



BANK OF ENGLAND

Staff Working Paper No. 545

Into the light: dark pool trading and intraday market quality on the primary exchange

James Brugler

September 2015

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Authority Board.



BANK OF ENGLAND

Staff Working Paper No. 545

Into the light: dark pool trading and intraday market quality on the primary exchange

James Brugler⁽¹⁾

Abstract

This paper uses regulator-provided transaction data to investigate how trading in dark pools affects intraday market quality on the limit order book of the primary exchange for members of the FTSE 100 index. Using trading patterns from execution algorithms as instrumental variables, I show that dark trading leads to improved liquidity on the primary exchange, both in absolute terms and relative to trading on the limit order book. Although these relationships differ across stocks of different sizes, dark trading does not lead to worse market quality at the intraday level for either small or large stocks during the sample period.

Key words: Dark pools, dark trading, market quality.

JEL classification: G10, G12, G14, G18.

(1) Bank of England and Faculty of Economics, University of Cambridge: Email: james.brugler@bankofengland.co.uk, jab244@cam.ac.uk.

I wish to thank Evangelos Benos, Carole Comerton-Forde, Jonathan Dark, Jane Cooley Fruehwirth, David Easley, Erik Hjalmarsson, Oliver Linton, Maureen O'Hara, Hashem Pesaran, Roberto Rigobon, Farzad Saidi, Pedro Saffi, Filip Zikes and seminar participants at the University of Cambridge Faculty of Economics, Cass Business School, Queen Mary University of London School of Economics and Finance, the Judge Business School, the University of New South Wales School of Banking and Finance, the University of Melbourne Department of Economics, the University of Melbourne Department of Finance, the Cambridge-INET Empirical Market Microstructure Workshop and the Bank of England FMI Seminar Series for their helpful comments. I also wish to thank the Bank of England and the Financial Conduct Authority for providing data for this project.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

1 Introduction

During the middle of the previous decade, securities regulators in both the US and the EU enacted a series of reforms with the intention of promoting competition in the provision of trading services for financial securities; Regulation NMS (RegNMS) in the US and the Markets in Financial Instruments Directive (MiFID) in the EU. A natural consequence of these reforms has been the proliferation of new trading venues for equities that compete with traditional exchanges. A notable subset within these new venues are “dark pools”, trading venues that limit pre-trade transparency. Though the specific rules of these venues can vary considerably, a common theme is that each trader’s interest remains hidden from the rest of the market unless a buyer can be matched with a seller, after which a trade occurs. Understanding how trading in these dark pools (“dark trading”), and competition between trading venues more generally, affect market quality is a key question for regulators tasked with promoting market integrity and stability, as well as for market practitioners concerned with optimal execution strategies. Indeed, regulators in both the US and the EU, citing concerns regarding transparency and fairness, are considering or have recently enacted rules to limit the amount of trading that can occur on dark venues relative to the primary exchange and other displayed venues (see e.g. European Commission (2014), Financial News (2014), Reuters (2014) and The Wall Street Journal (2014)).

A number of recent papers investigate the effect of dark trading on measures of market quality such as liquidity, volatility, transaction costs and price efficiency using data aggregated at the daily level or at lower frequencies (e.g. O’Hara and Ye (2011), Degryse, De Jong and Van Kervel (2014), Weaver (2014), Comerton-Forde and Putniņš (2014), Foley and Putniņš (2015) and Gresse (2015)). In this paper I use regulator-provided transaction data for constituents of the FTSE-100 index across multiple dark pools to fill a gap in this literature by estimating how dark pool trading affects liquidity and volatility on the primary exchange at the transaction level. These high frequency relationships are of direct importance for regulators tasked with promoting financial stability, due largely to the continued ascendancy of computer-based trading in increasingly fragmented markets and the associated risks of market failures manifesting over extremely short time frames (e.g. 2010’s “Flash Crash”). By analysing transaction-level data across dark pools and the primary exchange, this paper attempts to provide new empirical evidence regarding the contribution of dark trading to short-term market instability, something that previous work in this field using lower frequency data is unable to do. Understanding these

high frequency relationships also complements the existing studies by providing evidence about which are the most important channels through which dark trading affects market quality. By drilling down to the transaction level, I aim to disentangle the effect of dark trading on both liquidity and volatility, despite feedback effects between these variables that are important at lower frequencies. Such feedback effects may be due to interactions between market liquidity and the funding liquidity of market makers (Brunnermeier and Pedersen, 2009) or increased liquidity premia for fund managers during volatile periods in response to increased probabilities of redemptions (Vayanos, 2004).

A key challenge in analysing the relationship between dark trading and market quality is the endogeneity that arises from traders rationally choosing the venue and timing of trading based on expected future market conditions. This clearly can result in simultaneity from the outcome variables (e.g. volatility, spreads and depth) to the variable of interest (trading on dark venues). I address this identification issue by exploiting a novel source of exogenous variation in the probability of observing trades that comes from regular patterns in trade arrivals driven by unsophisticated or “low technology” execution algorithms. These algorithms (such as some of those that guarantee the time-weighted or volume-weighted average prices to the end-user) make trades at regular, pre-determined intervals throughout the trading day. This generates clustering of trade arrivals at these regular intervals compared with other periods during the day regardless of short-term expectations of market conditions.

I use these regular intraday patterns in the trading of such algorithms as instruments for trading in dark pools and trading on the limit order book of the primary exchange (“lit trading”) in order to estimate the causal impact of both types of trading on market quality. These patterns provide sources of high frequency variation in dark and lit trade arrivals that are independent of expected market conditions: on average there should be no fundamental changes occurring in the marketplace during the regularly-spaced timestamps when these algorithms are more likely to make trades, other than through the effect of increased trading activity on both kinds of venues. Importantly, by consistently estimating the effect of dark and lit trading together, I am able to compare the effect of both types of trading while also controlling for the effect of total volume (the sum of lit and dark volumes) on market quality.¹

¹While it may appear more natural to estimate how the *fraction* of dark trading affects the outcome variables rather than dark and lit volumes, at a very high frequency this is complicated by the presence of many periods with no trading, during which the fraction is undefined. This becomes even more problematic in a dynamic model when some trading must occur at each lag for the regressor to be defined which greatly reduces the number of useful observations, especially as the number of lags increases.

In order to identify the impact of both types of trading, I exploit the fact that the degree of clustering in trade arrivals differs across interval frequencies (e.g. one minute vs. five minute) and across venue types. In other words, the behaviour of the execution algorithms is not such that we see a proportional rise in volume across all venue types during the relevant intervals, which would invalidate the approach. Crucially, using standard tests for underidentification, I empirically validate that the clustering of trade arrivals is sufficiently different across venue types and intervals to identify the model. I also conduct a number of other robustness tests for the identification strategy. I examine the composition of trader types during the regular intervals compared with other time-stamps, I test whether changes to the composition of trader types affects liquidity and volatility at the frequency used in the analysis, and I test whether participation decisions by users of execution algorithms are affected by low frequency expectations about market quality. These robustness checks suggest that heterogeneous effects across trader types and endogenous low frequency trading decisions do not appear to invalidate the identification strategy.

My results indicate that the level of dark trading in the UK equity market during the sample period (September to December 2012) did not contribute to short-term market instability or lead to deterioration in market quality on the primary exchange, either in absolute terms or relative to trading on lit venues. In fact, trades on dark venues lead to a significant improvement in depth (shares available) at the best bid and offer on the primary exchange without significantly impacting the inside spread or volatility.² Using an autoregressive distributed lag (ARDL) framework with instruments for lit and dark trading based on the intraday patterns outlined above, I estimate that a dark trade for 5,000 shares leads to a contemporaneous improvement of approximately 2,250 shares in the depth at the best bid and offer. The associated long-run multiplier predicts that an increase of 10 shares in the expected dark volume traded per second leads to a statistically significant increase in depth of around 160 shares. In contrast, trades on the limit order book (LOB) of the primary exchange lead to a contemporaneous and long-lived increase in both the inside spread and volatility while having only a shorter and smaller positive impact on depth. A trade for 5,000 shares on the LOB would lead to an estimated increase in the inside spread of 1.7 basis points (bp) and of 2.1 bp in the absolute second-to-second midquote return. Depth in the LOB is estimated to increase by approximately 680 shares however this increased depth is at wider quoted prices and the long-run impact on depth

²The inside spread is the difference between the best bid and best offer in the limit order book, expressed as a % of the mid-price.

is insignificant.

I also investigate whether the effect of dark trading on intraday market quality differs across stocks of different sizes by estimating the model for three categories of stocks based on market capitalisation. Regressions based on data split across the three market capitalisation categories do provide some evidence that dark trading affects stocks of different sizes in different ways, however dark trading does not worsen market quality for large, mid-sized or smaller stocks relative to lit trading. For stocks in the largest category, dark trades do not appear to affect the state of the limit order book on the primary exchange in a meaningful way, other than via a long-run reduction in spreads. For mid-sized and smaller stocks, however, dark trades lead to short-run and long-run increases in depth that are significantly larger than the corresponding effects of trades on the LOB. For mid-sized stocks, dark trades do lead to increased short-run and long-run volatility, however this is significantly smaller in magnitude than the corresponding effect of LOB trades. For stocks in the smallest category, dark trades lead to wider spreads in the short-run and long-run, but these effects are not statistically distinguishable from those of LOB trades.

For the full sample, the improvement in liquidity on the primary exchange without a corresponding effect on volatility following dark trades is consistent with the theoretical predictions that dark venues attract and facilitate trading activity primarily from uninformed or liquidity traders as per Hendershott and Mendelson (2000) and Zhu (2014). Although both of these models suggest that this would lead to worse liquidity on the primary exchange, due to traders with private information clustering on the LOB of that market, a crucial feature of these models is that market makers do not trade in the dark pool. However, the transaction data used in this paper identify the counterparties to each trade and indicate that financial firms with large market making divisions (such as investment banks and high frequency traders) are very active in dark pools and account for a significant fraction of total trading in such venues. By trading in both dark and lit markets, these market making firms are able to still capture a fraction of this profitable uninformed order flow, decreasing their net adverse selection costs. Furthermore, when clearing the positions obtained by trading with uninformed traders in dark venues, market makers are not directly exposed to “picking-off” risk by informed traders, and can therefore post more competitive prices on the primary exchange.

The results from the regressions split across market capitalisation categories also provide some support for the predictions of Buti, Rindi and Werner (2014) that for larger stocks, dark pools attract a higher fraction of market orders to limit orders from

the LOB compared with smaller, less liquid stocks. This implies that spreads widen more for small stocks than for large stocks. The estimated long-run multipliers for large stocks and small stocks are consistent with these predictions, although the short-run parameters are not insofar as this parameter for large stocks is insignificant.

