Dark Pool Reference Price Latency Arbitrage*

Matteo Aquilina,[†] Sean Foley,[‡] Peter O'Neill[§] and Thomas Ruf[¶]

May 10, 2017

Abstract

Using proprietary order book data with participant-level message traffic and matching engine time stamps, we investigate the nature of stale reference pricing in dark pools. We document a sizeable proportion of stale trading which imposes large adverse selection on the passive side in dark pools. We are the first to document that HFTs almost never provide liquidity and instead frequently take liquidity in the dark, in particular in order to take advantage of stale quotes in the dark. Finally, we examine several market design interventions to mitigate stale trades, showing that only mechanisms to protect passive dark liquidity, such as random uncrossings, are effective.

Keywords: high-frequency trading, dark pools, latency arbitrage, reference prices

*We would like to thank Michael Aitken, Peter Andrews, Hao Ming Chen, Brian Eyles, Simon Hargreaves, Edwin Schooling Latter, Ted Macdonald, Albert Menkveld, Richard Payne, Talis Putnins, Jia Shao, Patrick Spens, Felix Suntheim, Martin Taylor, Carla Ysusi and Bart Z. Yueshen. We would also like to thank our colleagues at the FCA. We are also grateful to participants of the FCA/LSE conference on financial regulation for their comments. Any errors and omissions are our own. This article is an extension of the FCA Occasional Paper 21 'Asymmetries in Dark Pool Reference Prices.'

[†]Financial Conduct Authority (UK).

[‡]University of Sydney.

[§]Financial Conduct Authority (UK) and University of New South Wales.
e-mail:*peter.oneill@fca.org.uk* Peter O'Neill thanks the Capital Markets Cooperative Research Center for funding a period as a visiting researcher at the FCA during this study.

[¶]University of New South Wales.

1 Introduction

Financial markets have been subject to increasing competition on speed, which now governs investor trading costs and liquidity. In this new world, fast traders impose adverse selection costs on slower ones by acting on information signals faster than their competitors. These signals are not always clear, resulting in competition on interpreting information, as well as reacting to it (Budish et al., 2015). Where these signals are symmetrically interpretable, a race occurs on reaction times alone. These have been labeled 'toxic arbitrage' (Foucault et al., 2016) and 'latency arbitrage' (Bartlett and McCrary, 2016) and are argued to be harmful to liquidity in financial markets by increasing adverse selection to liquidity providers.

Dark pools are uniquely susceptible to latency arbitrage by relying on feeds from lit markets to price trades. Changes in lit prices create opportunities for fast participants to pick off stale quotes before they can update. Reference pricing creates a mechanical relationship, resulting in an 'old price' and a 'new price', a situation which does not exist on lit venues. While prices may vary between venues, there is a degree of uncertainty in their convergence or reversal. Dark pool reference prices thus create an ideal setting to observe latency arbitrage as signals have very low uncertainty.

Dark pools aim to lower the trading costs of institutional investors. The reliance on reference pricing can create latency derived adverse selection costs. We characterize the extent of stale trades in dark pools, and the impact it has on execution quality. We link this with the liquidity provision strategies of participants in dark pools, as the first study to provide orderbook information in dark pools using matching engine timestamps. Finally, we examine the effectiveness of several existing market design solutions that aim to address the problem of reference price latency arbitrage, finding batch auctions and random uncrossings are effective at reducing such conduct.

This paper has important implications for the design of modern markets in relation to the merits of 'slowing down' markets, such as the controversial IEX 'speedbump', batch auctions and minimum resting times. We confirm the predictions of theoretical models of high frequency arms races (Foucault et al., 2016). We show that HFT participants impose adverse selection on slower participants and market design interventions, which obstruct the speed race, reduce this adverse selection.

Dark pools differ from 'lit' markets in offering no pre-trade transparency, matching orders using reference prices from lit markets data feeds.[1] Two forms of delay or latency in referencing this primary market exist, resulting in potential costs to investors; processing latency and transmission latency. Where markets are in the same physical location or data center, a delay exists between the hardware and software responsible for calculating and disseminating the market data (processing latency).[2] When these two markets are in different physical locations, the time it takes to transmit price date creates an additional delay (transmission latency).[3]

These two sources of latency engender differences in speed amongst participants. To eliminate these, innovations have been proposed in modern markets.[4]In the US, the exchange and dark pool IEX has designed an 'inbound speed-bump' to prevent latency arbitrage arising from latency in its reference price calculation. The popularity of such innovations is evidenced by their recent introduction by other exchanges, such as the Chicago Stock Exchange, the Alpha Exchange in Canada[5] and recently, NYSE. IEX's speed-bump delays inbound and outbound orders by 350 microseconds (IEX, 2015). Importantly, it does not apply to the repricing of dark orders. Consequently, as long as the reference price feed is not stale by more than 350 microseconds, reference price latency arbitrage is prevented.[6]

[1]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' and external pricing sources, to enable dark trading.

[2]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.

[3]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 to travel 16 miles, a comparable distance to LSE and BATS in the UK. This significantly exceeds their 'theoretical minimum time' of 86 microseconds.

[4]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)

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

[6]This is acknowledged in competing exchange BATS' submission to the SEC regarding IEX's application to become an exchange: 'this 350 microsecond delay provides IEX the ability to update

The LSE has recently introduced a new Field Programmable Gate Array (FPGA)[7] product to reduce processing latency to 'sub-five microseconds'.[8] Several microwave networks now carry data between national and global exchanges, significantly reducing *transmission latency* compared with fiber optic networks.

While some latency (both processing and transmission) is unavoidable, in well-functioning markets both types of latency should be reduced to a minimum as latency can give rise to arbitrage opportunities. Foucault et al. (2016) model these 'asynchronous price adjustments which give rise to toxic arbitrage opportunities'. They show that dealer adverse selection risk, and thus liquidity, is a function of the 'arbitrage mix', the proportion of toxic vs non-toxic arbitrages, and the arbitrager or sniper's speed relative to the liquidity provider's. We show that the proportion of stale/toxic trades is significant, with aggressive snipers being predominantly HFT participants. This has an adverse impact on liquidity provision; we show that HFTs almost never provide marketable liquidity in dark pools.

According to publicly available information, the transmission latency between LSE's data center in central London and BATS's in Slough is about 320 microseconds and between the LSE and Turquoise (60 microseconds).[9] There is additional time to process incoming order messages, quoted by exchanges to be about 25 microseconds.[10] The latency of pre-trade risk checks, which must be performed on incoming orders, is also often quoted. A competitive market exists in minimizing this latency. LSE quotes a reduction of 2–3 microseconds to below 0.5 microseconds in a recent exchange upgrade (LSEG, 2013, p. 6). When messages spike, as is common within a millisecond, this infrastructure hits bandwidth/throughput constraints,

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' (Swanson, 2015, p. 1).

[7]FPGAs are a type of computer architecture in which the programming is in the hardware chip rather than the software, reducing processing latency significantly (LSEG, 2016).

[8]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.

[9]This is according to one market data vendorm Standard and Poors Capital IQ (2015). These figures will vary across providers, and technologies used. Microwave connections are reported to be 30 to 40% faster than fiber.

[10]LSEG (2013).

and messages get queued. This can increase latency by many multiples.[11]

The importance of speed in governing how financial markets operate has been demonstrated by recent actions by regulators and practitioners as well as academics. Latency in Goldman Sachs's US dark pool was significant enough to justify a fine of $800k levied by FINRA in 2014 (FINRA, 2014). It paid $1.2m to clients as compensation for losses stemming from 395,000 stale trades (Hope, 2014). IIROC, the Canadian securities regulator, recently published research showing that 4% of all dark pool trades in Canada occur at stale prices, documenting high variation across venues, with the overall proportion of latency-affected trades by value increasing over time, as well as in duration (Anderson et al., 2016). The Tabb Group consultancy analyzed ten months of trading data 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 (Alexander et al., 2015). 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.

