Banning Dark Pools: Venue Selection and Investor Trading Costs*

Christian Neumeier† Arie Gozluklu‡ Peter Hoffmann§

Peter O'Neill¶ Felix Suntheim‖

Abstract

In the context of a recent ban on dark pool trading (the 'Double Volume Cap (DVC)'), we analyze the relationship between transaction costs and venue choice. Using a novel proprietary transaction-level dataset from U.K. equity markets, we show that decisions to trade on venues with lower levels of pre-trade transparency are associated with a lower implementation shortfall — a measure of the cost of executing large orders. This reduction in transaction costs can also be observed for trading on alternative venues, such as 'periodic auctions', but only following the ban on dark trading. Using a Difference-in-Differences analysis, we find that the 'Double Volume Cap (DVC)' did not have a significant effect on the costs incurred for large trade executions in the U.K. equity market. There is also little evidence for heterogeneous effects across different types of end-investors.

Keywords: Dark trading, execution costs, implementation shortfall

JEL Classification: G10, G18

*This version: January 26, 2021. The views expressed herein are those of the authors and do not necessarily represent the views of the institutions to which they are affiliated.

†Financial Conduct Authority, RoZetta Institute Group, Justus-Liebig University, Macquarie University. Email: christian.neumeier@rozettainstitute.com.

‡Warwick Business School. Email: arie.gozluklu@wbs.ac.uk.

§European Central Bank (DG-Research). Email: peter.hoffmann@ecb.int.

¶Financial Conduct Authority. Email: peter.oneill@fca.org.uk.

‖International Monetary Fund. Email: FSuntheim@imf.org.

1 Introduction

Investors transacting in modern equity markets must select from a long menu of venues to execute their trades with varying degrees of market transparency. On the one extreme, the *'lit'* markets of public exchanges offer high degrees of pre- and post-trade transparency, while *'dark'* trading venues (such as dark pools) offer the least. Aside from these extremes, choices also include venues and trading mechanisms that cannot be classified as either dark or lit, such as periodic (batch) auctions or systematic internalizers (SIs). Given this complex net of routing choices, order exposure decisions have become increasingly more important for investors.

While the proliferation of different trading venues has received considerable attention in the academic literature, empirical evidence on the effects of routing decisions on trading costs remains scant at best (e.g. Anand et al. (2019); Battalio et al. (2018); Gomber et al. (2016)). This gap in the literature is largely due to the lack of detailed data which means it is not possible to assess trade execution performance in a setting with a large set of available venue types.

The effects of transparency on market participants is also an important regulatory issue since price discovery and fair access to financial markets are of major concern to regulators.[1] In response to the increasing market share of dark trading venues, and their effects on market quality, European regulators introduced the so-called 'double-volume cap' (DVC) on 12 March 2018. This policy banned trading on a set of trading venues with no pre-trade transparency (referred to as 'dark pools').

Against this backdrop, this paper aims to shed light on two important questions: First, what role does dark trading play in increasing or decreasing investors' trading costs? Second, what was the effect of the DVC?

First, we show that the use of trading venues with lower pre-trade transparency is associated with lower execution costs. Importantly, we obtain these results after controlling for a wide set of fixed effects at the investor, broker, and stock-day level. Interestingly, we find that the effects of dark and periodic auction (PA) trading are *very similar*, suggesting that these trading mechanisms are close substitutes in terms of their effect on execution costs. Dark pools offer less transparency than PA venues as they do not disclose any volume information while PAs reveal 'indicative auction uncrossing volumes'. However, as with dark pools, but unlike central limit order books, PAs do not disclose buying and selling interest at individual price levels. See

[1] See, e.g. the SEC's officially stated goals, available here or ESMA's MiFID II/MiFIR Review Report, available here.

Section 2 for a more detailed discussion of venue characteristics.

Previous research has focused on the impact of dark trading on standard measures of market quality (e.g. spreads and depth), with mixed evidence. Some papers (Conrad et al. (2003); Buti et al. (2016); Garvey et al. (2016); Gresse (2017)) find that dark trading improves market quality through increased liquidity from lower transactions costs for individual trades. Others show that the effect is either not significant (Comerton-Forde et al. (2018); Farley et al. (2018); Foley and Putnins (2016)) or even detrimental to the lit markets (Degryse et al. (2015)). Comerton-Forde and Putnins (2015) show that the effect on price discovery is non-linear, with adverse consequences of high levels of dark trading activity. An experimental study by Bloomfield et al. (2015) finds that changes to market opacity affect trading strategies of both informed and uninformed participants, but do not significantly affect liquidity. Because dark pools are designed to reduce the 'market impact' of large block executions, studies that only examine the bid-ask spreads of individual trade executions are incomplete. By contrast, we are able to analyze the implementation shortfall of complete parent orders. To our knowledge, we believe this is the first paper to examine the effects of dark trading in a multi-venue setting.

Similar to Menkveld et al. (2017), we find evidence for dark pools to be first in the venue selection 'pecking order' within the parent order life cycle. Trade executions in dark pools are more likely to occur earlier in the trading day, when there is a lower demand for immediacy.[2]

Second, we examine the causal impact of the DVC trading restriction in a Difference-in-Differences setting, both when it was introduced and when it was subsequently lifted. We find that the DVC did not have any significant impact on investors' execution costs. Investors that relied heavily on dark pool trading before the ban did not experience a change in their implementation shortfall, relative to a control group of investors that did not rely much on dark trading. The same is true for the subsequent removal of the DVC.

We provide evidence consistent with investor shifting their activity to alternative venues that were not affected by the policy. These venues offer trading with limited pre-trade transparency, such as periodic auctions. In other words, it is likely that the DVC policy did not affect transaction costs because substitutes like periodic auction based venues are available. Importantly, after the ban is lifted, we find participant volume moved back to dark pools, suggesting that

[2]There is a growing literature highlighting the importance of examining investor venue routing decisions, especially in the context of high frequency trading (Battalio et al. (2018); Chakrabarty et al. (2019); Hendershott et al. (2013); Sağlam et al. (2019); van Kervel and Menkveld (2019)). However, these studies do not examine the trade-off between pre-trade transparency and execution quality in fragmented markets.

participants prefer dark pools to periodic auctions when both options are available. However, this preference is not universal, with significant volumes remaining on PA venues.

Finally, we also find no evidence that different investors reacted differently to the DVC. To this end, we distinguish investors based on their informedness and size. Neither of these characteristics affected the estimated treatment effects.

Our findings on the effects of the DVC complement those of Johann et al. (2019), who show that the DVC did not affect lit market quality and led trading to move from dark venues to close substitutes. A recent study by Guagliano et al. (2020) extends the analysis by including the lifting of the ban, finding that market liquidity improves during the ban periods. The authors also highlight the increased use of periodic auctions. Our data allows us to identify counterparties, and so enables us to assess the policy's impact on individual market participants. We are also able to examine the spectrum of venue choices, demonstrating that trading in periodic auctions, in addition to large-in-scale dark and regular dark trading, reduces investor transaction costs.

The rest of the paper is organized as follows. Section 2 describes the regulatory environment. In Section 3 we provide the details of our dataset and descriptive statistics. Section 4 contains our empirical results. In Section 5 we conclude.

2 Regulatory Environment

The second iteration of the Markets in Financial Instruments Directive (MiFID II), a suite of new regulations for EU capital markets, came into force on January 3, 2018. As part of this set of rules, the so-called "double volume cap" (DVC) restriction on "dark pool trading", became effective a few months later, on March 12, 2018.

Under MiFID II, trades on regulated markets or multilateral trading facilities that are not pre-trade transparent must trade under at least one of four conditional "waivers": i) the "Large In Scale" (LIS) waiver, for trades that are sufficiently large (often termed "block trades"); ii) the "reference price waiver", for trades referencing a "widely regarded reference price" - typically dark pool MTFs (for example UBS MTF) referencing the primary market mid-quote; iii) the "negotiated trade waiver", for trades that are negotiated off-market but formalized on-market; or iv) the "order management facility" waiver, for trades that are held within the exchange, pending disclosure - in practice "iceberg" orders.

Under the DVC, all trading under reference price or the negotiated trade waivers in the

respective instrument is banned for a duration of six months if it exceeds any of two pre-defined thresholds. The two thresholds are: i) a market-wide cap, triggered if the total volume across EU dark pools exceeds 8% of the total traded volume in the preceding 12 months, and ii) a venue-specific cap that is triggered by a specific dark pool exceeding a share of 4% of the volume in the preceding 12 months.

Under MiFID I, dark trading could also occur on so-called "Broker Crossing Networks" (BCNs), such as Credit Suisse's "Crossfinder" venue. As these venues were unregulated, they did not require a pre-trade transparency waiver. They were banned under MiFID II - effective from 3 Jan 2018.

There are several trading mechanisms not subject to the DVC, which are potential substitutes to dark pools:

- Similar to dark pools, "Systematic Internalizers" (SIs) publish quotes based on primary or market-wide best-bid or offer prices. They were touted as alternatives to dark pools ahead of the ban, and several were created in anticipation of it, such as those operated by proprietary trading firms Virtu, Citadel and Hudson River.

- BATS Chi-X Europe (Now CBOE), was the first to develop a 'periodic batch auction'[3] mechanism where participants can submit orders with the option of pegging to the mid-point of the European Best Bid or Offer price (EBBO). As the EU allows trading to occur across different countries, the EBBO is analogous to a 'National Best Bid or Offer'(NBBO). These auctions are triggered on order entry, occur throughout the day and can be as frequent as several times a second. Periodic batch auctions provide some pre-trade transparency by disclosing an indicative uncrossing price and volume for the auction. But they do not disclose the buying and selling interest at each price level, as in a conventional lit market auction. Orders in the periodic auction can specify the price to reference the EBBO mid-price at the time of the auction, equivalent to dark pool MTF and BCN reference of the primary midpoint under the reference price waiver. So, batch auctions provide slightly more pre-trade transparency than dark pools, while retaining the functionality of hiding a given participant's order, and allowing reference pricing.

- Trades that are designated Off-book, or Over-the-Counter (OTC) are the outcome of bi-lateral negotiations with brokers or other participants - usually brokers source liquidity

[3] Also referred to as 'frequent batch auction'

on behalf of clients via internalization or their dealer networks.[4]

- Finally, dark trades that use the LIS waiver could potentially be considered a substitute for smaller reference price waiver dark trades, if traders are able to modify their execution strategies to aggregate child orders.

3 Data and Descriptive Statistics

We source transaction level trade data from the Financial Conduct Authority's (FCA) Market Data Processor (MDP) database. Importantly, these data allow for the identification of individual market participants.[5] We start by selecting all stocks that were a constituent of the FTSE 100 and FTSE 250 share index during the period January 2018 to October 2018. We then restrict our sample to the 327 stocks that were classified as 'liquid' by the European Securities and Markets Authority (ESMA).[6] Our sample period comprises of 80 days in the period February 12 to October 11, 2020. More specifically, it covers the 40-day event windows around each of the two events: the introduction of the DVC on March 12 and its lift on September 12 respectively. We complement this dataset with quote data from Refinitiv (Thomson Reuters) Datascope Select. Data to classify counterparties is taken from Orbis and internal FCA sources.

DVC effects on UK equity markets

ESMA publishes a monthly list of suspended and non-suspended stocks together with their share of dark trading volume.[7] With the implementation of the DVC on March 12th, dark trading was banned for 257 of our sample stocks.

Table 1 reports descriptive statistics. We label the period around the implementation of the DVC as BAN (from February 12 to April 10), and the period around its lift as LIFT (from

[4]These trades use the OTC or the Negotiated Trade waiver, or are executed on the LSE without a Central Clearing Counterparty. See Appendix 6.2

[5]The dataset has been anonymized by the Financial Conduct Authority before being handed over to the authors, so that identification of individuals is not possible. The lowest level of identification is the Legal Entity Identifier (LEI).

[6]In order to label a stock liquid/illiquid, ESMA calculates the Standard Market Size at the stock level, which is based on the average value of a transaction (see Article 11 and Annex II of COMMISSION DELEGATED REGULATION (EU) 2017/587).

[7]ESMA is applying a 12 months rolling window to calculate the share traded under the use of the reference price and the negotiated transaction waiver. The data can be found at https://www.esma.europa.eu/double-volume-cap-mechanism.

August 14 until October 11).[8] Stocks affected by the suspension have lower spreads, a higher number of transactions, and a large percentage of dark trading (Panel A). Panel B shows the variables for the 225 stocks after the ban was lifted and for the 67 stocks that were never subject to the ban. Naturally, the waiver percentage correlates with the stock's share of dark trading, but the ban has elicited a decrease after it is lifted. Interestingly, even unaffected stocks show lower activity in dark trading and smaller waiver usage during the second event observation period compared to the pre-ban period.

Table 1: Descriptive Statistics

This Table contains statistics for banned and lifted stocks in first observation period (February 12th - April 12th) and second observation period (August 14th - October 11th). We include stocks that are a constituent of the FTSE100 or FTSE250 index at any given month over the complete observation period and are classified as liquid, based on ESMA classification. Liquidity Measures and 'Dark trading %' for the respective groups are calculated in the pre-BAN period and the post-LIFT period respectively. Thereby, we exclude the event days when the ban commences and the lift occurs for the first time (March 12th and September 12th), as well as the quadruple witching dates (March 16th and September 21st). Waiver % (Dark trading %) gives the average reference price waiver usage (average dark pool share of trading) across stocks in the period when dark trading is allowed, i.e. before the ban commenced and after the suspension got lifted.). The effective spread is calculated as $effective\ spread = 2d(price - mid)/mid$, where d indicates a buy or sell order. Data for liquidity metrics is taken from Refinitiv and covers the LSE CLOB and have been winsorized at the 1% level. Trades have been signed with the Lee and Ready (1991) algorithm and if the trade executes at the mid, and an Institutional Investor is either buyer or seller (not both), we classify this trade as either buyer or seller initiated, depending on the side the Institutional Investors trades. 'Banned' indicates stocks that have been suspended from dark trading, 'Lifted' indicates stocks for which the suspension has been lifted again. Stocks that are labeled 're-suspended' by ESMA are not considered.

