

Competing for Dark Trades*

Paul J. Irvine Egle Karmaziene

March 2023

Abstract

We evaluate how the availability of trading venues affects the evolution of dark trading. We show that a regulatory restriction of dark trading at the most prominent platform has a detrimental effect on dark trading activity. Annual dark trading decreases by more than 50% over the half-a-year restriction period. Consistent with investors' sticky relationships with specific dark pools, our results suggest that substitution across dark pools is remarkably low. We also show that restricting dark trading on a dominant platform does not lead to migration from dark to lit markets, traders revert to the other less transparent markets. 12% of the decrease in ban-related trading volume migrates to the systemic internalizers. Our findings are consistent with the European regulation failing to guide dark traders to lit markets.

JEL: G12, G14, G18, D47.

Keywords: MiFID II, dark pool trading, competition.

*Irvine is at the TCU Neeley School of Business, p.irvine@tcu.edu; Karmaziene is at Vrije Universiteit Amsterdam, Swedish House of Finance and Tinbergen Institute, E.Karmaziene@vu.nl. Karmaziene gratefully acknowledges Nasdaq Nordic Foundation for financial support.

1 Introduction

Trading at dark pools demands examination due to its importance as an exchange mechanism. In January 2021, trading outside public stock exchanges reached 47.2% of total U.S. equity trading volume, up from 39.9% a year earlier.^{1,2} In general, buy-side investors favor dark to lit trading venues due to their lower information leakage and trading costs (Menkveld et al., 2017), despite the fact that these venues are slower and less likely to match buyers with sellers.³ European policymakers feared that the rise in dark pool popularity would reduce the transparency and price formation quality in existing lit markets. Therefore, as part of the MIFID-II regulation, they banned dark trading activity on leading venues if a stock's market share of volume rose above a 4%. Traders reaction to this restriction depends on the ease of substitution across dark trading platforms. If substitution is costless, then a regulatory restriction of dark trading at the dominant dark pool could decrease the level of dark trading concentration, but not the overall level of dark trading.⁵ Alternatively, the diversity of dark pools could isolate traders into particular dark pools and restrict the amount of substitution that occurs when trading at the dominant dark pool is restricted. According to Petrescu and Wedow (2017), dark pools differentiate their services along several dimensions, such as the size of the orders, diversity of instruments traded, mechanisms used to match orders, and order features. These services come at a cost; to avoid overpaying membership fees, investors may choose trading at a subset of dark pools, a practice that necessarily limits order migration across platforms. It remains an empirical question to evaluate the degree of substitution across dark pools.

We examine how an exogenous shock to the supply of potential trading venues affects the evolution of trading within the dark pool market. More specifically, to overcome the endogenous

¹Rosenblatt Securities report:

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

²This 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)).

³Public 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).⁴

⁵From 2018 to 2021, the number of venues reporting dark trades in Europe increased from 149 to 168, the number of dark traded stocks – from 3,893 to 4,538.

venue selection arising from traders' choice of where to place their orders, we examine a quasi-natural experiment – a new regulation in the EEA, Markets in Financial Instruments Directive (MiFID) II, the double-volume cap rule (DVC). Following the DVC regulation, the relevant regulator, generally referred to as the Relevant Competent Authority (RCA), suspends trading on the particular dark pool for six months if more than 4% of annual stock trading volume is transacted at a single dark platform. To facilitate the privacy needs of institutions, large trades are exempt from this regulation if their trade value exceeds a certain stock-specific value.⁶ 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 migration of liquidity away from lit markets to a concentrated dark market. Although Johann et al. (2019) argue that near substitutes for dark trading platforms such as traditional auctions or systematic internalizing systems may capture some of the activity on the banned trading platform, we focus on the dynamics of pure substitution: What happens to the trading structure in the dark pools that are still able to trade?

The underlying logic of our analysis is that restricting dark trading at the most active platform, *ceteris paribus*, exogenously increases to some degree the demand for trading at the other platforms. If migration is costless, and dark pools are perfect substitutes, then the pure substitution expectation is that trading would migrate proportionally to the remaining active dark pool platforms. But we show that is not the case, and explore what this experiment can teach us about trader decisions, platform choice, trading concentration, and the ease of entry into the dark pool platform market.

There are several reasons why suspending dark trading at the most active platform may not be an effective way to reduce dark pool trading. First, 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

⁶ESMA publishes monthly values of trading at a month-stock-platform and month-stock levels.

assets on the same platform and return fully once the suspension is lifted. Second, since some exchanges, such as Nordic, specialize in dark pool trading, their institutional clients are likely to maintain their relationship to access their execution expertise on these platforms. Third, the investors, especially the smaller ones, may find it costly to establish new or to maintain several relationships with different platforms. Dark pool membership fees reach up to 6,000 GBP/year (Petrescu and Wedow, 2017). Small investors would prefer having fewer costly relationship and trading at a platform that can transact in the largest number of stocks. These cost constraints would lead them to concentrate their trading in an active dark pool and use the lit alternative for restricted stocks.

