Journal of Financial Economics

Differential access to dark markets and execution outcomes ☆

James Brugler, Carole Comerton-Forde *

The University of Melbourne, 198 Berkeley St, Carlton, 3010, VIC, Australia

ARTICLE INFO

Keywords:
Dark trading
Execution outcomes
High frequency trading
Segmentation

ABSTRACT

Dark pools can restrict access for specific trader types. We compare execution outcomes between dark pools that restrict high frequency trader access and those that do not. We find that trades executed in dark pools with more access restrictions have less order flow information leakage, adverse selection risk and post-trade order imbalances than trades in less restricted pools. Evidence from exogenous dark pool closures demonstrates that these differences are causal. The ability to segment order flow can benefit investors because it allows them to make trade-offs between execution risk and information leakage across different dark venues.

1. Introduction

Dark pools are a ubiquitous feature of modern equities markets, accounting for around 13% of consolidated turnover in US equities and 7% in European equities.[1] Institutional investors use dark pools to minimize their trading footprint and price impact (Norges Bank Investment Management, 2015). Existing literature typically assumes dark pools are homogeneous, but in reality, they differ on a range of dimensions including how prices are set; the prices where trades can occur; whether dark orders interact with pre-trade transparent ("lit") orders; and who can access the venue. We examine access restrictions (i.e. restrictions on specific traders that can or cannot access a venue), an important form of heterogeneity not previously considered in the literature.

Regulation of access restrictions differs across jurisdictions. In the United States (US), Alternative Trading Systems (ATS) can impose access restrictions provided their market share is below 5%.[2] ATS must also make public disclosures about who can access their pool, their order handling rules and execution practices. In the European Union (EU), Broker Crossing Networks (BCN) allowed brokers to use access restrictions, however these trading venues were outlawed with the introduction of the Markets in Financial Instruments Directive II (MiFID II) in 2018. In Australia, the setting for this study, broker dark pools (equivalent to US ATS) can impose access restrictions but exchange dark pools, which are unique to Australia, cannot. These policy choices were made without empirical evidence on how access restrictions impact execution outcomes.

We provide this evidence by answering three questions. First, are there differences in execution outcomes between pools with restricted versus unrestricted access? Second, if there are differences in execution outcomes, are these causal? Third, can differences be attributed to variation in access by trader category across types of dark pools?

When building or exiting a position, institutional investors make choices about how to execute their orders to minimize the effect of trading on portfolio returns. They typically select an execution algorithm that splits larger "parent" orders into many smaller "child" orders. These algorithms decide the size of each child order, how frequently to submit them, and which venue type (e.g. lit vs. dark) to

☆ Nikolai Roussanov was the editor for this article. We thank an anonymous referee, Amber Anand, Jonathan Brogaard, Thomas Ernst, Sean Foley, Björn Hagströmer, Terry Hendershott, Patrick Kelly, Albert Menkveld, Richard Phillips, Dominik Rösch, Yazid Sharaiha, Andriy Shkilko, Yuanji Wen, Bart Zhou Yueshen, Zhuo Zhong and seminar participants at Hong Kong Polytechnic University, University of Melbourne, the Microstructure Exchange, the Microstructure Online Seminars Asia Pacific, Sydney Market Microstructure and Digital Finance Meeting and the TMX Markets Forum for helpful comments. We thank numerous industry participants for their insights about institutional arrangements. We thank Stanley Zhou for research assistance. We thank Cboe Australia for access to their broker-venue-trade data. Comerton-Forde is a Centre for Economic Policy Research Research Fellow, 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] Dark pool is a colloquial term used to refer to trading venues that do not offer pre-trade transparency. The market share statistics provided are from Rosenblatt Securities, Let there be light, September 2024, US and European editions.
[2] All ATS in the US have a market share well below 5%.
[3] In the remainder of the paper we use the term "order anticipation strategies" to refer to any strategy based on inferring a counterparty's trading intentions.

Received 3 April 2024; Received in revised form 1 May 2025; Accepted 5 May 2025
Available online 26 May 2025
0304-405X/

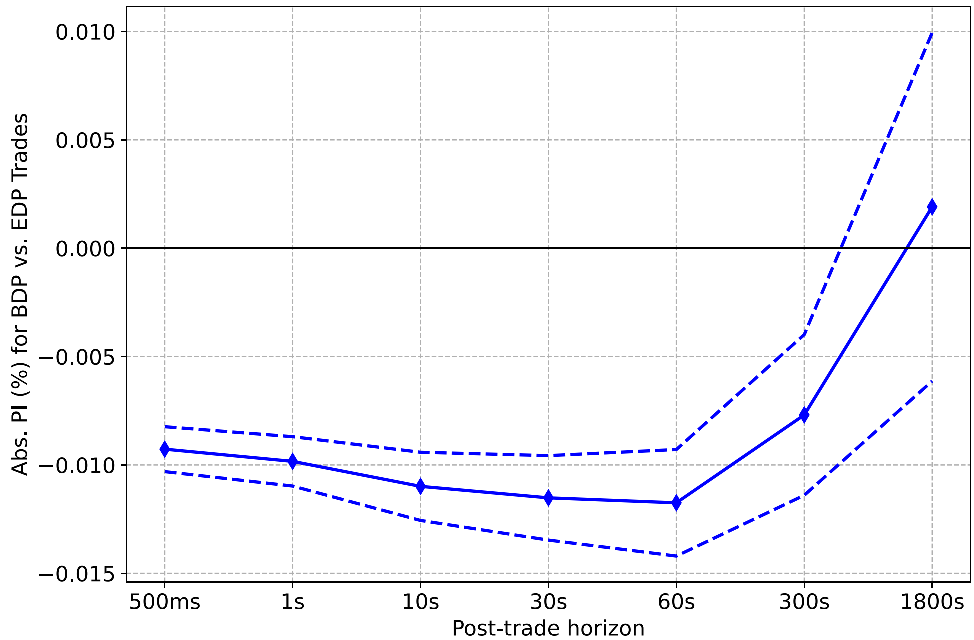

Fig. 1. Information leakage for broker dark pools trades vs. exchange dark pools trades.
Estimated differences in absolute price change following broker dark pool trades and exchange dark pool trades from panel regressions. Each point on the solid line represents the coefficient on a dummy for a trade taking place on a broker dark pool compared with an exchange dark pool, obtained from a panel regression of post-trade absolute price changes onto this variable, stock-day and trade-level controls and stock and date fixed effects. Negative coefficients indicate that broker dark pool trades have lower information leakage than otherwise comparable exchange dark pool trades. The model is estimated at horizons of 500ms, 1s, 10s, 30s, 60s, 300s and 1800s respectively. Dashed lines are 95% confidence intervals using standard errors that are clustered at the stock-level.

use. Access restrictions allow dark pool users to also control the types of counter-parties they interact with. When a dark pool trade involves a child order, the counter-party can learn information about current and future order flow, and therefore better predict future prices. This order flow information leakage can contribute to back-running strategies (Yang and Zhu, 2019), short-term directional strategies (Korajczyk and Murphy, 2018, van Kervel and Menkveld, 2019), order anticipation strategies (Sağlam, 2020 and Hirschey, 2021) and sniping (Malinova and Park, 2020). Institutional investors want to avoid trading with counter-parties that pursue these types of strategies.[3]

These strategies are usually attributed to high frequency traders (HFTs) and electronic liquidity providers (ELPs) and can adversely affect execution outcomes for institutional orders.[4] This is because as the counter-party to the institutional trade, the HFT/ELP, learns additional information about the order flow. Specifically, they learn about the direction of the institutional order flow, which allows them to anticipate future order flow better. Although other traders observe the dark execution, they do not observe trade direction because the trades typically occur at mid-point. Therefore, only the counter-party to the dark trade can pursue these strategies. Access restrictions typically apply to HFTs/ELPs or principal traders. Our analysis focuses on this order flow information leakage, rather than leakage of information about company fundamentals.[5] For simplicity, we refer to order flow information leakage as "information leakage".

The Australian market is an ideal setting for our research because (i) there are two types of dark pools that differ in terms of trader access and (ii) we are able to identify the venue type at an individual trade level.[6] We analyze all dark trades in a large sample of Australian stocks from January 1, 2017 to September 30, 2019. We categorize dark trades based on whether they take place in an exchange-operated dark pool ("exchange dark pool trades") or a broker-operated dark pool ("broker dark pool trades"). We construct measures of execution outcomes at the trade-level that capture information leakage and adverse selection risk. 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 microstructure variables such as 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.

To estimate information leakage, we compute the absolute log change in the NBBO mid-quote ("absolute price change") after every dark pool trade in our sample over horizons spanning 500ms to 30 min. Our logic is that, after controlling for relevant observable and unobservable factors that are correlated with the trading venue (exchange vs. broker dark pool), the average magnitude of price movements after a trade reflects information transmission from the trade to the rest of the market. We similarly compute the post-trade NBBO bid–ask spread at 500ms to 30 min horizons, and use this as a measure of adverse selection risk faced by liquidity providers.

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. Fig. 1 presents the estimated coefficient on venue type (broker vs. exchange dark pools) on absolute price change, conditional on controls and fixed effects at each horizon. The effect at the 60s horizon is approximately −1 bp, indicating that broker

[4] HFTs are professional traders that trade in a proprietary capacity and typically use extremely low latency technology to generate, route, and execute orders. A common high frequency strategy is two-sided liquidity provision, where they actively monitor the order book and make frequent revisions to their displayed limit orders. Traders executing this strategy are typically referred to as ELPs.