Panel A. First Event (BAN)		
	Not Banned	Banned
# Stocks	72	257
Average Daily Turnover (GBP mil)	32.96	46.68
Average Daily Trades	4051	6680
Waiver %	5.41	11.28
Dark trading %	3.71	6.51
Effective Spread	16.83	9.21
Panel B. Second Event (LIFT)		
	Not Lifted	Lifted
# Stocks	67	225
Average Daily Turnover (GBP mil)	34.13	47.71
Average Daily Trades	3971	6345
Waiver %	4.63	5.03
Dark trading %	2.48	4.62
Effective Spread	15.17	8.31

[8] We exclude the days of the ban and the lift (March 12 and September 12), as well as the quadruple witching dates (the third Friday of every March, June, September and December. On these days, the expiry of listed derivatives causes abnormal trading volume.)

Table 2 shows how the breakdown of trading activity across venue types changed during the BAN and LIFT periods, both for suspended and non-suspended stocks. We classify venues/trading mechanisms as follows: i) Auctions (traditional opening, midday and closing auctions on LSE), ii) Dark (trading venues that are exempt from pre-trade transparency applying the reference price waiver), iii) Dark (LIS) (trades that use the large-in-scale waiver), iv) Lit (fully pre- and post-trade transparent trading venues, i.e. regular exchanges and multilateral trading facilities), v) Off-book (bilateral trades or trades using the OTC waiver), vi) Periodic Auctions (Frequent Batch Auctions) and systematic internalizers.

Dark trading for suspended stocks ceased after the ban, with a small decrease for non-banned stocks.[9] Lit trading decreased significantly, but less in suspended stocks than in non-suspended stocks, with trading volume migrating towards periodic auctions, off-book, and to systematic internalizers. We formally test for these changes in venue share using an unbalanced Difference-in-Differences regression in Panel C, where suspended stocks are the treated and non-suspended stocks are the control group. Suspended stocks show a significant and positive coefficient for Auction, Lit and Periodic Auction trading, and a significant negative difference for dark trading.

With the end of the first suspension period on September 12th, dark trading volumes increased again. However, the increase was not of the same magnitude as the previous decrease after the ban was introduced. Similarly, Lit trading showed a significant decrease of about 5 percent. Within the affected stocks the increased trading volume went to all other trading venues, with Off-book trading taking the lion's share (Panel A). The coefficient for un-suspended stocks shows the opposite sign for Off-book trading, but the same direction for Periodic, SI and Dark (LIS). Notably, trading in periodic auction venues did not return to the pre-BAN level, but remained at a significantly higher level for both affected and unaffected stocks (column (8)) of Table 2). Hence, market participants continued to trade on alternative trading venues even after trading on traditional dark pools becomes available again.

Trader Types on UK Trading Venues

In order to assess how venue choice, and the DVC, affect execution costs, we focus on the trading activity of institutional investors. These are real-money investors that typically trade directionally and execute large blocks of shares over time in an effort to minimize transaction

[9] On September 24th and 25th we exclude trading activity in SKY, after Comcast announced they will acquire the company. Inclusion causes very large SI activity, most likely OTC trades.

Table 2: Changes in Share of Trading by Venue Type - Around Ban and Lift Events

This Table contains the change of trading of a stock on a particular venue around BAN and LIFT events. We include stocks that are a constituent of the FTSE100 or FTSE250 index at any given month over the complete observation period and are classified as liquid, based on ESMA classification. We report differences between the pre-BAN (12 February to 9 March) and post-BAN (13 March to 12 April) periods and pre-LIFT (14 August to 9 September) to post-LIFT (13 September to 11 October). Trading on a venue is calculated as the ratio of turnover in the respective venue to total turnover on stock level per period. We show the results of a t-test, with clustered standard errors on stock level. Panel A (B) reports the results for stocks (not) affected by trading suspension due to BAN and LIFT events. Levels and differences are shown in percentage points. Standard errors in brackets. The sample in Panel A includes 257 stocks that were banned in the first three columns, 225 stocks that were lifted, and in the next three columns and 216 that were subject to both (the overlapping sample). The sample in Panel A includes 72 stocks that were not banned in the first three columns, 67 stocks that were not lifted in the next three columns and 53 that were subject to both (the overlapping sample). Panel C reports estimates for a Difference-in-Differences model, where affected stocks are the treatment group and unaffected stocks are the control group. We estimate the following model: $venue\ share_{j,t} = \alpha_j + \gamma_t + \delta(treated_j \times post_t) + \varepsilon_{j,t}$.

	(%) pre	(%) post	post-BAN to pre-BAN	(%) pre	(%) post	post-LIFT to pre-LIFT
Panel A. Stocks that were banned, and stocks that were lifted						
# stocks			257			225
Auction	10.41	10.50	0.09	11.43	11.22	-0.22*
			(0.15)			(0.12)
Dark	6.00	0.00	-6.00***	0.00	4.46	4.46***
			(0.17)			(0.12)
Dark (LIS)	2.22	2.44	0.21	2.51	2.63	0.12
			(0.16)			(0.17)
Lit	36.74	36.24	-0.49**	33.88	32.34	-1.53***
			(0.25)			(0.23)
Off-book	20.12	24.16	4.04***	24.01	22.73	-1.28***
			(0.33)			(0.31)
Periodic Auction	0.53	2.05	1.52***	3.38	1.56	-1.82***
			(0.06)			(0.09)
SI	23.98	24.61	0.63**	24.78	25.05	0.27
			(0.27)			(0.27)
Panel B. Stocks that were not banned and not lifted						
# stocks			72			67
Auction	8.19	7.27	-0.92***	7.94	8.61	0.67***
			(0.26)			(0.25)
Dark	4.40	3.85	-0.55*	2.68	3.05	0.38
			(0.32)			(0.28)
Dark (LIS)	1.21	1.07	-0.14	0.82	1.09	0.27*
			(0.27)			(0.16)
Lit	25.38	23.39	-1.99***	24.81	25.88	1.06*
			(0.56)			(0.56)
Off-book	45.59	48.38	2.79***	45.31	42.69	-2.62***
			(0.77)			(0.63)
Periodic Auction	0.28	0.60	0.32**	1.42	1.39	-0.03
			(0.15)			(0.24)
SI	14.94	15.43	0.49	17.03	17.29	0.26
			(0.51)			(0.63)
Panel C. Unbalanced Difference-in-Differences with affected and un-affected stocks						
# stocks			329			292
Auction			1.09***			-0.85***
			(0.33)			(0.28)
Dark			-5.43***			4.09***
			(0.37)			(0.30)
Dark (LIS)			0.35			-0.16
			(0.32)			(0.24)
Lit			1.29**			-2.45***
			(0.64)			(0.61)
Off-book			1.38			1.14*
			(0.84)			(0.68)
Periodic Auction			1.19***			-1.78***
			(0.16)			(0.26)
SI			0.12			0.01
			(0.58)			(0.71)

costs. In the 20 days leading up to the introduction of the DVC, these investors account for 10.15 percent of the total trading activity.[10]

In order to assess the execution costs of institutional investors, we construct *parent orders* as the sum of all individual trade executions (henceforth referred to as 'child orders') on each side of a stock-day-participant-broker combination.[11] Our sample includes 58,437 parent orders, which we require to be of at least 100,000 GBP in total size, to consist of at least five child orders, and whose execution takes at least ten minutes.[12] Moreover, parent orders must have a directionality of at least 90%.

We measure transaction costs using the Implementation Shortfall (IS), developed by Perold (1988). It is defined as

$$IS = D \times \frac{p - p_0}{p_0},$$

where p_0 is the mid-quote at the time the trade starts (execution of first child order), p is the value-weighted execution price of the entire parent order, and D is a trade direction indicator ($D = 1$ for buy orders and $D = -1$ for sell orders). We winsorize the Implementation Shortfall at the 1%.

Panel A of Table 3 provides summary statistics on parent orders executed in periods when dark trading is not affected by the DVC. The average parent order has a value of 989 thousand GBP, consists of 166 child executions and needs almost four and a half hours to be fully executed. The average implementation shortfall is 12.13 basis points and increases in order size. For orders above 1 million GBP, the implementation shortfall is 13.87 bps, compared to 7.47 bps for parent orders below 1 million GBP. Similarly, the average execution time increases from 4.15 hours to 5.04 hours.[13] Panel A of Table 3 also provides the breakdown of parent orders

[10] While this appears a rather modest percentage, it is important to know that many institutional investors (especially hedge funds) trade in U.K. equities through derivatives such as Total Return Swaps, Contracts for Differences, and Spread-Bets. Accordingly, some of their trading is reflected by the hedging activity of broker-dealers in the cash market, but cannot be allocated to individual institutional investors. While some information about derivatives trading is available in the MDP database, it is subject to data quality issues and can thus only be interpreted in aggregate.

[11] The terms child trade and child order are used interchangeably. We exclude the 1% child transactions with the largest price deviations relative to the current mid-quote. More information about mapping parent and children orders can be found in Appendix 6.2.

[12] Additionally, we require the Volume Weighted Average Price (VWAP) of parent orders to deviate no more than 1bps from the 'true' parent order VWAP. The 'true' VWAP is the price reported in the MDP report. More details in Appendix 6.2

[13] By our definition, parent orders cannot span more than one business day.

Table 3: Descriptive Statistics of Investor Parent Orders

This table contains descriptive statistics of parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. Parent orders have been winsorized at the 1% level with respect to Implementation Shortfall. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Average value (GBP millions) is the average parent order size in GBP, average number of children is the average number of child trades per parent order, average duration is the average time difference of first trade to last trade in parent order per stock-day-participant in hours, average IS is the average volume weighted Implementation Shortfall per parent order in bps, Number of parent orders and number of children gives the total number of each. Additionally, we show the share of venue type usage across parent orders. Panel A shows parent order statistics during dark-trading periods (pre-BAN and post-LIFT). Panel B shows parent order statistics across periods. We include trades from stocks that are liquid and subject to suspension and lifting. Standard deviation for the respective measure is shown in brackets.

	All	Above 1 mln	100k to 1 mln
Panel A. Parent Order Characteristics during dark-trading periods			
Average value (GBP, millions)	0.99	3.12	0.35
	(2.15)	(3.73)	(0.23)
Average number of children	165.63	417.37	90.23
	(276.83)	(475.65)	(84.63)
Average duration	4.36	5.04	4.15
	(3.05)	(2.92)	(3.05)
Average IS (bps)	12.13	13.87	7.47
Number of parent orders	29,404	6,777	22,627
Number of children	4,870,239	2,828,539	2,041,700
Number of participants	840	499	789
Auction (%)	15.33	15.99	13.55
Dark (%)	16.19	16.34	15.78
Dark (LIS) (%)	11.10	14.73	1.41
Lit (%)	46.83	42.89	57.35
Periodic Auction (%)	2.92	2.62	3.72
SI (%)	2.41	2.16	3.07
Off-book (%)	5.23	5.26	5.13
Panel B. Parent Order Characteristics over full sample (dark and no-dark periods)			
Average value (GBP, millions)	0.95	3.07	0.35
	(2.18)	(3.95)	(0.23)
Average number of children	167.21	427.31	93.32
	(272.40)	(467.88)	(91.74)
Average duration	4.37	5.13	4.16
	(3.06)	(2.94)	(3.06)
Average IS (bps)	12.42	14.41	7.43
Number of parent orders	58,437	12,928	45,509
Number of children	9,771,014	5,524,247	4,246,767
Number of participants	989	632	931
Auction (%)	15.55	16.19	13.94
Dark (%)	8.53	8.76	7.97
Dark (LIS) (%)	11.72	15.67	1.81
Lit (%)	49.83	45.72	60.13
Periodic Auction (%)	5.69	5.13	7.10
SI (%)	2.42	2.18	3.05
Off-book (%)	6.25	6.35	6.00

across trading venues/mechanisms. On average, 46.83 percent of the total value is traded in lit venues, 16.19 percent in dark, 15.22 percent during auctions, 11.10 percent in large-in-scale Dark, 5.23 percent off-book, 2.92 percent in periodic auctions, and 2.41 percent on systematic internalizers. Large orders display a significantly lower share of Lit trading and, naturally, a significantly larger use of large-in-scale dark trading. Panel B shows the parent order statistics when combining dark and no-dark period observations. Our data covers a total of 989 unique investors, among which only 632 are engaged in the execution of orders larger than 1 million GBP.

Menkveld et al. (2017) develop a venue pecking order theory according to which market participants first attempt to trade in dark venues, and over time resort to more transparent trading mechanisms as order execution becomes more important. Figure 1 provides some visual evidence that is consistent with this view. The top Panel shows the breakdown of parent buy orders across venues/mechanisms over the order life cycle. The life-cycle is constructed by splitting the parent order in chronological order into quintiles. Hence, the first bar shows the venue distribution in the first 20% of the parent order life-cycle, the second bar shows the venue usage for 20% to 40%, and so on.

Importantly, Figure 1 only includes parent orders from pre-BAN and post-LIFT period, i.e. when dark trading is not subject to the DVC. We find that the share of dark and large-in-scale dark during the first 20% of the parent order is larger than during the remaining parent order. In the first quintile shown in Figure 1, dark and large in scale dark venues account for 41.9% of trade executions. During the lifetime of a parent order, the share of dark and large in scale dark trading decreases and eventually drops to 18.6%, yet during the last (20%) part of the parent order life cycle a large portion is executed in the closing auction, as the last bar in the upper panel indicates. The lower panel of Figure 1 shows a similar pattern over trading hours. However, the preference for dark venues during the early trading hours is not as salient as in the beginning of the parent orders' life-cycle.

This is confirmed by the results from a simple linear probability model in Table 4, where we regress dummy variables indicating the use of dark trading venues on a set of order life cycle dummy variables and fixed effects. The probability of choosing dark venues (and large in scale dark venues) decreases with the duration of the parent order. Interestingly, a similar pattern is observable for periodic auctions during times when dark trading is banned.[14]

[14](Tables A3 and A2) show results from a multinomial logit regression and provide additional insights on venue

Figure 1: Parent Order Venue Choice - by % Depleted and by Trading Hour
The figure below shows the parent order life-cycle. We include parent orders and their corresponding child executions when dark trading is possible, i.e. in the pre-BAN and post-LIFT period. Venue choice is reported for each quintile of original parent order value. Quintiles are calculated on parent order level (Parent orders are all child orders summed up by day, stock, broker, participant). Hence, the below life-cycle is an equally weighted average display of the depletion within parent orders. Parent orders must have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children.