Speed's importance is such that fines have been levied for misprepresenting it to clients. Barclays was fined by the SEC for claiming it was using faster direct feeds to price dark pool trades despite actually using slower 'SIP' feeds.[12] The Financial Conduct Authority of the UK (FCA)'s 2016 Thematic Review of Dark Pools expressed concern that some dark pool operators only monitored latency in pricing feeds irregularly, or only on a post-trade basis (FCA, 2016, p. 35).

Trading venues have begun to introduce features to address latency issues. The UK based Turquoise has introduced a 'random uncross' feature, which it says is beneficial for 'latency sensitive flow' (LSEG, 2017). Others, like Deutsche Bank's Super X, state that they will stop matching if orders are stale by more than a second

[11] 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, 2014, p. 7). This characteristic of latency 'spiking,' also called 'jitter' requires the measurement of latency in terms of percentiles, to capture the behavior 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 (LSEG, 2015, p. 5).

[12] (SEC, 2016a);(SEC, 2016b)

Deutsche Bank (2016). Regardless, many dark pool operators disclose information on latency and how it is managed.[13] The Australian securities regulator (ASIC) recently published a report on dark trading and HFT in Australia. They found that HFTs were on the 'winning side' of trades in dark pools that referenced stale prices 85% of the time compared with 31–32% for other users, (ASIC, 2015, p. 54).[14]

In Section 2 of provide the institutional details of dark pools. In Section 3 we detail our data and explain our method. In Section 4 we examine the prevalence of dark pool trades at stale reference prices, the cost to investors, whether this has changed over time and what causes stale reference prices. We then examine the impact of stale reference prices on adverse selection in Section 6 and we characterize liquidity provision by various classes of participants in dark pools. Finally, in Section 7 we examine several market design interventions aimed at combating adverse selection from stale reference prices.

2 Institutional Details

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 trader's intention is not revealed to the entire market. Another advantage is trading at a better price than is available on the lit market (price improvement). The main disadvantage of dark pools is execution uncertainty.[15] Specifically, it is impossible to know whether there is a willing counterparty, so one cannot know beforehand whether a trade will take place. Orders sent to

[13]For example, Goldman Sach's compares the timestamps in 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' (GSEC, 2016, p. 3).

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

[15]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.

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).[16] However, within the boundaries of 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 in lit markets. 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 typically 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 London Stock Exchange (LSE). The second option is to consider multiple (lit) venues. In the first case, dark pools use Best Bid/Offer (BBO) prices available on the LSE. In the second case, operators construct what is known as the European BBO (EBBO), which includes orders from all other venues.[17]

Having constructed the BBO or the EBBO, dark pools have to choose to match prices at the midpoint, the BBO, or both. 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 ask price (depending on the direction of the trade). Many jurisdictions have recently prohibited these non-price improving trades[18]. Under MiFID II, these will also be prohibited.[19] 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.

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

[17] 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).

[18] See Foley and Putniņš (2016) for an examination of this in Canada and Australia.

[19] MiFID II Article 4(1)(a) and (2).

Dark pools in the UK and Europe, from a regulatory perspective, are classed as either multilateral trading facilities (MTFs) or broker crossing networks (BCNs). MTFs fall within the scope of MiFID, require regulatory approval, and are in practice independent venues. This makes them subject to requirements of non-discriminatory access, with transparent rulebooks, and usually central clearing (CCPs). BCNs are not regulated under MiFID and are non-independent crossing networks under a broker's direct control. They are thus allowed to discriminate by participant types, with rules disclosed in bilateral contracts. MiFID II will ban BCN venues, and operators will be forced to register as MTFs, or Systematic Internalizers (SIs). In the US, the 'Alternative Trading System' (ATS) framework would apply to MTFs, but also BCNs.

We characterize dark pools into three subsets. MTFs that are operated by the lit exchanges Turquoise and BATS Europe.[20] which are integrated with the lit exchange infrastructure, and matched trades only at the midpoint. We have orders and trades with participant identifiers for these venues. Secondly, MTFs (such as Posit, SigmaX and UBS) that are operated by investment banks and other brokers and match trades at the BBO as well as the midpoint. We have less information on them in our data (only trades, not orders). Third, BCNs (such as 'Crossfinder' and 'SuperX') for which we cannot determine specific venues, as MiFID allows reporting simply as OTC (Over the Counter) trades. Therefore, these are excluded from our sample.

3 Data and Methodology

3.1 Data

We use two datasets in our analysis. Proprietary full order book data and transaction data from Thomson Reuters Tick History (TRTH). Our order book data was collected by the FCA directly from trading venues for market monitoring and research purposes. It includes all participant identified trades timestamped at the matching

[20]Turquoise is technically operated by an investment firm, Turquoise Global Holdings, but is is majority owned by London Stock Exchange Group (LSEG).

engine for four UK trading venues for five separate one week periods covering January 2014 to June 2015. The trading venues covered are the LSE, BATS, Chi-X[21] and Turquoise. The dataset includes all messages (entered/amended) and trades in both lit and dark orderbooks. These venues account for 99.56%[22] of all FTSE 350 on-exchange lit traded volume in the UK, and all exchange operated dark MTFs. This gives us a representative sample of lit and dark trading for UK listed stocks. Our data excludes smaller venues such as Equiduct and Aquis, but these are small enough that they are not included in the definition of the consolidated BBO. 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 unanonymized identity of each member of the trading venue behind each message. In other words, the order book is not anonymized.

For broker operated dark MTFs, and US dark pool data, we obtain millisecond time-stamped post-trade reported dark pool trades from Thomson Reuters Tick History. These include UBS MTF, Sigma X MTF, ITG Posit and Instinet Blockmatch. We exclude Liquidnet, as reference prices are determined through bilateral negotiation rather than the MiFID reference price waiver. 'Smartpool' and 'Blink MTF' are also excluded as they have de minimis volumes. Our remaining sample covers 93.97%[23] of dark MTF (exchange and broker operated) trading.

As BCNs are unregulated venues under MiFID I, post-trade reporting does not require these venues to be named so they are reported as 'OTC' trades. We exclude these from our analysis due to the inability to separate OTC trades organized on a BCN from other OTC trades.

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

[22]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 (2017).

[23]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, 2014).

3.2 Sample Composition

Our analysis uses a randomly selected sample of 57 stocks from the FTSE 100 and 57 from the FTSE 250. These two indices represent the FTSE 350, and are chosen to obtain stocks with both high and low liquidity. 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. These are the continuous five-day trading weeks starting: 13th of Jan, 31st of March, 16th of June and 1st of September in 2014 and the 22nd of June in 2015.

3.3 HFT Definition

Our data allows us to identify participants at the firm level, and not at the trading desk or client level. Traders are divided into 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. For example, an HFT engaged in predominantly liquidity consuming (aggressive) trading strategies, will have a lower order-to-trade ratio than an HFT engaged in liquidity providing (passive) market-making strategies.

Therefore, we use the FCA's internal supervisory knowledge, and the 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. (Baron et al., 2016, p. 37) list a subset of those we identify by name as Sweden allows non-anonymized exchange trading. In our sample, we observe 30 participants that we classify as HFTs.