The main conclusion of this paper, that the levels of dark trading in the UK equity market during the sample period did not pose a threat to the integrity of the market at a high frequency, contrasts with current regulatory trends surrounding dark trading, particularly in the EU. The latest revision of MiFID, to be implemented in January 2017, includes a cap on dark trading of 8% of total volume with no more than 4% on any single dark venue (European Commission, 2014). My results suggest that contrary to the concerns of regulators, dark trading at these proportions does not appear to destabilise markets or unduly affect the price formation process. This conclusion is consistent with related empirical work that examines the effect of dark trading at daily or monthly frequencies such as O'Hara and Ye (2011), Comerton-Forde and Putniņš (2014) and Foley and Putniņš (2015) where moderate amounts of dark trading are not found to be detrimental to market quality. My results suggest that the main high frequency channel through which dark trading leads to improved market quality at lower frequencies appears to be via liquidity (and specifically increased depth at the best bid and offer) rather than directly via volatility.

The rest of the paper proceeds as follows. In section 2, I present a brief summary of the related literature in this field, section 3 outlines the datasets used in this paper and section 4 presents summary statistics from these data that give a new and detailed picture of dark trading in the UK equity market. Section 5 contains the specification and results from the high frequency regressions that use intraday patterns in trade arrivals as instrumental variables, as well as a number of robustness checks for the identification strategy. Section 6 presents separate regressions for stocks split into three market capitalisation categories. Section 7 offers concluding remarks.

2 Related Literature

Since the introduction of competition-enhancing reforms adopted in the US in 2005 (RegNMS) and in the EU in 2007 (MiFID), there have been several empirical papers that examine how non-displayed trading on alternative trading facilities affects market conditions, for equity markets in the US as well as the EU and the Asia-Pacific region. A number of these papers employ an instrument variable strategy to identify the effect of

dark trading on market quality. The only other paper that I am aware of that considers dark trading in the UK equity market is Gresse (2015). Gresse (2015) uses the implementation of MiFID as a natural experiment affecting fragmentation and the prevalence of dark trading to estimate their effect on the liquidity of stocks in the FTSE-100 index, as well as stocks in the CAC-40 and the SBF-120 indices. Her analysis relies on comparing liquidity at the daily and monthly level before and after the implementation of MiFID (both comparing means and using a panel regression) as well as an IV approach that uses a number of other instruments for fragmentation such as market capitalisation, trade size, number of venues and average tick size. Her results suggest that competition between trading venues led to an improvement in spreads and depth and that dark trading has not harmed global liquidity. The implementation of MiFID coincided with a number of important other events that affected liquidity and volatility, such as the onset of the global financial crisis, which, as noted by Gresse (2015), somewhat complicates interpretation of results that rely on comparing market conditions before and after such events.

Other papers that use an IV or natural experiment approach to identify how dark trading affects market quality, or how market quality affects the fraction or volume of dark trading, include Degryse et al. (2014), Comerton-Forde and Putniņš (2014) and Foley and Putniņš (2015). Degryse et al. (2014) uses average order size and the number of limit orders to market orders as instruments for fragmentation, and average dark order sizes as an instrument for dark trading for their sample of 52 large Dutch stocks. They find that visible fragmentation is generally positive for global liquidity while dark trading has a detrimental effect with a one standard deviation increase in dark trading lowering global liquidity by 9%.³ Comerton-Forde and Putniņš (2014) use a number of changes to market structure and trading rules on the Australian Stock Exchange (ASX) as well as average dark trading in other stocks with similar market capitalisations as instruments to identify the effect of dark trading on price discovery in the Australian equity market. Their results suggest that low levels of dark trading are not harmful to price discovery but that as the share of dark trading rises, price discovery can suffer (they estimate that the threshold where dark trading becomes harmful is around 10% of total volume). Comerton-Forde and Putniņš (2014) also conclude that informed traders are less likely to execute trades on dark venues compared with lit traders. Foley and Putniņš (2015)

³Like O'Hara and Ye (2011) (as well as Comerton-Forde and Putniņš (2014), Korber et al. (2013) and others) the effect of fragmentation on market quality is found to be non-linear, with low levels found to be beneficial but the marginal improvement decreasing with the level of fragmentation. Only a linear term is included for the effect of dark trading.

use a natural experiment generated by a change in Canadian regulations that required dark trades under a certain size to provide an improvement of at least one tick over the prevailing national best bid and offer. This regulation change led to a very dramatic decrease in dark trading of Canadian stocks. Their results indicate that low levels of dark trading lead to reduced quoted, effective and realised spreads as well as improved price efficiency.

My approach departs from the previous literature primarily by using a novel identification strategy to identify the effect of dark trading on market quality at the intraday level. Whereas the studies mentioned above primarily use data aggregated at the daily level or higher, I believe this is the first paper to robustly identify this relationship at a very high frequency (1 second). These high frequency relationships are of growing importance in modern financial markets, where the rise of computerised trading has introduced new sources of short-term instability (Linton, O'Hara and Zigrand, 2012), and can also shed new light on the mechanisms that drive the lower frequency relationships from the studies above. The paper also uses an extremely broad and detailed proprietary dataset made available from the Financial Conduct Authority, the UK securities market regulator. This data allows me to identify transactions, including counterparties, volume, price and trade-time to one second precision by all EU-regulated trading firms on over 120 different trading venues.

3 Data

The data used in this paper cover the period from 3 September 2012 until 31 December 2012 and is obtained from three sources. Firstly, transaction reports for all trades executed by EU-regulated trading entities in the constituents of the FTSE-100 index are obtained from the Financial Conduct Authority's (FCA's) ZEN database, the FCA's surveillance and monitoring system. Under MiFID, all EU-regulated firms are required to submit data for reportable transactions via an Approved Reporting Mechanism or directly from the trading venue by the close of business the day after a trade is executed. Failure to do so can result in substantial fines (up to £5.6m in recent cases (Financial Conduct Authority, 2013)). Each entry in the ZEN database contains the security name, volume of shares, direction of the trade (buy or sell), transaction price and currency, total consideration of the trade, trading venue, trade time to the nearest second and the other counterparty to the trade (most commonly one of the clearing houses used by the trading venue). Stocks that had been removed from the index at the end of the period are

omitted from the analysis, as are Royal Dutch Shell A-Class and B-Class shares due to complications isolating transaction reports from certain trading institutions across these securities. I also remove trades between Christmas and New Year's Eve due to historically low activity on these days. This leaves transaction reports for a total of 92 stocks in the index across 80 trading days. The dataset contains approximately 182 million separate transaction reports from 1,215 unique trading entities reporting trades at 126 different venues, although the 30 most prevalent venues by transaction reports account for 99.9% of all reports in the database. Of these 30 venues, I identify seven dark-only venues that include non-displayed multi-lateral trading facilities (MTFs) and crossing networks owned by brokers and investment banks. For the purposes of this paper, I follow Mittal (2008) and refer to dark venues owned by brokers (i.e. those that do not allow proprietary order flow from the venue owner) as public crossing networks or PCNs, investment bank-owned non-displayed MTFs and crossing networks (i.e. those that allow proprietary order flow from the venue owner) as internalisation pools or IPs and exchange-owned dark trading venues as exchange-based pools or EBPs. I use this transaction data to construct high frequency (1 second) time-series of trade volumes on dark venues (in section 5) and summary statistics regarding trades by trader type on dark venues (figure 2 in section 4).

The ZEN database does not contain transaction reports for entities regulated outside of the EU. This implies that overseas trades by institutions with legal entities in multiple jurisdictions are not captured in the data. Conversations with market supervisors at the FCA suggest that the majority of such firms do choose to trade via their local branches and this issue should not substantially affect the coverage. The self-reported nature of the data also implies that reporting errors, such as incorrect amounts, prices or time-stamps, can be present. Regarding incorrect time-stamps, since the prices of trades in dark venues are derived from the bid and offer on the LOB of the primary exchange, the London Stock Exchange (LSE), I am able to filter out incorrect prices and time-stamps for trades on the dark venues by comparing the transaction price with the prevalent bid and offer on that LOB (obtained from Bloomberg, details below) and removing all trades that do not occur within these prices. Such trades outside of these prices constitute less than 5% of the sample. I also compare the reported consideration (£ value of the trade) with the price and volume to remove systematic errors in the volumes reported. Misreporting of trade details and lack of overseas coverage introduces a degree of measurement error into my high frequency time-series of dark transactions and this would downwardly bias OLS estimates. However, regressions using IVs based on trading patterns of execution algorithms will still yield consistent estimates of key parameters and standard errors

despite this measurement error under the standard IV requirements.

Regarding coverage across types of dark pools, the ZEN database does not identify the particular order book that a trade is executed on for venues with multiple order books. As such transactions in the category of EBP cannot be directly identified as these are not distinguished from trades that occur on the lit order book of the same venue. Accordingly, the dark transaction data used in this paper exclusively covers dark-only venues (those that run a single, non-transparent order book or matching program) and the results need to be interpreted with this in mind. However, since it is much easier for dark orders on EBPs to interact with the orders on that venue's lit order book through venue-provided interbook sweep orders, the degree to which venues with multiple order books contribute to the amount non-display liquidity is less than for PCNs or IPs. As such, these venues are arguably of less importance for analysing how dark trading affects market quality.

I complement the transaction data in the ZEN database with Bloomberg data containing the prices and volumes of all trades executed on the London Stock Exchange (LSE) over the sample period as well as the best bid and best offer in the LOB, and the volume of shares available at these quotes. These data contain all trades regardless of the legal domicile of the trading entity and identifies each trade by trade type (including automated trade, normal trade, contra-trades, late cancellations and so on). The counterparties to the trades are not identified in this database. I focus on trades with the classification automated trade as these occur on the LOB of the exchange. These data are used to construct information about lit trading throughout the trading day (in sections 4 and 5). I also obtain the daily total volume traded on the LSE as well as the three large alternative trading facilities Chi-X, BATS and Turquoise from Bloomberg.

Finally, I obtain trade volume data at the daily level for 10 non-displayed trading venues (four EBPs, four PCNs and two IPs) from Fidessa, a private firm providing investment, trading and information services to financial clients. Fidessa collates trading volume data for a wide variety of securities directly from trading venues themselves across the world. These data are representative of the entire trading community rather than just EU-regulated trading entities and so I use these for all daily-level analysis, combined with the Bloomberg data for the LSE, Chi-X, BATS and Turquoise (section 4).

4 Summary Statistics

The data collated at the daily level from Bloomberg and Fidessa as well as the transaction data from the ZEN database have not previously been studied in the context of dark trading and so a number of summary statistics that broadly describe the state of dark trading in the UK equity market are presented here.