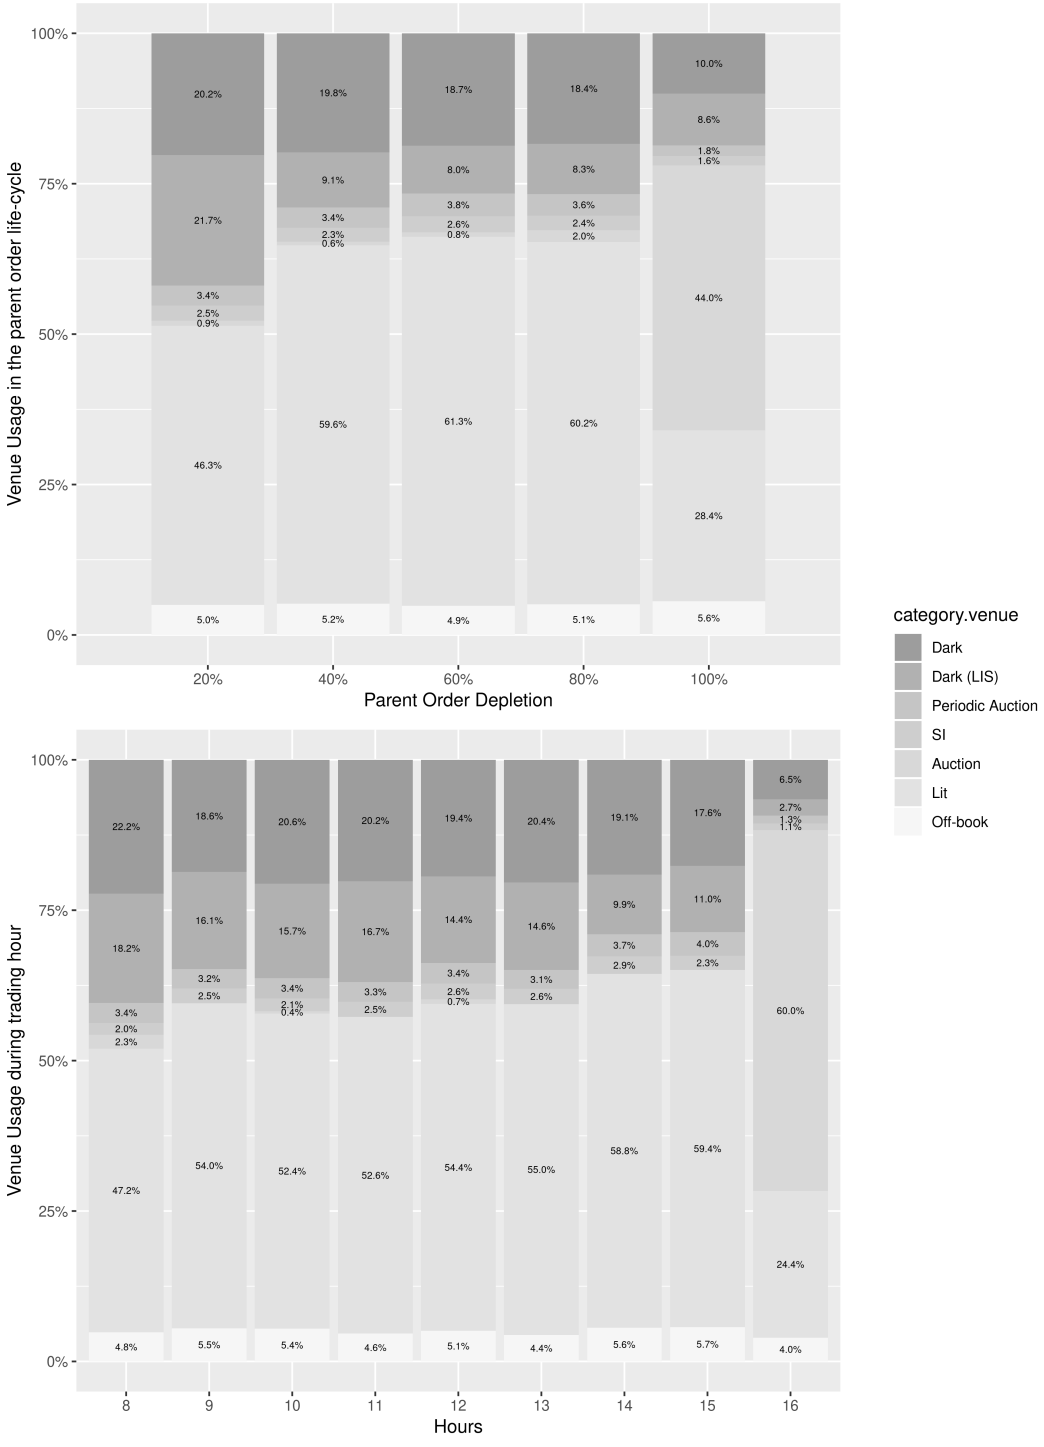

Table 4: Parent Order Venue Choice - Linear Probability Model
The table below shows the results from a linear probability model that shows the probability of a child being executed at a certain time in the life of the parent order on specific venue choices. We create a binary variable for each of the columns presented below that equals to 1 if the child order has been executed in a) Dark venues, b) Dark venues including Dark (LIS), c) Periodic Auction, d) Periodic Auction including Dark (LIS). The binary variable equal 0 if the child is not executed in the venues of interest. Columns (1) and (2) thereby consider periods when dark trading is allowed (pre-BAN and post-LIFT) whereas columns (3) and (4) consider periods when dark trading is prohibited (post-BAN and pre-LIFT). 'Depletion Bucket' indicates the time in the parent orders life in quintiles, e.g. 'Depletion Bucket 2' indicates child orders that are executed between 20% and 40% of the parent orders life cycle. The reference level in the below is the first bucket, i.e. the first 20% of the parent order. We include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level.

	Dependent variable:			
	Dark	Dark and Dark (LIS)	Periodic Auction	Periodic Auction and Dark (LIS)
	(1)	(2)	(3)	(4)
Depletion Bucket 2	-0.005^{***}	-0.005^{***}	-0.003^{**}	-0.003^{***}
	(0.002)	(0.002)	(0.001)	(0.001)
Depletion Bucket 3	-0.007^{***}	-0.007^{***}	-0.002^{**}	-0.003^{**}
	(0.002)	(0.002)	(0.001)	(0.001)
Depletion Bucket 4	-0.010^{***}	-0.011^{***}	-0.002	-0.003^{*}
	(0.002)	(0.002)	(0.002)	(0.002)
Depletion Bucket 5	-0.016^{***}	-0.016^{***}	-0.004^{***}	-0.004^{***}
	(0.002)	(0.002)	(0.001)	(0.001)
Stock-Day FE	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes
Period	pre-BAN and post-LIFT	pre-BAN and post-LIFT	post-BAN and pre-LIFT	post-BAN and pre-LIFT
Observations	4,851,067	4,851,067	4,885,124	4,885,124
R^2	0.315	0.318	0.170	0.174
Adjusted R^2	0.314	0.317	0.168	0.173
Note:				*p<0.1; **p<0.05; ***p<0.01

4 The Impact of Venue Choice on Investor Trading Costs

In this section we test how the choice of trading venues affects execution costs as measured by the implementation shortfall. Using data from the pre-BAN and post-LIFT periods where dark trading is allowed, we estimate the following regression:

$$IS_\tau = \alpha + \sum_{n=1}^{N-1} \beta_n \, PctVenue_{n,\tau} + \gamma_1 Size_\tau + \gamma_2 Execution\,time_\tau + FE + \varepsilon_\tau, \tag{1}$$

where $PctVenue_{n,\tau}$ is the share of parent order τ executed in venue n. The idea is to compare the implementation shortfall across trades with different levels of dark trading that are otherwise similar. To do so, we include a rich set of fixed effects (FE) such that we can compare trade executions that take place in the same stock on the same day, and we additionally control for observed and unobserved heterogeneity at the investor and broker level, which all may affect execution quality. We additionally control for the trade size ($Size_\tau$) in GBP, and the time of execution ($Execution\,time_\tau$), measured in hours. For the latter two we transform the variable using the natural logarithm.

The results in Table 5 focus on the period where dark trading is not subject to bans, showing that dark trading is associated with significantly reduced execution costs. For example, a 10 percent increase in the proportion of a parent order executed on a dark venue reduces implementation shortfall by 0.97 bps.

Interestingly, the effects of dark and large-in-scale dark trading are qualitatively similar, and we cannot reject the null hypothesis that they are equal. In contrast, lit trading and off-book trading are associated with significantly higher execution costs.[15] As expected, larger trades incur a higher implementation shortfall, while there is no effect of execution time.

Table 6 focuses on periods with constrained dark trading, the post-BAN and pre-LIFT periods, so does not contain 'Dark (%)'.[16] Results are similar to the unconstrained period, with lit trading showing a positive impact on transaction costs and large-in-scale dark trading showing a negative impact. Interestingly, during periods with no dark trading, the reduction of transaction costs associated with trading in periodic auction venues is similar extent to the effect of

choice.

[15] We obtain qualitatively similar results when comparing the pre-BAN and post-LIFT periods. See Appendix 6.1.

[16] We exclude a small number of trades that are reported to have occurred in dark venues for these periods, as they are likely erroneous transaction reports. They account for less than 0.01% of total trades.

Table 5: Effect of Venue Trading Share on Implementation Shortfall - Non-Ban Period
The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to regression specification 1. We include parent orders from both the pre-BAN and post-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:							
	Total IS (bps)							
Lit (%)	0.057^{***}							
	(0.020)							
Dark (%)		-0.073^{***}						-0.097^{***}
		(0.022)						(0.024)
Dark (LIS) (%)			-0.161^{***}					-0.203^{***}
			(0.036)					(0.042)
Periodic Auction (%)				-0.012				-0.037
				(0.041)				(0.041)
Auction (%)					0.015			-0.022
					(0.031)			(0.031)
SI (%)						0.023		0.007
						(0.063)		(0.065)
Off-book (%)							0.066^{**}	0.026
							(0.029)	(0.034)
Size	3.876^{***}	3.566^{***}	4.169^{***}	3.586^{***}	3.557^{***}	3.599^{***}	3.582^{***}	4.317^{***}
	(0.646)	(0.631)	(0.635)	(0.617)	(0.641)	(0.619)	(0.619)	(0.673)
Execution time	-0.368	-0.342	-0.395	-0.238	-0.244	-0.234	-0.235	-0.561
	(0.615)	(0.602)	(0.589)	(0.597)	(0.592)	(0.595)	(0.597)	(0.598)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	28,616	28,616	28,616	28,616	28,616	28,616	28,616	28,616
R^2	0.370	0.371	0.371	0.370	0.370	0.370	0.370	0.371
Adjusted R^2	0.106	0.106	0.106	0.105	0.105	0.105	0.105	0.107
Note:								*p<0.1; **p<0.05; ***p<0.01

Table 6: Effect of Venue Trading Share on Implementation Shortfall - Ban Period
The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to equation 1. We include parent orders from both the post-BAN and pre-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:						
	Total IS (bps)						
Lit (%)	0.091^{***}						
	(0.019)						
Dark (LIS) (%)		-0.122^{***}					-0.164^{***}
		(0.036)					(0.038)
Periodic Auction (%)			-0.084^{***}				-0.117^{***}
			(0.018)				(0.017)
Auction (%)				-0.036			-0.069^{**}
				(0.029)			(0.030)
SI (%)					0.042		-0.003
					(0.047)		(0.045)
Off-book (%)						0.003	-0.041
						(0.030)	(0.033)
Size	4.401^{***}	4.455^{***}	3.862^{***}	4.009^{***}	3.953^{***}	3.933^{***}	4.681^{***}
	(0.556)	(0.606)	(0.568)	(0.558)	(0.562)	(0.565)	(0.599)
Execution time	-0.103	-0.147	-0.033	0.083	0.057	0.052	-0.269
	(0.624)	(0.611)	(0.610)	(0.598)	(0.609)	(0.610)	(0.607)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	28,335	28,335	28,335	28,335	28,335	28,335	28,335
R^2	0.376	0.375	0.375	0.375	0.375	0.375	0.376
Adjusted R^2	0.109	0.108	0.108	0.108	0.108	0.108	0.110
Note:							*p<0.1; **p<0.05; ***p<0.01

dark pools in the unconstrained period, at 1.17 basis points for a 10 percent increase in proportion traded. We find that regular auction participation also reduces transaction costs in the constrained period.[17]

The similar effects we observe from periodic auction and dark trading mechanisms implies that they might be close substitutes. We cannot say whether the transaction cost benefits originate from the characteristics of the venue itself, or the trading flows directed to it, or both. The fact that we observe an improvement for periodic auction trading only after dark trading is constrained implies that flow is important, but of course, flow is a function of investor choices in response to venue characteristics. We do find evidence for a migration of flow in Table 2.

Next, we examine the effects of the DVC on execution costs in a Difference-in-Differences setting. Our approach is based on the idea that the policy change will have a larger effect on those institutional investors that tend to trade more in dark pools. Specifically, we estimate the following Difference-in-Differences regression at the market participant level

$$IS_{j,t} = \alpha_j + \gamma_t + \beta(Dark\ participant_j \times Post_t) + \varepsilon_{j,t}. \tag{2}$$

In this specification, α_j and γ_t are participant and day fixed effects, $IS_{j,t}$ denotes the volume-weighted implementation shortfall for participant j on date t, $Dark\ participant_j$ is a dummy variable equal to 1 for active dark pool users, and zero otherwise; $Post_t$ equals 1 during the time period after the event, taking the value of 1 for the post-BAN and post-LIFT period, and zero otherwise. Active dark pool users are defined as institutional investors above the cross-sectional median of dark pool usage prior to March 12, 2020. Figure 2 plots the distribution of volume weighted dark pool usage by participant during the pre-BAN and shows that the median dark usage is 9.8%.