3.4 Co-located Participants

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 non-HFT participants. To examine the role of these speed differences in determining trading outcomes, we divide non-HFT participants into fast and slow categories depending on whether they are co-located at the primary exchange.[24] 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 emanate from co-located servers. Many participants that are not HFTs are also co-located.

4 Extent of Stale Reference Prices

The methodology used to identify stale reference prices is based on that used by Anderson et al. (2016) and ASIC (2015) and is only applied to exchange operated venues for which we have accurate matching engine timestamps. To identify stale reference prices, we first look for quotes on the primary market that match the price of the dark pool trade.[25] As trades in our sample may only reference the midpoint, we look for quotes where the midpoint of the primary market matches the price of the trade.

This identification may give us multiple matches in our window around the dark pool trade. As we are unable to determine which was actually referenced, for conservatism, we assume it is the most recent quote, and thus we underestimate staleness. Our window of observation extends as far back as necessary to observe a match,

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

[25] Dark pools will round midpoint prices to four decimal places and we account for this in calculating matches, but we have no cases of such rounding in our sample.

but also one millisecond *after* the trade time, so as to allow for exchange clock non-synchronicity.[26] To determine a trade as stale, we also require at least one quote update before the dark trade time that does not match the dark trade price, the *intervening non-match*. This newer price allows us to determine that the dark pool is referencing the older, 'stale price'.[27] We ensure this quote update occurs after the matching quote, and is indeed intervening, by using the message sequencing number[28] from the primary market, which allows us to establish sequencing *within* milliseconds. To calculate the magnitude of latency for a stale trade we take the most recent timestamp of the quote that matched the dark trade subtracted by the timestamp of the dark pool trade, i.

$$Latency_i = TradeTime_i - MatchTime_i$$

We consider trades occurring with *latency* of two milliseconds and above as stale, recognizing that a one millisecond threshold would not allow for clock synchronicity and timestamp rounding effects. This methodology is similar to that of Anderson et al. (2016), except we require latency to be two milliseconds and above, rather than one, and the dark pools in our sample reference only the primary exchange, making identifying matches more certain and diminishing clock synchronization issues inherent in timestamps from different markets.

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

[26] This is because one millisecond reflects the upper bound of clock synchronization accuracy provided by the exchanges, and the minimum granularity of timestamps provided to us.

[27] This means that we underestimate stale reference prices because a dark pool may be referencing a quote update from many milliseconds ago, but we are unable to determine if it is stale in the absence of a newer intervening quote.

[28] 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.

exists only without price variation.

All dark trades referencing superseded prices are in theory, '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' 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. A price that is stale in practice, as well as in theory. We proxy for these stale trades we cannot observe by identifying cases where HFTs are on the 'winning side' of price movements around dark trades.

4.1 Costs of Stale Trades

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 counter factual of the reference price not being stale, assuming that the dark trade would have otherwise occurred. However, there is a possibility that the trade occurred precisely because it was at a stale price.

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

If latency is evenly distributed across participants, the percentage of trades for which a participant benefits or loses, should be random, with mean zero. That is, a participant is expected to be as likely to benefit or suffer from reference price latency. The buyer (seller) to the trade benefits from the reference price latency if the trade price is less (greater) than the prevailing mid.

$$Benefit_i = \begin{cases} Buyer, TradePrice_i < PrimaryMidPrice_i \\ Seller, TradePrice_i > PrimaryMidPrice_i \end{cases} \tag{1}$$

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.

4.2 Measures of Adverse Selection

Budish et al. (2015) document an inherent asymmetry between liquidity providers and liquidity takers on a conventional lit limit orderbook, where in the presence of new information, liquidity providers race liquidity demanders to cancel or hit their now stale quotes. To demonstrate this, they provide the example of price changes in the SPI-E-Mini in Chicago and markets based in New Jersey which trade the same underlying index, the SPI ETF. But this is not a setting with symmetric information. A jump in the E-Mini Future does not necessitate a jump in the SPI ETF if it consists of a jump in prices arising from a liquidity shock that subsequently reverts. But the presence of a dark pool reference price feed, creates a setting with symmetric information. The reference price feed allows participants to transact at a stale/old reference price, while observing (and transacting at) the new price without any necessary assumptions about the instrument convergence or reversion. This means that all dark orders are susceptible to 'toxic' arbitrage from sufficiently large price movements on the lit.

We calculate price impact, as it is a commonly used measure of execution quality, particularly for agency brokers and institutional investors assessing trading performance. This allows us to demonstrate the effect stale prices have on calculated price impact measures.

Traditionally, price impact is the extent to which prices respond to buying or selling pressure; with less price impact being desired by liquidity demanders. In the context of stale reference prices, the mechanism is not a response to orders, but a response to stale prices. If the primary midpoint moves down by £1 and a participant has a resting buy order in the dark pool that is still referencing the old price, a stale reference price arbitrageur may submit an aggressive IOC sell order at the old price, resulting in a virtually instantaneous £1 positive price impact, milliseconds later. However, with data feeds that are observable with zero delay, this is *not* price impact, but a trade with non-zero effective spread, where it would normally be zero. From the perspective of participants and academics that observe market data with non-zero latency, this is price impact, and from the participant's perspective it is adverse selection. Therefore, we treat them as such, but we do not ourselves have the timestamp granularity to separate the two to the fullest extent possible. We calculate price impact as:

$$PriceImpact_{it} = q_{it}(m_{i,t+m} - m_{i,t}/m_{i,t}) \tag{2}$$

We report results for the time period of 100 milliseconds, which we view as a time period which captures 'instantaneous adverse selection' effects, which result from latency arbitrage activity. Recently there has been appreciation of the immediacy at which adverse can occur in microstructure research. Chen et al. (2016) examines adverse selection over a 20 millisecond horizon, Wahal and Conrad (2016) examines the decreasing time horizons of price impact; down to 100 milliseconds, and Menkveld (2016) examines within millisecond trading interactions.

In Section 7, where conduct event studies on interventions to reduce reference price latency arbitrage, we calculate price impact (adverse selection) as the absolute value of the price change, removing q because we rely on public trade reports which do not have initiator flags. Calculating adverse selection in this manner is common amongst practitioners, such as ITG (Saraiya and Mittal, 2009) and IEX (Aisen, 2015). In academic research, Boni et al. (2013) examine volatility ahead of dark pool trades, and Nimalendran and Ray (2014) use changes in the spread.

5 Results: Extent and Nature of Stale Trading

5.1 Extent of Stale Trades

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. To estimate the costs associated with stale reference prices for all dark pools we scale up the estimates obtained for our subset using data for all venues, assuming stale trades remain a constant proportion.

Figure 1 details the percentage of stale trades across all stocks and markets for each date in our sample. 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 (Anderson et al., 2016). 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. In practice, 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 time 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. 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 a dark venue. Our timestamps do not allow us to observe these stale trades but we expect there to be a significant number. 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 (Standard and Poors Capital IQ, 2015).

[Figure 1 about here.]

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 increases our ability to observe stale reference prices by increasing the amount of price changes. Our methodology requires sufficient price variation prior to dark trades to observe stale reference prices.

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 about here.]

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 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.

5.2 Stale Trades 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. The results are reported in Table 1. 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 positive 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 messages, 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. This means that the level of stale dark trades is increasing over time, but this is explained by increasing message volumes and volatility. Future research with newer samples can confirm whether this trend persists.

[Table 1 about here.]

