Differential access to dark markets and execution outcomes

James Brugler and Carole Comerton-Forde*

March 28, 2022

Abstract

We compare dark pool trades across exchange-operated dark pools, where all trader types have equal access, and broker-operated dark pools, where brokers can restrict access to exclude certain traders, such as high frequency traders. Conditional on execution, trades on broker dark pools have less information leakage and adverse selection than trades on exchange-operated dark pools. Broker dark pools that do not allow high frequency traders have less information leakage than those that do. Differences in execution outcomes are concentrated in smaller trades. We conclude that the ability to segment order flow can improve execution outcomes for investors.

JEL Codes: G12, G14.

Keywords: Dark trading, execution outcomes, high frequency trading.

---

*We thank Cboe Australia for providing us with access to their broker-venue-trade data. We are grateful to Amber Anand, Jonathan Brogaard, Thomas Ernst, Terry Hendershott, Dominik Rösch, Yazid Sharaiha, Andriy Shkilko, Bart Zhou Yueshen and seminar participants at Hong Kong Polytechnic University and the Microstructure Online Seminars Asia Pacific for helpful comments. We also thank numerous industry participants for insights about institutional arrangements. Brugler is at the University of Melbourne (Email: james.brugler@unimelb.edu.au). Comerton-Forde is at University of Melbourne (Email: carole.comerton-forde@unimelb.edu.au). Comerton-Forde a Research Fellow at the Centre for Economic Policy Research, is an economic consultant for the Australian Securities and Investments Commission and an Academic Advisor to the Plato Partnership. Comerton-Forde gratefully acknowledges support from the Norwegian Finance Initiative. A previous version of this paper was titled "Differential access to dark markets and execution quality".

1 Introduction

Dark pools are a ubiquitous feature of modern equities markets. They account for 14.2% of trading in US equities and 6.7% in European equities.[1] Dark trading venues play an important role for institutional investors seeking to reduce information leakage and price impact (Norges Bank Investment Management, 2015). The academic literature typically assumes dark trading venues are homogeneous, however, in reality they differ on a range of dimensions including how prices are set, the prices where trades can occur, whether it is integrated with pre-trade transparent ("lit") order flow and arguably most importantly, who can access the venue.

The impact of heterogeneity in dark pools is relatively unexplored in the literature. A notable exception is Menkveld, Yueshen, and Zhu (2017) who consider differences in dark pools based on cost and immediacy, and demonstrate there is a pecking-order for execution venues with mid-point dark pools being at the top, non-mid-point dark pools in the middle, and lit markets at the bottom. We analyze another dimension of heterogeneity, namely access restrictions. We consider how execution outcomes are impacted by whether or not a venue is open to everyone or whether there are restrictions on the types of traders/investors that are permitted to trade in the pool.

We explore this issue in the context of the Australian equity market. Australian regulations restrict heterogeneity on some dimensions: dark trades can only be executed with price improvement (i.e. all venues are equivalent to Menkveld et al. (2017)'s mid-point venues). However, there are differences in access for exchange-operated vs. broker-operated dark trading venues. Dark trading on the two stock exchanges, the Australian Securities Exchange (ASX) and Cboe Australia, is accessible to all investor/trader types as exchanges are prohibited from imposing any access restrictions. In contrast, access to dark pools operated by brokers is typically only available to the customers of these brokers and the brokers may impose restrictions on the types of traders that have access to the pool. Some broker-

[1] Rosenblatt Securities, Let there be light, February 2022, US and European editions.

operated dark pools do not allow principal flow, high frequency traders (HFT) or electronic liquidity providers (ELP) to access their pool, while others allow customers to opt-out of interacting with these types of flow. Customers may choose to opt-out of this flow if they prefer to interact with only natural order flow and if they believe trading with HFT/ELP is more likely to result in information leakage and price impact. However, this choice reduces the liquidity available and the probability of execution.

Analysis by the Australian Securities and Investments Commission (ASIC) demonstrates that there are substantial differences in the level of HFT present in exchange- vs. broker-operated dark pools. For the first quarter of 2015 they report that that HFT accounted for 14.4% of turnover on the ASX dark pool, and 27.6% of hidden liquidity on Cboe. On average, HFT accounted for only 1.7% of turnover on broker-operated dark pools, ranging from 0.32% to 34% in individual pools.[2]

We explore the impact of differential access by asking three questions. First, are there observable differences in execution outcomes between exchange-operated dark pools with unrestricted access and broker-operated dark pools, where some trader categories cannot trade? Second, if there are observable differences, are these differences causal? Third, can any observable differences be attributed to variation in access by trader category across types of dark pools?

To answer these questions we analyze all dark trades in stocks in the ASX All Ordinaries Index from January 1, 2017 to September 30, 2019. The All Ordinaries Index is a market capitalization index of the 500 largest stocks trading on the ASX (S&P Global, 2020) and represents around 90% of the total value of securities trading on the ASX (Westpac, 2020). We categorize dark trades based on whether or not they take place in an exchange-operated dark pool ("exchange dark pool trades") or in a broker-operated dark pool ("broker dark pool trades") and estimate the execution outcomes of each trade in our sample. We do not observe unfilled orders so we measure execution outcomes conditional on execution.

[2]See ASIC (2015) for further details. ASIC also report that at the time their analysis was conducted three pool operators disclosed to their clients and to ASIC that there was no high-frequency trading in their pool.

Since the vast majority (92% in our sample) of dark pool trades take place at the mid-point of the National Best Bid and Offer (NBBO), trade direction cannot be identified and so standard market microstructure measures of execution outcomes (e.g. effective spreads, realized spreads or price impact) cannot be used. Without access to account level data, we also cannot compute other measures of execution outcomes such as implementation shortfall. We therefore propose four measures of execution outcomes for dark pool trades that can be applied to mid-point dark trades. These four variables capture information leakage, adverse selection risk, transitory price pressure and speed of price adjustment.

Using a panel regression approach with stock and date fixed-effects and controls for trade characteristics and the state of the limit order book, we show that broker dark pool trades have less information leakage and result in less adverse selection risk compared with trades on exchange dark pools. For price impact, statistically significant differences across venues are present from immediately after the trade until around five minutes later. For adverse selection, the effect is also present immediately after execution but is insignificant for horizons greater than one minute. We find evidence of lower transitory price pressure after broker dark pool trades but do not detect significant differences in the speed that prices adjust to the new efficient price. Together, these regressions suggest that broker dark pool trades have better average execution outcomes than exchange dark pool trades.

Figure 1 presents our information leakage results from these panel regressions in graph form. The y-axis is the estimated coefficient on the variable that captures the difference in absolute price impact between broker dark pool trades and exchange dark pool trades, conditional on controls and fixed effects. The x-axis is the horizon over which we estimate absolute price impact (from 500ms to 1800s). Broker dark pool trades have lower absolute price impact immediately after the trade, peaking at around 60s after the trade. This difference is statistically significant until at least five minutes after the trade takes place but is indistinguishable from zero at the 30 minute (1800s) horizon.

Figure 1 about here

Our panel approach conditions on a rich set of controls and fixed effects. However, this may not be sufficient to identify the causal effect of venue type on execution outcomes. Observed differences in average execution outcomes across exchange and broker dark pools may reflect strategic decisions about where to route different orders based on unobservable factors that vary at the trade-level, such as order informativeness or whether the order is part of a larger "parent" order.

Identifying the causal effect of venue type on execution outcomes requires a source of exogenous variation regarding where an order gets routed. Our solution exploits the fact that broker dark pools operated by three major broker-dealers cease operations during our sample. After their pool closes, these brokers have no choice but to execute dark pool trades on exchnage dark pools. Based on this, we assume that the sample of exchange dark pool trades from these brokers after pool closure contains some trades that would be executed on a broker dark pool if it was still operating. Using data that identifies the executing broker of all dark trades or, where possible, the trade venue of broker dark pool trades, we match broker dark pool trades from brokers who continue operating their dark pools with exchange dark pool trades from brokers who close their pools. We match trades within stocks across these categories based on trade and order book details, keeping only closely matched trades.

We compare average execution outcomes of the matched broker dark pool trades with the subset of the exchange dark pool trades from these brokers whose pools closed and are close matches for broker dark pool trades. The matching exercise confirms our main panel regression results: trades on broker dark pools have lower information leakage and result in less adverse selection risk for liquidity providers. We do not, however, find consistent evidence for differences in temporary price pressure or the speed of price adjustment.

Having documented significant differences in execution outcomes between exchange dark pools and broker dark pools, we next provide evidence that differential access does indeed explain these results. We do this in two ways: first by exploiting the heterogeneity in access across pools, and second by examining differences by trade size.

Broker dark pools display important heterogeneity in terms of the categories of traders allowed to access the pool. Some explicitly forbid trading from HFT/ELP, while others allow customers to "opt-out" of being matched with HFT/ELP order flow. If heterogeneity in access drives differences in execution outcomes between exchange dark pools and broker dark pools, then we also expect to observe better execution outcomes on broker dark pools that have the strongest restrictions on HFT/ELP flow. Our evidence directly supports this conjecture. Trades on pools that completely restrict HFT/ELP order flow have significantly lower information leakage and lower adverse selection risk than trades in pools that give customers the ability to opt-out of this flow. We also show that differences in information leakage are more pronounced for smaller trades (those in the bottom decile of size by stock-week) compared with larger trades. The types of traders whose activity we posit as being detrimental to overall execution outcomes (HFT and ELP) execute smaller trades on average. Our results split by trade size suggest that the differences in execution outcomes are concentrated in smaller trades and therefore are more likely to involve HFT/ELP counterparties.

Our results are consistent with the literature on how high frequency trading impacts institutional trading costs. Hirschey (2021) provides evidence that high frequency traders can anticipate order flow, and trade ahead of it. Korajczyk and Murphy (2018) and Van Kervel and Menkveld (2019) show that high frequency trading can increase the cost of trading for institutions when they trade in the same direction as institutions. Battalio, Hatch, and Saglam (2022) show that when marketable pieces of a parent order are routed to an ELP, parent orders exhibit larger implementation shortfalls. These results are consistent with our findings that limiting interactions with HFT/ELP improves execution outcomes.

Our results stand in sharp contrast to two recent papers on broker routing decisions. Anand, Samadi, Sokobin, and Venkataraman (2021) show that brokers that route more order flow to affiliated venues exhibit lower fill rates and higher implementation shortfalls. Battalio, Corwin, and Jennings (2016) find evidence that retail brokers route orders to maximize order flow payments, and that this is negatively associated with execution outcomes. Their results point to conflicts of interest in order routing decisions, while our results suggest that orders executed on broker-operated venues lead to better execution outcomes. Our results are not directly comparable to those in Anand et al. (2021) and Battalio et al. (2016) because our study considers execution outcomes conditional on orders being executed as we do not observe fill rates, while the earlier studies also observe unfilled orders. No similar data source exists for the Australian market, so a more direct comparison is not possible. Anand et al. (2021) examine differences between affiliated and unaffiliated pools, while we focus on differences in access restrictions. It is noteworthy that Anand et al. (2021) do not find substantial differences in their opacity measure, (which also proxies for access restrictions) between affiliated and unaffiliated pools.[3]

Our paper contributes to the nascent literature on heterogeneity in dark pools. Menkveld et al. (2017) is the only paper to consider heterogeneity in dark pools. Complementary to their study, our paper documents that for mid-point dark pools, access restrictions improve execution outcomes. It also builds on the literature on dark pools and segmentation of order flow. Comerton-Forde, Malinova, and Park (2018) find that a rule change that eliminated intermediation of retail orders in the dark, reduced order flow segmentation and enhanced lit liquidity. Hatheway, Kwan, and Zheng (2017) find that segmentation of uninformed order flow in dark pools harms liquidity providers in lit markets. In contrast, we show that segmentation can lead to improvements in execution outcomes for the investors trading in segmented venues. We attribute these improvements to the absence of, or reduction in, the likelihood of trading with high frequency traders/electronic liquidity providers.

We do not address the contentious issue of how "cream-skimming" and payment for order flow impact total market quality. Nor do we offer new insights on the impact of dark trading

[3]The opacity measure includes zero/one dummies for characteristics that make an ATS less opaque, such as whether the pool allows affiliated principal traders or external proprietary traders, and whether the pool is able to exclude counterparties.

on market quality, but rather complement the extensive literature already addressing these issues (Easley, Kiefer, and O'Hara, 1996, Bessembinder and Kaufman, 1997, Buti, Rindi, and Werner, 2011, Zhu, 2014, Comerton-Forde and Putniņš, 2015, and Brolley, 2020).

2 Institutional details

The Australian equity market is the fifteenth largest in the world, with approximately 2,200 listed companies and an average total market capitalization of between AUD 1.7 and 2.1 trillion between 2017 and 2019.[4] Trading activity is fragmented across two exchanges: the Australian Securities Exchange (ASX) and Cboe Australia (Cboe), and 13 broker-operated dark pools.[5] During our sample period ASX accounts for approximately 75.5% of trading by total dollar volume (including the opening and closing auctions), Cboe 10% and the broker dark pools 3.3%.[6] Trading is governed by the Australian Securities and Investments Commission (ASIC) Market Integrity Rules (MIRs).[7]

ASX operates two order books: TradeMatch which is a transparent limit order book and Centre Point which is a dark pool. TradeMatch operates on price-time priority, and Centre Point operates on time-priority with orders matched at the mid-point of the National Best Bid and Offer (NBBO) or on a specified tick size within the spread if the spread is greater than one tick. Traders can submit orders that sweep the two limit order books. Cboe operates a single electronic limit order book which allows both displayed and hidden orders to be submitted. Orders are matched based on price-display-time priority. All exchange order books are anonymous, but the brokerage firms associated with each trade are reported to the market on T+3. The minimum tick size varies with price. The tick is \$0.01 for stocks priced equal to or greater than \$2; \$0.005 from stocks priced equal to or greater than \$0.10

[4]Global market capitalization statistics come from the World Federation of Stock Exchange Fact Book and the ASX statistics from ASX (2021).
[5]The term "broker-operated dark pool" is not used in the Australian regulations. Instead the rules refer to a "crossing system" although colloquially they are referred to as broker-operated dark pools. We also use this term for consistency with the existing literature. ASIC provide a list of dark pools here.
[6]Off-exchange trading accounts for the remaining 10% of dollar volume.
[7]ASIC Market Integrity Rules (Securities Markets) 2017 are available here.

and less than $2.00; and $0.001 for stocks priced less than $0.001. Trading opens and closes with a call auction operated by ASX.

Unlike the US market, the Australian market does not have an order protection rule. Brokers have an obligation to provide best execution to their clients. These obligations differ for retail and institutional clients. For retail customers the broker must execute at the best price. For institutional customers brokers must have in place a best execution policy which defines the factors they consider when executing client orders which may include price, costs, speed, likelihood of execution or any other relevant outcome, or any combination of those outcomes. Australian regulations also prohibit all forms of payment for order flow including exchange rebates.