We run regression (2) for both BAN and LIFT periods separately to assess both the effects of the DVC's inception, and the lifting of the restriction. In addition, we also compare the pre-BAN and the post-LIFT periods, in order to check whether the effects of structural shifts in market shares across trading venues leads to an effect on average execution costs.

Table 7 presents the estimated treatment effects, where standard errors are clustered at the participant level. We observe that neither the introduction of the DVC, nor its suspension had a

[17] In Table A4 in the Appendix, we show the BAN led to a statistically significant decline in the execution costs associated with periodic auction trading for FTSE100 constituent stocks.

statistically significant effect. By and large, execution costs remained the same. This is also true when comparing the pre-BAN and the post-LIFT periods. Given our results, which show that alternative venues to dark pools provide similar execution cost benefits. A possible explanation is that banned trading flows migrated to these alternatives, which would mitigate the impact of the ban on investor trading costs.

Figure 2: Distribution of Participant Dark Pool Utilization - Pre-Ban Period
Distribution of participants according to average dark pool usage in the pre-BAN period (12 February to 9 March). We calculate the average dark pool usage as a volume weighted mean across all parent orders on participant level. We include parent orders that have at least 100,000 GBP in size, consist of at least five child transactions and last at least ten minutes. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP, parent orders must have at least a directionality of 90%. Parent Orders have to originate from stocks which are liquid and subject to the suspension or lifting. Bin size is 0.5%.

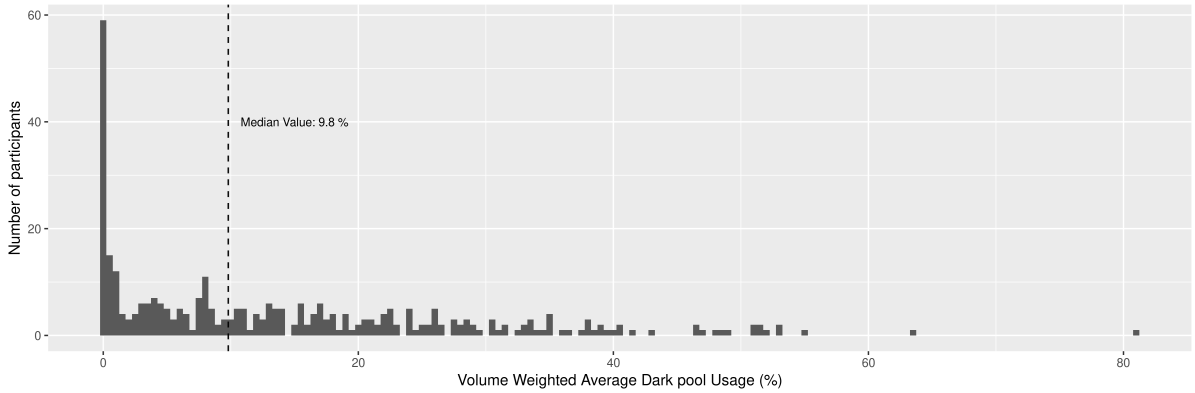

To examine the causes of the insignificant transaction cost effect in more detail, we investigate how dark participants' venue choice changed after the event using a Difference-in-Differences regression. Panel A. of Table 8 shows that dark users increase their share in Periodic Auctions compared to the non-dark users. They also increase their share in Lit venues by an even larger amount. When combined with evidence from Table 6, it could be argued that the increase in Lit markets has a negative impact on transaction costs, which offsets the positive impact of Periodic Auctions.[18]

[18] Additionally, Table A4 interacts the venue choice with the event (both Ban and Lift). This helps us to identify, if the venue choice impact on transaction cost is changing between periods. In column (1) the interaction between Periodic Auction (%) and Post is significantly negative, while the coefficient in column (2), which compares pre- and post-Lift remains insignificant. This indicates that the switch to Periodic Auctions initially has a positive impact on transaction costs. Table A4 also shows, that only the most liquid stocks (FTSE100) are affected, while less liquid stocks (FTSE250) do not show the same transaction cost benefit.

Table 7: Effect of Dark Pool Ban and Lift on Implementation Shortfall
The table below shows the baseline Difference-in-Differences estimates for three separate periods: pre-BAN (20 business days, 12 February to 9 March) to post-BAN (20 business days, 13 March to 12 April), and pre-LIFT (20 business days, 14 August to 9 September) to post-LIFT (20 business days, 13 September to 11 October) periods according to equation 2, as well as pre- BAN to post-LIFT periods. Observations are participant mean IS constructed from at least 10 parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. Participants are considered treated if they trade at or above the median value of dark trading across participants (are heavy users of dark venues and are thus impacted by the ban/lift), time post is one for the post-BAN and post-LIFT period. Standard errors are clustered by participant level.

	Dependent variable:		
	Total IS (bps)		
	BAN	LIFT	pre-BAN to post-LIFT
	(1)	(2)	(3)
Dark participant×Post	0.519	3.649	0.583
	(2.780)	(2.830)	(2.788)
Day FE	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes
Observations	6,199	5,546	5,778
R^2	0.106	0.112	0.097
Adjusted R^2	0.048	0.055	0.038
Note:	*p<0.1; **p<0.05; ***p<0.01		

4.1 Trader Heterogeneity and Venue Choice

Results in previous section are aggregated across all parent orders and will not capture unobserved behavior of (groups of) participants that execute their parent orders differently based on their level of trading informedness and size. Sağlam et al. (2019) show that trader abilities to forecast future returns (informedness) impacts their order size and venue choices. Therefore, the DVC may impact participants of different size or forecast ability differently.

To address this participant heterogeneity, we group participant types based on their forecasting precision, which we call informedness, and their size for participants within the category of Institutional Investors.

The informedness measure is calculated as the ability to predict a stock i's price movement between the closing price of day t and day $t+1$ ($return(t+1,t)_i$). We then run the following

Table 8: Effect of Dark Pool Ban and Lift on Participant Venue Choice

The table below shows estimates of Difference-in-Differences for two separate periods: pre-BAN (20 business days, 12 February to 9 March) to post-BAN (20 business days, 13 March to 12 April), and pre-LIFT (20 business days, 14 August to 9 September) to post-LIFT (20 business days, 13 September to 11 October) periods according to equation $PctVenue_{j,t} = \alpha_j + \gamma_t + \beta(Dark\ Participant_j \times \delta Post_t) + \varepsilon_{j,t}$, where $PctVenue_{j,t}$ is the share of participant j trading on each venue on day t, $Dark\ participant_j$ indicates active and non-active dark users, $Post_t$ equals 1 during the time period after the event (post-BAN and post-LIFT). Observations are participant mean venue shares constructed from at least 10 parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in 254 stocks that were subject to a Ban in the BAN period, 227 of those stocks which have their bans lifted in the LIFT period. Participants are considered treated if they trade at or above the median value of dark trading across participants (are heavy users of dark venues and are thus impacted by the ban/lift), time post is one for the post-BAN and post-LIFT period. After-hours are excluded. Standard errors are clustered by participant level.

	Dependent variable:					
Panel A. Ban Event						
	Periodic Auction	Lit	Dark (LIS)	Auction	SI	Off-book
	(1)	(2)	(3)	(4)	(5)	(6)
Dark participant×Post	0.068***	0.123***	0.011	0.019**	0.009**	0.001
	(0.009)	(0.018)	(0.007)	(0.009)	(0.004)	(0.013)
Day FE	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,199	6,199	6,199	6,199	6,199	6,199
R^2	0.250	0.438	0.217	0.250	0.562	0.455
Adjusted R^2	0.201	0.402	0.166	0.202	0.533	0.420
Panel B. Lift Event						
Dark participant×Post	−0.075***	−0.095***	0.005	−0.014	−0.015**	0.009
	(0.014)	(0.020)	(0.011)	(0.010)	(0.007)	(0.009)
Day FE	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,546	5,546	5,546	5,546	5,546	5,546
R^2	0.337	0.403	0.208	0.256	0.442	0.470
Adjusted R^2	0.295	0.364	0.157	0.208	0.406	0.436
Note:						*p<0.1; **p<0.05; ***p<0.01

regression at the participant level:

$$\begin{aligned} return(t+1,t)_i = \beta_0 + \beta_1 Trade\ Side_{i,\tau,p} + \beta_2 Size_\tau \\ + \beta_3 vola_i + \sum_{d=0}^{-4} \gamma_d\ Trade\ Side_{i,\tau,p} \times return(d,d-1)_i + \varepsilon_\tau, \end{aligned} \tag{3}$$

where $Trade\ Side_{i,\tau,p}$ is 1 for a buying parent order τ in stock i of participant p and -1 for a selling parent order, $Size_\tau$ is the natural logarithm of the trade size (measured in GBP) and $vola_i$ is the stocks intraday transaction price volatility. Additionally, we include five lagged one-day returns. We then follow Sağlam et al. (2019) and define participants as informed when they show a positive and significant sign (estimate of β_1) at the 10 percent level.[19]

Additionally, we categorize participants based on their (trade) size by taking the first and fifth quintile by total parent order size, where a higher quintile means larger size.

Next, Table 9 investigates the difference between the informed and matched control group in terms of venue choice. We use Propensity Score Matching to identify the matched control group.[20] The 30 informed participants use lit venues significantly more than the control group, both in times when dark trading is allowed and when dark trading prohibited, which is similar to the findings of Sağlam et al. (2019). During periods with no restrictions on dark trading, informed participants route fewer orders to dark venues and auctions (both regular and periodic), while during periods with restricted dark trading, also large-in-scale dark shows a significant negative sign. Additionally, the size of the coefficient for periodic auction increases in absolute terms.[21]

We augment the basic Difference-in-Differences approach by including an indicator if the participant is informed or large according to:

$$\begin{aligned} IS_{j,t} = \alpha_j + \gamma_t + \beta(Characteristic_j \times Dark\ participant_j) + \theta(Post_t \times Dark\ participant_j) \\ + \delta(Dark\ participant_j \times Characteristic_j \times Post_t) + \varepsilon_{j,t}, \end{aligned} \tag{4}$$

[19] We have too few observations with significant and negative coefficients to form the uninformed group.

[20] For the Propensity Score Matching we draw a matching sample out of the residual participants, where the binary dependent variable $informed$ equals 1 if the participant shows a positive and significant coefficient and 0 otherwise. We include the total parent order size, the average parent order size and the number of brokers used as explanatory variables.

[21] Table C17 applies a propensity score matching with different explanatory variables and still shows that informed participants prefer Lit venues in the pre-BAN and post-LIFT period, but coefficient looses significance in the post-BAN and pre-LIFT period.

Table 9: Venue Shares of Informed Participants - Around Ban and Lift Events

Average usage of venues between the informed group and matched group during periods of dark trading and periods of prohibited dark trading. Comparison of informed investors to matched control group. Investors are informed if the β_1 coefficient from equation 3 is positive and significant at the 10% level during the pre-BAN and post-LIFT period. Afterwards we match informed participants to a control sample based on trade size with a propensity score matching using a nearest neighbor algorithm (logit). Column 'Difference' shows the results of a regular t-test between the two groups. Column 'Difference (Fixed Effects)' shows the results of a regression of the form $venue(\%) = FE_{stock-day} + informed.dummy + \varepsilon$. Standard errors are clustered by stock-day. We use the same sample of 30 treated participants for both comparisons. Two participants from the control group are not active during the post-BAN and pre-LIFT period and we find two new matched participants based on a Propensity Score Matching (PSM) performed during the post-BAN and pre-LIFT period. The PSM is using 'Total Parent Order Size', 'Average Parent Order Size' and 'Average Number of Brokers'. Standard Errors in brackets.

	Informed Share (%)	Matched Share (%)	Difference	Difference (Fixed Effects)
Panel A. Period when dark trading is allowed (pre-BAN and post-LIFT)				
Number of Participants	30	30		
Total Parent Order Size (mln GBP)	5,879.15	4,211.61		
Average Parent Order Size (mln GBP)	0.96	1.33		
Average Number of Brokers	11.54	13.32		
Auction	12.80	17.97	-5.17***	-6.21***
			(0.53)	(1.30)
Dark	15.46	24.93	-9.47***	-8.12***
			(0.65)	(1.47)
Dark (LIS)	2.71	3.75	-1.05***	-1.07
			(0.31)	(0.66)
Lit	60.26	42.48	17.78***	16.86***
			(0.76)	(1.72)
Off-book	2.40	3.22	-0.82***	-0.49
			(0.30)	(0.69)
Periodic Auction	3.70	5.23	-1.54***	-1.13*
			(0.28)	(0.60)
SI	2.58	1.97	0.61***	0.45
			(0.21)	(0.46)
Panel B. Period when dark trading is prohibited (post-BAN and pre-LIFT)				
Number of Participants	30	30		
Total Parent Order Size (mln GBP)	4,442.08	3,356.23		
Average Parent Order Size (mln GBP)	0.84	1.05		
Average Number of Brokers	12.20	12.12		
Auction	15.51	18.54	-3.03***	-4.31***
			(0.60)	(1.36)
Dark (LIS)	3.30	4.57	-1.28***	-2.08**
			(0.37)	(0.88)
Lit	63.92	53.43	10.49***	14.46***
			(0.81)	(1.79)
Off-book	2.85	4.77	-1.93***	-0.96
			(0.38)	(1.06)
Periodic Auction	11.43	15.98	-4.55***	-7.21***
			(0.59)	(1.37)
SI	2.88	2.46	0.42*	0.31
			(0.24)	(0.49)

Table 10: Effect of Ban and Lift on Informed Dark Participant Implementation Shortfall
The table below shows the Difference-in-Differences estimates including informed and matched participants estimates for three separate periods: pre-BAN (20 business days, 12 February to 9 March) to post-BAN (20 business days, 13 March to 12 April), and pre-LIFT (20 business days, 14 August to 9 September) to post-LIFT (20 business days, 13 September to 11 October) periods according to 4. Observations are participant mean IS constructed from at least 10 parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in 254 stocks that were subject to a Ban in the BAN period, 227 of those stocks which have their bans lifted in the LIFT period. Investors are informed if the β_1 coefficient from equation 3 is positive and significant at the 10% level. Afterwards we match informed participants to a control sample based on trade size with a propensity score matching using a nearest neighbor algorithm (logit). Standard errors are clustered by participant level. We use the observation period pre-BAN and post-LIFT to create the treated group. Thereby, we identify 30 matches to the treatment group (i.e. 30 informed investors). Three participants from the control group are not active during the post-BAN and pre-LIFT period and we replace these with three new participants based on a PSM performed during the post-BAN and pre-LIFT period.