[5] Conrad and Wahal (2020) argue that information risk for liquidity providers is now defined by correlated liquidity-demanding trades over (short) inventory holding periods rather than the arrival of fundamental information.

[6] In contrast, in the US market individual dark pool trades cannot be identified and researchers must rely on weekly or bi-weekly FINRA dark trading reports.

dark pool trades have lower information leakage than comparable exchange dark pool trades. This is approximately one-sixth of the mean value in our sample. There are no statistically significant differences in execution outcomes 30 min after the trade. Similar patterns exist for adverse selection across various time horizons.

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 unobservable differences in the investor composition of brokers who operate dark pools and other brokers who can only execute dark trades on exchange dark pools.

Identifying this causal effect requires a source of exogenous variation regarding where an order gets routed. Our solution exploits the fact that three broker-operated pools cease operations during our sample. These closures are plausibly exogenous as they are motivated by broker concerns about regulatory and compliance risks rather than market share or execution outcomes.[7] After their pool closes, if their customers demand dark executions, these brokers have no choice but to execute dark orders on exchange dark pools. We assume that after a pool closure, the sample of exchange dark pool trades contains some trades that would have been executed on a broker dark pool if it were still operating. 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, matching trades within stocks based on trade and order book details. We keep only closely matched trades. The matching exercise confirms our panel regression results: trades on broker dark pools have lower information leakage and less adverse selection risk for liquidity providers.

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, we exploit the heterogeneity in access across broker pools. We show that trades in broker pools that give customers the ability to opt-out of trading with HFTs/ELPs have significantly higher information leakage and higher adverse selection risk than trades in broker pools that completely prohibit HFT/ELP order flow. Second, we examine trading activity and order imbalance variables after dark trades for evidence of differences in HFT/ELP participation across the pool types. Exchange dark pool trades, which are more likely to involve HFTs/ELPs, are followed by significantly more lit market trading activity, order imbalances and lower quote-to-trade ratios than broker dark pool trades. This is consistent with HFTs/ELPs learning from their interaction with orders in dark pools and adjusting their strategies away from providing liquidity to demanding liquidity — worsening trading outcomes for the counter-party.

Our results relate to execution outcomes, conditional on a trade taking place. Execution outcomes matter most for understanding information leakage, which is the main focus of our research. "Execution risk", the probability of execution conditional on receiving an order, is also an important component of trading costs. Access restrictions limit the possible counter-parties to any given order and reduces the amount of liquidity in restricted dark pools. Dark pools with different access restrictions offer different trade-offs between execution risk and execution outcomes, and in equilibrium, traders choose the venue (or execution algorithm) that balances their need for immediacy against the cost of price impact on future orders. We also examine relative market shares in broker and exchange dark pools over time. During periods of high volatility, when the costs of non-execution are greater, dark pool volumes fall relative to lit venues. However, broker dark pools maintain their aggregate market share relative to exchange dark pools, suggesting that time-variation in execution risk across dark pool categories is not of first-order importance for order routing decisions.

Contribution to regulatory policy: Regulators try to balance goals such as transparency, fairness, price discovery and liquidity. Access restrictions primarily trade-off fairness and two dimensions of liquidity: execution risk and execution outcomes. Permitting venues to restrict access to HFT/ELPs gives traders more choice between venues offering different information leakage and execution risk profiles. We offer insights into policy in other jurisdictions without sufficiently granular data to perform similar analyses. The US approach allowing access restrictions but mandating disclosures of execution practices and permitted trader types (Form ATS-N), and weekly disclosures of ATS volumes (FINRA Rule 4552) helps traders make more informed choices about their trading strategies. In contrast, the EU ban on BCNs removed this choice entirely, forcing traders to use venues with no control over their counterparties.

Contribution to literature: Our analysis of dark pool heterogeneity related to access restrictions complements (Menkveld et al., 2017) which considers other dimensions of heterogeneity — cost and immediacy. They demonstrate 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. In our setting, where all dark pools are mid-point dark pools, we show that conditional on execution, access restrictions improve execution outcomes.

Our research also builds on the literature on dark pools and order flow segmentation. Comerton-Forde et al. (2018) and Hatheway et al. (2017) find detrimental effects of order flow segmentation in dark pools on lit liquidity and liquidity providers in lit markets. In contrast, we show that segmentation can improve execution outcomes by reducing the likelihood of trading with HFT/ELPs

Our results are also consistent with the literature on high frequency trading and institutional trading costs. Hirschey (2021) finds 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. Van Kervel and Menkveld (2019) also show that when setting their trading intensity institutions trade-off higher speculative profits against increased risk of being detected by HFTs/ELPs. Our results complement these by considering how the trade-off between choosing a venue with or without HFTs/ELPs impacts execution outcomes such as information leakage. Battalio et al. (2024) show that child executions routed to an ELP can increase overall implementation shortfall. These results are consistent with our findings that limiting interactions with HFTs/ELPs improves execution outcomes.

Our results also relate to two papers on broker routing decisions. Anand et al. (2021) examine routing to affiliated vs. unaffiliated pools and Battalio et al. (2016) examine routing decisions related to payment for order flow. Both papers suggest that conflicts of interest broker routing decisions can lead to bad execution outcomes.

2. Institutional details

The Australian equity market is the fifteenth largest in the world with an average market capitalization of between AUD 1.7 and 2.1 trillion during our sample period (World Federation of Exchanges, 2019; ASX, 2021). Trading activity is fragmented across two exchanges: the Australian Securities Exchange (ASX) and Cboe Australia (Cboe), and 13 broker-operated dark pools. During our sample period ASX accounts for approximately 75.5% of trading by total dollar volume (including opening and closing auctions), Cboe 10% and broker dark pools 3.3%. Off-exchange trading accounts for the remaining 10% of dollar volume. The overall level of dark trading in Australia is approximately 13%, which is comparable to the 13% observed in the US and higher than the 7% in the Europe.

Trading is governed by the Australian Securities and Investments Commission (ASIC) Market Integrity Rules (MIRs). Orders must be pre-trade transparent unless they meet one of the exceptions set out in the

[7] It is important that order anticipation strategies are not prohibited by regulation and so variation in information leakage within dark pools does not contribute to compliance and regulatory risk.

Table 1
Broker dark pool classifications.

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 cannot 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 and of flow from aggregator algos	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 and of flow from aggregator algos	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 and of flow from aggregator algos	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 and of flow from aggregator algos	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 and of flow from aggregator algos	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 and of flow from aggregator algos	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 and of flow from aggregator algos	Yes, accepts orders from aggregator algos	No

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 HFTs/ELPs are explicitly excluded ("Restricted") or HFTs/ELPs are possibly present but other traders can opt-in to avoid executing against them ("Opt-in to restrictions").

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 National Best Bid and Offer (NBBO), which is the best bid and offer price set across the ASX and Cboe transparent limit order books. These trades must occur at the NBBO mid-point or a designated minimum price increment (tick) within the NBBO. The price improvement requirement is equivalent to what the US markets calls a trade-at rule. There are other pre-trade transparency exceptions for block and portfolio trades.[8]

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 within the NBBO. Traders can submit dark orders that interact with both the lit and dark order books, or choose to interact only with dark orders. This feature is unique to Centre Point. Cboe operates a single electronic limit order book that allows both displayed and hidden orders to be submitted. Cboe orders are matched based on price-display-time priority and market orders will automatically interact with dark liquidity. These are similar to the hidden order types used by US exchanges. Trades in broker dark pools must also satisfy the pre-trade transparency exceptions and must be immediately reported to either ASX or Cboe.

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 dark pools. Table 1 provides a summary of the access restrictions in each pool during our sample period.[9] Four pools completely prohibit access

[8] Block trades are negotiated away from the market. Thresholds for block trades are based on the stock's average daily volume: AUD 1 m 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 an aggregate transaction value of at least AUD 5 m, across at least 10 different stocks and each transaction must have a value of at least AUD 200k.

[9] Dark pool classifications are made based on regulatory disclosures and confirmed through discussions with dark pool operators.

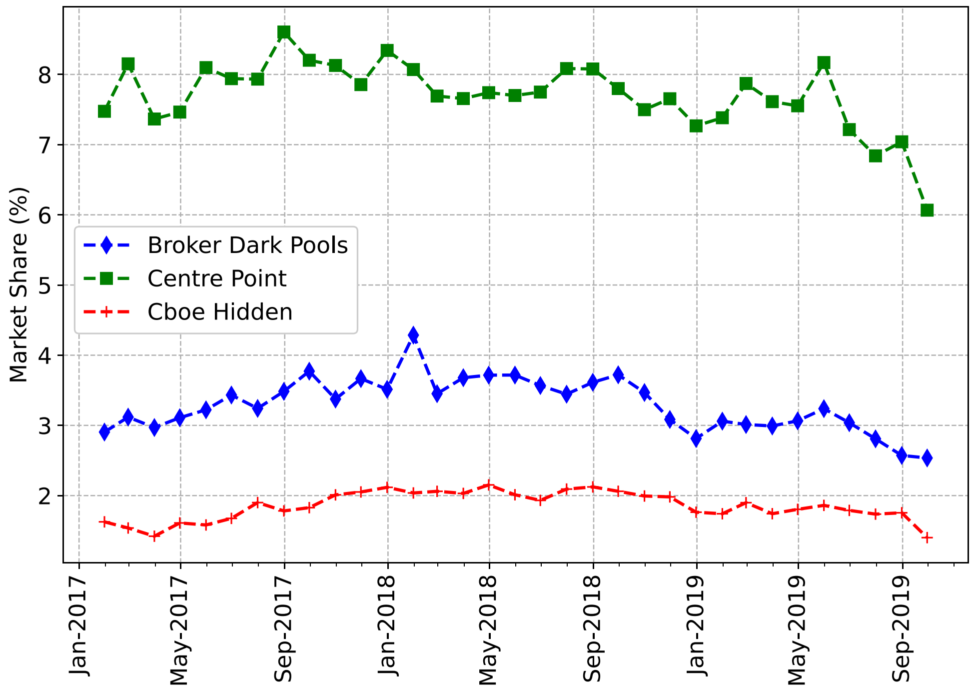

Fig. 2. Broker dark pools and exchange dark pools market shares.
Time-series of weekly market shares of total dollar volume traded on broker dark pools, Centre Point and Cboe Hidden orders. Market shares are calculated using trades in all stocks in our sample by week. Broker dark pool market share is summed across all broker dark pools in our sample.

to HFTs/ELPs and principal trading and nine pools allow customers to opt-in to restricting access by counter-party type. All broker-operated pools also offer minimum acceptable quantity (MAQ) functionality that allows traders to specify that the order can only execute if the MAQ is met (as do Centre Point and Cboe). Single Dealer Platforms (SDPs) like those that operate in the US do not exist in the Australian market; all broker dark pools match customer order flow.

