

Asymmetries in Dark Pool Reference Prices

September 2016

Occasional
Paper 21

Encouraging debate among academics, practitioners and policymakers in all aspects of financial regulation.

Financial Conduct Authority
www.fca.org.uk

FCA occasional papers in financial regulation

The FCA occasional papers

The FCA is committed to encouraging debate on all aspects of financial regulation and to creating rigorous evidence to support its decision-making. To facilitate this, we publish a series of Occasional Papers, extending across economics and other disciplines.

Occasional Papers contribute to the debate on specific issues relevant to the FCA's work. The main factor in accepting papers is that they should make substantial contributions to knowledge and understanding of financial regulation. If you want to contribute to this series or comment on these papers, please contact Peter Andrews or Kevin James at peter.andrews@fca.org.uk and kevin.james@fca.org.uk

Disclaimer

Occasional Papers contribute to the work of the FCA by providing rigorous research results and stimulating debate. While they may not necessarily represent the position of the FCA, they are one source of evidence that the FCA may use while discharging its functions and to inform its views. The FCA endeavours to ensure that research outputs are correct, through checks including independent referee reports, but the nature of such research and choice of research methods is a matter for the authors using their expert judgement. To the extent that Occasional Papers contain any errors or omissions, they should be attributed to the individual authors, rather than to the FCA.

Authors

Matteo Aquilina, Sean Foley, Peter O'Neill and Thomas Ruf.

Matteo Aquilina and Peter O'Neill work in the Chief Economist's Department of the Financial Conduct Authority.

Sean Foley is an Assistant Professor at the University of Sydney and Thomas Ruf is an Assistant Professor at the University of New South Wales.

Acknowledgements

We would like to thank Albert Menkveld for providing advice and reviewing the paper. We also thank Michael Aitken, Peter Andrews, Hao Ming Chen, Brian Eyles, Simon Hargreaves, Edwin Schooling Latter, Ted Macdonald, Richard Payne, Tālis Putniņš, Jia Shao, Patrick Spens, Felix Suntheim, Martin Taylor, Carla Ysusi, Bart Z. Yueshen and many other FCA colleagues. We are also grateful to participants of the FCA/LSE conference on financial regulation for their comments.

Peter O'Neill thanks the Capital Markets Cooperative Research Centre for funding a period as a visiting researcher at the FCA during this study.

Any errors and omissions are our own.

All our publications are available to download from www.fca.org.uk. If you would like to receive this paper in an alternative format, please call 020 706 60790 or email publications_graphics@fca.org.uk or write to Editorial and Digital Department, Financial Conduct Authority, 25 The North Colonnade, Canary Wharf, London E14 5HS.

Contents

1	Overview	6
	Purpose	6
	Key findings	7
2	Research context	9
	Dark Pools and How They Work	9
	Are Dark Pool Reference Prices Reliable?	10
	Is the Primary Market a Good Reference Price?	13
3	Method and approach	15
	Data	15
	Identifying Stale Reference Prices	17
	Dark Trading when the Primary Market has a Worse Price	19
4	Results	21
	Prevalence of Stale Reference Prices	21
	What Causes Stale Reference Prices?	27
	Dark Trading When Primary Market has Worse Price	27
5	Conclusions	32

Summary

A 'dark pool' is a trading venue with no pre-trade transparency. While in 'lit' venues market participants can observe the orders submitted by other participants, in dark pools, all orders are hidden. The main advantage of submitting an order to a dark pool is that the trade intention is not revealed to the entire market. Another potential advantage is getting a better price than that available on the lit market, as many dark pools match trades at the midpoint, allowing participants to save half the spread. The main disadvantage of dark pools is execution uncertainty. Specifically, it is impossible to know whether there is a willing counterparty, so one cannot know beforehand whether a trade will take place.

Dark pools have existed since the 1980s but have gained importance in recent years. Given their rise, academics and regulators have turned their attention to them and focused on how such venues effect overall market efficiency. The impact of the level of dark trading on price discovery and the informativeness of prices have been the main issues analysed in the literature.

However, conduct issues that focus on the reference price reliability in these venues are also important. First, it is pivotal to know whether the current market infrastructure delivers reliable prices and the level of detriment that may be present when they are not. Second, this is important for best execution considerations. Third, by analysing whether unreliable price effects are randomly distributed (or else) across market participants we can improve our understanding of speed's importance in modern financial markets. Finally, any lack of reference price reliability may be perceived as a deterioration in 'fairness' in modern markets. This could cause investors to reduce their participation in such markets with obvious implications for market quality and macroeconomic performance.

In this study, we analyse two important aspects of reference prices in dark pools. First, we examine the prevalence of trades at *stale* reference prices¹, their costs and their impact on different market participants. Second, we investigate questions concerning the *choice* of reference price: to what extent are participants implementing best execution practices when a dark pool references a worse price than the lit market? Is this influenced by conflicts of interest within dark pools and participant sophistication?

Our main findings when analysing stale reference prices are:

- Dark pool reference prices are sometimes stale in every dark pool in our sample. We find that 3.54% of all dark midpoint trades in our sample reference a stale price. This proportion is increasing over time, from 3.36% in 2014 to 4.05% in June 2015. This increase can be explained by increases in message volumes and volatility over the sample.
- We estimate that the cost of stale reference prices is approximately £4.2m per year across all UK dark venues. The figure does not appear to be economically significant. For comparison, the average daily order book equity value traded on the London Stock Exchange (LSE) in London is £4.9bn.
- All dark trades at stale reference prices are executed at a price that does not match the primary market midpoint² during the trade. One counterparty benefits from this, either paying less or receiving more for the trade than they would otherwise. If latency affects

¹ A stale reference price is not the most recent price. For dark pools, this means a reference price that has been superseded by a newer price that has not yet reached the dark pool. See our detailed definition, including time thresholds in Section 3 (pp. 17–18).

² The midpoint is the average (middle) price of the best ask (sell) price and the best bid (buy) price on a market.

participants equally, then we expect equal outcomes across participant types. This is not what we find: in 96% of cases, High Frequency Traders (HFT)³ are on the benefiting side of the trade.

Our main findings when analysing price dislocations are:

- A small percentage (0.57%) of dark trades occur that reference the LSE bid/ask when it is worse (from the trade initiator's perspective) than another lit market. This is much smaller than the average percentage of the time that dislocations are present on the lit market, which is around 33% of the time. Overall, participants must have smart order routers that observe and react to prices effectively most of the time.
- More midpoint trades (1.22%) occur when the price of the LSE midpoint is worse (from the trade initiator's perspective) than the Best Bid or Offer (BBO) of another lit market. This is roughly comparable to the percentage of the day we observe these dislocations, perhaps implying that participants are not as cautious with reference prices at the midpoint, assuming price improvement will occur regardless. The costs of this happening are very small, mainly because the price difference between the LSE and other venues is usually minor.
- When analysing the distribution of such costs, however, more sophisticated participants obtain better execution outcomes. Participants obtaining the best outcomes are venue operators themselves and HFTs. Non co-located participants execute 18 times as many trades at worse prices than those available in the lit venue than do dark venue operators. We note that the economic magnitude of this is small, at less than 3 basis points on average⁴.

Overall, we find asymmetric outcomes across participants when the reference price is stale, and when it is inferior to other prices available. This may result from participants' differing abilities to observe and manage latency, and differing abilities to engage in effective smart order routing in a fragmented market. These costs are more substantially borne by participant types that are less capable of managing them. However, it is likely that these outcomes are the result of individual participant decisions on the basis of their own analysis of costs and benefits of investment in reduction in latency. In addition, while the effects are highly statistically significant across participant types, the economic impacts are small.

Dark pools thus may still offer a valuable service to market participants, as in most cases they provide price improvement and in all cases allow investors not to show their hand to the market.

³ HFT is an acronym for High Frequency Traders, which refer to participants that use proprietary capital to generate returns using computer algorithms and low-latency infrastructure. See our detailed definition of HFT in Section 3.

⁴ See Table 4 on page 29.

1 Overview

Purpose

This section summarises our study's motivation and main results. Subsequent sections explain how trading in dark pools works, as well as details of the methodology and results.

Dark pools have existed since the 1980s but only recently have they comprised a significant share of the equities market. This has been steadily increasing.⁵ Dark pools differ from 'lit' markets in offering no pre-trade transparency, matching orders anonymously. Dark pools must⁶ reference prices from other venues to determine execution prices. Practically, this involves a primary market data feed from another market to the dark pool.

Two forms of delay or latency in referencing this primary market data exist, resulting in potential costs to investors; processing latency and transmission latency. Where markets are in the same physical location or data centre, a delay exists from the hardware and software processing times involved with calculating and disseminating the market data (*processing latency*). When these two markets are in different physical locations, the time it takes to transmit these data creates an additional delay (*transmission latency*).

In this study, we analyse two important aspects of reference prices in dark pools.

First, we examine reference price latency's prevalence. We attempt to answer the following questions: what is the probability of a dark trade occurring at a stale price? Has this changed over time? What are the causes of (processing) latency?

Second, we assess primary market choice (the LSE in our case) as the reference price's source. We analyse instances in which markets other than the LSE have a better price available.

For both aspects, we then measure the effect for different classes of market participants to analyse the role of participant speed and sophistication in driving outcomes in today's markets.

⁵ Dextrixhe, 2016, 'European Dark Pools Expand, Spiting Regulators' Ambitions,'

www.bloomberg.com/news/articles/2016-01-14/european-dark-pools-expand-in-face-of-rules-limiting-their-use

⁶ This refers to the vast majority of dark pools and dark pool trades by value which is not 'Large In Scale' so must rely on the 'Reference Price Waiver' to enable dark trading.

Key findings

Dark pool reference prices are sometimes stale

Dark pool reference prices are sometimes stale in every dark pool in our sample. 3.54% of dark midpoint trades in our sample across pools are referencing a stale price. The level of stale prices is 11.5% of trades in the dark pool with the highest prevalence.

90% of all stale reference prices are 6 milliseconds or less in duration, but the top 5% are above 20 milliseconds and the top 1% are above 217 milliseconds. All stale reference price events are long enough for an algorithm in the market to observe and act on, but the top 1% is long enough to be perceived by human traders.

We estimate the economic impact of stale reference prices as approximately £453,000 per year in the venues for which we have data, and over £4.2m across all dark venues.⁷ This figure does not appear to be economically significant. For comparison, the average daily order book equity value traded on the LSE in London is £4.9bn.

The prevalence of stale prices is increasing over time

The highest proportion of stale dark pool trades were in the recent sample period, June 2015. We find 4.05% of dark trades are stale in June 2015 compared with 3.36% in 2014.⁸

The age of the stale reference price is also increasing across the sample, from a median of 2 milliseconds to 3 milliseconds, with the oldest 25% rising from 10 milliseconds to 57 milliseconds, in 2014 and 2015 respectively. Increased message traffic and volatility are the most likely explanations.

We observe the amount of order messages on the UK lit markets and the primary market around stale reference prices. We find a statistically significant, positive relationship between trades at stale reference prices and increases in message levels.

Stale price costs fall disproportionately on higher latency participants

We find that HFT participants are on the profitable side of stale trades 96% of the time while co-located⁹ participants are on the losing side 88% of the time, and non co-located 91% of the time.

Higher latency participants execute at inferior reference prices more often

For trades executed with stale reference prices, we find that, in broker operated multilateral trading facilities (MTFs), the venue operators, HFTs, and co-located participants avoid executing when the LSE has a worse price than that available on the lit market far more often than do other participants.

Implications for connected markets

As markets have fragmented, the connections among them have become more important. We provide evidence of adverse outcomes for market participants when one set of connections is affected by latency. Our work demonstrates that a millisecond is a long time in modern markets and latency has a significant role in determining participants' outcomes.

⁷ In Canada, IIROC estimates the costs at \$748,188 CAD per year, approximately £438,000 at time of writing. As IIROC collected data from all Canadian dark venues, this is lower than our extrapolated figure of £4.2m across all UK dark venues. In Australia, ASIC measures the cost of dark pool reference price latency at around \$290,000 AUD a year, approximately £169,000. Although ASIC only measure the costs for a subset of dark trades (those that happen outside the NBBO).

⁸ We examine factors that may explain this increase over time, such as volatility, and messages volumes in Annex 2.

⁹ Co-location refers to the placement of a market participant's servers in close physical proximity to an exchange's to reduce transmission latency.

Some exchanges have proposed solutions, such as Turquoise's 'well formed market check' that suspends executions from stale reference price feeds, and IEX's inbound speed-bump which delays aggressive orders by 350 microseconds, but not the reference price feed.

Implications for exchange infrastructure resilience

We find a correlation between stale reference prices and increases in market-wide message traffic. This has implications for the resilience of market infrastructure: that is, the ability of markets to function when message traffic substantially increases. This also has implications for the reference price waiver, which allows the absence of pre-trade transparency with the rationale that the external reference price, instead, provides that transparency. MiFID requires that the reference price must be 'reliable'¹⁰ which is not the case if it is affected by latency.

Implications for MiFID II

MiFID II will require microsecond granularity and maximum timestamp divergence of 100 microseconds for venues with less than one millisecond gateway to gateway latency.¹¹ It would seem from our analysis that market data latency regularly exceeds this threshold.

Timestamps relating to a trade could be time-stamped at several locations within trading process, and each will be affected by latency differently, and result in different timestamps.¹² The location of timestamping is not specified in the latest draft MiFID II technical standards.

¹⁰ Commission Regulation (EC) 1287/2006 implementing Directive 2004/39/EC Article 18.

¹¹ This would apply to all major UK equities venues; MiFID II Draft 'Regulatory technical and implementing standards' – Annex I, RTS-25, www.esma.europa.eu/sites/default/files/library/2015/11/2015-esma-1464_annex_i_-_draft_rts_and_its_on_mifid_ii_and_mifir.pdf

¹² Corvil, 'Best Practices for Implementing RTS-25,' <http://corvil.com/content/05-resources/04-white-papers/09-ebook-mifid-ii-clock-synchronization-rts-25/ebook-mifid-ii-clock-synchronization-rts-25.pdf>, p. 8.

2 Research context

Dark Pools and How They Work

In simple terms, a dark pool is a trading venue with no pre-trade transparency. While in lit venues market participants can observe the orders submitted by other participants, in dark pools all orders are hidden. The main advantage of submitting an order to a dark pool is that the trade intention is not revealed to the entire market. As we will describe below, another potential advantage is getting a better price than that available on the lit market (price improvement). The main disadvantage of dark pools is execution uncertainty. Specifically, it is impossible to know whether there is a willing counterparty, so one cannot know beforehand whether a trade will take place.¹³

Orders sent to dark pools usually include a price limit – the maximum price at which a participant is willing to buy (or the minimum price at which a participant is willing to sell).¹⁴ However, within the boundaries set by these price constraints, the dark pool operator is responsible for determining the price at which trades take place. To determine such a price, and as a direct consequence of the absence of pre-trade transparency, dark pools have to rely on a reference price determined elsewhere.

There are two important aspects of how the reference price is determined: first, which venue (or venues) are used to calculate it; second, which specific price points are used to match trades.

Dark pool operators have two options to determine which venues to use to calculate the reference price. The first option is to rely on a single venue, usually the 'primary' market, which in our case is the LSE. The second option is to consider multiple (lit) venues. In the first case, dark pools use the BBO prices available on the LSE. In the second case, operators construct what is known as 'the European BBO'¹⁵ (EBBO), which includes orders from the other venues. MiFID II, however, will prohibit using the EBBO to determine the reference price.

Having constructed the BBO or the EBBO, dark pools have to choose whether to match prices at the midpoint or also at the bid and the ask prices. The dark pools currently operated by BATS/Chi-X and Turquoise use only the midpoint price (i.e. a price exactly half way between the best bid and the best ask). Other dark MTFs, such as ITG Posit, UBS MTF and Goldman Sachs Sigma X, also use the best bid or the best ask price (depending on the direction of the trade). MiFID II will prohibit non-price-improving trades.¹⁶ Therefore, once MiFID II is in force, it will not be possible for dark pools to execute at the best bid or the best ask.

Types of dark pools

Dark pools can be characterised in many ways. For our purposes, it makes sense to divide dark pools into three subsets, depending on who operates them and the applicable regulatory requirements.

¹³ This is because there may not be liquidity available at the desired time to trade, or the liquidity could be 'one-sided.' For example, at the midpoint there may be a resting sell order rather than buy orders to facilitate sells. In this paper, we examine whether there is also 'price uncertainty' for dark pool executions, arising from latency in reference prices.

¹⁴ Similar to lit markets, participants in dark pools may choose to submit orders without a price by using a 'market order' that executes at the prevailing BBO (or midpoint if a dark midpoint order). In practice, these are rarely used.

¹⁵ In practice, the EBBO used by some dark pools often excludes smaller lit markets such as Equiduct and Aquis, so is not a true EBBO, but these venues have de minimis volumes. Other countries, such as the US, refer to this composite as the NBBO (National Best Bid or Offer).

¹⁶ MiFID II Article 4(1)(a) and (2).

Exchange operated MTFs are multilateral trading facilities operated by Turquoise and BATS Europe in the UK¹⁷. They match trades at the midpoint.

Broker operated MTFs are multilateral trading facilities operated by investment banks and other brokers. We treat them separately from exchange operated MTFs for two reasons: first, they tend to match trades at either the best bid or the best ask rather than at the midpoint; second, we have less information on them in our data (only trades, not orders).

Broker crossing networks (BCNs) are dark venues subject to less regulation than MTFs. In our data, we cannot determine the specific venue on which a trade took place if it took place on a BCN, as all these trades are simply reported as OTC trades. Therefore, these are excluded from our sample.

Are Dark Pool Reference Prices Reliable?

Determining reference prices requires continuous market data feeds from venues used to calculate it.

Two types of delay (or latency) exist in this context: processing latency and transmission latency. Processing latency is the time needed for the hardware and software to process and disseminate the information generated by the various venues.¹⁸ Transmission latency is the time it takes this information to travel from the venue generating the feed to the dark pool.¹⁹

These two sources of latency are the reality of trading in high-frequency markets today, reflected in new products and features in markets.²⁰ In the US, the dark pool IEX has designed an 'inbound speed-bump' to prevent latency arbitrage arising from latency in its reference price calculation, a similar but not equivalent speed-bump has been introduced by Alpha Exchange in Canada.²¹ In Europe, UBS's MTF dark pool has recently introduced a reference price collar, which stops trades from happening if the dark pool price is outside some bounds. BATS Europe has implemented a 'look-back period' reference price feature for its intraday periodic auctions. These mechanisms all minimise latency's impact on how prices are calculated.

The LSE has recently introduced a new Field Programmable Gate Array²² market data dissemination product with 'sub-five microsecond' advertised latencies to reduce *processing latency*.²³ Several microwave networks now criss-cross the UK, where none existed until late 2013. Microwave networks reduce *transmission latency* significantly compared with fibre optic networks.

While some latency (both processing and transmission) is unavoidable, as processing and transmitting information requires some time, in well-functioning markets both types of latency should be reduced to a minimum. Latency can give rise to arbitrage opportunities. For instance, if

¹⁷ We note that Turquoise is technically operated by an investment firm, Turquoise Global Holdings, which is majority owned by London Stock Exchange Group (LSEG), but also several investment banks which are also brokers. However, when considering its majority ownership by LSEG and its significant integration with the Turquoise lit market, we consider it is appropriate to class it as "exchange operated".

¹⁸ Bartlett and McCrary (2016) find that the NYSE and NASDAQ SIPs take 450 and 750 microseconds on average, respectively, to process incoming quote updates from US exchanges.

¹⁹ Bartlett and McCrary (2016) compare proprietary co-located feeds at the exchanges to the US NBBO consolidated tape (SIP) finding that the NYSE SIP takes 9 microseconds on average to receive quote updates from its own exchange, despite being in the same building. Quotes from BATS take 999 microseconds on average and 523 microseconds as median to travel 16 miles, a comparable distance to LSE and BATS in the UK. This significantly exceeds their 'theoretical minimum time' of 86 microseconds.

²⁰ Latency and high-frequency concepts are embedded in new regulation such as microsecond timestamp precision requirements in MiFID II and the SEC's recent proposal on one millisecond tolerances for Reg. NMS quote dissemination (File No. S7-03-16)

²¹ See Chen et al. (2016) for an examination of the implementation of Alpha's speed-bump.

²² FPGAs are a type of computer architecture in which the programming is in the hardware chip rather than the software, reducing processing latency significantly. www.lseg.com/resources/media-centre/press-releases/lseg-launches-fpga-powered-market-data-dissemination-platform

²³ This reduces the processing time to disseminate market data by the exchange. HFT and other latency sensitive participants have also used FPGAs in their co-location servers for many years, and more recently, in their microwave networks.

some participants can observe a new reference price before the dark pool can, they can ‘win the race’ between their feed and the dark pool’s feed and pick off ‘stale’ orders in the dark pool.²⁴

According to some publicly available information, the transmission latency between LSE’s datacentre in central London and BATS’s in Slough is about 320 microseconds, the one between the LSE and Turquoise is 60 microseconds.²⁵ There is additional time to process incoming order messages, quoted by exchanges to be about 25 microseconds.²⁶ But when messages spike, as is common within a millisecond, this infrastructure hits bandwidth/throughput constraints, and messages get queued. This increases latency by many multiples.²⁷

Latency has recently been recognised as significant in size and prevalence by regulators, practitioners and academics. It was significant enough with Goldman Sachs’s US dark pool to justify a fine of \$800k levied by FINRA in 2014²⁸ for matching trades at inferior prices due to latency. Goldman Sachs paid \$1.2m to clients as compensation for losses stemming from 395,000 stale trades. IIROC, the Canadian securities regulator, recently published research²⁹ showing that on average, 4% of all dark pool trades in Canada are at stale prices, and high variation across venues, with the overall proportion of latency-affected trades by value increasing over time, as well as in duration. The Tabb Group consultancy analysed ten months of trading data³⁰ for a large buy-side firm, finding that midpoint trades were priced at the far touch or worse 11.19% of the time, averaged across 20 different dark pools. Academic research, such as Ding et al. (2014), has found frequent occurrences of ‘dislocations’ between the US NBBO reference feed (the SIP), and the NBBO constructed from proprietary feeds.