The DVC rule provides a clean setting to evaluate how the traders respond to a shock in dark pool supply. 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 a 2018-2022-year 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 and trading in general, respond to platform-level suspensions of a particular stock. Over the six-month restriction period, total annual dark trading in the restricted stock decreases over 50%, a high level of reliance on the leading platform given alternative dark trading platforms are available. This finding is consistent with traders establishing relationships with a subset of dark trading platforms. Once dark trading at one platform is restricted, the investors are unlikely to migrate their orders to another dark platform. This reluctance to substitute inhibits dark platform competition. Even 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 dominant platform, the probability of dark-restricted stock becoming available on a new platform increases by 12.4%. Thus, an absence of the alternative platforms does not explain stickiness to the dominating dark pool.

The exact purpose of the 4% rule is nebulous. To the best of our knowledge, its purpose 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? Without the 4% rule, a dark pool could achieve 5-7% of volume in a stock, and thus perhaps be a liquid dark alternative, because the dark pool has enough order flow to have a reasonable chance of finding a matching order.

We show that under the 4% rule dark trading suspensions significantly reduce dark pool volume. Trading does not migrate to other available dark platforms. The most likely cause of this finding 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 are 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 lowering the liquidity value of these alternatives (Pagano, 1989).

The lack of migration 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 show that the lack of migration 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. The DVC regulations effectively restricted the development of a small, but growing, network of dark platforms.

The absence of any serious migration across dark pools is an important result for institutional traders in light of Brugler and Comerton-Forde (2021) recent 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 the difficulty in substituting trade location 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. (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 all traders respond by tilting volume towards more immediate execution venues.

Costly switching, including the possibility that platforms raise dark pool access fees over time, could explain our findings. Unfortunately, we do not observe the actual dark trading fees to investigate the trading costs across changing levels of competition, which remains a limitation of our study. However, we show that the platform-level restriction has an equally strong effect on stocks with more and less developed dark trading markets. Trade migration across dark platforms is equally unlikely when dark traders already have active alternative platforms.

We investigate how a suspension in dark trading on one platform affects dark trading on the other platforms. Using the double differences approach, we evaluate how an increase in total trading demand coming from restricted traders migrates across the different available dark platforms. We find that the responses to this suspension 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 and hit a 4% ratio of platform dark trading to total trading at the EEA level. In general, dark trading in the restricted security does not migrate to the other venues. However, when prior to the ban there were only two active dark pools, 7% of the decreased dark trading in the restricted stock migrates to the second-largest platform. This migration is a rational response from traders. The ability of to meet at all only exists at the single remaining active platform. This finding highlights the importance of the 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 highest on the largest 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 migration among trading platforms. The past literature shows that the liquidity in a market 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 dominant 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); Fang et al. (2020); 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 - Association of French Financial Market Professionals, 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 via 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. In particular, 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 accounted for 1% and 8% of total trading volume in Europe in 2009 and 2016, respectively (Petrescu and Wedow, 2017)). A temporal ban on dark trading a particular stock decreases its trading in the dark, but the levels increase again after this 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.⁷

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 differs across platforms.⁸ 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 also more able to differentiate their services. According to Petrescu and Wedow (2017), 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 with the minimum order to be fulfilled 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

⁷For example, due to the inability to attract sufficient trading volumes, the NASDAQ OMX dark pool NEURO was forced to close in 2010.

⁸Petrescu 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 EUR 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.

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 decrease or stop submitting their dark orders. In particular, we test the following hypothesis:

Hypothesis 1 Total dark trading volume of the restricted stock decreases when trading at one platform is restricted.

The number of platforms that have permission to trade a particular stock in the dark varies within Europe. Some stocks can be traded in dark on one platform, other stocks can be traded on forty different platforms. Some securities are never traded in dark. Others are traded on several platforms. Using the differences in market development, we examine whether the ban on dark trading is more substantial in cases with less developed dark pool markets. When one platform dominates trading in a specific stock, traders are less likely to maintain multiple relationships with 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 when this platform is more active. We examine whether:

Hypothesis 2 When stock trading at one platform is restricted, total dark trading volume decreases by more in less developed dark pool environments.

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 Zhu (2014). Liquidity traders are attracted there by lower transaction costs Hendershott and Mendelson (2000). Despite lower immediacy, certain investors prefer dark to lit trading due to lower costs, and a lower risk of information leakage (Menkveld et al., 2017). 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.

Finally, we examine characteristics of dark trades that do migrate to another dark platform.

We determine whether there are any stock-specific characteristics, or alternative exchange characteristics that determine the level of substitution to alternative dark pools.

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, 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 the 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 for smaller trades on the popular platform. 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.⁹

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. Like the U.S. and Canadian markets, European equity markets are fragmented across exchanges and alternative trading systems. 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 (Petrescu and Wedow, 2017).

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. In particular, for each security, they reveal the aggregate volume transacted over the last 12-months, the fraction of this volume

⁹The large transactions threshold varies by asset, depending upon an annual average turnover measure. However, the 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/>

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. They start on from 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. We keep 141 platform-level suspensions. From ESMA website, we also obtain bi-annual security trading volume using systematic internalizers .

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 venues decreases from 11 to 10 and from 88 to 60, respectively.