5.3 Cause of Stale Trades

The resiliency of financial markets is now a major concern following major events such as the 2010 'flash crash'. Menkveld and Yueshen (2016) provides evidence that this crash was prompted by a deterioration in liquidity in Emini Futures contracts which arose from a rare breakdown in the arbitrage relationship with SPI contracts. They demonstrate that delays in market data feeds during this period may have contributed to this breakdown. The importance of processing latency in markets is also established by Kirilenko and Lamacie (2015). They find that processing latency predicts the volatility of asset prices, using timestamps from within an exchange matching engine. Exchanges have also recently been forthcoming about this relationship. The Eurex exchange states that large bursts of incoming messages can delay the dissemination of quote updates by many milliseconds, for a periods lasting over 3 milliseconds (Eurex, 2016, p. 28). The CME has recently introduced 'First in First Out' priority at the exchange gateway to ensure equitable outcomes amongst participants due to variances in processing time CME (2014). We provide evidence on the resiliency of exchange infrastructure by examining the relationship between market-wide message volumes around stale trades using probit regressions. The model specification is as follows: $StaleTrade_{isdv} = TotMktMessages_{isdt} + venue + date + \epsilon_{isdv}$ $StaleTrade_{isdv}$ is a dummy variable with value one for a dark trade which is stale, and zero for a trade which is not stale, for trade i, stock s, date d, venue v. $TotMktMessages_{isdt}$ is the change in the sum of the total number of messages across the 400 largest securities listed in the UK, for all of the markets in our sample for a discrete two millisecond time interval, relative to the previous time interval. This is calculated for trade i, stock s, date d and time t where t is fourteen consecutive two millisecond time intervals before and five intervals after the trade. We use two millisecond intervals to harmonize any timestamp differences between venues at the millisecond level. We subtract all messages that pertain to 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 **??**. We find statistically significant positive relationships with total market messages up to 28 milliseconds before stale trades, but find no relationships in the periods afterward. The strongest statistical significance occurs in the periods two milliseconds before, and in the same period as as the stale trade. This provides evidence that stale trades may be caused, at least in part, by exchange infrastructure not having enough capacity to process large influxes in messages without inducing latency.

[Table 2 about here.]

5.4 Cost of Stale Trades

Table 3 presents the costs of stale reference prices for the exchange operated dark pools in our sample being 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. As a comparison we consider the trading revenues of some of the largest HFTs operating in the UK. Knight Capital Group Europe's trading revenue for 2014 was $83.18m and Jump Trading's gross revenue for 2014 was $97.1m.[29] 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.[30]. 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.

[Table 3 about here.]

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

[29] International (2015) and Europe (2015).

[30] Deutsche Bank Europe's Super X broker crossing network (Deutsche Bank, 2016)

to the adverse selection costs of stale reference prices. We empirically examine this in Section 6.1.

5.5 Are costs borne 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 benefit from stale prices, losing only 3.6% of the time (due to negligible HFT-on-HFT trading).[31] These results are displayed in Figure 5, which report the proportion of a given participant group's trades that are on the benefit side, and thus do not add to 100% because each participant groups are involved in a subset of all trades. Our findings are consistent with Baron et al. (2016) who find that HFTs profit by 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.

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 and are unaffected by the latency associated with the reference price calculation. This could be because they are using a different (faster) feed from the primary market than the dark pool utilizes.

[Figure 3 about here.]

[31] 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. 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.

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 8 and 9 below demonstrate that for most executions at stale reference prices, the benefit side is that of the aggressor.

[Figure 4 about here.]

Figure 8 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 9 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 our examination of liquidity provision in dark pools in Section 6.1 where we show 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. We describe passive stale execution strategies in Section 6.

[Figure 5 about here.]

6 Execution Quality and Dark Liquidity

In this section, we examine the impact of stale reference prices on price impact as a measure of adverse selection costs. We follow the approach in Malinova and Park (2016) and carry out trade by trade regressions, utilizing similar controls. We run the following OLS regressions with standard errors clustered at the security and date level.

$$\begin{aligned} PriceImpact_{istd+m} = \alpha &+ \beta_1 Stale_{istd} + \beta_2 consideration_{istd} + \beta_3 momentum_{istd} \\ &+ \beta_4 aggrHFT_{istd} + \beta_5 aggrColo_{istd} + \beta_6 passHFT_{istd} + \beta_7 passHFT_{istd} \\ &+ \beta_8 passColo_{istd} + \beta_9 takebook_{istd} + \beta_{10} VFTSE_{istd} + \beta_{11} spread_{istd} \\ &+ \beta_{12} FE_{sd} + \epsilon_{istd} \end{aligned} \tag{3}$$

Where $PriceImpact_{isdt+m}$ measures the price impact for trade i at for stock s at time t on day d, for $m = 100$ milliseconds after the trade, in basis points. This is calculated as the difference between the trade price and the LSE mid price[32] 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

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

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 trade.[33] $spread_{istd}$ is the quoted spread of the EBBO at the time of the trade in basis points. We also use stock and date fixed effects.[34] 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 10 and 11. Figure 10 illustrates a buyer initiated dark trade at a stale reference price (see footnote for description of Figure). Because the stale midpoint price is much lower than the new midpoint 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 6 about here.]

Figure 11, 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.

[Figure 7 about here.]

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. Table 4 reports the results of

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

[34] We winsorize price impact at 1% and 99% but the results are qualitatively unchanged from non-winsorization.

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 $PriceImpact$, 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 $PriceImpact$, 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 realized spreads.[35] 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.

[Table 4 about here.]

6.1 Dark Liquidity Provision Effects

As we demonstrated in Section 5.5 and 6, HFTs seem to be able to avoid adverse selection in dark pools. To ascertain how they are able to achieve this, we examine their liquidity provision and of other participant classes, for the dark venues that we have orderbook data for. We show that HFTs virtually never provide resting marketable liquidity in the dark pools in our sample. This is quite remarkable given that they provide the majority of liquidity on lit venues.

Our study is, to our knowledge, the first to characterize liquidity provision in modern dark pools. In lit limit order book markets, measures of liquidity are widely accepted, 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.

[35] 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.

Unlike lit limit order markets, in dark pools 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' (liquidity unavailable in the dark). 'No sided' liquidity occurs when there are no dark orders in the order book, but more often dark orders exist, albeit with 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. In such cases standard market depth measures are not particularly informative. We propose a new measure of liquidity defined as the percentage of time there is a bid or ask order that available 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 primary market bid or ask quantities and for those which do not. We are being generous in our definition of liquidity supply in the context of midpoint dark pools since a participant with a marketable midpoint order may also be interpreted as a consumer of liquidity since they are willing to cross half the spread. Nonetheless, they are also a provider of liquidity: by resting passively on the order book, they facilitate executions. Without resting liquidity, participants with aggressive orders merely 'ping' dark pools without executing, like 'ships passing in the night.'[36] Figure 12 presents cumulative frequency histograms of our dark liquidity metrics calculated for a given stock-day and venue with each chart illustrating a distribution of 8550 observations (114 stocks x 25 days x 3 dark pools). This figure shows demonstrates that there are many stocks and/or dates, for which it is not possible for participants to access any dark liquidity. These charts are presented separately by participant class.

[Figure 8 about here.]

[Figure 9 about here.]

[36] 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). See BATS Europe (2017).

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 13. 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 reflects consistent repricing of resting orders to be non-marketable, in response to primary market movements. A potential rationale for this behavior is to take advantage of stale reference prices through passive executions, as described in Section 6 while also minimizing the adverse selection risks.