Table 1 contains mean and median volumes of trades (numbers of shares and £ value of shares) on the three categories of dark venues outlined in section 3 as well as these quantities for the limit order book on the LSE. Dark volumes traded are constructed from the Fidessa data, while the LSE LOB data are obtained from the Bloomberg quote and trade data. Total value traded is calculated using Bloomberg data containing all trades on the LSE, BATS, Chi-X and Turquoise and the total trading volume on the 10 dark venues according to Fidessa.

Table 1: Volumes traded by venue category

Mean and median daily volume of shares traded, £ volume traded and fraction of total trading by venue for stocks in the FTSE-100 index. The venues include the limit order book of the London Stock Exchange (LSE LOB) and the dark order books of four exchange-owned dark venues (EBPs), four broker-owned public crossing networks (PCNs) and two investment-bank owned internalisation pools (IPs). Fractions are of total volume of shares traded across the LSE, Chi-X, BATS, Turquoise and all dark venues.

Venue	Vlm (shrs)		Vlm (£)		Fraction	
	Mean	Median	Mean	Median	Mean	Median
LSE LOB	4,907,060	1,730,496	21,441,911	13,684,248	0.4033	0.4004
EBP	372,273	104,899	1,593,420	835,100	0.0304	0.0250
PCN	169,070	29,463	728,204	234,528	0.0141	0.0065
IP	190,112	51,525	874,543	411,191	0.0156	0.0121
All dark venues	731,455	215,487	3,196,168	1,704,004	0.0601	0.0499
All venues	12,168,022	4,322,027	53,169,447	34,177,313	-	-

Table 1 indicates that the average share of dark trading across these stocks is approximately 6%, which is lower than the equivalent share in the US equity market (estimated to be around 14% in 2012 (Degryse et al., 2013)). At approximately 3% of total trading, proportions of trading are larger in total on those venues that are operated by an exchange (EBPs) which may reflect the greater ease of routing orders to such venues that operate in parallel to other lit venues and so do not require separate arrangements for access. PCNs have the lowest share both in total and on average (there are four EBPs, four PCNs and two IPs used to construct these statistics) while IPs have a fractionally

higher share in total than PCNs and a comparable per venue share of trading with that of EBPs. The LOB of the LSE attracts approximately 40% of total order flow across these venues, which is a substantially higher share than the largest US trading venues attracted during the comparable period (according to Comerton-Forde (2013) only 25% of trading in NYSE-listed securities occurred at NYSE venues in 2012). That fragmentation has occurred to a greater degree in the US than in the UK and elsewhere in Europe partly reflects the later adoption of competition reforms in the EU.

Figure 1 contains time-series plots for the fraction of total trading accounted for by dark venues across all stocks on each trading day in the sample, split into the three venue categories, EBPs, PCNs and IPs. These plots indicate that the total share across all three categories of venues is relative stable across the time-frame I consider (i.e. no strong positive or negative trends). The shares of each venue within the total amount of dark trading that occurs are also relatively stable (i.e. peaks and troughs tend to occur simultaneously at all three venue categories).

Figure 1: Dark trades by venue category

Time-series for the fraction of total £ value traded in stocks from the FTSE-100 on four exchanged-owned non-displayed dark venues (EBPs), four broker-owned public crossing networks (PCNs) and two investment bank-owned internalisation pools (IPs). The fraction is calculated using all trades for all stocks on each day (total trading is defined as in table 1) and the series are stacked vertically such that the sum of all three represents the total fraction across all three venue categories.

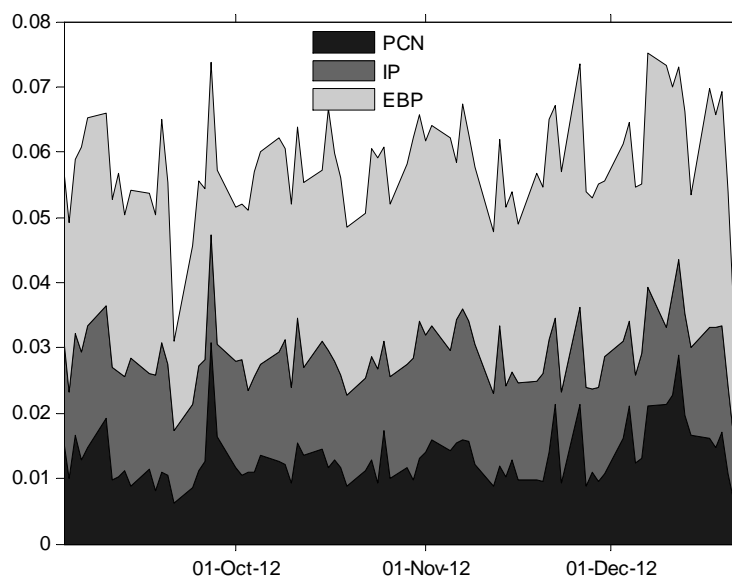


Table 2 present correlation coefficients constructed from the panel of stocks over the 80 trading days. Correlation coefficients are calculated for eight variables: the fraction

of trading occurring on dark venues; the total £ value of trading on dark venues; total £ value of trading on the LSE LOB; daily realised volatility (RV) calculated using the mid-price on the LSE LOB at five minute intervals; weighted inside spread measured as a percent of the prevailing mid-price (where weights are proportional to mean volume available at the bid and the offer at the point in the trading day); average depth available at the bid and offer; 5 minute/1 minute variance ratios using the mid-point on the LSE LOB; and the daily Amihud ratio (open-to-close absolute return divided by total £ volume traded).

The RV, spread and depth all display correlations with expected signs and significance. Higher volatility and spreads are positively correlated with each other and negatively correlated with depth. Both the spread and the depth have very high correlations (in absolute terms) with the Amihud ratio, which arguably reflects the extent to which this statistic is able to capture the state of liquidity in the market (as well as being a good proxy for price impact according to Goyenko, Holden and Trzcinka (2009)). This is relevant as this ratio can still be calculated when order book detail is not available.

Table 2: Dark trading correlation matrix

Correlation coefficients for £ volume of dark trading, £ volume of lit trading and fraction of dark trading with various market statistics across stocks from the FTSE-100 index. The correlation coefficients are calculated from the pooled sample across stocks and days. The market statistics are 5min realised volatility using midpoint returns, value weighted mean inside (BBO) spread on the LSE LOB, mean depth (average of bid and offer depth), 5min/1min variance ratio using midpoint of quotes on the LOB and daily Amihud ratios (open to close return/£ value traded). * indicates significance at the 1% level.

	Dark frac	Dark val (£)	Lit val (£)	RV	Sprd	Depth	VR	Ami
Dark frac	1.0000	0.3579*	-0.065*	-0.092*	-0.014	0.0832*	-0.035*	-0.018
Dark val (£)	-	1.0000	0.7772*	0.2103*	-0.411*	0.1585*	-0.026	-0.189*
Lit val (£)	-	-	1.0000	0.2431*	-0.546*	0.1305*	-0.017	-0.235*
RV	-	-	-	1.0000	0.1899*	-0.083*	0.0767*	0.0504*
Sprd	-	-	-	-	1.0000	-0.085*	0.0270	0.3437*
Depth	-	-	-	-	-	1.0000	-0.034*	-0.260*
VR	-	-	-	-	-	-	1.0000	0.0248
Ami	-	-	-	-	-	-	-	1.0000

While dark value and lit value display similar correlations with the other market statistics (positive and significant correlations with RV and depth, negative and significant correlations with the inside spread and the Amihud ratio - see table 2), correlations with the fraction of trading on dark venues differ in important ways.⁴ Indeed, the fraction of

⁴Since the denominator used to construct the Amihud ratio is the total value of trading, there is an

dark trading is negatively correlated with RV and the variance ratio and not significantly correlated with the spread. The fraction of dark trading is also positively correlated with the depth. That the fraction of trading is negatively correlated with RV is in line with the model of Zhu (2014) who predicts such an effect conditional on some informed traders trading in dark pools. In contrast, Ye (2012) predicts the opposite effect due to greater incentives for informed traders to hide their trades when volatility is higher.

Figure 2: Dark trades by trader type - all venues

Value traded in £m by trader type (left panel) and proportion of total trading accounted for by trader type (right panel) for all stocks in dark pools in the ZEN database. Traders are sorted into three types: investment banks (IB), high frequency traders (HFT) and brokers. Stocks are allocated into three groups by market capitalisation rank.




Figure 2 contains the total value (£m) traded in dark pools across three trader types: high frequency traders (HFTs); investment banks (IBs); and brokers. These three categories of firms account for over 99% of total dark pool trading in the database. Firms in the HFT category were identified by the FCA in consultation with venue operators and contains firms for whom the primary business is low-latency algorithmic trading (and so does not contain high-frequency trading desks within financial conglomerates).⁵ Investment banks are identified by own-company descriptions and include firms with both brokerage and proprietary order flows as well as providing traditional investment banking services (mergers and acquisitions, capital raising) and asset management. Brokers are identified by own-company descriptions and crucially do not submit orders to the market on a proprietary basis - all order flow from these firms is on the behalf of customers. I

element of mechanical correlation between total value traded and the Amihud Ratio.

⁵The classification of firms into the HFT category was made at the time the data were obtained and may not be an accurate reflection of the universe of high frequency traders operating today. Crucially though, the classification is accurate for the sample period. For confidentiality purposes, the firms used in the analysis in this section cannot be named.

calculate the total volume traded by each type of firm for three groups of stocks sorted by market capitalisation rank. The left-hand panel of each figure contains the total amount traded, allocated between the three types of firms, while the right-hand panel indicates the percentage shares of the total volume accounted for by each type of firm. A matching algorithm is used to identify all trades on dark venues between two EU-regulated firms (matching on price, volume, trade time, venue and opposing trade directions). This is done to ensure that such trades are not double-counted.

The proportions of volume accounted for by each trader type is relatively stable across the three market capitalisation categories, with HFTs making up approximately 10-15% of total trading, while IBs account for approximately 50% of all trades and brokers for the remaining 35%-40%. Crucially, these three categories of traders all have access to multiple trading venues as well as sophisticated order routing technology that is able to sweep across multiple venues in extremely short periods of time (e.g. within a second). In the case of IBs and HFTs, market makers within these firms (be they trading desks or market making algorithms) are also able to make markets across these different venues simultaneously. That these three types of traders constitute over 99% of total dark pool trading implies that these trading venues are less isolated from the primary exchange than they would be if dark pools were only populated by end-user firms with market makers confined to displayed venues. The trading of these three sophisticated trader types across the multiple venues can help integrate the multiple venues and facilitate the interaction of trading interest on separate venues. This feature of dark pools is also relevant for a number of recent theory papers that assume market makers do not have access to dark pools when modelling the interaction of such venues with lit markets (e.g. Hendershott and Mendelson (2000), Ye (2012), Zhu (2014)). In such an environment, concentration of informed trading on the displayed venue, combined with uninformed, liquidity traders using the displayed venue as a “market of last resort” (i.e. only in the event they cannot find a counterparty in the dark pool), can have very different consequences than in the case where liquidity providers are present on all types of venues.

