J., 1991. Inferring trade direction from intraday data. *Journal of Finance* 46, 733–746.
- Lo, A., MacKinlay, C., 1988. Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies* 1, 41–66.
- Nimalendran, M., Ray, S., 2014. Informational linkages between dark and lit trading venues. *Journal of Financial Markets* 17, 230–261.
- O'Hara, M., Saar, G., Zhong, Z., 2014. Relative tick size and the trading environment. Unpublished working paper. Cornell University.
- O'Hara, M., Ye, M., 2011. Is market fragmentation harming market quality? *Journal of Financial Economics* 100, 459–474.
- Ready, M.J., 2014. Determinants of volume in dark pool crossing networks. Unpublished working paper. University of Wisconsin.

- Rindi, B., 2008. Informed traders as liquidity providers: Anonymity, liquidity and price formation. *Review of Finance* 12, 497–532.
- Rösch, D.M., Subrahmanyam, A., van Dijk, M.A., 2013. An empirical analysis of co-movements in high- and low-frequency measures of market efficiency. Unpublished working paper. University of California, Los Angeles.
- Sarkar, A., Schwartz, R.A., 2009. Market sidedness: Insights into motives for trade initiation. *Journal of Finance* 64, 375–423.
- Stock, J.H., Yogo, M., 2005. Testing for weak instruments in linear IV regression. In: Andrews, D.W.K. (Ed.), *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rotenberg*. Cambridge University Press, New York, pp. 80–108.
- Zhu, H., 2014. Do dark pools harm price discovery? *Review of Financial Studies* 27, 747–789.

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 two months preceding the minimum price improvement rules (August 15, 2012 – October 14, 2012) and two months after (October 15, 2012 – December 15, 2012). The market share of Liquidnet / Instinet is obtained from IIROC statistics, whereas the other market shares are calculated from our data (TSX Composite Index constituents).

Venue	Lit / Dark	Pre-regulation market share	Post-regulation market share
TSX	Both	59.5%	56.4%
Chi-X	Both	14.1%	20.8%
Alpha (Lit)	Both (after Oct 15, 2012)	14.7%	13.0%
MatchNow	Dark	4.7%	4.9%
Alpha Intraspread	Dark	3.1%	0.3%
Pure	Lit	1.8%	1.8%
TMX Select	Lit	1.3%	1.6%
Omega	Lit	0.7%	0.9%
Liquidnet / Instinet	Dark (block)	0.2%	0.2%