Orders must be transparent unless they meet one of the pre-trade transparency exceptions set out in the MIRs. One exception is for *Trades with price improvement*, which may be used for trades of any size, provided they offer price improvement relative to the NBBO.[8] These trades must occur at mid-point, or at a designated tick within the NBBO. We refer to these trades as NBBO trades. There are other exceptions for block and portfolio trades.[9] All trades executed away from the exchanges under these pre-trade transparency exceptions must be immediately reported to either ASX or Cboe. Most NBBO trades are matched on a broker dark pool, however, a very small number are matched manually by a sales trader.

Centre Point, Cboe hidden orders and NBBO trades operate under these pre-trade transparency exceptions. Centre Point and Cboe allows users to specify a minimum acceptable quantity (MAQ). Centre Point also allows the user to specify a single fill MAQ, where the order executes only when the MAQ is satisfied by a single opposing order. Brokers may also opt-in to preferencing arrangements which enables them to prioritize matching with their own orders.

[8] This is equivalent to what US markets refer to as a trade-at rule.

[9] Block trades are negotiated away from the market. The thresholds for block trades are based on the stock's average daily trading volume: AUD 1m for the most active stocks; AUD 500k for the next most active stocks and AUD 200k for the least active stocks. *Portfolio trades* must have a transaction value of at least AUD 5m, across at least 10 different stocks and each individual stock is for at least AUD 200k.

Centre Point and Cboe hidden orders are required to offer unrestricted access to all trader/investor types. Broker dark pools however, are permitted to restrict access provided they do not unfairly discriminate between users. The extent of restrictions in place varies across different dark pools. Table 1 provides a summary of the access restrictions in each pool during our sample period.[10] Four pools completely restrict access to HFT/ELP and principal trading and nine pools allow customers to opt-in to restricting access by counter-party type. All broker operated pools also offer MAQ functionality, usually on an order-by-order basis.

Table 1 about here

3 Data description

We obtain limit orders at the NBBO and trades in all stocks in the ASX All Ordinaries Index over the period from January 1, 2017 to September 30, 2019 from Refinitiv's Datascope. Our focus is on dark trading, so we filter our data to identify all trades that utilize the pre-trade transparency price improvement exception. These include trades executed on Centre Point, hidden liquidity on Cboe and NBBO trades reported to both ASX and Cboe. During our sample period, they account for 7.6%, 1.8% and 3.3% of market share, respectively.

For every trade we observe the price, volume, the time to the nearest millisecond and a trade qualifier designating the type of trade. The Refinitiv data do not identify the venue where NBBO trades are executed and do not directly distinguish between broker dark pool trades and manual executions. The MIRs require brokers to report all off-exchange trades to either ASX or Cboe, immediately.[11] The brokers executing the trade and the execution venue where the trade occurred are not reported in real-time, but are included in the course of sales data on T+3. We obtain the broker and venue information for broker dark pool trades from two sources. For trades reported to Cboe we obtain data directly from Cboe

[10] Dark pool classifications are made based on regulatory disclosures, with missing documentation confirmed through discussions with dark pool operators.

[11] In our sample, 70% of NBBO trades and 55% NBBO dollar volume is reported to Cboe.

that identifies both the broker and the venue for the trade.

Two of the thirteen brokers operating dark pools report these trades to ASX. Data on the execution venue are available from ASX, but at considerable cost.[12] We therefore use a work-around where we match the Refinitiv trade records to broker trade data available from Rozetta. We identify broker dark pool trades using a combination of the NBBO trade identifier and the broker identifier. Although this classification process also includes a small number of manually executed trades, our proprietary Cboe data directly identifies the venue and so allows us to make inference about the relative prevalence of manual executions compared with true dark pool trades. These data show that less than 0.5% of NBBO trades are manual matches and therefore these trades do not meaningfully affect our results.[13] We also obtain both the buy and sell brokers responsible for each Centre Point execution from the Rozetta data. We refer to Centre Point and Cboe hidden trades and broker dark pool trades collectively as "dark pool trades".

Market shares for the three dark pool types are presented in Figure 2. Centre Point has the largest market share throughout our sample representing between 6-9% of total trading in All Ordinaries stocks (by dollar volume). Broker dark pool trades account for around 2.5-4.5% of total dollar volume traded. Cboe hidden trades are between around 1-2% of total trading.

Figure 2 about here

Figure 3 presents the average dollar trade size and the distribution of dollar trade size across the three dark pool trade categories in our sample. Panel (a) presents average trade sizes which follow quite similar downward trends across all three categories and are between $2,000 and $3,000 across most of the sample. Centre Point trades are slightly larger on

[12] The cost of accessing these data for our three year data sample is AUD 41,210.

[13] Manual matches are also substantially larger than broker dark pool trades on average. As a robustness check we exclude all NBBO trades on the ASX in excess of $50,000 from our sample, as these are more likely to be manual executions. Our results are not sensitive to this filter, nor the use of filters at lower thresholds (see Section 6 for details).

average than broker dark pool and Cboe hidden order trades. Panel (b) presents trade size distributions. Most trades are for \$5,000 or less in all three venues and less than 1% of all trades are for \$50,000 or more in all venues. Patterns across the venue categories are quite similar other than that there are around 10 percentage point fewer very small trades (\$100 or less) on Cboe hidden compared with Centre Point or broker dark pools and there are slightly more large trades (\$10,000 or above) for Centre Point compared with Cboe or broker dark pools.

Figure 3 about here

Our order data contains the price and depth available at the NBBO, time-stamped to the millisecond. We match these to the trade data to identify the mid-quote immediately prior to and at various intervals after the trade. We also calculate daily level liquidity summaries such as the daily time-weighted average bid-ask spread and dollar depth on the limit order book by stock-day, summarized in the next section.

3.1 Measuring execution outcomes for dark pool trades

Estimating execution outcomes for dark pool trades is complicated by the fact that many standard market microstructure variables (such as effective spreads, realized spreads and price impact) require trades to be assigned as either buyer-initiated or seller-initiated. For dark pool trades, assigning trade direction is problematic. The vast majority (92% in our sample) of dark pool trades take place at the mid-point, so neither counter-party is fully crossing the spread to trade.

Although dark pool trades differ from a standard execution on a limit order book, they can still affect the evolution of future prices and liquidity. Large institutional trades are often split into many sequential "child" orders that are submitted according to an execution algorithm (Van Kervel and Menkveld, 2019). When a dark pool trade involves a child order on one side of the trade, the counter-party to the trade can learn information about future order flow in the same direction, and therefore future prices. Evidence from Hirschey (2021) shows the presence of predatory traders (often high frequency traders) who "fish" for the presence of such institutional orders and then step ahead of the order by subsequently trading in the same direction. This raises the expected execution costs for the unfilled portion of the institutional order. Korajczyk and Murphy (2018) decompose trading costs into a spread component and price impact component and highlight that price impact is an important component of institutional order trading costs. In our context, this is particularly true because the spread component likely plays little role given that the majority of dark pool trades are executed at the prevailing mid-point of the NBBO in both types of dark pools. Further, Ye and Zhu (2020) show that informed traders have an incentive to route orders to dark pools to limit their direct and indirect execution costs.

The existence of a dark pool trade may convey information to other liquidity providers not involved in the trade that either an informed trader or an institutional-sized trade is present in the market, even if these liquidity providers are unable to discern which side of the trade is informed or large. These traders may adjust their perceptions of adverse selection or inventory risk subsequent to observing trades, leading to wider bid-ask spreads.

Our analysis focuses on dependent variables that capture the effect of dark pool trades on the evolution of future prices or the state of the limit order book. First, we calculate the absolute price impact of each trade as:

$$AbsPI_{ijt} = 100 \times \left|\log\left(M_{ijt}^{s+\tau}\right) - \log\left(M_{ijt}^{s}\right)\right| \tag{1}$$

where $M_{ijt}^{s+\tau}$ is the mid-point of the i^{th} trade in stock j at time s on day t after an interval of τ seconds after the trade time and M_{ijt}^{s} is the prevailing mid-point immediately prior to the trade. Our logic in constructing Equation (1) is that, on average, large values of $AbsPI$ suggest a higher degree of information revelation resulting from the mid-point execution. We therefore interpret this variable as capturing "information leakage" from dark pool trades, which constitutes a substantial cost for investors splitting a large parent order into many smaller "child" orders. We estimate Equation (1) over periods of 500ms, 1s, 10s, 30s, 60s, 300s and 1800s after the trade, which captures effects at both short horizons and long horizons. Importantly, these intervals contain the recommended maximum horizons for capturing price impact in large and small stocks in modern markets according to Conrad and Wahal (2020).

Second, for each trade in our sample, we calculate the percentage bid-ask spread at the same post-trade intervals as we use for absolute price impact. Larger values of bid-ask spreads after the trade takes place indicates relatively worse liquidity after execution and higher future trading costs. As above, wider bid-ask spreads after mid-point crosses can reflect updated expectations of information or liquidity risk for liquidity providers.

We also calculate a variable capturing the degree to which future prices are subject to reversals after a trade. $Reversal$ is an indicator variable taking the value of one if the sign of the mid-point returns from trade time to one minute after trade time is the opposite to that of the mid-point return from one minute to thirty minutes after trade time:

$$Reversal_{ijt} = \begin{cases} 1 \text{ if } sign(r_{ijt}^{s \to s+60s}) \neq sign(r_{ijt}^{s+60s \to s+1800s}) \\ 0 \text{ if } sign(r_{ijt}^{s \to s+60s}) = sign(r_{ijt}^{s+60s \to s+1800s}) \end{cases} \tag{2}$$

where $r_{ijt}^{s_1 \to s_2}$ is the mid-point return for the i^{th} trade in stock j on day t over the interval from s_1 to s_2. Intuitively, this variable captures temporary price pressure in one direction or another following a dark pool trade which we argue reflects temporary periods of poor liquidity and low levels of immediacy.

Finally, we calculate the speed of price adjustment by comparing the mid-point one minute after the trade with the mid-point thirty minutes after the trade:

$$AdjustmentSpeed_{ijt} = 100 \times \left| \frac{M_{ijt}^{s+1800s} - M_{ijt}^{s+60s}}{M_{ijt}^{s}} \right| \tag{3}$$

where M_{ijt}^{s} is defined as per Equation (1). Underpinning the logic of $AdjustmentSpeed_{ijt}$

is the assumption that the new equilibrium price following a trade is, on average, reached no later than 30 minutes after a trade is executed. Equation (3) therefore measures the percentage difference between the new equilibrium price and the price one minute after a trade. *AdjustmentSpeed* is fast when the mid-point price one minute after the trade is close to the new equilibrium price, and therefore its value is close to zero, but is slow when there is a large difference between these prices. Absolute price impact, bid-ask spreads and adjustment speed are all non-negative variables and are winsorized at the 99th percentile by week.[14]

4 Execution outcomes in exchange and broker dark pools

Our primary goal is to determine whether there are causal differences in execution outcomes for exchange and broker dark pool trades. Achieving this is complicated by the fact that venue choice is a strategic decision made by investors or their brokers. Observed differences in average execution outcomes across venue categories may not necessarily reflect actual differences in average execution outcomes but instead reflect differences in the average characteristics of orders that are submitted across venue categories.

For example, any trader can submit orders to exchange dark pools, but only a brokers' customers are able to trade on each broker dark pool. Trades from the customers of brokers who operate dark pools may differ systematically from the rest of the market. These trades may be more likely to form part of a large institutional trade, may be less likely to be from a retail trader, or may contain more information than the market average. These kinds of trades will likely have greater absolute price impact regardless of where they are executed. Differences in trading fees across venues may also drive differences in the types of orders submitted. For example, trader types that submit many small orders (such as high frequency traders) may choose to avoid trading on high cost venues.

We deal with this issue in two ways. First, we estimate the effect of execution venue in

[14]Specifically, for each week, we replace the largest percentile of values for each variable with the 99th percentile.

a panel regression that includes a rich set of controls and fixed effects. These regressions allow us to form inference regarding execution outcomes across venue types while controlling for observable order characteristics like trade size, trade price and liquidity at the time of execution, as well as unobservable components at the stock or date level via fixed effects.

These regressions cannot control for differences in unobservable order characteristics, such as whether the trade is part of a large institutional order or the order reflects a trader's private information. Our strategy for dealing with unobservable differences in order characteristics exploits the closure of three broker dark pools over our sample period. We view these pool closures as exogenous events that shift order flow from broker dark pools to exchange dark pools: dark orders sent to brokers whose pools have closed will necessarily execute in an exchange dark pool. We match these trades with broker dark pool trades from brokers whose pools remain in operation and test for differences in execution outcomes.

4.1 Execution quality panel regressions

We form a panel of all dark pool trades in our sample. For every trade we observe our execution outcomes variables described in Section 3, plus trade size and price in AUD, the bid-ask spread and depth at NBBO immediately prior to the trade and whether or not the trade is executed in an exchange dark pool (Centre Point or Cboe hidden) or a broker dark pool. At the trade level, the data generating process we are interested in estimating is described by:

$$y_{ijt} = \alpha_j + \gamma_t + \beta BDP_{ijt} + \rho' X_{ijt} + \varepsilon_{ijt} \tag{4}$$

where y_{ijt} is the execution outcome variable for dark pool trade i in stock j in day t, α_j is a stock fixed effect, γ_t is a date fixed effect, BDP_{ijt} is a dummy variable taking the value 1 if this trade takes place on a broker dark pool and 0 if it takes place in an exchange dark pool, X_{ijt} is a vector of other controls and ε_{ijt} is an error term. The parameter β captures the average difference in execution outcomes after controlling for controls and fixed effects.

Implementing a regression of the form in Equation (4) is complicated by the very large number of dark pool trades that take place in our universe of stocks over the sample period (approximately 185 million trades across all dark pools). To deal with this, we average the trade-level data by stock and day and run regressions on the averaged data:

$$\bar{y}_{jt} = \alpha_j + \gamma_t + \beta B\bar{D}P_{jt} + \rho' \bar{X}_{jt} + \bar{\varepsilon}_{jt} \tag{5}$$

where $\bar{v}_{jt} = \frac{1}{N}\sum_{i=1}^{N_{jt}} v_{ijt}$ is the average of the variable v_{ijt} across all i trades for stock j and day t. Note that $\bar{\alpha}_j = \alpha_j$ and $\bar{\gamma}_t = \gamma_t$. Equation (5) follows directly from taking the stock-day average of the data generating process in Equation (4), implying that we can recover the parameters of Equation (4) from a regression of the stock-day average of execution outcomes for all dark pool trades onto the fraction of dark pool that are taking place on broker dark pools and the stock-day averages of the control variables. In other words, replacing the trade venue dummy with the fraction of broker dark pool trades identifies β and the inference we get will be the same for large samples.