7 Market Design Solutions - Event Studies

There are a number of market design changes that could solve the problem we have characterized above; reference price latency which results in increased adverse selection for dark pool users and decreased dark liquidity from a reluctance to provide liquidity. We examine a number of market design changes that have been implemented in recent years to determine which interventions reduce adverse selection. We implement the following difference in differences model for each event study:

$$\begin{aligned} AdverseSelection_{sdv+m} &= \alpha + \beta_1 event_d + \beta_2 treated_{sv} + \beta_3 eventTreated_{sdv} \\ &+ \beta_4 Spread_{sd} + \beta_5 VIX_d + \beta_6 consideration_{sdv} + \beta_7 FE_{sd} + \epsilon_{sdv} \end{aligned} \tag{4}$$

Where $AdverseSelection_{sdv}$ measures the natural log of the consideration weighted average adverse selection for stock s date d, for exchange venue[37] v, $m = 100$ milliseconds before and after the trade, in basis points. This is the absolute value of

[37]This takes the value of one of dark venues in our UK event studies, or NASDAQ versus NYSE listed trades reported to FINRA

the difference between the NBBO midpoint m before the trade and m after the trade. We also calculate $AdverseSelection_{sdv}$ as an alternative measure which is the proportion of trades stock s, date d, venue v for which a change in the midpoint occurs over the time period m before and after the trade. This measure is useful as a measure of adverse selection as a significant proportion of dark trades occur with no price movement. $Event_d$ is a dummy variable which takes the value of one on and after the event in question, and zero otherwise. $Treated_{sv}$ is a dummy variable which takes the value of one for stocks traded on venues which are subject to the market design change, and zero otherwise. $EventTreated_{sdv}$ is the interaction term for venues which are subject to the treatment, after the event occurs. $Spread_{sd}$ is the stock's time weighted average quoted spread. $Consideration_{sdv}$ is the natural log of the sum of trading consideration in British Pounds or US dollars. We also use stock and date fixed effects. VIX_d is the natural log of the mean value of the VIX volatility index, or the VFTSE (for UK event studies), measured at 15 seconds intervals throughout the day.

7.1 Turquoise Random Uncross

The UK MTF 'Turquoise' contains a dark order book with an order type called 'random uncross.' This works by matching passive orders at random periods of between 5 and 45 seconds throughout the day, depending on the stocks underlying liquidity (Turquoise, 2016). As the uncrossing period cannot be foreseen, aggressive participants cannot race reference price update messages, reducing or eliminating adverse selection. While Turquoise has had this feature since October 4th, 2010, it was 'relaunched' in September 2013. Turquoise's CEO stated that volumes in uncross were 'up more than three times since the start of September 2013', (Times, 2013). Therefore, we use this event to determine if adverse selection decreases in relation to other comparable dark venues, using the same sample of stocks in a difference in differences model. The untreated sample are dark trades on the BATS and Chi-X dark pools. The sample period is the same 57 FTSE 100 and 57 FTSE 250 stocks from our previous analysis. We use data from the 1st of June 2013 to the 31st of August

2013 in our pre-event window and data from the 1st of October 2013 to the 30th of December 2013 in our post-event window. Table 5 sets out our results which show a statistically significant reduction in our adverse selection proxies after the event for our treatment sample. This is roughly a 12% reduction in adverse selection in basis points, and a 4.1% reduction in the proportion of dark trades which experience movements in prices around the trade. Our results indicate that this market design intervention appears to be effective at reducing adverse selection driven by reference price latency.

7.2 NASDAQ SIP Upgrade

On October 24th, 2016 NASDAQ, which manages one of the two largest Security Information Processors (SIPs) in the US, increased the speed at which it processes trades and quotes from around 500 microseconds to around 40. We use this event to examine the impact of speed upgrades to reference price feeds on adverse selection.

A number of dark pools utilize the SIP to price dark trades, including Barclays (SEC, 2016b), and Citadel (SEC, 2017). The NASDAQ SIP should decrease the extent of reference price latency for dark trades using the NASDAQ SIP.

We use the SIP upgrade event as a case study of one possible solution to reference price latency; speeding up reference price feeds. Ex ante, we do not expect this to diminish reference price latency arbitrage, and as such, do not expect a decrease in our adverse selection measure. This is because, while reference price feed latency is reduced, we do not expect it to be enough to allow dark pools to 'win the speed race' with HFT and other fast participants using direct exchange feeds. We expect many participants to be be faster than both the old and the new SIP.

We implement a difference in differences regression model on this event, with the untreated sample being dark pool trades reported to the NYSE SIP. As SIP usage is almost entirely a function of the stock's primary listing (NASDAQ v NYSE), we expect the NYSE to be unaffected by this change, and thus represents our 'untreated' sample group.

We utilize SIP quotes and trades from TRTH (Thomson Reuters Tick History).

Our population of dark trades comes from FINRA reported trades, which, as noted in Menkveld et al. (2014) comprises the majority of dark trading. We only include normal trades, excluding flags identifiers such as sweep, average, basket, burst, odd lots, next day and bunched trades. We exclude the 15 minutes after open and before close and calculate adverse selection as the movement in the NBBO mid around the trade using the NBBO TAQ file, ignoring flags with any slow quotes and any crossed markets. We note that despite allegations of slow SIP reporting of dark trades, we find that the majority (more than 99%) of FINRA trades occur within the NBBO at the time of reporting, implying that these are insignificant. Regardless, we allow for a 100 millisecond window before and after the trade, and do not expect, or observe a change in this ratio in response to a 450microsecond reduction in SIP processing times. Our sample period is the 52 trading days before and after the event date, starting from the 2nd of September 2016, to 15th of December 2016. Our sample of stocks is the S&P500, but we remove a random selection of NYSE stocks to balance the panel, as there is a higher proportion of NYSE stocks in the index. This leaves us with 246 stocks.

Unlike our previous DID models, rather than examining the same sample of stocks across different venues, we are examining a different sample of stocks across different venues. This means that if NASDAQ stocks experience a period of volatility in our post-treatment period, (after October 24th) this may confound our results. As the 2016 US presidential election occurs after the SIP upgrade, and the election impacted technology stocks more significantly than industrial stocks, this is a potential confounding event. To address this, we express our dependent variables as a proportion of their averages across all non-dark venues. For example, for each stock we express the dependent variable as: $AdverseSelection_{DarkFinra}/mean(AdverseSelection_{lit}) * 100$. This controls for idiosyncratic factors specific to a stock being a NYSE or NASDAQ listing and allows us to look at changes in adverse selection that are specific to dark trading, rather than information.

Table 5 presents the results of our difference in differences regression for our two measures of adverse selection. Our results show that the *event* variable has no statistically significant association, perhaps implying that the denominator is effectively controlling for information. As *treated* is associated with higher adverse selection, this implies that NASDAQ listed stocks are more volatile than NYSE ones, which is to be expected. As predicted, the event does not decrease adverse selection, $treated_event$ has no statistical significance for price impact, but for the percentage of trades which experience price movements, the effect is positive. However, the coefficient is only 1.6%, this presents an increase in ratio of % price movements in the dark versus the lit of 1.6%, this is not a significant change.

[Table 5 about here.]

7.3 BATS Intraday Periodic Batch Auctions

BATS Europe introduced a new order type on the 19th of October 2015, which is a separate orderbook that performs intra-day auctions with randomized uncross periods throughout the day. It has price, time and size priority, but is only partially pre-trade transparent. There is an indicative uncrossing price and executable volume published in the call period prior to the uncross, but this is only published if the full minimum acceptable quantity of orders is fully met, (BATS Europe, 2017).

Trades in this auction are flagged in our data from TRTH, and we examine whether the random uncross feature is effective at reducing adverse selection in comparison to other dark orders. Given the unpredictability of the uncross prevents race conditions against lit market updates, we expect lower adverse selection.