Table 2
Descriptive statistics on trading activity

This table reports descriptive statistics on market-wide dark and lit trading activity in Canada during the two months preceding the minimum price improvement rules (August 15, 2012 – October 14, 2012) and two months after (October 15, 2012 – December 15, 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 in Canada during the two months preceding the minimum price improvement rules (August 15, 2012 – October 14, 2012) and two months after (October 15, 2012 – December 15, 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. $Depth_{it}$ is the natural log of the time-weighted quoted dollar depth at the NBBO. $Autocorrelation_{it}$ and $HFVolatility_{it}$ are ten-second midquote return absolute autocorrelations and standard deviations, respectively. $VarianceRatio_{it}$ measures the variance ratio of the standard deviation of ten-second and one-minute midquote returns. $Delay_{it}$ measures intraday midquote return predictability using lagged market returns. $\$Volume_{it}$ is the natural log of traded dollar volume. $MarketCap_{it}$ is the natural log 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								
<i>QuotedSpread_{it}</i>	12.69	9.66	11.03	13.26	10.11	12.37	0.56	(2.66)***
<i>EffectiveSpread_{it}</i>	10.44	7.64	9.93	11.28	8.08	11.94	0.84	(3.58)***
<i>RealizedSpread_{it}</i>	2.29	1.56	11.65	3.49	2.15	13.38	1.21	(4.46)***
<i>ILLIQ_{it}</i>	1.74	1.51	1.21	1.80	1.53	1.30	0.06	(2.03)**
<i>Depth_{it}</i>	10.86	10.67	0.93	10.82	10.64	0.92	-0.05	(-2.17)***
Panel B. Informational inefficiency variables								
<i>Autocorrelation_{it}</i>	0.05	0.04	0.04	0.06	0.04	0.05	0.01	(4.96)***
<i>VarianceRatio_{it}</i>	0.13	0.11	0.11	0.15	0.12	0.12	0.02	(5.62)***
<i>HFVolatility_{it}</i>	2.88	2.51	1.66	3.1	2.71	1.82	0.22	(3.05)***
<i>Delay_{it}</i>	86.12	91.94	15.50	87.01	92.63	14.86	0.89	(1.82)*
Panel C. Control variables								
<i>\$Volume_{it}</i>	15.76	15.63	1.40	15.73	15.63	1.46	-0.03	(-0.58)
<i>MarketCap_{it}</i>	21.81	21.53	1.19	21.79	21.52	1.19	-0.01	(-1.68)*
<i>Volatility_{it}</i>	196.29	157.99	142.77	193.03	158.91	137.15	-3.26	(-0.53)
<i>Price_{it}</i> (\$)	27.42	22.02	29.67	27.31	22.29	29.31	-0.10	(-0.52)
<i>Constrained_{it}</i> (%)	58.80	65.69	35.33	59.78	71.41	36.90	0.98	(1.24)

Table 4**First-stage regressions of the impact of minimum price improvement rules on dark trading**

This table reports estimates from first-stage instrumental variables regressions in which the endogenous variable for which we instrument is the level of dark trading (measured as a fraction of total dollar volume), $Dark_{it}$:

$$Dark_{it} = \alpha_i + \beta_1 D_t^{Post} + \beta_2 Dark_{i,t-1} + \sum_{j=1}^5 \gamma_j Control_{j,it} + \varepsilon_{it}$$

The instrumental variables are D_t^{Post} (a dummy variable that takes the value one after the minimum price improvement rules come into effect and zero before), and $Dark_{i,t-1}$ (the lagged level of dark trading). The set of control variables includes the following. $Time_t$ is a trend variable that starts at zero at the beginning of the sample period and increments by one 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. This table reports results using quoted spreads in $US Mean_t$. $\$Volume_{it}$ is the natural log of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by 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. Specification (2) is identical to specification (1) with the addition of stock 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. The F-statistic tests the null hypothesis that the instruments do not affect the level of dark trading. The sample comprises TSX Composite Index stocks during the period August 15, 2012 to December 15, 2012.

Variable	(1)	(2)
<i>Intercept</i>	0.19 (21.23)***	0.40 (23.42)***
D_t^{Post}	-0.03 (-14.37)***	-0.02 (-12.08)***
$Dark_{i,t-1}$	0.03 (3.86)***	0.23 (19.92)***
$Time_t$	-0.00 (-0.05)	0.00 (1.45)
$US Mean_t$	0.00 (3.05)***	0.00 (0.21)
$\$Volume_{it}$	-0.01 (-19.04)***	-0.02 (-23.12)***
$Volatility_{it}$	-0.50 (-14.87)***	0.02 (0.45)
$Constrained_{it}$	0.00 (0.38)	0.01 (2.33)**
Observations	19,233	19,233
Adjusted R ²	8%	30%
F-statistic	1,312	826
Fixed effects	None	Stock

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

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

$$y_{it} = \alpha_i + \beta_1 \widehat{Dark}_{it} + \sum_{j=1}^5 \gamma_j Control_{j,it} + \varepsilon_{it}.$$

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. $Depth_{it}$ is the natural log of the time-weighted quoted dollar depth at the NBBO. 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 zero at the beginning of the sample period and increments by one 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 log of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by 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. Adjusted R^2 s do not include the variance explained by the fixed effects (α_i). 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 sample comprises TSX Composite Index stocks during the period August 15, 2012 to December 15, 2012.