A key difference between exchange and broker dark pools is that orders submitted to exchange dark pools can interact directly with displayed liquidity on the corresponding lit order books. For Cboe this is because the lit and dark books are integrated, so marketable orders submitted to Cboe automatically execute against any price-improving hidden liquidity unless the trader specifically declines to do this. On the ASX, this occurs through the use of sweep orders which first execute against any dark liquidity resting in the Centre Point order book and then executes any unfilled portion against TradeMatch. Broker dark pools are not connected to a lit limit order book in the same way. Therefore, it is possible that differences in execution outcomes between broker and exchange dark pool reflect the fact that many observed exchange dark pool trades are actually sweeps of the lit book, and may move the best price. To ensure our results are not biased in such a way, we filter the trades in our panel analysis comparing exchange and broker dark pool trades to eliminate any dark pool trade that has a lit order that executes in the exact same millisecond.[15]

Table 2 contains summary statistics for our stock-day panel formed from all dark pool trades in our sample. The unit of observation is the stock-day level and every stock-day with at least one dark pool trade is included in the sample. The average dark pool trade size is AUD 2,030 but, as expected, the distribution of trade size is heavily right skewed. The median trade is for around one-third of this amount, AUD 1,000, while the largest trade in our sample is for AUD 3.4m. The average bid-ask spread at the time of the trade is 0.44%, indicating that an average dark trade saves around 0.22% to both parties in direct trading costs compared with both traders crossing the spread. The absolute percent change in mid-points from immediately before to after a dark pool trade ranges from one bp at the 500ms horizon, six bps at the one minute horizon to 33 bps at the 30 minute horizon. On average, the bid-ask spread widens by one bp 500ms after a dark pool trade to 0.45%, by three bps one minute after the trade to 0.47% and another 21 bps to 0.68% thirty minutes after the trade. Twelve percent of trades exhibit price reversals in the thirty minutes after the execution for the average stock-day in our sample, as defined in Equation (2). On average, around 69% of the mid-point price movement in the thirty minutes after the trade is realized within one minute of the trade.

Table 2 about here

Table 3 contains results from our stock-day panel regressions for absolute price impact across the different time horizons spanning 500ms to 1800s. These regressions include stock and day fixed effects as well as stock-day averages of log trade size, log price, log number of dark pool and lit trades, log of total volume traded, the prevailing bid-ask spread and log of depth as controls. All standard errors are clustered at the stock level.

[15]Results that do not eliminate sweeps are presented with other robustness tests in Section 6.

Table 3 about here

Broker dark pool trades have lower absolute price impact at all horizons from 500ms and 300s. The size of the effect is around -1 bp over the first 60s before falling to around -0.8 bps at 300s. All effects are significant at better than the 1% level over these horizons. At the 1800s horizon, we detect no statistically significant difference between broker dark pool and exchange dark pool trades. Figure 1 presents these estimates and upper and lower confidence intervals by time horizon graphically. The size of the difference in absolute price impact between broker dark pool and exchange dark pool trades is monotonically increasing from the shortest horizon (500ms) until one minute after the trade, after which the effect attenuates.

Abstracting momentarily from possible endogeneity that is not addressed by fixed effects alone, these regressions demonstrate that there is substantially less information leakage from trades on broker dark pools compared with exchange dark pools in the period immediately after the trade takes place up until at least five minutes after the trade. The fact that we detect no significant difference at the 30 minute horizon is important because it suggests that, in the long-term, there are no differences in the total amount of information contained in trades across the two venues, just the speed at which this information is impounded into prices.

Regarding other variables in our preferred specification, average dark pool trade price impact is generally higher for stocks-days with lower depth, higher average bid-ask spreads, more total trading volume, lower prices, more trading activity in the limit order book and less trading activity in dark pools, depending on the time horizon. Absolute price impact is decreasing in average dark pool trade size, conditional in all other controls and fixed effects.

Table 4 contains analogous regressions to Table 3 but where the dependent variable is the bid-ask spread at the various horizons after dark pool trade execution. Similar to our results for absolute price impact, we detect a statistically significant reduction in bid-ask spreads from 500ms to 60s after broker dark pool trades compared with exchange dark pool trades.

The size of the effect is smaller than is estimated for absolute price impact, both absolutely, and relative to sample-wide standard deviations. Further, the effect is insignificantly different from zero by 300s, while the effect on absolute price impact remains significant at that horizon. Nevertheless, evidence from bid-ask spreads is consistent with that of absolute price impact insofar as less information leakages maps closely to less perceived adverse selection risk from market makers. At the 30 minute horizon, the bid-ask spreads after broker dark pool trades are larger than for exchange dark pool trades. Noting that, at the same horizon, we find insignificant differences in absolute price impact, wider spreads after broker dark pool trades is consistent with market makers learning more slowly from these trades compared with those on exchnage dark pools.

Table 4 about here

Results for the speed of adjustment and price reversals are presented in Table 5. For the speed of price adjustments, the point estimate on the broker dark pool variable is positive, indicating that price adjustments to the new equilibrium price are slower for broker dark pool trades, however the effect is insignificant at the 10% level. Price reversals are slightly less common for broker dark pool trades with an estimated coefficient of -0.0069 and a t-statistic of -3.65. Economically, the magnitude of this effect is quite small, with broker dark pool trades approximately 0.7 percentage points less likely to result in price reversals.

Table 5 about here

Our stock-day panel regressions deliver a new and important insight regarding execution outcomes of broker dark pool trades compared with exchange dark pool trades: broker dark pool trades have significantly less information leakage than exchange dark pool trades from immediately after the trade takes place up until five minutes after the trade is executed. In the long-term there are no discernible differences. The reaction of liquidity providers to broker dark pool trades and exchange dark pool trades respectively, reflect this slower transmission of information. Spreads are relatively wider in the first 60s after exchange dark pool trades compared with the same period for broker dark pool trades. By 30 minutes, there is some evidence that the effect reverses. Finally, prices are slightly less likely to exhibit reversals after broker dark pool trades compared with exchange dark pool trades.

4.2 Matching trades around broker dark pool closure

The stock-day panel analysis in Section 4.1 suffers from one important limitation: we cannot rule out the existence of endogeneity between the error term ε_{jt} and our regressors of interest. Although our fixed effects account for time-invariant endogeneity at the stock level, or endogenous market-wide shocks that affect all stocks at a given date, we cannot consistently estimate our coefficients in the presence of endogeneity that varies within stocks over time. For example, brokers may prefer to route institutional trades to broker dark pools compared with retail traders. Trades from these groups may differ in ways that are not captured by our fixed effects and controls.

To deal with this issue, we require a source of exogenous variation in whether a dark pool trade is executed on a broker dark pool or an exchange dark pool. Our solution exploits the closure of three broker dark pools over our sample period by Bank of America Merrill Lynch (March 6, 2017), UBS (April 1, 2019) and Citigroup (July 1, 2019). Each closure represents an event where a dark pool ceases executing trades. After these dates, the broker who previously operated a dark pool has no choice but to route non-displayed orders to an exchange dark pool. A subset of exchange dark pool trades from these brokers in the period after pool closure would previously have been executed on a broker dark pool.

Figure 4 plots the trading activity (number of trades) in all three broker dark pools over our sample period. Each pool executes a significant number of trades up until the closure date, after which trading activity falls to zero, as expected. The lack of a downward trend in activity ahead of each closure suggests that customers appear to route a similar number of orders to these brokers despite the impending closure of the dark pool. Discussions with broker dark pool operators confirmed that institutional clients invariably leave the details of the execution strategy of an order to the discretion of the broker and typically do not specify a preference for execution to take place in a particular venue. Further, the presence of a broker dark pool is not considered a relevant factor in the institutions deciding which broker receives an order. Importantly for our purpose, brokers indicated that, order flow that previously would have been routed to the broker dark pool is routed to exchange dark pools following a broker dark pool closure.[16] Our matching approach is designed to isolate these trades from brokers after their dark pool closes and compare these with actual broker dark pool trades.

Figure 4 about here

We form a sample capturing dark pool trading in each of the three one month periods following each pool closure. The sample contains all broker dark pool trades from brokers whose pool is still operating combined with all Centre Point trades from brokers whose pools are closing in each event (a "closed-pool broker"). We focus on Centre Point due to its similarity to broker dark pools in terms of the way that orders interact with liquidity on the main order book and again eliminate sweeps from the analysis. We then perform a matching exercise where each broker dark pool trade is matched to a Centre Point trade in the same stock from a closed-pool broker and use these to estimate a treatment effect for the effect of execution outcomes in broker dark pools vs. exchange dark pools. Table 6 presents summary statistics for the sample of these trades, pooled across all three events. As well as the means, standard deviations and percentiles presented in Table 2, Table 6 additionally presents means split by trades on broker dark pool vs. Centre Point trades from closed-pool brokers.

Table 6 about here

[16] One broker noted that any liquidity provision in the pool would likely disappear from the market when the pool closed.

There are approximately 6.9 million trades in total in our matching sample, comprising of 5.7 million broker dark pool trades from all brokers whose dark pools continue to operate and 1.2 million Centre Point trades from the three brokers with recently closed dark pools. The average trade size in the matching sample is approximately $1,800 though trades executed on broker dark pools tend to be around $1,000 smaller than those executed on Centre Point. There are only small differences in average trade price, total daily dollar value traded, absolute price impact, reversals or price adjustment by category, though the bid-ask spread prior to and after a Centre Point trade from a closed-pool broker is approximately 2 to 3 bps wider on average, compared with a broker dark pool trade. Compared with the stock-day panel summary statistics presented in Table 2, bid-ask spreads and absolute price impact are both significantly smaller while average dollar value traded is larger. This largely reflects the fact that the summary statistics in our matching sample (Table 6) are constructed at the trade level whereas the stock-day summary statistics are computed at the stock-day level.[17]

We match each broker dark pool trade to a Centre Point trade from a closed-pool broker via propensity score matching. For each stock and event in our sample, we estimate the propensity score for the broker dark pool trade execution dummy variable where the explanatory variables are the log of dollar trade size, log of execution price, the bid-ask spread preceding the trade, the log of dollar depth available prior to the trade, the total dollar value executed on the day of the trade, the number of days from pool closure (as an integer) and its square and the time of day as a expressed as decimal and its square.[18] We then find the nearest neighbor for each broker dark pool trade from the set of Centre Point trades from closed-pool brokers, matching with replacement. We keep only trades that can be matched using a caliper of 0.25 standard deviations of the estimated propensity scores. By restricting our sample in this way, we aim to construct a control group as the subset of Centre Point

[17]Trade-level summary statistics are presented separately for each event in Tables IA.1-IA.3 in the Internet Appendix that can be downloaded here.

[18]The date from closure integer variable controls for a general time trend. The time-of-day variable records the number of seconds in the day from 12:00AM onward and controls for intraday patterns in execution outcomes.

trades from brokers whose pools close that plausibly would have executed on a broker dark pool if the broker dark pool still operated. We then estimate the difference in means for the broker dark pool trades and the matched Centre Point trades.

Our matching method closely follows Arpino and Cannas (2016) who present propensity score matching methods for clustered data. Formally, let N denote dark pool trades indexed by $i = 1, 2, 3, ..., n_j$ within J stocks indexed by $j = 1, 2, 3, ..., J$ in the one month period following each pool closure event. The binary variable T_{ij} indicates whether the trade took place on a broker dark pool ($T_{ij} = 1$) or Centre Point ($T_{ij} = 0$) and Y_{ij} denotes the outcome variable (e.g. absolute price impact of the trade). $Y_{ij}(t)$ denotes the potential outcome if trade ij takes place on the treated or control venue corresponding to $t \in 0, 1$. The average treatment effect on the treated (ATT) is given by the usual form:

$$ATT = \mathbb{E}\left[Y_{ij}(1) - Y_{ij}(0)|T_{ij} = 1\right]. \tag{6}$$

The ATT requires the standard assumptions of unconfoundedness and overlap to be identified, which Rosenbaum and Rubin (1983) show allows for matching on the propensity score (for example as estimated using a logit model for T estimated using the set of controls X), assuming this model is appropriately specified. In our case, we are concerned with appropriately accounting for potential unobserved effects at the stock level (clustering) as well as allowing for heterogeneity in the parameters in the propensity score across stocks.

We estimate a logit model for trade venue (T_{ij}) separately for each stock in our sample for each event. For each broker dark pool trade $Y_{ij}(1)$ we then find its nearest neighbor, $\tilde{Y}_{ij}$:

$$\tilde{Y}_{ij} = \left\{ kj \in I_0 : \hat{e}_{kj} = \min_{kj \in I_0} \left[\hat{e}_{ij} - \hat{e}_{kj}\right] < 0.25\hat{\sigma}_{\varepsilon_j} \right\} \tag{7}$$

where I_0 represents the set of all Centre Point trades, $\hat{e}_{kj}$ is the fitted propensity score for trade k obtained from the logit model estimated for stock j and $\hat{\sigma}_{\varepsilon_j}$ is the standard deviation of the fitted propensity scores for stock j. We then form the matched dataset, M, containing all broker dark pool trades that have a matched Centre Point trade within the caliper and their respective matches:

$$M = \left\{ ij : \tilde{Y}_{ij} \neq \emptyset \right\} \cup \left\{ \bigcup_{ij} \tilde{Y}_{ij} \right\} \tag{8}$$

and estimate the ATT via:

$$A\hat{T}T = \frac{1}{N_m} \sum_{ij \in I_1 \cap M} \left(Y_{ij} - \tilde{Y}_{ij} \right) \tag{9}$$

where N_m is the number of elements in the set M and I_1 is the set of all broker dark pool trades. We only include trades in stocks that record at least 100 dark pool trades in the month following pool closure and standard errors are clustered at the stock level. We estimate Equation (9) for each event separately.

By estimating propensity scores for each stock separately, and then matching trades within stocks based on these propensity scores, we account for the existence of group-level unobserved effects and heterogeneity in the determinants of trade venue across stocks. Limiting our set of control trades to those with a broker with a recently closed pool on one side, helps control for strategic order placement at the broker level and recover the unconfoundedness assumption. In addition, our matching analysis gives equal weight to all broker dark pool trades (so long as an adequate match can be found in the control sample) whereas our panel analysis gives equal weight to each stock-day combination. Estimates of the ATT for each event are presented in Table 7. We present results using a horizon of 60s, corresponding to the horizon with the largest effect for price impact from our panel regressions.

Table 7 about here

Our matching analysis confirms that trades on broker dark pools have significantly less absolute price impact and adverse selection risk in the period after trade execution. The ATT for price impact is between -0.31 bps and -0.84 bps across the three events and the magnitude of the t-statistics are well above the 1% threshold in all three cases. Under our panel analysis, we estimate the effect of a trade on a broker dark pool of -1.17 bps. For bid-ask spreads, our ATT estimates are between -0.16 bps and -2.64 bps and are significant at the 1% level for all three events, compared with -0.53 bps at the same horizon in our panel analysis. Our matching analysis does not generate strong evidence regarding price adjustments and price reversals. For price adjustments, the effect is positive and significant for the BAML event, negative and insignificant at the 10% level for the UBS event and negative and significant at the 10% level for Citi. Only the BAML event is significant at the 10% level or better for price reversals, though the effect is positive.

