



Competing for dark trades

Paul J. Irvine^{a,*}, Egle Karmaziene^b

^a TCU Neeley School of Business, Fort Worth, Texas, USA

^b Vrije Universiteit Amsterdam and Tinbergen Institute, the Netherlands

ARTICLE INFO

JEL:
G12
G14
G18
D47

Keywords:
MiFID II
dark pool trading
competition

ABSTRACT

We use recent European restrictions to evaluate how traders substitute across available dark pools. Our findings suggest that restricting dark trading at the most prominent platform has a detrimental effect on dark trading activity. Annual dark trading in a restricted stock decreases by more than 50 % over the six-month restriction period. Consistent with investors' sticky relationships with specific dark pools, our results suggest that substitution across dark pools is remarkably low. Despite the availability of alternative dark pools, traders are unwilling to trade elsewhere. Our study provides evidence that dark trading is not a market of exchanges, but rather a collection of independent silos. This fact has implications for the vulnerability of dark trading to the introduction of an HFT into the pool, and sharpens our understanding of how the pecking order theory of trading actually functions.

1. Introduction

In January 2021, trading outside public stock exchanges reached 47.2 % of total U.S. equity trading volume.^{1,2} A portion of this off-exchange trading occurs in dark pools, private exchanges where trades can be executed without providing pre-trade bid and ask prices to the market. As dark trading grew in Europe, policymakers became concerned that the transparency and quality of price formation on the lit exchanges could deteriorate. In reaction to this growth, European regulators introduced two new rules, known jointly as the double volume cap (DVC), that capped the amount of a stock's trading volume that could be executed in a single dark pool (4 %), and all dark pools (8 %). Violation of the 4 % rule banned trading in that stock on the contra-vening platform for six months. This sudden restriction in the number of dark platforms available to traders provides us with a clean setting to study where traders send their trades when a popular dark platform is

closed. Do they keep their trades in the dark pool market, sending them to alternative dark platforms? Examining this question allows us to learn about how traders use the dark pool market.

Traders' reaction to the dark pool ban depends on the ease of substitution across dark trading platforms. In general, buy-side investors favor dark to lit trading venues due to their lower information leakage and trading costs (Menkveld et al., 2017) and therefore may choose another dark pool.³ If substitution is costless, then inability to trade on one pool should not affect the overall level of dark trading. Alternatively, the diversity of dark pools could isolate traders into particular dark pools and limit the amount of substitution. We examine how the shock to the supply of potential trading venues affects trading in the dark pool market. The 4 % platform-specific ban provides us with a clean setting to investigate the degree of substitution across dark pools.

To facilitate the privacy needs of institutions, large trades are exempt from this regulation if their trade value exceeds a certain stock-specific

Karmaziene gratefully acknowledges Nasdaq Nordic Foundation for financial support. We are grateful for comments received from Sabrina Buti, participants at the AFA annual meeting (2024), the Baltic Economic Conference (2021), the Eastern Finance Association meeting (2024), the German Economic Association annual meeting (2023), seminar participants at the Bank of Lithuania, Dutch Financial Markets Authority, and VU Amsterdam.

* Corresponding author.

E-mail addresses: p.irvine@tcu.edu (P.J. Irvine), E.Karmaziene@vu.nl (E. Karmaziene).

¹ Rosenblatt Securities report:

<https://www.rblt.com/news/gamestop-mania-highlights-shift-to-dark-trading>.

² This fraction underestimates the importance of off-exchange trading since not all submitted dark orders are executed. In fact, the execution probability in these pools is lower than 4.13% for NYSE, NASDAQ, and ITG Europe stocks Gresse (2006); Ye (2010).

³ Public stock exchanges, which compete for orders with dark venues, highlight the negative features of dark pools, such as harming price discovery (Callahan from NYSE Euronext) or not contributing to price formation and market transparency (Greifeld, NASDAQ).

<https://doi.org/10.1016/j.jbankfin.2025.107509>

Received 22 May 2023; Accepted 22 June 2025

Available online 28 June 2025

0378-4266/© 2025 Elsevier B.V. All rights reserved, including those for text and data mining, AI training, and similar technologies.

value.⁴ However, most European dark pools have small average trade sizes (Petrescu and Wedow, 2017) and such trades require a waiver from normal lit pre-trade transparency to be allowed to transact in a dark pool. It is this ‘trading under the waiver’ that the regulatory authority intends to limit, preventing widespread movement of liquidity away from lit markets to dark pools. Johann et al. (2019) argue that near substitutes for dark pools exist such as traditional auctions or systematic internalizing systems. They find these alternatives capture some of the activity formerly executed in the dark during a total ban on dark trading in the stock. Unlike the 8 % all-platform ban analyzed in Johann et al. (2019), the 4 % single-platform ban still allows dark trading on other platforms. What happens to the trading structure in the dark pools that are still able to trade? By analyzing the degree of substitution to the remaining active dark pools, we can gain insight into the unexamined question of how institutions use different dark pools to trade.

There are several facts that could affect the degree of substitution across dark pools after the 4 % ban goes into effect. First, large trades can still be transacted and do not require a waiver. Second, the suspension is narrow, affecting the platform’s trading in a single asset. Because the suspension is so limited, traders are unlikely to break their relationship with the trading platform. Traders can continue trading other assets on the same platform and return fully once the suspension is lifted. Third, since some exchanges, such as Nordic, specialize in dark pool trading their listed stocks, their institutional clients are likely to maintain their relationship to access their execution expertise on these platforms. Fourth, it is not expensive for investors to switch trading platforms. Although dark pool membership fees can reach up to £10,000, institutional traders inform us that larger brokers maintain access to all platforms so that these are available to their clients.⁵ Since the effective date of the DVC in March 2018, dark trading was banned at multiple platforms in multiple stocks. Our empirical analysis uses a sample of 141 platform-asset-level suspensions on 17 different platforms over the 2018–2022 period. This restriction allows us to evaluate a cross-section of stocks that differ in their primary listing markets, trading prevalence, and liquidity.

We start by evaluating how dark trading responds to platform-level suspensions of a particular stock. Over the six-month restriction period, total annual dark trading in the restricted stock decreases by over 50 %, a high level of reliance on the leading platform given alternative dark trading platforms are available.⁶ This finding is consistent with traders establishing relationships with a subset of dark trading platforms. When the leading dark platform is restricted, alternative platforms are not used to a significant degree. The availability of alternative platforms is not the reason for the low level of substitution. Every stock in our sample can be traded at multiple alternative dark venues. In fact, in response to venue-level restrictions at the leading platform, the probability of dark-restricted stock becoming available on a new platform increases by 12.4 %. Thus, the absence of alternative platforms does not explain the failure of trades to migrate away from the leading dark pool.

The most likely cause of the post-ban volume reduction is that any orders migrating to alternative platforms are unlikely to find a match

and execute. The volume on these alternative platforms *must* be less than 4 % by regulation, and is typically much lower. Following the suspension, orders may migrate to multiple venues however, due to the dispersion of trading across multiple alternative platforms, the probability of matching a buyer with a seller is remote, lowering the liquidity value of these alternatives (Pagano, 1989)

The lack of substitution across dark pools is an important result for market structure. Effectively, these regulations restricted the development of the nascent dark pool market, in preference to systematic internalizers such as Jane Street. Under the DVC regulations dark pool growth has stagnated, and execution volumes by internalizers increased. We find that the paucity of substitution is robust over time, across countries, and stocks. The results are consistent before and after Britain exits our sample, for the subsamples of venues supervised by German, British and other European authorities, for more and less liquid stocks, and for stocks with overall lower and higher dark trading demand.⁷ The DVC regulations effectively restricted the development of a small, but growing, network of dark platforms in Europe.

The absence of substitution across dark pools is an important result for institutional traders in light of Brugler and Comerton-Forde (2021) observation that the introduction of a high-frequency trader (HFT) into an existing dark pool increases the information leakage of trades. If substitution across pools is economically difficult, and our evidence suggests that it is, then this finding can explain why traders are, to an extent, ‘sitting ducks’ for the introduction of an HFT into an existing dark pool.

Our finding that traders treat dark pools like silos, and rarely substitute across platforms also has implications for the interpretation of Menkveld et al.’s (2017) ‘pecking order theory’. Menkveld et al. (2017) find that when shocks hit the market, trading volume tends to substitute away from the low-cost, low-immediacy dark markets, toward higher cost, higher immediacy dark and lit alternatives. With negligible substitution across dark markets, this theory should not be interpreted as individual traders ‘sorting’ dark pools. Instead, consistent with Menkveld et al. (2017) arguments, the shock changes the relative cost of low-immediacy (dark) markets, and traders respond by tilting volume towards more immediate execution venues.