Variable	<i>Quoted Spread_{it}</i>	<i>Effective Spread_{it}</i>	<i>Realized Spread_{it}</i>	<i>ILLIQ_{it}</i>	<i>Depth_{it}</i>
\widehat{Dark}_{it}	-3.98 (-2.61)***	-13.78 (-6.42)***	-13.93 (-2.61)***	-0.65 (-2.66)***	-0.10 (-0.74)
$Time_t$	0.00 (3.36)***	0.01 (3.34)***	0.01 (3.18)***	0.00 (0.82)	-0.00 (-5.12)***
$US Mean_t$	0.14 (5.82)***	0.05 (1.61)	0.12 (1.19)	0.12 (3.77)***	0.48 (13.93)***
$\$Volume_{it}$	-0.84 (-15.61)***	-0.38 (-4.13)***	0.70 (2.68)***	-0.48 (-50.68)***	0.15 (23.92)***
$Volatility_{it}$	47.58 (17.72)***	56.39 (12.77)***	-191.27 (-13.44)***	8.24 (19.40)***	-5.87 (-19.26)***
$Constrained_{it}$	-12.66 (-51.55)***	-9.14 (-25.73)***	-5.39 (-5.61)***	-0.20 (-5.78)***	0.88 (34.54)***
Observations	19,233	19,233	19,233	19,233	19,233
Adjusted R^2	30%	9%	2%	25%	21%
Fixed effects	Stock	Stock	Stock	Stock	Stock

Table 6

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

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

$$y_{it} = \alpha_i + \beta_1 \widehat{Dark}_{it} + \sum_{j=1}^5 \gamma_j Control_{j,it} + \varepsilon_{it}.$$

The dependent variables, y_{it} , are estimates of informational efficiency for each stock-day. $Autocorrelation_{it}$ and $HFVolatility_{it}$ are absolute autocorrelations and standard deviations of ten-second midquote returns, respectively. $VarianceRatio_{it}$ is the variance ratio of ten-second and one-minute midquote returns. $Delay_{it}$ 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 zero at the beginning of the sample period and increments by one 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 log of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by 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. Adjusted R^2 s do not include the variance explained by the fixed effects (α_i). Standard errors are clustered both by stock and date, t-statistics are reported in parentheses. ***, **, and * indicate statistical significance at 1%, 5%, and 10% levels, respectively. The sample comprises TSX Composite Index stocks during the period August 15, 2012 to December 15, 2012.

Variable	$Autocorrelation_{it}$	$VarianceRatio_{it}$	$HFVolatility_{it}$	$Delay_{it}$
\widehat{Dark}_{it}	-0.06 (-2.83)***	-0.16 (-3.01)***	-1.49 (-4.27)***	-2.11 (-0.34)
$Time_t$	0.00 (2.02)**	0.00 (3.60)***	0.00 (3.87)***	0.01 (1.79)*
$US Mean_t$	0.18 (4.33)***	0.00 (0.14)	0.24 (14.82)***	22.13 (7.86)***
$\$Volume_{it}$	-0.00 (-1.13)	-0.00 (-0.07)	0.22 (17.09)***	0.58 (2.55)**
$Volatility_{it}$	-0.05 (-1.35)	-0.03 (-0.41)	64.75 (54.30)***	-69.69 (-6.30)***
$Constrained_{it}$	-0.02 (-7.70)***	-0.03 (-4.49)***	-1.70 (-28.65)***	-2.32 (-2.56)**
Observations	19,205	19,219	19,233	19,205
Adjusted R^2	3%	4%	57%	6%
Fixed effects	Stock	Stock	Stock	Stock