	Dependent variable:		
	Total IS (bps)		
	BAN	LIFT	pre-BAN to post-LIFT
	(1)	(2)	(3)
Informed × Post × Dark participant	−2.770	4.812	4.763
	(8.206)	(10.401)	(10.618)
Informed × Post	−4.914	−4.559	−7.055
	(5.988)	(9.137)	(6.559)
Post × Dark participant	−0.957	−6.189	−9.309
	(6.245)	(7.372)	(7.890)
Day FE	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes
Observations	1,196	1,122	1,161
R^2	0.096	0.137	0.094
Adjusted R^2	0.011	0.054	0.008
Note:	*p<0.1; **p<0.05; ***p<0.01		

where α_j and γ_t are participant and time fixed effects, $Dark\ participant_j$ indicates active and non-active dark users, $Post_t$ equals 1 during the time period after the event (post-BAN and post-LIFT), $Characteristic_j$ indicates either large vs small or informed vs uninformed (matched) investors. Within each quintile, we identify the median value for dark trading and assign participants to the treatment group, i.e. active dark traders, if they trade above the median value within each quintile.[22] The assignment of participants to their $Characteristic_j$ variable is carried out during the pre-BAN period when comparing the first event and the comparison between pre-BAN and post-LIFT. When examining the second event only, the assignment is based on post-LIFT observations.

In Table 10 we show the results of the estimation results of the modified Difference-in-Differences equation 4 and find no evidence that informed participants exhibit a significantly different transaction cost impact compared to their matched control group. In Table A11, we conduct a similar exercise where we differentiate participants based on their size (total trading activity) and also find no significant differences. In sum, we cannot find evidence for a heterogeneous impact of the DVC across different participants.

5 Conclusion

This paper provides evidence that investors can reduce their execution costs by selecting venues with less pre-trade transparency, such as dark pools or venues with similar characteristics. We find venue selection decisions matter. By analyzing 58,437 parent orders from 989 distinct market participants, we find that the higher the proportion of dark or large-in-scale dark executions in the parent order, the lower its implementation shortfall. We also find that periodic batch auctions reduce implementation shortfall when dark pools are banned.

We also find that banning one venue type (dark pools) does not affect investor trading costs when similar alternatives exist. We provide evidence that investors reallocate trading flows in response to a ban on dark pool trading, and these reallocations do not fully reverse after the ban is lifted. We do this by examining the MiFID II DVC mechanism, introduced on 12 March 2018 with the aim of increasing pre-trade transparency by banning dark pool trading in individual stocks. Most UK stocks were subject to the ban. We also examine the lifting of the ban, and

[22] $Dark\ participant_j \times Characteristic_j$ equals 0 if the participant is not an active dark pool user and a member of the lowest quintile, i.e. uninformed or small.

find no impact on investor trading costs for either event. Yet, we do observe a substantial reversal towards dark pools after the lift, indicating that investors exhibit a preference for dark pools over periodic auctions. We also find that the dark pool ban or lift does not affect investors of varying size or informedness differently.

While previous research has examined the impact of dark pool and low transparency venues on measures of liquidity and measures of trading costs at the individual trade level (such as effective spreads), we examine their impact on a more complete measure of investor execution costs – implementation shortfall. Individual trade executions within a parent order are not independent, earlier executions can impact subsequent executions. This means the venue composition of the parent order matters for determining its overall cost. Consistent with this, we show that investors choose venues in a sequence of increasing transparency over the life of the parent order.

While we demonstrate the importance of pre-trade transparency in venue selection decisions, we are unable to examine other important factors that determine trading costs: the use of passive versus aggressive limit orders, execution algorithm design and broker skill. These are important directions for future research.

References

Anand, A., Samadi, M., Sokobin, J., and Venkataraman, K. (2019). Institutional order handling and broker-affiliated trading venues. *SSRN Electronic Journal*.

Battalio, R. H., Hatch, B. C., and Saglam, M. (2018). The cost of exposing large institutional orders to electronic liquidity providers. In *Sixth Annual Conference on Financial Market Regulation*.

Bloomfield, R., O'Hara, M., and Saar, G. (2015). Hidden liquidity: Some new light on dark trading. *The Journal of Finance*, 70(5):2227–2274.

Buti, S., Rindi, B., and Werner, I. M. (2016). Diving into dark pools. *Working Paper*, (2010-10).

Chakrabarty, B., Hendershott, T., Nawn, S., and Pascual, R. (2019). Order exposure in high frequency markets. *Available at SSRN 3074049*.

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

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

Conrad, J., Johnson, K. M., and Wahal, S. (2003). Institutional trading and alternative trading systems. *Journal of Financial Economics*, 70(1):99–134.

Degryse, H., de Jong, F., and van Kervel, V. (2015). The impact of dark trading and visible fragmentation on market quality*. *Review of Finance*, 19(4):1587–1622.

Farley, R., Kelley, E. K., and Puckett, A. (2018). Dark trading volume and market quality: A natural experiment. In *13th Annual Mid-Atlantic Research Conference in Finance (MARC) Paper*.

Foley, S. and Putnins, T. J. (2016). Should we be afraid of the dark? dark trading and market quality. *Journal of Financial Economics*, 122(3):456–481.

Garvey, R., Huang, T., and Wu, F. (2016). Why do traders choose dark markets? *Journal of Banking & Finance*, 68:12–28.

Gomber, P., Sagade, S., Theissen, E., Weber, M. C., and Westheide, C. (2016). Spoilt for choice: Order routing decisions in fragmented equity markets.

Gresse, C. (2017). Effects of lit and dark market fragmentation on liquidity. *Journal of Financial Markets*, 35:1–20.

Guagliano, C., Guillaumie, C., Reiche, P., Spolaore, A., and Zanon, A. (2020). Dvc mechanism: The impact on eu equity markets. *ESMA Working Paper*, (3).

Hendershott, T., Jones, C. M., and Menkveld, A. J. (2013). Implementation shortfall with transitory price effects. *Appeared in David Easley, Marcos López de Prado, and Maureen O'Hara*.

Johann, T., Putnins, T. J., Sagade, S., and Westheide, C. (2019). Quasi-dark trading: The effects of banning dark pools in a world of many alternatives. *SSRN Electronic Journal*.

Lee, C. M. C. and Ready, M. J. (1991). Inferring trade direction from intraday data. *The Journal of Finance*, 46(2):733–746.

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

Perold, A. F. (1988). The implementation shortfall: Paper versus reality. *Journal of Portfolio Management*, 14(3):4.

Sağlam, M., Moallemi, C. C., and Sotiropoulos, M. G. (2019). Short-term trading skill: An analysis of investor heterogeneity and execution quality. *Journal of Financial Markets*, 42:1–28.

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

6 Appendix

6.1 Appendix A: Changes in Participant Venue Choices

Table A1 shows that the two events related to DVC has affected venue choices of different market participants.

Table A1: Changes in Share of Trading by Venue Type and Trader Type - Around Ban and Lift Events

The table below shows the change in trader type participation on each trading venue between event windows. Participation is measured by the ratio of trader type turnover to total turnover per day. We show the results of a t-test to compare trader type participation between event windows, with clustered standard errors on day level, and report standard errors in parentheses. Event windows are pre-BAN (February 12th to March 9th), post-BAN (March 13th to April 12th); pre-LIFT (August 14th to September 9th) and post-LIFT (September 13th to October 11th). Values are shown in percentage points.

Metric	Auction	Dark	Dark (LIS)	Lit	Off-book	Periodic Auction	SI
Panel A. BAN (pre- vs. post-event)							
Banks	0.06*		0.01**	-0.08***	-0.13	0.02***	-0.06
	(0.03)		(0.01)	(0.02)	(0.38)	(0.00)	(0.11)
Broker-Dealer	0.43		0.05	0.30	2.71***	0.90***	-0.17
	(0.45)		(0.04)	(0.51)	(0.55)	(0.04)	(1.34)
Prop Trader - HFT	0.04		-0.03***	-0.89*	0.13***	0.08***	0.12*
	(0.05)		(0.01)	(0.48)	(0.04)	(0.01)	(0.07)
Institutional	-0.03		-0.03	0.17**	0.58***	0.16***	0.25
	(0.09)		(0.10)	(0.06)	(0.13)	(0.01)	(0.31)
Other	0.01		0.06	-0.03	0.47***	0.03***	0.06
	(0.03)		(0.06)	(0.02)	(0.15)	(0.01)	(0.13)
Panel B. LIFT (pre- vs. post-event)							
Banks	-0.01		-0.00	-0.03*	-0.21	-0.02***	-0.50**
	(0.01)		(0.00)	(0.02)	(0.14)	(0.00)	(0.24)
Broker-Dealer	-0.44		-0.08	-1.30***	-0.29	-0.96***	0.54
	(0.40)		(0.09)	(0.46)	(0.51)	(0.05)	(0.79)
Prop Trader - HFT	-0.05		0.01	-0.04	-0.31*	-0.09***	0.16*
	(0.06)		(0.02)	(0.46)	(0.16)	(0.01)	(0.08)
Institutional	-0.15		0.06	-0.13	0.07	-0.20***	0.21
	(0.13)		(0.11)	(0.10)	(0.09)	(0.02)	(0.27)
Other	0.02		0.11***	0.06**	-0.11	-0.04***	-0.23*
	(0.02)		(0.03)	(0.03)	(0.08)	(0.01)	(0.13)
Panel C. Pre-BAN vs. post-LIFT							
Banks	-0.01	-0.02**	0.00	-0.17***	-0.99***	0.01***	-0.08
	(0.01)	(0.01)	(0.00)	(0.02)	(0.29)	(0.00)	(0.13)
Broker-Dealer	0.80*	-0.61***	0.14*	-2.46***	1.74***	0.63***	0.23
	(0.42)	(0.12)	(0.08)	(0.40)	(0.44)	(0.03)	(0.70)
Prop Trader - HFT	0.03	-0.40***	0.00	0.40	0.05	0.17***	0.79***
	(0.04)	(0.03)	(0.01)	(0.31)	(0.05)	(0.01)	(0.07)
Institutional	-0.14	-0.17***	0.14	-0.78***	0.21***	0.07***	0.09
	(0.09)	(0.03)	(0.10)	(0.09)	(0.06)	(0.01)	(0.30)
	(0.04)	(0.02)	(0.04)	(0.06)	(0.03)	(0.01)	(0.06)
Other	0.02	-0.06***	0.08***	0.13***	0.02	0.02***	0.11
	(0.02)	(0.01)	(0.03)	(0.02)	(0.06)	(0.00)	(0.13)

Panel A shows that Broker-Dealer significantly increase trading in PA venues and off-book during the BAN period. A similar pattern can be seen for HFTs, however their economic increase is smaller compared to Broker-Dealers. Institutionals increase their share in lit venues and off-book. However, informed institutional investors show a significantly reduction in lit venues. Although Panel B shows a decrease for both Broker-Dealer and HFTs in PA venues, Panel C demonstrates that there is an overall trend towards these venues, as the positive coefficient when comparing post-LIFT to pre-BAN is strongly significant. Institutionals significantly reduce their trading activity from pre-BAN to post-LIFT in both dark and lit venues, but show a significant increase in both PA and off-book transactions.

6.2 Appendix B: Duplicate Transaction Reports

For a single trade MDP will typically contain at least two transaction report, one from each leg of the transaction. We remove redundant transaction reports as follows: For transactions taking place at the same venue the trading venue's transaction identification code is used to link buy and sell transaction reports. For transactions not sharing the same trading venue transaction identification code, we combine the transaction legs chronologically. Typically, transactions involve several intermediaries, such as central counterparties (CCPs) and brokers providing direct market access. To find the ultimate buyer and seller to a transaction we first eliminate all central counterparties and link both legs of the transaction. Second, direct market access brokers may report transactions with other (direct market access) brokers. In this scenario, instead of the CCP, the broker is the intermediary between the client and the other broker. Additionally, trades where no CCP is involved, may also be a double report by the involved brokers. In these cases, we identify the ultimate client by eliminating either double reports or removing the intermediary broker. Differentiating between on-market, off-market and OFF-book trades is used to mark possible parent orders. We distinguish OFF-BOOK, OTC and off-market trades. OFF-BOOK trades are bilateral agreements between two parties. If a trade is not specifically flagged (identified with a corresponding waiver as laid out in the Annex of RTS 22) as Over-The-Counter (OTC) or where a trade is executed according to the rules of the venue, we label it OFF-BOOK.[23] Thereby, according to the rules of the venue means that an OFF-BOOK trade will be identified when there is a single transaction report, without any other reported leg and no

[23] See Guidelines 5.16.1.3. We combine OFF-BOOK and OTC to off-book trading.

Table A2: Parent Order Venue Choice - Multinomial Logit Model

The table below shows the results from a multinomial logit model that shows the probability of a child being executed at a certain time in the life of the parent order on a specific venue. The dependent variable, or discrete choice, thereby indicates on which venue the child order is executed. Column (1) considers periods when dark trading is allowed (pre-BAN and post-LIFT) whereas column (2) considers periods when dark trading is prohibited (post-BAN and pre-LIFT). 'Depletion Bucket' indicates the time in the parent orders life in quintiles, e.g. the first bucket indicates child orders that are executed between 20% and 40% of the parent orders life cycle. The reference level for the model is the Lit market.