5 High frequency IV regressions

In this section, I use the transaction data available in the ZEN database to examine how trading in dark pools affects contemporaneous and future market conditions at a very high frequency. Arguably the most challenging aspect of identifying these effects is dealing with the endogeneity between current and future market conditions with both

timing of trades and venue selection. Traders have discretion over both when they decide to place an order and which type of venue they choose to place that order in. It is reasonable to expect these decisions to be made in a way that optimises some objective function (such as perhaps minimising expected price impact or trading costs including exchange or venue fees). Sophisticated traders may manage this process themselves and others might choose to delegate these decisions either on an agency basis or through direct market access facilities. For large orders that cannot be filled with a single trade, traders must also make these decisions in a dynamic framework where current actions can affect future market conditions. While a fraction of investors may find the costs of market monitoring or third-party order management systems prohibitively expensive to justify strategic order placement, the optimising behaviour of at least some traders is sufficient to ensure that aggregated volumes across different venues will be simultaneously determined with market conditions.

Thus, in order to identify the effect of trading volume at dark and lit venues on current and subsequent market conditions, I rely on instruments generated by regular intraday patterns in trading volumes. Specifically, I identify a positive and significant spike in trading volumes across the stocks in the sample on both dark and lit venues at the turn of every minute (i.e. at 9:00:00AM, 9:01:00AM etc. compared with 9:00:01AM-9:00:59AM), as well as relatively larger spikes in volume on both kinds of venues during times that fall on even 5 minute intervals (9:20:00AM, 9:25:00AM) compared with times between the even 5 minute periods.⁶ Figure 3 displays the volume spikes against times within the minute and times within even 5 minute intervals.

5.1 Intraday patterns in trade arrivals

Intraday patterns similar to those in figure 3 have also been identified by Easley, De Prado and O'Hara (2012) in the E-Mini S&P500 futures contract. The authors attribute the pattern to execution algorithms that operate in clock-time and trade at regular chronological intervals (usually one minute slots). Two typical benchmarks for such execution algorithms are the time-weighted average price (TWAP) and the volume-weighted average price (VWAP). The VWAP is defined as the average price over a given time period

⁶These spikes have been validated not only in the ZEN transaction data containing trades identified by counterparties, but also with data obtained from Bloomberg market services, and in trade data obtained directly from the LSE covering transactions in FTSE-100 stocks from 1st July 2011 to 31st August 2011 used in Brugler and Linton (2014). The latter verification is the most relevant as the time-stamps in these data come from the exchanges own clocks, rather than being self-reported or being reported through through a trade reporting facility.

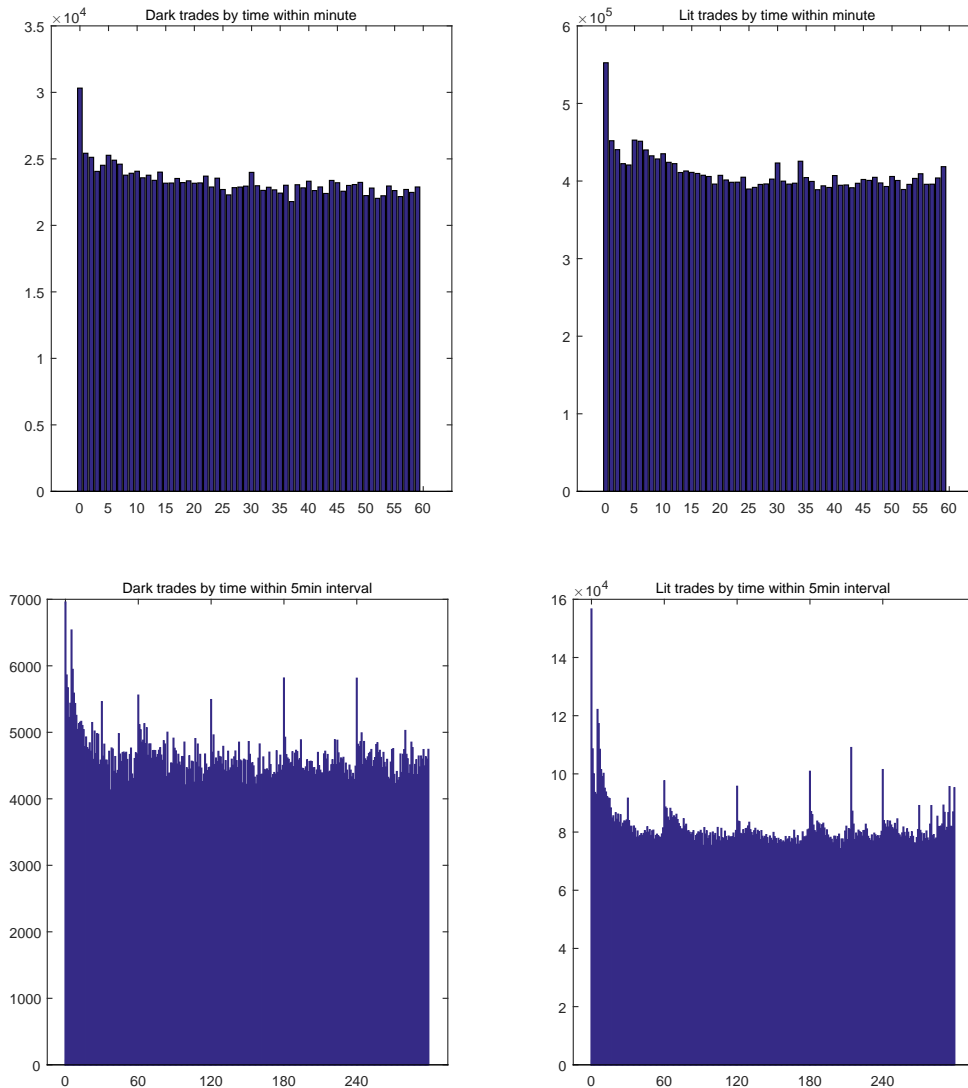
weighted by the volume traded at each pre-specified interval within that period. The TWAP is defined simply as the average price over a given period measured at regular intervals.⁷ While other more sophisticated benchmark prices exist, VWAP and TWAP execution algorithms still account for a significant proportion of institutional investor trading activity. According to a survey of 750 institutional investors conducted by the *TRADE* magazine in 2012, VWAP-benchmarked execution strategies were used by approximately 60% of respondents while TWAP-benchmarked execution strategies were used by around 25% of these traders (The *TRADE* Magazine, 2012).

In order to trade as close to these benchmark prices as possible, investors can choose between a number of strategies as outlined in Madhavan (2002). For example, they may choose to engage in a forward VWAP/TWAP cross whereby buyers and sellers are matched and execute together at the benchmark price at the end of the day. They may instead choose to give the entire trade to a broker-dealer upfront who then guarantees the benchmark price in exchange for a per-share commission. Lastly, the investors themselves could actively trade throughout the day in order to achieve or improve upon the benchmark price. Using the last two options, investors, or their brokers acting in either a principal or agency capacity, will typically break the total order up over the day or other pre-designated period in order to participate either proportionally to that interval's forecasted volume (in the case of a VWAP benchmark) or evenly throughout the day or period (in the case of the TWAP benchmark) (Ibid). In both cases, and especially in the case of the TWAP benchmark, orders hitting the market at regular time-intervals is responsible for generating the pattern seen in figures 3.

⁷The TWAP can also be calculated using the opening or the average of the high and low prices in an interval, or any combination of these and the closing price of that interval.

Figure 3: Trades by time-stamp within even 1 minute and 5 minute intervals

Total number of trades across all stocks in the sample indexed by time within the minute (0s-59s) and by time within even 5 minute intervals (0s-299s). Times designated with 0s refer to the beginning of even 1 or 5 minute periods (e.g. 9:00:00AM, 9:01:00AM or 9:00:00AM, 9:05:00AM respectively). The spike at 0s is driven by unsophisticated, “low technology” execution algorithms that trade at regular intervals, including some that target the VWAP and TWAP. All times corresponding to pre-specified data releases from the Office of National Security and the Monetary Policy Committee of the Bank of England (9:30AM and 11:00AM) as well as time-stamps corresponding to US data releases (12:30PM) and the opening of the US markets (2:30PM) are excluded.



For the purpose of identifying the effect of trading on dark and lit venues, the regular intervals at which trades by these execution algorithms occur provides a source of exogenous variation in the probability of observing trading based on specific time-stamps within the day: trades are more likely to be observed at time-stamps that fall on the turn of the minute or the turn of the even 5 minute interval, regardless of past, current or expected

ted market conditions. As such, these time-stamps can be exploited as instruments for trading on dark and lit venues, respectively. The relevance condition for the instruments is satisfied by the fact that trades are more likely to occur on these intervals (as can be seen in figure 3 and in the first-stage regressions in table 5). The untestable exogeneity condition requires that no other fundamental changes occur in the market place on these time-stamps, other than through the effect of the increased probability of trading. This requires that certain times are removed that are coincidental with pre-specified data releases from the Office of National Security and the Monetary Policy Committee of the Bank of England (9:30AM and 11:00AM) as well as time-stamps corresponding to US data releases (12:30PM) and the opening of the US markets (2:30PM) as these are periods when new information arrives in the marketplace and thus affects market conditions outside of the effect of increased volumes alone.

While this paper is mainly concerned with estimating the effect of dark trading on liquidity and volatility, the instruments are clearly also correlated with lit trading, and so they cannot be used to consistently estimate the effect of dark trading in isolation. In any case, estimating the effect of lit trading on liquidity and volatility is important so as to provide a benchmark for comparison. In order to instrument for both endogenous variables (dark and lit volumes) I rely on the fact that the volume spike that occurs on the even 5 minute intervals is greater in magnitude compared with that which occurs on the turn of minute but does not coincide with an even 5 minute period (see figure 3), and that the relative increase in the volume spike on even 5 minute periods is larger on lit venues than dark venues. In other words the effect on the endogenous variables differs sufficiently across the two instruments, as even greater clustering occurs at the turn of every 5 minute period compared with the turn of the minute on lit venues compared with dark venues. I formally test whether or not these distinct periods have sufficiently different effects on dark and lit volume to identify the model by testing for underidentification and, most importantly, that the $\mathbb{E}[Z'X]$ matrix in the just-identified 2SLS estimator is full rank.