Table 7
Robustness tests

This table reports coefficients of the key independent variable in the second-stage instrumental variables regressions, for a variety of different specifications. The independent variable to which the coefficient estimates correspond is the fitted value of a stock-day's dark dollar volume as a fraction of the stock-day's total dollar volume (\widehat{Dark}_{it}). 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 for the sake of comparison (the specification reported in Tables 5 and 6, using TSX Composite Index stocks during the period August 15, 2012 to December 15, 2012). Specification (2) is identical to specification (1) except that it uses a single instrumental variable, the post-regulation dummy variable. Specification (3) is identical to specification (1) except that it includes cross-listed securities only. Specification (4) is identical to specification (1) except that the first-stage regression is estimated on each stock separately. Specification (5) is identical to specification (1) omitting two weeks either side of the introduction of the minimum price improvement rules. Specification (6) uses the US side of cross-listed stocks as the control group. Specifications (7) and (8) are estimated on the largest 121 stocks and smallest 120 stocks, respectively, using the same variables as in specification (1). Specification (9) is identical to specification (1) except that the two stages are estimated simultaneously using maximum likelihood. Specifications (10) and (11) are estimated on the 68 stocks that are constrained by the tick size at least 90% of the time pre-regulation, and the 174 stocks that are constrained less than 90% of the time, 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)	(10)	(11)
<i>QuotedSpread_{it}</i>	-3.98***	-6.21***	-3.33*	-4.24***	-6.75***	-4.03***	-1.73	-4.47**	-7.99***	-2.01	-4.26**
<i>EffectiveSpread_{it}</i>	-13.78***	-18.55***	-11.66***	-11.67***	-17.57***	-13.78***	-5.85***	-18.97***	-19.38***	-12.22***	-13.82***
<i>RealizedSpread_{it}</i>	-13.93***	-24.71**	-11.17*	-4.52	-22.61***	-13.91***	-4.47	-19.31**	-24.87***	-5.73	-15.61**
<i>ILLIQ_{it}</i>	-0.65***	-2.57***	-0.73***	-0.67***	-0.32	-0.68***	-0.41*	-0.81**	-0.80***	-0.57	-0.61**
<i>Depth_{it}</i>	-0.10	0.81***	-0.07	-0.11	0.22	0.19	-0.08	-0.20	0.31***	-0.48	0.07
<i>Autocorrelation_{it}</i>	-0.06***	-0.37***	-0.05*	-0.00	-0.04	-0.08***	-0.05*	-0.07**	-0.09***	0.00	-0.09***
<i>VarianceRatio_{it}</i>	-0.16***	-0.72***	0.05	0.01	-0.16**	-0.16***	-0.16**	-0.16**	-0.27***	0.03	-0.23***
<i>HFVolatility_{it}</i>	-1.49***	-1.48**	-1.60***	-0.44**	-1.77***	-2.53***	-1.03***	-1.50***	-2.23***	-1.52**	-1.79***
<i>Delay_{it}</i>	-2.11	-24.55*	-4.02	0.76	-5.79	-2.03	-0.69	-0.77	-8.91*	6.61	-5.50
Observations	19,233	19,233	14,385	19,233	14,999	19,233	9,784	9,449	19,233	5,605	13,628
First stage	Pooled	Pooled	Pooled	By stock	Pooled	Pooled	Pooled	Pooled	Pooled	Pooled	Pooled
Estimation method	2SLS	2SLS	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:

$$y_{it} = \alpha_i + \beta_1 \widehat{TwoSided}_{it} + \beta_2 \widehat{OneSided}_{it} + \sum_{j=1}^5 \gamma_j \text{Control}_{j,it} + \varepsilon_{it}.$$

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 each stock-day using hourly return and volume observations. $Depth_{it}$ is the natural log of the time-weighted quoted dollar depth at the NBBO. $Autocorrelation_{it}$ and $HFVolatility_{it}$ are absolute autocorrelations and standard deviations of ten-second midquote returns, respectively. $Delay_{it}$ measures intraday midquote return predictability using lagged market returns. The key independent variables, $\widehat{TwoSided}_{it}$ and $\widehat{OneSided}_{it}$, are the fitted values of the fraction of dollar volume executed by dark trades with fractional price improvement and dark trades at the midquote, respectively. The set of control variables is as follows. $Time_t$ is a trend variable that starts at zero at the beginning of the sample period and increments by one 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 log of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by 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. Adjusted R's do not include the variance explained by the fixed effects (α_i). 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 sample comprises TSX Composite Index stocks during the period August 15, 2012 to December 15, 2012.