Overall, our treatment effects estimated under our matching procedure are consistent with our stock-day panel regarding information leakage and adverse selection risk. We obtain consistent estimates when using a focused sample that concentrates on the one month period after dark pool closures, weights all trades equally and uses a matching method that allows for substantial heterogeneity across stocks. This gives us a high degree of confidence in our main finding from Section 4.1: information leakage and adverse selection risk from trades on broker dark pools are less than for exchange dark pool trades.

4.3 Do broker trades on exchnage dark pools change after pool closure?

One way to test if order flow from customers changes in non-random ways after pool closure is to examine whether execution outcomes for exchange dark pool orders from brokers change in systematic ways after their dark pool is closed. Significant post-pool closure differences in execution outcomes, for example, changes to average price impact for exchange dark pool trades after a broker's dark pool is closed, could suggest that certain traders tend to route fewer orders to the broker after the pool closes. Alternatively, such differences in execution outcomes after pool closure could reflect broker preferences regarding where they route different customers' orders when they previously had the option to choose between a broker dark pool or an exchange dark pool. Our matching approach is robust to these selection issues if selection is sufficiently well predicted by observable order characteristics, but is less robust if there is a substantial unobserved component.

To test this, we need to isolate the causal effect of a broker's pool closure on execution outcomes for that broker's trades on Centre Point. We achieve this via a difference-in-differences regression estimated using (i) Centre Point trades from the broker whose dark pool closes and (ii) Centre Point trades from all brokers who continuously operate a dark pool over one month windows around the pool closure date. Trade outcomes for the other brokers operate as a control sample for the closing broker. We again focus on Centre Point rather than Cboe hidden liquidity because of its similarity in matching process with a typical broker dark pool. The regression is given by:

$$y_{ijbt} = \alpha_j + \gamma_t + \mu_b + \beta\tau_{bt} + \rho' X_{ijbt} + \varepsilon_{ijbt} \tag{10}$$

where y_{ijbt} is the 60s absolute price impact or bid-ask spread for trade i in stock j by broker b on date t, α_j is a stock fixed effect, γ_t is a date fixed effect, μ_b is a fixed effect for each of the brokers closing their dark pools, τ_{bt} is a treatment status indicator taking the value of one if the trade is from a closing broker after their dark pool closes, X_{ijbt} is a vector of controls and ε_{ijbt} is an error term. The controls are identical to those used in our stock-day regressions and all standard errors are clustered at the stock level and we estimate Equation (10) separately for each event.

Intuitively, our difference-in-differences approach categorizes exchange dark pool trades from brokers whose dark pools close in our sample period as the "treated" category and trades from all brokers who can execute trades in an exchange dark pool or their own dark pool during our sample period as the "control" category. The parameter β in Equation (10) then estimates the relative change in execution outcomes for brokers after their pools close compared with other brokers whose pools continue to operate. If this parameter is significant, it suggests that the trades now being routed to an exchange dark pool by brokers whose dark pools have closed are affected by pool closure in ways that other broker trades do not.[19] Table 8 contains these parameter estimates for each event.

Table 8 about here

For two of our three closure events, UBS and Citi, the treatment effect is insignificant at the 5% level for both absolute price impact and bid-ask spreads. Only one treatment effect is significant at the 10% level for these two closures — bid-ask spreads for the UBS event. For our other event, BAML, we detect a statistically significant decrease in absolute price impact and an increase in bid-ask spreads. This suggests that the closure of this pool resulted in changes in that broker's order flow that may invalidate our identification assumptions for that event. However, our matching results are qualitatively consistent across all three events, including those for which there is only very marginal evidence, if any, for an endogeneity concern. Our conclusions would be unchanged even if we were to completely exclude the BAML event from our matching analysis.

In addition to our empirical analysis of exchange dark pool trades from closing brokers and brokers who continue to operate dark pools, several institutional features make it unlikely that traders and investors substantially change their choices of brokers in response to a pool closure, especially in the short term. Institutions usually rely on a panel of brokers, and the choice of brokers on the panel and the market share traded with each is based on the institutions best execution policy which gets reviewed only periodically. Our conversations with dark pool operators indicated that while institutions will have preferences about lit vs dark trading, they generally defer the choice of dark venue to the broker and their order routers.[20] Finally, the pool closure events that we study are unrelated to market quality outcomes. Instead they are informed by the cost of operating the pool vs the revenue

[19] If a broker whose pool closes loses customer order flow to other brokers in non-random ways, then the trade outcomes for the control group could conceivably be affected by the pool closure. If this were the case, then our treatment effect estimator captures the sum of the effect on the closing broker and this "spillover" effect for the control group. This is not problematic for our purposes, since a significant total effect suggests a potential violation of our identifying assumptions. In reality, the control group of trades is much larger than the treated group, so these "spillover" affects would be small on average.

[20] Some institutions may "white-list" or "black-list" specific venues, but venue choices on an individual order level are typically made by brokers rather than institutions.

benefits gained, and the perceived regulatory risk of operating a pool.[21] Brokers' decisions to close their pool indicates that the benefits of running the pool no longer exceeded the costs.

To verify whether these anecdotes are supported by the data we plot the market shares of all broker dark pools that are continuously operating before and after pool closure, presented in Figure 5. If institutions are routing orders away from brokers who do not continue operating their dark pool, then we expect to see the market shares of the remaining pools increase. There is no clear trend in any direction in these market shares for any of the three closure events in our sample.

Figure 5 about here

5 Information leakage and high frequency trading

Why are execution outcomes worse for exchange dark pool trades compared with broker dark pool trades? A key difference between the venue categories is that brokers are able to limit access of certain traders to their dark pools. In contrast, in exchange dark pools any trader who can submit orders to an exchange's limit order books can also submit orders to the respective exchange dark pool.

A potential explanation for why execution outcomes are worse in exchange dark pools is that these markets have a higher proportion of HFT/ELP who place small orders with the intention of detecting the presence of an institutional order. These traders then trade in the same direction as the detected order to take advantage of any price pressure from the institutional order. Such short-term directional strategies (order anticipation in Sağlam, 2020 and Hirschey, 2021, back-running in Yang and Zhu, 2019 and sniping in Park and

[21] During the sample period, ASIC issued a number of infringement notices to brokers operating dark pools for failures to comply with Market Integrity Rules relating to regulatory reporting requirements. For example, some brokers incorrectly labeled trades as agency when they were executed as principal, included incorrect information about the origin of an order or incorrect venue reporting information.

Malinova, 2020) result in the HFT/ELP firm subsequently consuming liquidity on the same side of the book as the institutional order. Van Kervel and Menkveld (2019) show that when HFT trade against institutional order flow it increases the short-term order imbalance, resulting in both the higher absolute price impact and a wider bid-ask spread subsequent to the dark trade execution.

5.1 Evidence from differences in trader access across broker dark pools

Broker dark pools differ by how accessible they are to high frequency traders. Our first test for the importance of HFT/ELP trading is to examine the difference in execution outcomes across broker dark pools split by whether or not the pool permits HFT/ ELP trades.

Table 1 groups broker dark pools into two groups based on the level of restrictions. The first group, labeled *restricted*, limit access the most. They prohibit order flow from the firms' principal trading desk, HFT and ELP. These pools do not accept order flow from other pools, nor do they send their orders to other pools. These pools therefore include only *natural* liquidity.[22] The second group, labelled *opt-in to restrictions*, allow customers to opt-out of interacting with principal or ELP flow either entirely or on an order-by-order basis. These pools may also interact with other pools or send/receive orders from dark aggregators, but customers are given the choice of whether or not to participate in this flow. Therefore, while these pools comprise diverse order flow, customers can opt-in to interacting with predominantly natural liquidity. Our sample comprises four restricted access pools and nine that allow customers to opt-in to restrictions. The choice to opt-in or -out of flow represents a trade-off for institutions. Opting into this flow increases the pool of available liquidity and the probability of execution, but likely increases information leakage and adverse selection risk.

Approximately 94% (87%) of broker dark pool trades (dollar volume) in our sample

[22] Unlike other pools, one of our restricted pools, Liquidnet, does not operate a dark limit order book. Instead customers send indications of interest which become actionable when counter-party liquidity is found. Liquidnet does allow liquidity partners, but the minimum order size is AUD 100,000, which implicitly excludes HFT firms. Therefore we classify it as a *restricted* pool.

occur in opt-in-to-restrictions pools, which is not surprising given that (i) these pools are more numerous; (ii) two of the largest restricted broker dark pool, UBS PIN and Citi Match, close during our sample period; and (iii) that another restricted broker dark pool, Liquidnet, has a minimum order size of $100,000.[23] Consequently, our main results so far largely reflect differences in execution outcomes between broker dark pools with opt-in-restrictions and exchange dark pools. However, if HFT/ELP activity is truly responsible for the differences in execution outcomes between exchange dark pools and broker dark pools, then we also expect to see differences in execution outcomes between broker dark pools that completely or partially restrict access to this kind of activity.

To test this, we form a sample of all trades on broker dark pools and run panel regressions on the stock-day averages of execution outcomes onto fixed effects, controls and the stock-day average of a dummy variable taking the value of 1 if the trade is on a broker that does not permit any HFT/ELP activity ("HFT-restricted Pool") and 0 otherwise:

$$\bar{y}_{jt} = \alpha_j + \gamma_t + \beta Restr\bar{i}cted_{jt} + \rho' \bar{X}_{jt} + \bar{\varepsilon}_{jt} \tag{11}$$

where all variables are again defined as per Equation (5) other than $Restr\bar{i}cted_{jt}$ which is the stock-day average of the dummy variable for the trade taking place in a pool that does not permit HFT/ELP. The key parameter in Equation (11) is β which captures average execution outcomes for trades on broker dark pools that restrict HFT/ELP broker dark pools vs. other broker dark pools. Table 9 presents the estimates from these regressions where absolute price impact and bid-ask spreads are measured at the 60s horizon.

Table 9 about here

Our results from Table 9 are consistent with HFT/ELP influencing dark pool execution

[23]Trades in restricted dark pools represent more than twice as much of total broker dark pool dollar volume (13%) as they do total broker dark pool trades by number of executions.

outcomes.[24] Trades in pools that do not permit HFT/ELP activity have significantly lower absolute price impact than do trades in pools that allow HFT/ELP activity. The size of the effect is -0.8 bps, compared to our difference in absolute price impact of around -1.2 bps between broker dark pool trades and exchange dark pool trades under the same specification in Table 3. We find that post-trade bid-ask spreads are -0.4 bps lower for trades on restricted broker dark pools (t-statistic of -1.69) vs. -0.5 bps (-2.93) in the main panel analysis. Reversals are less likely after trades on these pools and price adjustments are faster.

5.2 Evidence from small trades

Our second approach to testing the relevance of this channel is to re-run our panel regressions from Section 4.1 while splitting the effect of trade location by trade size. We create two new dummy variables that take the value 1 when trade size is below the lowest 10[th] percentile by stock-week and above this percentile respectively. Since HFTs tend to trade in smaller sizes than other trader categories, if HFT trading strategies are responsible for the patterns in execution outcomes across venue types, we expect to see greater differentials in execution outcomes for smaller sized trades. The regression model is:

$$\bar{y}_{jt} = \alpha_j + \gamma_t + \beta_0 \bar{D}_{jt}^{size \leq \bar{v}} + \beta_1 B\bar{D}P_{jt}^{size \leq \bar{v}} + \beta_2 B\bar{D}P_{jt}^{size > \bar{v}} + \rho' \bar{X}_{jt} + \bar{\varepsilon}_{jt} \tag{12}$$

where all variables are defined as per Equation (5) other than $\bar{D}_{jt}^{size \leq \bar{v}}$ which is the stock-day average of a dummy variable taking the value of 1 if the trade size is below the lowest 10[th] percentile and 0 otherwise, $B\bar{D}P_{jt}^{size \leq \bar{v}}$ which is the stock-day average of a dummy variable taking the value of 1 if the trade takes place on a broker dark pool and trade size is below the lowest 10[th] percentile, and $B\bar{D}P_{jt}^{size > \bar{v}}$ which is the stock-day average of a dummy variable taking the value of 1 if the trade takes place on a broker dark pool and trade size exceeds

[24]The sample size and number of stocks are smaller than our main regressions because the main regressions also include the exchange dark pool trades.

this percentile. The key parameters of Equation (12) are β_1 and β_2 which are estimates of the difference in execution outcomes between broker dark pool trades and exchange dark pool trades for small and large sized trades respectively. Table 10 presents the parameter estimates from these regressions.

Table 10 about here

Again consistent with the order anticipation channel, the reduction in price impact for broker dark pool trades vs. exchange dark pool trades is greater for smaller sized trades relative to larger trades. The average absolute price impact for a small trade on a broker dark pool is -1.9 bps (t-statistic of -4.52). For a larger trade, this difference is still negative and significant, but around half the size. Similarly, prices adjust more slowly for large trades on broker dark pools compared with small trades. Together, these suggest that information is being impounded relatively more quickly after a small trade in an exchange dark pool compared with a broker dark pool than is the case for large trades. The patterns in bid-ask spreads and reversals are less supportive of the hypothesis — for smaller trades the point estimate on bid-ask spreads is -0.0062 and insignificant while for large trades, it is -0.0052 and significant. There are no significant differences in the reversal probability for small trades on broker dark pools compared with exchange dark pools, while for large trades, broker dark pool trades are slightly more likely to exhibit reversals. Since our hypotheses directly relate to information leakage, our results on absolute price impact are the most relevant for assessing the role of access.

Although we cannot observe "minimum acceptable quantity" (MAQ) instructions, their use may also contribute to this result. MAQs allow traders to specify that the order can only execute if the MAQ is met. Therefore, orders with a MAQ are expected to be larger on average than those that without a MAQ. Table 10 shows that, conditional on execution, larger orders have lower information leakage on average, although this will come at the cost of a lower probability of execution.[25]

6 Robustness

6.1 Long horizons price impacts and spreads

Our reversals indicator and price adjustment variable each rely on the argument that the new equilibrium price is, on average, reached no later than 30 minutes after the trade. Examining the effect on absolute price impact 30 minutes after the trade is executed can also be considered a test for whether or not our results reflect differences in the average informativeness of orders that are routed to broker dark pools vs. exchange dark pools. If this were the case, it would confound our estimates of the effect of venue choice on execution outcomes as we would not be comparing like-with-like in our stock-day panel regressions.

Column (7) of Table 3 is consistent with our interpretations as there are no significant differences in absolute price impact that can be attributed to trading on broker dark pools vs. exchange dark pools once we control for stock and date fixed effects and other control variables. Table IA.4 of the Internet Appendix contains treatment effects for the same horizon obtained under our matching approach. Here, we find qualitatively similar results — we can reject the null of equal absolute price impact over these horizons at the 5% level for only one of the three events while for two of the three events, there is evidence that bid-ask spreads are wider.