We find that the responses to a 4 % platform ban are quite homogeneous across the remaining open platforms. Mechanically due to the ban, trading drops at the platform that was the most active in dark-trading the particular stock. In general, dark trading in the restricted security does not migrate to the other venues. We show that the platform-level restriction has an equally strong effect on stocks with more and less developed dark trading markets. Substitution across dark platforms is equally unlikely when dark traders already have several active alternative platforms.

However, when prior to the ban there were only two active dark pools, 7 % of the decreased dark trading in the restricted stock moves to the second-largest platform. This substitution is a rational response from traders. The ability for a buyer to meet a seller only exists at the single remaining active platform. This finding highlights the importance of dark pool liquidity for our results. When traders are more certain about the identity of the platform attracting the most trades, they submit their orders to this platform, because the probability of two parties meeting is higher at the remaining active platform. The results of this test are important for market structure. If the opportunity costs of failure to execute are high, then new platforms face a hidden cost of entry. Specifically, without the volume to reduce this opportunity cost, they are likely to have difficulty attracting the scale necessary to become a profitable platform.

We contribute to research examining substitution among trading platforms. The past literature shows that the liquidity in a market

⁴ ESMA publishes monthly values of trading at a month-stock-platform and month-stock levels.

⁵ Connectivity fees vary widely. Liquidnet’s £10,000 fee is the highest we found, but connectivity fees range down to £500 at Aquis. A variable charge based on amount executed is also charged by platforms. Smaller brokers who might not maintain access to all platforms will add a platform at the client’s request for a nominal fee.

⁶ Under the 8% total dark trading ban there can be only one platform that reaches the 4% platform trading ban at a particular time. Mathematically, other active platforms must trade less than 4%, and in practice they trade considerably less. Hence, we define the 4% violating platform to be the leading platform.

⁷ These latter results are untabulated, but available on request from the authors.

increases with competition among dealers (Degryse et al. (2015); Foerster and Karolyi (1998); Werner and Kleidon (1996)), and when traders move from over-the-counter (OTC) to lit markets, but liquidity decreases if they move to trading with systemic internalizers (Comerton-Forde et al. (2018); Gomber et al. (2018)). Other studies find that this migration does not significantly affect the cost of trading, price efficiency (Farley et al. (2018)), or increase informational efficiency (Brogaard and Pan (2019)). Our results are consistent with Bekaert et al. (2011), suggesting that stock market development affects its segmentation. In this case, the nascent development level of the European dark pool market appears to inhibit the use of alternative dark pools when the leading dark pool is restricted from trading.

We also add to studies examining the effects of MiFID II and the Markets in Financial Instruments Regulation (MiFIR) on stock markets. One branch of past research examines the quality of information production following the MiFID II requirement to unbundle research costs from transactions. Even though the traders are relatively inelastic to costs of information (Di Maggio et al., 2021), the new requirement decreased information quantity but increased its quality (Fang et al. (2020); Guo and Mota (2021); Lang et al. (2021)). Another branch of studies examines responses to a market-wide EU-level dark trading suspension. This research concludes that only a tiny part of dark trading volume migrates to lit markets Johann et al. (2019), AMAFI (2019). Suspension at a single platform level allows us to measure how traders allocate their dark trades across the other platforms, how this allocation evolves over time, what the effect is on market concentration, and whether alternative platforms competitively respond to these changes. We contribute to the literature by showing that in expectation of more trading transparency, the bans on dark trading decrease trading concentration only through the decrease in dark volume on the most active platforms. With these results, we also add to broader interdisciplinary literature on platform competition by responding to the (Rietveld and Schilling, 2021) invitation to examine platform competition in areas outside the technology industry in the United States.

2. Hypotheses

In this section, we define our predictions. We discuss the demand for dark trading and the European dark pool market structure in subSections 2.1. and 2.2, respectively, and formally introduce our hypotheses in subSection 2.3.

2.1. The demand for dark trading

The levels of dark to total trading are substantial. Before MiFID II, they were increasing over time in the U.S. (Menkveld et al., 2017) and in Europe (Comerton-Forde, 2017). Dark trading in Europe rose from only 1 % of total trading volume in 2008 to 8 % of total trading volume in 2016 (Petrescu and Wedow, 2017). The 8 % rule, a complete ban on dark trading 'under the waiver' in a particular stock, decreases its dark volume, but the levels increase again after the ban is lifted (Guagliano et al., 2020). The result is in line with the enduring demand for dark trading.

2.2. A multiple platform market

In the EU, multiple platforms engage in dark trading, as they do in the United States. As in the U.S. the same stock can be traded on several different dark platforms. The nature of demand and supply for asset trading and, especially dark pool trading, benefits from oligopolistic design (Petrescu and Wedow, 2017). The benefits of competition across platforms are balanced by the difficulties of matching buyers and sellers in fragmented markets with low order volume.⁸ The European dark pool

⁸ For example, due to the inability to attract sufficient trading volumes, the NASDAQ OMX dark pool NEURO was forced to close in 2010.

market developed with many platforms competing for trader's orders, every stock in our sample has at least one alternative active platform, and many have several alternatives. Traders generally ignore these alternatives and concentrate their trading with the largest dark pool. The interesting twist is that the largest dark pool in a particular stock is usually a completely different dark pool than the largest dark pool for a second stock. Fig. 1 presents a histogram plotting the frequency of stocks that trade the most at each particular dark pool every month. The histogram presents the number of stock-months different platforms do the most trading in our sample of 4 % ban stocks. While some platforms are more active than others, there is no single platform that dominates trading in all the banned stocks.

If all platforms traded all assets, the costs were equal across platforms, and there were no likely trading disruptions, it would be optimal for traders to maintain a relationship with a single platform: a conclusion established by Pagano (1989) and Mendelson (1987). In practice, even though the trading prices are calculated similarly across dark pools, the subset of assets that can be traded in dark pools differs across platforms.⁹ Also, platforms provide different services for order placement and processing (Petrescu and Wedow, 2017). This differentiation in platform characteristics allows platforms to co-exist. Therefore, unlike lit markets, trading in dark pools is more fragmented.

Compared to public exchanges, dark pools are better able to differentiate their services. Dark pools differentiate their services along four lines - size of the orders, diversity of instruments traded, mechanisms used to match orders (e.g., choosing scheduled, non-continuous, matching may prevent traders from crossing with predatory algorithms), and order features (basic or complex, such as minimum order size to be filled to avoid information leakage). The diversity of dark pools may affect traders' self-selection and stickiness to a particular pool.

2.3. Hypotheses

We assume that due to substantial dark pool trading features, membership fees, per-order fees, time costs related to maintaining relationships with multiple platforms and the limited number of stocks that most platforms trade, investors choose to trade using only a subset of available platforms. If market participants do not treat platforms as perfect substitutes, then restricting dark trading of the stock on a particular platform decreases aggregate dark trading. Being unable to dark trade at the platform of their choice, agents could decrease their dark trading activity. In particular, we test the following hypothesis:

Hypothesis 1. Total dark trading volume of a restricted stock decreases when trading in that stock is banned at one platform.

The number of platforms that have permission to trade a particular stock in the dark varies within Europe. Some stocks can be traded in the dark on one platform, other stocks can be traded on forty different platforms. Some securities never dark trade. Using the differences in the number of active platforms as a proxy for market development, we examine whether the ban on dark trading is more severe when dark pool

⁹ (Petrescu and Wedow, 2017) explain that the clients experience financial costs of dark pool trading related to several factors. First, it relates to the mid-point execution price, which is commonly calculated using bid and ask prices on the lit exchange (but may be based on more factors like BlockMatch or Société Generale Alpha Y platforms). Second, clients' costs depend on whether and how much the venue compensates the liquidity providers. Third, some venues charge membership fees up to £10,000 per month (potentially providing other services for the fee), and all venues charge per-order fees. The per-order fees range from 0.1 to 1 basis point depending on the venue, the size of the order and monthly executed volume. Fourth, the execution price is related to the probability and speed of execution, which depend on pool liquidity. The unmatched dark pool orders are canceled at the close of the trading day in Europe.




Fig. 1. Frequency distribution of largest volume dark pools.

The graph presents the frequency distribution of the largest volume dark pools for the stock months in the 4 % suspension sample.

markets are less developed. When one platform dominates trading in a specific stock, traders are less likely to maintain multiple relationships with other platforms and are not used to switching across platforms. However, when the dark pool market is more widely developed for a particular stock, dark trading volume is more likely to migrate to another platform since the alternative platforms are also active. We examine whether:

Hypothesis 2. Total dark trading volume of a restricted stock decreases by more in stocks with fewer active platforms.