Both exchanges charge execution fees on both sides of a trade. On the ASX, trade fees are higher for dark trades compared with displayed orders. A typical trade that executes on the ASX's TradeMatch is charged a fee of 0.15 bps per side, capped at AUD 75 while a standard Centre Point trade is charged a fee of 0.5 bps per side (ASX, 2022). Cboe implements maker-taker pricing where the liquidity supplier is charged 0.06 bps per trade and the liquidity demander is charged 0.12 bps per trade, regardless of whether the trade is hidden or displayed (Cboe, 2023a, Cboe, 2023b). Brokers pay trade reporting fees of 0.04 bps per side for trades matched in their dark pools, but these fees are capped at AUD 1000 per month for each broker. The lower fees for trade reporting vs. other trade types provide an incentive for brokers to execute in their dark pools.

Unlike the US market, the Australian market does not have an order protection rule and best execution rules are less prescriptive. In Australia, like in Europe, best execution obligations differ for retail vs. institutional customers. For retail customers, the broker must execute at the best price. For institutional customers, brokers must have in place a best execution policy that defines the factors they consider when executing client orders which may include price, costs, speed, the likelihood of execution or any other relevant outcome, or any combination of those outcomes.[10] While these principles-based obligations are less prescriptive than what exists in North America, principles-based best execution regulations are used in most jurisdictions outside of North America. Australian regulations also prohibit all forms of payment for order flow including exchange rebates.

[10] See Part 3.8 of the ASIC MIR for details.

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 LSEG's Datascope Select database.[11] 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 broker dark pool trades executed within the NBBO (NBBO trades).

For every trade, we observe the price, volume, the time to the nearest millisecond and a trade qualifier designating the type of trade. The qualifiers are sufficient to identify dark trades executed on Centre Point or Cboe directly but do not identify the venue of broker dark pool trades. The MIRs require brokers to immediately report all off-exchange trades to either ASX or Cboe. The broker executing the trade and the execution venue are then made available in the ASX and Cboe course of sales data three days after the trade.[12] We use the course of sales data to identify the brokers executing trades on Centre Point, hidden trades on Cboe and the specific broker dark pool for each NBBO trade. Further details about these data are provided in Appendix A. 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 Fig. 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.[13]

[11] The All Ordinaries Index is a market capitalization index of the 500 largest stocks trading on the ASX and represents around 90% of the total value of securities trading on the ASX.

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

[13] The mean and distribution of dollar trade size across the three dark pool trade categories are presented in Figure A.1 in Appendix A. Average trade sizes

Broker dark pools are operated by institutional brokers and typically service only institutional customers. In contrast, Centre Point and Cboe hidden orders can be used by both retail and institutional traders. In practice the retail share of exchange dark pool trading is low, relative to the US market. One reason why retail brokers avoid Centre Point is the higher execution fees.[14]

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 before 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.

3.1. Measuring execution outcomes for dark pool trades

We hypothesize that traders can learn about other traders' future order flow by participating in a dark pool. Large institutional trades are often split into many sequential "child" orders that are submitted according to an execution algorithm (van Kervel and Menkveld, 2019). Counter-parties to dark child orders learn about other traders' current order flow, which gives them an advantage in anticipating future order flow and in turn future prices. In addition to knowing about the trade itself, which all traders observe, counterparties can also observe the direction of the order flow. Dark executions, therefore, can contribute to order flow information leakage.[15] We are interested in determining whether there are systematic differences in this type of information leakage across dark pool categories.

Most dark trades (92% in our sample) take place at the midquote and trade direction cannot be defined because neither counter-party is fully crossing the spread to trade. However, the assignment of trade direction is not important for studying information leakage in our context. Large institutions execute child orders using a range of aggressiveness from most aggressive (via market orders) to most passive (via a limit order). When a large institution uses a dark pool, information leakage regarding future order flow can be inferred by the counterparty regardless of the aggressiveness of the trader on the other side of the trade. What matters most is whether an execution occurs between an institution and another trader who is trying to infer future order flow, not whether the institution would be willing to cross the spread on that particular child order.[16]

To measure information leakage on a trade-by-trade basis, we compute the absolute (midquote) price change following each trade across a range of time intervals:

$$Abs\Delta P_{ijt} = 100 \times \left| \log\left(M_{ijt}^{s+\tau}\right) - \log\left(M_{ijt}^{s}\right) \right| \quad (1)$$

where $M_{ijt}^{s+\tau}$ is the mid-point of the ith 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 before the trade. We estimate Eq. (1) over periods of 500ms, 1s, 10s, 30s, 60s, 300s and 1800s after the trade. These capture effects at both short horizons and long horizons, helping us to delineate between temporary and permanent effects. These intervals also contain the recommended maximum horizons for capturing price impact in large and small stocks in modern markets according to Conrad and Wahal (2020).

Larger average values of absolute price change for one venue type compared with another indicate a larger market reaction to trades in that venue type. If we can adequately control for confounding factors that are potentially correlated with dark pool type and future price movements, then systematic differences in market reactions across venues should reflect differences in the average quantity of information transmission.

Crucial to this interpretation is that we adequately control for confounding factors, both observable and unobservable. Potential observable confounding factors include trade size, the state of the order book, and the level of trading activity. Unobservable factors include not only stock and time components that can be dealt with via fixed effects but also possibly more complex endogeneity factors such as selection effects at the trader and broker level, requiring a careful identification strategy.

For each dark pool trade in our sample, we also calculate the percentage bid–ask spread at the same post-trade intervals as we use for absolute price change. The interpretation of bid–ask spreads is relatively straightforward. Larger values of bid–ask spreads after the trade takes place indicate relatively worse liquidity after execution and higher future trading costs.[17] After adequately controlling for both observed and unobserved factors that are potentially correlated with venue type, we interpret wider bid–ask spreads to reflect updated expectations of information or liquidity risk for liquidity providers.

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 reflects strategic decisions made by investors, who choose which brokers to send parent orders to, and brokers, who offer different suites of execution algorithms to their investors. While brokers do not have discretion to overrule the execution algorithm chosen by the investor, and so cannot "cream-skim" orders with less expected price impact (e.g. small orders, the last slices of a large order), observed differences in average execution outcomes across venue categories may reflect differences in the average characteristics of orders that are submitted across venue categories.

For example, orders from all traders can execute in exchange dark pools, but only a brokers' customers can trade on a 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 or may contain more information than the market average. These kinds of trades will likely have a larger effect on the future midquote price regardless of where they are executed. Differences in trading fees across venues may also drive differences in the types of orders submitted.

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

trend down across all three categories, while the distributions are relatively similar. Most trades are for $5000 or less in each venue.

[14] For example, we compute the market share of retail trading on Centre Point using two proxies. First, online retail brokers, where we know with certainty the brokers service only retail clients (the lower limit of retail market share), and second, full-service retail brokers who may also service some small institutional customers. Market share is defined as the number of trades with at least one retail broker as a counterparty, scaled by 2× the total number of Centre Point trades. Figure A.2 in Appendix A plots this market share by week. Retail brokers account for between 3% and 7% of trades in Centre Point on average in our sample.

[15] Recent empirical work documents the importance of information about future order flow in determining future prices and market quality (see e.g. van Kervel and Menkveld, 2019, Sağlam, 2020 and Hirschey, 2021). O'Hara (2015) notes the complex interaction between information and order flow in modern computer-based markets. She argues that the distinction between inventory and information components of the spread is no longer meaningful.

[16] To establish trade direction, one would need dark order submission information including order sequence and limit price instructions. Similar to parent order level data, such data are highly sensitive as they would reveal institutional trading strategies. As a result, these data are rarely available to academic researchers.

[17] Absolute price change and bid–ask spreads are winsorized at the 99th percentile by week.

Table 2
Stock-day summary statistics.