	Dependent variable:	
	Choice	
	(1)	(2)
Auction (intercept)	-6.094^{***}	-5.919^{***}
	(0.012)	(0.011)
Dark (intercept)	-2.052^{***}	
	(0.004)	
Dark (LIS) (intercept)	-7.003^{***}	-6.942^{***}
	(0.042)	(0.041)
Off-book (intercept)	-2.874^{***}	-2.585^{***}
	(0.005)	(0.004)
Periodic Auction (intercept)	-3.567^{***}	-2.936^{***}
	(0.007)	(0.005)
SI (intercept)	-4.064^{***}	-4.120^{***}
	(0.009)	(0.009)
Auction×Depletion Bucket	0.696^{***}	0.651^{***}
	(0.003)	(0.003)
Dark×Depletion Bucket	-0.028^{***}	
	(0.001)	
Dark (LIS)×Depletion Bucket	-0.079^{***}	-0.122^{***}
	(0.013)	(0.013)
Off-book×Depletion Bucket	0.014^{***}	0.005^{***}
	(0.002)	(0.001)
Periodic Auction×Depletion Bucket	0.008^{***}	-0.005^{***}
	(0.002)	(0.002)
SI×Depletion Bucket	-0.003	-0.027^{***}
	(0.003)	(0.003)
Period	pre-BAN and post-LIFT	post-BAN and pre-LIFT
Observations	4,851,067	4,883,551
R^2	0.011	0.013
Log Likelihood	−3,747,356.000	−2,897,123.000
LR Test	$79{,}991.410^{***}$ (df = 12)	$75{,}833.090^{***}$ (df = 10)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table A3: Parent Order Venue Choice - Marginal Effects
Marginal Effects of multinomial regressions presented in Table A2.

Lit	Auction	Dark	Dark (LIS)	Off-book	Periodic Auction	SI
Panel A. Periods of dark trading (pre-BAN and post-LIFT)						
-0.025	2.061	-0.109	-0.261	0.019	-0.002	-0.035
Panel B. Periods of no dark trading (post-BAN and pre-LIFT)						
-0.031	1.921		-0.398	-0.015	-0.045	-0.112

CCP involved, but reported with the venue market identifier code (MIC) of the trading venue, that is not XOFF.[24] Off-market trades are disregarded and not added to off-book trades; they are reports occurring in the XOFF venue that can not be classified in any category mentioned above.

Mapping parent and child orders

A broker can execute a trade on behalf of a client either as an agency trade or as a principal trade. In MDP agency trades, i.e. when a broker directly executes a client order on-market, the reported trading capacity is 'MTCH' (Matched Principal Trading Capacity) or 'AOTC' (Any Other Trading Capacity). We label both cases as direct agency trading (*DAT*). A principal trade occurs if a broker performs a transaction against its own books and the trading capacity 'DEAL' is reported. Yet, it is not necessary that the 'DEAL' capacity reflects a proprietary trade. There may occur a parent order in the 'DEAL' capacity in two cases:

i) principal trading: it is the sum of executed child orders of the same executing entity (i.e., either a dealer or prop HFT). In principal trading, we can identify the child orders as those coming from the same dealer over a specified trading horizon (e.g., the regular trading hours) and back out the (synthetic) parent order by aggregating the child orders using the trade direction. Thus, by construction we have 100% coverage ratio.

ii) indirect agency trading (*IAT*): it is the sum of executed child orders of the same executing entity (i.e., dealer) trading on behalf of a client (any Institutional Investor). In this case, we identify the parent order as any order that is recorded in the XOFF venue where the

[24]The venue MIC is a code used to identify trading venues.

Table A4: Impact of Event on Implementation Shortfall

The table below shows the impact of the event on implementation shortfall (IS) according to equation $IS_\tau = \alpha + (\beta_n \sum_{n=1}^{N-1} PctVenue_{n,\tau}) \times \delta Post_t + \gamma_1 Size_\tau + \gamma_2 Execution\,time_\tau + FE + \varepsilon_\tau$, where we add a time variable $Post_t$ which is 1 for the post event in each period, i.e. post-Ban or post-Lift. We include parent orders from both the pre-BAN and post-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the parent order size, *Size*, and execution time in hours, *Execution time*. We standardize the explanatory variables. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded. We only include stocks if they are traded on periodic auction venues (at least 1% turnover on periodic auction venues) and exclude trades from participants who trade less than 10 days in each period.

	Dependent variable:			
	IS_total_bps			
	Pre vs Post-Ban	Pre vs Post-Lift	Pre vs Post-Ban	Pre vs Post-Lift
Dark (%)	−0.069**	−0.055	−0.094	−0.151**
	(0.032)	(0.034)	(0.074)	(0.069)
Dark (LIS) (%)	−0.127**	−0.122**	−0.367***	−0.294***
	(0.063)	(0.052)	(0.109)	(0.103)
Periodic Auction (%)	0.105	−0.088***	−0.245	−0.135*
	(0.094)	(0.029)	(0.180)	(0.082)
Auction (%)	−0.056	−0.059	0.001	0.028
	(0.036)	(0.067)	(0.091)	(0.098)
SI (%)	0.106	0.058	−0.304	−0.108
	(0.090)	(0.053)	(0.186)	(0.150)
Off-book (%)	0.062	0.008	0.051	0.040
	(0.052)	(0.060)	(0.123)	(0.143)
Post				
	(0.000)	(0.000)	(0.000)	(0.000)
Size	4.098***	3.916***	5.834***	8.394***
	(0.580)	(0.682)	(1.799)	(2.107)
Execution time	−0.381	−1.065*	1.040	−2.113
	(0.548)	(0.598)	(1.546)	(2.026)
Dark (%)×Post				
	(0.000)	(0.000)	(0.000)	(0.000)
Dark (LIS) (%)×Post	−0.016	−0.049	0.130	0.064
	(0.076)	(0.060)	(0.121)	(0.170)
Periodic Auction (%)×Post	−0.214**	0.086	0.029	−0.078
	(0.098)	(0.059)	(0.196)	(0.128)
Auction (%)×Post	0.010	0.030	−0.072	−0.073
	(0.038)	(0.071)	(0.117)	(0.132)
SI (%)×Post	−0.109	−0.031	0.332	−0.136
	(0.077)	(0.070)	(0.241)	(0.342)
Off-book (%)×Post	−0.049	−0.024	−0.025	−0.001
	(0.044)	(0.060)	(0.149)	(0.193)
Index	FTSE100	FTSE100	FTSE250	FTSE250
Stock-Day FE	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes
Observations	17,287	16,190	10,126	7,876
R^2	0.252	0.262	0.506	0.524
Adjusted R^2	0.055	0.052	0.047	0.017
Note:				*p<0.1; **p<0.05; ***p<0.01

Table A5: Effect of Venue Trading Share on Implementation Shortfall - Non-Ban Period, Standardized

The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to regression specification 1. We include parent orders from both the pre-BAN and post-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the parent order size, *Size*, and execution time in hours, *Execution time*. We standardize the explanatory variables. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:							
	Total IS (bps)							
Lit (%)	1.461**							
	(0.719)							
Dark (%)		−2.003***						−2.305***
		(0.596)						(0.661)
Dark (LIS) (%)			−1.487***					−1.918***
			(0.502)					(0.573)
Periodic Auction (%)				−0.221				−0.377
				(0.455)				(0.461)
Auction (%)					0.648			−0.0002
					(0.692)			(0.707)
SI (%)						0.099		0.034
						(0.688)		(0.713)
Off-book (%)							1.330**	0.704
							(0.572)	(0.665)
Size	1.600*	1.413	1.811**	1.471*	1.446*	1.482*	1.473*	1.810*
	(0.881)	(0.875)	(0.903)	(0.864)	(0.876)	(0.866)	(0.865)	(0.924)
Execution time	0.125	0.056	0.101	0.177	0.137	0.183	0.175	−0.071
	(0.826)	(0.820)	(0.815)	(0.818)	(0.803)	(0.813)	(0.816)	(0.813)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	28,616	28,616	28,616	28,616	28,616	28,616	28,616	28,616
R^2	0.369	0.369	0.369	0.369	0.369	0.369	0.369	0.369
Adjusted R^2	0.103	0.104	0.103	0.103	0.103	0.103	0.103	0.104
Note:								*p<0.1; **p<0.05; ***p<0.01

Table A6: Effect of Venue Trading Share on Implementation Shortfall - Ban Period, Standardized
The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to equation 1. We include parent orders from both the post-BAN and pre-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the parent order size, *Size*, and execution time in hours, *Execution time*. We standardize the explanatory variables. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:						
	Total IS (bps)						
Lit (%)	2.711^{***}						
	(0.676)						
Dark (LIS) (%)		-1.050^{*}					-1.550^{***}
		(0.546)					(0.574)
Periodic Auction (%)			-1.956^{***}				-2.494^{***}
			(0.391)				(0.391)
Auction (%)				-0.623			-1.243^{*}
				(0.661)			(0.681)
SI (%)					0.330		-0.127
					(0.582)		(0.561)
Off-book (%)						0.115	-0.754
						(0.708)	(0.761)
Size	2.261^{**}	2.276^{**}	1.998^{**}	2.097^{**}	2.083^{**}	2.078^{**}	2.308^{**}
	(0.932)	(0.958)	(0.876)	(0.885)	(0.881)	(0.881)	(0.981)
Execution time	0.744	0.686	0.620	0.825	0.767	0.760	0.604
	(0.838)	(0.843)	(0.834)	(0.805)	(0.834)	(0.836)	(0.816)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	28,335	28,335	28,335	28,335	28,335	28,335	28,335
R^2	0.374	0.374	0.374	0.373	0.373	0.373	0.374
Adjusted R^2	0.107	0.106	0.106	0.106	0.105	0.105	0.107
Note:							*p<0.1; **p<0.05; ***p<0.01

Table A7: Effect of Venue Trading Share on Implementation Shortfall - Pre-Ban Period Only
The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to regression specification 1. We include parent orders from both the pre-BAN and post-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:							
	Total IS (bps)							
Lit (%)	0.039							
	(0.029)							
Dark (%)		-0.058^{*}						-0.078^{**}
		(0.032)						(0.035)
Dark (LIS) (%)			-0.186^{***}					-0.217^{***}
			(0.056)					(0.064)
Periodic Auction (%)				0.039				0.009
				(0.088)				(0.093)
Auction (%)					0.017			-0.008
					(0.039)			(0.040)
SI (%)						0.040		0.020
						(0.116)		(0.119)
Off-book (%)							0.085	0.061
							(0.063)	(0.065)
Size	4.030^{***}	3.857^{***}	4.409^{***}	3.854^{***}	3.813^{***}	3.845^{***}	3.873^{***}	4.555^{***}
	(0.790)	(0.761)	(0.785)	(0.755)	(0.772)	(0.754)	(0.751)	(0.815)
Execution time	0.026	0.023	-0.025	0.104	0.113	0.113	0.093	-0.177
	(0.761)	(0.748)	(0.749)	(0.755)	(0.752)	(0.752)	(0.751)	(0.749)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	15,440	15,440	15,440	15,440	15,440	15,440	15,440	15,440
R^2	0.387	0.387	0.387	0.387	0.387	0.387	0.387	0.388
Adjusted R^2	0.108	0.108	0.109	0.108	0.108	0.108	0.108	0.109
Note:								*p<0.1; **p<0.05; ***p<0.01

Table A8: Effect of Venue Trading Share on Implementation Shortfall - Post-Ban Period Only
The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to equation 1. We include parent orders from both the post-BAN and pre-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:						
	Total IS (bps)						
Lit (%)	0.100^{***}						
	(0.023)						
Dark (LIS) (%)		-0.120^{***}					-0.161^{***}
		(0.044)					(0.045)
Periodic Auction (%)			-0.128^{***}				-0.160^{***}
			(0.030)				(0.030)
Auction (%)				-0.041			-0.075^{**}
				(0.035)			(0.036)
SI (%)					0.026		-0.025
					(0.061)		(0.061)
Off-book (%)						0.033	-0.0004
						(0.049)	(0.050)
Size	4.135^{***}	4.158^{***}	3.604^{***}	3.717^{***}	3.661^{***}	3.661^{***}	4.377^{***}
	(0.741)	(0.740)	(0.698)	(0.728)	(0.713)	(0.709)	(0.758)
Execution time	0.798	0.742	0.857	0.998	0.965	0.957	0.597
	(0.669)	(0.679)	(0.681)	(0.686)	(0.691)	(0.692)	(0.664)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	15,646	15,646	15,646	15,646	15,646	15,646	15,646
R^2	0.396	0.395	0.395	0.394	0.394	0.394	0.396
Adjusted R^2	0.123	0.122	0.123	0.122	0.122	0.122	0.124
Note:							*p<0.1; **p<0.05; ***p<0.01

Table A9: Effect of Venue Trading Share on Implementation Shortfall - Pre-Lift Period Only
The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to equation 1. We include parent orders from both the post-BAN and pre-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:						
	Total IS (bps)						
Lit (%)	0.080^{***}						
	(0.030)						
Dark (LIS) (%)		-0.112^{**}					-0.156^{***}
		(0.055)					(0.058)
Periodic Auction (%)			-0.058^{**}				-0.091^{***}
			(0.023)				(0.026)
Auction (%)				-0.031			-0.059
				(0.044)			(0.046)
SI (%)					0.082		0.025
					(0.055)		(0.058)
Off-book (%)						-0.059	-0.101^{**}
						(0.048)	(0.050)
Size	4.422^{***}	4.481^{***}	3.891^{***}	4.067^{***}	4.018^{***}	3.988^{***}	4.749^{***}
	(0.945)	(1.032)	(0.990)	(0.945)	(0.970)	(0.976)	(0.997)
Execution time	-1.437	-1.455	-1.358	-1.290	-1.300	-1.296	-1.552
	(1.025)	(0.993)	(0.994)	(0.985)	(0.995)	(0.994)	(0.999)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,689	12,689	12,689	12,689	12,689	12,689	12,689
R^2	0.378	0.378	0.378	0.377	0.377	0.377	0.378
Adjusted R^2	0.094	0.094	0.094	0.093	0.093	0.093	0.095
Note:							*p<0.1; **p<0.05; ***p<0.01