The SEC fined Barclays and Credit Suisse in early 2016 for various violations and misrepresentations in the management of their dark pools.³¹ In particular, Barclays was fined for claiming it was pricing dark pool trades off fast direct feeds from exchanges, but was actually using slower ‘SIP’ feeds for many exchanges, including NYSE. The FCA’s 2016 Thematic Review of Dark Pools³² noted that some dark pool operators only monitored pricing feeds irregularly, or only on a post-trade basis.

Trading venues have begun to introduce features to address latency issues. In 2013, the US dark pool IEX, which has recently received approval as a regulated exchange, launched with a speed-bump that delays inbound and outbound orders by 350 microseconds.³³ Importantly, the inbound speed-bump does not apply to the dark pool reference price feed, which means that as long as the reference price feed is not stale by more than 350 microseconds, reference price latency

²⁴ Given the vast investment by HFT in high speed processing hardware such as FPGAs and transmission hardware, such as microwaves, it is highly likely they will consistently win this race.

²⁵ This is according to one market data vendor, S&P Capital IQ. These figures will vary across providers, and technologies used. Microwave connections are reported to be 30%–40% faster than fibre. For example, see LSE’s own offering: NexxCom. Sources: www.spcapitaliq-realtime.com/wp-content/uploads/2015/06/SP-CIQ-Real-Time-Solutions-Global-Network-Diagram-Q2-2015.pdf www.londonstockexchange.com/products-and-services/connectivity/hosting/wirelessfactsheetoctober2013.pdf

²⁶ LSE ‘Connectivity,’ www.lseg.com/sites/default/files/content/documents/LSEG_Connectivity_Full_Brochure.pdf. The latency of pre-trade risk checks, which must be performed on incoming orders, is also often quoted. A competitive market exists in minimising this latency. LSE quotes a reduction of 2–3 microseconds to below 0.5 microseconds in a recent exchange upgrade. www.londonstockexchange.com/products-and-services/technical-library/technical-user-group/tuglonov15.pdf, page 6.

²⁷ Corvil, a firm that provides latency management solutions for exchanges and market participants, states ‘most attention gets paid to minimum or average latencies, whereas it is usually the maximum latency or the high percentiles of the latency distribution that are most important.’ Corvil, ‘White Paper: Electronic Trading System Performance,’ 2014, <http://corvil.com/content/05-resources/04-white-papers/03-electronic-trading-system-performance/wp-electronic-trading-system-performance.pdf>, p.7. This characteristic of latency ‘spiking,’ also called ‘jitter’ requires the measurement of latency in terms of percentiles, to capture the behaviour at the upper-end of the distribution, not reflected in an average figure. For example, LSE quotes an improvement in 99th percentile latency following a hardware upgrade in 2015, www.londonstockexchange.com/products-and-services/technical-library/technical-user-group/tuglonov15.pdf, p.5

²⁸ FINRA, ‘FINRA Fines Goldman Sachs Execution & Clearing, L.P. \$800,000 for Failing to Prevent Trade-Throughs in its Alternative Trading System,’ 01/07/14, www.finra.org/newsroom/2014/finra-fines-goldman-sachs-execution-clearing-lp-800000-failing-prevent-trade-throughs

²⁹ Anderson, Devani, and Zhang (2016)

³⁰ Tabb Forum, ‘Dark Pool Execution Quality: A Quantitative View,’ 26/08/15, <http://tabbforum.com/opinions/dark-pool-execution-quality-a-quantitative-view>

³¹ (SEC 2016a; SEC 2016b)

³² FCA. 2016. “UK equity market dark pools – Role, promotion and oversight in wholesale markets.” TR16/05. Thematic Review. Financial Conduct Authority. <https://www.fca.org.uk/static/documents/thematic-reviews/tr16-05.pdf>, page 34.

³³ IEX Exchange’s Form 1 Application to the SEC: www.sec.gov/rules/other/2015/investors-exchange-form-1-exhibits-a-e.pdf#page=571

arbitrage is prevented.³⁴ The UK's Turquoise has introduced a 'random uncross' feature, which it says is beneficial for 'latency sensitive flow'.³⁵ Others, like Deutsche Bank's Super X, state that they will stop matching if orders are stale by more than a second.³⁶ Regardless, many dark pool operators disclose information on latency and how it is managed.³⁷

The Australian securities regulator (ASIC) recently published a report on dark trading and HFT in Australia.³⁸ They found that, on-exchange operated dark pools, HFTs were on the 'winning side' of trades that took place at stale prices 85% of the time compared with 31–32% for other users.³⁹ ASIC find that less than 1% of trades in Australian dark pools occurred outside the BBO reference price.⁴⁰

In this **study's first half**, we examine the prevalence of dark pool trades at stale reference prices. We attempt to answer the following questions: what is the probability of a dark trade occurring at a stale price? Has this changed over time? What causes latency?

³⁴ This is acknowledged in BATS' submission to the SEC regarding IEX's application to become an exchange: 'this 350 microsecond delay provides IEX the ability to update the prices of resting orders that are pegged... before market participants with faster access to market data can access those now stale prices on IEX.' www.sec.gov/comments/10-222/10222-9.pdf, page 1.

³⁵ Turquoise Dark Midpoint Orderbook, www.lseg.com/areas-expertise/our-markets/turquoise/equities/tq-dark-midpoint-order-book, accessed 28/04/2016.

³⁶ Deutsche Bank SuperX EMEA Brochure, https://autobahn.db.com/microSite/docs/SuperXEMEA_QuickGuide_Jan2016_v1.pdf, accessed 28/04/2016.

³⁷ For example, Goldman Sach's disclosure about its US dark pool: 'In compliance with Regulation NMS, GSEC monitors the latency in the market data used by SIGMA X in real time. This process works by comparing the time stamps accompanying market data received from the source which is primarily direct market data feeds to the time that a quote is received by SIGMA X (based on GSEC's internal clock). If this process identifies a latency greater than a defined threshold, SIGMA X will automatically suspend crossing functionality in the relevant security.' www.goldmansachs.com/what-we-do/securities/gset/equities/liquidity-access/sigma-x-us-faqs.pdf, p.3

³⁸ ASIC, 'Report 452 – Review of High-Frequency Trading and Dark Liquidity' <http://download.asic.gov.au/media/3444836/rep452-published-26-october-2015.pdf>, page 54

³⁹ We find similar results in our analysis (HFT:96% Co-located:12% Non co-located:9%).

⁴⁰ We find similar results for the UK.

Is the Primary Market a Good Reference Price?

Most dark pools choose the primary market to determine the reference price. This may be rational because most trading occurs in such markets, but it raises the question of whether market participants could get better prices if dark pools referred to a price constructed with the EBBO rather than relying solely on the LSE.

Most dark MTFs reference the LSE BBO (PBBO), but some, such as Instinet's Blockmatch, reference the EBBO. Some have a feature that prevents trades from occurring if the PBBO is inferior to the EBBO (UBS MTF).⁴¹ BATS Europe has introduced a new intraday periodic auction, which allows trades to reference the EBBO.⁴²

For BCNs, the reference price is not revealed publicly, making it difficult to draw conclusions on Primary BBO versus EBBO use as a reference price. However, composites of national market prices are common in US dark pools.⁴³

One European venue, Deutsche Bank's Super X, references the BBO, for example.⁴⁴ As stated above, MiFID II will prohibit using the EBBO⁴⁵, so we aim to examine the impact of this by quantifying the prevalence of executions at inferior prices due to the PBBO reference price.

'Best execution' rules in MiFID 1 create an 'obligation to execute orders on terms most favourable to the client.'⁴⁶ This means that a broker routing to a dark pool referencing a primary market price worse than that available on another lit market would require a reasonable justification.⁴⁷ These justifications were found lacking in the FCA's 2014 thematic review on best execution,⁴⁸ which found that *'firms which relied heavily on internalisation or on executing orders through connected parties were often unable to evidence whether this delivered best execution.'*⁴⁹ The thematic review also found that best execution failures were *'likely to fall disproportionately on less sophisticated clients'* unable to monitor their brokers. The FCA's 2016 Thematic Review of Dark Pools⁵⁰ found that the ability of users to effectively monitor and understand dark pools "was generally quite limited"⁵¹ and the need for dark pool operators to manage potential conflicts of interest. Therefore, we examine whether participants in dark pools are executing at worse prices

⁴¹ UBS, 'Well Formed Market Check,' and UBS MTF Rulebook, Section 6.6c(ii). Although this may be suspended for latency reasons in section 6.5c(i)).

www.ubs.com/microsites/multilateral-trading-facility/en/shared/memberinfo/jcr_content/par/assetlist.403901049.file/PS9jb250ZW50L2RhbS91YnMvbWljcm9zaXRlcy9tdGYybWVtYmVyaW5mb3JtYXRpb24vVUJTLU1URi1XZWxsLUZvcml1ZC1NYXJrZXQtLUNoZWNoZW5wZGY=/UBS-MTF-Well-Formed-Market-Checks.pdf

⁴² BATS Europe, 'Guidance Note Periodic Auctions,' 19/10/15,

http://cdn.batstrading.com/resources/participant_resources/BCE-GuidanceNote_Periodic_Auctions_Final.pdf

⁴³ Credit Suisse's Crossfinder, and Goldman Sachs's SigmaX in the US matches at the NBBO, www.credit-suisse.com/media/sites/aes/doc/form-ats-crossfinder-amd.pdf, www.goldmansachs.com/what-we-do/securities/gset/equities/liquidity-access/sigma-x-us-ats.pdf

⁴⁴ Deutsche Bank's Super X dark pool references the PBBO,

https://autobahn.db.com/microSite/docs/SuperXEMEA_QuickGuide_Jan2016_v1.pdf

⁴⁵ MiFID Implementing Directive Article 44(1).

⁴⁶ Under MiFID 1, firms were required to take 'all reasonable steps' to achieve best execution. In MiFID II, firms will be required to take 'all sufficient steps'; which is a more stringent standard. This makes the considerations in this paper more relevant under the future regime.

⁴⁷ In Europe, the principles approach to best execution allows for execution strategies that do not only rely purely on price, as in the US and Canada, which have 'order protection' rules. These require exchanges to route incoming orders to other exchanges with better prices. The FCA Handbook defines execution factors as, 'price, costs, speed, likelihood of execution and settlement, size, nature or any other consideration relevant to the execution of an order.' www.handbook.fca.org.uk/handbook/glossary/G2383.html but expects price to have 'high relative importance' (COBS 11.2.9)

⁴⁸ FCA. 2014. 'Best Execution and Payment for Order Flow.' TR14/13. Thematic Review. Financial Conduct Authority. www.fca.org.uk/static/documents/thematic-reviews/tr14-13.pdf.

⁴⁹ Ibid.

⁵⁰ FCA. 2016. "UK equity market dark pools – Role, promotion and oversight in wholesale markets." TR16/05. Thematic Review. Financial Conduct Authority. <https://www.fca.org.uk/static/documents/thematic-reviews/tr16-05.pdf>

⁵¹ FCA. 2016. "UK equity market dark pools – Role, promotion and oversight in wholesale markets." TR16/05. Thematic Review. Financial Conduct Authority. <https://www.fca.org.uk/static/documents/thematic-reviews/tr16-05.pdf>, p.20

than those available on the lit market, and whether less sophisticated participants are doing this more often.

An inherent conflict of interest exists where brokers operate dark pools, while having discretion over venue routing decisions for their clients' orders. These conflicts were recognised in the FCA's 2014 best execution review, which stated firms could not outline '*how they were managing potential conflicts of interest.*' The SEC's Rule 606 requires brokers to publish detailed metrics on their venue selection decisions each quarter. These metrics reveal that every major broker directs a disproportionately large number of orders to its own venue, with competing broker venues ranked towards the bottom, if they receive any orders at all.⁵² This demonstrates that brokers have strong links with their own dark pools, and are unwilling to route to competing brokers' dark pools. This implies that these venues do not appear to generate the same amount of network externality and liquidity aggregation effects as independent exchanges. A potential reason for this might be that competitive asymmetries exist within the dark pool that favour operators⁵³. We examine whether this is the case by comparing execution outcomes of the dark pool venue operator, in relation to other participants.

In the **study's second half**, we therefore investigate three questions in relation to the *choice* of reference price: what impact will MiFID II's compulsory primary reference price have on execution quality? To what extent are participants implementing best execution practices when a dark pool references a primary market price that is worse than that available on another lit market? Is this influenced by conflicts of interest with dark pools and participant sophistication?

⁵² For example, for NYSE Euronext Securities in Q1 2015, Deutsche Bank routed 24.86% of its orders in to Deutsche Bank's ATS in Q1 2015, with the next venue, BATS at 9.81%. No competing broker venues are listed. Credit Suisse routed 27.39% of its orders in NYSE Euronext Securities to its own venue. JP Morgan routed 15.52% to its own venue, followed by 8.67% to NYSE Arca. It also routed 4.98% to Deutsche Bank's dark pool and 4.63% to Credit Suisse,' one of the few brokers that routed to competing dark pools. Merrill Lynch routed 5.61% of its orders to its own dark pool, but no other broker dark pools. Goldman routed 4.5% and 3.7% to its own dark pools but no other broker dark pools. Non-investment bank brokers appear to be access a wider variety of dark pools, such as ITG and KCG.

Note: These figures concern 'non-directed flow,' where the client has not 'specifically instructed the broker-dealer to route to a particular venue for execution.' This varies by broker, but most orders are non-directed, consistent with agency broking still being dominant.

Sources: <http://vrs.vista-one-solutions.com/sec606rule.aspx>,
<https://web.archive.org/web/20150823023702/www.goldmansachs.com/compliance/Rule606>,
www.itg.com/order_routing/SEC_RULE_606_2015Q1.pdf

⁵³ This was identified as a "poor practice" in the FCA's 2016 "Thematic Review" of dark pools, wherein an in-house trading desk received a latency advantage over clients in its BCN, resulting from infrastructure differences, p.34.

3 Method and approach

Data

We use three datasets in our analysis. Order book data, transaction data from Thomson Reuters Tick History (TRTH) and transaction reporting data from the FCA's Zen dataset.

Order book data

This data was collected by the FCA directly from trading venues for market monitoring and research purposes to enhance our understanding of UK markets.

This is the most detailed and accurate dataset to examine dark pools in the UK to date.⁵⁴ It includes all the information recorded by the matching engine of four UK trading venues, at the millisecond level. The trading venues covered are the LSE, BATS, Chi-X and Turquoise.⁵⁵ Both lit and dark trades and orders are part of our dataset. These venues account for approximately 99.56%⁵⁶ of all FTSE 350 on-exchange (Lit) traded volume in the UK. This gives us a representative sample of lit trading, and an accurate calculation of the EBBO.⁵⁷

For these data, we observe all order submissions, amendments and cancellations, as well as executions. It is time-stamped with millisecond granularity, with buyer or seller initiated flags, price, quantity and information on the order type. We observe the identity of the member of the trading venues behind each event. In other words, the order book is not anonymised.⁵⁸

The timespan covered by our data is January 2014 to June 2015.⁵⁹

TRTH Data

We obtained millisecond time-stamped post-trade data from MiFID post-trade reporting repositories via Thomson Reuters Tick History.

Our sample of dark trading MTFs includes UBS MTF, Sigma X MTF, ITG Posit and Instinet Blockmatch as well as Liquidnet. After excluding Liquidnet, as reference prices are determined through bilateral negotiation⁶⁰ rather than the MiFID reference price waiver, our sample reflects 93.97%⁶¹ of overall dark MTFs trading.

⁵⁴ Previous studies in the UK on dark pools mostly include trade data. Brugler (2014) uses FCA transaction reports at second granularity and dark pool volume composition data from Fidessa. Outside the UK, Foley and Putniņš (2016) examine the introduction of minimum price improvement rules for dark venues on market quality in Canada, Comerton-Forde and Putniņš (2015) examine the effect of dark volumes on price discovery in the lit market in Australia.

⁵⁵ BATS and Chi-X are part of the same legal entity, having merged in 2012, but they maintain separate order books.

⁵⁶ Estimates were calculated for the period 1/1/14 to 30/06/15 for the FTSE 100 and FTSE 250 Index using information from Fidessa, fragmentation.fidessa.com

⁵⁷ This excludes smaller lit markets such as Equiduct and Aquis, but these are not included in the definition of the EBBO by dark pool venues.

⁵⁸ However, we do not know the underlying client if an order has been executed on an agency basis.

⁵⁹ It is the same data used in Aquilina and Ysusi (2016) but we focus on 2014–15 because we have the Turquoise exchange for this period, but not in 2013.

⁶⁰ Liquidnet Europe comprises two MTFs, 'Liquidnet Negotiation MTF' and 'Liquidnet Europe H2O.' While H2O does reference the midpoint price of the primary market BBO, it is not differentiated from negotiation in our post-trade data from TRTH, which it obtains from BATS Chi-X Trade Reporting. Rosenblatt estimates H2O to be 8% of Liquidnet Negotiation MTF volume ('Monthly Dark Liquidity Tracker – European Edition' – as at October 28, 2014 report for September 2014).

⁶¹ Liquidnet accounts for 5.42% of dark trading in the FTSE 100 and 8.65% in the FTSE 250. Estimates from Fidessa. We exclude Smartpool and Blink MTF as they are de minimis. Note that Fidessa does not include Goldman Sach's MTF, SigmaX in its estimates, but this is included in our sample. Rosenblatt estimates SigmaX MTF to be typically 5.7% of dark volume ('Monthly Dark Liquidity Tracker – European Edition' – as at October 28, 2014 report for September 2014).

Unregulated BCNs are not included in our sample. As these are unregulated venues under MiFID I, post-trade reporting does not require venue code reporting, merely being reported as 'OTC' venue trades. We exclude these from our analysis due to the inability to separate OTC trades organised on a BCN from other OTC trades. We estimate that this leaves us with a sample of approximately 51.2% of all UK dark trading.⁶²

Zen Data

Zen is the FCA's surveillance and monitoring system and includes transaction reporting data. By matching these transaction reports with those in TRTH we can have information on the counterparties to the trades reported in Zen but not in TRTH.⁶³

Sample composition

Our analysis uses a random sample of 57 stocks from the FTSE 100 and 57 from the FTSE 250. These two indices are chosen to obtain stocks representative of the whole market, which includes high and low liquidity stocks.

We exclude opening and closing auction periods in both samples as they are not relevant to dark trading.

The full order book data covers all of 2014 and half of 2015; but we restrict our analysis to five weeks, approximately two-and-a-half months apart, for computational reasons.⁶⁴ Unfortunately, for TRTH and Zen we could only access data for 2014 so the analysis is limited to four weeks in 2014 for this subset of the data. Annex 3 gives further detail of the data.

HFT Definition

We divide the traders in our sample in three categories, HFTs, co-located participants that are not HFTs, and non co-located participants that are not HFTs traders.

We follow the approach in Aquilina and Ysusi (2016) in identifying HFT participants. Our list is essentially the same, except for additions arising from our more recent sample. Our criterion for defining HFTs is that they are a subset of algorithmic trading participants that use proprietary capital to generate returns using computer algorithms and low-latency infrastructure.

Objective measures of HFTs have been proposed by Hagströmer and Nordén (2013) and Kirilenko et al. (2015), such as high order-to-trade ratios and inventory mean reversion. These measures aim to proxy for characteristics that latency sensitive participants may demonstrate, but do not guarantee these participants are truly latency sensitive, nor that others do not exhibit these characteristics.⁶⁵

Therefore, we use our internal supervisory knowledge, as well as knowledge of the platforms from which the original list was obtained, as the most accurate means of identification. Many of these firms now identify publicly as HFTs and established their lobby group, 'The Modern Markets Initiative,' in 2013.

In our sample, we observe 30 participants that we classify as HFTs. Our data identifies exchange participants at the firm level and not at the trading desk level.⁶⁶

⁶² Source: 'Monthly Dark Liquidity Tracker – European Edition' – as at October 28, 2014 (report for September 2014). This figure factors in non-reporting BCN estimates.

⁶³ For more details on matching methodology, see Annex 3.

⁶⁴ The periods covered are, the continuous five-day trading weeks starting: 13/01/14, 31/03/14, 16/06/14, 1/09/14, 22/06/15.

⁶⁵ For example, an HFT engaged in predominantly liquidity consuming (aggressive) trading strategies, will have a low order-to-trade ratio than an HFT engaged in liquidity providing (passive) market-making strategies. Brogaard et al (2015) identifies HFT participants using these measures, obtaining 43 with Hagströmer and Nordén (2013) and 7 with Kirilenko et al. (2015).

⁶⁶ If a participant has several accounts on the same venue or is a member of several trading venues, we consider all the activity of these accounts together as the activity of the firm. The accounts are likely to include the activity of many trading desks and we are unable to separate the activity of each desk.

Speed

While HFTs are acknowledged to be participants that rely on superior speed as part of their business model, there are also significant differences in speed amongst other participants. To examine the role of these speed differences in determining trading outcomes, we divide participants into degrees of latency sensitivity by which they are co-located at the primary exchange. Co-location refers to the placement of a market participant's servers in close physical proximity to an exchange to reduce transmission latency. This information is obtained from FCA supervisors. The vast majority of HFTs in our sample are co-located, and 99.84% of all dark trades by HFTs are co-located HFTs. Many participants that are not HFTs are also co-located.

Studies have also used co-location as a proxy for participant speed. These include Brogaard et al. (2015), which examines the optional take-up of co-location services by individual participants on NASDAQ OMX Stockholm, as well as Conrad, Wahal, and Xiang (2015), which examines its introduction on the exchange, and numerous others.⁶⁷

Identifying Stale Reference Prices

The methodology to identify stale reference prices is uncontroversial and based on the one developed by Anderson, Devani and Zhang (2016) and ASIC (2015).

To identify stale reference prices, we first look for quotes that match the dark trade price on the primary market.⁶⁸ That is, we look for the quote the dark pool used in calculating its reference price. Specifically, we look for quotes where the midpoint of the primary market matches the price of the trade for exchange operated MTFs, while for broker operated MTFs we look for quotes that match either the midpoint, the bid or the ask price.