This paper also investigates two other questions relating to the single dark platform ban. Academic literature suggests several ways that lit and dark trading platforms can simultaneously coexist. Informed investors self-select to lit platforms and contribute to price formation by trading, while uninformed investors move to trading in dark platforms. Liquidity traders are attracted there by lower transaction costs. Despite lower immediacy, certain investors prefer dark to lit trading due to lower costs, and a lower risk of information leakage. Since lit and dark trading are used for different trades and by different traders, we test whether the limited single-stock single platform restriction has any effect on the overall level of trading in the dark pool network.

3. Institutional setting

We exploit a quasi-natural experiment – a new regulation in the European Economic Area MiFID II. Among its other requirements, it introduces the most restrictive regulation in history (Comerton-Forde et al., 2018), aiming to restrict the use of dark pools. Unlike the previous controls, the restriction is applied based on trading volume, rather than trading price, as in Canada or Australia.

From January 2018 on, the regulation sets a DVC on dark trading. Following Johann et al. (2019), we define dark pool trading as trading without pre-trade transparency. Trading without pre-trade transparency is generally forbidden unless the regulatory authority grants a waiver. For this reason, dark pool trading in Europe is often referred to as ‘trading under the waiver’. These waivers allow market operators not to disclose pre-trade information when: 1) the prices are based on a lit-market reference price, 2) the prices are uncertain due to illiquidity

or other reasons, 3) the prices lie within the bid-ask spread, 4) the orders are large, 5) or orders held in an order management facility of the trading platform that are pending disclosure. The specific part of the rule we examine is that if more than 4 % of annual stock trading volume is transacted at a single dark pool, then trading on this pool in the particular stock is suspended for six months. The rule has been binding since January 2018, and the first suspension took place in March 2018.

The 4 % single-asset volume cap on platform-based dark pool trading appears as a warning or a break on dark pool activity. If the total asset volume on all dark pool platforms reaches 8 % then trading under the waiver is restricted in that stock across all dark platforms. This latter restriction, the complete ban, is the rule studied by Johann et al. (2019) and Guagliano et al. (2020). Some exceptions to the trading ban are still allowed. In particular large transactions can still be permitted since they are allowed under a separate large trade waiver that is not subject to the double volume cap rules.¹⁰

What happens to the stock’s dark pool trading after the violation of the 4 % platform limit? In this instance, the asset cannot be traded ‘under the waiver’ on the violating dark pool for six months. As opposed to the complete shutdown of dark pool activity studied by Johann et al. (2019), the 4 % rule still allows dark trading activity in the asset on other available dark platforms.

Petrescu and Wedow (2017) show that the European dark pool market is continent-wide and not country specific. Although some dark pools specialize in trading stocks from a certain region (e.g., Nordic), the largest pools facilitate trading in a broad range of assets. As a result, the dark pool market is fragmented across the continent. Most European dark pools are not limited to transacting large block trades but process orders of any size. Such market arrangement is similar to the U.S., where the largest trades constitute less than 3 % of total trading in the five most active dark pools.

There is a regulatory decision taken prior to January 2018 that turns

¹⁰ The large transactions threshold varies by asset, depending upon an annual average turnover measure. However, practitioners suggest that these thresholds are too high and may reduce the liquidity for many less-traded stocks. <https://www.thetradenews.com/> updated-MiFID-rules-slash-large-in-scale-thresholds.

out to have a significant influence on our empirical results. Comerton-Forde (2017) discusses that MiFID-II dark pool regulations, prior to the introduction of the double-volume cap, prohibited bank-owned dark pools from continuing to operate under the MiFID-II regulatory regime. Many of these bank-owned dark pools converted to systematic internalizers through a loophole in the regulations.¹¹ In the US dark pool market bank-owned dark pools generally have higher dark volumes than exchange-owned or independent dark pools. Appendix Table B1 reports that in July 2021, for example, 7 of the top 10 dark volume markets, are bank-owned pools. Two are independent dark pools, and only one is exchange owned (CBOE). Banks have a stream client of order flow constantly coming into their brokerage operations. Some of this order flow makes its way into their dark pools. This resting order flow increases the chance of finding a counterparty in the bank pool, and is the likely reason bank-owned dark pools tend to be high-volume dark pools. In the EU, these high-volume dark pools exist only as SI's, lowering the possibility that migrating orders will be executed in alternative independent and exchange-owned dark pools, and increasing the likelihood that traders will migrate to SI's to execute there.

4. Data and descriptive statistics

This section details the sources of the data used in our study (SubSection 4.1) and discusses summary statistics (SubSection 4.2)

4.1. Data

Following Johann et al. (2019), we obtain data on trading under the waiver and suspensions from the European Securities and Markets Authorities (ESMA) reports. The monthly reports detail aggregate trading under the waiver over the last 12 months at a month-platform-stock level for the period from December 2017 to June 2022.¹² In particular, for each security, they reveal the aggregate volume transacted over the last 12-months, the fraction of this volume traded in the dark in total and in each platform. Moreover, ESMA reports detail the start and end dates of dark trading suspensions and their level of severity (at platform or EEA level).

The suspensions file lists all dark trading suspensions at the EEA. Suspensions begin on March 12, 2018. We apply several data filters to make our sample more homogeneous. We start with 790 platform-level suspensions that were never revoked. In line with Johann et al. (2019) and based on stock's RCA, we keep stocks listed in 11 European markets with a largest number of stocks, leaving us with 486 suspensions. The DVC regulation applies for liquid stocks only. To avoid illiquid stocks, for which this regulation does not apply, we eliminate stocks with total-EEA dark trading larger than 9 % or venue-level dark trading larger than 6 % of total EEA-level trading over the last 12 months in any month since March 2018. The high dark volume in these stocks clearly indicating that the DVC does not apply.¹³ We keep 141 platform-level suspensions. From ESMA website, we also obtain bi-annual stock trading volume at systematic internalizers.

¹¹ The Financial Times 'Brussels moves to tighten dark pool trading rules', June 21, 2017.

¹² Due to the restricted data provided, our ability to enhance robustness through alternative volume benchmarks is constrained. Each monthly change we report is a comparison of month t , to month $t-12$.

¹³ ESMA Liquidity thresholds can vary over time (Appendix A). To ensure all sample stocks are eligible to be banned at all times during the sample period, we remove stocks that may have once been subject to a 4% ban, but did not stay an RCA liquid stock throughout the sample period. We do this by applying the above restrictions that clearly indicate the stock did not meet the current liquidity threshold.

4.2. Descriptive statistics

Over our sample period of 54 months, under our sample restrictions 122 unique stocks were banned from dark trading at a single platform at some point during our sample period. Ninety-nine platforms from 11 EEA countries provide 110,876 asset-platform-month level reports on dark trading in these stocks. However, due to Brexit, after March 2021, we do not observe securities that are admitted to trading or traded on British trading platforms. After Brexit, the number of reporting countries and platforms decreases from 11 to 10 and from 88 to 60, respectively.

Table 1 shows that the mean dark trading volume in sample stocks constitutes 2.1 % of total trading volume. On average, an asset can be dark-traded on 19 venues but is actually traded on only 2.4. The statistic is in line with limited competition in dark-pool trading. Panel B in Table 1 demonstrates that the average dark platform trades just 0.12 % of annual stock trading volume. However, the dark volume is heavily skewed – only 13 % of available venues have positive values.

In this paper, we investigate variation in dark trading volume for stocks that experienced venue-level dark trading suspensions. Over our sample period 18 % of our stock-months fall during a six-month venue-level suspension, and 1 % occur during an EU-level complete suspension period.

The mean level of trading under the waiver is well below ESMA's 4 % trading ban limit. Yet, our sample still contains 141 security-specific platform-stock level suspensions, spread over 17 different platforms. The suspensions are quite concentrated in the most active platforms, which cumulatively transacted 82 % of total trading under the waiver at the EEA level during our sample period. We are confident that our analysis does not miss any important market participants. The largest three platforms experienced over 71 % of sample suspensions – London Stock Exchange (35 %) in London, Investment Technology Group Limited (21 %) in Dublin, and CBOE Europe in London (14 %). Given that the restriction applies when platform volume in a particular asset is considered to be too concentrated, it is natural that the bulk of the restrictions occur at high-volume platforms.