	Mean	SD	Min	25%	50%	75%	Max
Trade size (AUD '000s)	2.02	13.3	0.00	0.44	1.00	2.10	3,403
Total dollar value (AUD 'm)	13.8	33.0	0.00	0.51	2.99	13.0	1,458
Price (AUD)	8.64	18.8	0.00	1.40	3.38	7.76	241
Daily average dollar depth (AUD '000s)	101	227	1.38	12.5	33.2	95.8	2,126
Daily average bid–ask spread (%)	0.57	0.78	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	0.51	3.92
Abs. 500ms price change (%)	0.01	0.03	0.00	0.00	0.01	0.01	0.31
Abs. 10s price change (%)	0.03	0.05	0.00	0.01	0.02	0.03	0.48
Abs. 60s price change (%)	0.06	0.08	0.00	0.02	0.04	0.06	0.84
Abs. 1800s price change (%)	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
60s dollar volume (AUD '000s)	39.2	75.2	0.00	3.72	13.2	42.7	2,018
60s Abs. order imbalance (AUD '000s)	18.0	28.1	0.00	2.68	8.16	21.5	650
60s Quote-to-trade ratio	3.96	1.35	0.06	3.09	3.91	4.72	28.0

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. All variables are averaged across all dark pool trades by stock-day unless otherwise stated. Trade size is the dollar volume of the trade in thousands of dollars. Total dollar value is the total amount traded summed across all venues for that stock-day in millions of dollars. Price is the average trade price in dollars. Daily Average Bid–ask Spread is the time-weighted average log bid–ask spread at the NBBO for that stock and day. Daily Average Dollar Depth is the time-weighted average depth available at the NBBO for that stock and day 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, (i.e. it is the average proportion of dark pool trades on broker dark pools). Pre-Cross Bid–ask Spread is the log bid–ask spread at the NBBO at the time of the trade. Absolute τ price change 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 at the NBBO τ seconds after the trade, expressed in percent. 60s dollar volume and abs. order imbalance are computed across all exchanges over the 60s after the dark pool trade, in thousands of dollars. 60s quote-to-trade ratio is the number of NBBO quote changes divided by number of trades over the interval.

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 some order flow from broker dark pools to exchange dark pools. 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 outcomes 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 before the trade and whether 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:

$$y_{ijt} = \alpha_j + \gamma_t + \beta BDP_{ijt} + \rho' X_{ijt} + \varepsilon_{ijt} \tag{2}$$

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 Eq. (2) 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:

$$\bar{y}_{jt} = \alpha_j + \gamma_t + \beta B\bar{D}P_{jt} + \rho' \bar{X}_{jt} + \bar{\varepsilon}_{jt} \tag{3}$$

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 $\bar{\gamma}_t = \gamma_t$. Eq. (3) follows directly from taking the stock-day average of the data-generating process in Eq. (2), implying that we can recover the parameters of Eq. (2) from a regression of the stock-day average of execution outcomes for all dark pool trades onto the fraction of dark pool trades 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 using sweep orders which first execute against any dark liquidity resting in the Centre Point order book and then execute 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 same millisecond.[18]

Table 2 contains summary statistics for our stock-day panel formed from our sample of dark pool trades. The unit of observation is the stock-day level and every stock-day with at least one dark pool trade is included. The average dark pool trade size is AUD 2020 but, as

[18] Analysis that does not eliminate sweeps shows consistent results. These results are presented with other robustness tests in Section 7.

Table 3
Information leakage panel regressions.

	(1)		(2)		(3)		(4)		(5)		(6)		(7)	
	500ms		1s		10s		30s		60s		300s		1800s	
Ln(Dollar trade size)	−0.0009	(−5.19)	−0.0008	(−4.21)	−0.0005	(−2.10)	−0.0006	(−2.06)	−0.0007	(−1.86)	−0.0020	(−3.74)	−0.0046	(−4.49)
Ln($N_{dark pool}$)	−0.0083	(−33.4)	−0.0094	(−33.1)	−0.0131	(−35.6)	−0.0167	(−36.9)	−0.0203	(−36.4)	−0.0315	(−36.3)	−0.0478	(−27.8)
Ln(N_{Lit})	0.0056	(17.9)	0.0064	(17.6)	0.0111	(21.9)	0.0158	(25.7)	0.0215	(27.4)	0.0472	(33.2)	0.1026	(34.8)
Ln(Price)	−0.0001	(−0.18)	0.0002	(0.23)	−0.0009	(−0.69)	−0.0043	(−2.50)	−0.0092	(−4.05)	−0.0406	(−9.32)	−0.1422	(−13.9)
Ln(Dollar volume)	0.0032	(14.5)	0.0035	(13.6)	0.0062	(16.1)	0.0098	(19.7)	0.0138	(20.1)	0.0317	(22.6)	0.0720	(23.2)
Pre-cross bid–ask spread	0.0069	(5.95)	0.0098	(7.33)	0.0192	(9.81)	0.0262	(10.3)	0.0367	(10.8)	0.0798	(12.8)	0.1808	(15.6)
Ln(Depth)	−0.0047	(−19.2)	−0.0056	(−20.1)	−0.0102	(−25.0)	−0.0144	(−27.7)	−0.0189	(−27.9)	−0.0360	(−29.3)	−0.0699	(−27.5)
Broker dark pool	−0.0093	(−17.5)	−0.0098	(−16.9)	−0.0110	(−13.7)	−0.0115	(−11.6)	−0.0117	(−9.39)	−0.0077	(−4.08)	0.0019	(0.47)
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}	242,840		242,840		242,840		242,840		242,840		242,840		242,840	
N_{stocks}	626		626		626		626		626		626		626	

This table contains estimates from regressions of the stock-day average post-trade absolute price change 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, bid–ask spread and log depth available at the NBBO at the time of the 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.

Table 4
Adverse selection panel regressions.

	(1)		(2)		(3)		(4)		(5)		(6)		(7)	
	500ms		1s		10s		30s		60s		300s		1800s	
Ln(Dollar trade size)	−0.0010	(−3.38)	−0.0012	(−4.07)	−0.0017	(−5.67)	−0.0023	(−6.89)	−0.0025	(−5.89)	−0.0042	(−4.78)	−0.0123	(−5.82)
Ln($N_{dark pool}$)	−0.0004	(−1.08)	−0.0010	(−2.39)	−0.0024	(−4.97)	−0.0041	(−7.45)	−0.0058	(−8.52)	−0.0098	(−8.47)	−0.0065	(−2.72)
Ln(N_{Lit})	−0.0043	(−6.49)	−0.0045	(−6.60)	−0.0052	(−6.67)	−0.0058	(−6.60)	−0.0068	(−6.10)	−0.0153	(−6.38)	−0.0304	(−5.87)
Ln(Price)	−0.0198	(−8.68)	−0.0210	(−8.83)	−0.0233	(−8.25)	−0.0257	(−7.55)	−0.0293	(−5.84)	−0.0622	(−5.09)	−0.1546	(−6.00)
Ln(dollar volume)	0.0018	(3.27)	0.0024	(4.49)	0.0034	(6.02)	0.0041	(6.51)	0.0047	(5.98)	0.0066	(4.07)	0.0125	(3.72)
Pre-cross bid–ask Spread	0.9209	(178)	0.9169	(164)	0.9073	(141)	0.9065	(126)	0.9251	(108)	0.9475	(55.5)	0.8661	(31.7)
Ln(Depth)	−0.0046	(−7.55)	−0.0051	(−8.11)	−0.0059	(−8.56)	−0.0067	(−7.12)	−0.0058	(−3.84)	−0.0018	(−0.51)	0.0027	(0.42)
Broker dark pool	−0.0024	(−2.42)	−0.0033	(−2.82)	−0.0049	(−3.58)	−0.0066	(−4.17)	−0.0053	(−2.96)	−0.0052	(−1.41)	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		0.81		0.75		0.55		0.21	
N_{obs}	242,840		242,840		242,840		242,840		242,840		242,840		242,840	
N_{stocks}	626		626		626		626		626		626		626	

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, bid–ask spread and log depth available at the NBBO at the time of the 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.

expected, the distribution of trade size is heavily right-skewed. The median trade is for around one-third of this amount, AUD 1000, 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 min 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% 30 min after the trade.

Table 3 contains results from our stock-day panel regressions for absolute price change 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.

Broker dark pool trades result in lower absolute price changes 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. Fig. 1 presents these estimates and upper and lower confidence intervals by time horizon graphically. The size of the difference in absolute price change between broker dark pool and exchange dark pool trades increases 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 min horizon is important because it suggests that there are no long-term 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.[19]

Regarding other variables in our preferred specification, the average dark pool trade price change 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 change 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 change, 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 change, both absolutely and relative to sample-wide standard deviations. Further, the effect is insignificantly

[19] Like standard definitions of price impact, price changes over a particular interval can contribute to the estimation of price impact for multiple trades in our sample. This introduces serial correlation into the absolute price change residual that would, if our regressions were estimated at the trade-level, bias unadjusted standard errors. Estimating our regressions at the stock-day level rather than the trade level automatically corrects for this serial correlation in the standard errors.

different from zero by 300s, while the effect on absolute price change remains significant at that horizon. Nevertheless, evidence from bid–ask spreads is consistent with that of absolute price change insofar as less information leakages maps closely to less perceived adverse selection risk from market makers. At the 30 min 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 total information leakage, wider spreads after broker dark pool trades are consistent with market makers learning more slowly from these trades compared with those on exchange dark pools.

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 min, there is some evidence that the effect reverses.

4.2. Analysis of broker dark pool closures

The stock-day panel analysis in Section 4.1 suffers from one important limitation: we cannot rule out endogeneity between the error term ε_{jt} and our regressors of interest. Although 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 more subtle forms of endogeneity. For example, brokers who operate dark pools may also have a greater fraction of large institutional customers with typically long holding periods, such as mutual or pension funds, compared with other brokers who service a greater fraction of traders with shorter horizons (such as hedge funds). 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). These closures were driven by compliance risk rather than execution outcomes.