To be conservative, we assume this is the most recent match. We also look for quotes that occur at most one millisecond *after* the trade time,⁶⁹ so as to allow for exchange clock non-synchronicity.

For a stale reference price to be identified, there must be at least one quote update before the dark trade occurs that *does not* match the dark trade price, the *intervening non-match*. This allows us to observe that the dark pool is referencing an older, 'stale price.'

We ensure this quote update occurs after the match by using the message sequencing number⁷⁰ from the primary market, which is reliable within a millisecond.⁷¹

To calculate the size of a stale reference quote latency we take the most recent timestamp of a quote that matched the dark trade subtracted by the timestamp of the dark pool trade, i .

$$Latency_i = MatchTime_i - TradeTime_i$$

We consider trades that are two milliseconds and above as stale, recognising that a one millisecond threshold would not allow for clock synchronicity and timestamp rounding effects. This methodology is similar to that in Anderson, Devani, and Zhang (2016), except we require latency to be two milliseconds and above, rather than one, and we use the message sequencing number in determining whether the *intervening non-match* is indeed intervening.

This methodology eliminates the risk of misidentifying stale trades⁷², but risks understating the extent of stale reference prices.⁷³ For example, if a trade is referencing a price two milliseconds

⁶⁷ Aitken, Cumming, and Zhan (2014), Boehmer, Fong, and Wu (2015) and Frino, Mollica, and Webb (2014).

⁶⁸ Dark pools will round midpoint prices to four decimal places and we appreciate this in calculating matches, but we have no cases of this in our sample.

⁶⁹ One millisecond reflects the upper bound of clock synchronisation accuracy provided by the exchanges.

⁷⁰ Message sequence numbers are ascending integers applied by exchanges to all messages to record the sequence with which an exchange processes incoming orders. As exchanges process orders in series, this enables event sequencing in historical data.

⁷¹ This is to allow for any clock synchronisation issues inherent in timestamps from different markets. We also require this *intervening quote* to occur after the 'stale price,' but we do not require at least one millisecond here as we can rely on message sequencing within the primary market data.

earlier, and there is a matching quote update at zero milliseconds, our methodology will assume this quote is not stale. Of course, in this example there is no cost (or benefit) to the counterparties of the stale price, but this exists only without price variation.

Measuring the cost of stale reference prices

To measure the cost of stale reference prices, we multiply the absolute value of the difference between the trade price and the LSE midpoint price at the time of the trade, by the volume of the dark trade.

For trade i , we calculate the cost as:

$$Cost_i = | (TradePrice_i - LSEMidatTradeTime_i) | * Volume_i$$

This reflects the cost relative to the counterfactual of the reference price not being stale. This assumes that the dark trade would have still occurred otherwise. There is a possibility that the trade occurred precisely *because* it was at a stale price.⁷⁴

Are the costs of stale reference prices borne equally?

For every stale dark trade, one counterparty loses out and one gains on the transaction, buying at a price better than the prevailing midpoint during the trade. We explore whether these costs are equally shared by participant types. Our expectation is that they are not, as participants experience different levels of latency (Hasbrouck, 2015).⁷⁵

If latency is evenly distributed across participants, the percentage of trades for which a participant is on the benefit side, versus the loss side, should be random, with mean zero. That is, a participant is expected to just as likely benefit or suffer from reference price latency.

The buyer to the trade benefits from the reference price latency if the trade price is less than the prevailing mid. The seller benefits from the trade if the trade price is greater than the prevailing mid at the time of the trade.

$$Benefit_i = \begin{cases} Buyer, TradePrice_i < PrimaryMidPrice_i \\ Seller, TradePrice_i > PrimaryMidPrice_i \end{cases}$$

We then calculate the proportion of stale dark trades in which a participant is on the benefit side by whether a participant is HFT and, separately, co-located at the primary market.

Causes of latency

Within exchanges, a relationship is recognised between higher latency and higher processing requirements. The LSE discloses strategies it uses to mitigate this effect, stating *[we use] sophisticated techniques to reduce latency and jitter*.⁷⁶ Exchanges like LMAX advertise their low-latency engine as a key sales feature.⁷⁷ There has also been an increase in academic research on the impact of exchange speed on market quality, such as Menkveld and Zoican (2016) and provoking discussion of market design more broadly in Budish, Cramton and Shim (2015).

⁷² See Annex 1B for robustness testing of clock synchronisation effects that may result in false positives.

⁷³ All dark trades referencing superseded prices are technically 'stale trades,' given the time required to transmit information with a theoretical lower bound of the speed of light. A more useful theoretical definition of a 'stale trade' (and one we would use if we had more accurate timestamps) would be a stale trade referencing a price superseded by a quote update transmitted slow enough to the dark venue for a participant to observe and react to it, thus having real practical implications for market participants.

⁷⁴ We also measure the cost of trades that occur at reference prices outside the non-stale primary market BBO, which are significant, but a subset of the total cost. In this case, there is a clear opportunity cost, as the participant could have aggressed the lit market and obtained a better price, assuming liquidity was available in the lit market.

⁷⁵ 'Latency depends on proximity to the market, status (retail vs. institutional, or subscriber/member vs. public customer), and technology.' Hasbrouck (2015), p.4.

⁷⁶ www.londonstockexchange.com/products-and-services/technical-library/technical-user-group/technicalusergrouparchive/may2013londontug.pdf, page 29.

⁷⁷ www.lmax.com/trading-tech

We provide empirical evidence of sources of latency, and in turn, stale reference prices by examining the relationship between stale reference prices and message traffic using regression models. See Appendix 2C for a detailed description of our methodology.

Dark Trading when the Primary Market has a Worse Price

For brevity, we refer to periods in which the primary market does not have the best price at a given price point as 'dislocations.' We refer to trades that occur when there is a better price elsewhere (during dislocations) as 'suboptimal.' We examine this in our full order book sample of midpoint-only venues and our sample of trades from broker operated dark venues that execute at the BBO as well as the midpoint.

Identification of dislocation periods

The primary market (the LSE in the UK) does not always have the best price compared to competing lit markets.

This means that the primary market may have a worse bid price, ask price or both (i.e. a wider spread) than other lit markets and therefore dark pools that reference these prices will also execute at worse prices than available elsewhere.

We examine two dislocation types. First, we call dislocations affecting trades that reference the best bid or the best ask, 'BBO Worse.' Second, we call dislocations affecting trades that reference the midpoint, where the LSE midpoint is worse than (outside the) EBBO, 'Mid Worse.' Dislocation periods are identified using the full order book data, and we require them to persist for longer than a millisecond.

Annex 3 details the prevalence of these dislocations in the market. The proportion of the day, averaged across stocks, with 'BBO Worse' prices is 32.8% and 'Mid Worse' is 2.68%.

Identification of suboptimal trades

To identify a trade as suboptimal, we require it to have occurred appreciably within a dislocation period. This means, in practice, for broker operated dark pools we apply timestamp tolerances which require the trade to occur within 40 milliseconds of the start and end of the dislocation to guarantee that the trade did indeed occur within a dislocated period. This period of 40 milliseconds is determined via sensitivity analysis detailed in Annex 3. To check the robustness of our results, we also re-performed our calculations with two and four times the threshold (80 and 160 milliseconds) and our key results are qualitatively unchanged. In all our examples, we require the number of shares at the better price to match or exceed the size of the trade to make sure that liquidity would have been available had the trade happened somewhere else.

For exchange operated dark pools, we require dislocations to start at least two milliseconds before trades to guarantee they did happen during a market dislocation. The difference between the treatment of exchange operated MTFs and broker operated ones is due to different timestamp accuracy. We remove trades at stale reference price identified in the previous section from this analysis. As we want to explore the decision to execute at a suboptimal price, we require settings where the price is correct.

Calculating the cost of suboptimal trades

For every suboptimal trade one party gains and one party loses. If the trade happens at a higher price the seller gains and the buyer loses and vice versa. In this dimension, the net cost is therefore zero by construction. However, to understand how markets function and calculate a gross measure of such costs we use the following metric:

$$ReferencePriceLossBBO_{t,v} = (DarkTradePrice_{t,v} - EBBO_t) * TradeVolume_t$$

DarkTradePrice is the price of dark trade, t , on venue v (which matches the LSE best bid, ask or midpoint, depending on dark trade type). *EBBO* is the European best bid or ask price, depending on whether the trade is referencing the bid or the ask for BBO pegged trades and for midpoint trades, depending on whether the European best bid or ask is superior to the primary mid. If LSE prices are equal to, or better than⁷⁸, all other venues that contribute to the EBBO, the trade is not suboptimal, and the reference price loss will be nil.

We do not consider trading fees in our analysis and acknowledge they are an important consideration in best execution and order routing decisions. However, we note that fee differences between venues are small in comparison to the size of our dislocations. LSE aggressive execution fees range from 0.45 to 0.15 basis points⁷⁹, BATS fees are from 0.15 basis points on the lit and dark books and BATS' Chi-X order book charges 0.20 to 0.30 basis points for UK securities on its lit order book and 0.15 to 0.30 basis points for its dark order book.⁸⁰ Turquoise lit and dark fees are 0.30 basis points⁸¹ and The UBS MTF dark pool is 0.10 basis points⁸², SigmaX is 0.10 basis points.⁸³ At their largest, fee differences would be 0.30 basis points, which is the fee amount of the largest non-primary lit market, which assumes that the broker does not charge their client for execution fees in its MTF. The majority of dislocations in our sample exceed this size, so it is unlikely to impact our findings.

We also do not consider price impact and adverse selection considerations in venue routing decisions for best execution which we also acknowledge as an important factor. For example, buyer initiated trades in dark pools at higher primary market referenced prices than other lit venues may be rational if subsequent executions are required to complete parent orders. Dark executions may have lower price impact, improving execution quality of the entire order.

⁷⁸ When discussing markets, 'better than' means that on the LSE, the ask price would be lower and the bid price would be higher.

⁷⁹ The LSE has several pricing tiers that depend on a participant's total execution values and liquidity taker scheme participation. Source: London Stock Exchange, Trading Services Price List, 2015, www.web.archive.org/web/20151022131129/http://www.lseg.com/sites/default/files/content/documents/Trading%20Services%20Price%20List%2020150708_final.pdf

⁸⁰ The Chi-X book had tiered pricing which depended on subscription participation. Dark execution fees depend on whether the dark order was resting (Immediate or Cancel IOC orders).

Source: BATS Europe, 'Trading Pricing', 2014,

www.web.archive.org/web/20141209130142/http://cdn.batstrading.com/resources/participant_resources/BATSEuro_Pricing.pdf

⁸¹ London Stock Exchange Group, "Turquoise Equities - Tariff Schedule", 2014, www.lseg.com/sites/default/files/content/documents/Turquoise%20tariff%20schedule%20-%20206.6%20final.pdf

⁸² UBS MTF "Fee Schedule", 2016, www.ubs.com/global/en/investment-bank/multilateral-trading-facility/legal-and-tariffs.html

⁸³ SigmaX MTF, "Fee Schedule", 2014, web.archive.org/web/20140422071207/http://qset.gs.com/Sigmamxmtf/Public/GetDocument/a861503b-620a-4c41-b4e7-4423fdd16447?compName=ParticipantDocs

4 Results

Prevalence of Stale Reference Prices

In this section, we report results on the prevalence of dark pool trades at stale reference prices on our sample of data from the exchange operated dark pools. Given the clock synchronisation issues present for broker operated dark pools, we cannot estimate with any precision the prevalence of stale prices on these venues. To estimate the costs associated with stale reference prices, however, we scale up the estimates obtained using only exchange operated dark pools using data on all venues, assuming a constant proportion, which is unlikely to be the case. Proportions in some pools may be worse than those in our sample.

Figure 1 details the percentage of stale trades (across all stocks and markets on a trading day). If a trade references a price two milliseconds or more before the trade, then we consider it stale.

This averages 3.5% over the entire sample, similar to IIROC's figure of 4% for Canada.⁸⁴ This appears to be trending upwards over time, averaging 3.36% in 2014 and 4.05% in 2015. As discussed in Section 3, these figures represent a lower bound on the true rate of stale reference prices in the market. This is because we only classify prices as stale if they are 'older' than two milliseconds. But any price that is older than the minimum practical transmission time of participants in the market is stale. For example if the fastest possible transmission time⁸⁵ from the LSE to a dark pool is 350 microseconds, participants may successfully race the quote update to the venue if prices are more than 350 microseconds old. Our timestamps do not allow us to observe these stale trades but we expect there to be a significant number⁸⁶.

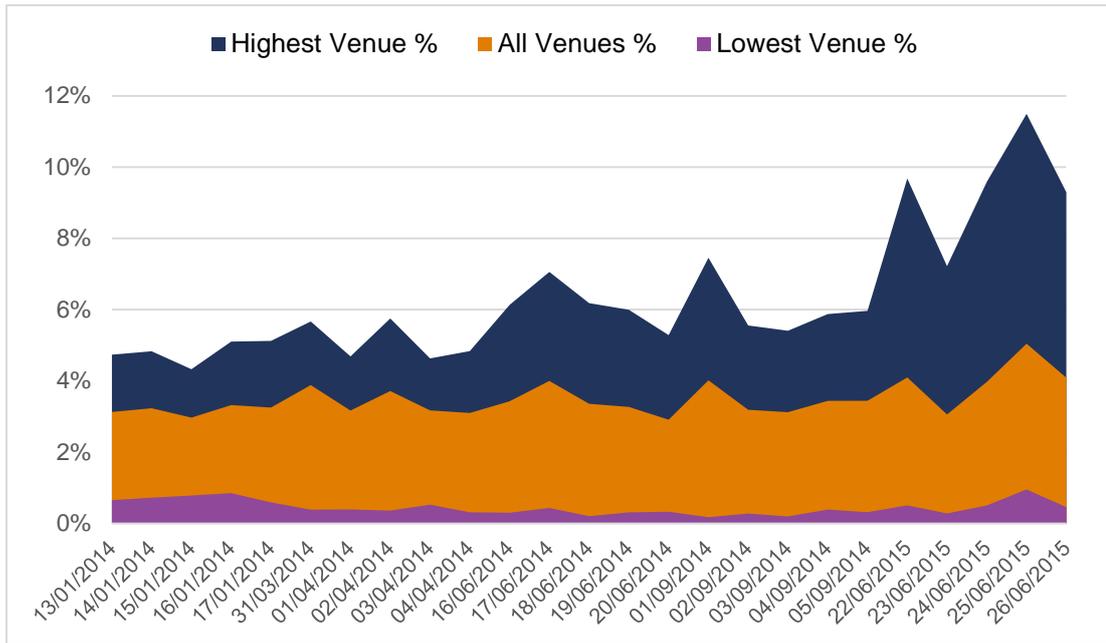
The upward trend appears to be explained by increases in volatility and message traffic over the sample, such that when we perform regressions (see Annex 2) controlling for these factors, we do not observe a time trend in most venues. This means that the level of stale dark trades is increasing over time, but this is explained by increasing message volumes. Future research with newer samples can confirm whether this trend persists.

⁸⁴ Anderson, Devani, and Zhang (2016)

⁸⁵ This fastest possible transmission time would be the sum of the speed of light over the geographic distance for the best route currently available, plus processing time for a participant to observe the LSE quote and transmit an order to do a dark venue.

⁸⁶ This is because our distribution of the age of stale trades exhibits exponential decay properties after 2 milliseconds and we expect the minimum transmission time to be less than 500 microseconds for the dark venues in our sample.

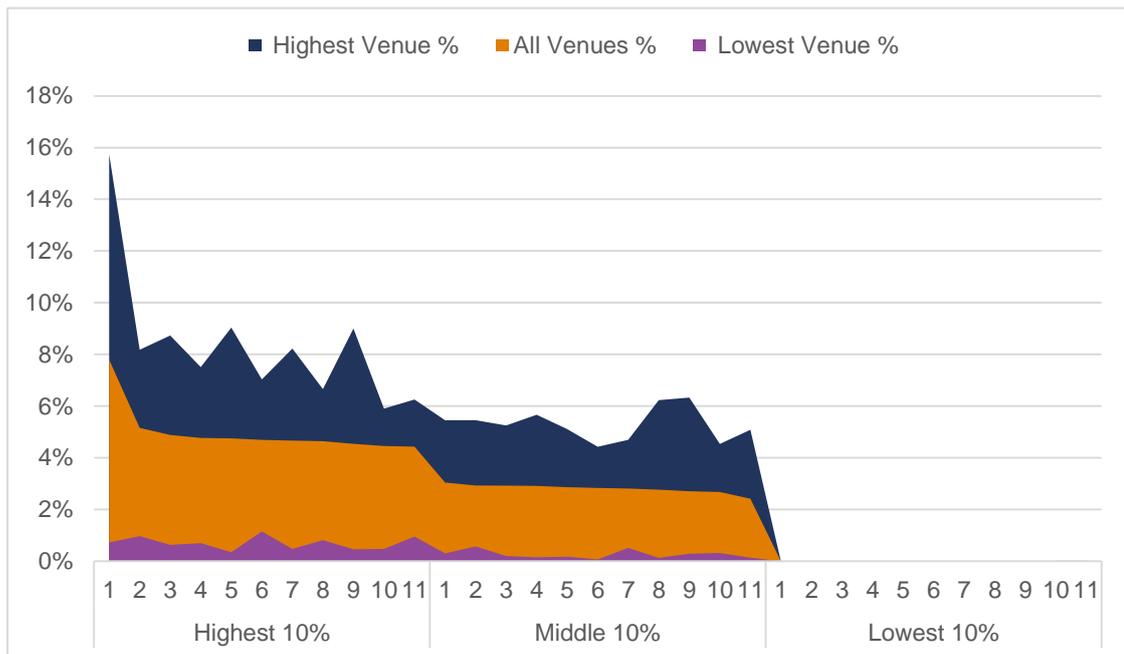
Figure 1: Proportion of Dark Trades at Stale Reference Prices (%)



The proportion of dark trades varies significantly by security. Across the entire venue date period, the highest proportion is 7.8% and the lowest is 0%. The proportion is highly correlated with the scale of the stock price. Larger stock prices allow a greater amount of price variation, which in turn reduces our ability to observe stale reference prices without price changes.⁸⁷

Figure 2 reports the proportion of stale dark trades within individual securities over the entire sample period. We report stocks with the highest proportion (top 10%), stocks right of the median (50–60%) and the lowest proportion (lowest 10%). We also report figures for the highest and lowest venue for that stock, as well as all venues. There is a large amount of intra-stock variation, from 15.7% for the highest stock and the highest venue for that stock, to 0% for the lowest.

Figure 2: Proportion of Dark Trades at Stale Reference Prices by Stock (%)



⁸⁷ Our methodology of detecting stale reference prices requires sufficient price variation prior to dark trades to observe stale reference prices.

Figure 3 shows the size of latency, measured as the time interval between the trade and the most recent match. We present metrics along three intervals of the distribution of latency times in seconds: the median, the top 25% and the top 10%.

While being relatively constant throughout 2014, the duration of latency appears to increase significantly in 2015, from a median of 2 milliseconds to 3, and the top 25% being above 3 milliseconds in 2014 and above 11 milliseconds in 2015.

Figure 3: Dark Trades at Stale Reference Prices – Latency in Seconds

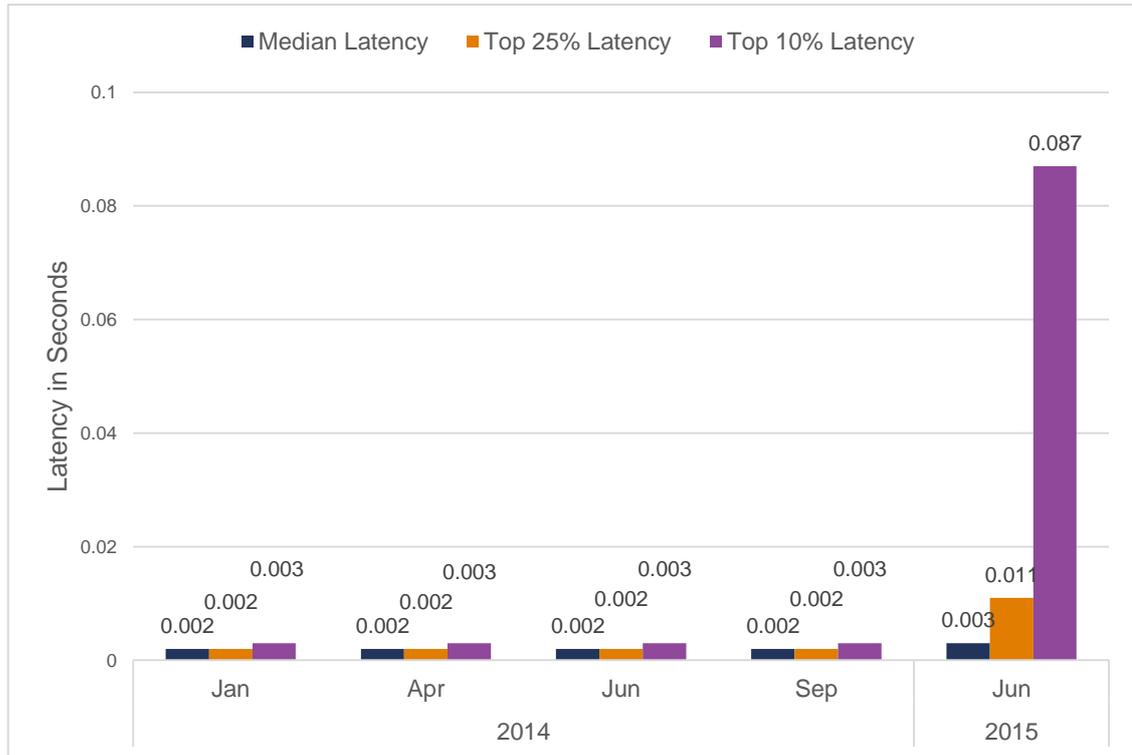
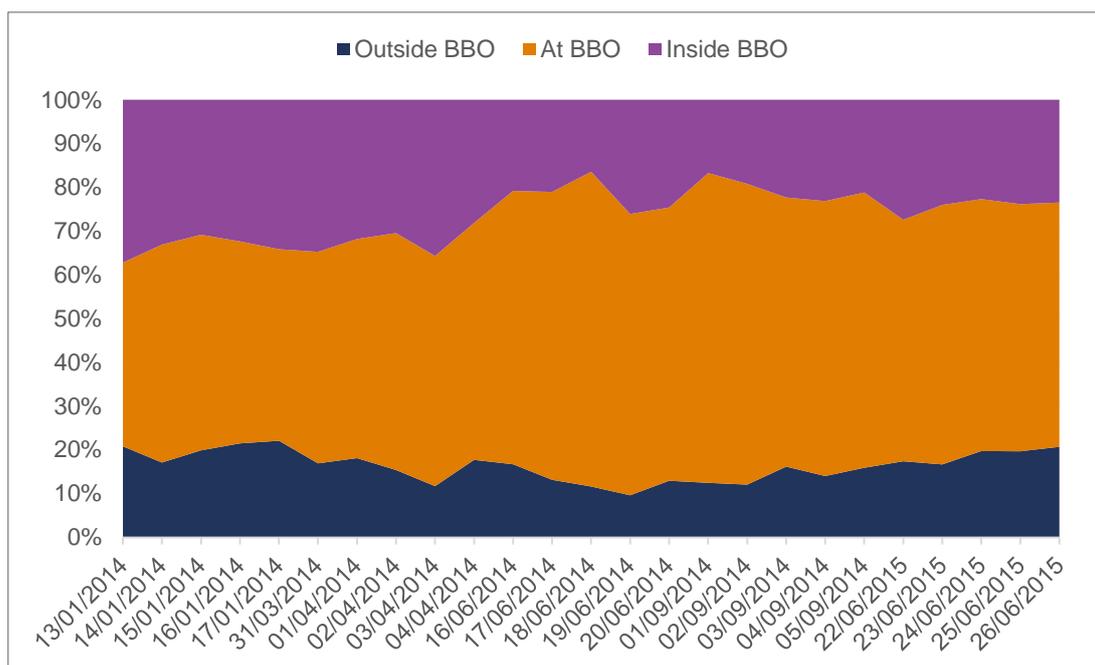


Figure 4 presents the proportion of dark trades at different price points in relation to the BBO: inside the BBO (26%), at the BBO (57%) and outside the BBO (16%). Therefore, most stale trades do not offer price improvement over the lit market. Prices referencing a stale price outside the primary market BBO represent risk-free arbitrage opportunities for participants able to buy (sell) at the stale reference midpoint when the current best bid (best ask) is higher (lower) than the stale midpoint price. Prices at the BBO represent cases where a participant does not receive any price improvement over the lit market. Figure 4 presents the proportion of stale trades by date across all venues and stocks in the sample.