Table 1 shows that the mean dark trading volume in sample stock constitutes 2.1% of total trading volume. On average, an asset can be dark-traded on 19 venues, but is actually traded on 2.4. The statistic is in line with quite limited competition in dark-pool trading. Next, we zoom in stock-month-venue level dataset. Panel B in Table 1 demonstrates that the average dark platform trades just 0.12% of annual stock trading volume. However, the variable distribution is very 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. We follow trading in suspended stocks over time. 18% of our observations fall on the six-month venue-level suspensions and 1% on the EU-level blank suspensions 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. 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.

The 141 bans in the sample were imposed at 17 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.

Next, we investigate the population of 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 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 these stocks, which is a low level compared to 76.6% at XHEL. Dark trading at 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. Figure 1 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.

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 (Figure 2). 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. Figure 3 shows, that in half of the cases, prior to the ban, the entire dark trading volume was transacted at a single venue.

Table 3 compares the monthly average platform-level trading over months $[-6,-3]$ and months $[-1, 1]$ around the ban. The numbers suggest that problems with risk management at the platform level were unlikely to trigger the ban – there was no sudden change in trading around the ban period. The values of total trading, dark trading and trading at the main venue did 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:

$$\Delta Y_{s,t} = \sigma_s + \theta_t + \beta Ban_{s,t} + \gamma \Delta X_{s,t} + \epsilon_{s,t} \quad (1)$$

where $\Delta Y_{s,t}$ is a measure of trading in month t stock s . We use several different measures of trading. Our primary measure is 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. $Ban_{s,t}$ 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 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 in Johann et al. (2019) report or (iii) execute at a systematic internalizer such as Jane Street, (iv) execute 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 with five lags.

5.2 Dark pool trading migration after the ban

Next, we investigate how the dark trading flows redistribute across platforms in response to restricted, dark trades on the largest platform. In particular, we estimate flows across platforms using the following regression equation:

$$\Delta Y_{s,v,t} = \sigma_s + \theta_t + \eta_v + \beta_i \sum_{v=0}^8 Ban_{s,v,t} + \gamma d.X_{s,v,t} + \epsilon_{s,v,t} \quad (2)$$

where $\Delta Y_{s,t}$ is a measure of trading under the waiver at a platform v in a year to month t in security s (in percent). $Ban_{s,v,t}$ 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 eight actively trading platforms v . The platforms are ranked from the most active (first) to the least active (last) and are related to dummies Ban_1 to Ban_7 . 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 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 on Table 4 present the Equation (1) regression results.

To ensure that the total trading does not drive our estimates, we control for the time-fixed effects and for time- and stock-varying trading turnover. Depending on the set of controls or fixed-effects employed the fraction of dark trading to total trading reduces 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, Column (5), does not materially affect the *Ban* coefficient. Column (6) examines the evolution of trading throughout the 6-month ban period.¹⁰ 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.

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

¹⁰As time under the ban passes, the dark pool volume slowly increases from an initial shock of -0.23 percentage point to a decrease of 0.41 in month 6, the final month of the ban (Table 4, Column 6). The start of the ban always falling between the 10th and the 20th day of the month explains the lower coefficient on the first month of the ban (*1.BanMonth*).

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 was equal 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, the dark trading decreased by between 2.6 to 2.7 percent of total trading. In an alternative regression specification, we replace a dummy variable indicating a ban with six variables taking values of one for each month of the ban. The combination of the six coefficients suggests a 2.77 percentage point reduction in dark pool trading over the six-month ban period (Table 4, Column 6). These findings support Hypothesis 1 and shows 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 trading was able to increase lit trading as the regulator intended. In Table 5 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 5, 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 SIs. Johann et al. (2019) show that a blank EEA ban on dark trading leads to an increased trading volumes at SIs.

The magnitude of the migration to lit venues is similar but the effect is not statistically significant (Table 5, 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.

We next examine Hypothesis 2, whether the 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 at 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 6 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 effect on total dark trading volume for assets with more and less developed dark pool markets. The difference in the effect between the two markets is not significant (Table 6, Columns 1-2). Securities, which are 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 4). 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 5). 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 6, Column 6). 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 7, we 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, 2-3 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. However, 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 7, Columns 1-4). The ban's negative effect on trading at the largest platform is first estimated controlling for time-invariant stock-specific characteristics by including 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 (Table 7, 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.

Trading at the less active dark platforms does not respond to the ban. Only in the more stringent regression specifications (Table 7, 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 they can trade in dark. We expect platforms with lower dark trading activity to experience a lower surge in dark trading once it is restricted on the largest platform. We estimate the regressions in subsamples with more or less active pre-existing platforms. We find that the migration effect is similar for the both types of platforms (Table 7, Columns 4 - 5).