4.2.1. Economics of broker dark pool closures

For pool closures to be a source of exogenous variation in venue choice the decision to close must be unrelated to execution outcomes in the pool. The pool closures in our sample were largely motivated by concerns about regulatory and compliance risks related to operational issues, not execution outcomes. This is supported by the fact that ASIC actively monitors and takes regulatory actions against broker dark pools that breach the Market Integrity Rules and that these rules do not prohibit traders using order anticipation strategies that potentially lead to information leakage.

Between 2017 and 2019, ASIC issued three infringement notices related to the UBS and BAML broker pools, all related to operations of the pools not execution quality.[20] In each instance, the problems were initially self-reported by the broker operating the pool. Issues raised by ASIC include trade reporting errors relating to trade capacity (i.e. principal vs. agent) and trading venue (i.e. exchange or broker pool), not allowing customers to opt-out of types of flow, and allowing unfilled orders to lose time priority. ASIC is explicit in the infringement notices that these issues did not lead to bad customer outcomes nor did the pool operators gain a material benefit from their errors or deliberately contravene market rules. The problems identified in these notices relates to trading between December 2013 and April 2016, well before our sample period.

The heightened regulatory scrutiny and increased compliance costs resulted in some brokers determining that the risks of operating a dark pool exceeded the benefits for reasons unrelated to the use of order anticipation strategies at the extensive or intensive margin.[21] While Citi was not subject to an infringement notice, ASIC's focus on the operational practices of dark pools was well-known to the market. The importance of compliance risk in the decisions to shut dark pools is also supported by the fact that these brokers continued operating dark pools in other jurisdictions with less regulatory scrutiny.

4.2.2. Matched differences in execution outcomes

After the closure dates, if a customer demands dark liquidity, 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.

Fig. 3 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 route a similar number of orders to these brokers despite the impending closure of the dark pool.

Discussions with broker dark pool operators revealed that the presence of a broker dark pool is not considered a relevant factor in the institutions deciding which broker receives an order. They also confirmed that institutional clients invariably use the execution strategy offered by the broker's execution algorithms and typically do not specify a preference for executions to take place in particular venues. Importantly for our purpose, brokers indicated that following a broker pool closure, order flow that previously would have been routed to the broker pool is routed to exchange dark pools.

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 5 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 5 presents means split by trades on broker dark pool vs. Centre Point trades from closed-pool brokers.

There are approximately 6.6 million trades in total in our matching sample, comprising of 5.6 million broker dark pool trades from all brokers whose dark pools continue to operate and 1.1 million Centre

[20] Infringement notices are reported in the ASIC Gazette and to the media. See Market Disciplinary Panel 06/17, 07/17 and 11/17.

[21] Minimizing compliance risk requires fixed-cost investments in technology and processes to minimize operational failures. Dark pools with larger market share are more likely to make these investments due to economies of scale. Of the three pools that close during our sample period, two are restricted. This category naturally attracts less order flow and has lower market shares than unrestricted pools.

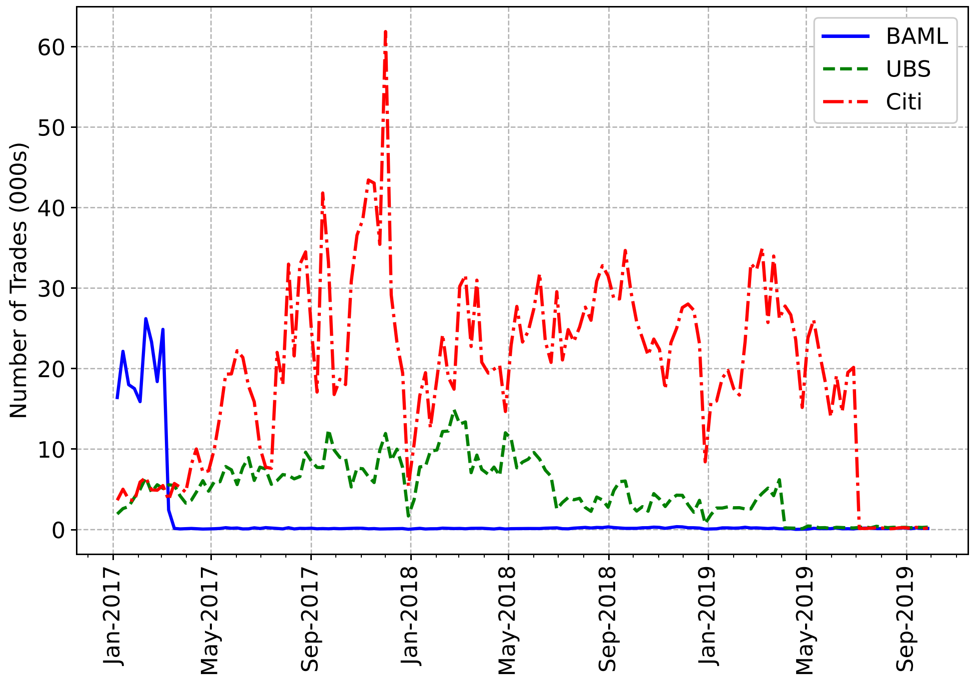

Fig. 3. Number of broker dark pool trades for closing pools.
Time-series of the number of broker dark pool trades for each of the three brokers whose pools are closed during our sample period: Bank of America Merrill Lynch (BAML), UBS and Citigroup (Citi). The respective closure dates are March 6, 2017 (BAML), April 1, 2019 (UBS) and July 1, 2019 (Citi).

Table 5
Trade-level summary statistics around pool closure.

	Mean	SD	Min	25%	50%	75%	Max	Mean BDP	Mean CP
Trade size (AUD '000s)	1.79	29.8	0.00	0.08	0.31	1.14	22,025	1.62	2.70
Price (AUD)	20.2	33.0	0.06	4.03	8.73	21.0	232	20.6	18.8
Total dollar value (AUD 'm)	40.6	54.0	0.00	9.07	21.3	46.2	441	40.6	41.0
Pre-cross bid–ask spread (%)	0.18	0.18	0.00	0.06	0.12	0.25	1.87	0.17	0.19
Abs. 500ms price change (%)	0.00	0.02	0.00	0.00	0.00	0.00	0.25	0.00	0.01
Abs. 10s price change (%)	0.01	0.05	0.00	0.00	0.00	0.00	0.42	0.01	0.02
Abs. 60s price change (%)	0.04	0.08	0.00	0.00	0.00	0.04	0.66	0.04	0.04
Abs. 1800s price change (%)	0.24	0.30	0.00	0.03	0.15	0.32	2.79	0.24	0.23
500ms Bid–ask spread (%)	0.18	0.18	0.00	0.06	0.12	0.25	1.87	0.17	0.20
10s Bid–ask spread (%)	0.18	0.18	0.00	0.06	0.12	0.26	1.87	0.17	0.20
60s Bid–ask spread (%)	0.18	0.19	0.00	0.06	0.12	0.26	1.90	0.18	0.21
1800s Bid–ask spread (%)	0.30	0.48	0.00	0.07	0.16	0.32	5.33	0.30	0.30
60s Dollar volume (AUD '000s)	96.5	198	0.00	5.38	24.4	91.9	1,867	94.7	105
60s Abs. order imbalance (AUD '000s)	39.4	84.7	0.00	2.28	10.2	36.6	906	38.7	43.0
60s Quote-to-trade Ratio	3.92	3.15	0.00	2.25	3.04	4.36	28.0	3.95	3.77

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 the 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 is the log bid–ask spread at the NBBO at the time of the trade. Absolute τ price change 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. 60s dollar volume and abs. order imbalance are computed across all exchanges over the 60s after the dark pool trade, in thousands of dollars. 60s quote-to-trade ratio is the ratio of the number of NBBO quote changes and the number of trades over the same interval. The final two columns report the average of these variables for trades on broker dark pools ("BDP") and on Centre Point ("CP").

Point trades from the three brokers with recently closed dark pools. The average trade size in the matching sample is approximately $1800 though trades executed on broker dark pools tend to be around $1000 smaller than those executed on Centre Point. There are only small differences in average trade price, total daily dollar value traded, absolute price change, reversals or price adjustment by category, though the bid–ask spread before 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 change are both significantly smaller while average total dollar value traded, as well as post-trade dollar value and absolute order

Table 6
Matching regression.

	(1) BAML	(2) UBS	(3) Citi
Abs. Price change (60s)	−0.0083	−0.0037	−0.0025
	(−8.54)	(−7.84)	(−3.92)
Bid–ask spread (60s)	−0.0264	−0.0016	−0.0023
	(−13.6)	(−3.94)	(−4.53)

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 have recently closed. For each stock and each of the one-month periods after the three pool closures described in Table 5, 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, 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 outcomes are measured using the same variables as defined in Table 2 and Table 5. Standard errors are clustered at the stock level and *t*-statistics are in parentheses.

imbalance are larger. This largely reflects the fact that the summary statistics in our matching sample (Table 5) are constructed at the trade level whereas the stock-day summary statistics are computed at the stock-day level.[22]

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 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 before 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 expressed as a decimal and its square.[23] We then find the nearest neighbor for each broker dark pool trade from the set of Centre Point trades from closed-pool brokers in the same stock, matching with replacement. We keep only trades that can be matched using a caliper of 0.25 standard deviations of the estimated propensity scores.

Our control group is a matched subset of Centre Point 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. 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.[24] More details regarding the matching process are presented in Section B of Appendix A.