Figure 4: Proportion of Stale Dark Trades Relative to BBO (%)

Costs of stale reference prices

For the exchange operated dark pools in our sample, the costs of stale reference prices are approximately £453,000 per year. Assuming the prevalence of stale prices was similar in broker operated dark pools the total yearly cost would be £4.2m across all dark venues. These figures do not appear to be economically significant. As a comparison we consider the trading revenues of some of the largest HFTs operating in the UK for an estimate of the total rewards to latency sensitive trading strategies, such as profiting from stale reference prices. Knight Capital Group Europe's trading revenue for 2014 was \$83.18m and Jump Trading's gross revenue for 2014 was \$97.1m.⁸⁸

However, we have assumed a constant level of reference price latency across our exchanges in our calculations. This is unlikely to be the case across all dark pools as we see considerable variation in our sample. In particular, some broker dark pools allow stale reference prices of up to a second in duration.⁸⁹ So the total cost of reference price latency could be much higher.

These costs include trades inside the BBO as well as outside the BBO. Only trades outside the BBO reflect 'real' opportunity costs because the losing counterparty could almost certainly have obtained a better price on the lit market due to the resting liquidity. This is not the case with inside BBO stale trades, for which we assume the dark trade would have still occurred had the reference price not been stale.

If measured in basis points per trade, costs are also relatively modest at 1.73bps. The agency broker ITG reports average broker commission costs of 9.4bps in the UK and implementation shortfall costs of 40.3bps.⁹⁰ Our figure is very similar to Hasbrouck's (2015) prediction of 1.83bps lost to fast traders by slower traders in a broader lit market setting of high-frequency quote volatility.

⁸⁸ Source: Public 2014 Financial Statements Lodged.

Jump Trading International Limited: <https://beta.companieshouse.gov.uk/company/05976015/filing-history>

Knight Capital Group Europe: <https://beta.companieshouse.gov.uk/company/03632121/filing-history>

⁸⁹ Deutsche Bank Europe's Super X broker crossing network.

https://autobahn.db.com/microSite/docs/SuperXEMEA_QuickGuide_Jan2016_v1.pdf

⁹⁰ ITG Peer Analysis, Global Cost Review Q3/2014, www.itg.com/assets/Q3_2014_2015.pdf

Although the costs we find are reasonably small, two observations are worth making. First, if activities like latency arbitrage contribute to the perception of a deterioration in ‘fairness’ in modern markets, this could cause investors to reduce their participation in such markets (Guiso, Sapienza, and Zingales (2008)). Second, our results are only based on UK stocks and UK venues. Latency arbitrage for stocks traded in UK-based dark venues and other European lit markets could be considerably higher, given the physical distance between the venues.

Costs may also be economically significant if liquidity providers are dissuaded from providing liquidity in dark venues in response to adverse selection costs of stale reference prices. We empirically examine this in Annex 2D, finding some evidence for this.

Table 1: Costs of Stale Reference Prices

Calculation	Total	Inside BBO	Outside BBO
Average bps per Trade ⁹¹	2.36	1.97	4.31
Total Measured Cost ⁹²	£44,793	£30,915	£13,878
Scaled Per Year ⁹³	£453,000	£313,000	£140,000
Scaled to FTSE 350 ⁹⁴	£1.4m	£928,000	£417,000
Scaled to all UK Dark Venues ⁹⁵	£4.2m	£2.9m	£1.3m

Are the costs shared equally?

All dark trades at stale reference prices are executed at a price which does not match the primary market BBO at the time of the trade. One counterparty benefits from this: they pay less or receive more for the trade than they would otherwise. If latency affects participants equally, then we expect equal outcomes across participant types. This is not what we find. We find that 96% of an HFT participant’s trades at stale prices are on the side of the trade which benefits, and on the losing side only 3.6% of the time (due to negligible HFT-on-HFT trading).⁹⁶ We report results for aggressive and passive benefit trades separately in Annex 2.

This finding is consistent with Baron, Brogaard and Kirilenko (2014) which find that HFTs profit through using aggressive market orders at the expense of other participants. This result may also be explained by HFT willingness to subscribe to faster market data feeds rather than just faster processing and order submission capabilities.

⁹¹ Bps = basis points (a hundredth of a percentage point). Calculated as stale reference price cost / trade consideration * 10000.

⁹² For the ‘Total’ column, this is the sum of ‘Outside BBO’ costs (real opportunity costs) and inside BBO costs (assumed opportunity costs).

⁹³ Assuming 253 trading days in the year, and 25 days in our sample.

⁹⁴ Assuming equal proportionality across stocks. We have 114 stocks in our sample.

⁹⁵ Assuming constant proportions across venues. According to Rosenblatt’s European Dark Liquidity report for September 2014, the venues in our full order book sample have 32.41% of total dark liquidity (MTF and BCN).

⁹⁶ When we calculate ‘Loss rates,’ they are the reciprocals of the benefit proportions within 1% due to negligible within-category trading. Results are also consistent when we split by aggressive vs passive trades, see Annex 2. The counterparty that initiates a midpoint trade does not have the same level of significance that a non-midpoint trade does, such as on the lit market: at the non-midpoint one counterparty inherently pays the half spread to demand liquidity and one earns it by providing liquidity.

Figure 5: Participant on Benefit side of Stale Dark Trade

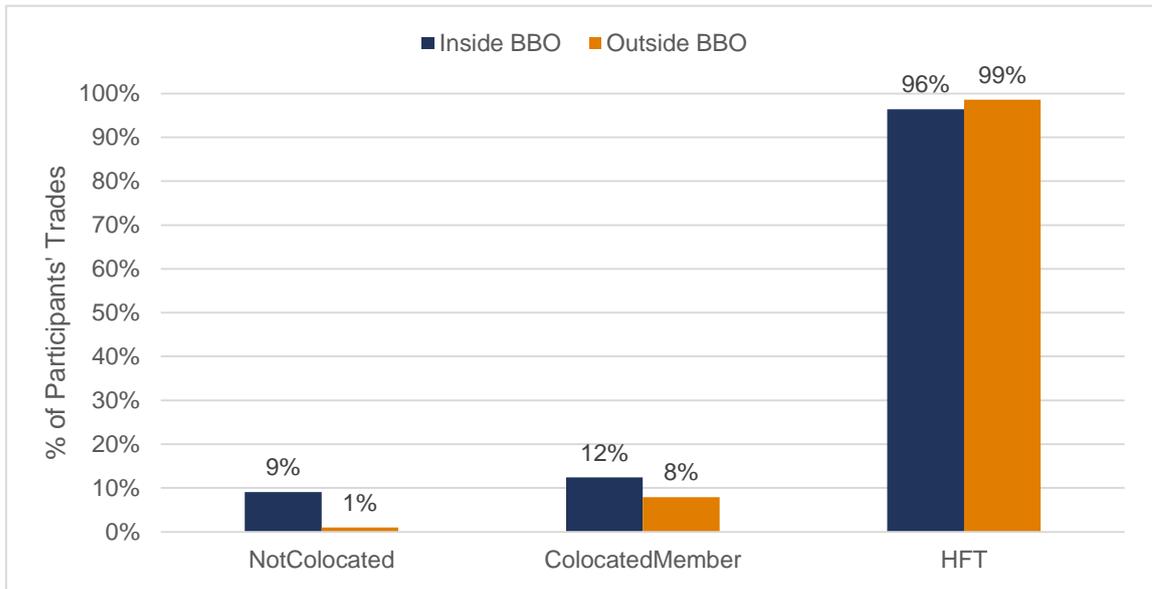
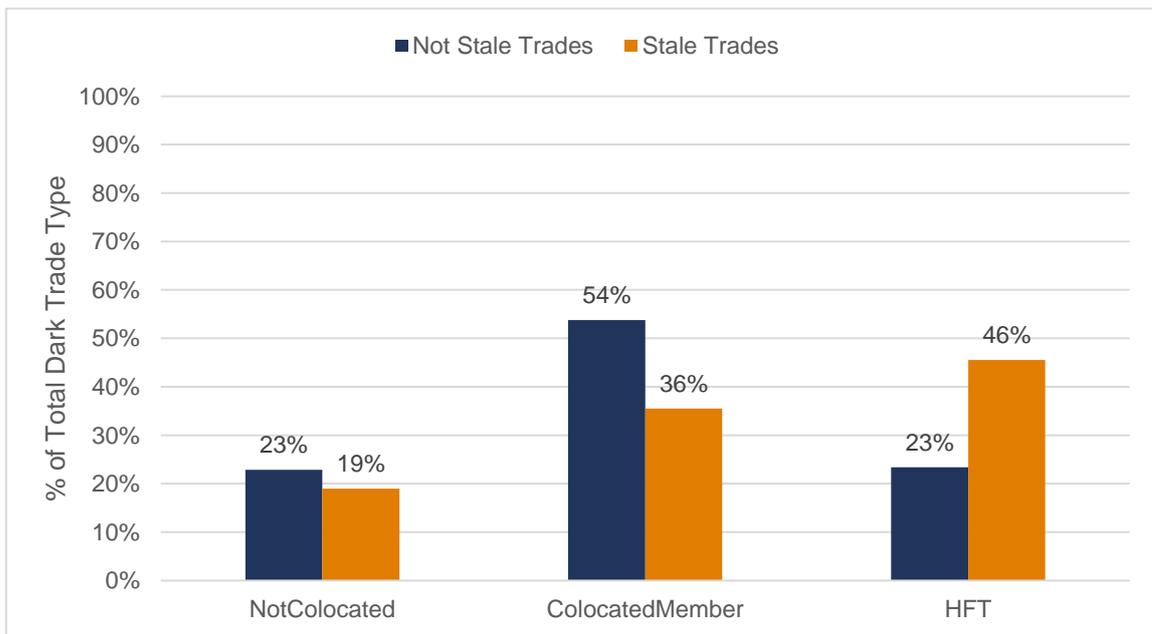


Figure 6 shows that HFT participation in these trades seems to increase substantially from their participation in non-stale trades. This implies that they are able to identify latency-affected periods and act on them to their advantage. They are able to observe the stale reference price, as well as the ‘true’ current price. This implies that they are unaffected by the latency associated with the reference price calculation. This could be because they are using a different or faster feed from the primary market than the exchange does.

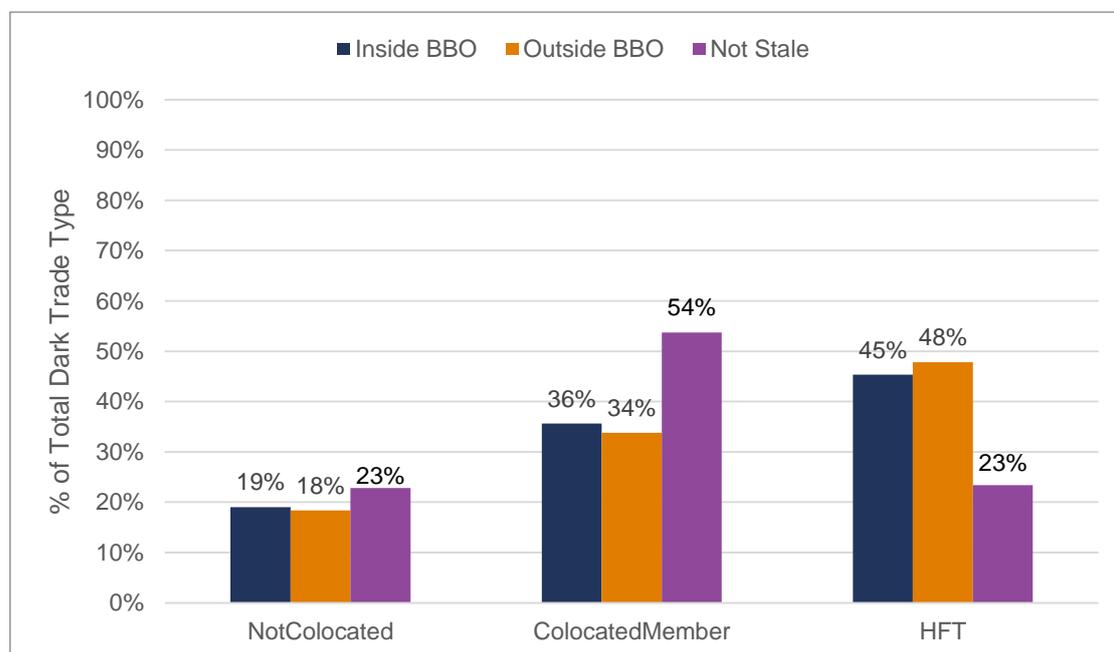
Figure 6: Dark Trades Participant on Either Side – Not Stale Vs Stale



Whether the stale reference price trade is inside or outside the BBO seems to have little impact on HFT willingness to participate. Figure 7 sets out the participation rates for stale trades inside and outside the spread. There is an insignificant difference in HFT participation rates between inside BBO trades (45%) and outside BBO trades (48%). Although stale trades inside or at the BBO are not as beneficial as pure-arbitrage opportunities, they are still beneficial as market-making strategies. They earn the half spread (or fractional, e.g. quarter spread, as the case may

be) and are executed at the top of the queue. Figures A3 and A4 in Annex 2 demonstrate that for most executions at stale reference prices, the benefit side is that of the aggressor.

Figure 7: Stale Dark Trades Participant on Either Side – Not Stale Vs Stale



What Causes Stale Reference Prices?

We attempt to determine factors associated with the level of stale reference prices. We do this by examining the changes in message volumes around individual trades (probit regressions) and by examining the proportion of stale trades in individual stock-days across the sample against factors that may explain it (OLS panel regressions).

In our probit regressions (see Annex 2), we find positive and statistically significant relationships with increases in message volumes across the entire market around our sample.⁹⁷

In our OLS regressions per stock day (see Annex 2), we find that the proportion of stale messages has a positive and statistically significant relationship with message volumes and market-wide volatility.

To model the causes of stale reference price more accurately, we require intraday historical data on the performance of the exchange infrastructure. For example, there may be other factors that are driving message latency, such as bandwidth consumed by feeds to other European markets by data vendors. Further, our message volume measure does not include exchange messages which do not reach the matching engine.

Dark Trading When Primary Market has Worse Price

In Table 2, we present the percentage of trades at inferior reference prices. These trades occur during periods where a better price is available on the lit market for the price point it is referencing (bid, ask or mid). We present this separately for BBO trades and midpoint trades.

We report results by venue type. 'Broker Operated Dark' refers to UBS MTF, Sigma X MTF and ITG Posit MTF. 'Exchange Operated Dark' refers to BATS Dark, Chi-X Dark and Turquoise Dark. The distinction here is that broker operators have a strong relationship with their venue, reflected in their significant participation in them (see Table 5).

⁹⁷ We do not report findings of individual market message traffic for confidentiality reasons.

Table 2 reports that a small number (0.57%) of dark trades occur when the respective LSE bid or ask price is worse than another lit market ('BBO Worse'). This is much smaller than the average percentage of the time that dislocations are present on the lit market, around 3.5% (see Annex 2). This demonstrates that, overall, participants must have smart order routers that observe and react to prices effectively most of the time. A larger proportion (1.22%) of midpoint trades occur when the price of the LSE midpoint is worse than the BBO of another lit market. This is roughly comparable to the percentage of the day we observe these dislocations, perhaps implying that participants are not as cautious with reference prices at the midpoint, assuming price improvement will occur regardless. For exchange operated dark MTFs, the percentage of suboptimal trades is much smaller at 0.31%.

Table 2: % Trades that Occur During Dislocations – By Venue Type

Venue	BBO Worse	Mid Worse
	BBO	Midpoint
Dark Broker MTF	0.57%	1.22%
Dark Exchange MTF	n/a	0.31%

This table contains the % of trades in the dark venue category that occur when the primary market bid or offer is worse than another lit venue for a respective bid or offer pegged trade. It also contains the % of trades where the primary mid is worse than another lit venue bid or offer for midpoint trades.

Table 3 reports the cost from the suboptimal trade as a percentage of the total consideration executed in the venue type in basis points. This is very small compared with the value traded, at 0.004 of a basis point for broker BBO trades and 0.001 for midpoint trades. This is similar for exchange operated venues at 0.033 basis points. The reason this is so small is that the magnitude of price differences between markets, when they exist, is not very significant.

Table 3: Total Suboptimal Cost as a Proportion of Total Consideration in Basis Points

Venue	BBO Worse	Mid Worse
	BBO	Midpoint
Dark Broker MTF	0.004	0.001
Dark Exchange MTF	n/a	0.033

This table presents the total cost of executing at a worse price for all identified suboptimal trades as a proportion of total trade consideration in basis points. This is presented by venue category by worse type (BBO worse or midpoint worse) and by trade type (BBO or Midpoint).

Table 4 reports the mean cost of suboptimal trades in basis points by venue. This is relatively small in size at 2.92 basis points for broker MTF BBO trades and 1.1 and 1.83 bps for midpoint dark trades on broker and exchange operated dark pools respectively. These represent minor costs compared with average broker commission costs of 9.4bps.⁹⁸ However, these costs do not appear to be evenly shared across participants, which we examine in Tables 7–9. For example, the worst performing category has 5.37% of their BBO trades in broker venues as suboptimal, which is nine times the average across *all* participants.

⁹⁸ ITG Peer Analysis, Global Cost Review Q3/2014, www.itg.com/assets/Q3_2014_2015.pdf

Table 4: Mean Dislocation Loss in BPs Per Trade – By Venue

Venue	BBO Worse	Mid Worse
	BBO	Midpoint
Dark Broker MTF	2.92	1.10
Dark Exchange MTF	n/a	1.83

This table sets out, for identified suboptimal trades, the average cost of executing at a worse price as a % of the trade consideration in basis points. This is presented by venue category by worse type (BBO worse or midpoint worse) and by trade type (BBO or Midpoint).

Table 5 details the composition of participants trading in each venue type, calculated with respect to both sides of a trade. For broker operated dark pools, the most active participant is the venue operator himself, e.g. UBS if it is UBS MTF, ITG if it is ITG Posit, etc. The next largest participant type is HFT participants, but they are significantly more active at the BBO than the midpoint. Co-located participants are the next most active participant in broker operated dark pools, followed by non co-located. We cannot match 22% of participants in these venues. These participants are likely to be either HFT or co-located or non co-located participants as this is driven by firm-level reporting requirements. In exchange operated dark pools, the most active participants are co-located participants, followed by HFT and non co-located firms. HFT participation is much higher in dark exchange operated MTFs than broker MTFs at the midpoint.

Table 5: Venue Trades by Participant (%)

Participant Group	Dark Broker MTF		Dark Exchange MTF
	BBO	Midpoint	Midpoint
Own Venue	46%	59%	n/a
Unmatched	27%	15%	n/a
HFT	22%	5%	24%
Co-located	4%	10%	53%
Non Co-located	2%	11%	24%

This table presents proportion of total trades by participant type, reported by venue type and trade type. This is calculated with respect to both counterparties to a trade. For example, a participant on both sides of every trade in a given venue type would score 100%.