5.2 Panel Instrument Variable Regressions

I construct a panel at the 1 second level over the 80 trading days for each stock in the sample. The high frequency relationship between market quality and dark and lit volumes is modelled using a panel autoregressive distributed lag framework:

$$y_{it} = \alpha_i + \sum_{p=0}^{p^{lit}} \theta_p^{lit} litv_{it-p} + \sum_{p=0}^{p^{dark}} \theta_p^{dark} darkv_{it-p} + \sum_{p=1}^Q \phi_p y_{it-p} + \beta' x_{it} + \varepsilon_{it} \quad (1)$$

where y_{it-p} is the market condition variable (inside spread measured in %, depth at the best bid and offer (BBO) or absolute 1s midquote return in %) for stock i and time $t - p$, $litv_{it-p}$ and $darkv_{it-p}$ are the lit and dark volumes traded for stock i and time $t - p$, x_{it} is a vector of exogenous control variables (time of day and its square, date and its square, the return over the five minutes preceding time t and the log of market capitalisation), α_i is a stock-specific fixed effect and ε_{it} is an error term associated with the realisation of y_{it} . Equation (1) can be written compactly using lag notation:

$$\phi(L)y_{it} = \alpha_i + \theta^{lit}(L)litv_{it} + \theta^{dark}(L)darkv_{it} + \gamma'x_{it} + \varepsilon_{it} \quad (2)$$

where $\phi(L) = 1 - \phi_1L - \dots - \phi_QL^Q$ and $\theta^i(L) = \theta_0^i + \theta_1^iL + \dots + \theta_P^iL^P$ for $i = \{lit, dark\}$ and L is the lag operator. The dynamic multipliers that represent how y_{it} responds to a shock to lit and dark volume in a single period is contained in the corresponding terms of the combined lag polynomials

$$\gamma^i(L) = \phi(L)^{-1}\theta^i(L) \quad (3)$$

while the long-run multiplier is given by

$$\begin{aligned} \gamma_{LR}^i &= \gamma^i(1) \\ &= \phi(1)^{-1}\theta^i(1) \\ &= \frac{\theta_0^i + \theta_1^i + \dots + \theta_P^i}{1 - \phi_1 - \dots - \phi_Q} \end{aligned} \quad (4)$$

for $i = \{lit, dark\}$. In this application, the long-run multiplier can be thought of as the estimated change in market quality for a given change in the *expected* volume of trading on dark or lit venues during each period. In order to identify both the dynamic and long-run multipliers, it is necessary to identify the parameters θ_0^{lit} and θ_0^{dark} that correspond to the contemporaneous effects of lit and dark volume on y_{it} . However, the simultaneity issue described above implies that the error term in (1) will be correlated with current volumes traded on either kind of venue and OLS of (1) will be inconsistent. Accordingly, I define two instruments for $litv_{it}$ and $darkv_{it}$ from the time-of-day seasonalities generated by unsophisticated execution algorithms discussed in section 5.1:

$$D_t^{1min} = \begin{cases} 1 & \text{if time is exact 1min interval} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

$$D_t^{5min} = \begin{cases} 1 & \text{if time is exact 5min interval} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

and use these instruments to estimate (1) via FE-2SLS.⁸ Due to computational constraints associated with the very large number of observations in the panel (~211m), the lag lengths P^{lit} , P^{dark} and Q are fixed at 30, with separate parameters for lagged dark volumes, lit volumes and the left-hand side variable for each second over the previous half minute (98 parameters per regression in total).⁹ Standard errors for the parameters are clustered at the stock-date level. By clustering at this level we can obtain parameter covariance matrices that are robust to any form of error correlation in a given day within a particular stock and that are full rank and invertible. Clustering at a coarser level (e.g. at the stock or date level, or two-way clustering on these dimensions) would guarantee that the covariance matrix of parameters will not be full rank which prohibits joint significance tests across a large number of the parameters. However, t -test for parameter significance in models with fewer lags using standard errors clustered at stock level, at the date level or two-way on these dimensions are largely consistent with those obtained clustered at the stock-date level. Table 3 contains means and standard deviations for the three dependent variables as well as lit and dark volumes traded.

Table 3: Means and standard deviations for regression variables

Mean of the quoted spread, depth at the BBO, absolute 1s midquote return, dark volume and lit volume per period. Standard deviations are in parenthesis. Mean trade size for periods with at least one trade on lit and dark venues are contained in the final two rows of the table. These statistics are calculated across all stocks in the sample as well as across the stocks split into three market capitalisation categories.

Variable	Large	Mid-sized	Small	All stocks
Spread (%)	0.060 (0.021)	0.102 (0.036)	0.114 (0.041)	0.092 (0.033)
Depth	9,602 (5,810)	8,386 (5,306)	6,409 (4,066)	8,138 (5,066)
Abs. 1s rtn (%)	0.011 (0.012)	0.016 (0.019)	0.018 (0.020)	0.015 (0.017)
Dark volume	18.52 (351.0)	4.48 (131.5)	3.04 (119.2)	8.59 (199.1)
Lit volume	327.8 (2017)	62.18 (717.1)	50.53 (619.7)	145.0 (1109)
Mean dark trade size	1,313	853.3	955.8	1,036
Mean lit trade size	1,280	838.6	888.4	998.8

⁸The large- T dimension in the data ensures the FE-2SLS does not suffer from Nickell bias.

⁹The model was also estimated using a range of lag lengths (between 1 and 30) and the results are consistent across these different specifications.

Consistent estimation of (1) requires that the variables y_{it} and x_{it} are either stationary or co-integrated. The Levin, Lin and Chu (2002) (LLC) stationarity tests contained in table 4 reject the null that the panel data contain a unit root at the 1% level for each dependent variable as well as both lit and dark volume. The LLC test corresponds to the whether the AR parameter in the standard Augmented Dickey Fuller (ADF) test pooled across the i dimension of a given panel is equal or less than zero. It therefore generates a single test statistic that corresponds to whether or not each of the individual time-series in each panel either do or do not contain a unit root and thus whether standard regression techniques are valid. For this application, such a test is advantageous compared with panel stationarity tests where the alternative is that some subset of the time-series in the panel are stationary. Since the dataset has a large- T dimension, it is also possible to conduct individual ADF tests for each stock in the sample. Results from these tests are consistent with table 4 but are omitted in favour of the more easily interpreted LLC test results.

Table 4: Panel unit root tests

t -statistics and p -values from Levin, Lin and Chu (2002) panel stationarity tests for inside spread (%), depth at the BBO, absolute 1s midquote return (%), dark volume and lit volume. The panel for each variable is constructed using the 1s time-series over the 80 trading days in the sample across the stocks in the FTSE-100 index. The null hypothesis is that each individual time-series in a particular panel contains a unit root and the alternative is that each individual time-series is stationary. The test statistics are calculated without a constant or time trend, corresponding to the null hypothesis of a random walk.

Variable	Spread	Depth	Abs. rtn	Dark vlm.	Lit vlm.
t -statistic	-281.2	-418.6	-1,728.9	-2,005.4	-2,365.5
p -value	0.00	0.00	0.00	0.00	0.00
N (stocks)	92				
T (1s)	2,352,080				

Summarised results from the first-stage regressions in table 5 indicate that the instruments satisfy the rank and relevance conditions discussed in section 5.1 for both endogenous regressors, $litvlm_{it}$ and $darkvlm_{it}$. Estimated F -statistics for the joint significance of the instruments in first-stage regressions for both endogenous regressors are well in excess of the 1% critical values for each market capitalisation category. The null that the rank condition is not satisfied is rejected at the 1% level for each regression using the Kleibergen-Paap LM test while Kleibergen-Paap Wald statistics are greater than the 10% maximum size critical values of Stock and Yogo (2005). The parameter estimates on

Table 5: First stage IV regression results

Tests for individual instrument significance, F -tests for joint instrument significance, Kleibergen-Paap RK statistics for underidentification (deficient rank of $E[Z'X]$) and Kleibergen-Paap Wald statistics for weak identification in the first stage regressions of the model $y_{it} = \alpha_i + \sum_{p=0}^{p^{lit}} \theta_p^{lit} litv_{it-p} + \sum_{p=0}^{p^{dark}} \theta_p^{dark} darkv_{it-p} + \sum_{p=1}^Q \phi_p y_{it-p} + \beta' x_{it} + \varepsilon_{it}$ where the contemporaneous variables $litv_{it}$ and $darkv_{it}$ are instrumented using dummy variables, one for whether the time-stamp falls on the turn of the minute (9:00:00AM, 9:01:00AM, 9:02:00AM etc.), and another for time-stamps that fall on even 5 minute intervals (10:05:00AM, 10:10:00AM etc.). The variable y_{it-p} is the market condition variable (inside spread, depth at the BBO and absolute 1s midquote return) for stock i and time period $t-p$ (measured at the 1s frequency across the 80 days in the sample), $litv_{it-p}$ and $darkv_{it-p}$ are the lit and dark volumes traded for stock i and time $t-p$, x_{it} is a vector of exogenous control variables (time of day and its square, date and its square, the return over the five minutes preceding time t and the log of market capitalisation), α_i is a stock-specific fixed effect and ε_{it} is an error term associated with the realisation of y_{it} . Time-stamps that coincide with US data releases (12:30PM), the open of the US market (2:30PM), ONS release times (9:30AM) and Bank of England MPC minutes releases (9:30AM and 11:00AM monthly) are not used in the instrument list. Parameter covariance matrices are clustered at the stock-date level.

Endog. regressor	Dependent variable						
	Spread		Depth		Abs. Rtn		
	Coef.	p-val	Coef.	p-val	Coef.	p-val	
Lit volume	D_t^{lmin}	22.56	0.00	22.35	0.00	22.07	0.00
	D_t^{5min}	68.79	0.00	68.33	0.00	68.33	0.00
	Joint (F -stat)	-	0.00	-	0.00	-	0.00
Dark volume	D_t^{lmin}	15.91	0.00	15.90	0.00	15.88	0.00
	D_t^{5min}	6.78	0.10	6.78	0.10	6.77	0.10
	Joint (F -stat)	-	0.00	-	0.00	-	0.00
Identification tests		Wald		Wald		Wald	
KP RK statistic		211.91		141.91		300.98	
KP Wald statistic		423.82		283.82		601.96	
N (stocks):	92	T (1s.):		2,352,080			

the D_t^{1min} and D_t^{5min} dummy variables suggest that there is a large, statistically significant spike in trade arrivals on the LOB of the primary exchange on the turn of even five minute intervals compared with even one minute intervals, while the same spike is not significant on the dark venues. In other words, there appears to be correlation between preferences of traders using execution algorithms for time slicing at higher frequencies and for sending orders to venues outside the primary exchange.