Next, we investigate the dark trading reports that these 17 platforms submit during the 2017–2022 period. When investigating ESMA reports, we see that platforms have different characteristics (Table 2). They may choose to assist with dark trading in a large number of listed stocks. For example, Liquidnet Europe Equities (LIQU) was able to dark-trade 27,847 securities over our sample period. In contrast, other platforms trade more selectively and specialize in particular securities. Nasdaq Helsinki (XHEL) was approved to dark-trade only 205 stocks. However, having the longest menu of stocks does not necessarily lead to the highest dark trading levels. LIQU dark-traded 10.3 % of their eligible stocks, which is a low level compared to 76.6 % at XHEL. Dark trading at

Table 1
Summary statistics.

	count	mean	sd	p25	p50	p75
DarkVol (%)	5726	2.10	Panel A 2.15	0.00	1.46	3.76
CountVenues	5726	19.05	9.91	11.00	19.00	26.00
CountNon0	5726	2.39	2.54	0.00	1.00	4.00
			Panel B			
DarkVen (%)	85,892	0.12	0.58	0.00	0.00	0.00
Ban	85,892	0.18	0.39	0.00	0.00	0.00
BanEU	85,892	0.01	0.11	0.00	0.00	0.00

This table reports descriptive statistics of the venues trading under the waiver. Panel A describes data at month-stock level. Panel B describes data at month-stock-venue level. *DarkVol* and *DarkVen* represent the average dark trading volume at EEA level and at a reporting venue level over the last 12 months. The measures are standardized by dividing by total trading volume at the EEA level and presented in %. *CountVenues* and *CountNon0* count the number of venues reporting and reporting non zero dark trading levels over the last 12 months. *Ban* and *BanEU* are indicators for stock-venue level suspensions due to the breach of DVC 4 % and 8 % rules, respectively.

Table 2
Sample venues and their characteristics.

		Observations	Stocks	Traded in dark	Dark
AIMX	London Stock Exchange AIM MTF	46,021	1168	4.20	0.02
BCXE	CBOE Europe London	230,905	6249	61.48	0.68
CEUD	CBOE Europe Amsterdam	109,021	4207	48.78	0.26
LIQU	Liquidnet Europe Limited	1024,104	27,847	10.28	0.02
MERK	Euronext Oslo	3396	170	13.53	0.37
MTAA	Euronext Milan	14,747	367	54.77	0.15
TQEM	Turquoise Europe	46,788	2637	67.96	0.42
TRQM	Turquoise Plato	190,913	4778	67.50	0.54
XCSE	Nasdaq Copenhagen	34,798	818	29.58	0.44
XHEL	Nasdaq Helsinki	8273	205	76.59	0.44
XLOM	London Stock Exchange Non-AIM MTF	304,182	7379	14.65	0.01
XLON	London Stock Exchange	182,769	8588	23.72	0.21
XMLI	Euronext Access Paris	13,234	400	0.75	0.01
XOSL	Oslo Bors	11,822	305	76.39	0.24
XPOS	POSIT Dark	599,968	17,356	27.10	0.66
XSTO	Nasdaq Stockholm	37,151	1642	40.68	0.35
XUBS	Aquis Exchange PLC	159,104	3868	78.26	0.45

This table lists summary statistics for 17 venues at which trading was restricted in our sample. *Observations* and *Stocks* count the number of stock-month reports and stocks for which these venues provide over our sample period. *Traded in dark* is a fraction of venue's stocks ever traded in dark, in %. *Dark* is the mean ratio of dark trading volume at the particular exchange to total dark trading volume at the EEA level, in %.

XHEL was also more active than at LIQU; it constituted 0.44 % and 0.02 % of total annual trading in the reported stocks, respectively. These patterns suggest that investors see some platforms as experts in dark trading particular stocks. Fig. 2 corroborates these ideas – when dark trading is restricted at the largest platform, the total dark trading in the restricted stock decreases as the remaining platforms do not absorb the restricted dark trading.

Fig. 3 graphs the evolution of dark volume at the suspended platform. It is apparent that much of the overall increase in trading in the dark reported in Fig. 2 comes from the banned venue alone. From month –12 until the ban there is a gradual accumulation of trading volume at the banned platform that eventually broaches the 4 % level. Thus, even though dark volumes are increasing in the stock, this trading is not broadly developed, but rather concentrated in the banned platform as dark traders appear to be continuously attracted to the platform until this accumulation triggers a ban.

The lack of prior relationships may explain the result. In particular, most of the banned stocks could have been dark traded on anywhere from 5 to 30 platforms (Fig. 4). So, restricting trading under the waiver on a single platform should not have eliminated so much trading volume if substitution occurs. Assuming the platforms were close substitutes, traders could have traded on any of the remaining 4 to 29 platforms. Yet, the platforms were not perfect substitutes. Fig. 5 shows that in half of the cases, prior to the ban, the entire dark trading volume was transacted at a single venue. Problems with risk management at the platform level were unlikely to trigger the ban. Untabulated results show that there is no sudden change in trading volume around the ban period. The values of total trading, dark trading, and trading at the main venue do not increase significantly.

5. Methodology

5.1. Total stock trading after the single-platform ban

Platform-level suspension directly affects only one platform, while

the other active platforms are allowed to trade without restriction. Such exclusive targeting of particular platforms provides a good environment to examine how actual investors perceive the degree of substitutability between alternative dark pools. In the first analysis we examine how investors react to single-platform trading restrictions using the following regression specification:

$$\Delta \text{DarkVol}_{s,t} = \sigma_s s + \theta_t + \beta \text{Ban}_{s,t} + \gamma \Delta X_{s,t} + \epsilon_{s,t} \quad (1)$$

where $\Delta \text{DarkVol}_{s,t}$ is a measure of dark trading in month t stock s , it is calculated as the ratio of dark trading in the restricted stock to total trading (in percent). Since this measure is provided only over a 12-month period, we use the first difference of the values to measure the monthly effects. $\text{Ban}_{s,t}$ is a dummy variable taking values of one in months t when trading under the waiver in the stock s is restricted at any platform. In some specifications, we add a control for the changing demand in dark trades, the first difference in logarithmic value of trading volume at the EEA level. We also include an indicator for the asset-specific 8 % EU-level restriction on dark pool trading.

We assume that when trading is restricted on one dark platform, investors choose between four options (i) executing a trade on a lit market, (ii) executing on a lit crossing alternative as Johann et al. (2019) report, (iii) executing at a systematic internalizer such as Jane Street, (iv) executing a trade on an alternative dark pool. Investors may be more likely to keep trades in particular stocks in a dark pool, such as stocks with high information asymmetry, or stocks that are more likely to reveal a trading strategy. To mitigate omitted variable bias arising for stock-specific reasons, we augment the regression with stock fixed effects (σ_s) using the securities ISIN numbers. We also add month fixed effects to accommodate time trends, such as changes in the number of trading platforms (θ_t). Unobserved characteristics may cause a particular and autocorrelated growth in trading under the waiver. To mitigate this issue, we cluster standard errors using the Newey and West (1987) technique with five lags.

5.2. Dark pool trading migration after the ban

Next, we investigate how the dark trading flows migrate across platforms in response to the ban at the largest platform. In particular, we estimate flows across platforms using the following regression equation:

$$\Delta \text{DarkVenue}_{s,v,t} = \sigma_s + \theta_t + \eta_v + \beta_i \sum_{v=0}^4 \text{Ban}_{s,v,t} + \gamma X_{s,v,t} + \epsilon_{s,v,t} \quad (2)$$

where $\text{DarkVenue}_{s,v,t}$ is a measure of trading under the waiver at platform v in a year to month t in security s (in percent). Δ is the first difference transformation. $\text{Ban}_{s,v,t}$ is a dummy variable taking values of one in months when trading under the waiver in asset s is not allowed in month t in each of four active trading platforms v . In this test, platforms are ranked from the most active (first) to the least active (last) and are related to dummies Ban_1 to Ban_4 .¹⁴ These controls include a ratio of total dark trading volume at EEA level to trading volume, a log-number of platforms reporting their dark trades and an indicator for the asset-specific EU-level restriction on dark pool trading. We limit concerns regarding the security-level informational environment, time trends, and platform-level characteristics by adding stock, time, platform fixed effects and using first difference transformations of the variables.

6. Results

In this section, we discuss how restricting dark trading on one exchange affects the total trading (subSection 6.1) and competition in

¹⁴ In the few cases where there are more than four active alternatives, these platforms are consolidated into Ban_4 .




Fig. 2. Evolution of dark trading around the venue-level suspension. The graph represents the monthly average annual total dark trading for stocks 12 months prior and post the ban venue-level suspension. The line indicates mean dark pool trading levels. The vertical lines refer to the beginning and the end dates of the restrictions.