Variable	<i>Quoted Spread_{it}</i>	<i>Effective Spread_{it}</i>	<i>Realized Spread_{it}</i>	<i>ILLIQ_{it}</i>	<i>Depth_{it}</i>	<i>Autocorr- elation_{it}</i>	<i>HF Volatility_{it}</i>	<i>Delay_{it}</i>
<i>TwoSided_{it}</i>	-8.98 (-8.34)***	-16.60 (-9.24)***	-17.22 (-3.76)***	-1.03 (-4.77)***	-0.25 (-1.93)*	-0.10 (-5.14)***	-1.09 (-3.61)***	-4.25 (-0.69)
<i>OneSided_{it}</i>	0.16 (0.10)	-5.94 (-2.69)***	-4.69 (-0.85)	0.15 (0.62)	-0.18 (-1.35)	0.04 (1.69)*	-1.25 (-3.52)***	4.75 (0.81)
<i>Time_t</i>	-0.00 (-0.07)	0.00 (0.37)	0.01 (1.31)	-0.00 (-1.93)*	-0.00 (-4.88)***	-0.00 (-1.28)	0.00 (4.03)***	0.01 (0.84)
<i>US Mean_t</i>	0.10 (3.71)***	0.04 (1.20)	0.11 (1.09)	0.10 (3.02)***	0.49 (13.73)***	0.16 (3.88)***	0.24 (14.74)***	21.75 (7.67)***
<i>\$Volume_{it}</i>	-0.76 (-17.34)***	-0.14 (-1.65)*	0.95 (3.86)***	-0.47 (-57.23)***	0.15 (25.65)***	0.00 (0.97)	0.24 (21.38)***	0.65 (3.35)***
<i>Volatility_{it}</i>	47.42 (17.32)***	53.28 (11.90)***	-194.05 (-13.65)***	8.24 (18.95)***	-5.95 (-19.26)***	-0.04 (-1.09)	64.21 (53.51)***	-68.34 (-6.08)***
<i>Constrained_{it}</i>	-12.57 (-51.37)***	-9.13 (-25.90)***	-5.34 (-5.59)***	-0.19 (-5.33)***	0.87 (34.64)***	-0.02 (-6.45)***	-1.73 (-29.15)***	-2.18 (-2.37)**
Observations	19,233	19,233	19,233	19,233	19,233	19,205	19,233	19,205
Adjusted R ²	30%	9%	2%	25%	21%	4%	57%	6%
Fixed effects	Stock	Stock	Stock	Stock	Stock	Stock	Stock	Stock