Table 6: Second stage IV regression results

Second stage coefficient estimates, test statistics and p -values for the the model described in section 5.2 and table 5. The parameter θ_0^i is the contemporaneous (short-run) effect of volume on venue type $i = \{lit, dark\}$ on the relevant dependent variable while γ_{LR}^i is the long-run multiplier following a sustained increase in expected volume on venue type i on the relevant dependent variable, defined in equation (4). The short-run coefficients are scaled up to correspond to a single trade for 5,000 shares, while the long-run multipliers are scaled up to correspond to a change in the expected volume traded per period of 10 shares. Standard errors and the parameter covariance matrices are clustered at the stock-date level and are computed via the delta method for the long-run multiplier.

	Dep. variable								
	Spread			Depth			Abs. Rtn		
	Coef.	t -stat	p -val	Coef.	t -stat	p -val	Coef.	t -stat	p -val
θ_0^{dark}	0.0033	0.87	0.38	2,248	3.06	0.00	0.0004	0.10	0.92
θ_0^{lit}	0.0171	11.19	0.00	680	2.45	0.01	0.0214	13.64	0.00
γ_{LR}^{dark}	-0.0001	0.12	0.90	156.77	6.71	0.00	-0.00002	1.30	0.19
γ_{LR}^{lit}	0.0018	22.74	0.00	-20.60	0.14	0.89	0.00009	45.99	0.00
Hypothesis tests	Wald	p -val		Wald	p -val		Wald	p -val	
$\theta_0^{dark} = \theta_0^{lit}$	7.28	0.01		2.54	0.11		17.04	0.00	
$\gamma_{LR}^{dark} = \gamma_{LR}^{lit}$	6.92	0.01		2.45	0.12		16.66	0.00	

Table 6 contains the contemporaneous parameters and long-run multipliers estimated from the second stage regressions of equation (1) using all 92 stocks in our sample. Figure 4 contains the estimated dynamic multipliers and associated confidence intervals following a single trade on dark or lit venues to the three market condition variables (equation (3)). The short-run coefficients and dynamic multipliers are scaled to represent the immediate and dynamic effects of a single trade for 5,000 shares on either the lit or dark venue on the three market condition variables. The long-run multipliers are scaled to represent the effect on the market condition variables of an increase of 10 shares on the expected volume of trading per period (calculated using equation (4)). Standard errors for the dynamic and long-run multipliers are calculated via the delta method with the clustered

parameter covariance matrices.

The second stage results in table 6 indicate that trading on lit venues has a statistically significant positive, long-lived impact on spreads, depth and volatility. A trade for 5,000 shares on the primary exchange leads to an immediate increase of 1.7 bp on the bid-offer spread and generates an immediate increase of 2.1 bp on the absolute 1s return. Following a 10 share increase in the long-run or expected volume traded per period on the lit venue, spreads increase by approximately 0.2 bp while the absolute return increases by 0.009 bp. There is also a significant contemporaneous effect of lit trading on depth (an increase of 680 shares in depth at the BBO following a trade for 5,000 shares) however the long-run multiplier is negative and insignificant. The dynamic multipliers in figure 4 suggest that a single lit trade has a relatively short-lived effects on depth and volatility - statistically indistinguishable from zero after around 10 seconds - whereas the increase on spreads following a single lit trade is much longer lived.

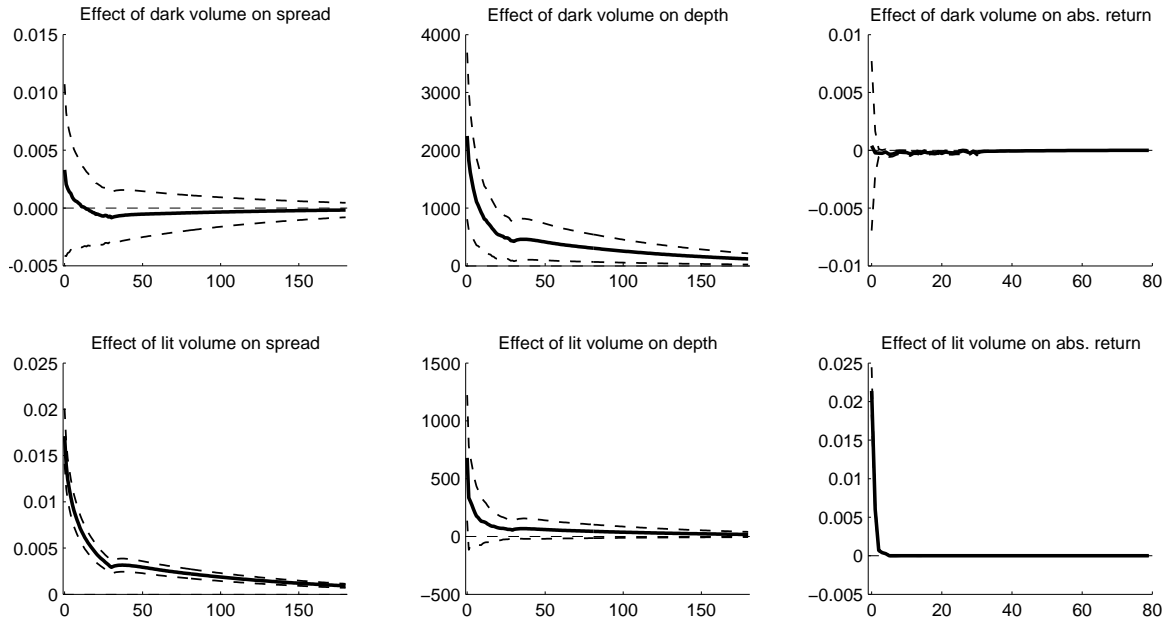
The fact that lit trading consumes liquidity and increases absolute returns (at least in the short-term) accords well with what might be considered reasonable prior beliefs about how trading on limit order books should affect market conditions. Trades on these venues usually involve one party taking liquidity from another (with a market order hitting a resting limit order in the LOB). This in turn reduces the total amount of short-term liquidity available in the market leading to higher spreads. The increase in absolute returns could be driven by greater bid-offer bounce due to higher spreads, or serial correlation in trade arrivals.

In contrast, trades on dark venues do not lead to increased spreads or volatility in the short-run or the long-run but do lead to a sustained, significant increase in depth available at the best bid and offer at the primary exchange. Following a dark trade for 5,000 shares, depth on the primary exchange increases by approximately 2,250 shares (approximately three times more than for the equivalent lit trade), while a 10 share increase in the long-run or expected volume traded per period leads to long-run increase of approximately 160 shares in the depth on the primary exchange. The dynamic multiplier for depth following a single dark trade (figure 4) remains positive and significant for more than 120 seconds. The corresponding dynamic multipliers for lit trading is statistically indistinguishable from zero within 10 seconds of an equivalent lit trade. The data reject the null hypothesis that the short-run and long-run effects of trading on lit venues and dark venues on spreads and absolute returns are equal at the 1% level although equivalent tests for depth do not reject the null of equal effects at the 10% level.

Taken together, these results suggest that levels of dark trading in the UK equity

Figure 4: Dynamic multipliers

Dynamic multipliers derived from the second stage coefficient estimates for the the model described in section 5.2 and table 5. The multipliers are scaled up to correspond to the contemporaneous and dynamic effect of a single trade for 5,000 shares on either dark or lit venue on the dependent variables. Standard errors for the multipliers are derived using the delta method with the parameter covariance matrix clustered at the stock-date level.



market observed over the sample period were not obviously harmful to market quality and integrity. Trades on dark venues lead to a sustained improvement in depth at the best bid and offer without affecting spreads or volatility in the short-run or long-run, indicating that liquidity on the primary exchange actually improves following dark trades in absolute terms. That dark trades lead to improved liquidity in the primary exchange without affecting volatility is consistent with the hypothesis that uninformed, liquidity traders are more likely to be present on these venues compared with traders with informational advantages, as predicted by the theoretical models of both Hendershott and Mendelson (2000) and Zhu (2014). Although in these models “cream skimming” of uninformed order flow to dark venues and traders using the lit venue as a “market of last resort” raises costs for market makers on the primary exchange, section 4 demonstrated that the majority of large market making institutions such as investment banks are actually very active across the dark venues in the sample, in contrast to the assumptions of these models. Therefore, despite the migration of uninformed traders away from the primary exchange, market makers are still able to realise the benefits of trading with uninformed traders, such as reduced overall adverse selection costs, by trading simultaneously on both lit and dark venues. Market makers may also infer that the probability of informed

trading is lower, *ceteris paribus*, following periods of more active trading in dark venues. Furthermore, market makers would not be directly exposed to adverse selection costs when clearing positions obtained by trading with uninformed traders in dark venues.

5.3 Intraday trading patterns as instrumental variables

Although the instruments used in this section do deal with simultaneity from expected market conditions to trade participation decisions, they are not truly randomised sources of variation in trading volumes. As such there are a number of potential issues that need to be considered when interpreting the results, aside from the need to remove time-stamps that coincide with data releases and other regular events during the day from the instruments. For example, if the patterns are driven by a small group of traders with similar characteristics, then the extent to which the results above can be generalised is not clear. The results may unduly reflect idiosyncrasies in the trading behaviour of this subset of firms and not be easily applied to the wider trading population. The decision to trade by the participants using the low technology algorithms may also be based on low frequency expectations about future market conditions. For example, these traders may make the decision to participate in trading on a given day based on their expectations about volatility or liquidity for that day. If this is the case, the instrument might unduly weight days with worse liquidity and higher volatility or vice versa. These two issues are addressed below.

Table 7: Trader compositions on even vs. uneven time-stamps

Mean fraction of trading accounted for by three trader types (HFTs, IBs and broker) on even one minute time-stamps (9:00:00AM, 9:00:01AM etc.) and uneven time-stamps (9:00:01AM, 9:00:59AM etc.). The averages are calculated across each stock-day for the constituents of the FTSE-100. The tests for equal means do not assume equal variance.