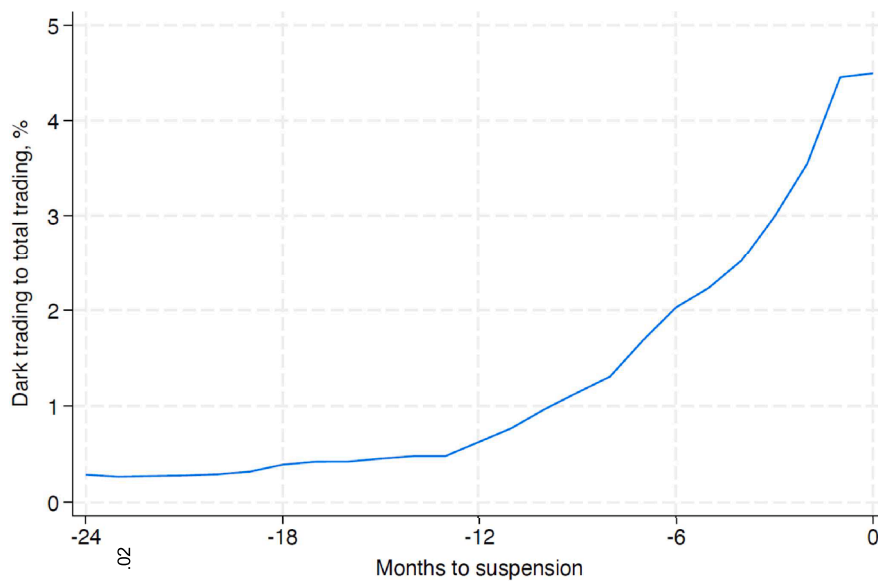


Fig. 3. Trading at the banned platform. This Figure presents the monthly average annual dark trading, as a percentage of total trading, at the banned platform only. Dark volume starting 24 months prior to the platform-level suspension is presented.

trading (subSection 6.2) between platforms.

6.1. Trading levels

We test Hypothesis 1 by examining what happens to total dark trading volume when trading at one platform is restricted. The

regulators restrict dark pool trading when such trading at the most active platform reaches 4 % of total annual trading. The first five columns in Table 3 present the Eq. (1) regression results.

To ensure that the total trading does not drive our estimates, we control for the time-fixed effects and for time-varying and stock-varying trading turnover. Depending on the set of controls or fixed effects

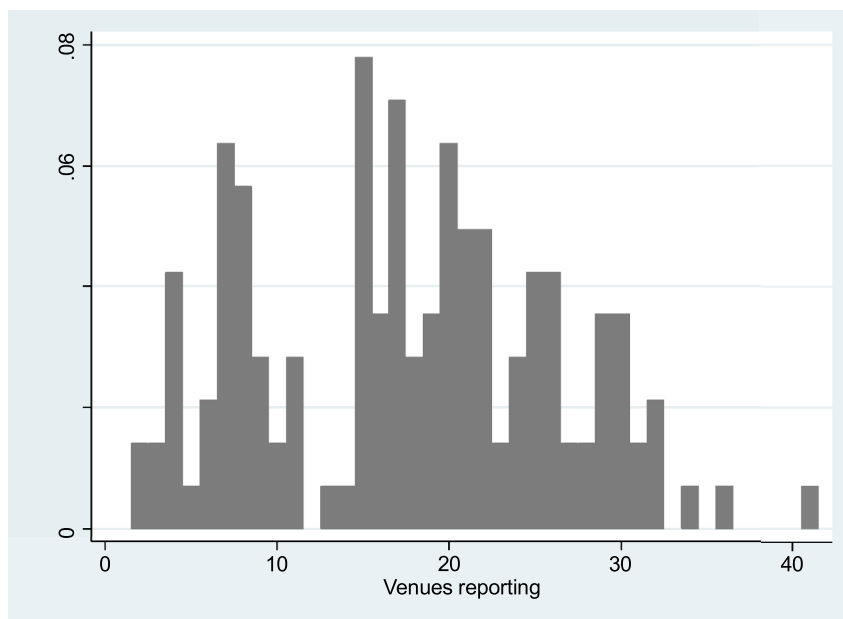


Fig. 4. Venues that can dark-trade the suspended stock. The histogram represents the distribution of the number of venues that could dark-trade the suspended stock before the suspension.

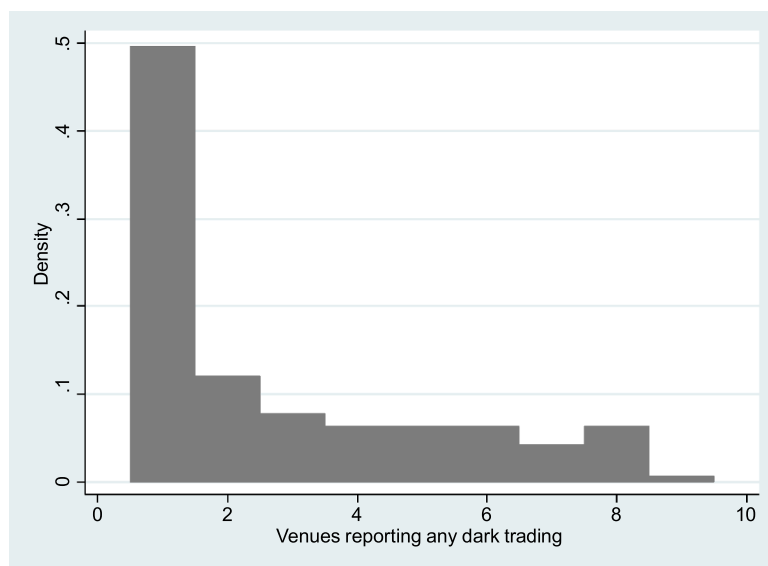


Fig. 5. Venues dark-trading the suspended stock. The histogram represents the distribution of the number of dark pools that have dark-traded the suspended stock in the year prior to the suspension.

employed the fraction of dark trading to total trading falls between 0.43 and 0.45 percentage points per month (Columns 1–4). The effect is significant at the 1 % level. Adding a control for the change in total demand for dark trading in Column (5), does not materially affect the *Ban* coefficient.

Naturally the ban restricts dark trading at the restricted platform, but also markedly decreases overall dark trading at the EEA level. The result is consistent with the limited level of substitution across dark platforms. When the leading platform is restricted, dark traders do not migrate to alternative venues, suggesting considerable barriers to platform substitution in the dark pool market.

At the time of the restriction, the median dark trading in the restricted stocks at all dark platforms was 5.23 % of total trading. This measure fell to 2.79 % at the end of the ban, suggesting a 2.44 percentage point raw decrease in dark trading during the six-month ban

period. Controlling for the time trends and stock characteristics, we estimate that throughout the ban, dark trading decreased by between 2.6 to 2.7 percent of total trading. Column (6) examines the evolution of trading throughout the 6-month ban period.¹⁵ In this regression, we replace the single dummy variable indicating a ban with six variables taking values of one for each month of the ban. The decline does not differ much across the remaining months of the ban. These findings indicate that time does not materially increase dark pool migration. The combination of the six coefficients suggests a 2.77 percentage point reduction in dark pool trading over the six-month ban period. The

¹⁵ The start of the ban always falls between the 10th and the 20th day of the month. This dating explains the lower coefficient on the first month of the ban (1.*BanMonth*).

Table 3
Regression results.

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ DarkVol	Δ DarkVol	Δ DarkVol	Δ DarkVol	Δ DarkVol	Δ DarkVol
Ban	-0.440*** (0.026)	-0.427*** (0.027)	-0.437*** (0.027)	-0.454*** (0.028)	-0.454*** (0.028)	
BanEU		-0.462*** (0.084)	-0.466*** (0.090)	-0.533*** (0.093)	-0.533*** (0.093)	-0.548*** (0.094)
Δ logTrade					0.005 (0.023)	0.005 (0.022)
1.BanMonth						-0.225*** (0.051)
2.BanMonth						-0.541*** (0.081)
3.BanMonth						-0.523*** (0.068)
4.BanMonth						-0.502*** (0.072)
5.BanMonth						-0.567*** (0.097)
6.BanMonth						-0.413*** (0.058)
Observations	5726	5726	5726	5726	5726	5726
R-squared	0.073	0.079	0.096	0.122	0.122	0.122
Time FE	No	No	No	Yes	Yes	Yes
Stock FE	No	No	Yes	Yes	Yes	Yes

DarkVol is the ratio of trading under the waiver to total trading over the last 12 months, in percent. Δ is the first difference transformation. *Ban* (*BanEU*) is an indicator variable taking values of one for the months when the stock was restricted from trading at a single venue (all venues) and zero otherwise. *1.BanMonth*–*6.BanMonth* indicate the first to sixth months of the ban. *logTrade* is the logarithmic value in EUR of trading a particular issue on European exchanges over the last 12 months. Regressions 3–6 include asset-fixed effects, regressions 4–6 include time-fixed effects. The table does not report for the coefficients of the constant. Regressions are ordinary least squares, the standard errors are Newey-West with five lags. All observations are at a month-stock level.

results in this table support Hypothesis 1 and show that dark platforms are not close substitutes, and a trading ban at a single platform leads to a significant total decrease in dark trading. Given the aim of the Double Volume Cap (DVC) regulation is to ensure that most volume stays in the lit markets, it appears to accomplish this aim if restricting single-stock trading drives a considerable fraction of dark trading in the stock away from dark pools.