Estimates of the average treatment effect on the treated (ATT) for each event are presented in Table 6. We use the 60s horizon, corresponding to that with the largest absolute price change effect in our panel regressions.

Our matching analysis confirms that trades on broker dark pools have significantly less information leakage and adverse selection risk in the period after trade execution. The ATT for absolute price change is between −0.25 bps and −0.83 bps across the three events and the size 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 events, compared with −0.53 bps at the same horizon in our panel analysis.[25]

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.

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 can 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. Evidence of these differences is presented in ASIC (2015) which documents substantial differences in the level of high frequency trading present in exchange- vs. broker-operated dark pools. High frequency trading accounts for 14.4% of turnover on the ASX dark pool and 27.6% of hidden liquidity on Cboe in the first quarter of 2015. On average, high frequency trading accounts for only 1.7% of turnover on broker-operated dark pools over the same period.

A potential explanation for why execution outcomes are worse in exchange dark pools is that these markets have a higher proportion of HFTs/ELPs who place orders intending to detect 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

22 Trade-level summary statistics are presented separately for each event in Tables A.1-A.3 in Appendix A.

23 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.

24 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.

25 Table A.4 in Appendix A contains treatment effects for the 30 min horizon obtained under our matching approach. We find qualitatively similar results to our panel results — we can reject the null of equal absolute price change over these post-trade 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. The absolute price change effects are most relevant as they help demonstrate there are no clear long-term differences in the total amount of information contained in trades across the two venues.

and Zhu, 2019 and sniping in Malinova and Park, 2020) result in the HFT/ELP firms subsequently consuming liquidity on the same side of the book as the institutional order. Van Kervel and Menkveld (2019) show that when HFTs trade in the same direction as institutional order flow, it increases the short-term order imbalance, resulting in higher subsequent price impact. Hendershott and Madhavan (2015) show that requests for quotes on an electronic platform for trading corporate bonds can lead to costly information leakage by revealing trading intentions to potential counterparties.

5.1. Evidence from differences in trader access across broker dark pools

Broker dark pools differ by how accessible they are to HFTs. ASIC (2015) also reports that in the first quarter of 2015 the level of HFT activity ranged from 0.32% to 34% across individual pools. Our first test for this explanation is to examine the difference in execution outcomes across broker dark pools split by whether the pool permits HFTs/ ELPs.

Table 1 groups broker dark pools into two groups based on the level of restrictions. The first group, labeled *restricted*, limits access the most. They prohibit order flow from the firms' principal trading desk, HFTs and ELPs. 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.[26]

The second group of pools, labeled *opt-in to restrictions*, allow customers to opt-out of interacting with principal, HFT 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 also given the choice of whether to participate in this flow. When orders are received by aggregators the counter-party type is not known — so a customer wanting to interact only with natural flow would choose not to engage with aggregator flow. Therefore, while these opt-in to restriction pools comprise more diverse order flow, customers can choose to interact only with natural liquidity.[27]

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 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 pools, UBS PIN and Citi Match, close early in our sample period; and (iii) another restricted broker dark pool, Liquidnet, has a minimum order size of \$100,000. 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.

[26] 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 high frequency trading firms. Therefore we classify it as a *restricted* pool.

[27] ASIC's active monitoring of the broker dark pools and the fact that all but one of the infringement notices related to activity before our sample period give us confidence that pools disclosing access restrictions actually do restrict access during our sample period. As a robustness check, we exclude the pool whose infringement notice included activity during our sample period and the results remain consistent (Table A.5 in Appendix A).

Table 7
Stock-day regression for broker dark pool trades split by HFT/ELP access.

	(1)		(2)	
	Abs. ΔP		Spread	
Ln(Dollar trade size)	0.0029	(10.8)	−0.0003	(−1.37)
Ln($N_{Darkpool}$)	−0.0081	(−23.2)	−0.0033	(−10.1)
Ln(N_{Lit})	0.0171	(21.6)	−0.0060	(−6.78)
Ln(Price)	−0.0115	(−4.80)	−0.0380	(−10.5)
Ln(Dollar volume)	0.0116	(18.8)	0.0033	(5.87)
Pre-cross bid–ask spread	0.0465	(7.86)	0.8744	(115)
Ln(Depth)	−0.0192	(−30.6)	−0.0074	(−9.48)
HFT-restricted pool	−0.0080	(−3.94)	−0.0037	(−1.69)
Fixed effects	$N\&T$		$N\&T$	
R^2	0.10		0.77	
N_{obs}	192,068		192,068	
N_{stocks}	563		563	

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/ELP 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 change and bid–ask spreads measured at the 60s post-trade horizon), $\bar{X}_{jt}$ is the stock-day average of a vector of controls including: log of dollar trade size, log of trade price, 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 permit HFT/ELP 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.

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 pool that does not permit any HFT/ELP activity ("HFT-restricted Pool") and 0 otherwise:

$$\bar{y}_{jt} = \alpha_j + \gamma_t + \beta \bar{Restricted}_{jt} + \rho' \bar{X}_{jt} + \bar{\varepsilon}_{jt} \tag{4}$$

where all variables are again defined as per Eq. (3) other than $\bar{Restricted}_{jt}$ which is the stock-day average of the dummy variable for the trade taking place in a pool that does not permit HFTs/ELPs. The key parameter in Eq. (4) is β which captures average execution outcomes for trades on broker dark pools that prohibit HFTs/ELPs vs. broker dark pools that allow customers to opt-in to HFT/ELP flow.

Table 7 presents the estimates from these regressions where absolute price change and bid–ask spreads are measured at the 60s horizon. Our results from Table 7 are consistent with HFTs/ELPs influencing dark pool execution outcomes. Trades in pools that do not permit HFT/ELP activity result in significantly lower absolute price changes than those that allow HFT/ELP activity. The size of the effect is −0.8 bps, compared to our difference in absolute price change 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.

We also consider differences in access restrictions across the three pools that closed during our sample period. Two of the pools (UBS PIN and Citi Match) operated as restricted broker dark pools while the third (BAML InstinctX) allowed HFT/ELP activity. Our matching analysis in

Section 4.2 compares execution outcomes of broker dark pool trades with execution outcomes of matched Centre Point trades from each broker whose pools closed, in a short period after each pool closure. For each event, we estimate negative and significant coefficients, indicating that trades in the remaining broker dark pools have less information leakage than the matched Centre Point trades of brokers whose pools closed.

The presence of HFT/ELP in the BAML broker dark pool suggests that the post-BAML closure Centre Point trades will contain more HFT/ELP activity than the other closing pools. If HFT/ELP are important determinants of information leakage, then we would expect to observe larger (more negative) coefficients on absolute price change and adverse selection for the BAML event compared with the other events. Table 6 confirms that this is indeed the case. For absolute price change, the BAML coefficient is −0.8 bps compared with approximately −0.3 and −0.4 bps for UBS and Citi. For bid–ask spreads, the BAML coefficient is −2.6 bps compared with approximately −0.2 bps for the other events.

The difference in matching results across brokers closing restricted vs. unrestricted pools complements our panel results which compare execution outcomes across trades in restricted vs. unrestricted broker dark pools. In both cases, trades that ex-ante involve more HFT/ELP also have more information leakage and adverse selection risk.

5.2. Evidence from trading activity and quote-to-trade ratios

Our second approach to testing the relevance of the HFT channel examines quoting and trading activity on the lit market following trades in exchange and broker dark pools. When an HFT trades in a dark pool, they learn that there is trading interest on the other side of the market. For example, learning that there is buying interest in a stock may cause an HFT to begin buying the same stock, in anticipation of future buying interest from the counter-party to their original dark pool trade (see e.g. van Kervel and Menkveld, 2019, Yang and Zhu, 2019, Sağlam, 2020, Malinova and Park, 2020 and Hirschey, 2021). To test the presence of such a mechanism, we look for evidence of increased liquidity demand and higher short-term order imbalance after exchange dark pool trades compared with broker dark pool trades, signaling that, on average, this type of information leakage is more prevalent in dark pool categories where HFTs/ELPs are more prevalent.

We construct three new dependent variables: total dollar volume traded in thousands of dollars, absolute order imbalance in thousands of dollars and the quote-to-trade ratio, and estimate the effect of execution venue for these variables in our panel regression framework. All variables are computed using all lit trades and quotes in either the ASX or Cboe limit order books, over a 60s horizon that follow a dark pool trade. Absolute order imbalance is the absolute value of the difference in buy dollar volume and sell dollar volume over that interval. The quote-to-trade ratio is the number of changes in the NBBO (price or size) divided by the number of trades over the interval. Table 8 presents these results.

These regressions show that exchange dark pool trades are followed by significant increases in lit market trading activity and order imbalance (negative coefficients on the BDP dummy variable). The sizes of the effects are between 10% and 15% of the mean value of these variables (see Table 2). Van Kervel and Menkveld (2019) find that HFTs increase their net order imbalance when trading "with the wind" of an institutional order and that institutional trades have higher overall price impact when trading in the same direction as an HFT. Our results provide support that this mechanism can explain our observed differences in execution outcomes following trades in exchange vs. broker dark pools.