	Dark venues			Lit venues		
	HFT	IB	Brokers	HFT	IB	Brokers
Mean fraction: even time-stamps	0.0928	0.6623	0.2449	0.1868	0.7555	0.0577
Mean fraction: uneven time-stamps	0.1132	0.5333	0.3535	0.1900	0.7510	0.0589
<i>t</i> -stat	-7.31	25.59	-22.86	-2.15	2.53	-1.13
<i>p</i> -value	0.00	0.00	0.00	0.03	0.01	0.26

Firstly, I examine the composition of the counterparties to trades that fall on even one minute intervals with the composition of traders that trade on other time-stamps within the minute. Using the same three trader classifications as section 4 (HFT, IBs

and brokers), I calculate the fraction that each group accounts for of total trades on the turn of the minute and on other time-stamps across all stock-days used for the regressions in this section. The average fraction of trading by each group on both dark and lit venues are contained in table 7.

For trades on either type of venue, investment banks account for a relatively greater proportion of trades on even time-stamps compared with uneven time-stamps. For lit venues, this difference is less than 0.5%, as are the differences in the corresponding fractions of trading executed by HFTs and brokers. However for dark venues, investment banks account for approximately 13% more of the trading that occurs on even time-stamps compared with uneven time-stamps, while brokers account for approximately 10% less trading during these periods. Considering the diverse range of trading interests that investment banks represent, such as client order flow from their brokerage businesses, proprietary trading, asset management etc., it is arguably of less concern to have such a group over-represented on even time-stamps compared with the other trader categorisations.

Table 8: Effect of dark trader compositions on outcome variables

Parameter estimates for regressions of the outcome variables (spreads, depth and absolute 1s return) on the fraction of broker trading and fraction of IB trading on dark venues lagged by one period, as well as all control variables from the model described in section 5.2 and table 5 (parameter estimates for the additional control variables are omitted from the table but are available upon request). The sample for these regressions contains all periods where at least one dark trade occurred in the previous period (so as the fraction variables are defined) across all 92 stocks. Standard errors are unadjusted for heteroskedasticity and clustering and as such can be considered estimates of the lower bounds of the true standard errors (similar results were obtained using heteroskedasticity-robust and clustered standard errors at the stock level).

	Dep. variable								
	Spread			Depth			Abs. Rtn		
	Coef.	<i>t</i> -stat	<i>p</i> -val	Coef.	<i>t</i> -stat	<i>p</i> -val	Coef.	<i>t</i> -stat	<i>p</i> -val
Broker fraction	0.00002	0.52	0.61	16.32	1.28	0.20	0.00002	0.76	0.45
IB fraction	0.00002	0.85	0.40	16.31	1.44	0.15	-0.00001	-0.5	0.61
Joint test <i>p</i> -val	0.69			0.33			0.24		

To more rigorously discern whether these changes unduly influence the results in section 5.2, I test whether the outcome variables are sensitive to the composition of traders in dark pools at the frequency of the analysis used in this paper. To do so, I regress the outcome variables on the first-order lag of the fraction of dark trading accounted for by IBs and brokers for all periods and stocks where this fraction is defined

(i.e. those with at least one dark trade occurring in the previous period).¹⁰ I also include all exogenous variables that are present in the first and second stage regressions (lags of lit and dark volume, lags of the endogenous variables and the control variables in x_{it}) and stock-specific fixed effects, and calculate unadjusted standard errors as a lower bound on the true standard errors. Even using these conservatively small standard errors, these regression demonstrate that the fraction of dark trading by either brokers or IBs is insignificant for predicting the next period outcome variables at the 1s frequency, even at the 10% level. Joint tests of the null that both parameters equal zero also fail to reject, even at the 10% level. So although IBs (brokers) are more (less) prevalent on even time-stamps, these changes in relative activity do not have a significant effect on the outcome variables. In other words, the results in section 5.2 do not appear to be driven by changes in trader composition, and related heterogeneous effects, such as possible changes in the average information content of trades on even vs. uneven time-stamps.

Table 9: Parameter estimates for low frequency participation decision

Parameter estimates for the regression model $evenfrac_{it}^j = \alpha_i + \beta_1^j \hat{sprd}_{it} + \beta_2^j \hat{depth}_{it} + \beta_3^j \hat{rv}_{it} + \varepsilon_{it}^j$ where $evenfrac_{it}^j$ is the daily fraction of trades that occur on even one minute time-stamps (9:00:00AM, 9:01:00AM etc.) compared with all other time-stamps for stock i , day t and venue type $j = dark, lit, both$ and \hat{sprd}_{it} , \hat{depth}_{it} and \hat{rv}_{it} are the fitted values for daily frequency AR(1) regressions of quoted spreads, log of depth and 5min realised volatility respectively (with a stock-specific intercepts). Standard errors are unadjusted for heteroskedasticity or clustering and as such can be considered estimates of the lower bounds of the true standard errors. Separate regressions are run for the fraction of even trades on dark venues, lit venues and both venues combined.

	Dark venues			Lit venues			Both venues		
	Param.	Std. Err	<i>t</i> -stat	Param.	Std. Err	<i>t</i> -stat	Param.	Std. Err	<i>t</i> -stat
\hat{sprd}_{it}	-0.0517	0.0551	-0.94	-stat	0.0254	-0.81	-0.0176	0.0245	-0.72
\hat{depth}_{it}	-0.0003	0.0002	-1.24	-0.0002	0.0001	-1.86	-0.0002	0.0001	-1.77
\hat{rv}_{it}	-0.0069	0.1038	-0.06	-0.0205	0.0478	-0.43	-0.022	0.0461	-0.48

Secondly, I examine whether daily participation decisions by traders who use low technology algorithms are affected by expectations about future market conditions. I estimate a simple AR(1) model with stock-specific intercepts for the three dependent variables (spreads, depth and realised volatility) using daily data. The fitted values from these regressions are then used as estimates of expected market conditions for the three dependent variables, and I test whether or not these forecasts have explanatory power for

¹⁰The fraction of trades accounted for by IBs, brokers and HFTs sums to one for more than 99% of the observations, and so to avoid collinearity issues, the fraction of HFT trades is omitted from the regression. As expected, replacing the broker or IB fraction with the HFT fraction does not change the interpretation of these results.

the total fraction of trading that occurs at even time-stamps at the stock-day level. The results from these regressions including detailed regression specifications are contained in table 9. These regressions indicate that the predicted values for the market condition variables are insignificant for the fraction of trading that occurs at even time-stamps, both on dark and lit venues, and also for trading on all venues. The standard errors for these regressions are the standard OLS estimates and hence most likely underestimate the true standard errors, due to clustering in either the stock or time dimensions. It follows that the t -statistics most likely overestimate the significance of the predicted market condition variables in the regressions, or in other words are upper bounds on the actual t -statistics. As such, table 9 suggests that low frequency expectations about future market conditions do not play a significant role in determining the participation decisions of users of low technology algorithms.

6 Dark trading, market quality and stock size

The results presented in the preceding section indicate that for the sample pooled across all stocks in the FTSE-100 index, trading in dark pools does not lead to a short-term deterioration of market quality on the primary exchange. However, this relationship may differ between large, liquid, actively traded stocks and for smaller, less frequently traded stocks. For example, the theory papers of Degryse et al. (2009) and Buti et al. (2014) predict that dark trading has different impacts on market quality for more and less liquid stocks while O’Hara and Ye (2011) and Degryse et al. (2014) both find empirical evidence that overall market fragmentation (across lit and dark venues) affects large and small stocks in different ways.

To investigate whether the relationship between dark trading and intraday market quality on the primary exchange differs across stocks of different size, I split the sample into three groups by market capitalisation. Group 1 (“Large”) contains the 30 largest stocks by market capitalisation in the sample. Group 2 (“Mid-sized”) uses the next 32 largest stocks by market capitalisation. Group 3 (“Small”) contains the 30 smallest stocks in the sample by market capitalisation. The panel-IV regressions from section 5.2 are run for each of the three groups and the second stage estimates as well as the first stage identification tests are presented in table 10.

Table 10: Second stage parameters and identification tests by stock size

Summarised results for the model described in section 5.2 and table 5 estimated for large stocks (1-30 by market cap.), mid-sized stocks (31-62 by market cap.) and small stocks (63-92 by market cap.). Parameter definitions, scaling factors and computation of standard errors are all described in table 6. The “RK” and “Wald” identification tests refer to the Kleibergen-Paap RK statistic and the Kleibergen-Paap Wald test for weak identification respectively.

		Dep. variable								
		Spread			Depth			Abs. Rtn		
Category		Coef.	<i>t</i> -stat	<i>p</i> -val	Coef.	<i>t</i> -stat	<i>p</i> -val	Coef.	<i>t</i> -stat	<i>p</i> -val
Large	θ_0^{dark}	-0.0035	-0.89	0.37	28	0.02	0.98	-0.0016	-0.42	0.67
	θ_0^{lit}	0.0108	7.62	0.00	1,042	2.61	0.01	0.0128	8.83	0.00
	γ_{LR}^{dark}	-0.0004	2.35	0.02	-6.17	0.01	0.99	-0.00002	1.30	0.20
	γ_{LR}^{lit}	0.0009	13.50	0.00	61.35	1.62	0.11	0.00006	19.67	0.00
Mid-sized	θ_0^{dark}	0.0069	1.09	0.28	4,326	3.91	0.00	0.0124	2.63	0.01
	θ_0^{lit}	0.0278	9.45	0.00	510	1.14	0.26	0.0300	13.10	0.00
	γ_{LR}^{dark}	0.0000	0.00	1.00	545.42	12.04	0.00	0.00003	4.57	0.00
	γ_{LR}^{lit}	0.0032	24.66	0.00	-107.04	1.03	0.30	0.00011	78.35	0.00
Small	θ_0^{dark}	0.0343	3.70	0.00	5,700	4.40	0.00	0.0101	1.30	0.19
	θ_0^{lit}	0.0262	5.68	0.00	-85	-0.13	0.90	0.0407	8.94	0.00
	γ_{LR}^{dark}	0.0041	12.00	0.00	999.15	14.52	0.00	0.00002	0.81	0.42
	γ_{LR}^{lit}	0.0027	5.44	0.00	-375.11	3.06	0.00	0.00015	40.05	0.00
Hypothesis tests		Wald	<i>p</i> -val		Wald	<i>p</i> -val		Wald	<i>p</i> -val	
Large	$\theta_0^{dark} = \theta_0^{lit}$	7.62	0.01		0.43	0.51		7.88	0.00	
	$\gamma_{LR}^{dark} = \gamma_{LR}^{lit}$	6.95	0.01		0.44	0.51		7.82	0.01	
Mid-sized	$\theta_0^{dark} = \theta_0^{lit}$	5.54	0.02		6.65	0.01		7.20	0.01	
	$\gamma_{LR}^{dark} = \gamma_{LR}^{lit}$	5.19	0.02		6.48	0.01		7.23	0.01	
Small	$\theta_0^{dark} = \theta_0^{lit}$	0.38	0.54		9.66	0.00		6.87	0.01	
	$\gamma_{LR}^{dark} = \gamma_{LR}^{lit}$	0.39	0.53		8.84	0.00		7.05	0.01	
Identification tests		RK	Wald		RK	Wald		RK	Wald	
Large		94.48	188.96		207.35	414.71		76.50	153.00	
Mid-sized		55.69	111.38		163.46	326.92		71.53	143.06	
Small		115.58	231.15		249.65	499.30		129.22	258.45	

The results in table 10 demonstrate that the model is still well identified for each subgroup, and that the effect of dark trading on intraday market quality does differ by stock size, with improvement in depth following dark trades most apparent for mid-sized and small stocks. For larger stocks, trades on dark venues lead to a statistically significant reduction in spread of approximately -0.04 bp following an increase in the expected volume traded on dark venues of 10 shares per period. The corresponding long-run parameter for lit trades is a statistically significant increase of 0.09 bp. For this cohort of stocks, dark trading does not have a significant impact on depth or the absolute return either in the short-run or the long-run.