Next, we investigate whether the ban-related drop in dark pool

Table 4
Economic effects.

	(1)	(2)	(3)	(4)
	Δ Dark	Δ Dark	Δ Lit	Δ SI
Ban	-23.964*** (8.101)	-18.028** (7.625)	2.002 (1.430)	2.077* (1.061)
BanEU	-23.144 (14.318)	-20.093 (12.659)	-1.563 (1.194)	3.480 (3.261)
Δ logTrade	6.780** (3.430)	6.727** (3.412)	3.557*** (1.242)	0.275* (0.151)
L.Dark		-0.042* (0.025)		
L.Lit			-0.024 (0.019)	
L.SI				-0.065*** (0.019)
Observations	4040	4040	4040	4040
R-squared	0.069	0.090	0.099	0.077
Time FE	Yes	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes	Yes

Dark (*Lit*) is the value of the average monthly dark (lit) trading volume over the last 12 months, in million EUR. *SI* is the value of the monthly trading volume using systematic internalizers, in million EUR. Δ is the first difference transformation. *L* indicates a lagged value. *Ban* (*BanEU*) is an indicator variable taking values of one for the months when the stock was restricted from trading at a single venue (all venues) and zero otherwise. *logTrade* is the logarithmic value in EUR of trading a particular issue on European exchanges over the last 12 months. The regressions include time- and asset-fixed effects. Regressions are ordinary least squares, the standard errors are Newey-West with five lags. All observations are at a month-stock level.

trading was able to increase lit trading as the regulator intended. In [Table 4](#) we show that an asset, banned from dark pool trading on a single dark platform, experienced a monthly dark trading reduction of at least 18 million EUR ([Table 4](#), Column 1–2). Concurrently, trading using systematic internalizers (SI) increased by 2.1 million EUR (Column 4). These findings are consistent with 12 % of the decrease in dark trading volume migrating to the SI's. [Johann et al. \(2019\)](#) show that the complete 8 % EEA ban on dark trading also leads to increased trading volumes at SIs.

There exists a large literature on opaque trading. The paper most relevant to our study of trade migration is [Degryse et al. \(2021\)](#). They study the relationship between dark pool trading and hidden order submission strategies using Dutch order book data from the pre-MiFID II period. They report that lit-venue hidden orders and dark pool trading are substitutes. Given their results, it is possible that hidden order submission should increase on the stock's lit trading venue. We do not have limit order data to identify hidden orders, we can only observe this substitution pattern when the hidden orders execute on the lit venue. We report the magnitude of migration to lit venues, perhaps through hidden order submission. This magnitude is similar to the SI migration level, but the effect is not statistically significant ([Table 4](#), Column 3). We conclude that regulators' initial aim to attract trading volume from dark pools to the lit markets using a 4 % rule was not successful, even if hidden orders are employed as an alternative opaque trading strategy.

We next examine Hypothesis 2, testing whether the degree of development of the dark market in a particular stock affects the impact of the ban. Using heterogeneity across stocks in the level of the development of competition, we examine whether the ban has a more substantial effect on stocks with less developed dark trading environments. In particular, we assume that securities, which are commonly dark traded on at least two platforms, have sufficiently developed markets. The idea is simple, does the existence of other active dark platforms affect the degree of substitution?

[Table 5](#) presents regressions investigating this issue. Column 1 evaluates the incremental impact of having a more developed existing dark pool market on the degree of substitution. The ban has a negative

Table 5
Dark trading restrictions and competition.

	(1)	(2)	(3)	(4)	(5)
	Δ DarkVol	Δ DarkVol	New	Δ Venues	Δ HHI
Ban	-0.459*** (0.049)	-0.460*** (0.049)	0.124* (0.067)	0.000 (0.000)	-0.044*** (0.014)
Ban_Developed	-0.014 (0.059)	0.008 (0.060)			
Ban_Venues	0.400** (0.176)	0.401** (0.175)			
Developed	0.005 (0.023)	0.005 (0.023)	0.085** (0.034)	0.000 (0.000)	0.011 (0.010)
Δ logTrade			0.137 (0.228)	0.000 (0.000)	0.025 (0.032)
BanEU		-0.534*** (0.093)			
Observations	5726	5726	5639	5726	5726
R-squared	0.115	0.122	0.307	0.375	0.019
Time FE	Yes	Yes	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes	Yes	Yes

DarkVol is the ratio of trading under the waiver across all venues to total trading over the last 12 months, in percent. *New* is an indicator variable taking the value of one if the number of venues, which report dark trading in a particular stock (zero or non-zero), increased over the previous month. *HHI* is a Herfindahl - Hirschman Index (HHI) of market concentration in dark trading the restricted stock. Δ is the first difference transformation. *Ban* (*BanEU*) is an indicator variable taking values of one for the months, when the stock was restricted from trading at a single venue (all EEA venues) and zero otherwise. *Developed* takes the value of one if a stock was dark-traded at least two venues before the ban, and zero otherwise. *logTrade* is the logarithmic value in EUR of trading a particular issue on European exchanges over the last 12 months. All regressions include time- and security-fixed effects. OLS regressions (in Columns 1–2, 4–5) have Newey and West (1987) errors with five lags. Regression in Column 3 is a probit regression. The table does not report for the coefficients of the constant. All observations are at a month-stock level.

Table 6
Migration of dark pool trading across the venues.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Δ DarkVenue	Δ DarkVenue	Δ DarkVenue	Δ DarkVenue	Δ DarkVenue	Δ DarkVenue	Δ DarkVenue
1.BanV	-0.367*** (0.005)	-0.368*** (0.005)	-0.464*** (0.027)	-0.440*** (0.047)	-0.485*** (0.032)	-0.540*** (0.034)	-0.387*** (0.039)
2.BanV	-0.011 (0.007)	-0.012* (0.007)	-0.015 (0.016)	-0.025 (0.022)	-0.001 (0.024)	0.036** (0.017)	-0.043** (0.021)
3.BanV	-0.010 (0.008)	-0.011 (0.008)	-0.008 (0.011)	-0.007 (0.020)	-0.005 (0.010)		-0.011 (0.013)
4.BanV	0.001 (0.005)	0.001 (0.005)	0.008 (0.006)	0.001 (0.006)	0.019* (0.010)		0.005 (0.006)
0.BanV	-0.001 (0.002)	-0.002 (0.002)	0.004*** (0.001)	0.006*** (0.002)	0.003* (0.001)	0.005*** (0.002)	0.004** (0.001)
BanEU	-0.021*** (0.005)	-0.025*** (0.005)	-0.025*** (0.005)	-0.029*** (0.010)	-0.021*** (0.007)	-0.029*** (0.008)	-0.027*** (0.006)
L.DarkVolume	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.001)	-0.002*** (0.001)	-0.001 (0.001)	-0.002** (0.001)	-0.002** (0.001)
logVenues	-0.015*** (0.002)	-0.007** (0.003)	-0.007 (0.005)	-0.009 (0.007)	-0.004 (0.008)	-0.006 (0.008)	-0.003 (0.007)
Observations	85,892	85,892	85,503	32,994	52,407	51,178	33,834
R-squared	0.071	0.073	0.217	0.138	0.281	0.254	0.233
Stock FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Venue FE	No	No	Yes	Yes	Yes	Yes	Yes
Venue-time FE Market Venue	No	No	Yes	Yes	Yes	Yes Active2	Yes Activemore2
				Large	Small		

DarkVenue (*DarkVol*) is the ratio of trading under the waiver per venue (total) to total trading over the last 12 months, in percent. Δ is the first difference transformation. *iBanV* is an interaction variable between an indicator that a stock is restricted from trading under a waiver at one venue and an indicator that at the restriction announcement, the specific venue was ranked number *i* by the importance of dark trading at it. *i* = 0 means that the venue was not dark trading this particular stock, *i* = 4 indicates the fourth and all other less active venues. Regressions 1–7 include stock-fixed effects, 2–7 – time-fixed effects, regressions 3–7 – venue-fixed effects. The estimates in Column 4 (5) report coefficients estimated on a subsample of observations at the venues that are able to dark-trade below (above) month-level median number of stocks. The estimates in Column 6 (7) report coefficients estimated on a subsample of observations for the stocks commonly traded in at most (more than) two dark pools before the ban. The table does not report for the coefficients of the constant. All regressions are ordinary least squares with errors clustered at the issue level. All observations are at a venue-month-stock level.

effect on total dark trading volume for assets with both more and less developed dark pool markets. The difference in the effect between the two markets is not significant (Table 5, Columns 1–2). Stocks that dark-traded at multiple platforms are not less sensitive to platform-level prohibitions. The result does not support our second hypothesis. Despite the presence of alternative platforms, the low level of volume at these alternatives makes it likely that traders do not use them because their orders are unlikely to find a match and execute.