We use a seven month period in 2016 (June to December inclusive) where volumes in the BATS Batch auctions are higher than in 2015. Because we have trade identifiers, we compare adverse selection for batch trades versus dark non-batch trades. We use the dummy variable $batch_{sdv}$ which is coded one for the consideration weighted mean of adverse selection for the stock, date, venue observation. Our sample includes the same 114 stocks for dark trades in Chi-X and batch and dark trades in BATS.[38]. Our results are reported in Table 6. There are statistically significant negative relationships between $batch_{sdv}$ and our adverse selection measures. Our model estimates

[38]Our results are qualitatively unchanged when our sample includes just BATS dark and batch trades alone

a 37% reduction in adverse selection in basis points, and a 16.6% reduction in the proportion of midpoint changes 100 milliseconds before and after the trade for batch auction trades in comparison to dark trades on UK dark venues.

[Table 6 about here.]

8 Conclusion

We examine the prevalence of stale reference prices, their effect on adverse selection and participant strategies with respect to stale dark trades for a representative sample of UK exchange operated midpoint dark pools. We find trades are often at stale reference prices, and this results in asymmetric outcomes across participant types, wherein faster (HFT) participants impose adverse selection costs on slower participants. We show that stale reference prices play a significant role in increasing adverse selection costs for passive dark orders, increasing them by 0.88 basis points. We present descriptive statistics of liquidity provision in dark pools in the UK by participants, finding that HFTs virtually never provide resting (marketable) liquidity. We provide reasons why, such as the lack of spread at the midpoint, and adverse selection from reference price latency. We also examine market design interventions by three separate dark pools to resolve the problem of adverse selection from reference price latency. We find that removing the ability for aggressive participants to race the market data feed by making dark executions unpredictable (through random uncross mechanisms) is effective in reducing adverse selection in comparison to other dark orders. Efforts to increase reference price feeds are not effective at removing adverse selection, as they do not eliminate the market data race.

9 Bibliography

References

Aisen, D. 2015. The Genesis of an Order Type. Blog. URL https://blog.quantopian.com/the-genesis-of-an-order-type-by-daniel-aisen/.

Alexander, J., L. Giordano, and D. Brooks. 2015. Dark Pool Execution Quality: A Quantitative View. URL http://tabbforum.com/opinions/dark-pool-execution-quality-a-quantitative-view.

Anderson, L., B. Devani, and Y. Zhang. 2016. The Hidden Cost: Reference Price Latencies. Trading Review And Analysis, Investment Industry Regulatory Organization of Canada. URL http://www.iiroc.ca/Documents/2016/f1e06451-0f2b-4af8-a973-81fa881a6958_en.pdf.

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

ASIC. 2015. Review of high-frequency trading and dark liquidity. Tech. Rep. 452, Australian Securities and Investments Commission.

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

Baron, M. D., J. Brogaard, B. Hagströmer, and A. A. Kirilenko. 2016. Risk and Return in High-Frequency Trading. SSRN Scholarly Paper ID 2433118, Social Science Research Network, Rochester, NY. URL http://papers.ssrn.com/abstract=2433118.

Bartlett, R. P., and J. McCrary. 2016. How Rigged Are Stock Markets? Evidence from Microsecond Timestamps. SSRN Scholarly Paper ID 2812123, Social Science

Research Network, Rochester, NY. URL http://papers.ssrn.com/abstract=2812123.

BATS Europe. 2017. BATS Europe Guidance Note: Periodic Auctions Book. URL http://cdn.batstrading.com/resources/participant_resources/BCE-GuidanceNote_Periodic_Auctions_Final.pdf.

BATS Europe. 2017. Trading Price List. Tech. rep., BATS Europe. URL http://cdn.batstrading.com/resources/participant_resources/BATSEuro_Pricing.pdf.

Boni, L., D. C. Brown, and J. C. Leach. 2013. Dark pool exclusivity matters. *Available at SSRN 2055808* URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2055808.

Brogaard, J., B. Hagströmer, L. Nordén, and R. Riordan. 2015. Trading Fast and Slow: Colocation and Liquidity. *Review of Financial Studies* p. hhv045. URL http://rfs.oxfordjournals.org/content/early/2015/09/09/rfs.hhv045.

Budish, E., P. Cramton, and J. Shim. 2015. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics* 130:1547–1621. URL http://qje.oxfordjournals.org/lookup/doi/10.1093/qje/qjv027.

Chen, H., S. Foley, M. A. Goldstein, and T. Ruf. 2016. The Value of a Millisecond: Harnessing Information in Fast, Fragmented Markets. SSRN Scholarly Paper ID 2860359, Social Science Research Network, Rochester, NY. URL https://papers.ssrn.com/abstract=2860359.

CME. 2014. CME Group Customer Forum Q2 2014. URL http://www.cmegroup.com/globex/resources/files/spring-2014-cust-forum-seoul.pdf.

Corvil. 2014. Electronic Trading System Performance. White Paper, Corvil. URL http://corvil.com/content/resources/

04-white-papers/03-electronic-trading-system-performance/wp-electronic-trading-system-performance.pdf.

Deutsche Bank. 2016. SuperX EMEA: Order type, matching logic and operational summary. Quick Guide. URL https://www.autobahn.db.com/microSite/docs/SuperXEMEA_QuickGuide_Jan2016_v1.pdf.

Ding, S., J. Hanna, and T. Hendershott. 2014. How Slow Is the NBBO? A Comparison with Direct Exchange Feeds. *Financial Review* 49:313–332. URL http://onlinelibrary.wiley.com.ezproxy1.library.usyd.edu.au/doi/10.1111/fire.12037/abstract.

Eurex. 2016. Insights into trading system dynamics: Eurex Exchange's T7. URL http://www.eurexchange.com/blob/238346/a98f43ac813a2555871bc6b85a155ad7/data/presentation_insights-into-trading-system-dynamics_en.pdf.

Europe, K. 2015. KCG Europe Limited 2015 Financial Statements. Financial Statements, KCG Europe. URL https://beta.companieshouse.gov.uk/company/05976015/filing-history.

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

Fidessa. 2017. Fidessa Fragmentation Index. URL http://fragmentation.fidessa.com.

FINRA. 2014. FINRA Fines Goldman Sachs Execution & Clearing, L.P. $800,000 for Failing to Prevent Trade-Throughs in its Alternative Trading System. URL http://www.finra.org/newsroom/2014/finra-fines-goldman-sachs-execution-clearing-lp-800000-failing-prevent-trade-throu

Foley, S., and T. J. Putniņš. 2016. Should We Be Afraid of the Dark? Dark Trading and Market Quality. SSRN Scholarly Paper ID 2279719, Social Science Research Network, Rochester, NY. URL http://papers.ssrn.com/abstract=2279719.

Foucault, T., R. Kozhan, and W. W. Tham. 2016. Toxic Arbitrage. SSRN Scholarly Paper ID 2409054, Social Science Research Network, Rochester, NY. URL https://papers.ssrn.com/abstract=2409054.

Glosten, L. R., and P. R. Milgrom. 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14:71–100. URL http://www.sciencedirect.com/science/article/pii/0304405X85900443.

GSEC. 2016. SIGMA X FAQ. Tech. Rep. Version 8, Goldman Sachs Execution & Clearing, L.P. ("GSEC"). URL http://www.goldmansachs.com/what-we-do/securities/gset/equities/liquidity-access/sigma-x-us-faqs.pdf.

Guiso, L., P. Sapienza, and L. Zingales. 2008. Trusting In Stock Markets. *The Journal of Finance* 63:2557–2600. URL http://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.2008.01408.x/full.