Table 6 sets out the proportion of a participant's trades that are liquidity taking (aggressive) and liquidity providing (passive). Passive trades are calculated as the remainder of aggressive trades. Most participants in broker dark pools are liquidity consumers at the BBO with the exception of HFTs. Therefore, HFTs are providing the overwhelming majority of liquidity in these venues. A likely explanation for this is the attractiveness of dark pools as a means of 'queue jumping' the lit BBO. The shorter queue at the BBO in dark pools enables liquidity providers to circumvent time-priority constraints, as discussed in Kwan, Masulis, and McNish (2015) and Foucault and Menkveld (2008). Unfortunately, these metrics are unavailable for midpoint trades in broker operated dark pools as we do not have buyer or seller initiator flags. Regardless, the concept of liquidity provision at the midpoint is fundamentally different from the BBO: both sides give up half the spread by executing, and are thus neither liquidity demanders nor providers. Rather, this is more a reflection on the means with which a participant type goes about initiating a trade. Does it have a tendency to initiate, or wait for trades to occur at the midpoint?

Therefore, the high rate of aggressive trades by HFTs (94%) may merely reflect a tendency not to provide resting liquidity or a significant amount of order amendment activity. The exception is of course stale trades, in which a counterparty, *in actuality*, pays the majority of the spread, and the other counterparty receives it.

Table 6: Venue Trades by Participant that are Aggressive (%)

Participant Group	Dark Broker MTF	Dark Exchange MTF
	BBO	Midpoint
Co-located	94%	43%
Own Venue	88%	n/a
Non Co-located	76%	23%
Unmatched	24%	n/a
HFT	3%	94%

This table contains the proportion of a participants' liquidity-demanding (aggressive) trades, reported by venue type and trade type. This is presented by venue type and trade type.

Table 7 details the percentage of trades which are suboptimally executed, split by participant type, for broker operated venues. This shows a clear trend of more sophisticated participants obtaining better execution outcomes. Participants obtaining the best outcomes are venue operators themselves, and HFTs. Non Co-located participants execute 18 times as many trades at worse prices than those available in the lit venue than do venue operators trading in their venues.

Own venue traders may obtain good outcomes because they have the most accurate view of the reference price that the venue is *actually* obtaining, or at least are the most familiar with how its reference price feed behaves from heavy usage. Another explanation is that own venue operators are the most mindful of issues for best execution concerning the routing of orders to their venue, given the inherent conflict of interest considerations.

Table 7: % of a Participant Group's Trades – Suboptimal – Broker Operated

Participant Group	Primary Bid or Offer Worse	Primary Mid is Worse
	BBO	Midpoint
Own Venue	0.29%	0.59%
HFT	0.57%	0.40%
Unmatched	0.94%	0.62%
Co-located	1.86%	0.61%
Non Co-located	5.37%	0.87%

This table presents the % of a participant's trades in broker operated dark pools that occur when the primary market bid or offer is worse than another lit venue for a respective bid or offer pegged trade. It also presents the % of a participants' trades where the primary mid is worse than another lit venue's bid or offer for midpoint trades (i.e. midpoint trades executed outside the EBBO).

Table 8 reports a breakdown of suboptimal executions in exchange operated venues by participant group. This table also shows a clear trend of sophistication (or ability to observe prices accurately) being associated with improved execution outcomes. HFT participants rarely execute suboptimally, while co-located participants execute ten times as much as HFTs. Non co-located participants execute suboptimally almost twice as much as those not co-located. Annex 3 Tables A1 and A2 report the same results but for the worst performing participants (90% distribution cut-off).

Table 8: % of a Participant Group's Trades – Suboptimal – Exchange Operated

Participant Group	Primary Mid is Worse Midpoint
HFT	0.04%
Co-located	0.42%
Non Co-located	0.78%

This table contains the % of a participant's trades in exchange operated dark pools which occur when the midpoint of the primary market bid or offer is worse than another lit venue bid or offer for midpoint trades (i.e. midpoint trades executed outside the EBBO).

5 Conclusions

The rationale behind any reference price is that the price is transparent and reliable, resulting in symmetric outcomes across participant types.

We find asymmetric outcomes across participants when the reference price is stale, and when it is inferior to other prices available. This may result from participants' differing abilities to observe and manage latency, and differing abilities to engage in smart order routing effectively in a fragmented market. While the effects are highly statistically significant across participant types, the economic impacts are small.

Some amount of latency in reference prices is unavoidable. While we have observed a significant amount of stale trades on dark MTFs in our sample, the vast majority of dark trades are not stale, and thus the vast majority provide price improvement over the lit market for participants in dark pools.

Further, this study finds that latency is significant and persistent in modern markets. This demonstrates the need for reliable timestamps in understanding modern markets, as mandated in MiFID II.

This study also examines the extent to which participant classes achieve 'best execution' requirements when using dark pools, finding less sophisticated participants achieving poorer outcomes.

Annex 1: Further Descriptions and Robustness Tests

Annex 1A: Description of the full order book dataset

Our full order book dataset consists of detailed order book data from LSE, BATS, Chi-X and Turquoise and covers a sample of five weeks from the 2014 calendar year and 2015. These venues account for approximately 99% of all FTSE on-exchange traded volume in the UK. The sample is made up of 57 stocks from the FTSE 100 and 57 stocks from the FTSE 250 index. They are randomly selected stocks present across the time series. We observe all order submissions, amendments and cancellations, as well as trades for these instruments. The data include information on the instrument, date, time (at the millisecond level), side (buy or sell), price, total quantity, disclosed quantity, undisclosed quantity and consideration. Our data also include some further information on the order: type (e.g. limit, market, pegged, iceberg), time validity (e.g. IOC, Day), and if it derives from a sweep order (we cannot see if an order is a sweep order unless the sweep actually takes place). There is also information on the opening and closing auctions, but we exclude them from our analysis, as they are not relevant to our research question.

We also exclude periods where the best price on the LSE is comprised entirely of hidden orders, including icebergs, to avoid any confounding effects. Given that the timestamps are available at the millisecond level, we cannot know the exact sequence of messages across exchanges within the millisecond.

Importantly, our data are not anonymised. We can observe the member of the trading venue who submits the order.⁹⁹ This allows us to classify the participants as HFTs, co-located or non co-located. However, we do not know in which capacity the order is entered – as a principal or as an agency (i.e. on behalf of a client). If entered on behalf of a client, we do not know who the client is. As such, our classifications are based on supervisory knowledge of the business model of firms.

Annex 1B: Robustness test on clock synchronisation effects:

The results of this paper concerning stale reference prices are reliant on accurate timekeeping by the exchanges. Our data are historical data with timestamps generated by separate exchange clocks rather than real-time data captured at one source with one clock. Therefore, if the clocks in our sample data are not synchronised, the stale reference prices we observe may be ‘false positives’ driven by differences in timekeeping, rather than delays in market data dissemination and processing.

All clocks tend to drift over time. Marouani et al. (2008) quotes typical clock drift rates of one microsecond per second (one millisecond every 16.7 minutes). Therefore, we would expect exchanges to synchronise clocks intraday, at least every 16 minutes, or possibly continuously.

If this is not so, we would expect to see ‘jumps’ in timestamps. Empirically, these are only observable when a clock is running fast,¹⁰⁰ and must corrected by subtracting time.

When our order book data are sorted by message sequence numbers, this would create the appearance of time ‘going backward.’ Bartlett and McCrary (2016) perform this analysis on their 100 microsecond granularity direct feed US data and find that 0.88% of quote updates experience these negative timestamp intervals.

⁹⁹ Each message includes the IDs of the member of the trading venue that submits the order. If a participant has several accounts on the same venue, or is a member of several trading venues, we consider all the activity of these accounts together as the activity of the firm.

¹⁰⁰ Sommer and Wattenhofer (2008) provide evidence that drift has symmetric properties, so clocks are as likely to run faster or slower than a stationary reference clock.

We find no instances of this occurring in our data. This implies that the exchanges in our sample are synchronising their clocks at least often enough to correct for divergences smaller in magnitude than our minimum timestamp granularity of one millisecond. These results seem consistent given our higher granularity and fewer exchanges compared with the US.

Bartlett and McCrary (2016) detail strong efforts to provide highly accurate clock synchronisation in the US, among the largest FINRA¹⁰¹ regulated exchanges and firms. They cite the introduction of 'High Precision Time' by Perseus Telecom in 2014, which purports to provide timestamps at 'sub-nanosecond accuracy',¹⁰² as evidence of this.

Therefore, we conclude that the stale dark trades we identify are not caused by clock synchronisation issues.

¹⁰¹ Financial Industry Regulatory Authority of the USA. A self-regulatory organisation that regulates exchanges and brokerage firms.

¹⁰² Waters Technology, Equinix Deploys Perseus "High Precision Time" Timestamping Service,' 2014, www.waterstechnology.com/inside-market-data/news/2364536/equinix-deploys-perseus-high-precision-time-timestamping-service

Annex 2: Further Results on Reference Price Latency

In this Annex, we report separately for aggressive and passive benefit stale trades the proportion of trades on the benefit side by participant class. This is unlike the previous results in Figure 5, which present results for aggressive and passive benefit trades together.

Figure A1 shows that virtually all of HFT aggressive trading at stale reference prices is on the benefit side, and co-located firms can capture almost a third of outside BBO arbitrage opportunities, while non co-located firms capture very few. Figure A2 shows that only HFTs seem able to execute at stale reference prices on the benefit side consistently when the benefit side is passively executed. This is a likely result of the large amount of resting liquidity they provide at non-marketable prices, as demonstrated in Annex 2D, Figure A9.

Figure A1: Aggressive Participation on Benefit Side of Stale Dark Trade

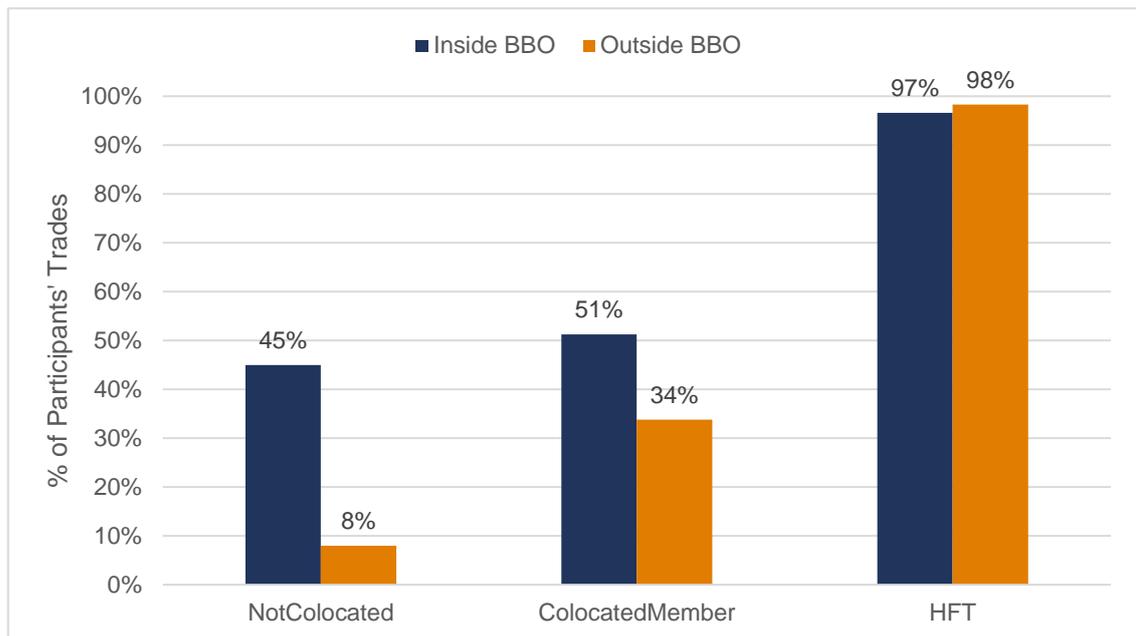


Figure A2: Passive Participation on Benefit Side of Stale Dark Trade

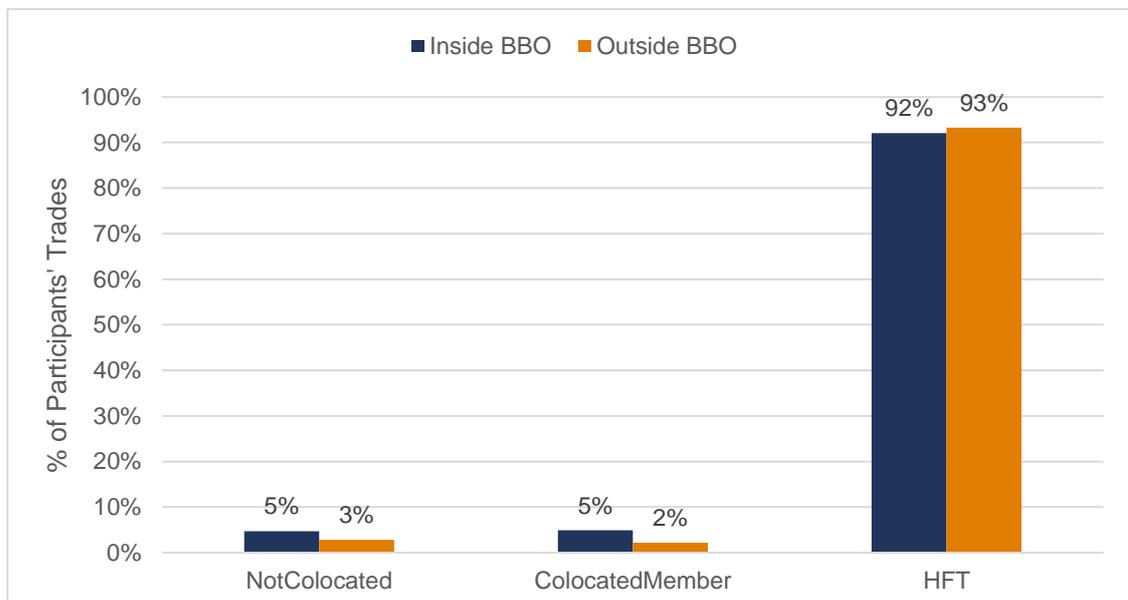


Figure A3 details what percentage of total dark pool trades the participant class is on the aggressive side. This is presented separately for stale and not stale trades, within each category the figures sum to 100%. The figure shows that HFTs are on the aggressive side of stale trades in the significant majority of cases (83%), perhaps crowding out the other participant classes. Non co-located firms are on the aggressive side of stale trades only 5% of the time. For non-stale trades, HFTs and co-located firms are equally likely to initiate trades, whereas non co-located firms rarely initiate (aggressively execute) trades.

Figure A3: Participation on Aggressive side of Dark Trades (% of all trades)

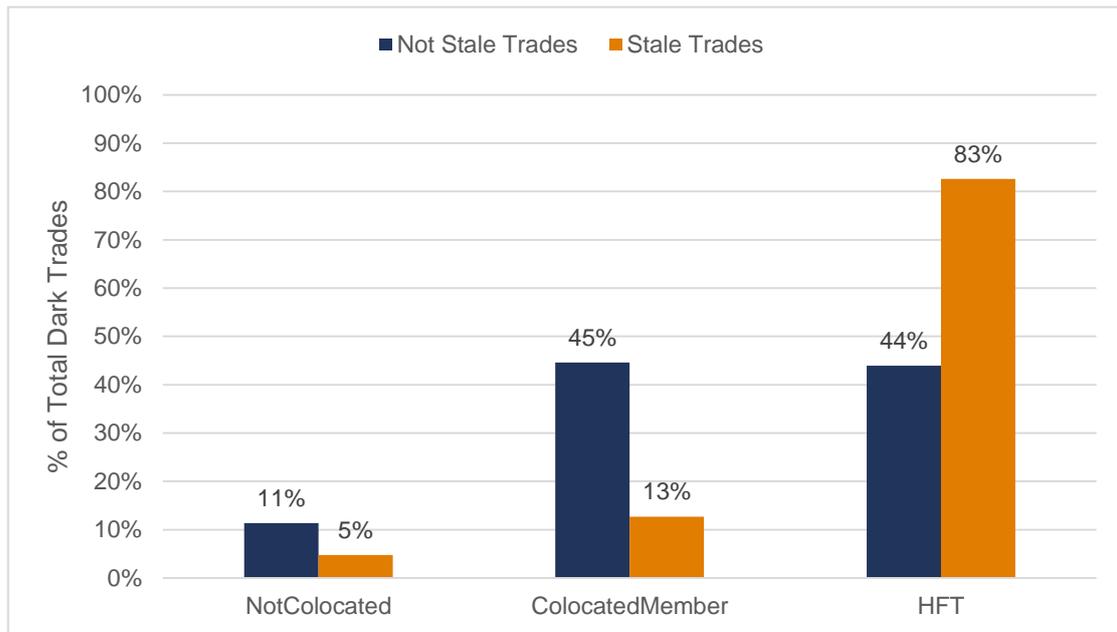
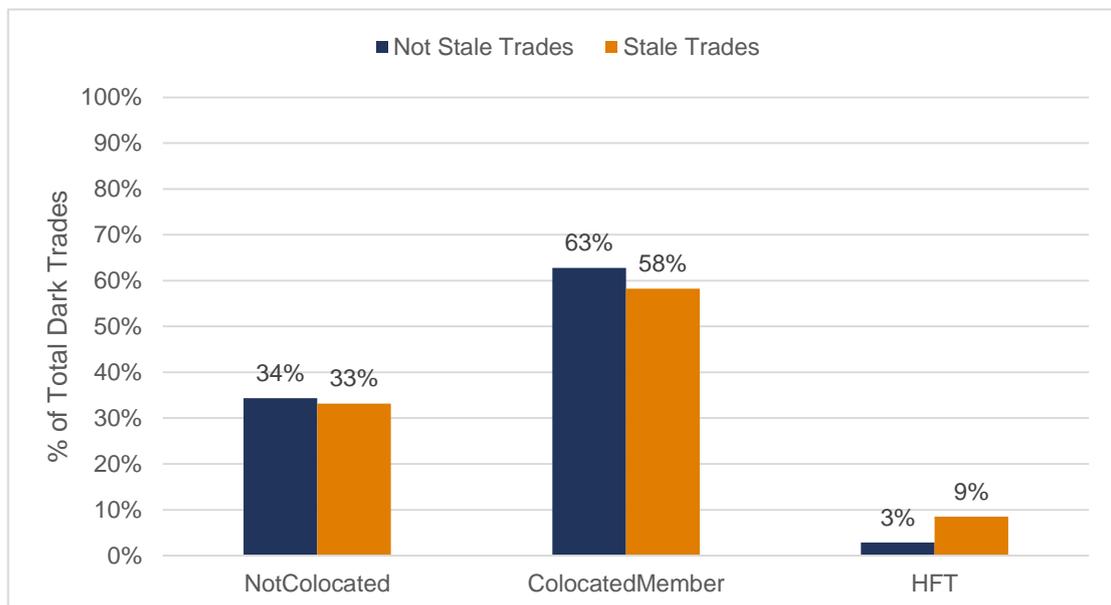


Figure A4 details what percentage of total dark pool trades the participant class is on the passive side. This is presented separately for stale and not stale trades. Within each category, the figures sum to 100%. The figures show that most dark pool trades are initiated with non-HFT participants on the passive side. This reconciles with Annex 2D, Figure A8, which shows that these participants provide the vast majority of resting liquidity in dark pools. When HFTs do execute passively, they predominantly do so when the reference price is stale.

Figure A4: Participation on Passive side of Dark Trades (% of all trades)



Annex 2A: Relationship of stale prices with adverse selection

In this section, we examine the impact of stale reference prices on price impact as a measure of adverse selection costs.

We focus on price impact, but results are also valid for realised spreads. We follow the approach in Malinova and Park (2016) and carry out trade by trade regressions, utilising similar controls. We run the following OLS regressions with standard errors corrected for clustering at the security and date level.

$$\begin{aligned}
 PriceImpact_{istd+m} &= \alpha + \beta_1 stale_{istd} + \beta_2 consideration_{istd} + \beta_3 momentum_{istd} + \beta_4 aggrHFT_{istd} \\
 &+ \beta_5 aggrColo_{istd} + \beta_6 passHFT_{istd} + \beta_7 passColo_{istd} + \beta_8 takebook_{istd} \\
 &+ \beta_9 VFTSE_{td} + \beta_{10} spread_{istd} + \beta_{11} stock_s + \beta_{12} date_d + \varepsilon_{istd}
 \end{aligned}$$

Where $PriceImpact_{istd+m}$ measures the price impact for trade i at for stock s at time t on day d , for $m = 100$ milliseconds, 5 seconds and 1 minute after the trade, in basis points. This is calculated as the difference between the trade price and the LSE mid price¹⁰³ multiplied by the trade direction (+1 for buyer initiated trades and -1 for sells).

$stale_{istd}$ refers to a dummy variable with the value of one if the dark trade is deemed to be stale. We have included various controls which aim to proxy for information, liquidity shocks, participant and stock specific factors. $consideration_{istd}$ is the natural log of the value of the trade in British Pound Sterling, $momentum_i$ is the midpoint return in the second prior to the trade, multiplied by the trade direction. $aggrHFT_{istd}$ and $aggrColo_{istd}$ are dummy variables representing if the aggressive side of the trade is an HFT or a co-located participant, respectively. $passHFT_{istd}$ and $passColo_{istd}$ are dummy variables representing if the passive side of the trade is an HFT or a co-located participant respectively. The rationale here is that HFT or co-located participants may be expected to infer information from the dark trade and cause price impact on the lit market in profiting from it. $takebook_{istd}$ is a dummy variable with a value of one for cases in which the aggressor counterparty of the dark trade also aggressively executes more than the available liquidity on the LSE within a 2 millisecond period before and after the dark trade. This aims to capture price impact relating to liquidity shocks from participants accessing multiple markets at the same time. $VFTSE_i$ is the natural log of the value of the FTSE 100 volatility index in the 15 seconds prior to the trade¹⁰⁴. $spread_{istd}$ is the quoted spread of the EBBO at the time of the trade in basis points. We also use stock and date fixed effects.¹⁰⁵

If in a trade that references a stale reference price it is the aggressor who benefits from the stale reference price, we would expect the trade to have a positive effect on price impact. If, however, the passive counterparty benefits from the stale price, we expect price impact to be negative. This is illustrated in Figure A5 and A6.

Figure A5 illustrates a buyer initiated dark trade at a stale reference price (see footnote for description of Figure).¹⁰⁶ Because the stale midpoint price is much lower than the new midpoint

¹⁰³ We also run results with the EBBO mid-price and they are qualitatively unchanged.