We also examine whether the single stock ban affects the degree of competition among dark pool markets. We show that under the restriction, the probability that a security is admitted to trading in another dark pool increases by 12 % (Column 3). Platforms respond to restrictions at other platforms by entering the market for dark trading in the banned stock. However, despite being available for trading on more platforms, a security is not actually traded on a larger number of venues (Column 4). These results are consistent with a lack of alternative platforms not being a barrier to substitution across dark platforms.

Finally, we show that the market concentration in dark trading the restricted stock, measured with the Herfindahl–Hirschman Index (HHI), decreases by 0.04 each month (Table 5, Column 5). In the following subsection, we show that the decrease in dark pool trading on the restricted platform is driving this result.

6.2. Competition among dark trading platforms

We examine competition by first showing how the single platform restriction affects investors' choice of where to transact their dark trades. The pure competition model suggests that trading will migrate evenly to all other potential platforms, but we already show that total dark trading severely declines after the single-platform ban. Therefore, we should expect that the pure competition null will not hold. In Table 6, we estimate Eq. (2) to examine whether the intensity of using the other platforms affects the trading volume that remains in the dark after the

ban. We rank all platforms by level of preexisting trade in the banned stock from 1 to 4. 1 represents the restricted venue, 2 and 3 represent the second and third rank of other platforms by their share of dark trading the stock over the previous year; 4 indicates the fourth and smaller active platforms. Not all alternative platforms have preexisting trade for every banned stock. $0.BanV$ indicates a platform which had no preexisting trade but positive levels of trading during the ban.

Not surprisingly, trading on the largest platform declines. Large-size trades are still permitted, so dark pool trading in the stock on the restricted platform still exists, but small-trade waivers are not granted during the trading ban. However, dark trading during the ban generally does not significantly increase at any alternative platform, regardless of the amount of existing active platforms available. Yet, there are two particular exceptions to this conclusion.

Specifically, in line with expectations, dark trading decreases for the most active platform. After reaching a 4 % platform dark-to-total trading level, the suspension leads to a 0.37–0.46 percentage point monthly decrease in trading volume (Table 6, Columns 1–4). The ban's negative effect on trading at the largest platform is first estimated controlling for time-invariant stock-specific characteristics using security fixed-effects (Column 1). When we add additional controls by augmenting the regression with month fixed-effects (Column 2), or when we control for platform-specific characteristics by augmenting the regression with platform fixed-effects (Column 3) and platform-month fixed effects (Column 4), the measured decline remains. The coefficient on $1.BanV$ is negative and significant, suggesting that the ratio of dark pool trading at the largest platform over the last 12 months to total dark trading decreases from 4 % to 1.2 % over the six-month ban period. This decrease shows that all of the decline in total dark volume is due to the decline in trading under the waiver at the banned platform.¹⁶

Trading at the less active dark platforms does not respond to the ban. Only in the more stringent regression specifications (Table 6, Columns 3 and 4), do we observe that new entrants trading in that security secure additional significant volume. In these regressions the coefficient on the new entrants, $0.BanV$, is positive and statistically significant, but economically very small.

Finally, we investigate whether investors migrate their dark trades to lesser-known platforms. To evaluate this prediction, we split our sample into two groups by the number of stocks available to trade on a dark platform. We expect platforms with lower dark trading activity to experience a lower level of trade migration after trading is restricted on the leading platform. Table 6 Columns (4) and (5) report that the migration activity is similar across more or less active preexisting platforms, rejecting Hypothesis 2. Both platforms exhibit an economically small amount of dark trading going to new entrants, but in both cases this migration is a small percentage of the lost volume.

The most active migration we observe occurs when a particular security was traded at only two platforms before the ban. In this case we observe migration to the second platform, 7 % of the decrease in the banned dark pool volume drifts to the second platform ($2.BanV / 1.BanV = 0.036 / (-0.54)$) (Table 6, Columns 6–7). We interpret the ban-related overall decrease in dark pool trading as the regulation causing a decline in investor interest in dark pools. However, in our dataset, we observe the matched orders that execute, rather than the orders themselves. Thus, the lower dark trading volume levels could indicate less frequent matches. After the suspension, orders may migrate to multiple venues looking for a counterparty, but due to the low level of preexisting volume, the probability of matching a buyer with a seller is small. In Column (6) traders know there is only one active alternative platform, and a modest increase in dark pool volume is observed. This result indicates

two effects. First, the 7 % level of substitution confirms the idea that institutional traders tend to cluster in particular dark pool silos, and that switching platforms is costly enough to minimize the use of alternative dark pools. Second, the problem of finding a match when many alternative platforms are available, indicates that even if the institution is willing to switch, a coordination problem exists since it is difficult to predict the platform that other traders have migrated to.

Columns (6) and (7) also report that stocks with only two active platforms is the more common case (51,178 observations against 33,834 for stocks with trading on multiple platforms). For stocks that traded on multiple platforms prior to the ban, there is no commensurate increase in trading on the second platform, and actually a small decline during the ban period. The more developed market seems to exacerbate the matching problem. Not only are the alternative platforms small, and thus matching is unlikely, but also there is a coordination problem, even if a trader wants to stay in the dark, it is more difficult to find an off-setting order. When dark buyers are less certain to meet dark sellers, they trade less in dark.

Overall, our findings suggest that when dark trading at the most active platform is restricted, little trading volume migrates to the other pools. We do observe small substitution effects in cases when prior to the ban only two dark pools were active.

7. Conclusion

The European regulatory authorities introduced wide-ranging restrictions intended to inhibit the growth of dark trading, to protect the quality of lit market price discovery. Two of the rules involved trading prohibitions. One restriction involved the total prohibition of dark trading 'under the waiver' if a stock's dark trading volume crossed the threshold of 8 % of total volume. Johann et al. (2019) report that this regulation is ineffective at migrating trading volume back to lit venues. Instead, traders looked for alternative execution mechanisms such as internalization.

A second part of the double volume cap regulation is the restriction of trading a particular stock on a particular platform when dark trading volume on that platform reaches 4 %. In this instance, the regulation appears neither to increase lit trading volume, nor competition across dark pools. Total dark trading in the stock declines by over 50 percent over the period the ban is in place, and the ban also appears to significantly limit the growth in dark pool volume for that security. 12 % of the decline in dark pool trading migrates to the systematic internalizers, but there is no evidence of economically significant migration to the lit markets or other dark pools. The draconian nature of these restrictions appears to have inhibited the growth of the dark pool market structure at this early stage of its development.

The lack of substitution across dark pools as well as the relative infrequency of 8 % bans over our sample period, indicates that only a single dark pool restriction is necessary to accomplish ESMA's regulatory aims. Perhaps recognizing this reality, regulators have recently proposed a single 7 % limit on dark trading to be introduced in late 2025.¹⁷ Based on the results in this paper, if this rule is implemented, we predict a market structure with only a single prominent dark pool will still obtain, only with a less restrictive cap. We leave it to future research to confirm or deny our conjecture.

Aside from the success or failure of the regulation, this regulatory experiment permits us to study the degree of substitution across dark pools when the ban is in place at the former market leading platform. The degree of substitution is important in determining the market structure of the dark pool market. We find substitution across platforms to be very weak due to two main reasons. First, the concentration of liquidity, and thus the likelihood of a match upon submission to the dark

¹⁶ Dark pools in Europe are invariably exchange owned or independent (ITG, Liquidnet), unlike the U.S. dark market where brokerage-owned dark pools are significant traders. Therefore, pre-existing broker ties do not influence our results, as they might in the U.S. market.

¹⁷ Banking Risk and Regulation, "In the shadows: Regulators eye dark pool reforms," October 9, 2024.

pool, appears to be important for dark pool substitution. Traders are too unsure of execution at alternative venues, and do not substitute trade across venues. The single instance of some migration is the setting of two previously active venues with dark trading restricted on one of them. Then, the traders see a clear alternative. Even in this case, traders migrate only 7 % of the decreased volume to an alternative dark pool.

The second reason substitution is weak is related to trader motivation to execute on multiple venues. The benefits of dark trading in the single banned security are not strong enough to encourage traders to invest in the relationship with alternative platforms. The low level of substitution implies a good deal of stock-specific specialization since the bans are imposed on several different platforms, rather than a single

always-dominant platform. Therefore, new dark pool entrants are unlikely to be successful as a broad-based entity. Finally, the lack of substitution across dark pools suggests that institutional investors are unlikely to switch platforms should an HFT enter their preferred dark pool.