Table 9

Further robustness tests separating two-sided and one-sided dark trading

This table reports coefficients of the key independent variables in the second-stage instrumental variables regressions, for a variety of different specifications. The independent variables to which the coefficient estimates correspond are the fitted value of a stock-day's two-sided and one-sided dark dollar volume ($\widehat{TwoSided}_{it}$ in Panel A and $\widehat{OneSided}_{it}$ in Panel B). 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 for the sake of comparison (the specification reported in Table 8, using TSX Composite Index stocks during the period August 15, 2012 to December 15, 2012). Specification (2) is identical to specification (1) except as the set of instrumental variables, it uses interactions of the post-regulation dummy variable and dummy variables for each stock. Specification (3) is identical to specification (1) except that it includes only cross-listed securities. Specification (4) is identical to specification (1) omitting two weeks either side of the introduction of the minimum price improvement rules. Specification (5) uses the US side of cross-listed stocks as the control group. Specifications (6) and (7) are estimated using the largest 121 and smallest 120 stocks, respectively, using the same variables as specification (1). Specifications (8) and (9) are estimated on the 68 stocks that are constrained by the tick size at least 90% of the time pre-regulation, and the 174 stocks that are constrained less than 90% of the time, 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)
Panel A: Two-sided dark trading coefficient estimates									
$QuotedSpread_{it}$	-8.98***	-13.06***	-9.35***	-15.57***	-9.03***	-4.83***	-11.04***	-3.85***	-11.96***
$EffectiveSpread_{it}$	-16.60***	-19.95***	-14.94***	-23.26***	-16.60***	-4.65***	-27.78***	-13.20***	-18.47***
$RealizedSpread_{it}$	-17.22***	-12.20***	-12.33**	-25.38***	-17.21***	-11.92***	-21.67***	-11.46	-20.42***
$ILLIQ_{it}$	-1.03***	-1.34***	-0.90***	-0.88***	-1.06***	-0.39**	-1.77***	-0.99***	-0.91***
$Depth_{it}$	-0.25*	-0.29***	-0.05	0.05	0.07	0.02	-0.58***	-0.40*	-0.23
$Autocorrelation_{it}$	-0.10***	-0.08***	-0.09***	-0.10***	-0.11***	-0.09***	-0.11***	-0.01	-0.18***
$HFVolatility_{it}$	-1.09***	-1.57***	-1.24***	-0.91**	-2.56***	-0.51*	-1.37***	-0.99*	-1.84***
$Delay_{it}$	-4.25	-5.52	-8.84	-5.88	-4.18	-1.02	-5.47	2.06	-7.70
Panel B: One-sided dark trading coefficient estimates									
$QuotedSpread_{it}$	0.16	1.49	0.38	0.61	0.19	1.51	-0.25	-0.64	0.45
$EffectiveSpread_{it}$	-5.94***	-4.99***	-6.09***	-4.56*	-5.95***	-2.98	-6.94*	-8.42**	-5.21**
$RealizedSpread_{it}$	-4.69	2.63	-6.69	-4.55	-4.61	5.57	-11.41	-8.06	-3.72
$ILLIQ_{it}$	0.15	-0.44**	-0.20	0.33	0.16	-0.10	0.38	0.91*	0.05
$Depth_{it}$	-0.18	-0.02	-0.24	-0.28*	-0.23*	-0.55***	0.10	-1.85***	0.12
$Autocorrelation_{it}$	0.04*	0.06***	0.04*	0.07***	0.04*	0.05*	0.03	0.07	0.03
$HFVolatility_{it}$	-1.25***	-0.62**	-1.26***	-1.62***	-0.77**	-0.78**	-1.43**	-2.69**	-1.15***
$Delay_{it}$	4.75	2.37	4.01	6.29	4.77	10.19	1.71	20.02	2.25
Observations	19,233	19,233	14,385	14,999	19,233	9,784	9,449	5,605	13,628

Table 10

The impact of dark trading on market quality in Australia

This table reports estimates for Australian stocks (ASX 200 Index constituents) for two months before and after the introduction of the minimum price improvement regulation on May 27, 2013. Estimates are from second-stage instrumental variables regressions:

$$y_{it} = \alpha_i + \beta_1 \widehat{Dark}_{it} + \sum_{j=1}^4 \gamma_j Control_{j,it} + \varepsilon_{it}.$$

The dependent variables, y_{it} , are estimates of informational efficiency for each stock-day. $Autocorrelation_{it}$ and $HFVolatility_{it}$ are absolute autocorrelations and standard deviations of ten-second midquote returns, respectively. $Delay_{it}$ 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 zero at the beginning of the sample period and increments by one every trading day. $\$Volume_{it}$ is the natural log of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by 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. Adjusted R^2 s do not include the variance explained by the fixed effects (α_i). Standard errors are clustered both by stock and date, t-statistics are reported in parentheses. ***, **, and * indicate statistical significance at 1%, 5%, and 10% levels, respectively.