¹⁰⁴ This index is similar to the VIX in the US. We use the close price of 15 second intraday intervals.

¹⁰⁵ We winsorise price impact at 1% and 99% but the results are qualitatively unchanged from non-winsorisation.

¹⁰⁶ The charts in Figures A5 and A6 are visual representations of a lit market order book over time. The green shaded section represents liquidity at the Best Ask, the red shaded section represents liquidity at the Best Bid, and the unshaded section in the middle, represents 'the spread.' Therefore, the best ask is the lowest edge of the red shaded area, and the best bid is the upper most edge. The spread represents prices for which market participants are unwilling to place resting limit orders to buy or sell, or *unable* to place prices due to minimum spread requirements. The shaded circle represents a trade on a dark pool. Normally, on the lit market, trades would execute at the uppermost and lowermost edges of the best bid and ask. Because midpoint dark pools reference the midpoint of the lit market, they should execute in the middle of the shaded area. Because the trade in question is referencing a stale price, the price is within

price, the buyer side, and thus the *aggressive* side, benefits. When calculating price impact relative to the midpoint at the time of the trade, it is immediate and positive.

Figure A5 – Stale Dark Trade Example (Aggressive Benefit)

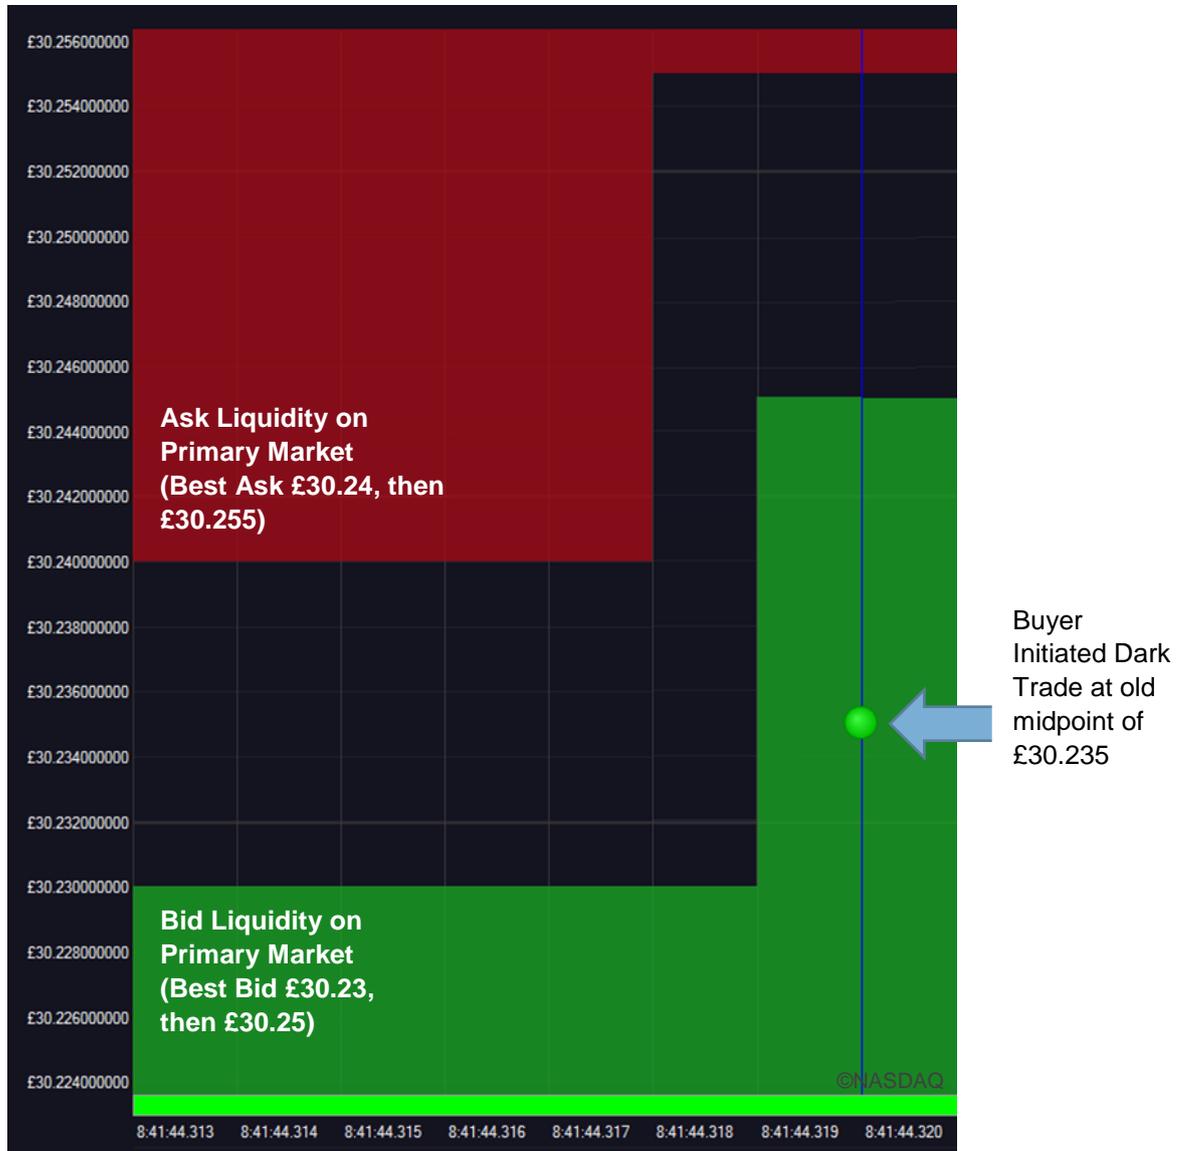


Figure A6, illustrates the opposite case: a buyer initiated midpoint trade that is referencing an older, higher, midpoint. Therefore, the benefit side is on the sell side, and thus the *passive* side benefits. In this case, price impact calculated with respect to the trade price and the midpoint at the time of the trade, is immediate and negative.

resting bid prices (for Figure A5) or resting ask prices (for figure A6) at the time of the trade. This is because it is referencing the *old* midpoint towards the left of the figure.

Figure A6 – Stale Dark Trade Example (Passive Benefit)

To isolate these effects, we perform separate regressions which include only stale trades where the benefit side is aggressive, and stale trades where the benefit side is passive, to isolate and test these opposite effects.

Regression Results:

Table A1 reports the results of the estimated model on three different samples for 100 millisecond price impact, but results are qualitatively the same at 5 seconds and 1 minute. The first column reports estimates from the full sample of trades, stale and non-stale. The second reports all non-stale trades, but only stale trades in which the aggressive counterparty benefits, and the third reports all non-stale trades but only stale trades in which the passive counterparty benefits. In the first column, we find a highly statistically significant and positive relationship between stale trades and price impact: aggressive benefit stale trades are more numerous than passive, and therefore overall their effect dominates.

The model estimates positive overall price impact of midpoint dark pool trades, of 2.4 basis points (the value of the constant in the first column). Stale trades increase this by 0.88 basis points for aggressive benefit stale trades (the coefficient on the variable *stale* in the second column). This effect is larger in size than the effect on price impact if the aggressor to a dark pool trade also executes against the full LSE best bid or ask (0.50 basis points, the coefficient on the variable *takebook* in the first column). Therefore, stale price effects seem to be larger than short term liquidity effects.

The variables in our model that control for participants involved in the trade, *aggrHFT* and *aggrColo* demonstrate a statistically significant positive relationship with price impact. Two interpretations are possible for this. First, HFT and co-located participants can react faster to information than non co-located participants. Second, they react faster to stale reference prices than non co-located participants, wherein the *aggrHFT* variable acts as a proxy for stale trades that we are unable to observe due to timestamp limitations. To provide evidence for this, we implement the same model, except with the dependent variable as stale and estimate it as a probit regression. These results are reported in column four, which show a strong relationship between stale trades and aggressive and passive HFT, as would be expected from our previous univariate statistics.

passHFT is strongly related with lower price impact. This may imply that HFT are adept at avoiding adverse selection, as Brogaard et al. (2015) finds for participants that take-up co-location. Spread is positively related to price impact, as it magnifies the effect of bid-ask bounce. We also see *takebook* is strongly correlated with higher price impact, as would be expected for trades that consume all BBO liquidity on the LSE at the same time as the dark trade. Spread is positively related to price impact, as it magnifies the effect of bid-ask bounce.

When we split the stale trades by whether the aggressive or passive counterparty benefits from it, we see that the aggressive benefit has a stronger positive relationship between stale prices and price impact than the first column, demonstrating that the opposing effects are masked in the aggregate. It also demonstrates participants with resting limit orders (passive initiators) to stale trades are facing higher adverse selection costs, measured as positive price impact. Conversely, stale trades can also allow aggressive trade initiators to face adverse selection costs, through negative realised spreads.¹⁰⁷ The passive benefit sample in column three shows a negative relationship with price impact, demonstrating that aggressive initiators of stale trades face adverse selection costs.

¹⁰⁷ Traditional market microstructure theoretical models such as Glosten and Milgrom (1985) model adverse selection as a cost that liquidity providers, or marketmakers, face. For midpoint dark pools, both counterparties are arguably providing liquidity, as they must forgo/pay half the spread to execute. Therefore, we can view the initiator to a trade as facing adverse selection costs despite the initiators to trades being viewed as liquidity demanders, rather than providers in traditional markets. We argue the systematic component of adverse selection occurs from selective liquidity provision from the passive benefit counterparty.

Table A1: Regression of Price Impact and Stale Trades

VARIABLES	(1)	(2)	(3)	(4)
	Price Impact (100 Millisecond)			StaleTrade
	Full Sample	Aggressive Benefit	Passive Benefit	Full Sample
Stale	0.649*** (24.806)	0.883*** (30.042)	-1.095*** (-15.603)	
Consideration	-0.039*** (-9.150)	-0.039*** (-9.068)	-0.039*** (-8.920)	0.046*** (10.015)
Spread	0.016*** (7.754)	0.016*** (7.766)	0.015*** (7.347)	-0.018*** (-10.539)
VFTSE	-0.783*** (-2.668)	-0.775*** (-2.645)	-0.755*** (-2.570)	-0.735*** (-2.622)
Momentum	0.479 (1.278)	0.509 (1.349)	0.526 (1.413)	-0.030 (-0.152)
AggressiveHFT	1.511*** (53.381)	1.492*** (53.101)	1.499*** (53.563)	0.829*** (35.762)
AggressiveColo	0.466*** (24.989)	0.466*** (24.993)	0.462*** (25.362)	-0.085*** (-3.831)
PassiveHFT	-1.444*** (-31.638)	-1.264*** (-27.414)	-1.266*** (-28.727)	1.017*** (33.991)
PassiveColo	0.058*** (4.132)	0.057*** (4.059)	0.063*** (4.460)	0.011 (0.847)
Takebook	0.503*** (5.631)	0.515*** (5.801)	0.484*** (5.361)	-0.321*** (-5.113)
Constant	2.425*** (3.288)	2.402*** (3.257)	2.357*** (3.189)	-0.720 (-1.049)
Observations	723,979	720,570	698,862	723,979
R-squared	0.135	0.134	0.126	0.107
Date Fixed Effects	Yes	Yes	Yes	Yes
Stock Fixed Effects	Yes	Yes	Yes	No

Robust t-statistics in parentheses, clustered by date and stock.

*** p<0.01, ** p<0.05, * p<0.1

Robustness:

We carry out our analysis using the change in VIX in the 15 second period prior to the trade, momentum as a continuous variable, *takebook* to exceed all EBBO liquidity, and price impact calculated from the LSE midpoint, rather than the EBBO midpoint and results are qualitatively unchanged.

Annex 2B: Examination of trends in stale prices over time

In this section, we aim to determine whether the proportion of dark trades is increasing over time, as it appears in the univariate trends, or if this is merely being driven by other factors. We do this by regressing the proportion of stale dark trades for a given stock day against our explanatory variables. In our first model, we include only a time trend, venue fixed effects and stock fixed effects, finding that the time trend is statistically significant and positive¹⁰⁸ consistent with our univariate trend.

In our second, third and fourth models we add controls for various factors that may have a causal relationship with stale prices, such as increases in messages¹⁰⁹, and factors which increase our measurement of it, or the opportunities for it to occur (price changes/volatility). Once we do this, the time trend is not statistically significant. Therefore, volatility and message counts explain the variance in stale trades far more effectively and happen to be increasing over time. Therefore, on average, the increase in stale trades we observe is related to increases in message volumes and volatility.

We investigate further the effects of message volumes on stale trades in the next section of this Annex.

Table A2: Regression of Stale Trades over Time

VARIABLES	(1)	(2)	(3)	(4)
	PercStale	PercStale	PercStale	PercStale
Time Trend	0.000** (2.613)	-0.000 (-0.211)	-0.000 (-0.104)	0.000 (0.632)
Spread		-0.581*** (-4.856)	-0.605*** (-5.237)	-0.598*** (-4.778)
PriceVolatility		0.000 (0.270)	0.000 (0.375)	0.000 (0.110)
VFTSE		0.030*** (2.772)		0.045*** (4.979)
MessageCount		0.002** (2.086)	0.004*** (4.095)	
Constant	0.033*** (19.849)	-0.049** (-2.067)	0.017*** (2.950)	-0.074*** (-3.287)
Observations	4,144	4,144	4,144	4,144
R-squared	0.437	0.445	0.444	0.444
Venue Fixed Effects	Y	Y	Y	Y
Stock Fixed Effects	Y	Y	Y	Y

Robust t-statistics in parentheses

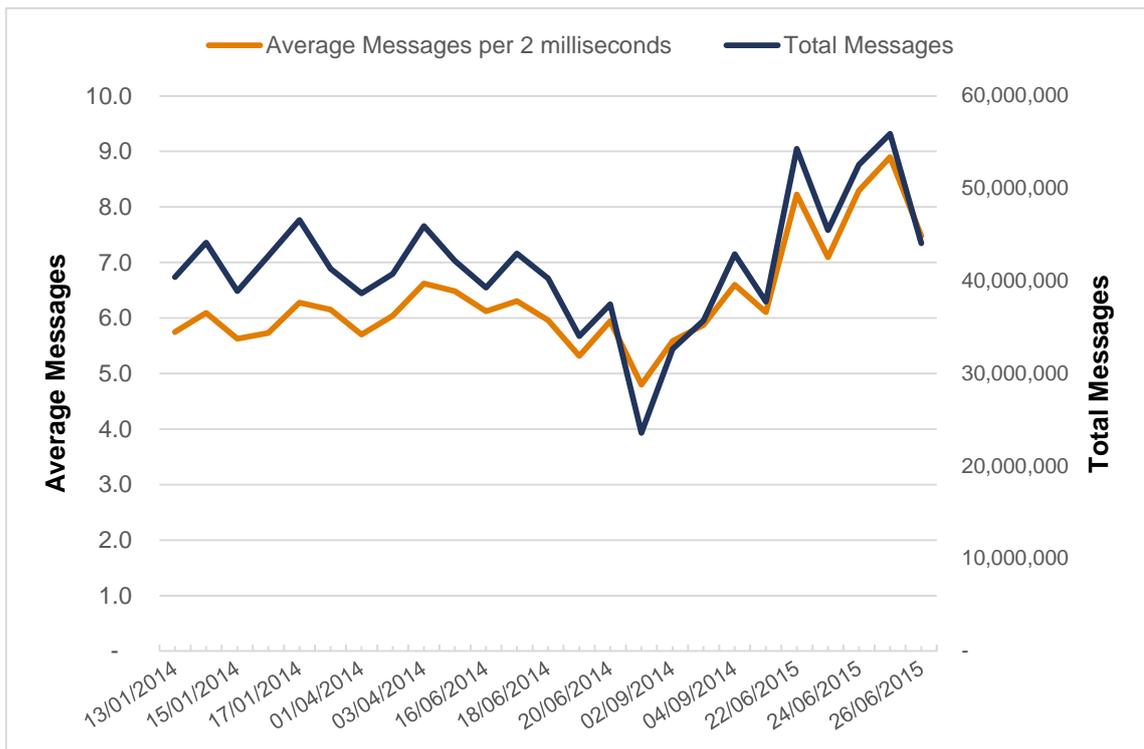
*** p<0.01, ** p<0.05, * p<0.1

T is our time trend variable, *spread* is the quoted spread in pence, *pricevol* is the % change in a stock's opening and closing price. *lnvftse* is the natural log of the FTSE Volatility Index. *Count_mean* is the average number of messages per two millisecond bucket across the top 400 stocks on all markets and dark pools on that day.

¹⁰⁸ We cluster standard errors at the stock level.

¹⁰⁹ We take the sum of all messages in all of our full order book markets for every stock in the FTSE 350, rather than our sample of 114 stocks.

Figure A7: Total Message Counts Across all Markets – FTSE 350



Annex 2C: Probit regressions of stale trades and changes in message volumes

We first sum the total number of messages across the 350 largest securities in all of the markets¹¹⁰ in our sample in discrete two millisecond buckets to harmonise any timestamp differences between venues at the millisecond level. We then calculate the change in message volume in the total market from period x to $x-1$.

We run probit regressions on individual trades where the dependent variable is a dummy variable that takes the value one if the trade is stale. We regress it against the change in message volume in 14 millisecond periods before the trade, and five periods after, removing all messages for the stock relating to the trade to mitigate any endogeneity. We also include venue and date fixed effects and cluster standard errors at the date level. This is reported in Table A3 below.

We find positive relationships with total market messages with the strongest statistical significance in the two millisecond period before stale trades, but also in the same two milliseconds as the stale trade. In unreported results, there are also positive relationships up to 28 milliseconds before stale trades. As expected, these relationships appear not to be significant after the dark trade occurs.

Table A3: Probit Regression of Stale Trades Against Market Messages

VARIABLES	StaleTrade
TotalMarketMessages – 14	0.005*** (2.640)
TotalMarketMessages – 13	0.004* (1.911)
TotalMarketMessages – 12	0.006*** (2.610)
TotalMarketMessages – 11	0.009*** (4.175)
TotalMarketMessages – 10	0.006*** (3.947)
TotalMarketMessages – 9	0.010*** (4.636)
TotalMarketMessages – 8	0.006*** (3.436)
TotalMarketMessages – 7	0.008*** (4.295)
TotalMarketMessages – 6	0.006*** (3.202)
TotalMarketMessages – 5	0.008*** (3.599)
TotalMarketMessages – 4	0.005** (2.569)
TotalMarketMessages – 3	0.004** (2.568)
TotalMarketMessages – 2	0.005*** (4.611)
TotalMarketMessages – 1	0.003*** (6.382)

¹¹⁰ In contrast to our previous sample of 57 stocks from the FTSE100 and 57 stocks from the FTSE250, we sum all stocks from the FTSE350 so that we can proxy for the aggregate market message traffic as accurately as possible.

TotalMarketMessages – 0	0.002***
	(5.934)
TotalMarketMessages + 1	0.000
	(0.439)
TotalMarketMessages + 2	-0.000
	(-0.039)
TotalMarketMessages + 3	-0.001
	(-0.824)
TotalMarketMessages + 4	0.001
	(0.634)
TotalMarketMessages + 5	0.000
	(0.232)
Constant	-1.691***
	(-48.356)
Pseudo R ²	0.1012
Observations	1,041,340

Robust z-statistics in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Annex 2D: Liquidity provision in dark pools

Given that stale reference prices impose costs on participants who rest passively in dark pools, we would expect increases in stale reference prices to have a negative impact on liquidity provision by increasing adverse selection risks to liquidity providers. In this section, we discuss liquidity provision in dark pools. We first provide some descriptive statistics on dark liquidity and then analyse the effects of stale reference prices on liquidity provision using an instrumental variable approach.

Descriptive Statistics

Our study is, to our knowledge, the first to characterise liquidity provision in modern dark pools. In traditional limit order book markets, measures of liquidity are commonly accepted and widely used in industry and academia. These include measures such as quoted spreads, effective spreads, and market depth. Given that the dark pool prices in our sample are mainly fixed at the midpoint, the first two measures are unusable. We consequently focus on depth related measures of liquidity. However, we note that unlike lit limit order markets, there are significant periods of time where liquidity is 'one sided' (liquidity is only available at the bid or the ask, but not both) as well as 'no sided' (no liquidity at the bid or the ask). 'No sided' liquidity occurs when there are no dark orders in the order book, but more often the dark orders have non-marketable prices. That is, they are buy orders with limit prices set at less than the prevailing midpoint or sell orders with limit prices set higher than the prevailing midpoint. Standard market depth measures are therefore not particularly informative in dark pools.

Therefore, we propose a measure of liquidity defined as the percentage of time there is a bid or ask order that is at a marketable price. In practice, this means that our measure of liquidity indicates periods in which people could trade in a dark pool if they were aware of the order's presence.

We calculate our measure of liquidity both for orders which exceed the respective primary market bid or ask quantities and for those which do not. We note that we are being liberal in our definition of liquidity supply in the context of midpoint dark pools. A participant with a marketable order at the midpoint price in the dark pool, may also be interpreted as a consumer of liquidity because they are willing to cross half the spread. Nonetheless, they are also a provider of liquidity: by resting passively on the order book, they allow executions to occur. Without resting liquidity, participants with aggressive orders merely 'ping' dark pools without executing, like 'ships passing in the night.'¹¹¹

Figure A8 presents cumulative frequency histograms of our dark liquidity metrics calculated for a given stock, date and market. Therefore, each chart illustrates a distribution of 8550 observations (114 stocks x 25 days x 3 dark pools). This was chosen to demonstrate the variability of dark liquidity by stock, date and market. This also demonstrates that there are many stocks and/or dates, in which it is not often possible for participants to access any dark liquidity. These charts are presented separately by the participant classes we have defined previously.