CRediT authorship contribution statement

Paul J. Irvine: Writing – original draft, Project administration, Methodology, Formal analysis. **Egle Karmaziene:** Writing – original draft, Methodology, Formal analysis, Data curation.

Appendix A

Technical details 4 % rule

In 2007, European policymakers adopted Markets in Financial Instruments Directive (MiFID), which introduced pre-trade transparency waivers. The four waivers allowed not publishing bid and offer prices by the trading venue before an order was executed and thereby introducing dark pools.

There are four waivers (Guagliano, 2020):

- Reference price waiver (RPW): Systems matching orders based on the midpoint within the current bid and offer process of the trading venue where that financial instrument was first admitted to trading or the most relevant market in terms of liquidity.
- Negotiated trade waiver (NTW): Systems that formalize negotiated transactions.
- Large in scale (LIS): Orders that are large in scale compared with normal market size.
- Order management facility (OMF): Orders held in an order management facility of the trading venue pending disclosure.

According to MiFID II the LIS orders are defined based on the average daily turnover. Table A1 defines the average daily turnover (EUR) and minimum order sizes qualifying for LIS when MiFID II was introduced (Comerton-Forde, 2017). Yet, MiFID II is an active regulatory arrangement the thresholds change over time. They are announced every spring and apply for the next year.

To ensure implementation of these waivers across markets, in 2018 double-volume cap (DVC) mechanism was introduced with the MiFID II. It is a combination of 4 % and 8 % rules. Dark trading in equity and equity-like instruments is restricted for six months on a particular venue if a ratio of trading under the waiver on this venue to trading in those financial instruments across

Table A1
Large-in-scale threshold definitions.

Average daily turnover (e)	Minimum order size qualifying as LIS
< 50k	15,000
50k - 100k	30,000
100k - 500k	60,000
500k - 1m	100,000
1 m - 5m	200,000
5 m - 25m	300,000
25 m - 50m	400,000
50 m - 100m	500,000
> 100m	650,000

all EU trading venues is higher than 4 %. Dark trading in equity and equity-like instruments is restricted for six months on all venues if a ratio of trading under the waiver in those financial instruments across all EU trading venues is higher than 4 %.

In the beginning of each month, the European Securities and Markets Authority (ESMA) announces the annual levels of total and under the waiver trading across all European venues as well as under the waiver trading levels per each venue. When the DVC is violated, the National Competent Authorities (NCAs) suspend trading under the waiver in these instruments for the next six months. The suspensions come into force within the following two days. The rules apply only to instruments having liquid markets. Specifically, if the stock has:

- A free float at least €100 million.
- At least 250 average daily transactions.
- And an average turnover of at least €1 million.

The exact purpose of the 4 % rule is nebulous. To the best of our knowledge, its goal is never mentioned apart from the 8 % dark pool volume restriction. The purpose of these simultaneously introduced regulations is to ensure that the bulk of executed volume occurs on the lit market. Yet, the 8 % rule seems sufficient, by itself, to accomplish this aim. Why also introduce the 4 % rule on individual platforms? Perhaps because without the 4 % rule, a dark pool could achieve 5–7 % of volume in a stock, and thus become a liquid dark alternative, since, at this size, the dark pool has enough order flow to have a reasonable chance of finding a matching order. Under the 4 % rule, dark platforms are constrained to be small volume platforms that do not possess a large population of resting orders. As we will show, this lack of liquidity regulated into the European dark pool market severely limits substitution to alternative dark platforms.

Appendix B

Table B1
US dark pool data.

<u>Platform Name</u>	<u>Owner</u>	<u>Market share</u>
UBS ATS	UBS	18.00 %
Sigma X2	Goldman Sachs	11.20 %
Crossfinder	Credit Suisse	9.10 %
JPM-X	JP Morgan	7.40 %
LEVEL ATS	Luminex	6.40 %
MS Pool	Morgan Stanley	6.30 %
IBKR ATS	Interactive Brokers	4.90 %
Intelligent Cross	Imperative Execution	4.60 %
Barclays ATS	Barclays	4.30 %
BIDS ATS	Cboe	3.90 %

U.S. market share of dark pool trading data for July 2021. Data is collected from FINRA ATS. *Platform Name* is the reported dark pool name on FINRA ATS. The *Owner* is the ultimate owner of each particular dark pool. *Mkt Share* is the market share of total U.S. dark pool trading in July 2021.

Data availability

the data is publicly available from European regulators. Code will be available on request.

References

- AMAFI, 2019. Impact of the MiFID II Volume Cap Mechanism On the Microstructure of European equity Markets, 19. AMAFI.
- Bekaert, G., Harvey, C.R., Lundblad, C.T., Siegel, S., 2011. What segments equity markets? *Rev. Financ. Stud.* 24 (12), 3841–3890.
- Brogaard, J., Pan, J., 2019. Dark trading and the fundamental information in stock prices. Unpublished working paper. Available at <https://papers.ssrn.com/sol3/papers.cfm>.
- Brugler, J. and Comerton-Forde, C. (2021). Differential access to dark markets and execution quality. Available at SSRN 3910952.
- Comerton-Forde, C., 2017. Shedding light on dark trading in Europe. In: CEPR-Imperial-Plato Inaugural Market Innovator (MI3) Conference.
- Comerton-Forde, C., Malinova, K., Park, A., 2018. Regulating dark trading: order flow segmentation and market quality. *J. Financ. Econ.* 130 (2), 347–366.
- Degryse, H., De Jong, F., Kervel, V.v., 2015. The impact of dark trading and visible fragmentation on market quality. *Rev. Finance* 19 (4), 1587–1622.
- Degryse, H., Karragiannis, N., Tombeur, G., Wuyts, G., 2021. Two shades of opacity: hidden orders and dark trading. *J. Financ. Intermediation* 47.
- Di Maggio, M., Egan, M., Franzoni, F., 2021. The value of intermediation in the stock market. *J. Financ. Econ.* 145 (2), 208–233.
- Fang, B., Hope, O.-K., Huang, Z., Moldovan, R., 2020. The effects of MiFID II on sell-side analysts, buy-side analysts, and firms. *Rev. Account. Stud.* 25 (3), 855–902.
- Farley, R., Kelley, E.K., Puckett, A., 2018. Dark trading volume and market quality: a natural experiment. In: 13th Annual Mid-Atlantic Research Conference in Finance (MARC) Paper.
- Foerster, S.R., Karolyi, G.A., 1998. Multimarket trading and liquidity: a transaction data analysis of Canada–US interlistings. *J. Int. Financ. Mark. Inst. Money* 8 (3–4), 393–412.
- Gomber, P., Clapham, B., Lausen, J., Panz, S., 2018. The impact of MiFID II/MiFIR on European market structure: a survey among market experts. *J. Trading* 13 (2), 35–46.
- Gresse, C., 2006. The effect of crossing-network trading on dealer market's bid-ask spreads. *Eur. Financ. Manag.* 12 (2), 143–160.
- Guagliano, C., Guillaumie, C., Reiche, P., Spolaore, A., Zanon, A., 2020. DVC mechanism: The impact On EU Equity Markets, 3. ESMA Working Paper.
- Guo, Y., Mota, L., 2021. Should information be sold separately? Evidence from MiFID II. *J. Financ. Econ.* 142 (1), 97–126.
- Johann, T., Putnins, T., Sagade, S., Westheide, C., 2019. Quasi-dark trading: The effects of Banning Dark Pools in a World of Many Alternatives. SAFE Working Paper, p. 253.
- Lang, M.H., Pinto, J., and Sul, E. (2021). MiFID II unbundling and sell side analyst research. Available at SSRN 3408198.
- Mendelson, H., 1987. Consolidation, fragmentation, and market performance. *J. Financ. Quant. Anal.* 22 (2), 189–207.
- Menkveld, A.J., Yueshen, B.Z., Zhu, H., 2017. Shades of darkness: a pecking order of trading venues. *J. Financ. Econ.* 124 (3), 503–534.
- Newey, W., West, K., 1987. A simple positive semi-definite heteroscedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55 (3), 703–708.
- Pagano, M., 1989. Trading volume and asset liquidity. *Q. J. Econ.* 104 (2), 255–274.
- Petrescu, M., Wedow, M., 2017. Dark pools in European equity markets: emergence, competition and implications. *ECB Occas. Pap.* 193.
- Rietveld, J., Schilling, M.A., 2021. Platform competition: a systematic and interdisciplinary review of the literature. *J. Manag.* 47 (6), 1528–1563.
- Werner, I.M., Kleidon, A.W., 1996. UK and US trading of British cross-listed stocks: an intraday analysis of market integration. *Rev. Financ. Stud.* 9 (2), 619–664.
- Ye, M. (2010). Non-execution and market share of crossing networks. Available at SSRN 1719016.