	<i>Quoted Spread_{it}</i>	<i>Effective Spread_{it}</i>	<i>Realized Spread_{it}</i>	<i>ILLIQ_{it}</i>	<i>Depth_{it}</i>	<i>Autocorr- elation_{it}</i>	<i>HF Volatility_{it}</i>	<i>Delay_{it}</i>
\widehat{Dark}_{it}	-21.44 (-3.97)***	-12.32 (-1.22)	5.37 (0.42)	-1.63 (-6.84)***	0.85 (6.18)***	-0.01 (-0.38)	-9.90 (-3.14)***	3.79 (0.85)
$Time_t$	0.06 (7.40)***	0.05 (3.56)***	0.04 (2.25)**	-0.00 (-4.94)***	0.26 (26.67)***	-0.00 (-1.64)	-0.01 (-2.52)**	-0.00 (-0.58)
$\$Volume_{it}$	-3.65 (-8.30)***	-6.03 (-7.32)***	-2.73 (-2.32)**	-0.26 (-13.88)***	-4.03 (-15.34)***	0.01 (3.86)***	0.63 (2.26)**	0.59 (2.15)**
$Volatility_{it}$	163.62 (12.24)***	263.10 (13.49)***	-461.96 (-12.65)***	12.04 (20.05)***	0.70 (18.49)***	-0.45 (-10.41)***	152.16 (20.22)***	-13.86 (-2.07)**
$Constrained_{it}$	-66.06 (-29.51)***	-41.87 (-9.80)***	-23.81 (-4.64)***	-1.45 (-18.06)***	0.00 (1.10)	-0.07 (-9.05)***	-17.49 (-11.55)***	0.23 (0.23)
Observations	18,048	18,048	18,048	18,048	18,048	17,958	18,048	17,887
Adjusted R^2	26%	10%	6%	20%	22%	1%	16%	2%
Fixed effects	Stock	Stock	Stock	Stock	Stock	Stock	Stock	Stock

Table 11

The impact of different types of dark trading on market quality in Australia

This table reports estimates for Australian stocks (ASX 200 Index constituents) for two months before and after the introduction of the minimum price improvement regulation on May 27, 2013. Estimates are from second-stage instrumental variables regressions:

$$y_{it} = \alpha_i + \beta_1 \widehat{TwoSided}_{it} + \beta_2 \widehat{OneSided}_{it} + \sum_{j=1}^4 \gamma_j \text{Control}_{j,it} + \varepsilon_{it}.$$

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 each stock-day using hourly return and volume observations. $Depth_{it}$ is the natural log of the time-weighted quoted dollar depth at the NBBO. $Autocorrelation_{it}$ and $HFVolatility_{it}$ are absolute autocorrelations and standard deviations of ten-second midquote returns, respectively. $Delay_{it}$ measures intraday midquote return predictability using lagged market returns. The key independent variables, $\widehat{TwoSided}_{it}$ and $\widehat{OneSided}_{it}$, are the fitted values of the fraction of dollar volume executed by dark trades with fractional price improvement and dark trades at the midquote, respectively. The set of control variables is as follows. $Time_t$ is a trend variable that starts at zero at the beginning of the sample period and increments by one every trading day. $\$Volume_{it}$ is the natural log of traded dollar volume. $Volatility_{it}$ is the stock-day's high-low price range divided by 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. Adjusted R^2 s do not include the variance explained by the fixed effects (α_i). 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.