Yet, when a particular security was traded at only two platforms before the ban, we observe dark volume migration to the less active pool. In particular, 7% of the decrease in the banned dark pool volume drifts to the other platform ($2.BanV/1.BanV=0.036/(-0.54)$) (Table 7, Columns 6-7). We interpret the ban-related overall decrease in dark pool trading as a decrease in dark order volume. However, in our dataset, we observe the matched trading, but we do not observe the submitted orders. Thus, the lower dark trading volume levels could indicate less frequent matches. Following the suspension, orders may migrate to multiple venues however, due to the low level of pre-existing volume, the probability of matching a buyer with a seller is small. In Column 6 when traders know there is only one active alternative platform, a modest increase in dark pool volume is observed. This result indicates two effects. First, the 7% level of substitution confirms the idea that institutions 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 where counterparties can be found.

The effect is opposite for securities, which, before the restriction, were traded on multiple platforms. Dark trading volume decreases not only on the first, but also on the second largest platform. 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 observe some substitution effects in cases

when prior to the ban only two dark pools were active.

7 Further exploration of the banned dark pool market

Despite covering securities listed in 11 countries, our sample securities are biased towards stocks having Great Britain (21%) and Germany (32%) as their RCAs. To investigate whether dark trading in stocks with British or German regulators responds differently to restrictions, we evaluate our regression results for the subsamples of securities with British, German, and the remaining regulators (Table 8, Columns 1-3). Our results are robust to ban having severe detrimental effects on dark trading for all three sets of stocks. We also show, that Britain leaving the EU, did not change drastically the way traders respond to dark pool restrictions. After March 2021, when we do not observe British securities traded solely on British venues, the coefficient on *Ban* is slightly smaller but in the range of our other estimates (Table 8, Columns 4-5).

Next, we examine whether the restriction has a different effect depending on the characteristics of the banned stocks. First, we evaluate whether the more liquid securities overcome the restriction easier and their dark trading volume decreases less. Regressions in Table 9 separately estimate the change in dark volume for stocks that have a greater (Column 1) or lower than average level of overall trading (Column 2). These results indicate that the ban has a similar effect on dark trading activity for both heavily-traded and lightly-traded stocks.

Table 9 also reports a test of whether the ban had a lower impact for the stocks with higher dark trading demand. In particular, if the demand for dark trading is high, we expect traders to search for ways to transact a security and dark trade on another platform. For these stocks the restriction could have a limited effect, since stocks with higher dark trading demand could cause traders to spend more effort searching for a dark alternative to the banned platform. We define securities with higher demand for dark pool trading by revealed preference. Specifically, the stocks that traded more than 8% in dark pools at any time during our sample period, and thus were completely banned from dark trading by the EEA. We show that the securities

with both 4% and 8% bans had a larger dark trading volume in total, more trading on non-dominant venues. Trading was not only possible on more venues (19 vs 12.9), trades were actually executed on more venues (5 vs 1.8).¹¹ These statistics suggest not only higher demand for dark trading for these stocks, but also the dark pool market was more developed to satisfy this demand.

We show that the platform-level ban had an economically 15% lower impact on the total dark trading volume of the securities with a higher dark trading demand (0.069/0.475). It appears as though traders make an effort to keep certain stocks in the dark, however, the difference between the two coefficients is not statistically significant.

In Table 10, we investigate whether the subset of restricted stocks is different from the overall population of stocks traded at the EEA. Using the population of stocks from ESMA DVC reports, we compare the trading levels over the first month the stock appears in the EEA reports and report the t-tests of the difference between never restricted and restricted stocks. The trading of restricted stocks, on average, is greater, but the difference in trading volume between the two subsets is not statistically significant. Yet, the restricted stocks are traded more in the dark, so they must have characteristics that lead investors to trade them on dark pools. The level of dark trading for banned stocks is 1.28 percentage points higher during the first month of reporting. In addition, the restricted stocks have more developed dark trading markets. On average, they can be traded on 14.6 platforms, whereas the stocks that are never platform-restricted stocks have trade at only 6 dark platforms.

Restricted securities are dark-traded on significantly more platforms than the unrestricted securities, 2.3 and 0.6 respectively. The large difference in the number of potential platforms (6-14.6) and actually used platforms (0.6-2.3) suggests that very few dark pools successfully attract trade in a stock.

Finally, we investigate dark trading a month after it was restricted at the most active platform. In particular, for such events with all data available, we compare 54% of cases when dark trading migrated to other platforms with the remaining 46% cases when dark-trading at

¹¹These differences are reported in Panel B of 9.

unrestricted platforms did not increase after the ban. We report the p-values of t-tests in Table 11.

Stocks whose dark trading migrated from restricted to unrestricted venues did not have significantly different trading volumes, numbers of venues reporting zero or positive dark trading just before the suspension. Yet, trading and dark trading volumes, dark-trading share at non-largest venue, are significantly lower for securities that experience volume migration to non-restricted dark pools. These last two tables suggest that the lack of platform migration after the ban is more likely to be related to the lack of development of the dark pool market, rather than for any stock-specific reasons.

8 Conclusions

The European regulatory authorities introduced wide-ranging restrictions intended to inhibit the growth of dark trading, to protect of 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 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 almost 50 percent over the period the ban is in place, and fails to migrate to alternative available platforms. The single-platform ban 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 statistically 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 nascent level of its development.

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 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 entities. 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.