6.2 Excluding large ASX-reported broker dark pool trades

As discussed in Section 3, we cannot directly distinguish between broker dark pool trades reported to the ASX and manual matches from brokers with dark pools. The latter are a very small component of overall trading activity (less than 0.5% according to our Cboe data that directly identifies these kinds of trades) however to ensure that these do not unduly

[25] MAQ conditions can also be added to orders on exchange dark pools and so the use of MAQs cannot explain the differences in execution outcomes between exchange and broker dark pools observed in Section 4.

affect our results, we rerun our main panel regressions where we exclude all NBBO trades reported to the ASX for values of $50,000 or more. The logic is that manual executions are usually much larger than typical broker dark pool trades and filtering trades that are more likely to be manual matches.[26] We replicate our main results using these data and obtain qualitatively similar findings, though the effect on adjustment speed becomes significant at the 5% level using these data. These results are presented in Table IA.5 in the Internet Appendix.

6.3 Excluding Cboe hidden trades

Orders sent to exchange dark pools can interact with lit orders on the corresponding exchange. For ASX trades, submitting such an order is an "opt-in" process: a lit order will not interact with hidden liquidity on Centre Point unless the trader explicitly chooses to. On Cboe, all marketable orders will by default execute against any available hidden liquidity as Cboe hidden orders at that time, though traders can "opt-out" of executing against hidden orders. Our analysis does not distinguish between Cboe hidden and Centre Point trades because we want to make comparisons across the the entire landscape of dark pool trading. However because Cboe hidden liquidity is more integrated with the lit order book, we run our main panel regressions with only Centre Point trades to ensure that our results do not simply reflect differences between broker dark pool trades and the integration of lit and dark liquidity on Cboe. These results are consistent with our main analysis, although we find no evidence of differences in reversal probability and some evidence of slower adjustment speed that is significant at the 10% level (see Table IA.6 in the Internet Appendix).

[26] Again, Cboe data allows us to test this claim and a cutoff of $50,000 results in mis-classifying around 0.5% of trades, in total. Though this a similar fraction of mis-classified trades as using no filter, the composition of erroneous trades changes and this helps test the sensitivity of our results to manual matches.

6.4 Intermarket sweep orders

When we estimate differences in execution outcomes between exchange and broker dark pools in Section 4.1, we first remove dark pool trades that occur at the exact same time as a lit order book trade, according to time-stamps in our data. We do this because traders have the ability to send sweep orders to the exchange dark pools that can directly execute against both hidden and displayed liquidity without latency. Sweep orders can potentially generate differences in execution outcomes that are independent of any information effects. We also generate our main panel results where we do not filter out these trades. Results including sweeps are contained in Table IA.7 in the Internet Appendix. As expected, we detect stronger effects for price impact and spreads when sweeps are included.

6.5 Tick constraints and venue selection

When the minimum tick size is a binding constraint for the bid ask spread, the depth the NBBO tends to be large, and the cost of demanding liquidity is also high. Therefore, trading within the spread, on dark venues becomes more attractive (Kwan, Masulis, and McInish, 2015; Comerton-Forde, Grégoire, and Zhong, 2019; O'Hara, Saar, and Zhong, 2019). Tick constraints should affect order routing across both exchange and broker dark pools, however, our analysis could potentially be biased if one dark pool category was favored over the other in a tick-constrained environment.

To check this, we compute the ratio of the time-weighted dollar spread for each stock-day to the minimum tick size for that stock.[27] Stocks with low (high) ratios of time-weighted spread to minimum tick size are relatively more tick constrained (unconstrained). We then sort stocks in each month into two categories, tick constrained and tick unconstrained, based on whether or not this ratio is below that month's 25$^{\text{th}}$ percentile or above that month's 75$^{\text{th}}$ percentile of the spread-to-tick-size ratio, leaving the middle 50$^{\text{th}}$ percent unassigned.

[27] Minimum tick sizes in the Australian market are 0.1c for stocks priced below 10c, 0.5c for stocks priced greater than or equal to 10c and below \$2, and 1c for stocks priced \$2 and above.

We use this dummy variable in regressions where the left hand side variable is (i) the total fraction of trading that takes place in dark pools and (ii) the fraction of dark trading that takes place on broker dark pools, including stock and day fixed effects and controls used in our other stock-day regressions. Results from the literature suggest the coefficient on the tick-constrained dummy should be positive and significant for regression of (i). Our main interest is in the tick-constrained coefficient for the regression of (ii), where statistical and economic significance could indicate a problem with our inference. Results from these regressions are presented in Table IA.8. These show that tick constrained stocks have greater fractions of total volume traded in dark pools, but more importantly, that there is no significant effect of tick constraints on the amount of dark pool volume executed on broker dark pools relative to exchange dark pools. Therefore, tick constraints are unlikely to affect our results.

7 Conclusion

Dark pools are an important part of the trading ecosystem in most developed markets. To date, the academic literature largely treats dark pools as homogeneous, but in practice they differ on a number of important dimensions. We examine differences in the extent to which dark pools restrict access to certain groups of traders, and how these differences affect execution outcomes. We focus on differences between exchange dark pools which are open to all traders and broker dark pools which can exclude HFT, ELP and other principal flow.

Our results from a panel analysis, with a rich set of controls and fixed effects, as well as a matching analysis based on brokers that close their pools during our sample, show that broker dark pool trades have less information leakage and less adverse selection risk than exchange dark pools. We find some evidence that broker dark pool trades have less transitory price pressure than exchange dark pool trades. Comparing execution outcomes in broker dark pools that restrict HFT/ELP access to those that allow customers to opt-out of interacting with HFT/ELP, and across broker and exchange dark pools by trade size suggest that the presence of HFT/ELP order flow is a key driver of the main results.

These findings are consistent with analysis undertaken by ASIC using account level data. They classify traders into agency (trading on behalf of a client), principal (the broker trading for their house book) and HFT and show that there are no obvious winners and losers by counter-party type in broker dark pools. However, on exchange dark pools, agency counter-parties are on the losing side of the trade around 68% of the time, while HFT counter-parties are on the winning side 95% of the time.

Our results are relevant in other settings given that other jurisdictions also allow dark pools to provide differential access to different types of customers. For example, in the U.S. exchanges must provide "fair access" but Alternative Trading Systems (ATS) can segment order flow if their market share is below 5%.[28] However, analysis similar to ours is not possible in the US because there is no trade-level attribution to dark pools. In Europe, regulators banned the operation of Broker Crossing Networks (BCN) which allowed brokers to restrict access, without a cost-benefit analysis of this decision. If our results translate in the European context, this ban may have had a negative impact on investors.

Although exchange dark pools are not able to restrict access by trader type, investors can indirectly influence the type of flow they interact with through the use of Minimum Acceptable Quantities (MAQ). MAQs allow traders to specify their minimum trade size, therefore reducing the risk of trading with HFT/ELP. Unfortunately, we are unable to examine the impact of MAQ on execution outcomes as these data are non-public. However, if these data could be secured it would be an interesting area for future research.

Our study also shines light on the absence of regulations related to order routing disclosure in Australia. SEC Rule 606(b)(3) has recently updated requirements for these disclosures in the US. The widespread support for these rules in the US market suggests that these types of routing data help buy-side traders to engage with their brokers and make better decisions. The adoption of similar rules would likely benefit Australian investors. Standardized public disclosures would also facilitate better industry wide analysis and independent research.

[28] All ATS in the US have market share well below 5%.

References

Anand, A., M. Samadi, J. Sokobin, and K. Venkataraman (2021). Institutional order handling and broker-affiliated trading venues. *Review of Financial Studies 34*(7), 3364–3402.

Arpino, B. and M. Cannas (2016). Propensity score matching with clustered data. An application to the estimation of the impact of caesarean section on the Apgar score. *Statistics in Medicine 35*(12), 2074–2091.

ASIC (2015). Report 452: Review of high-frequency trading and dark liquidity. https://asic.gov.au/media/3444836/rep452-published-26-october-2015.pdf. Accessed: 2021-06-01.

ASX (2021). Historical market statistics. https://www2.asx.com.au/about/market-statistics/historical-market-statistics. Accessed: 2021-07-22.

Battalio, R., S. A. Corwin, and R. Jennings (2016). Can brokers have it all? On the relation between make-take fees and limit order execution quality. *Journal of Finance 71*(5), 2193–2238.

Battalio, R. H., B. C. Hatch, and M. Saglam (2022). The cost of exposing large institutional orders to electronic liquidity providers. *https://ssrn.com/abstract=3281324*.

Bessembinder, H. and H. M. Kaufman (1997). A cross-exchange comparison of execution costs and information flow for NYSE-listed stocks. *Journal of Financial Economics 46*(3), 293–319.

Brolley, M. (2020). Price improvement and execution risk in lit and dark markets. *Management Science 66*(2), 863–886.

Buti, S., B. Rindi, and I. M. Werner (2011). Diving into dark pools. *Charles A. Dice Center Working Paper* (2010-10).

Comerton-Forde, C., V. Grégoire, and Z. Zhong (2019). Inverted fee structures, tick size, and market quality. *Journal of Financial Economics 134*(1), 141–164.

Comerton-Forde, C., K. Malinova, and A. Park (2018). Regulating dark trading: Order flow segmentation and market quality. *Journal of Financial Economics 130*(2), 347–366.

Comerton-Forde, C. and T. J. Putniņš (2015). Dark trading and price discovery. *Journal of Financial Economics 118*(1), 70–92.

Conrad, J. and S. Wahal (2020). The term structure of liquidity provision. *Journal of Financial Economics 136*(1), 239–259.

Easley, D., N. M. Kiefer, and M. O'Hara (1996). Cream-skimming or profit-sharing? The curious role of purchased order flow. *Journal of Finance 51*(3), 811–833.

Hatheway, F., A. Kwan, and H. Zheng (2017). An empirical analysis of market segmentation on U.S. equity markets. *Journal of Financial and Quantitative Analysis 52*(6), 2399–2427.

Hirschey, N. (2021). Do high-frequency traders anticipate buying and selling pressure? *Management Science 67*(6), 3321–3345.

Korajczyk, R. A. and D. Murphy (2018). High-frequency market making to large institutional trades. *Review of Financial Studies 32*(3), 1034–1067.

Kwan, A., R. Masulis, and T. H. McInish (2015). Trading rules, competition for order flow and market fragmentation. *Journal of Financial Economics 115*(2), 330–348.

Menkveld, A. J., B. Z. Yueshen, and H. Zhu (2017). Shades of darkness: A pecking order of trading venues. *Journal of Financial Economics 124*(3), 503–534.

Norges Bank Investment Management (2015). Sourcing liquidity in fragmented markets 01-2015. https://www.nbim.no/en/publications/asset-manager-perspectives/2015/sourcing-liquidity-in-fragmented-markets/. Accessed: 2021-07-22.

O'Hara, M., G. Saar, and Z. Zhong (2019). Relative tick size and the trading environment. *Review of Asset Pricing Studies 9*(1), 47–90.

Park, A. and K. Malinova (2020). "Sniping" in Fragmented Markets. *Available at SSRN 3534367*.

Rosenbaum, P. R. and D. B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika 70*(1), 41–55.

Sağlam, M. (2020). Order anticipation around predictable trades. *Financial Management 49*(1), 33–67.

S&P Global (2020). All Ordinaries. https://www.spglobal.com/spdji/en/indices/equity/all-ordinaries/#overview. Accessed: 2020-11-30.

Van Kervel, V. and A. J. Menkveld (2019). High-frequency trading around large institutional orders. *Journal of Finance 74*(3), 1091–1137.

Westpac (2020). The All Ordinaries Index. https://www.westpac.com.au/personal-banking/share-trading/the-share-market/asx-all-ords/. Accessed: 2020-11-30.

Yang, L. and H. Zhu (2019). Back-running: Seeking and hiding fundamental information in order flows. *Review of Financial Studies 33*(4), 1484–1533.

Ye, M. and W. Zhu (2020). Strategic informed trading and dark pools. *Available at SSRN 3292516*.

Zhu, H. (2014). Do dark pools harm price discovery? *Review of Financial Studies 27*(3), 747–789.

Figure 1: Difference in execution outcomes for trades on broker dark pools and exchange dark pools

This figure presents the estimated difference in absolute price impact for broker dark pool trades and exchange dark pool trades obtained from stock-day panel regressions. Each point on the solid line represents the estimated difference in absolute price impact for a trade taking place on a broker dark pool compared with an exchange dark pool at horizons of 500ms, 1s, 10s, 30s, 60s, 300s and 1800s respectively, conditional on stock and time fixed effects and control variables. Dashed lines are 95% confidence intervals using standard errors that are clustered at the stock-level.

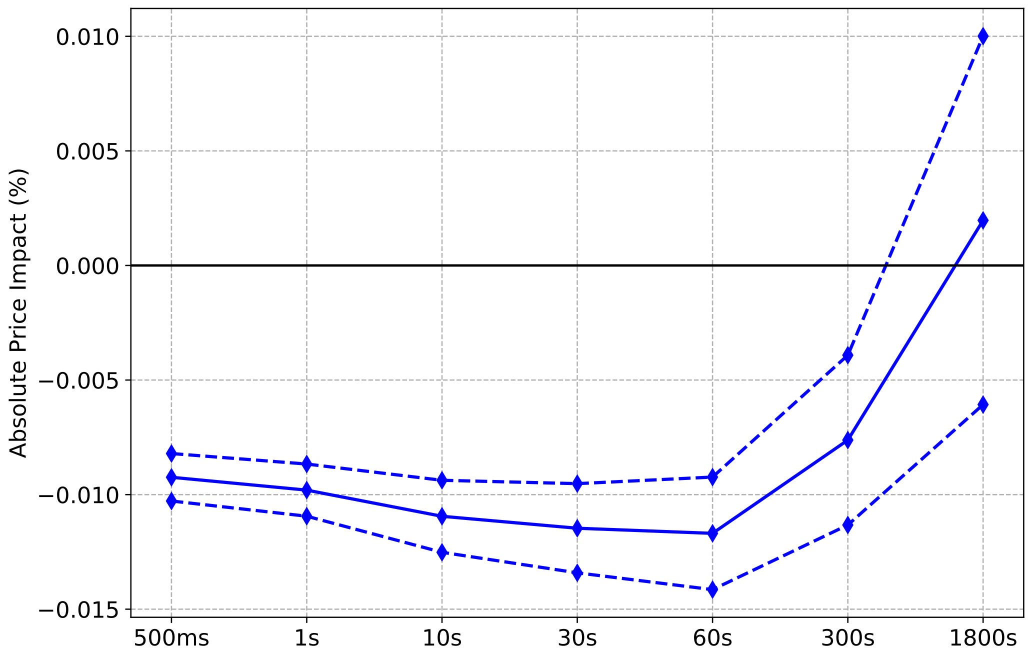

Figure 2: Broker dark pools and exchange dark pools market shares

Time-series of weekly market shares of total dollar volume traded on broker dark pools, Centre Point and Cboe Hidden orders. Market shares are calculated using trades in all stocks in our sample by week. Broker dark pool market share is summed across all broker dark pools in our sample.