For mid-sized stocks, dark trading does not have the same significant effect on spreads but does lead to a significant increase in depth in both the short-run and long-run, with the size of the effect of similar magnitude to the results from the full sample in table 6. For this group of stocks, dark trading has a greater impact on the subsequent price process (absolute return) than the full sample, although this effect is still significantly smaller than the corresponding responses to lit trading.

For small stocks, dark trades also lead to significant increases in depth in both the short-run and the long-run, with the magnitude of this effect larger than for mid-sized stocks. However, since the average trade size for smaller stocks is typically smaller than for large or mid-sized stocks, it is perhaps not surprising that a trade for a given number of shares has a larger total impact for this category. Unlike mid-sized stocks, there is no significant relationship between absolute returns and dark trading, however dark trading does lead to a significant increase in spreads in the short-run and long-run. The effect of dark trading on spreads is not statistically distinguishable from the effect of lit trading.

The results in table 10 confirm those obtained using the same sample. For all three categories of stocks, dark trading does not lead to a statistically significant worsening of market quality on the primary exchange. For large stocks, the effect of dark trading is negligible, other than some evidence that increased equilibrium trading in dark pools may lead to tighter spreads on the main exchange. For mid-sized and small stocks, dark trading has either no effect on spreads and volatility, a significantly smaller effect to that of lit trading, or an effect that is indistinguishable from lit trading. Dark trading does however lead to improved depth both in the short-run and the long-run, in absolute terms and relative to the effect of lit trading. In no cases do the results suggest that dark trades have a detrimental impact on market quality compared with lit trading.

These regressions provide some support for the predictions of Buti et al. (2014) whereby competition between a crossing network and a LOB widens spreads more for

less liquid stocks compared with more liquid stocks. This is because dark pools attract relatively more market orders to limit orders from the LOBs of more liquid stocks, compared with less liquid stocks, leading to relatively wider spreads in the latter category. The long-run estimates of the effect of dark trading on spreads for both large and small stocks are consistent with these predictions, although the short-run impact of dark trading on spreads is insignificant for the largest stocks.

7 Conclusion

Equity trading on dark pools is a high priority issue for securities market regulators both in the US and EU. This paper contributes to the nascent literature examining how dark trading affects market quality and market integrity by estimating this relationship at the transaction level. The analysis in this paper indicates that the scale of dark trading during the sample period did not pose a threat to the short-term market quality of the UK equity market. Using transaction data from the FCA's ZEN database, I show that trades on dark pools lead to improved high frequency liquidity (in terms of depth at the best bid and offer), both in absolute terms and relative to trades on the limit order book of the primary exchange, without affecting volatility of midquote returns or the quoted spread. These results are consistent with the hypothesis that uninformed, liquidity traders cluster in the dark pool where lower trading costs for end-users can help to facilitate new trading activity. Importantly, by using intraday trade patterns generated by unsophisticated execution algorithms, I am able to make inference about the effect of dark trading on market conditions despite endogeneity generated by rational traders using expectations about future market conditions to make trading decisions.

The key contribution of this paper is the use of a novel identification strategy with a very detailed proprietary dataset covering multiple dark pools to identify the effect of dark trading in the UK equity market at a high frequency. Previous papers in this field use data aggregated to the daily level or higher, whereas the rich detail in my data allows me to drill down to a much higher frequency (1 second) and to examine how market conditions respond to dark trading at these frequencies. These high frequency relationships are of growing importance in modern, fragmented financial markets where computer-based trading is responsible for the majority of activity (Linton, O'Hara and Zigrand, 2012). To the best of my knowledge, this is the first paper to use intraday patterns in trade arrivals as an instrument for trading volume and this technique could be applied to a wide variety of other datasets in a similar manner.

The main conclusion of this paper, that the levels of dark trading in the UK equity market from September to December 2012 did not pose a threat to short-term market quality, is consistent with empirical work examining dark trading and fragmentation at lower frequencies including O’Hara and Ye (2011), Foley and Putniņš (2015), Korber et al. (2013) and Comerton-Forde and Putniņš (2014). My results suggest that the main high frequency channel through which dark trading can lead to improved market quality is via liquidity rather than directly through volatility. That the majority of evidence from the academic literature tends to find a positive or insignificant effect of low levels of dark trading (less than 10% of total trading) on equity market quality contrasts with current trends by regulators. This especially is the case in Europe, where transparency is being promoted as a key component of financial market integrity (see e.g. Barnier (2014) and Financial News (2014)). Indeed, the results of this paper, as well as those of the majority of empirical papers in this field, conclude that low levels of dark trading do in fact lead to improvements in market-wide liquidity without harming price efficiency. With this in mind, it may be worthwhile for regulators and academics to engage more closely on this issue to ensure that the evidence being generated by the latter directly addresses the concerns of the former and helps to form evidence-based policy recommendations.

References

- Michel Barnier. New rules on markets in financial instruments (MiFID II): Statement by Commissioner Barnier following the agreement in trilogue, January 2014. http://europa.eu/rapid/press-release_MEMO-14-15_en.htm.
- James Brugler and Oliver B Linton. Single stock circuit breakers on the london stock exchange: do they improve subsequent market quality? *Available at SSRN 2379029*, 2014.
- Markus K Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.
- Sabrina Buti, Barbara Rindi, and Ingrid M Werner. Dark pool trading strategies, market quality and welfare. *Available at SSRN 1630499*, 2014.
- Carole Comerton-Forde. Current issues in market design. *Market Microstructure in Emerging and Developed Markets: Price Discovery, Information Flows, and Transaction Costs*, pages 177–212, 2013.
- Carole Comerton-Forde and Tālis J Putniņš. Dark trading and price discovery. *Available at SSRN 2183392*, 2014.
- Hans Degryse, Mark Van Achter, and Gunther Wuyts. Dynamic order submission strategies with competition between a dealer market and a crossing network. *Journal of Financial Economics*, 91(3):319–338, 2009.
- Hans Degryse, Geoffrey Tombeur, and Gunther Wuyts. Dark trading. *Market Microstructure in Emerging and Developed Markets: Price Discovery, Information Flows, and Transaction Costs*, pages 213–230, 2013.
- Hans Degryse, Frank De Jong, and Vincent Van Kervel. The impact of dark trading and visible fragmentation on market quality. *Review of Finance*, 2014.
- David Easley, Marcos Lopez de Prado, and Maureen O’Hara. The volume clock: Insights into the high frequency paradigm. *The Journal of Portfolio Management*, (Fall, 2012) *Forthcoming*, 2012.
- European Commission. Markets in Financial Instruments Directive (MiFID II): Frequently Asked Questions, 2014. http://europa.eu/rapid/press-release_MEMO-14-305_en.htm?locale=en.

Financial Conduct Authority. Royal Bank of Scotland fined £5.6m for failing to properly report over a third of transactions, 2013. <http://www.fca.org.uk/news/rbs-fined>.

Financial News. Esma's Mifid II takeaways, May 2014. <http://www.efinancialnews.com/story/2014-05-22/esma-mifid-ii-key-takeaways-dark-pools-hft-data-fees-open-access>.

Sean Foley and Tālis J Putniņš. Should we be afraid of the dark? Dark trading and market quality. *Available at SSRN 2279719*, 2015.

Ruslan Y Goyenko, Craig W Holden, and Charles A Trzcinka. Do liquidity measures measure liquidity? *Journal of Financial Economics*, 92(2):153–181, 2009.

Carole Gresse. Effects of lit and dark market fragmentation on liquidity. *Available at SSRN 1918473*, 2015.

Terrence Hendershott and Haim Mendelson. Crossing networks and dealer markets: Competition and performance. *The Journal of Finance*, 55(5):2071–2115, 2000.

Lena Korber, Oliver Bruce Linton, and Michael Vogt. The effect of fragmentation in trading on market quality in the UK equity market. *Available at SSRN 2315484*, 2013.

Andrew Levin, Chien-Fu Lin, and Chia-Shang James Chu. Unit root tests in panel data: Asymptotic and finite-sample properties. *Journal of Econometrics*, 108(1):1–24, 2002.

Oliver Linton, M O'Hara, and JP Zigrand. Economic impact assessments on mifid ii policy measures related to computer trading in financial markets. *Foresight, Government Office for Science Working Paper*, 2012.

Ananth N Madhavan. VWAP strategies. *Trading*, 2002(1):32–39, 2002.

Hitesh Mittal. Are you playing in a toxic dark pool? A guide to preventing information leakage. *The Journal of Trading*, 3(3):20–33, 2008.

Maureen O'Hara and Mao Ye. Is market fragmentation harming market quality? *Journal of Financial Economics*, 100(3):459–474, 2011.

Reuters. Exclusive: SEC eyes test that may lead to shift away from "dark pools", April 2014. <http://www.reuters.com/article/2014/04/11/us-sec-darkmarkets-idUSBREA3A0CP20140411>.

James H Stock and Motohiro Yogo. Testing for weak instruments in linear IV regression. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 2005.

The TRADE Magazine. The 2012 Algorithmic Trading Survey, 2012. www.thetradenews.com/uploadedFiles/T31%20Algo%20survey.pdf.

The Wall Street Journal. SEC Set to Spur Exchange Trading, May 2014. <http://online.wsj.com/news/articles/SB10001424052702304811904579586300158560642>.

Dimitri Vayanos. Flight to quality, flight to liquidity, and the pricing of risk. *NBER Working Paper*, (w10327), 2004.

Daniel G Weaver. The trade-at rule, internalization, and market quality. *Available at SSRN 1846470*, 2014.

Mao Ye. Price manipulation, price discovery and transaction costs in the crossing network. *Working Paper, University of Illinois*, 2012.

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