	<i>Quoted Spread_{it}</i>	<i>Effective Spread_{it}</i>	<i>Realized Spread_{it}</i>	<i>ILLIQ_{it}</i>	<i>Depth_{it}</i>	<i>Autocorr- elation_{it}</i>	<i>HF Volatility_{it}</i>	<i>Delay_{it}</i>
<i>TwoSided_{it}</i>	-15.25 (-3.32)***	-2.65 (-0.26)	2.71 (0.25)	-0.84 (-4.28)***	0.53 (4.72)***	0.00 (0.07)	-3.56 (-1.39)	0.60 (0.16)
<i>OneSided_{it}</i>	8.84 (1.18)	-1.90 (-0.20)	2.24 (0.12)	-1.74 (-5.59)***	0.54 (3.13)***	-0.03 (-0.98)	-8.10 (-2.09)**	14.61 (2.87)***
<i>Time_t</i>	0.06 (7.51)***	0.07 (3.59)***	0.04 (1.94)*	-0.00 (-2.05)**	0.29 (38.31)***	-0.00 (-1.13)	-0.00 (-0.85)	-0.01 (-1.48)
<i>\$Volume_{it}</i>	-4.56 (-12.80)***	-6.56 (-9.89)***	-2.52 (-2.52)**	-0.31 (-21.03)***	-4.51 (-18.46)***	0.01 (4.85)***	0.25 (1.11)	0.68 (3.57)***
<i>Volatility_{it}</i>	182.18 (14.33)***	272.43 (14.75)***	-465.43 (-13.22)***	12.78 (22.31)***	0.65 (15.83)***	-0.45 (-11.50)***	157.56 (22.18)***	-12.65 (-2.23)**
<i>Constrained_{it}</i>	-66.79 (-28.61)***	-41.68 (-9.63)***	-24.04 (-4.46)***	-1.28 (-15.28)***	-0.00 (-0.65)	-0.07 (-8.80)***	-16.73 (-10.76)***	-1.11 (-1.01)
Observations	18,048	18,048	18,048	18,047	18,048	17,957	18,047	17,886
Adjusted R ²	26%	10%	6%	20%	23%	1%	16%	2%
Fixed effects	Stock	Stock	Stock	Stock	Stock	Stock	Stock	Stock




Fig. 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 TSX Composite Index constituents from August 15, 2012 to December 15, 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 October 15, 2012.





Fig. 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 rules in Canada (indicated by the vertical bar). The sample comprises TSX Composite Index constituents from August 15, 2012 to December 15, 2012. *Two-sided* price improvement consists of trades executed with 10% and 20% price improvement relative to the prevailing NBBO (these trades are executed on MatchNow and Alpha Intraspread, respectively). *One-sided* consists of trades executed at the midpoint of the NBBO, and could be executed on any of the dark venues. *NBBO Trades* refers to trades larger than 50 Standard Trading Units or \$100,000 executed at the NBBO. Such trades only became valid after the introduction of the minimum price improvement rules.

Panel A: Distribution of dark trading (as % of total volume) across stock-days



Panel B: Distribution of changes in dark trading (as % of total volume) around the regulation





Fig. 3. Distribution of dark trading and changes in dark trading around the regulation. Panel A shows the distribution of dark trading (as a % of total dollar volume), $Dark_{it}$, across stock-days. The sample comprises TSX Composite Index constituents from August 15, 2012 to December 15, 2012. Values of $Dark_{it}$ greater than 50% have been aggregated into the 50% bucket. Panel B shows the distribution of changes in dark trading around the minimum price improvement regulation (the average of $Dark_{it}$ before the regulation minus the average after the regulation for each stock) across stocks.

Panel A: Distribution of changes in two-sided dark trading around the regulation



Panel B: Distribution of changes in one-sided dark trading around the regulation




Fig. 4. Distributions of changes in the two types of dark trading around the regulation. Panels A and B show the distributions of changes in the two types of dark trading around the minimum price improvement regulation (the average of $TwoSided_{it}$ (Panel A) and $OneSided_{it}$ (Panel B) before the regulation minus the average after the regulation for each stock) across stocks. The sample comprises TSX Composite Index constituents from August 15, 2012 to December 15, 2012.