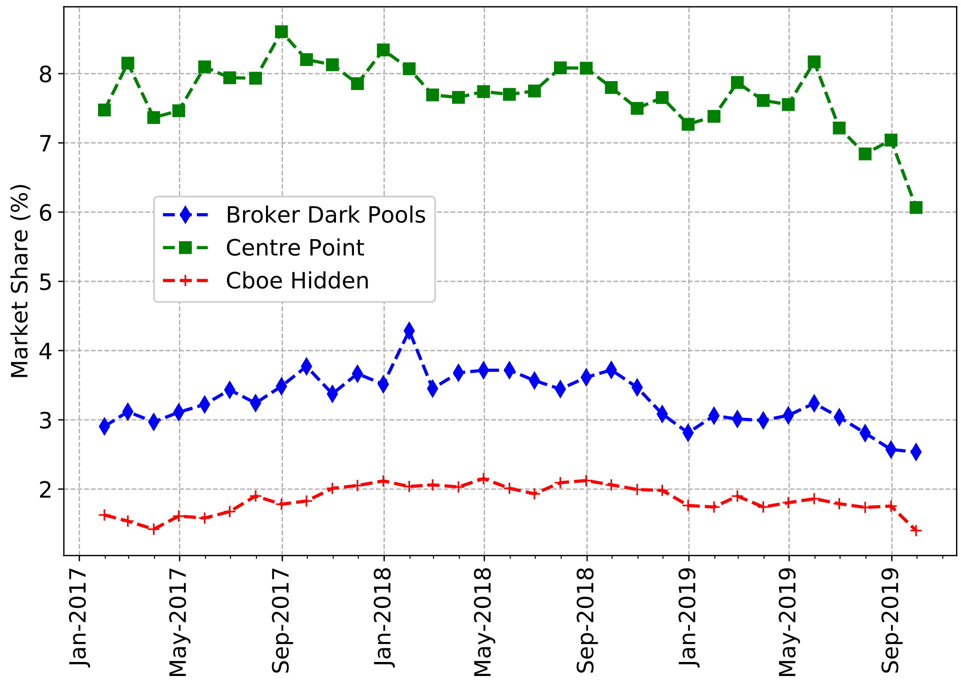

Figure 3: Average trade size time-series and trade size distributions by dark pool category

Panel (a) presents average dollar trade size by stock-week for broker dark pools, Centre Point and Cboe Hidden trades. Panel (b) presents the fraction of trades that are for $100 or less, $100 to $500, $500 to $1,000, $1,000 to $5,000, $5,000 to $10,000, $10,000 to $50,000, and above $50,000 for the three venue categories.

(a) Average trade size

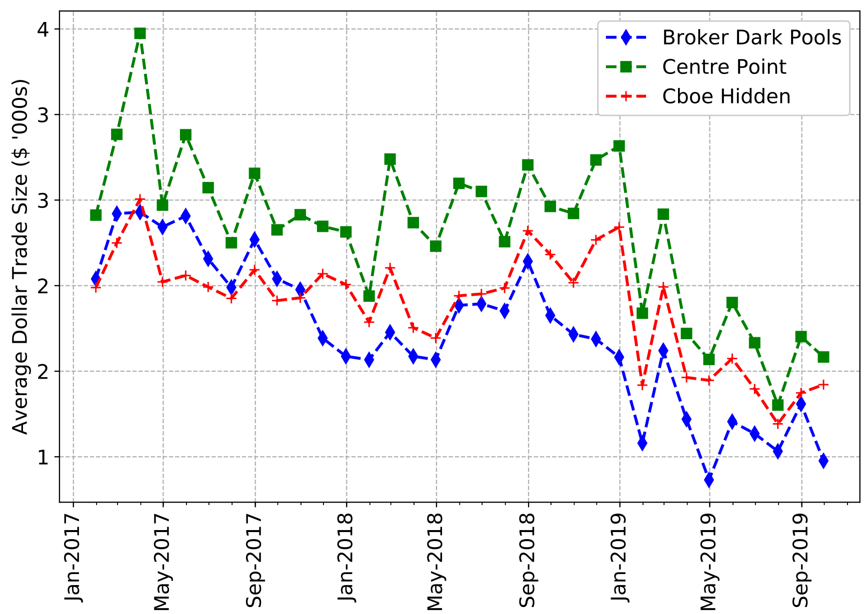

(b) Trade size distributions

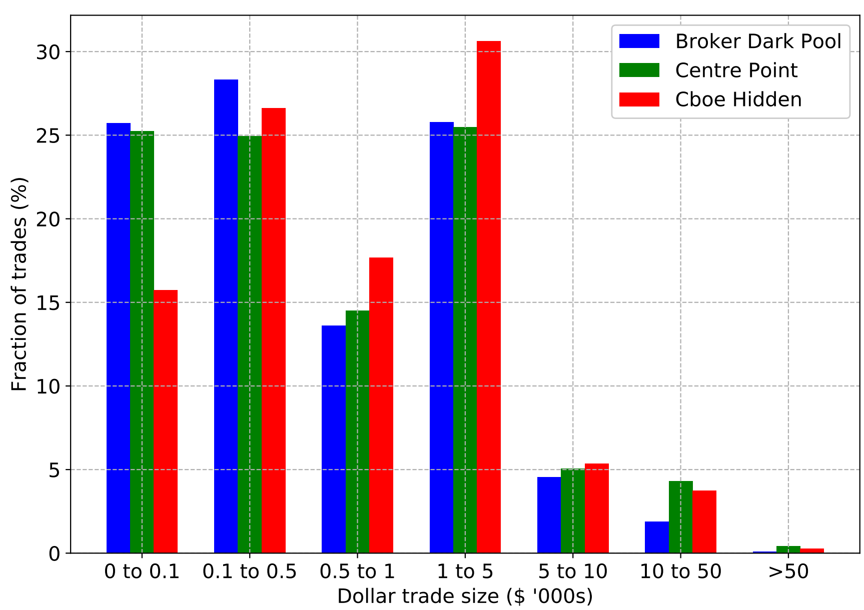

Figure 4: Number of broker dark pool trades for closing pools

Time-series of the number of broker dark pool trades for each of the three brokers whose pools are closed during our sample period: Bank of America Merrill Lynch (BAML), UBS and Citigroup (Citi). The respective closure dates are March 6, 2017 (BAML), April 1, 2019 (UBS) and July 1, 2019 (Citi).

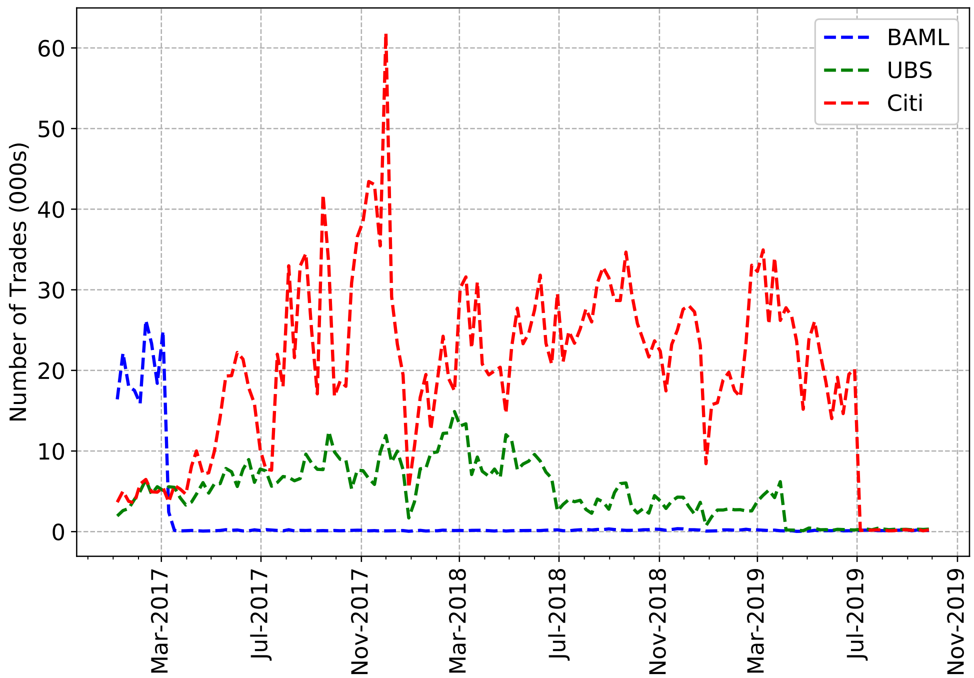

Figure 5: Market share of continuously operating broker dark pool around closures

Time-series of the total market shares of all continuously operating broker dark pools one month before and after the three pool closure events in our sample period: Bank of America Merrill Lynch (BAML), UBS and Citigroup (Citi). The respective closure dates are March 6, 2017 (BAML), April 1, 2019 (UBS) and July 1, 2019 (Citi). Market share is on the y-axis and is defined as the dollar volume traded on broker dark pools that remain in operation summed across stocks and days, divided by total dollar volume traded summed across stocks and days, expressed in percent. The x-axis is the number trading days until and since the closure date ($t = 0$).

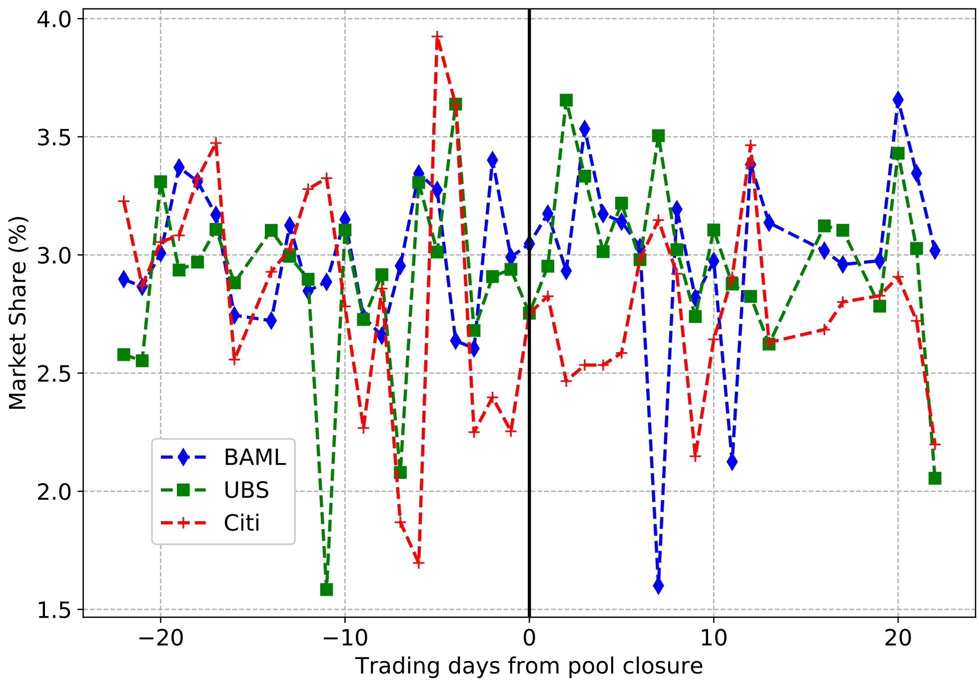

Table 1: Broker dark pool classifications

This table summarizes access restrictions and other relevant trading rules by broker dark pool. We classify broker dark pools into two “Pool type” categories based on whether or not HFTs are explicitly excluded (“Restricted”) or HFTs are possibly present but other traders can opt-in to avoid executing against them (“Opt-in to restrictions”).

Pool type	Pool name	Operator	Launch month	Closure month	Matching rules	Allows HFT, principal or ELP	Can clients opt-out of specific flow	Receives orders from other pools	Sends orders to other pools
Restricted	UBS PIN	UBS Securities	August 2005	March 2019	Price-time	No	-	No	No
Restricted	Citi Match	Citigroup	July 2013	July 2019	Price-time	No	-	No	No
Restricted	CLSA Match	CLSA	October 2012	-	Price-time	No	-	No	No
Restricted	Liquidnet	Liquidnet	February 2008	-	Negotiated or when automated volume split equally	Liquidity partners may place principal orders, but can not negotiate directly and must meet minimum order size of $100,000 and minimum average daily order flows and average order resting time requirements	No	No, but aggregator algorithms provides access to other pools	On an order-by-order basis or by default clients can give instructions to place an order on external venues, including aggregation algorithms
Opt-in to restrictions	Crossfinder	Credit Suisse	April 2006	-	Price-time	Yes, all order flow is accepted, but toxicity checks in place with potential for excluding customers that fail checks	Yes, may opt-out by counterparty type	Yes, accepts orders from aggregator algos	No
Opt-in to restrictions	MAQX	Macquarie Securities	September 2010	-	Price-time	Yes, allows ELP and prop	Yes, may opt-out by counterparty type	Yes, accepts orders from aggregator algos	No
Opt-in to restrictions	SuperX	Deutsche Securities	June 2011	March 2020	-	Yes, all order flow is accepted	Yes, may opt-out by counterparty type	Yes, accepts orders from aggregator algos	No
Opt-in to restrictions	BLX	Instinet	April 2011	-	Price - pro-rata	No principal flow, but no restrictions on client types	No	No, but clients may access orders from other crossing systems through aggregator algo	No, but clients may access Instinet’s aggregator algo to send orders to other pools
Opt-in to restrictions	JPM-X	J. P. Morgan	October 2015	-	-	Yes, all order flow is accepted	Yes, may opt-out by counterparty type	Yes, accepts orders from aggregator algos	No
Opt-in to restrictions	MS Pool	Morgan Stanley	March 2010	-	-	Yes, all order flow is accepted	Yes, may opt-out by counterparty type	Yes, accepts orders from aggregator algos	No
Opt-in to restrictions	POSIT	Virtu ITG	May 2010	-	Price - pro-rata	No HFT but allows liquidity providers, other participants and third-party brokers, and orders from other crossing system operators (including principal orders)	Yes, may opt-out by counterparty type	Yes, accepts orders from aggregator algos	No, but clients may access orders from other crossing systems through POSIT Marketplace
Opt-in to restrictions	Sigma X	Goldman Sachs	January 2010	-	Price-time	No orders from liquidity providers, market makers of HFT, but allows orders from GS equity-linked businesses	Yes, may opt-out of flow from aggregator algos	Yes, accepts orders from aggregator algos	No
Opt-in to restrictions	InstinctX	BAML	August 2010	March 2017	Price-time	Yes, all order flow is accepted	Yes, may opt-out by counterparty type	Yes, accepts orders from aggregator algos	No

Table 2: Stock-day summary statistics

This table contains stock-day summary statistics for dark pool trades in stocks in the ASX All Ordinaries Index from the period Jan 1, 2017 to Sep 30, 2019. All stock-days that record at least one trade either on a broker dark pool or an exchange dark pool are included in the sample. Trade size is the dollar volume of the trade in thousands of dollars, averaged across all dark pool trades by stock-day. Total dollar value is the total amount traded across all venues for that stock-day. Price is the average trade price in dollars. Daily Average Bid-ask Spread is the time-weighted average log difference between the national best ask and bid for that stock and day. Daily Average Dollar Depth is the average depth available at the national best bid and offer for that stock and day, measured in thousands of dollars. Broker Dark Pool is a dummy variable taking the value one if a trade is on a broker dark pool and zero if the trade is in an exchange dark pool, averaged across all dark pool trades by stock-day (i.e. it is the average proportion of dark pool trades on broker dark pools). Pre-Cross Bid-ask Spread is the bid-ask spread at the time of the dark pool trade, averaged across all dark pool trades by stock-day. Absolute τ price impact is the absolute log difference between the mid-quote τ seconds after the trade and the prevailing mid-point at the time of the trade, averaged across all dark pool trades by stock-day and expressed in percent. τ bid-ask spread is the log bid-ask spread τ seconds after the trade, averaged across all dark pool trades by stock-day, expressed in percent. Reversals Indicator is an indicator variable taking the value one when the sign of the one minute mid-quote return is opposite to the sign of the return from one minutes to thirty minutes after the trade. Price Adjustment is the absolute difference between the mid-point thirty minutes and one minute after the trade time, scaled by the prevailing mid-point and expressed in percent, averaged across all dark pool trades by stock-day.