If HFTs/ELPs are rotating their strategies from liquidity provision to liquidity demanding (e.g. when they learn of the presence of an institutional order and begin trading "with the wind"), then the quote-to-trade ratio in lit markets should fall when they detect an institutional order. HFT/ELP liquidity provision naturally involves frequent updating of quotes as the traders process and react to new information with very low latency. While this ratio, or related measures such as messages or cancels to trade ratios, are often used as a proxy for overall algorithmic trading (Hendershott et al., 2011; Conrad et al., 2015; Boehmer et al., 2021; Rosu et al., 2021), they arguably capture liquidity provision more than liquidity demand, since the measures naturally increase with quoting activity.[28] We find a statistically and economically significant reduction in the lit market quote-to-trade ratio following exchange dark pool trades vs. broker dark pool trades, consistent with reduced liquidity provision following dark trades that are more likely to involve an HFT/ELP counter-party.

Overall, these results document higher lit trading activity, more unidirectional lit trading and less active algorithmic quoting activity (consistent with less liquidity supplying) following exchange dark pool trades compared with comparable broker dark pool trades, consistent with HFTs learning about future order flow from trading in dark pools.[29]

6. Execution risk in exchange and broker dark pools

Our analysis has shown that average execution outcomes are worse in dark pools with fewer access restrictions. A natural question that follows from this result is what supports cross-sectional variation in access restrictions across dark pools in equilibrium? If institutional investors understand that execution outcomes are better in restricted dark pools, why do these investors route orders to unrestricted pools?[30]

Key to answering this question is the fact that execution outcomes such as information leakage and adverse selection risk only capture part of the cost of trading. Another component of the cost, which varies across venues, is the probability of order execution, or execution risk. High frequency traders who place orders in dark pools are necessarily supplying liquidity, even if a fraction of these traders are also trying to detect institutional order flow. A consequence of restricting access to these traders is that the amount of liquidity in the pool and the probability that any single order executes reduces.

Dark pools with different access restrictions can be thought of as offering different trade-offs between execution risk and execution quality, analogous to the trade-off investors face when choosing between dark pools or crossing networks vs. lit venues studied in Hendershott and Mendelson (2000), Zhu (2014), Menkveld et al. (2017), Ye and Zhu (2020) and Ye (2024), among others. In our setting, we make comparisons across dark pools offering similar price improvement mechanisms but different access restrictions. Traders rationally compare execution probability against price impact (information leakage), which affects costs of *future* child order executions. In equilibrium, traders choose the venue (or execution algorithm) that balances their need for immediacy against the cost of price impact on future orders. Under different conditions, traders may make different trade-offs.

Combining our results on market share and execution outcomes, we find that the least restricted pool (Centre Point) has the highest execution probability (proxied by market share — see Fig. 2) and the largest absolute price change (information leakage — see Table 3). The

[28] Yao and Ye (2018) provide evidence that the association between message-to-trade ratios and HFT liquidity provision is primarily driven by within-security time-series variation, while the cross-sectional correlation between average message-to-trade ratios and stock-level HFT liquidity provision is weak. Our focus is exclusively on within-stock variation in these ratios, where the correlation is well-established and frequently used in the literature.

[29] To check whether these regressions suffer from the endogeneity problems discussed in Section 4.2, we also conduct our matching analysis for these dependent variables. The results are broadly consistent, although we obtain an insignificant coefficient for absolute order imbalance and quote-to-trade ratios for one of the three events respectively (see Table A.6 in Appendix A).

[30] We thank an anonymous referee for posing this question.

Table 8
Stock-day regression for trading and quoting activity.

	(1)		(2)		(3)	
	Dollar volume		Abs. order imbalance		Quote to trade	
Ln(Dollar trade size)	3.2084	(15.9)	1.8281	(18.9)	0.0253	(5.11)
Ln($N_{darkpool}$)	−1.8471	(−8.79)	−1.3563	(−16.2)	−0.4029	(−39.8)
Ln(N_{Lit})	7.4024	(10.3)	2.1553	(9.91)	0.3492	(21.3)
Ln(Price)	1.5064	(0.85)	0.0490	(0.08)	0.3613	(6.91)
Ln(Dollar volume)	10.597	(15.3)	5.2233	(18.6)	−0.0649	(−7.06)
Pre–cross bid–ask spread	10.687	(6.47)	3.0117	(5.19)	0.1305	(2.17)
Ln(Depth)	1.7662	(4.54)	1.3657	(9.00)	0.0318	(2.19)
Broker dark pool	−4.9707	(−5.76)	−1.9637	(−7.64)	0.5856	(26.5)
Fixed effects	$N\&T$		$N\&T$		$N\&T$	
R^2	0.13		0.17		0.09	
N_{obs}	237,180		237,180		237,180	
N_{stocks}	624		624		624	

This table contains estimates from regressions of the stock-day average of trading and quoting activity 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 total dollar volume traded in thousands of dollars (Column 1), the absolute order imbalance (Column 2) in thousands of dollars, or the quote-to-trade ratio (Column 3) 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, bid–ask spread and log depth available at the NBBO at the time of the 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. All dependent variables are measured in the 60s interval following a dark pool trade. 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.

most restricted pools (broker dark pools with access restrictions) attract the least volume and have the lowest information leakage.

A related question is whether execution risk affects order routing between dark pools in different market conditions? Without information at the parent order level, it is not possible to directly test for differences in execution risk across venue types and market conditions. Instead, we test how execution risk changes with market conditions by examining market shares across dark pool types during periods of heightened volatility. Holding all else equal, the cost of non-execution increases with the volatility, and empirically, dark pool market share falls when volatility rises (Kwan et al., 2015; Menkveld et al., 2017; Buti et al., 2022). Understanding the association between volatility and broker dark pool market share sheds light on whether these venues still attract order flow when costs of non-execution are high, relative to exchange dark pools.

We estimate our main stock-day panel regressions with two differences. First, the left-hand side variable is either (i) the fraction of *total* trading in all dark pools or (ii) the fraction of *dark* trading in broker dark pools (i.e. broker dark pool volume divided by the sum of all broker dark pool and exchange dark pool volume). Second, the right-hand side variable of interest is an indicator for whether stock-day realized volatility (RV) of five-minute returns is above the upper quartile computed across all stock-days in that month. All other details are as per our main stock-day regressions in Eq. (3).

We expect the coefficient on the RV indicator to be negative and significant for total dark pool market share. The sign and significance of the RV coefficient for broker dark pool market share captures whether traders avoid these venues during periods with higher costs of non-execution. Table 9 presents these results.

As expected, we observe a negative and significant effect of RV on total dark market-share. A stock-day with volatility in that month's upper quartile has around 3 to 4 percentage points lower dark volume share, conditional on fixed effects and controls (Columns (1) and (2)). For the fraction of dark volume in broker dark pools, (Columns (3) and (4)), we estimate a *positive* and significant (insignificant) association between volatility and broker dark pool market share when using fixed effects only (fixed effects and controls). The fact that broker dark pools maintain their share of overall dark volume during periods of high volatility suggests that time-variation in the costs of non-execution do not drive order routing decisions between broker dark pools and exchange dark pools.

7. Robustness

Our main empirical results are supported by a large number of robustness tests. We briefly describe these tests here and present details in Appendix A. Our first set of robustness tests estimate our main panel regressions using various alternative sample constructions where we include sweep trades, exclude Cboe hidden trades, exclude ASX-reported NBBO trades that are possible manual matches, and sequentially exclude each broker from the sample (ensuring that our results are not driven by a single broker). All results are robust and conclusions are unchanged. Details are in Section C.1 of Appendix A.

Next, to complement our tests of the HFT channel, in Section C.2 of Appendix A we estimate our main panel regressions where the effect of venue is estimated by trade size. HFTs/ELPS tend to trade in smaller sizes than other trader categories (see e.g. Davis et al., 2014, Benos et al., 2017 and Brogaard et al., 2019). We show that the reduction in absolute price change for broker dark pool trades vs. exchange dark pool trades is nearly twice as large for smaller-sized trades as for larger trades. There are smaller and statistically insignificant differences in bid–ask spreads across trade sizes.

We also test whether binding tick constraints affect our results. When the minimum tick size is a binding constraint for the bid–ask spread, the depth at the NBBO tends to be large, and the cost of demanding liquidity is high. Therefore, trading within the spread on dark venues becomes more attractive (Kwan et al., 2015; Comerton-Forde et al., 2019; O'Hara et al., 2019). Our analysis could potentially be biased if one dark pool category is favored over the other in a tick-constrained environment. In Section C.3 of Appendix A we construct regressions to test this and find no evidence that relative market shares of either dark pool category are affected by tick constraints (although as expected, dark pools in total receive more orders when tick constraints bind).

Last, in Section C.4 of Appendix A we construct a test for whether brokers lose order flow in non-random ways when their dark pool closes, if for example, a broker pool is an important factor for customers choosing brokers or if brokers strategically choose where to route certain customers' orders. We do this by comparing execution outcomes on Centre Point for trades from brokers whose pool closes compared with other brokers via a difference-in-differences regression. For two of our three closure events, UBS and Citi, the treatment effect is insignificant

Table 9
Dark pool trading volume and execution risk.