Hagströmer, B., and L. Nordén. 2013. The diversity of high-frequency traders. *Journal of Financial Markets* 16:741–770. URL http://www.sciencedirect.com/science/article/pii/S1386418113000256.

Hasbrouck, J. 2015. High frequency quoting: Short-term volatility in bids and offers. *Available at SSRN 2237499* URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2237499.

Hope, B. 2014. Goldman Agrees to $800,000 Fine Over Dark Pool. *Wall Street Journal* URL http://www.wsj.com/articles/finra-fines-goldman-dark-pool-trading-unit-800-000-1404224624.

IEX. 2015. Investors Exchange Form 1 Submission. Filing to Register as Exchange 3235-0017, Investors' Exchange LLC. URL https://www.sec.gov/rules/other/2015/investors-exchange-form-1-exhibits-a-e.pdf.

International, J. T. 2015. Jump Trading International 2015 Financial Statements. Financial Statements, Jump Trading International. URL https://beta.companieshouse.gov.uk/company/05976015/filing-history.

ITG. 2015. Global Cost Review Q3 2014. Report 120710-29252, Investment Technology Group, Investment Technology Group. URL http://www.itg.com/assets/Q3_2014_2015.pdf.

Kirilenko, A. A., A. S. Kyle, M. Samadi, and T. Tuzun. 2015. The Flash Crash: The Impact of High Frequency Trading on an Electronic Market. SSRN Scholarly Paper ID 1686004, Social Science Research Network, Rochester, NY. URL http://papers.ssrn.com/abstract=1686004.

Kirilenko, A. A., and G. Lamacie. 2015. Latency and Asset Prices. SSRN Scholarly Paper ID 2546567, Social Science Research Network, Rochester, NY. URL http://papers.ssrn.com/abstract=2546567.

LSEG. 2013. Connectivity: World Leading Access to Markets. Brochure PM0248, London Stock Exchange. URL http://www.lseg.com/sites/default/files/content/documents/LSEG_Connectivity_Full_Brochure.pdf.

LSEG. 2013. Technical User Group. URL http://www.londonstockexchange.com/products-and-services/technical-library/technical-user-group/technicalusergrouparchive/may2013londontug.pdf.

LSEG. 2015. Technical User Group. URL http://www.londonstockexchange.com/products-and-services/technical-library/technical-user-group/tuglonnov15.pdf.

LSEG. 2016. LSEG launches FPGA-powered market data dissemination platform. URL http://www.lseg.com/resources/media-centre/press-releases/lseg-launches-fpga-powered-market-data-dissemination-platform.

LSEG. 2017. Dark Midpoint Order book. URL http://www.lseg.com/areas-expertise/our-markets/turquoise/equities/tq-dark-midpoint-order-book.

Menkveld, A. J. 2016. High-Frequency Trading as Viewed Through an Electronic Microscope. SSRN Scholarly Paper ID 2875612, Social Science Research Network, Rochester, NY. URL https://papers.ssrn.com/abstract=2875612.

Menkveld, A. J., and B. Z. Yueshen. 2016. The Flash Crash: A Cautionary Tale about Highly Fragmented Markets. SSRN Scholarly Paper ID 2243520, Social Science Research Network, Rochester, NY. URL https://papers.ssrn.com/abstract=2243520.

Menkveld, A. J., B. Z. Yueshen, and H. Zhu. 2014. Shades of Darkness: A Pecking Order of Trading Venues. SSRN Scholarly Paper ID 2510707, Social Science Research Network, Rochester, NY. URL http://papers.ssrn.com/abstract=2510707.

Nimalendran, M., and S. Ray. 2014. Informational linkages between dark and lit trading venues. *Journal of Financial Markets* 17:230–261. URL http://www.sciencedirect.com.ezproxy1.library.usyd.edu.au/science/article/pii/S1386418113000128.

Rosenblatt. 2014. Monthly Dark Liquidity Tracker European Edition. Tech. Rep. September 2014, Rosenblatt Securities Limited.

Saraiya, N., and H. Mittal. 2009. Understanding and Avoiding Adverse Selection in Dark Pools. URL http://www.atmonitor.co.uk/content/researchimages/AdverseSelectionDarkPools_113009F.pdf.

SEC. 2016a. ADMINISTRATIVE PROCEEDING File No. 3-17077. URL https://www.sec.gov/litigation/admin/2016/33-10010.pdf.

SEC. 2016b. ADMINISTRATIVE PROCEEDING File No. 3-17079. URL https://www.sec.gov/litigation/admin/2016/33-10014.pdf.

SEC. 2017. Citadel Securities LLC - Order Instituting Administrative Cease and Desist Proceedings. Administrative Proceeding File No. 3-17772. URL https://www.sec.gov/litigation/admin/2017/33-10280.pdf.

Standard and Poors Capital IQ. 2015. Real Time Solutions Global Network Map. URL http://web.archive.org/web/20150702051320/http://www.spcapitaliq-realtime.com/wp-content/uploads/2015/06/SP-CIQ-Real-Time-Solutions-Global-Network-Diagram-Q2-2015.pdf.

Swanson, E. 2015. Investors' Exchange, LLC ("IEX"); Notice of Filing of Application, as Amended, for Registration as a National Securities Exchange under Section 6 of the Securities Exchange Act of 1934; Release No. 34-75925. SEC Filing 10-222., BATS. URL https://www.sec.gov/comments/10-222/10222-9.pdf.

Times, A. S. 2013. Trading Platforms: Azure waters ahead. *Asset Servicing Times* URL http://assetservicingtimes.com/astimes/ASTimes_issue_78.pdf.

Turquoise. 2016. Turquoise Plato Uncross Factsheet. URL http://www.lseg.com/sites/default/files/content/documents/TURQUOISE_UNCROSS_FACTSHEET_2016_AW.pdf.

Wahal, S., and J. S. Conrad. 2016. The Term Structure of Liquidity Provision. SSRN Scholarly Paper ID 2837111, Social Science Research Network, Rochester, NY. URL https://papers.ssrn.com/abstract=2837111.

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

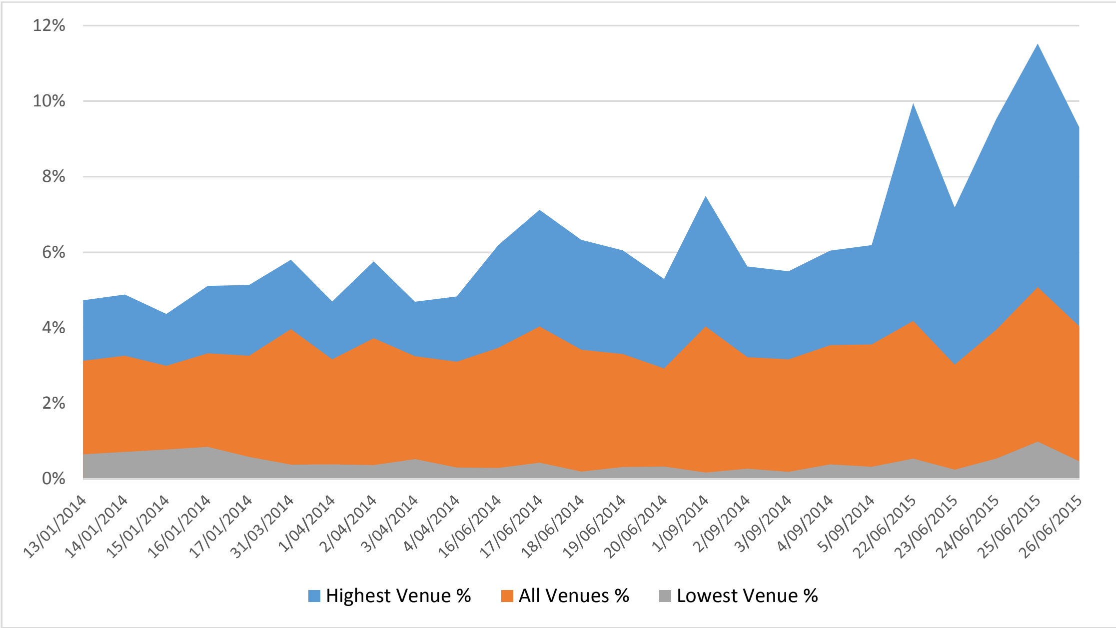

Figure 2: Proportion of Stale Dark Trades by Stock (%)