	Mean	SD	Min	25%	50%	75%	Max
Trade Size (AUD '000s)	2.03	13.4	0.00	0.44	1.00	2.10	3,402
Total Dollar Value (AUD 'm)	13.8	33.0	0.00	0.51	2.98	13.0	1,458
Price (AUD)	8.63	18.8	0.00	1.40	3.38	7.76	241
Daily Average Dollar Depth (AUD '000s)	101	226	1.38	12.5	33.2	95.8	2,126
Daily Average Bid-ask Spread (%)	0.57	0.79	0.02	0.16	0.34	0.61	14.0
Broker Dark Pool	0.27	0.24	0.00	0.04	0.25	0.43	1.00
Pre-Cross Bid-ask Spread (%)	0.44	0.43	0.00	0.14	0.31	0.52	3.92
Abs. 500ms Price Impact (%)	0.01	0.03	0.00	0.00	0.01	0.01	0.31
Abs. 10s Price Impact (%)	0.03	0.05	0.00	0.01	0.02	0.03	0.48
Abs. 60s Price Impact (%)	0.06	0.08	0.00	0.02	0.04	0.06	0.84
Abs. 1800s Price Impact (%)	0.33	0.31	0.00	0.15	0.24	0.40	3.89
500ms Bid-Ask Spread (%)	0.45	0.43	0.00	0.15	0.31	0.52	3.92
10s Bid-Ask Spread (%)	0.45	0.44	0.00	0.15	0.32	0.52	3.92
60s Bid-Ask Spread (%)	0.47	0.47	0.00	0.16	0.33	0.54	3.92
1800s Bid-Ask Spread (%)	0.68	0.71	0.00	0.24	0.45	0.80	5.89
Price Adjustment (%)	0.31	0.29	0.00	0.14	0.23	0.37	3.79
Reversal Indicator	0.12	0.14	0.00	0.02	0.07	0.17	1.00

Table 3: Stock-day panel regression — Absolute price impact all horizons

This table contains estimates from regressions of the stock-day average absolute price impact at various horizons after a dark pool trade onto stock-day level controls, fixed effects and the fraction of all dark pool trades that occur on a broker dark pool. The regression model is $\bar{y}_{jt} = \alpha_j + \gamma_t + \beta B\bar{D}P_{jt} + \rho' \bar{X}_{jt} + \bar{\varepsilon}_{jt}$ where α_j is a stock fixed effect, γ_t is a date fixed effect, $\bar{y}_{jt}$ is the stock-day average of the absolute impact after a trade on either an exchange dark pool or a broker dark pool for stock j and day t, $\bar{X}_{jt}$ is the stock-day average of a vector of controls including log of dollar trade size, log of trade price, the national best bid-ask spread at the time of the dark pool trade, the log of depth available at the national best bid and ask at the time of the dark pool trade, the log of total dollar volume traded across all trades and venues, and the log of total number of dark pool and lit trades, $B\bar{D}P_{jt}$ is the stock-day average of a dummy variable that takes the value 1 if the trade occurs on a broker dark pool and 0 otherwise (i.e. the fraction of broker dark pool trades out of all dark pool trades) and $\bar{\varepsilon}_{jt}$ is an error term. We estimate the model for horizons of 500ms, 1s, 10s, 30s, 60s, 300s and 1800s using all stock-days from Jan 1, 2017 to Sept 30, 2019. Reported R^2 values relate to the within variation in the dependent variables. Standard errors are clustered at the stock level and t-statistics are in parentheses.

	(1)		(2)		(3)		(4)		(5)		(6)		(7)	
	500ms		1s		10s		30s		60s		300s		1800s	
Ln(Dollar Trade Size)	-0.0009	(-5.11)	-0.0008	(-4.13)	-0.0005	(-2.01)	-0.0006	(-1.97)	-0.0007	(-1.78)	-0.0020	(-3.65)	-0.0046	(-4.44)
Ln($N_{DarkPool}$)	-0.0083	(-33.4)	-0.0094	(-33.2)	-0.0131	(-35.6)	-0.0166	(-36.9)	-0.0202	(-36.3)	-0.0315	(-36.2)	-0.0477	(-27.8)
Ln(N_{Lit})	0.0056	(17.9)	0.0064	(17.6)	0.0111	(22.0)	0.0158	(25.7)	0.0215	(27.5)	0.0471	(33.1)	0.1025	(34.8)
Ln(Price)	-0.0001	(-0.19)	0.0002	(0.22)	-0.0009	(-0.70)	-0.0043	(-2.51)	-0.0092	(-4.05)	-0.0406	(-9.31)	-0.1420	(-13.9)
Ln(Dollar Volume)	0.0032	(14.6)	0.0035	(13.7)	0.0062	(16.2)	0.0098	(19.7)	0.0138	(20.2)	0.0317	(22.6)	0.0721	(23.3)
Pre-cross Bid-ask Spread	0.0069	(6.02)	0.0099	(7.40)	0.0193	(9.89)	0.0263	(10.3)	0.0369	(10.9)	0.0801	(12.9)	0.1813	(15.7)
Ln(Depth)	-0.0047	(-19.3)	-0.0056	(-20.1)	-0.0102	(-24.9)	-0.0144	(-27.7)	-0.0189	(-27.9)	-0.0360	(-29.3)	-0.0700	(-27.6)
Broker Dark Pool	-0.0092	(-17.5)	-0.0098	(-16.9)	-0.0109	(-13.6)	-0.0115	(-11.5)	-0.0117	(-9.35)	-0.0076	(-4.04)	0.0020	(0.48)
Fixed Effects	$N\&T$		$N\&T$		$N\&T$		$N\&T$		$N\&T$		$N\&T$		$N\&T$	
R^2	0.08		0.08		0.09		0.10		0.11		0.13		0.13	
N_{obs}	242,825		242,825		242,825		242,825		242,825		242,825		242,825	
N_{stocks}	626		626		626		626		626		626		626	

Table 4: Stock-day panel regression — Bid-ask spreads all horizons

This table contains estimates from regressions of the stock-day average bid-ask spread at various horizons after a dark pool trade onto stock-day level controls, fixed effects and the fraction of all dark pool trades that occur on a broker dark pool. The regression model is $\bar{y}_{jt} = \alpha_j + \gamma_t + \beta B\bar{D}P_{jt} + \rho' \bar{X}_{jt} + \bar{\varepsilon}_{jt}$ where α_j is a stock fixed effect, γ_t is a date fixed effect, $\bar{y}_{jt}$ is the stock-day average of the bid-ask spread after a trade on either an exchange dark pool or a broker dark pool for stock j and day t, $\bar{X}_{jt}$ is the stock-day average of a vector of controls including log of dollar trade size, log of trade price, the national best bid-ask spread at the time of the dark pool trade, the log of depth available at the national best bid and ask at the time of the dark pool trade, the log of total dollar volume traded across all trades and venues, and the log of total number of dark pool and lit trades, $B\bar{D}P_{jt}$ is the stock-day average of a dummy variable that takes the value 1 if the trade occurs on a broker dark pool and 0 otherwise (i.e. the fraction of broker dark pool trades out of all dark pool trades) and $\bar{\varepsilon}_{jt}$ is an error term. We estimate the model for horizons of 500ms, 1s, 10s, 30s, 60s, 300s and 1800s using all stock-days from Jan 1, 2017 to Sept 30, 2019. Reported R^2 values relate to the within variation in the dependent variables. Standard errors are clustered at the stock level and t-statistics are in parentheses.

	(1)		(2)		(3)		(4)		(5)		(6)		(7)	
	500ms		1s		10s		30s		60s		300s		1800s	
Ln(Dollar Trade Size)	-0.0010	(-3.27)	-0.0011	(-3.92)	-0.0016	(-5.56)	-0.0023	(-6.78)	-0.0025	(-5.80)	-0.0041	(-4.75)	-0.0123	(-5.82)
Ln($N_{DarkPool}$)	-0.0004	(-1.02)	-0.0009	(-2.29)	-0.0024	(-4.91)	-0.0041	(-7.43)	-0.0058	(-8.51)	-0.0097	(-8.46)	-0.0065	(-2.71)
Ln(N_{Lit})	-0.0043	(-6.66)	-0.0045	(-6.77)	-0.0053	(-6.78)	-0.0058	(-6.69)	-0.0068	(-6.16)	-0.0153	(-6.42)	-0.0305	(-5.88)
Ln(Price)	-0.0198	(-8.73)	-0.0210	(-8.92)	-0.0233	(-8.29)	-0.0257	(-7.58)	-0.0293	(-5.85)	-0.0622	(-5.09)	-0.1546	(-6.00)
Ln(Dollar Volume)	0.0018	(3.40)	0.0024	(4.64)	0.0034	(6.12)	0.0041	(6.57)	0.0047	(6.01)	0.0066	(4.09)	0.0126	(3.74)
Pre-cross Bid-ask Spread	0.9214	(179)	0.9173	(165)	0.9077	(142)	0.9068	(127)	0.9255	(109)	0.9481	(55.5)	0.8668	(31.7)
Ln(Depth)	-0.0047	(-7.67)	-0.0051	(-8.23)	-0.0060	(-8.66)	-0.0067	(-7.19)	-0.0059	(-3.88)	-0.0019	(-0.53)	0.0026	(0.40)
Broker Dark Pool	-0.0024	(-2.38)	-0.0032	(-2.78)	-0.0049	(-3.55)	-0.0065	(-4.14)	-0.0053	(-2.93)	-0.0051	(-1.39)	0.0128	(1.72)
Fixed Effects	$N\&T$		$N\&T$		$N\&T$		$N\&T$		$N\&T$		$N\&T$		$N\&T$	
R^2	0.88		0.87		0.85		0.82		0.76		0.55		0.21	
N_{obs}	242,825		242,825		242,825		242,825		242,825		242,825		242,825	
N_{stocks}	626		626		626		626		626		626		626	

Table 5: Stock-day regression — price reversals and adjustments

This table contains estimates from regressions of the stock-day average of the speed of price adjustment and the price reversals indicator variable after a dark pool trade onto stock-day level controls, fixed effects and the fraction of all dark pool trades that occur on a broker dark pool. The general form of the regression model is $\bar{y}_{jt} = \alpha_j + \gamma_t + \beta B\bar{D}P_{jt} + \rho'\bar{X}_{jt} + \bar{\varepsilon}_{jt}$ where α_j is a stock fixed effect, γ_t is a date fixed effect, $\bar{y}_{jt}$ is the stock-day average of the speed of price adjustment (Column 1) or price reversals (Column 2) for stock j and day t, $\bar{X}_{jt}$ is the stock-day average of a vector of controls including log of dollar trade size, log of trade price, the national best bid-ask spread at the time of the dark pool trade, the log of depth available at the national best bid and ask at the time of the dark pool trade, the log of total dollar volume traded across all trades and venues, and the log of total number of dark pool and lit trades, $B\bar{D}P_{jt}$ is the stock-day average of a dummy variable that takes the value 1 if the trade occurs on a broker dark pool and 0 otherwise (i.e. the fraction of broker dark pool trades out of all dark pool trades) and $\bar{\varepsilon}_{jt}$ is an error term. The speed of price adjustment is defined as the absolute difference between the mid-point thirty minutes and one minute after the dark pool trade time scaled by the prevailing mid-point and expressed in percent. The price reversals indicator takes the value one if the sign of the one minute mid-point return after the trade is opposite to the sign of the one to thirty minute mid-point return. We estimate the model using all stock-days from Jan 1, 2017 to Sept 30, 2019. Reported R^2 values relate to the within variation in the dependent variables. Standard errors are clustered at the stock level and t-statistics are in parentheses.

	(1)		(2)	
	Price Adjustment		Reversal Indicator	
Ln(Dollar Trade Size)	-0.0025	(-2.53)	0.0009	(1.88)
Ln($N_{DarkPool}$)	-0.0363	(-23.2)	-0.0194	(-28.5)
Ln(N_{Lit})	0.0977	(35.3)	0.0301	(24.8)
Ln(Price)	-0.1366	(-14.4)	0.0326	(7.80)
Ln(Dollar Volume)	0.0691	(24.1)	0.0173	(19.4)
Pre-cross Bid-ask Spread	0.1572	(14.9)	0.0470	(10.6)
Ln(Depth)	-0.0648	(-27.3)	-0.0298	(-28.6)
Broker Dark Pool	0.0053	(1.34)	-0.0069	(-3.65)
Fixed Effects	$N\&T$		$N\&T$	
R^2	0.13		0.10	
N_{obs}	242,825		242,825	
N_{stocks}	626		626	

Table 6: Trade-level summary statistics around pool closure

This table contains summary statistics for dark pool trades in stocks in the ASX All Ordinaries Index over three one month periods corresponding to the month after closure of three broker dark pools, operated by Merrill Lynch (March 6, 2017), UBS (April 1, 2019) and Citigroup (July 1, 2019) respectively. Dark pool trades from any remaining broker are included as are trades on Centre Point from the broker whose pool has recently closed. There are approximately 6.9 million trades in the sample pooled across these three windows, 5.7 million of which are on broker dark pools and the remaining 1.2 million on Centre Point from a broker whose pool has recently closed. Trade size is the dollar volume of the trade (measured in thousands). Price is the trade price in dollars. Total dollar value is the daily total dollar volume by stock and day across all venues. Pre-Cross Bid-ask Spread the bid-ask spread at the time of the dark pool trade. Absolute τ price impact is the absolute log difference between the mid-quote τ seconds after the trade and the prevailing mid-point at the time of the trade expressed in percent. τ bid-ask spread is the log bid-ask spread τ seconds after the trade expressed in percent. Price adjustment is the absolute difference between the mid-point thirty minutes and one minute after the dark pool trade time scaled by the prevailing mid-point and expressed in percent. Reversal Indicator takes the value one if the sign of the one minute mid-point return after the trade is opposite to the sign of the one to thirty minute mid-point return. The final two columns report the average of these variables for trades on broker dark pools ("BDP") and on Centre Point ("CP").