Discussion of results:

The largest providers of resting liquidity in dark pools appear to be co-located participants that are not HFTs, such as investment bank brokers. Most of this liquidity is for orders that are at a smaller size than that available on the primary lit market best bid or ask. Non co-located participants provide less liquidity, but more than HFTs, who provide almost no significant resting liquidity. Interestingly, HFTs provide a significant amount of resting liquidity that is *not* marketable, as illustrated in Figure A9. These orders can be interpreted as 'stop' orders, wherein they only become executable at a given price. The disparity between the small marketable, but high non-marketable resting orders by HFT can be explained by consistent repricing of resting orders to be

¹¹¹ This is a common analogy used in midpoint dark pools, crossing networks and dark aggregators. (Banks 2014)

To mitigate this effect, BATS Europe's dark pools both provide lower execution fees for orders which rest on the order book (Non-IOC orders). http://cdn.batstrading.com/resources/participant_resources/BATSEuro_Pricing.pdf

non-marketable, in response to primary market movements. A potential rationale for this behaviour is to take advantage of stale reference prices through passive executions, as described in Annex 2A. But this could also be explained by attempts to minimise adverse selection risks on the dark pool.

Figure A8 – Histograms of Marketable Orders in Midpoint Dark Pools by Participant

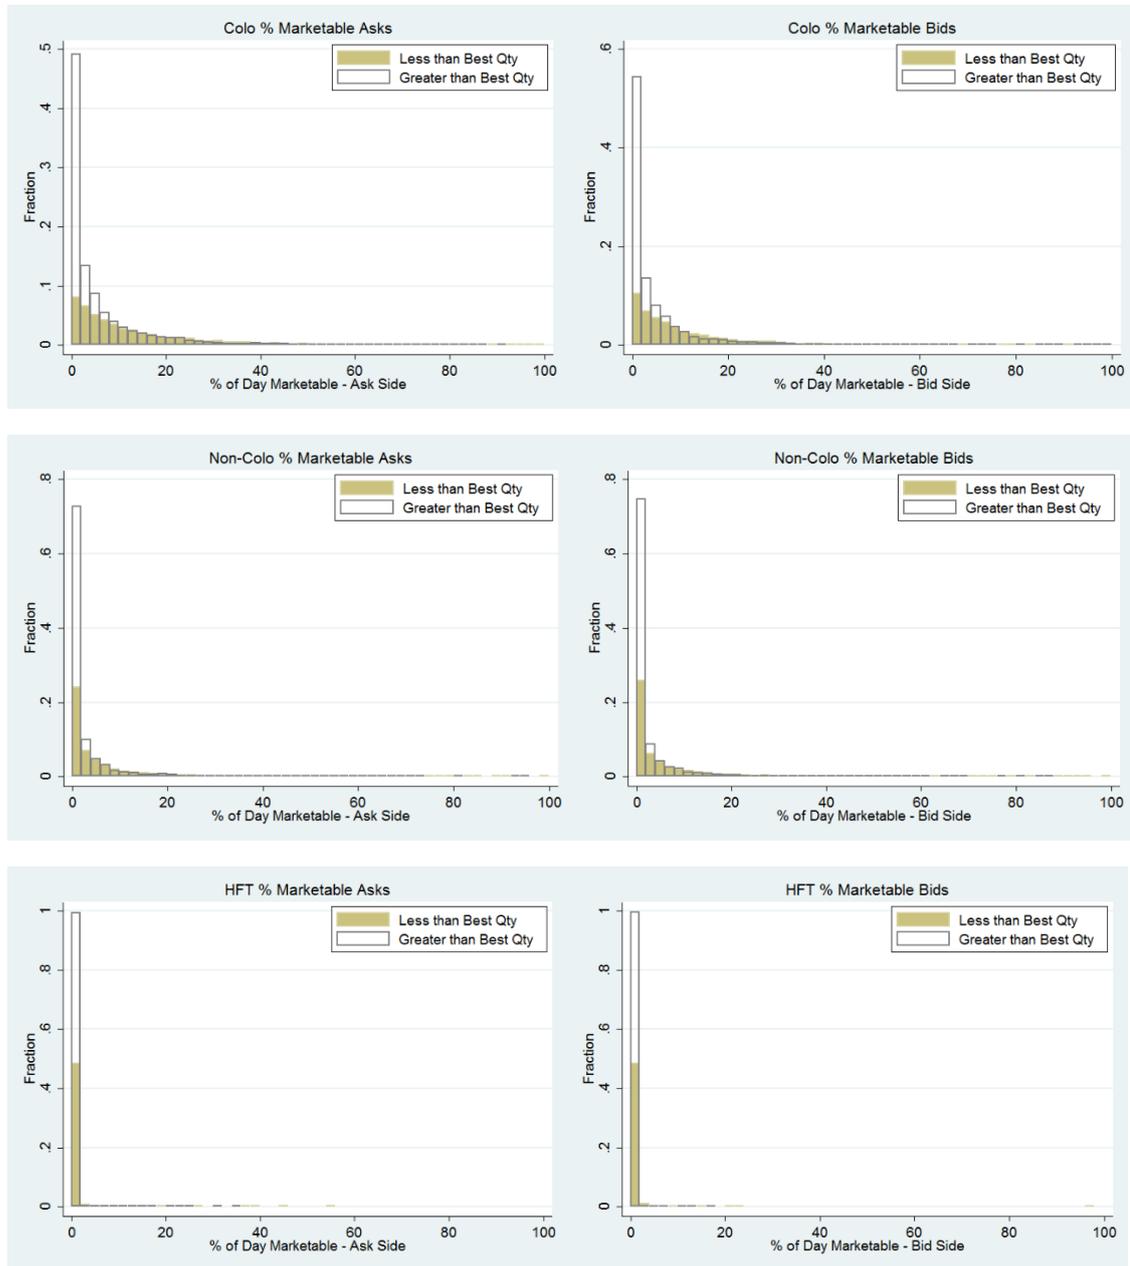
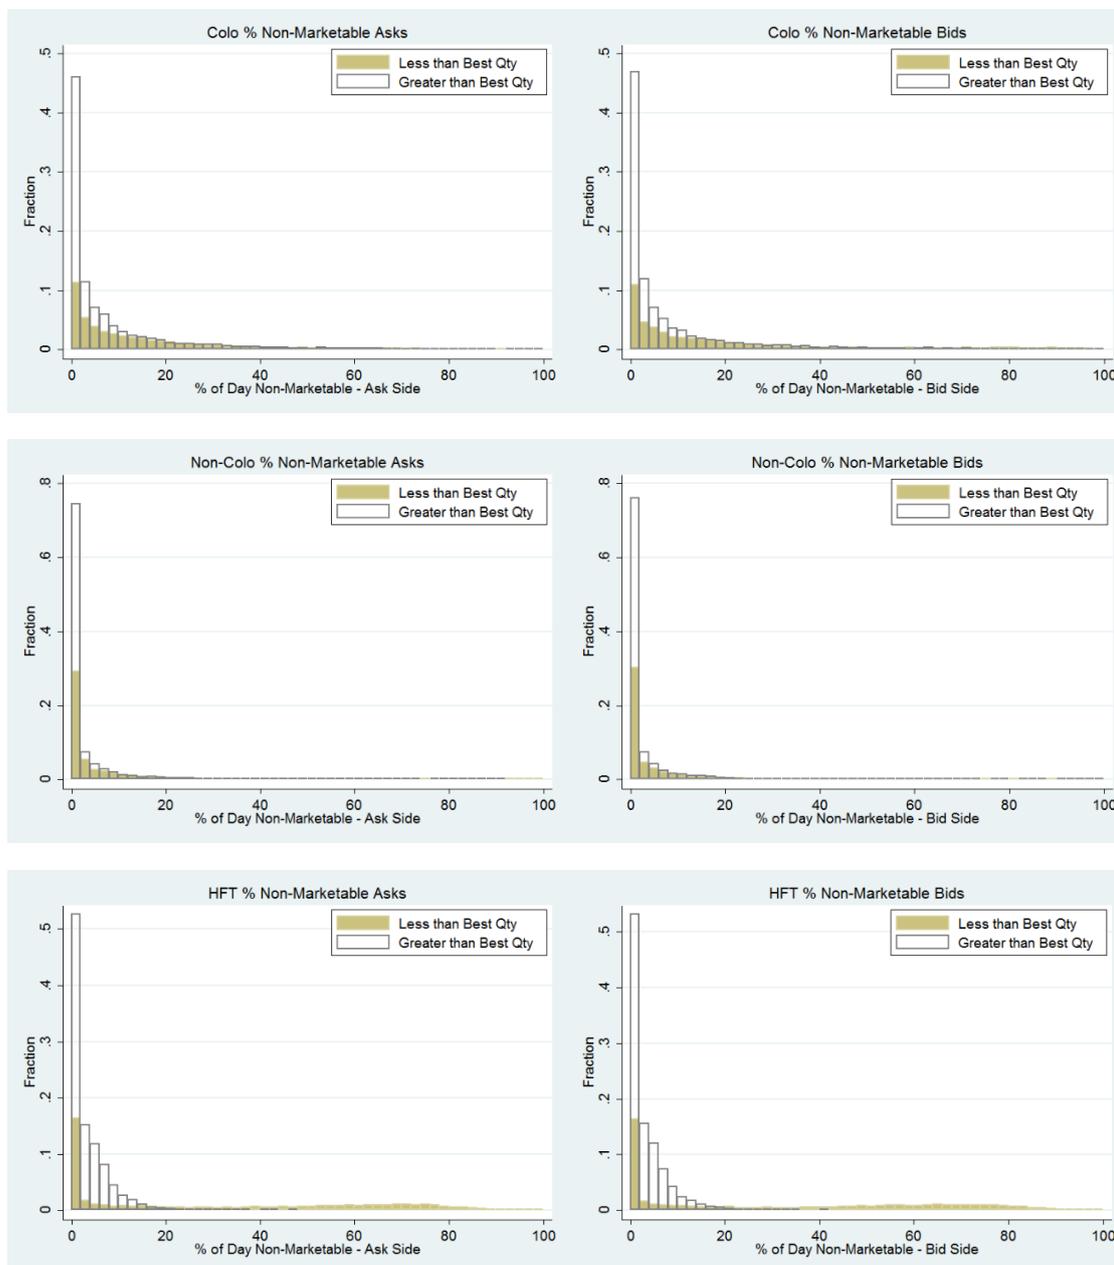


Figure A9 – Histograms of Marketable Orders in Midpoint Dark Pools by Participant



Models of Liquidity Provision:

To study the effect of stale reference prices on dark liquidity provision, we need to take into account our measure of stale reference prices, which requires dark trading, which requires dark liquidity, is thus likely to be endogenously determined with liquidity provision itself.

We therefore need to find an instrument correlated with our measure of stale prices but not correlated with our measure of liquidity provision. Our instrument is a continuous measure of latency within the marketplace calculated across three LSE market data dissemination security groups¹¹², intraday in ten minute intervals.

To construct our instrument, we match all LSE quote updates in our order book data to TRTH quote updates and calculate the difference in timestamps in milliseconds. We only examine

¹¹² Market data is disseminated by the LSE in three channels: Channel A (FTSE100 Channel Group A), Channel B (FTSE100 Channel Group B) and all others in Channel C. We obtain channel groups for the LSE from FCA supervisors but these are made public to market participants. The logic here is that any transmission and processing latency will be correlated within these groups, following Ye, Yao, and Gai (2013)'s examination of NASDAQ channel assignments.

updates to the BBO, and exclude any updates which have the same price and quantity within 100 milliseconds to prevent any matching errors. We calculate timestamp differences in 100 millisecond intervals for all stocks in our sample throughout the day, and match them to all dark trades in our sample that occur in the same 100 millisecond interval. As can be seen in Figure A10, this latency measure exhibits significant intraday variation, in particular in response to the US market open in the afternoon. Given that latency is so variable on an intraday basis, for our instrument to be effective, it must only instrument market conditions around dark trades themselves. As we have constructed our measure of stale dark trades conservatively, by including only trades referencing a price over two milliseconds as stale, we must similarly measure latency for our instrument by only measuring significant latency. Therefore, we take the 99% cut-off of the distribution of differences in each 100 millisecond interval, to observe only significant latency spikes around quote updates. We then average these matched time differences in ten minute intervals to form our latency instrument.

For the instrument to be a valid one it must satisfy the exclusion restriction; there must be no direct effect of market wide latency on dark liquidity. It is possible that in extreme periods, high volatility and messages may increase latency and coincide with changes in liquidity. However, we argue that in normal periods the idiosyncrasies of the individual market data channels should dominate in determining latency. We provide some evidence of low correlation between latency and dark liquidity in Figure A11.

We then employ a Two Stage Least Squares (2SLS) approach to estimating the model.

The first stage of the model is:

$$\begin{aligned} PropStaleTrades_{g,t,d,v} &= \alpha + \beta_1 LatencyInstrument_{g,t,d} + \beta_2 consideration_{g,t,d,v} + \beta_3 return_{g,t,d} \\ &+ \beta_4 VFTSE_{t,d} + \beta_5 spread_{g,t,d} + \beta_6 group_g + \beta_7 date_d + \beta_8 time_t + \beta_9 venue_v \\ &+ \varepsilon_{g,t,d,v} \end{aligned}$$

Where g represents our market data stock group, t represents our ten minute interval, d represents date, and v represents the dark venue. $PropStaleTrades_{g,t,d,v}$ measures the proportion of stale trades in the market data group for a given venue in the ten minute interval. $consideration_{g,t,d,v}$ refers to the natural log of the total value of dark trades in the same grouping. $return_{g,t,d}$ refers to the log returns over the time bucket, averaged over the stock grouping. $VFTSE_{t,d}$ is the natural log of the value of the FTSE 100 volatility index at the beginning of the ten minute interval. $spread_{g,t,d}$ refers to the time-weighted average spread in basis points of the market data stock group over the time interval.

The second stage of the model is:

$$\begin{aligned} Liquidity_{g,t,d,v} &= \alpha + \beta_1 PropStaleTrades_{g,t,d,v} + \beta_2 consideration_{g,t,d,v} + \beta_3 return_{g,t,d} + \beta_4 VFTSE_{t,d} \\ &+ \beta_5 spread_{g,t,d} + \beta_6 group_g + \beta_7 date_d + \beta_8 time_t + \beta_9 venue_v + \varepsilon_{g,t,d,v} \end{aligned}$$

Where $Liquidity_{g,t,d,v}$ is the percentage of the time interval there is marketable liquidity for dark bid and ask orders for group g , date d in venue v .

Figure A10 – Intraday Timestamp Differences to Thomson Reuters

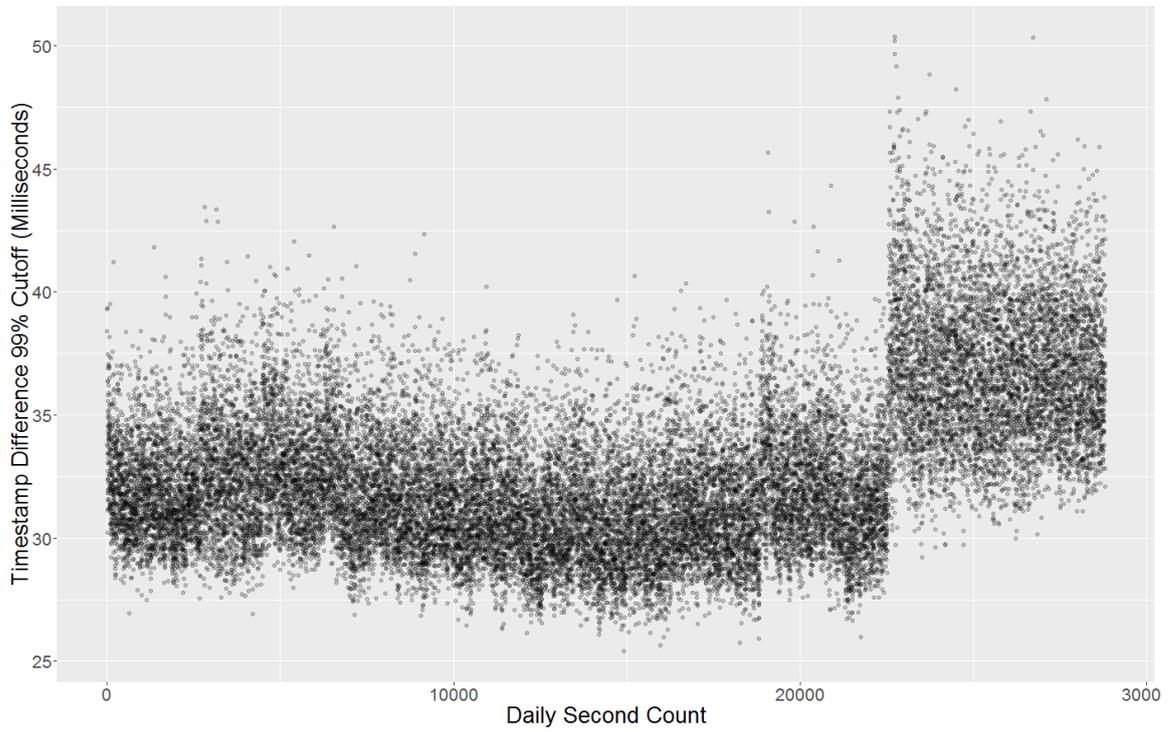
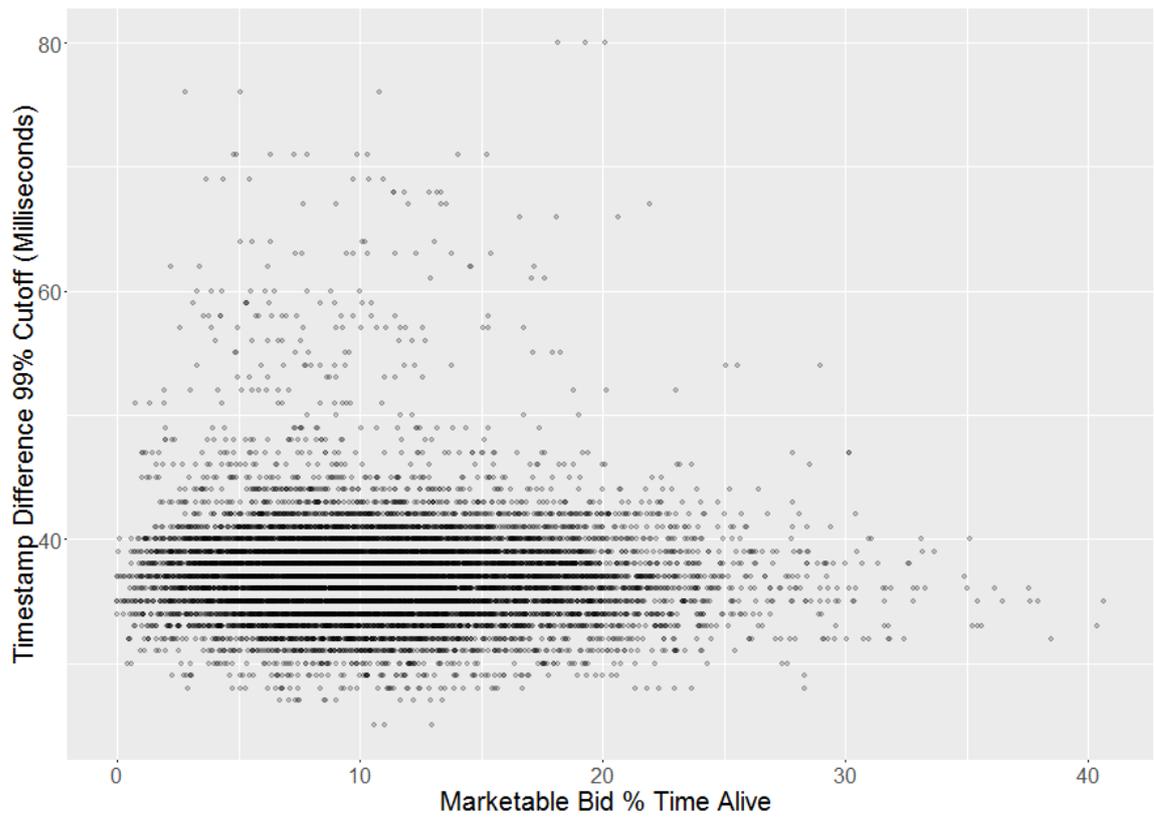


Figure A11 – Scatter Plot of Intraday Timestamp Differences and Marketable Bid Liquidity



Results

Results of the first stage regression are detailed in Table A4 below. The proportion of stale trades is highly correlated with our latency instrument, with a t -statistic of 10. This reassures us about our instrument's usefulness. We also provide evidence of the instrument's exogeneity, calculating correlation coefficients with our dependent variable liquidity measures of between -4.3% and -5.2%.