	Dark fraction of total				BDP fraction of all dark			
	(1)		(2)		(3)		(4)	
Ln(Dollar trade size)	–	–	0.0622	(41.4)	–	–	−0.0208	(−15.0)
Ln($N_{darkpool}$)	–	–	0.0747	(57.6)	–	–	0.0092	(6.08)
Ln(N_{Lit})	–	–	−0.0873	(−50.4)	–	–	0.0104	(4.68)
Ln(Price)	–	–	−0.0025	(−1.04)	–	–	0.0513	(7.82)
Ln(Dollar volume)	–	–	0.0069	(5.87)	–	–	0.0033	(2.54)
Pre-cross bid–ask spread	–	–	0.0016	(0.53)	–	–	0.0221	(2.78)
Ln(Depth)	–	–	−0.0244	(−20.4)	–	–	−0.0041	(−2.64)
Upper quartile RV	−0.0391	(−23.7)	−0.0289	(−24.0)	0.0043	(2.00)	0.0006	(0.32)
Fixed effects	$N\&T$		$N\&T$		$N\&T$		$N\&T$	
R^2	0.01		0.48		0.00		0.02	
N_{obs}	242,854		242,854		242,854		242,854	
N_{stocks}	626		626		626		626	

This table contains estimates from regressions of the fraction of total trading that takes place in all dark pools and the fraction of dark pool trading that takes place on broker dark pools onto stock-day level controls, fixed effects and a dummy variable indicating whether stock-day realized volatility is above the upper quartile computed across all stock-days in the month. The general form of the regression model is $y_{jt} = \alpha_j + \gamma_t + \beta RV_{jt} + \rho' X_{jt} + \varepsilon_{jt}$ where α_j is a stock fixed effect, γ_t is a date fixed effect, y_{jt} is the ratio of dark pool volume traded to total volume traded during the continuous trading period, excluding auctions and block trading, (Columns 1 and 2) or the ratio of broker dark pool trading to total dark pool trading (Columns 3 and 4) for stock j and day t, X_{jt} is the stock-day average of a vector of controls including: log of dollar trade size, log of trade price, bid–ask spread and log depth at NBBO at the time of the 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; RV_{jt} is a dummy variable taking the value 1 if the volatility for stock j on day t is equal to or above that month's upper quartile (computed across all stocks) and ε_{jt} is an error term. We estimate the model using 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.

at the 5% level for both absolute price change and bid–ask spreads. For our other event, BAML, we estimate a statistically significant decrease in absolute price change and an increase in bid–ask spreads. The closure of this pool may have resulted in confounding changes in that broker's order flow, however, our conclusions would be unchanged even if we were to exclude BAML event from our matching analysis.

8. Conclusion

Dark pools are an important part of the trading ecosystem in most developed markets. The academic literature largely treats dark pools as homogeneous, but in practice, they differ on several important dimensions. We examine differences in the extent to which dark pools restrict access to some 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 principal flow. Our results show that broker dark pool trades have less information leakage and less adverse selection risk than exchange dark pools. The presence of HFT/ELP order flow is a key driver of the main results. These findings are consistent with account-level analysis undertaken by ASIC. They show 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 (ASIC, 2015).

Our results are relevant in other settings given that other jurisdictions also allow differential access to different trader types. For example, in the US, Alternative Trading Systems (ATS) can segment order flow if their market share is below 5% and are required to provide disclosures regarding execution practices, permitted trader types and volumes. Our evidence suggests that these disclosures provide traders with valuable information to help them assess the trade-offs between execution risk and execution outcomes. Transaction-level data, such as the data available in Australia may further enhance the execution tool kit for traders. In contrast, European regulators banned 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 harmed institutional investors.

While our results show that regulations that allow access restrictions in dark pools can enhance execution outcomes for investors we do not evaluate their impact on overall market quality. Segmentation of order flow in this manner may potentially reduce overall market quality. This is an interesting question to explore in future research.

Traders can indirectly influence the type of flow they interact with in unrestricted pools through the use of MAQ, which allow traders to specify their minimum trade size, reducing the risk of trading with HFTs/ELPs. Detailed data capturing the use of MAQ are currently non-public, however, their effect on execution outcomes 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 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.

CRediT authorship contribution statement

James Brugler: Writing – review & editing, Writing – original draft, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Carole Comerton-Forde:** Writing – review & editing, Writing – original draft, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The first author has nothing to disclose.

Cboe Australia provided the second author with complimentary access to the broker-venue data used in the paper. The second author received compensation exceeding $10,000 during the past three years for consulting and advisory services provided to the Australian Securities and Investments Commission, Cboe Australia and the Plato Partnership. The second author also received financial support from the Norwegian Finance Initiative for research activities, including this paper. These institutions may or may not have an interest but did not have the right to review this paper prior to its circulation. The author has no other potential conflicts to disclose.

References

Anand, A., Samadi, M., Sokobin, J., Venkataraman, K., 2021. Institutional order handling and broker-affiliated trading venues. Rev. Financ. Stud. 34 (7), 3364–3402.

ASIC, 2015. Report 452: Review of high-frequency trading and dark liquidity. Accessed: 2021-06-01.

ASX, 2021. Historical market statistics.

ASX, 2022. ASX trade: Markets participant and trading schedule of fees. https://asxonline.com/content/dam/asxonline/public/documents/schedule-of-fees/asx-trade-markets-participant-and-trading-schedule-of-fees.pdf. (Accessed 18 December 2023).

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

Battalio, R.H., Hatch, B.C., Saglam, M., 2024. The cost of exposing large institutional orders to electronic liquidity providers. Manag. Sci. 70 (6), 3381–4615.

Benos, E., Brugler, J., Hjalmarsson, E., Zikes, F., 2017. Interactions among high-frequency traders. J. Financ. Quant. Anal. 52 (4), 1375–1402.

Boehmer, E., Fong, K., Wu, J.J., 2021. Algorithmic trading and market quality: International evidence. J. Financ. Quant. Anal. 56 (8), 2659–2688.

Brogaard, J., Hendershott, T., Riordan, R., 2019. Price discovery without trading: Evidence from limit orders. J. Financ. 74 (4), 1621–1658.

Buti, S., Rindi, B., Werner, I.M., 2022. Diving into dark pools. Financ. Manag. 51 (4), 961–994.

Cboe, 2023a. Cboe Australia fee schedule. https://www.cboe.com/au/equities/participation/fee_schedule/. (Accessed 18 December 2023).

Cboe, 2023b. Maximise price improvement opportunities with cboe mid point orders. https://cdn.cboe.com/resources/au/participant_resources/CXA_MidPeg_Orders.pdf. (Accessed 18 December 2023).

Comerton-Forde, C., Grégoire, V., Zhong, Z., 2019. Inverted fee structures, tick size, and market quality. J. Financ. Econ. 134 (1), 141–164.

Comerton-Forde, C., Malinova, K., Park, A., 2018. Regulating dark trading: Order flow segmentation and market quality. J. Financ. Econ. 130 (2), 347–366.

Conrad, J., Wahal, S., 2020. The term structure of liquidity provision. J. Financ. Econ. 136 (1), 239–259.

Conrad, J., Wahal, S., Xiang, J., 2015. High-frequency quoting, trading, and the efficiency of prices. J. Financ. Econ. 116 (2), 271–291.

Davis, R.L., Van Ness, B.F., Van Ness, R.A., 2014. Clustering of trade prices by high-frequency and non–high-frequency trading firms. Financ. Rev. 49 (2), 421–433.

Hatheway, F., Kwan, A., Zheng, H., 2017. An empirical analysis of market segmentation on U.S. equity markets. J. Financ. Quant. Anal. 52 (6), 2399–2427.

Hendershott, T., Jones, C.M., Menkveld, A.J., 2011. Does algorithmic trading improve liquidity? J. Financ. 66 (1), 1–33.

Hendershott, T., Madhavan, A., 2015. Click or call? Auction versus search in the over-the-counter market. J. Financ. 70 (1), 419–447.

Hendershott, T., Mendelson, H., 2000. Crossing networks and dealer markets: Competition and performance. J. Financ. 55 (5), 2071–2115.

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

van Kervel, V., Menkveld, A.J., 2019. High-frequency trading around large institutional orders. J. Financ. 74 (3), 1091–1137.

Korajczyk, R.A., Murphy, D., 2018. High-frequency market making to large institutional trades. Rev. Financ. Stud. 32 (3), 1034–1067.

Kwan, A., Masulis, R., McInish, T.H., 2015. Trading rules, competition for order flow and market fragmentation. J. Financ. Econ. 115 (2), 330–348.

Malinova, K., Park, A., 2020. "Sniping" in fragmented markets. Available At SSRN 3534367.

Menkveld, A.J., Yueshen, B.Z., Zhu, H., 2017. Shades of darkness: A pecking order of trading venues. J. Financ. Econ. 124 (3), 503–534.

Norges Bank Investment Management, 2015. Sourcing liquidity in fragmented markets 01-2015. Accessed: 2021-07-22.

O'Hara, M., 2015. High frequency market microstructure. J. Financ. Econ. 116 (2), 257–270.

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

Rosu, I., Sojli, E., Tham, W.W., 2021. Quoting activity and the cost of capital. J. Financ. Quant. Anal. 56 (8), 2764–2799.

Sağlam, M., 2020. Order anticipation around predictable trades. Financ. Manag. 49 (1), 33–67.

World Federation of Exchanges, 2019. Market statistics — December 2019. https://focus.world-exchanges.org/issue/december-2019/market-statistics. (Accessed 15 June 2021).

Yang, L., Zhu, H., 2019. Back-running: Seeking and hiding fundamental information in order flows. Rev. Financ. Stud. 33 (4), 1484–1533.

Yao, C., Ye, M., 2018. Why trading speed matters: A tale of queue rationing under price controls. Rev. Financ. Stud. 31 (6), 2157–2183.

Ye, L., 2024. Understanding the impacts of dark pools on price discovery. J. Financ. Mark. 68, 100882.

Ye, M., Zhu, W., 2020. Strategic informed trading and dark pools. Available At SSRN 3292516.

Zhu, H., 2014. Do dark pools harm price discovery? Rev. Financ. Stud. 27 (3), 747–789.