References

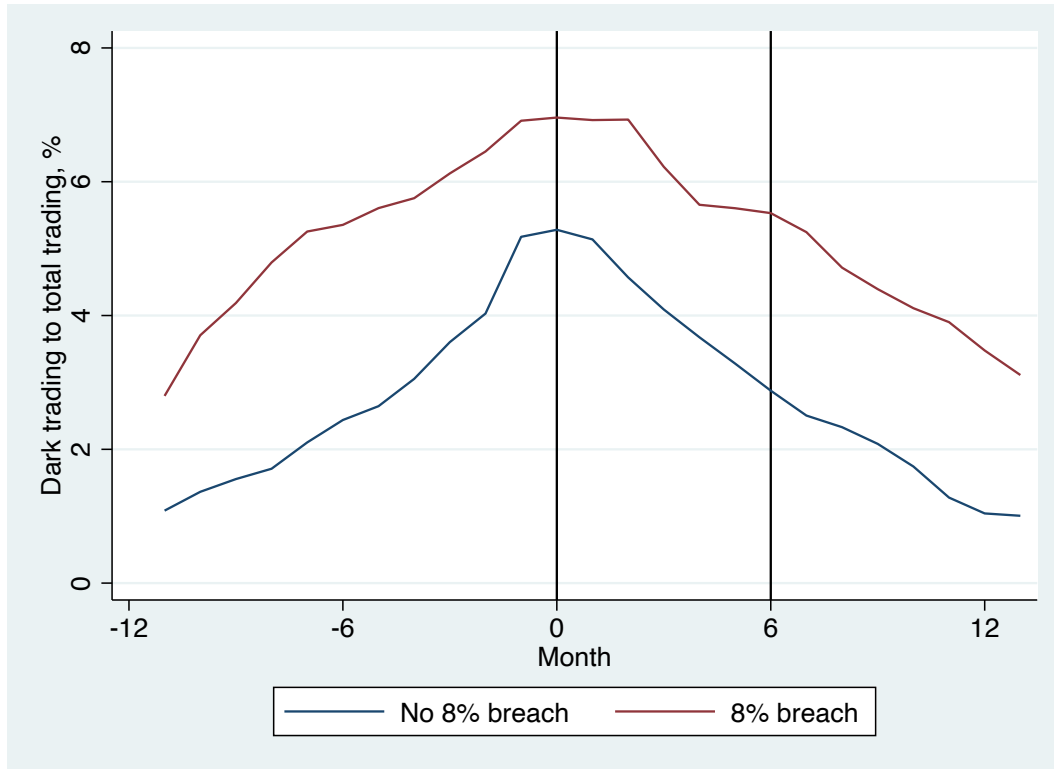
- Bekaert, G., Harvey, C. R., Lundblad, C. T., and Siegel, S. (2011). What segments equity markets? *The Review of Financial Studies*, 24(12):3841–3890.
- Brogaard, J. and Pan, J. (2019). Dark trading and the fundamental information in stock prices. *Unpublished working paper. Available at: <https://papers.ssrn.com/sol3/papers.cfm>.*
- 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., and Park, A. (2018). Regulating dark trading: Order flow segmentation and market quality. *Journal of Financial Economics*, 130(2):347–366.
- Degryse, H., De Jong, F., and Kervel, V. v. (2015). The impact of dark trading and visible fragmentation on market quality. *Review of Finance*, 19(4):1587–1622.
- Di Maggio, M., Egan, M., and Franzoni, F. (2021). The value of intermediation in the stock market. *Journal of Financial Economics*, 145(2):208–233.
- Fang, B., Hope, O.-K., Huang, Z., and Moldovan, R. (2020). The effects of mifid ii on sell-side analysts, buy-side analysts, and firms. *Review of Accounting Studies*, 25(3):855–902.
- 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*.
- Foerster, S. R. and Karolyi, G. A. (1998). Multimarket trading and liquidity: a transaction data analysis of canada–us interlistings. *Journal of International Financial Markets, Institutions and Money*, 8(3-4):393–412.

- Gomber, P., Clapham, B., Lausen, J., and Panz, S. (2018). The impact of mifid ii/mifir on european market structure: a survey among market experts. *The Journal of Trading*, 13(2):35–46.
- Gresse, C. (2006). The effect of crossing-network trading on dealer market’s bid-ask spreads. *European Financial Management*, 12(2):143–160.
- Guagliano, C., Guillaumie, C., Reiche, P., Spolaore, A., and Zanon, A. (2020). Dvc mechanism: The impact on eu equity markets. *ESMA Working Paper*, 3.
- Guo, Y. and Mota, L. (2021). Should information be sold separately? evidence from mifid ii. *Journal of Financial Economics*, 142(1):97–126.
- Hendershott, T. and Mendelson, H. (2000). Crossing networks and dealer markets: Competition and performance. *The Journal of Finance*, 55(5):2071–2115.
- Johann, T., Putnins, T., Sagade, S., and Westheide, C. (2019). Quasi-dark trading: The effects of banning dark pools in a world of many alternatives. *SAFE Working Paper*, 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. *Journal of Financial and Quantitative Analysis*, 22(2):189–207.
- 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.
- Pagano, M. (1989). Trading volume and asset liquidity. *The Quarterly Journal of Economics*, 104(2):255–274.
- Petrescu, M. and Wedow, M. (2017). Dark pools in european equity markets: emergence, competition and implications. *ECB Occasional Paper*, 193.