	Mean	SD	Min	25%	50%	75%	Max	Mean BDP	Mean CP
Trade Size (AUD '000s)	1.77	29.6	0.00	0.07	0.31	1.12	22024	1.63	2.66
Price (AUD)	19.7	32.6	0.06	3.89	8.39	20.0	231	20.1	18.3
Total Dollar Value (AUD 'm)	39.5	53.6	0.00	8.45	20.4	44.7	440	39.7	39.9
Pre-Cross Bid-ask Spread (%)	0.18	0.19	0.00	0.06	0.13	0.26	1.87	0.18	0.20
Abs. 500ms Price Impact (%)	0.00	0.02	0.00	0.00	0.00	0.00	0.25	0.00	0.01
Abs. 10s Price Impact (%)	0.01	0.04	0.00	0.00	0.00	0.00	0.42	0.01	0.01
Abs. 60s Price Impact (%)	0.04	0.08	0.00	0.00	0.00	0.04	0.66	0.04	0.04
Abs. 1800s Price Impact (%)	0.24	0.30	0.00	0.03	0.15	0.32	2.79	0.24	0.23
500ms Bid-Ask Spread (%)	0.19	0.19	0.00	0.06	0.13	0.26	1.87	0.18	0.21
10s Bid-Ask Spread (%)	0.19	0.19	0.00	0.06	0.13	0.26	1.87	0.18	0.21
60s Bid-Ask Spread (%)	0.19	0.20	0.00	0.06	0.13	0.26	1.90	0.18	0.21
1800s Bid-Ask Spread (%)	0.30	0.48	0.00	0.07	0.16	0.33	5.33	0.30	0.31
Price Adjustment (%)	0.20	0.27	0.00	0.00	0.13	0.28	2.42	0.20	0.20
Reversal Indicator	0.24	0.43	0.00	0.00	0.00	0.00	1.00	0.24	0.22

Table 7: Matching regression

This table contains average treatment effects obtained from comparing execution outcomes for broker dark pool trades with matched dark pool trades from Centre Point from brokers whose dark pools had recently closed. For each stock and each of the one month periods after the three pool closures described in Table 6, we match each broker dark pool trade to a Centre Point trade from the closing broker, using propensity score matching on log trade size, log trade price, log total dollar volume traded, best bid-ask spread and log depth at NBBO at the time of trade, date, time of day, and quadratic terms for date and time of day, keeping only trades that can be matched using a caliper of 0.25 standard deviations of the estimated propensity scores. We estimate the average treatment effect as the difference in means of the matched broker dark pool trades and Centre Point trades. Treatment effects by event are contained in Columns 1 — 3 with t-statistics from a test that the average effect across stocks is equal to zero in parenthesis below the estimated effect. Execution quality is measured using the same variables as defined in Table 2 and Table 6. Standard errors are clustered at the stock level and t-statistics are in parentheses.

	(1)	(2)	(3)
	BAML	UBS	Citi
Abs. Price Impact (60s)	-0.0084	-0.0044	-0.0031
	(-9.01)	(-8.99)	(-4.46)
Bid-ask Spread (60s)	-0.0264	-0.0016	-0.0023
	(-13.9)	(-4.06)	(-4.56)
Price Adjustment (%)	0.0051	0.0023	-0.0041
	(2.07)	(1.35)	(-1.72)
Reversal Indicator	0.0060	-0.0024	0.0011
	(2.00)	(-0.97)	(0.39)

Table 8: Center Point trading difference-in-differences around broker pool closures

This table contains estimates from difference-in-differences regressions for trades on an exchange dark pool around closures of broker dark pools. Centre Point trades from brokers closing their dark pools are the treated category while Centre Point trades from brokers that continue to operate are the control group. The regression model is $y_{ijbt} = \alpha_j + \gamma_t + \mu_b + \beta\tau + \rho' X_{ijbt} + \varepsilon_{ijbt}$ where y_{ijbt} is the absolute price impact or bid-ask spread at the 60s horizon for trade i in stock j by broker b on date t, α_j is a stock fixed effect, γ_t is a date fixed effect, μ_b is a fixed effect for each of the brokers closing their dark pools, τ is a treatment status indicator taking the value of one if the trade is from a closing broker after their broker dark pool closes, X_{ijbt} is a vector of controls and ε_{ijbt} is an error term. The controls include log of dollar trade size, log of trade price, best bid-ask spread and log depth at NBBO at the time of the trade, log of total dollar volume traded across all trades and venues and the log of total number of dark pool and lit trades, the time of day and its square. The model is estimated separately for each event including all trades in a one month window either side of the three Dark Pool closures (BAML, UBS and Citigroup). Reported R^2 values relate to the within variation in the dependent variables. Standard errors are clustered at the stock level and t-statistics are in parentheses.

	BAML				UBS				Citi			
	Abs. Price Impact		Bid-ask Spread		Abs. Price Impact		Bid-ask Spread		Abs. Price Impact		Bid-ask Spread	
Ln(Dollar Trade Size)	0.0012	(11.0)	-0.0014	(-9.59)	0.0012	(16.7)	0.0001	(1.21)	0.0012	(16.2)	0.0003	(6.25)
Ln(Price)	-0.0238	(-3.80)	-0.0555	(-1.44)	-0.0188	(-1.89)	-0.1053	(-5.31)	-0.0131	(-1.56)	-0.0787	(-4.52)
Pre-cross Bid-ask Spread	0.0051	(1.06)	0.6882	(14.6)	0.0360	(4.84)	0.5990	(31.5)	0.0302	(3.09)	0.5815	(27.7)
Ln(Depth)	-0.0161	(-17.0)	-0.0187	(-16.4)	-0.0110	(-13.5)	-0.0115	(-15.0)	-0.0118	(-15.2)	-0.0100	(-15.1)
Ln(Dollar Volume)	0.0179	(18.9)	0.0050	(4.04)	0.0149	(12.6)	0.0027	(3.74)	0.0180	(13.3)	0.0004	(0.60)
Time of day	-0.3645	(-31.8)	-0.3490	(-22.7)	-0.3426	(-31.6)	-0.1840	(-19.7)	-0.3620	(-32.1)	-0.1022	(-13.6)
Time of day (square)	0.0373	(31.9)	0.0389	(23.2)	0.0346	(31.9)	0.0202	(20.8)	0.0365	(32.6)	0.0112	(14.6)
$D_{ClosedPool}$	0.0029	(5.97)	0.0016	(2.80)	0.0024	(4.83)	0.0013	(3.21)	0.0005	(1.05)	0.0020	(5.70)
Treatment Effect	-0.0024	(-4.02)	0.0052	(5.36)	0.0006	(0.99)	-0.0011	(-1.80)	0.0000	(0.05)	0.0005	(1.13)
Fixed Effects	$N\&T$		$N\&T$		$N\&T$		$N\&T$		$N\&T$		$N\&T$	
R^2	0.06		0.26		0.06		0.22		0.07		0.27	
N_{obs}	3,990,056		3,990,056		5,135,189		5,135,189		6,428,942		6,428,942	
N_{stocks}	449		449		462		462		433		433	

Table 9: Stock-day regression for broker dark pool trades split by HFT access

This table contains estimates from regressions of stock-day averages of execution outcomes after a broker dark pool trade onto stock-day level controls, fixed effects and the fraction of trades that take place on broker dark pools that do not permit HFT activity. The regression model is $\bar{y}_{jt} = \alpha_j + \gamma_t + \beta \bar{Restricted}_{jt} + \rho' \bar{X}_{jt} + \bar{\varepsilon}_{jt}$ where α_j is a stock fixed effect, γ_t is a date fixed effect, $\bar{y}_{jt}$ is the stock-day average of execution outcomes for dark pool trades in stock j and day t (as defined in Table 2 with absolute price impact and bid-ask spreads measured at the 60s horizon), $\bar{X}_{jt}$ is the stock-day average of a vector of controls including log of dollar trade size, log of trade price, best bid-ask spread and log depth at NBBO at the time of the dark pool trade, the log of total dollar volume traded across all trades and venues, and the log of total number of dark pool and lit trades, $\bar{Restricted}_{jt}$ is the stock-day average of a dummy variable that takes the value 1 if the trade takes place on a broker dark pool that does not permits HFT activity and 0 otherwise and $\bar{\varepsilon}_{jt}$ is an error term. We estimate the model using trades on broker dark pools covering all stock-days from Jan 1, 2017 to Sept 30, 2019 including stock and date fixed effects and controls. Reported R^2 values relate to the within variation in the dependent variables. Standard errors are clustered at the stock level and t-statistics are in parentheses.

	(1)		(2)		(3)		(4)	
	Abs. PI		Spread		Price Adjustment		Reversal Indicator	
Ln(Dollar Trade Size)	0.0029	(10.8)	-0.0003	(-1.37)	0.0020	(2.58)	0.0029	(6.17)
Ln($N_{DarkPool}$)	-0.0081	(-23.2)	-0.0033	(-10.1)	-0.0059	(-5.97)	-0.0040	(-6.77)
Ln(N_{Lit})	0.0171	(21.6)	-0.0060	(-6.78)	0.0806	(28.9)	0.0368	(23.0)
Ln(Price)	-0.0115	(-4.80)	-0.0380	(-10.5)	-0.1144	(-11.8)	0.0367	(6.69)
Ln(Dollar Volume)	0.0116	(18.8)	0.0033	(5.87)	0.0499	(21.0)	0.0195	(17.6)
Pre-cross Bid-ask Spread	0.0465	(7.86)	0.8744	(115)	0.1829	(12.6)	0.0722	(8.29)
Ln(Depth)	-0.0192	(-30.6)	-0.0074	(-9.48)	-0.0539	(-24.7)	-0.0416	(-31.4)
HFT-restricted Pool	-0.0080	(-3.94)	-0.0037	(-1.69)	-0.0236	(-4.25)	-0.0099	(-2.40)
Fixed Effects	$N\&T$		$N\&T$		$N\&T$		$N\&T$	
R^2	0.10		0.77		0.13		0.09	
N_{obs}	192,068		192,068		192,068		192,068	
N_{stocks}	563		563		563		563	

Table 10: Stock-day regression split by trade size

This table contains estimates from regressions of stock-day averages of execution outcomes after a dark pool trade onto stock-day level controls, fixed effects, the fraction of trades below the 10$^{\text{th}}$ percentile in terms of dollar volume by stock-week, and the fraction of all dark pool trades that occur on a broker dark pool that are above or equal to and below this threshold respectively. The regression model is $\bar{y}_{jt} = \alpha_j + \gamma_t + \beta_0 \bar{D}^{size\leq\bar{v}}_{jt} + \beta_1 B\bar{D}P^{size\leq\bar{v}}_{jt} + \beta_2 B\bar{D}P^{size>\bar{v}}_{jt} + \rho'\bar{X}_{jt} + \bar{\varepsilon}_{jt}$ where α_j is a stock fixed effect, γ_t is a date fixed effect, $\bar{y}_{jt}$ is the stock-day average of execution outcomes for dark pool trades in stock j and day t (as defined in Table 2 with absolute price impact and bid-ask spreads measured at the 60s horizon), $\bar{X}_{jt}$ is the stock-day average of a vector of controls including log of dollar trade size, log of trade price, the best bid-ask spread and log depth at NBBO at the time of the trade, log of total dollar volume traded across all trades and venues, and the log of total number of dark pool and lit trades, $\bar{D}^{size\leq\bar{v}}_{jt}$ is the stock-day average of a dummy variable that takes the value 1 if trade size is at or below the lowest 10$^{\text{th}}$ percentile for that stock-week, $B\bar{D}P^{size\leq\bar{v}}_{jt}$ is the stock-day average of a dummy variable that takes the value 1 if the trade occurs on a broker dark pool and trade size is at or below the lowest 10$^{\text{th}}$ percentile, $B\bar{D}P^{size>\bar{v}}_{jt}$ is the stock-day average of a dummy variable that takes the value 1 if the trade occurs on a broker dark pool and trade size is above the lowest 10$^{\text{th}}$ percentile and $\bar{\varepsilon}_{jt}$ is an error term. We estimate the model using all stock-days from Jan 1, 2017 to Sept 30, 2019 including stock and date fixed effects and controls. Reported R^2 values relate to the within variation in the dependent variables. Standard errors are clustered at the stock level and t-statistics are in parentheses.

	(1)		(2)		(3)		(4)	
	Abs. PI		Spread		Price Adjustment		Reversal Indicator	
Ln(Dollar Trade Size)	-0.0008	(-1.94)	-0.0023	(-5.14)	-0.0023	(-2.20)	0.0006	(1.18)
Ln($N_{DarkPool}$)	-0.0203	(-36.4)	-0.0057	(-8.54)	-0.0362	(-23.1)	-0.0195	(-28.7)
Ln(N_{Lit})	0.0215	(27.3)	-0.0068	(-6.04)	0.0977	(35.4)	0.0300	(24.5)
Ln(Price)	-0.0092	(-4.03)	-0.0294	(-5.86)	-0.1367	(-14.4)	0.0327	(7.82)
Ln(Dollar Volume)	0.0139	(20.2)	0.0046	(5.94)	0.0690	(24.1)	0.0174	(19.5)
Pre-cross Bid-ask Spread	0.0368	(10.8)	0.9250	(108)	0.1574	(14.9)	0.0470	(10.6)
Ln(Depth)	-0.0189	(-27.9)	-0.0058	(-3.84)	-0.0648	(-27.3)	-0.0297	(-28.6)
$D^{size\leq\bar{v}}$	-0.0001	(-0.02)	0.0044	(0.90)	0.0071	(0.84)	-0.0069	(-1.62)
Broker Dark Pool $\times D^{size\leq\bar{v}}$	-0.0188	(-4.52)	-0.0062	(-0.89)	-0.0119	(-1.01)	-0.0077	(-1.21)
Broker Dark Pool $\times D^{size>\bar{v}}$	-0.0101	(-7.43)	-0.0052	(-2.47)	0.0093	(1.94)	-0.0066	(-3.20)
Fixed Effects	*N&T*		*N&T*		*N&T*		*N&T*	
R^2	0.11		0.75		0.13		0.10	
N_{obs}	242,844		242,844		242,844		242,844	
N_{stocks}	626		626		626		626	