Table A4: 1st Stage IV Regression

VARIABLES	PropStaleTrades
LatencyInstrument	0.121*** (10.292)
Consideration	-0.092*** (-3.276)
Returns	0.430 (1.229)
Spread	0.003 (1.221)
VFTSE	-4.551** (-2.199)
Constant	16.592*** (3.227)
Observations	10,800
R-squared	0.334
Venue Fixed Effects	Yes
Time Fixed Effects	Yes
Date Fixed Effects	Yes
Stock Fixed Effects	Yes

Robust t -statistics in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

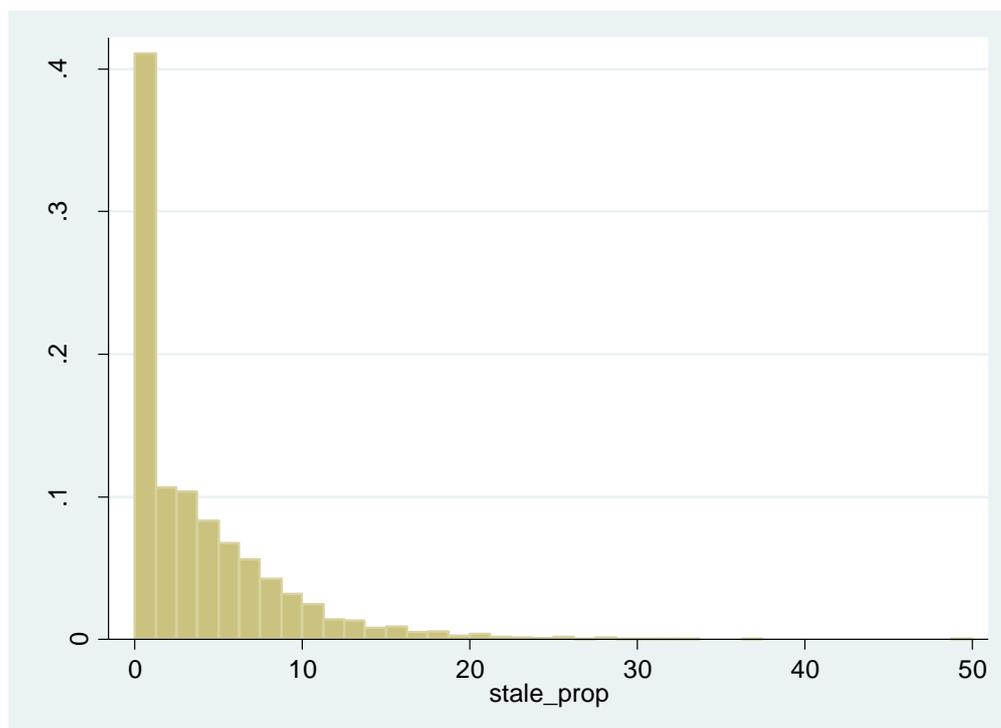
Table A5 reports the results of the second stage of the model. We estimate liquidity provision by our three participant categories for four measures of liquidity. These measures were reported previously in our descriptive statistics section, but are calculated over the ten minute buckets, rather than over the day. They measure the proportion of that ten minute period in which at least one member of the participant class has a resting marketable dark order for a given stock, date, and dark venue. This is calculated separately for bid and ask, and when this marketable order is at least as large as the total resting liquidity at the respective LSE best bid or ask, (this is denoted 'Lg Qty').

Most statistically significant results are present only for co-located participants, which could be explained by their higher sophistication, and thus their ability to detect changes in latency unperceivable to non co-located participants. There are some statistically significant results for HFT firms concerning stale trades, but the coefficients are insignificant. This may be explained by the fact that HFTs provide insignificant resting liquidity (see Figure A8).

Results show that for co-located participants, a statistically significant relationship exists between higher proportions of stale dark trades and lower liquidity provision for all measures. This implies that co-located participants may be able to observe latency in reference price feeds and respond by reducing their liquidity provision. However, the fact that co-located participants are on the losing side of 88%-92% of stale dark trades (Figure 5) implies that not all participants within this measure are able to detect this, or the aggregate level of stale dark trades we identify is already diminished by some protective measures.

Our estimates show that a 10% increase in stale trades would result in an (at most) 5.9% decrease in liquidity provision. From our earlier descriptive statistics of stale trades, wherein the proportion of stale trades ranges from 3.36% to 4.05%, an increase of this magnitude would be rare. Although our measure of stale trades, with its conservative time threshold, means it likely represents a significant subset of the true¹¹³ level of stale trades. With respect to this *true* proportion, a 10% increase may be more reasonable. But we would expect a larger correlation with our subset of *highly* stale trades than those we do not identify. Further, a significant amount of intraday variation is apparent across time, venue, group, date observations, which we illustrate in Figure A12 below. Therefore, an increase of 10% on an intraday basis is reasonable.

Figure A12: Histogram of Proportion of Stale Trades by Time, Venue, Group, Date



We also find a positive relationship with liquidity provision and the average spread of the stock group. This can be explained by the nature of midpoint dark pools in providing price improvement that is relatively more valuable with higher spreads.

¹¹³ All dark trades referencing superseded prices are technically 'stale trades,' given the time required to transmit information with a theoretical lower bound of the speed of light. A more useful theoretical definition of a 'stale trade' (and one we would use if we had more accurate timestamps) would be a stale trade referencing a price superseded by a quote update transmitted slow enough to the dark venue for a participant to observe and react to it, thus having real practical implications for market participants.

Table A5: 2nd Stage IV Regression

Participant:	Co-located				Non Co-located				HFT			
	-1	-2	-3	-4	-5	-6	-7	-8	-9	-10	-11	-12
VARIABLES	Bid	Bid (Lg Qty)	Ask	Ask (Lg Qty)	Bid	Bid (Lg Qty)	Ask	Ask (Lg Qty)	Bid	Bid (Lg Qty)	Ask	Ask (Lg Qty)
Prop Stale Trades (Instrumented)	-0.594*** (-4.570)	-0.347*** (-4.218)	-0.494*** (-3.274)	-0.158* (-1.730)	-0.051 (-0.388)	-0.116* (-1.739)	-0.016 (-0.142)	-0.011 (-0.164)	0.007 (0.438)	-0.021*** (-3.137)	0.036* (1.827)	0.024** (2.277)
Consideration	0.230*** (5.922)	0.163*** (6.756)	0.369*** (8.081)	0.289*** (9.504)	-0.176*** (-4.455)	0.018 (0.906)	0.260*** (6.199)	0.170*** (5.789)	0.005 (1.075)	0.003 (1.107)	0.011* (1.842)	0.006 (1.640)
Returns	-0.036 (-0.065)	-0.764** (-2.088)	1.036** (2.027)	0.991*** (2.895)	-0.003 (-0.005)	-0.439 (-1.381)	0.802 (1.491)	-0.163 (-0.435)	-0.033 (-0.446)	-0.012 (-0.380)	0.238** (2.371)	0.120** (2.316)
Spread	0.007* (1.762)	0.001 (0.546)	0.014*** (5.689)	0.007*** (5.017)	0.011*** (4.007)	0.005*** (4.095)	0.006*** (4.164)	0.004*** (5.085)	0.001*** (3.178)	0.000 (0.985)	0.001*** (3.505)	0.001* (1.861)
VFTSE	14.716*** (4.844)	10.473*** (5.376)	-13.814*** (-4.453)	-7.061*** (-3.841)	12.179*** (4.857)	6.567*** (4.386)	-4.521* (-1.685)	-2.163 (-1.340)	1.225*** (3.472)	0.340** (2.313)	-0.759* (-1.646)	0.288 (1.177)
Constant	-23.110*** (-2.943)	-19.492*** (-3.881)	47.272*** (5.930)	22.041*** (4.733)	-18.440*** (-2.792)	-13.667*** (-3.515)	10.671 (1.514)	3.648 (0.845)	-2.847*** (-3.166)	-0.755** (-1.987)	1.792 (1.535)	-0.911 (-1.506)
Observations	10,800	10,800	10,800	10,800	10,800	10,800	10,800	10,800	10,800	10,800	10,800	10,800
R-squared	*	0.129	0.271	0.341	0.262	0.265	0.138	0.161	0.102	0.023	0.097	0.082
Venue Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stock Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Robust t-statistics in parentheses. *R-Squared is negative in this model, but is unreliable for 2SLS.

*** p<0.01, ** p<0.05, * p<0.1

Annex 3: Further Detail on Primary Market Dislocations

Prevalence of Dislocations

We first examine how frequently dislocations occur on the lit market in our sample. We calculated these by observing lit market quotes throughout the day. The length of each discrete dislocation observation is recorded in milliseconds, which we sum, excluding dislocations one millisecond in length or less. We divide by the length of the day in milliseconds to obtain the percentage of the day a stock is dislocated. We then calculate an average for each of the 114 stocks across the 20 day sample. We report descriptive statistics of the distribution of stock-days in Table A6 below for 'BBO worse' and 'Mid worse' dislocations respectively. This demonstrates a significant amount of variation between stocks in the sample. While BBO dislocations are relatively frequent, mid-worse dislocations are relatively rare.

Table A6: Dislocation Statistics – Average of Averages by Stock

Descriptive Statistic	% of Day Dislocated BBO Worse	% of Day Dislocated Mid Worse
Mean	32.82%	2.68%
Standard Deviation	24.52%	2.25%
Minimum	3.28%	0.21%
Maximum	95.44%	11.98%

Further Results on Participant Outcomes

The following tables extend the results of Tables 7 and 8, but instead of presenting proportions of suboptimal executions across participant classes, we present the figures which characterise the worst performing participants in the participant class.

We do this by first calculating proportions for each participant-venue, and then calculating the 90% distribution cut-off (i.e. the proportion at which the worst 10% of participants exceed). We exclude firms with fewer than 100 executions.

This demonstrates a similar trend to the averages.

Table A7: % of a Participant Groups' Trades – Suboptimal – Broker Operated (90% cut-off)

Participant Group	Primary Bid or Offer Worse	Primary Mid is Worse
	BBO	Midpoint
Own Venue	0.19%	1.05%
Unmatched	1.12%	0.69%
HFT	1.13%	0.62%
Co-located	1.92%	1.00%
Non Co-located	12.56%	2.18%

This table contains the 90% distribution cut-offs of proportions of participant trades which occur at worse prices in a venue.

Table A8: % of a Participant Groups' Trades – Suboptimal – Exchange Operated (90% cut-off)

Participant Group	Primary Mid is Worse Midpoint
HFT	0.69%
Co-located	1.08%
Non Co-located	1.51%

This table contains the 90% distribution cut-offs of proportions of participant trades which occur at worse prices in a venue.

Annex 4: Technical Annex

In this Annex, we give detailed definitions of the different statistics used in this paper and explain in more detail the methodologies we implemented.

Signing trades in non-full order book data

To calculate by participant, we must sign the trades to ascertain which participant receives a benefit and which receives a loss from the reference price dislocation. Signing trades allows us to label the aggressive and passive party to each trade. We exclude midpoint trades on non-exchange operated dark pools, such as broker operated dark pools, as we cannot reliably sign trades at the midpoint without data on who initiated the trade. These flags are available for Turquoise, BATS and Chi-X dark pools for which we have full order book data.

Dark trades in dark MTFs may occur at either the best bid, best ask or midpoint on a given reference price venue. Midpoint trades are termed 'price improving' as the liquidity taker receives a price better than the BBO on the lit market. However, the liquidity taker only receives a better price if the reference exchange's mid is not outside the BBO of another market. While most dark MTFs use the LSE alone as a reference prices, Instinet Blockmatch uses what it refers to as the 'European Best Bid or Offer' (EBBO). This is the consolidated BBO prices of the major UK-based lit MTFs and primary markets.¹¹⁴

To sign trades for which we do not have full order book data (venues other than Turquoise, Chi-X and BATS Dark) we examine a window of 40 milliseconds before and after the trade and record all potential matches to the three potential price points (Bid, Ask and Mid) in the window. If the reference market contains multiple price changes in this period, there will be more than one possible match for the dark trade. For example, a dark trade may match both a midpoint and a best ask price. We are able to uniquely match 83% of broker dark pool trades to a single price.¹¹⁵ Those trades that do not uniquely match to a single price are excluded from our analysis.

Further details on data and matching methodology

Data for our broker operated dark pools comes from post-trade reports from trade reporting firms: LSE, Markit BOAT and BATS Chi-X. Historical data is then provided via TRTH.

This establishes our population of dark trades, time-stamped to the millisecond and containing venue field identifiers for dark MTFs such as UBS MTF and Goldman's Sigma X. Transaction reporting data from the FCA's Zen database include all equity market transactions for which the FCA is the relevant competent authority. This includes both counterparties to the trade in most cases.

For trades on venues not included in our full order book sample, we match Reuters trades to Zen trades to fill in participant information, from both sides where available, ignoring central clearing counterparties. We match by date, instrument, price, volume and time within 60 seconds. We remove instances of multiple trades at the same date, instrument, price, volume and 60 second rolling windows to prevent misattribution. This numbers less than 2% of trades. Our Zen coverage rates are set out in Table 5 of the main text.

As we are focused on the role of dark trading in continuous trading periods, we exclude trades within 15 minutes of opening and closing auctions. We also remove trades that are eligible for the delayed reporting regime to prevent inaccuracies in assessing price impact between reported and executed times.

¹¹⁴ Their definition of the EBBO includes: the LSE, BATS, Chi-X and Turquoise markets. This comprises around 99% of the total value traded for LSE listed stocks.

¹¹⁵ The minimum trades matched uniquely to a price for any venue is 78.3%.

Most dark trading in Europe occurs under the MiFID I reference price waiver. The primary market (the LSE) in the UK is predominantly used as the reference price, with the exception of Instinet's Blockmatch, which uses the EBBO. In referencing these markets, trades occur at the Midpoint of the BBO or at the BBO itself (referred to as midpoint pegs or BBO pegs respectively). Midpoint trades are often referred to as 'Price Improving' trades because they save liquidity-demanding participants half the quoted spread.

Annex 5: References

ASIC (2015). Review of high-frequency trading and dark liquidity. Report 452.

Aitken, Michael, Douglas Cumming, and Feng Zhan. 2014. 'Trade Size, High-Frequency Trading, and Colocation around the World.' *The European Journal of Finance* 0 (0): 1–21. doi:10.1080/1351847X.2014.917119.

Aquilina, Matteo and Ysusi, Carla. 2016. 'Are high-frequency traders anticipating the order flow? Cross-venue evidence from the UK market.' FCA Occasional Papers. www.fca.org.uk/publication/occasional-papers/occasional-paper-16.pdf.

Anderson, Lisa, Baiju Devani, and Yifan Zhang. 2016. 'The Hidden Cost: Reference Price Latencies.' Investment Industry Regulatory Organization of Canada. Investment Industry Regulatory Organization of Canada. www.iroc.ca/Documents/2016/f1e06451-0f2b-4af8-a973-81fa881a6958_en.pdf.

Banks, E. 2014. *Dark Pools: Off-Exchange Liquidity in an Era of High Frequency, Program, and Algorithmic Trading*. Springer.

Budish, Eric, Peter Cramton, and John Shim. 2015. 'The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response.' *The Quarterly Journal of Economics* 130 (4): 1547–1621. doi:10.1093/qje/qjv027.

Bartlett, Robert P., and Justin McCrary. 2016. 'How Rigged Are Stock Markets? Evidence from Microsecond Timestamps.' SSRN Scholarly Paper ID 2812123. Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=2812123>.

Bekaert, Geert, and Guojun Wu. 2000. 'Asymmetric Volatility and Risk in Equity Markets.' *Review of Financial Studies* 13 (1): 1–42. doi:10.1093/rfs/13.1.1.

Boehmer, Ekkehart, Kingsley Y. L. Fong, and Juan (Julie) Wu. 2015. 'International Evidence on Algorithmic Trading.' SSRN Scholarly Paper ID 2022034. Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=2022034>.

Brogaard, Jonathan, Björn Hagströmer, Lars Nordén, and Ryan Riordan. 2015. 'Trading Fast and Slow: Colocation and Liquidity.' *Review of Financial Studies*, August, hhv045. doi:10.1093/rfs/hhv045.

Chen, Hao Ming, Sean Foley, Michael Goldstein, and Thomas Ruf. 2016. 'The Value of a Millisecond: Structural Segmentation of Uninformed Order Flow.' Working Paper. www.sec.gov/comments/10-222/10222-412.pdf.

Conrad, Jennifer, Sunil Wahal, and Jin Xiang. 2015. 'High-Frequency Quoting, Trading, and the Efficiency of Prices.' *Journal of Financial Economics* 116 (2): 271–91. doi:10.1016/j.jfineco.2015.02.008.

Foucault, Thierry, and Albert J. Menkveld. 2008. 'Competition for Order Flow and Smart Order Routing Systems.' *The Journal of Finance* 63 (1): 119–58. doi:10.1111/j.1540-6261.2008.01312.x.

Frino, Alex, Vito Mollica, and Robert I. Webb. 2014. 'The Impact of Co-Location of Securities Exchanges' and Traders' Computer Servers on Market Liquidity.' *Journal of Futures Markets* 34 (1): 20–33. doi:10.1002/fut.21631.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. 'Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders.' *Journal of Financial Economics* 14 (1): 71–100. doi:10.1016/0304-405X(85)90044-3.

- Hagströmer, Björn, and Lars Nordén. 2013. 'The Diversity of High-Frequency Traders.' *Journal of Financial Markets, High-Frequency Trading*, 16 (4): 741–70. doi:10.1016/j.finmar.2013.05.009.
- Hasbrouck, Joel. 2015. 'High Frequency Quoting: Short-Term Volatility in Bids and Offers.' Available at SSRN 2237499. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2237499.
- Kirilenko, Andrei A., Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. 2015. 'The Flash Crash: The Impact of High Frequency Trading on an Electronic Market.' SSRN Scholarly Paper ID 1686004. Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=1686004>.
- Kwan, Amy, Ronald W. Masulis, and Thomas H. McInish. 2015. 'Trading Rules, Competition for Order Flow and Market Fragmentation.' *Journal of Financial Economics* 115 (2).
- Malinova, Katya, and Andreas Park. 2016. "Modern' Market Makers.' Working Paper. <http://firn.org.au/wp-content/uploads/2016/05/Modern-Market-Makers-Park-Malinova.pdf>.
- Marouani, Hicham, Michel R. Dagenais, Hicham Marouani, and Michel R. Dagenais. 2008. 'Internal Clock Drift Estimation in Computer Clusters, Internal Clock Drift Estimation in Computer Clusters.' *Journal of Computer Networks and Communications, Journal of Computer Networks and Communications* 2008, 2008 (May): e583162. doi:10.1155/2008/583162, 10.1155/2008/583162.
- Menkveld, Albert J. and Zoican, Marius A., 'Need for Speed? Exchange Latency and Liquidity', 2016. Tinbergen Institute Discussion Paper 14-097/IV/DSF78. Available at SSRN: ssrn.com/abstract=2442690
- SEC. 2016a. 'ADMINISTRATIVE PROCEEDING File No. 3-17077.' www.sec.gov/litigation/admin/2016/33-10010.pdf.
- . 2016b. 'ADMINISTRATIVE PROCEEDING File No. 3-17079.' www.sec.gov/litigation/admin/2016/33-10014.pdf.
- Sommer, Philipp, and Roger Wattenhofer. 2008. 'Symmetric Clock Synchronization in Sensor Networks.' In *Proceedings of the Workshop on Real-World Wireless Sensor Networks*, 11–15. ACM. <http://dl.acm.org/citation.cfm?id=1435477>.
- Ye, Mao, Chen Yao, and Jiading Gai. 2013. 'The Externalities of High Frequency Trading.' SSRN Scholarly Paper ID 2066839. Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=2066839>.

Glossary

Definitions in this glossary are provided solely for the convenience of readers of this report. They are not presented as approved regulatory definitions or to be used for any other purpose.

Aggregator – A service operator that decides which dark pool or other trading venue through which to route an order on behalf of a client.

Algorithm – A specific set of clearly defined instructions programmed into a computer to execute a trade in a certain manner.

Broker crossing network (BCN) – A subset of an investment bank operator's electronic platform where third-party orders can be matched anonymously using reference prices taken from selected lit markets. Under MiFID, trading under a BCN would fall under OTC trading. OTC is defined in relation to a transaction in an investment, not on-exchange.

Child order – A subsection of a parent order, sent to market at a particular time.

Co-location – The practice of placing a market participant's servers in close physical proximity to an exchange's to reduce transmission latency.

Dark pool market or venue – A trading platform with no pre-trade transparency, wherein all resting liquidity is hidden with respect to price and volume.

Direct market access (DMA) – Direct electronic access to an exchange provided to clients using a broker-dealer's IT infrastructure.

EBBO – The 'European Best Bid and Offer' is a composite of the best prices available for buying or selling a stock from a selected number of European trading venues.

High-frequency trading (HFT) – Market participants that use proprietary capital to generate returns using computer algorithms and low-latency infrastructure. This description is not to be confused with the definition in the delegated acts underpinning MiFID II published by the European Commission on 25 April 2016.

Latency – The time that elapses from when a signal is sent to when it is received. Lower latency means lower delays in transmission.

Lit market or venue – Where the order book is visible to all members, so that traders can see the amount of liquidity available on the bid and offer. Examples include the London Stock Exchange and the order books of BATS and Turquoise that have pre-trade transparency (lit order books).

MiFID / MiFIR – The Markets in Financial Instruments Directive is the framework of EU legislation for the organised trading of financial instruments, and MiFIR is the related regulation. MiFID was first implemented in 2007 and is being comprehensively revised (MiFID II), with the changes expected to take effect from January 2018.

Multilateral trading facility (MTF) – A multilateral system, operated by an investment firm or a market operator, which brings together multiple third-party buying and selling interests in financial instruments (in the system and in accordance with non-discretionary rules) in a way that results in a contract in accordance with the provisions of Title II of MiFID.

Operator – The sponsor or business owner of a dark pool or platform. **Parent order** – A larger order from which a number of child orders are split and routed separately to be executed in the market.

PBBO – The ‘Primary Best Bid and Offer’ is the best price available for buying or selling a stock from an individual European primary trading venue.

Price impact – the tendency of share prices to react in the direction of a trade in response to liquidity and information effects.

Principal/proprietary flow – In the context of an operator, this refers to order flow arising from its internal activity, such as hedge unwinds, central risk book or proprietary trade positions.

Reference price waiver – A waiver from pre-trade transparency whereby a system satisfies the criteria that ‘they must be based on a trading methodology by which the price is determined in accordance with a reference price generated by another system, where that reference price is widely published and is regarded generally by market participants as a reliable reference price.’

Regulated market – A multilateral system operated and/or managed by a market operator, which brings together or facilitates the bringing together of multiple third-party buying and selling interests in financial instruments (in the system and in accordance with its non-discretionary rules) in a way that results in a contract, in respect of the financial instruments admitted to trading under its rules and/or systems, and which is authorised and functions regularly and in accordance with the provisions of Title III of MiFID. In the UK, a regulated market can only be operated by an RIE.

Resting order – A non-executed order sitting on the order book.

Resting time – The period of time an order is left on an order book before being executed, automatically expiring or being withdrawn.

Smart order router (SOR) – A computer – or algorithm-assisted process – used in electronic trading to send order instructions to an exchange or trading market following a defined set of rules.

Stale reference price – A reference price that is not the most recent price. For dark pools, this means a reference price superseded by a newer price that has not yet reached the dark pool.

Stale trade – A trade in a dark pool that occurs at a stale reference price.

Transaction cost analysis (TCA) – the practice of measuring the effectiveness of trades. TCA provides analysis of how a trade has performed when compared to a particular benchmark and may include adverse price movements during the timeframe taken to complete a trade.