- Rietveld, J. and Schilling, M. A. (2021). Platform competition: A systematic and interdisciplinary review of the literature. *Journal of Management*, 47(6):1528–1563.
- Werner, I. M. and Kleidon, A. W. (1996). Uk and us trading of british cross-listed stocks: An intraday analysis of market integration. *The Review of Financial Studies*, 9(2):619–664.
- Ye, M. (2010). Non-execution and market share of crossing networks. *Available at SSRN 1719016*.
- Zhu, H. (2014). Do dark pools harm price discovery? *The Review of Financial Studies*, 27(3):747–789.

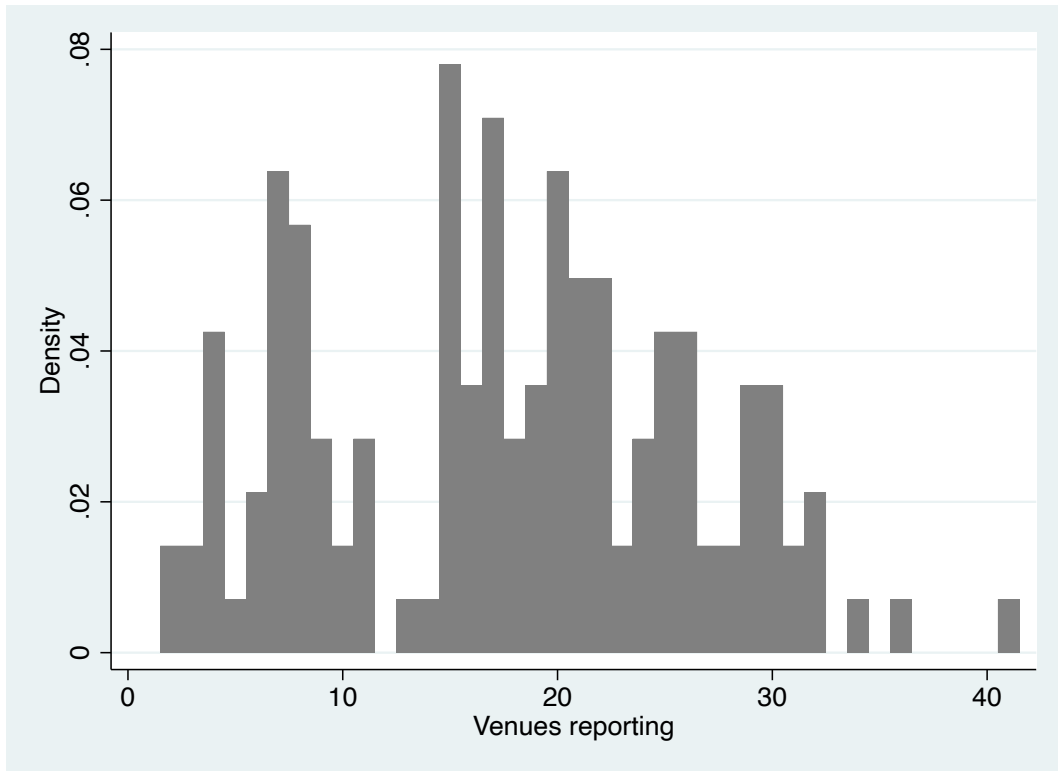
Figures and Tables

Figure 1: Evolution of dark trading around the venue-level suspension.



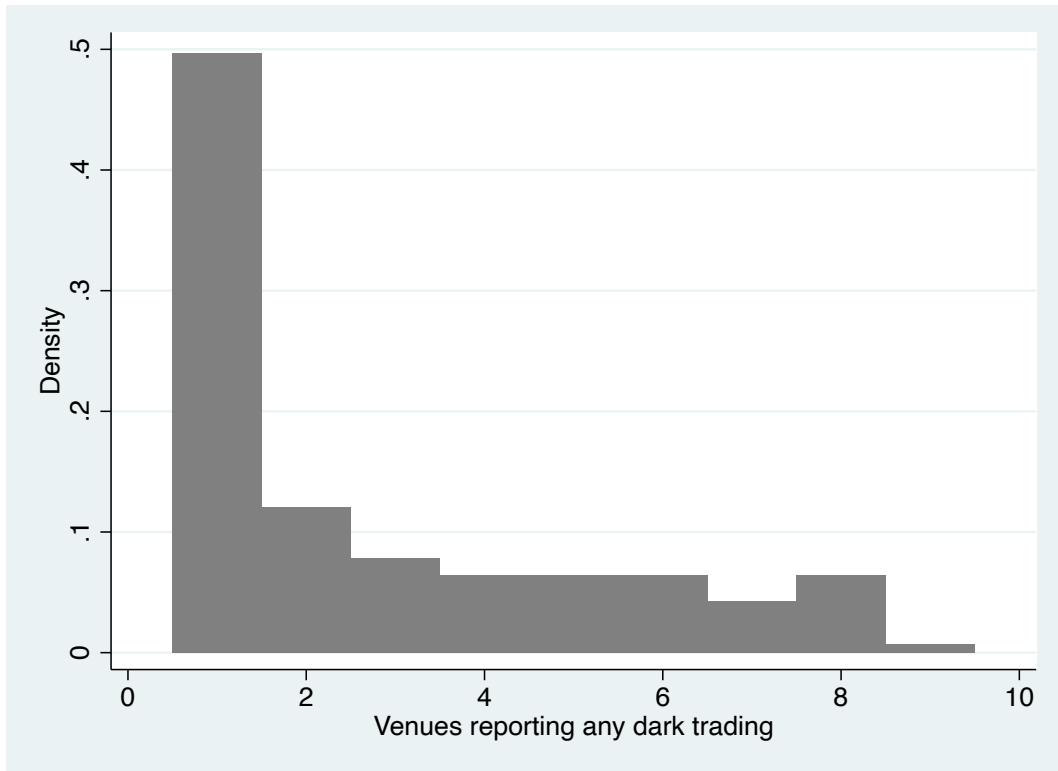
The graph represents the monthly average annual total dark trading for stocks 12 months prior and post the ban venue-level suspension. The red (blue) line indicates mean dark pool trading levels for stocks that have (have not) experienced EEA-level ban over our sample period. The vertical lines refer to the beginning and the end dates of the restrictions.