Table A10: Effect of Venue Trading Share on Implementation Shortfall - Post-Lift Period Only
The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to regression specification 1. We include parent orders from both the pre-BAN and post-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:							
	Total IS (bps)							
Lit (%)	0.069**							
	(0.029)							
Dark (%)		−0.108***						−0.136***
		(0.026)						(0.032)
Dark (LIS) (%)			−0.167***					−0.223***
			(0.048)					(0.056)
Periodic Auction (%)				−0.002				−0.037
				(0.056)				(0.056)
Auction (%)					0.038			−0.009
					(0.043)			(0.046)
SI (%)						0.019		−0.006
						(0.077)		(0.077)
Off-book (%)							0.091*	0.041
							(0.047)	(0.052)
Size	3.694***	3.216***	4.030***	3.310***	3.203***	3.321***	3.281***	4.134***
	(0.900)	(0.871)	(0.884)	(0.837)	(0.882)	(0.850)	(0.851)	(0.962)
Execution time	−0.339	−0.292	−0.399	−0.183	−0.210	−0.176	−0.184	−0.611
	(0.826)	(0.825)	(0.796)	(0.813)	(0.805)	(0.803)	(0.808)	(0.811)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,176	13,176	13,176	13,176	13,176	13,176	13,176	13,176
R^2	0.378	0.379	0.378	0.378	0.378	0.378	0.378	0.380
Adjusted R^2	0.107	0.108	0.107	0.106	0.107	0.106	0.107	0.109
Note:								*p<0.1; **p<0.05; ***p<0.01

Table A11: Effect of Ban and Lift on Large Dark Participants Implementation Shortfall
The table below shows the Difference-in-Differences estimates including large and small participants estimates for three separate periods: pre-BAN (20 business days, 12 February to 9 March) to post-BAN (20 business days, 13 March to 12 April), and pre-LIFT (20 business days, 14 August to 9 September) to post-LIFT (20 business days, 13 September to 11 October) periods according to equation 4. Size is measured with the participants trading volume. Observations are participant mean IS constructed from at least 10 parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. Size equals 0 for lowest quintile and 1 for the highest quintile, the middle quintiles are disregarded. Within each quintile, participants are considered treated if they trade at or above the median value of dark trading across participants (are heavy users of dark venues and are thus impacted by the ban/lift), time post is one for the post-BAN and post-LIFT period. Standard errors are clustered by participant level.

	Dependent variable:		
	Total IS (bps)		
	BAN	LIFT	pre-BAN to post-LIFT
	(1)	(2)	(3)
Size × Post	1.278	3.730	−4.551
	(6.467)	(6.454)	(8.063)
Post × Dark participant	−4.407	4.067	−14.965
	(9.813)	(10.336)	(9.861)
Size × Post × Dark participant	10.189	1.372	15.446
	(10.453)	(10.974)	(10.903)
Day FE	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes
Observations	2,764	2,578	2,593
R^2	0.135	0.153	0.110
Adjusted R^2	0.076	0.097	0.049
Note:			*p<0.1; **p<0.05; ***p<0.01

dealer acts in deal capacity with an Institutional Investor. However, in this case, it is not certain that recorded transaction in the XOFF venue is a parent order, because it can be in principal also a prop trade or misreporting.

In a principal or agency trade a parent order is the sum of all executed child orders (at least two) of the same client via an intermediary (i.e., a broker) in the same direction over a regular trading day. The client order can originate from any market participant classified as Institutional Investor. In such cases the on-market child transactions are ultimately followed by the parent order between the ultimate client and the broker. We identify such parent orders as orders that are recorded in off-venue and where the dealer acts in the deal capacity with an Institutional Investor. To map child orders to parent orders we create a rolling sum of one-directional broker transactions, beginning with the Opening Auction, until the corresponding off-venue parent order is filled. The rolling sum is reset, once a parent order is filled. In the case of multiple clients per broker, it is not clear the mapping of child to parent orders is not necessarily unique, e.g. if the broker mixes market execution to fill orders simultaneously. We use information about the investment decision person, i.e. the trader, the desk trader or the ultimate beneficiary, to overcome this problem. This approach will still have an issue if there are several client orders per broker-trader ID combination, but will reduce the noise in the initial case, where only the broker information is used. Generally, the reporting of the client order execution in the XOFF venue contains the volume weighted average price (VWAP) the client is paying or receiving for its parent order. Comparing this 'true' order price to the VWAP of the constructed parent is a first quality constraint. We impose that the difference between the two VWAPs must not exceed 1 basis points (bps). However, when applying this narrow comparison criteria, we lose a significant amount of potential transactions from the *DAT*. Further analysis is loosening the assumption about the difference in basis points. Afterwards, we combine the *IAT* and *DAT* trades to a synthetic parent order with a maximum execution period of one day. We exclude any parent order with only one child trade and impose a trade directionally of 90%. Directionality is calculated by dividing the absolute difference between buying and selling volume by the sum of both sides. Identification of Market Participant Categories are done via the mapping of the MDP Legal Entity Identifier information to ORBIS and then using the fields 'Peer Group Description' and 'Specialization' from ORBIS to group market participants into aggregate categories.

6.3 Appendix C

Table C1: Participant Category Trading Share
The table below reports trading value in percentage of each trader type to overall trading around the first BAN (February 12th - April 12th) and second LIFT event (August 14th - October 11th). The time interval before and after the event in each period are 20 business days. We exclude the event days when the ban commences and the lift occurs for the first time (March 12th and September 12th), as well as the quadruple witching dates (March 16th and September 21st).

Trader Type	pre-BAN	post-BAN	pre-LIFT	post-LIFT
Broker-Dealer	61.27	61.82	63.24	63.35
Prop Trader - HFT	22.00	20.82	22.44	23.02
Institutional	10.27	11.14	9.31	8.89
Banks	3.00	2.43	1.66	1.41
Agg. Client Account	1.42	1.54	1.42	1.28
Other	1.12	1.21	1.10	1.30
Retail	0.91	1.04	0.84	0.74

Table C2: Proportion of Buying Counter Party
Proportion of trader type being on the buying side of the transaction on the Lit market in each observation period across all transactions.

Category	Pre-BAN	Post-BAN	Pre-LIFT	Post-LIFT
Prop Trader - HFT	49.57	48.73	49.29	49.68
Other	54.84	53.07	57.45	54.09
Banks	52.09	48.64	52.33	51.43
Broker-Dealer	50.24	50.52	50.68	50.16
Institutional	49.99	53.05	46.37	50.89

Table C3: Herfindahl-Hirschman-Index Around Ban and Lift Events
Market Concentration measured with Herfindahl-Hirschman-Index (HHI), where the index is the sum of squared market shares of each trading venue with SIs included and a second time without SIs.

Period	HHI	# trading venues	HHI without SI	# trading venues without SI
pre-BAN	17.79	101	29.92	61
post-BAN	19.69	100	32.55	61
pre-LIFT	18.63	103	31.39	62
post-LIFT	16.85	100	28.29	59

Table C4 shows that the average price impact measured on the ten second level varies across trader type. Prop Traders have the most price impact in the short term with a 10-sec price impact

of 5.42 bps, followed by Broker-Dealers (3.59 bps), Institutional (2.41 bps) and Banks (1.08 bps) in the BAN period. The price impact does not change much in the post-BAN period. In the pre-LIFT period Prop Traders and Broker-Dealers are still the top two categories based on short term price impact (5.12 bps and 3.14 bps, respectively). Institutional investors and Banks have a similar price impact of 2.03 bps and 1.90 bps. Again, the values remain similar in the post-LIFT period.[25]

Table C4: 10 Second Price Impact of Trader Type
The table below shows the 10-sec price impact (in bps) per trader type in each period. INTC stands for the Aggregate Client Account. Trades have been signed with the Lee and Ready (1991) algorithm and if the trade executes at the mid, and an institutional investor is either buyer or seller (not both), we classify this trade as either buyer or seller initiated, depending on the side the institutional investor trades.

Trader Type	coefficient	std. dev.	coefficient	std.dev.
	Pre-period		Post-period	
Panel A. BAN event				
Prop Trader - HFT	5.42	11.31	5.46	11.62
Broker-Dealer	3.59	14.23	3.58	16.29
Institutional	2.41	14.65	2.98	13.92
Other	1.42	12.60	1.03	10.19
Banks	1.08	9.94	1.68	31.83
Panel B. LIFT event				
Prop Trader - HFT	5.12	9.11	5.23	10.83
Broker-Dealer	3.14	11.8	2.59	11.51
Institutional	2.03	11.76	1.72	10.75
Other	0.67	7.74	1.54	20.38
Banks	1.90	11.08	1.56	10.06

[25]For alternative price impact intervals, please refer to Table C5, Table C6 and Table C7.

Table C5: 5 Second Price Impact of Trader Type
The table below shows the 5-sec price impact (in bps) per trader type in each period. INTC stands for the Aggregate Client Account. Trades have been signed with the Lee and Ready (1991) algorithm and if the trade executes at the mid, and an institutional investor is either buyer or seller (not both), we classify this trade as either buyer or seller initiated, depending on the side the institutional investor trades.

Trader Type	coefficient	std. dev.	coefficient	std. dev.
	Pre-period		Post-period	
Panel A. BAN event				
Prop Trader - HFT	5.32	9.59	5.39	9.53
Broker-Dealer	3.45	11.11	3.47	13.19
Institutional	2.31	12.14	2.89	13.67
Banks	1.23	8.16	1.53	18.44
Other	0.67	8.35	0.4	7.25
Panel B. LIFT event				
Prop Trader - HFT	5.04	7.82	5.13	9.48
Broker-Dealer	3.00	10.45	2.41	10.19
Institutional	1.91	10.5	1.62	9.86
Banks	1.71	9.66	1.68	8.48
Other	0.38	5.45	1	16.78

Table C6: 1 Minute Price Impact of Trader Type
The table below shows the 1-min price impact (in bps) per trader type in each period. INTC stands for the Aggregate Client Account. Trades have been signed with the Lee and Ready (1991) algorithm and if the trade executes at the mid, and an institutional investor is either buyer or seller (not both), we classify this trade as either buyer or seller initiated, depending on the side the institutional investor trades.

Trader Category	coefficient	std. dev.	coefficient	std. dev.
	Pre-period		Post-period	
a. BAN				
Prop Trader - HFT	5.3	19.96	5.25	19.74
Other	4.15	28.6	3.2	22.17
Broker-Dealer	3.72	24.19	3.79	24.85
Banks	2.88	19.64	1.12	46.57
Institutional	2.71	28.09	3.01	26.26
b. LIFT				
Prop Trader - HFT	4.97	14.94	5.19	17.74
Broker-Dealer	3.13	17.84	2.67	18.3
Institutional	2.48	18.68	2.06	18.01
Other	2.4	18.45	3	25.47
Banks	1.61	19.36	1.7	17.12

Table C7: 5 Minute Price Impact of Trader Type

The table below shows the 5-min price impact (in bps) per trader type in each period. INTC stands for the Aggregate Client Account. Trades have been signed with the Lee and Ready (1991) algorithm and if the trade executes at the mid, and an institutional investor is either buyer or seller (not both), we classify this trade as either buyer or seller initiated, depending on the side the institutional investor trades.

Trader Category	coefficient	std. dev.	coefficient	std. dev.
	Pre-period		Post-period	
a. BAN				
Prop Trader - HFT	5.25	37.68	4.99	37.92
Other	4.6	56.89	5.02	44.5
Banks	3.59	36.1	-0.93	69.54
Institutional	3.38	53.86	3.21	45.72
Broker-Dealer	2.76	53.79	4.09	43.05
b. LIFT				
Prop Trader - HFT	4.91	28.08	5.21	35.11
Other	3.77	35.73	5.09	52.49
Broker-Dealer	3.01	31.99	3.21	34.91
Institutional	2.47	32.97	3.32	36.05
Banks	1.08	34.43	1.51	33.31

Table C8: Effect of Venue Trading Share on Implementation Shortfall - Non-Ban Period FTSE100 Only

The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to regression specification 1. We include parent orders from both the pre-BAN and post-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:							
	Total IS (bps)							
Lit (%)	0.031							
	(0.020)							
Dark (%)		-0.063^{***}						-0.072^{***}
		(0.023)						(0.025)
Dark (LIS) (%)			-0.124^{***}					-0.147^{***}
			(0.040)					(0.044)
Periodic Auction (%)				0.029				0.015
				(0.042)				(0.043)
Auction (%)					0.004			-0.015
					(0.029)			(0.030)
SI (%)						0.058		0.054
						(0.055)		(0.059)
Off-book (%)							0.068^{**}	0.046
							(0.033)	(0.037)
Size	3.610^{***}	3.459^{***}	3.844^{***}	3.487^{***}	3.463^{***}	3.487^{***}	3.461^{***}	3.940^{***}
	(0.645)	(0.624)	(0.643)	(0.612)	(0.632)	(0.615)	(0.617)	(0.674)
Execution time	-0.425	-0.450	-0.470	-0.379	-0.386	-0.371	-0.381	-0.539
	(0.534)	(0.529)	(0.521)	(0.525)	(0.519)	(0.523)	(0.524)	(0.526)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	18,861	18,861	18,861	18,861	18,861	18,861	18,861	18,861
R^2	0.254	0.255	0.255	0.254	0.254	0.254	0.254	0.255
Adjusted R^2	0.050	0.050	0.050	0.050	0.049	0.050	0.050	0.051
Note:								*p<0.1; **p<0.05; ***p<0.01