Figure 2: Venues that can dark-trade the suspended stock



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

Figure 3: 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.

Tables

Table 1: Summary statistics

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 to 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.

	count	mean	sd	p25	p50	p75
Panel A						
DarkVol (%)	5,726	2.10	2.15	0.00	1.46	3.76
CountVenues	5,726	19.05	9.91	11.00	19.00	26.00
CountNon0	5,726	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

Table 2: Sample venues and their characteristics

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 %.

		Observations	Stocks	Traded in dark	Dark
AIMX	London Stock Exchange AIM MTF	46,021	1,168	4.20	0.02
BCXE	CBOE Europe London	230,905	6,249	61.48	0.68
CEUD	CBOE Europe Amsterdam	109,021	4,207	48.78	0.26
LIQU	Liquidnet Europe Limited	1,024,104	27,847	10.28	0.02
MERK	Euronext Oslo	3,396	170	13.53	0.37
MTAA	Euronext Milan	14,747	367	54.77	0.15
TQEM	Turquoise Europe	46,788	2,637	67.96	0.42
TRQM	Turquoise Plato	190,913	4,778	67.50	0.54
XCSE	Nasdaq Copenhagen	34,798	818	29.58	0.44
XHEL	Nasdaq Helsinki	8,273	205	76.59	0.44
XLOM	London Stock Exchange Non-AIM MTF	304,182	7,379	14.65	0.01
XLON	London Stock Exchange	182,769	8,588	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	1,642	40.68	0.35
XUBS	Aquis Exchange PLC	159,104	3,868	78.26	0.45

Table 3: Trading prior to the restriction

This table compares platform-level trading over months [-6,-3] (*Pre-ban*) and months [-1, 1] (*Ban*) around the ban. *Trading, M* is the total trading at the restricted venue over 12 months, in million EUR. *DarkVol* and *DarkVenue* are ratios of total dark trading and dark trading at the restricted venue to total trading at the EEA level, in %.

	Pre-ban	Ban	Difference	P-value
Trading, M	910.86	901.20	-9.66	0.97
DarkVol, M	44.71	48.08	3.37	0.82
DarkVenue, M	33.45	37.51	4.06	0.71

Table 4: Regression results

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.

	(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	5,726	5,726	5,726	5,726	5,726	5,726
Time FE	No	No	No	Yes	Yes	Yes
Stock FE	No	No	Yes	Yes	Yes	Yes

Table 5: Economic effects

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.

	(1) Δ Dark	(2) Δ Dark	(3) Δ Lit	(4) 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	4,040	4,040	4,040	4,040
Time FE	Yes	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes	Yes

Table 6: Dark trading restrictions and competition

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 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.

	(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					
Developed	0.400** (0.176)	0.401** (0.175)			
Δ logTrade	0.005 (0.023)	0.005 (0.023)	0.085** (0.034)	0.000 (0.000)	0.011 (0.010)
BanEU		-0.534*** (0.093)	0.137 (0.228)	0.000 (0.000)	0.025 (0.032)
Observations	5,726	5,726	5,639	5,726	5,726
Time FE	Yes	Yes	Yes	Yes	Yes
Stock FE	Yes	Yes	Yes	Yes	Yes

Table 7: Migration of dark pool trading across the venues

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. *i.BanV* 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.