Table C9: Effect of Venue Trading Share on Implementation Shortfall - Ban Period FTSE100 Only
The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to equation 1. We include parent orders from both the post-BAN and pre-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:						
	Total IS (bps)						
Lit (%)	0.091^{***}						
	(0.025)						
Dark (LIS) (%)		-0.114^{***}					-0.153^{***}
		(0.038)					(0.040)
Periodic Auction (%)			-0.075^{***}				-0.110^{***}
			(0.018)				(0.022)
Auction (%)				-0.058			-0.085^{**}
				(0.036)			(0.038)
SI (%)					0.067		0.019
					(0.050)		(0.047)
Off-book (%)						-0.010	-0.051
						(0.031)	(0.034)
Size	4.003^{***}	3.996^{***}	3.564^{***}	3.720^{***}	3.640^{***}	3.611^{***}	4.236^{***}
	(0.588)	(0.647)	(0.618)	(0.589)	(0.609)	(0.615)	(0.616)
Execution time	-0.702	-0.786	-0.747	-0.553	-0.645	-0.657	-0.792^{*}
	(0.491)	(0.491)	(0.490)	(0.452)	(0.484)	(0.486)	(0.471)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	18,358	18,358	18,358	18,358	18,358	18,358	18,358
R^2	0.274	0.273	0.273	0.273	0.272	0.272	0.274
Adjusted R^2	0.069	0.067	0.067	0.067	0.067	0.067	0.069
Note:							*p<0.1; **p<0.05; ***p<0.01

Table C10: Effect of Venue Trading Share on Implementation Shortfall - Non-Ban Period FTSE250 Only

The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to regression specification 1. We include parent orders from both the pre-BAN and post-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:							
	Total IS (bps)							
Lit (%)	0.131**							
	(0.052)							
Dark (%)		−0.109**						−0.168***
		(0.050)						(0.056)
Dark (LIS) (%)			−0.229***					−0.342***
			(0.079)					(0.089)
Periodic Auction (%)				−0.175*				−0.238**
				(0.097)				(0.095)
Auction (%)					0.072			−0.033
					(0.078)			(0.085)
SI (%)						−0.213		−0.241
						(0.207)		(0.204)
Off-book (%)							0.139	0.042
							(0.106)	(0.115)
Size	6.091***	4.894**	6.539***	4.937**	4.777**	4.939**	5.015**	7.018***
	(2.240)	(2.183)	(2.169)	(2.174)	(2.259)	(2.185)	(2.166)	(2.303)
Execution time	−0.188	0.371	0.155	0.753	0.800	0.671	0.685	−0.525
	(1.801)	(1.689)	(1.614)	(1.668)	(1.701)	(1.657)	(1.661)	(1.773)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,755	9,755	9,755	9,755	9,755	9,755	9,755	9,755
R^2	0.521	0.521	0.521	0.521	0.521	0.521	0.521	0.523
Adjusted R^2	0.045	0.044	0.045	0.044	0.044	0.044	0.044	0.047
Note:								*p<0.1; **p<0.05; ***p<0.01

Table C11: Effect of Venue Trading Share on Implementation Shortfall - Ban Period FTSE250 Only

The table below shows the impact of dark trading on implementation shortfall (IS) for each venue type individually and combined according to equation 1. We include parent orders from both the post-BAN and pre-LIFT period. In the combined specification (last column) the lit venue is excluded to prevent perfect multicollinearity. We include parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from trades in the 257 stocks that were subject to a suspension in the BAN period and 225 of those stocks which have their suspension lifted in the LIFT period. The variables of interest are the trading percentages of each parent order on a particular venue. Additional control variables are the (natural logarithm of the) parent order size, *Size*, and execution time in hours (also log), *Execution time*. The specifications include stock-day, participant and broker fixed effects. Standard errors are clustered by participant level. After-hours trading is excluded.

	Dependent variable:						
	Total IS (bps)						
Lit (%)	0.097*						
	(0.051)						
Dark (LIS) (%)		−0.154*					−0.197**
		(0.085)					(0.097)
Periodic Auction (%)			−0.095**				−0.125**
			(0.046)				(0.053)
Auction (%)				0.045			−0.013
				(0.061)			(0.068)
SI (%)					−0.040		−0.095
					(0.102)		(0.109)
Off-book (%)						0.002	−0.050
						(0.101)	(0.111)
Size	5.708***	6.145***	4.627***	4.687***	4.812***	4.830***	6.242***
	(1.725)	(1.788)	(1.558)	(1.620)	(1.560)	(1.561)	(1.938)
Execution time	1.119	1.173	1.736	1.920	1.792	1.783	0.941
	(1.799)	(1.798)	(1.756)	(1.775)	(1.770)	(1.769)	(1.817)
Stock-Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Broker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,977	9,977	9,977	9,977	9,977	9,977	9,977
R^2	0.501	0.501	0.501	0.501	0.500	0.500	0.501
Adjusted R^2	0.016	0.016	0.016	0.015	0.015	0.015	0.017
Note:							*p<0.1; **p<0.05; ***p<0.01

Table C12: Effect of Dark Pool Ban and Lift on Implementation Shortfall for FTSE100 Only
The table below shows the baseline Difference-in-Differences estimates for three separate periods: pre-BAN (20 business days, 12 February to 9 March) to post-BAN (20 business days, 13 March to 12 April), and pre-LIFT (20 business days, 14 August to 9 September) to post-LIFT (20 business days, 13 September to 11 October) periods according to equation 2, as well as pre- BAN to post-LIFT periods. Observations are participant mean IS constructed from at least 10 parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders must have at least a directionality of 90%. Parent orders are constructed from stocks that are a FTSE100 constituent and subject to the suspension and lifting. Participants are considered treated if they trade at or above the median value of dark trading across participants (are heavy users of dark venues and are thus impacted by the ban/lift), time post is one for the post-BAN and post-LIFT period. Standard errors are clustered by participant level.

	Dependent variable:		
	Total IS (bps)		
	BAN	LIFT	pre-BAN to post-LIFT
	(1)	(2)	(3)
Dark participant × Post	−0.421	2.082	−2.281
	(2.680)	(3.089)	(2.875)
Day FE	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes
Observations	4,609	4,183	4,429
R^2	0.121	0.093	0.093
Adjusted R^2	0.057	0.028	0.027
Note:			*p<0.1; **p<0.05; ***p<0.01

Table C13: Effect of Dark Pool Ban and Lift on Implementation Shortfall for FTSE250 Only
The table below shows the baseline difference-in-difference (diff-in-diff) estimates for three separate periods: pre-BAN (20 business days, February 12th to March 9th) to post-BAN (20 business days, March 13th to April 12th), and pre-LIFT (20 business days, August 14th to September 9th) to post-LIFT (20 business days, September 13th to October 11th) periods according to equation 2, as well as pre- BAN to post-LIFT periods. Observations are participant mean IS constructed from at least 10 parent orders that have a value of at least 100,000 GBP, last ten minutes or longer, and consist of five or more children. The VWAP of parent orders is not to deviate more than 1bps from 'true' parent order VWAP and parent orders have to have at least a directionality of 90%. Parent orders are constructed from stocks that are a FTSE250 constituent and subject to the suspension and lifting. Participants are considered treated if they trade at or above the median value of dark trading across participants (are heavy users of dark venues and are thus impacted by the ban/lift), time post is one for the post-BAN and post-LIFT period. Standard errors are clustered by participant level.

	Dependent variable:		
	Total IS (bps)		
	BAN	LIFT	pre-BAN to post-LIFT
	(1)	(2)	(3)
Dark participant × Post	1.623	9.336	0.080
	(4.847)	(6.133)	(5.898)
Day FE	Yes	Yes	Yes
Participant FE	Yes	Yes	Yes
Observations	2,961	2,484	2,580
R^2	0.112	0.097	0.097
Adjusted R^2	0.044	0.027	0.029
Note:			*p<0.1; **p<0.05; ***p<0.01

Table C14: Mean Comparison Around Ban and Lift Events for Institutional Traders
The table below shows the results for a test of means pre and post: pre (20 business days, February 12th to March 9th) to post (20 business days, March 13th to April 12th) BAN and from pre (20 business days, August 14th to September 9th) to post (20 business days, September 13th to October 11th) LIFT period accordingly. Time variable is 1 for after event period (i.e. post DVC or post LIFT). Stocks are included if they are labeled as liquid and are suspect to the suspension.

	Dependent variable:		
	BAN	LIFT	pre-BAN to post-LIFT
a. Effective Spread			
	(1)	(2)	(3)
Pre to post	−0.302	−1.514	0.159
	(0.647)	(2.488)	(0.891)
Participant FE	Yes	Yes	Yes
Observations	5,373	4,805	5,148
R^2	0.517	0.543	0.494
Adjusted R^2	0.491	0.519	0.466
b. Realized Spread			
	(1)	(2)	(3)
Pre to post	−0.265	−1.496	0.286
	(0.644)	(2.503)	(0.903)
Participant FE	Yes	Yes	Yes
Observations	5,373	4,805	5,148
R^2	0.516	0.544	0.494
Adjusted R^2	0.491	0.520	0.466
Note:			*p<0.1; **p<0.05; ***p<0.01

Table C15: Mean Comparison Around Ban and Lift Events for High Frequency Traders
The table below shows the results for a test of means pre and post: pre (20 business days, February 12th to March 9th) to post (20 business days, March 13th to April 12th) BAN and from pre (20 business days, August 14th to September 9th) to post (20 business days, September 13th to October 11th) LIFT period accordingly. Time variable is 1 for after event period (i.e. post DVC or post LIFT). Stocks are included if they are labeled as liquid and are suspect to the suspension.

	Dependent variable:		
	BAN	LIFT	pre-BAN to post-LIFT
a. Effective Spread			
	(1)	(2)	(3)
Pre to post	−1.133	−5.683	−1.320
	(2.397)	(4.020)	(3.252)
Participant FE	Yes	Yes	Yes
Observations	7,769	7,135	7,532
R^2	0.152	0.121	0.143
Adjusted R^2	0.149	0.118	0.140
b. Realized Spread			
	(1)	(2)	(3)
Pre to post	−1.227	−5.728	−1.439
	(2.413)	(4.054)	(3.252)
Participant FE	Yes	Yes	Yes
Observations	7,769	7,135	7,532
R^2	0.159	0.125	0.151
Adjusted R^2	0.156	0.122	0.148
Note:			*p<0.1; **p<0.05; ***p<0.01

Table C16: Mean Comparison Around Ban and Lift Events for Broker-Dealer
The table below shows the results for a test of means pre and post: pre (20 business days, February 12th to March 9th) to post (20 business days, March 13th to April 12th) BAN and from pre (20 business days, August 14th to September 9th) to post (20 business days, September 13th to October 11th) LIFT period accordingly. Time variable is 1 for after event period (i.e. post DVC or post LIFT). Stocks are included if they are labeled as liquid and are suspect to the suspension.

	Dependent variable:		
	BAN	LIFT	pre-BAN to post-LIFT
a. Effective Spread			
	(1)	(2)	(3)
Pre to post	−1.221	−0.546	−0.398
	(0.971)	(1.592)	(1.809)
Participant FE	Yes	Yes	Yes
Observations	17,456	15,618	16,617
R^2	0.321	0.271	0.267
Adjusted R^2	0.318	0.267	0.264
b. Realized Spread			
	(1)	(2)	(3)
Pre to post	−1.015	−0.666	−0.111
	(0.975)	(1.608)	(1.815)
Participant FE	Yes	Yes	Yes
Observations	17,456	15,618	16,617
R^2	0.333	0.287	0.282
Adjusted R^2	0.330	0.284	0.278
Note:			*p<0.1; **p<0.05; ***p<0.01

Table C17: Venue Shares of Informed Participants with Alternative Propensity Score Matching - Around Ban and Lift Events

Average usage of venues between informed group and matched group during periods of dark trading and periods of prohibited dark trading. Comparison of informed investors to matched control group. Investors are informed if the *beta_1* coefficient from equation 3 is positive and significant at the 10% level during the pre-BAN and post-LIFT period. Afterwards we match informed participants to a control sample based on trade size with a propensity score matching using a nearest neighbor algorithm (logit). Column 'Difference' shows the results of a regular t-test between the two groups. Column 'Difference (Fixed Effects)' shows the results of a regression of the form $venue(\%) = FE_{stock-day} + informed.dummy + \varepsilon$. Standard errors are clustered by stock-day. We use the same sample of 29 treated participants for both comparisons. Four participants from the control group are not active during the post-BAN and pre-LIFT period and we find new matched participants based on a PSM performed during the post-BAN and pre-LIFT period. The PSM is using 'Total Parent Order Size', 'Average Parent Order Size', 'Number of Parent Orders' and 'Average Number of Traded Stocks'. Standard Errors in brackets.

	Informed Share (%)	Matched Share (%)	Difference	Difference (Fixed Effects)
Panel A. Period when dark trading is allowed (pre-BAN and post-LIFT)				
Number of Participants	29	29		
Total Parent Order Size (mln GBP)	4,852.96	3,729.05		
Average Parent Order Size (mln GBP)	1.12	1.26		
Number of Parent Orders	412.11	440.84		
Average Number of Traded Stocks	100.70	124.63		
Auction	10.88	15.50	-4.62***	-4.27***
			(0.51)	(1.40)
Dark	16.84	17.17	-0.33	0.14
			(0.63)	(1.61)
Dark (LIS)	4.22	2.88	1.34***	2.35**
			(0.37)	(0.92)
Lit	57.03	47.89	9.15***	8.77***
			(0.84)	(2.20)
Off-book	5.09	9.79	-4.70***	-6.44***
			(0.54)	(1.48)
Periodic Auction	3.31	4.36	-1.05***	0.25
			(0.26)	(0.67)
SI	2.49	2.40	0.09	-0.82
			(0.24)	(0.53)
Panel B. Period when dark trading is prohibited (post-BAN and pre-LIFT)				
Number of Participants	29	29		
Total Parent Order Size (mln GBP)	3,727.42	4,018.99		
Average Parent Order Size (mln GBP)	1.02	1.10		
Number of Parent Orders	406.74	604.95		
Average Number of Traded Stocks	92.29	145.87		
Auction	14.46	17.30	-2.84***	-1.84
			(0.59)	(1.61)
Dark (LIS)	5.46	3.16	2.30***	1.97*
			(0.41)	(1.03)
Lit	59.94	53.60	6.35***	3.23
			(0.88)	(2.30)
Off-book	7.02	11.26	-4.24***	-3.24**
			(0.62)	(1.58)
Periodic Auction	10.74	11.70	-0.96*	0.45
			(0.55)	(1.28)
SI	2.25	2.71	-0.46*	-0.38
			(0.27)	(0.61)