	(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)
Constant	0.052*** (0.007)	0.028*** (0.009)	0.028* (0.016)	0.036* (0.021)	0.021 (0.026)	0.025 (0.025)	0.018 (0.022)
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	No	No	Yes	Yes	Yes	Yes	Yes
Market						Active2	Activemore2
Venue				Large	Small		

Table 8: Robustness tests

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 EEA venues) and zero otherwise. *logTrade* is logarithmic value in EUR of trading a particular issue on European exchanges over the last 12 months. Regressions in Columns 1, 2, and 3 include subsamples of stocks under supervision of the British, German, and other authorities. Regression in Columns 4 and 5 include subperiods before and after British stocks traded solely on British venues disappeared from our sample. All regressions include time and stock-fixed effects. All regressions are ordinary least squares, Newey-West standard errors are with five lags. All observations are at a month-stock level.

	(1)	(2)	(3)	(4)	(5)
	Δ DarkVol	Δ DarkVol	Δ DarkVol	Δ DarkVol	Δ DarkVol
Ban	-0.609*** (0.051)	-0.430*** (0.065)	-0.380*** (0.045)	-0.506*** (0.030)	-0.358*** (0.102)
BanEU	-0.619*** (0.106)	-0.467** (0.209)	-0.626*** (0.099)	-0.545*** (0.117)	-0.976*** (0.144)
Δ logTrade	0.120 (0.134)	-0.026 (0.026)	-0.007 (0.016)	0.015 (0.027)	-0.043 (0.032)
Constant	-0.141 (0.090)	0.225 (0.168)	-0.040 (0.102)	-0.141*** (0.039)	0.178 (0.208)
Observations	1,193	1,848	2,685	4,269	1,457
Country	GB	DE	Other		
Brexit				Before	After

Table 9: Sample split tests

PANEL A

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 EEA venues) and zero otherwise. *EU-banned-stock* indicates if a stock was banned from dark trading at the EU level any time during our sample period. *logTrade* is logarithmic value in EUR of trading a particular issue on European exchanges over the last 12 months. A regression in Column 1 (2) includes a subsample of stocks, which had annual trading levels below (above) the sample median. All regressions are ordinary least squares with Newey-West standard errors at five lags, include time and stock-fixed effects. All observations are at a month-stock level.

	(1)	(2)	(3)
	Δ DarkVol	Δ DarkVol	Δ DarkVol
Ban	-0.419*** (0.029)	-0.507*** (0.051)	-0.475*** (0.028)
BanEU	-0.666*** (0.088)	-0.394** (0.165)	
Ban \times EU-banned-stock			0.069 (0.114)
Δ logTrade	-0.266*** (0.100)	0.015 (0.023)	0.005 (0.023)
Observations	2,983	2,743	5,726
Trading	High	Low	

PANEL B

This table compares the characteristics of two types of securities. The first group is restricted at a venue level, the second group experienced restriction at a venue and EEA level. We compare their characteristics on the first month they appear in our sample. *Volume, M* (*Dark volume, M*) is the trading (dark trading) volume at the EEA level, in million EUR over the last 12 months. *DarkVol* is a ratio of annual total dark trading to total trading at the EEA level, in %. *NonFirst* is a ratio of dark trading at other than the largest venues to total trading at the EEA level, in %. *Venues* and *Venues non zero* count the number of venues submitting any or non-zero dark trade reports on a particular stock.

	Venue	Venue and EEA	Difference	P-value
Volume, M	860.96	1,704.09	843.13	0.19
Dark volume, M	23.30	87.26	63.96	0.05
DarkVol, %	1.87	4.21	2.34	0.00
NonFirst, %	0.43	2.54	2.12	0.00
Venues	12.92	19.00	6.08	0.02
Venues non zero	1.77	5.00	3.23	0.00

Table 10: Bans

This table compares the characteristics of stocks that are restricted at the venue level against stocks that are never restricted. The comparison data is the first month the stock appears on the EEA banned list. The comparison is in the first month the stocks appear in the reports. *Trading, M* is total annual trading in million EUR. *DarkVol* is a ratio of total dark trading (to total trading at the EEA level, in %. *Venues reporting* and *Venues non zero* count the number of venues that submit the under the waiver trading reports to ESMA and venues that have transacted the stock in the dark over the previous year.

	Not restricted	Restricted	Difference	P-value
Trading, M	419.56	936.98	517.42	0.22
DarkVol	0.80	2.08	1.28	0.01
Venues reporting	6.00	14.57	8.56	0.00
Venues non zero	0.55	2.29	1.74	0.00

Table 11: Migration of dark orders

This table compares the characteristics of stock which, upon dark trade suspension on the largest venue, migrate their trades to the other venues and the ones that do not. *Volume, M* (*Dark volume, M*) is the trading (dark trading) volume at the EEA level, in million EUR over the last 12 months. *DarkVol* is a ratio of annual total dark trading to total trading at the EEA level, in %. *NonFirst* is a ratio of dark trading at other than the largest venues to total trading at the EEA level, in %. *Venues* and *Venues non zero* count the number of venues submitting any or non-zero dark trade reports on a particular stock.

	Not migrated	Migrated	Difference	P-value
Volume, M	1,203.35	543.99	659.36	0.07
Dark volume, M	71.26	30.09	41.16	0.05
DarkVol, %	5.44	5.41	-0.03	0.90
NonFirst, %	1.04	0.60	0.44	0.03
Venues	18.62	16.63	1.99	0.18
Venues non zero	3.19	2.78	0.41	0.34