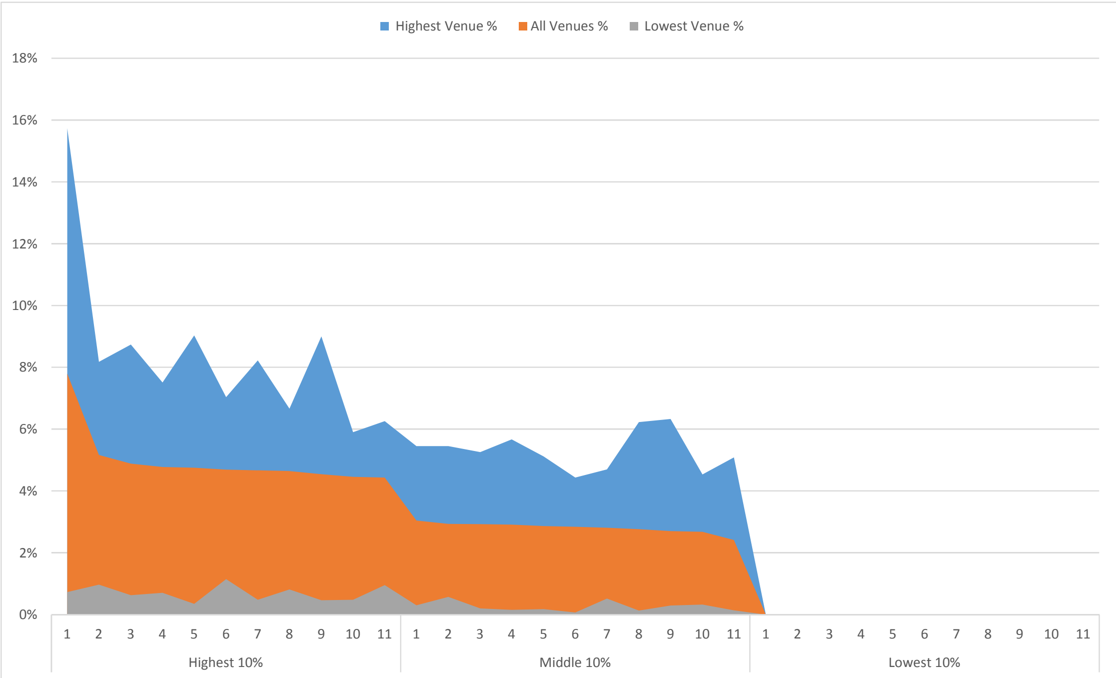

Figure 3: Dark Trades at Stale Reference Prices (%)

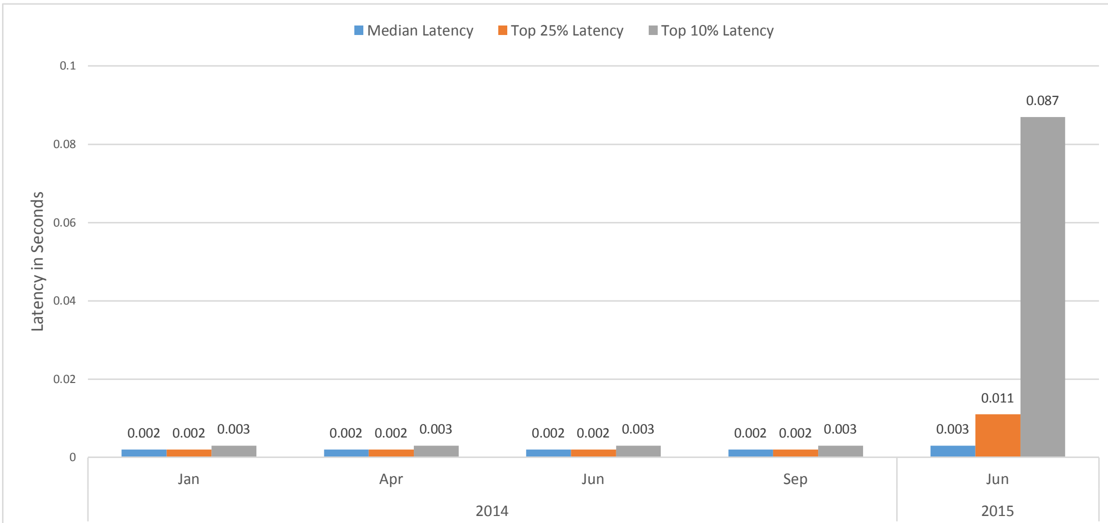

Figure 4: Proportion of Stale Dark Trades Relative to BBO (%)

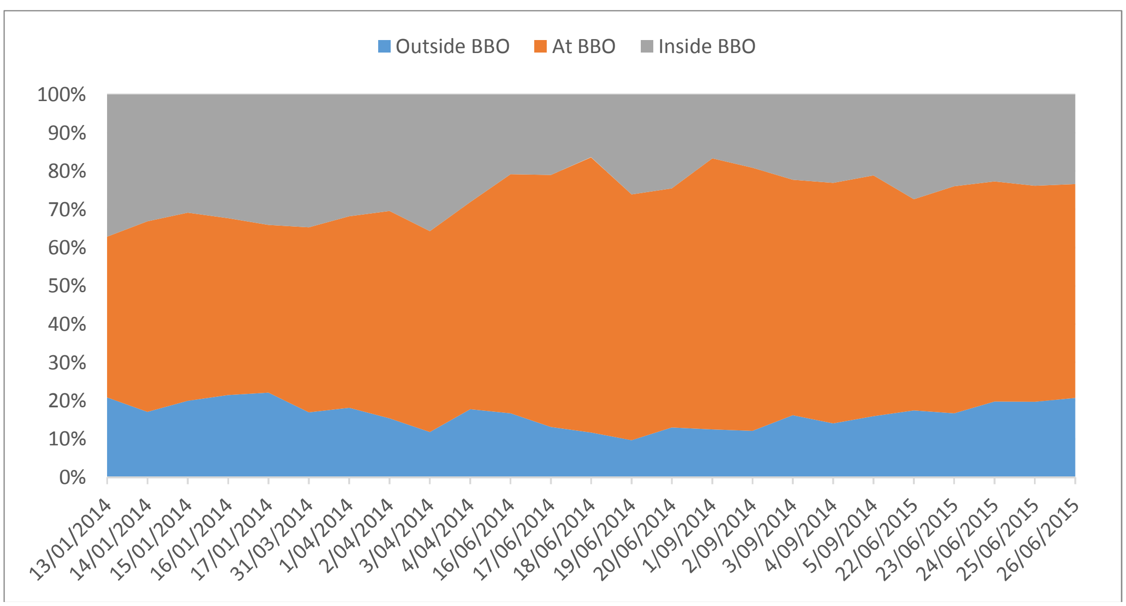

Figure 5: Proportion of Participants' Trades on Benefit Side of Stale Dark Trade (%)

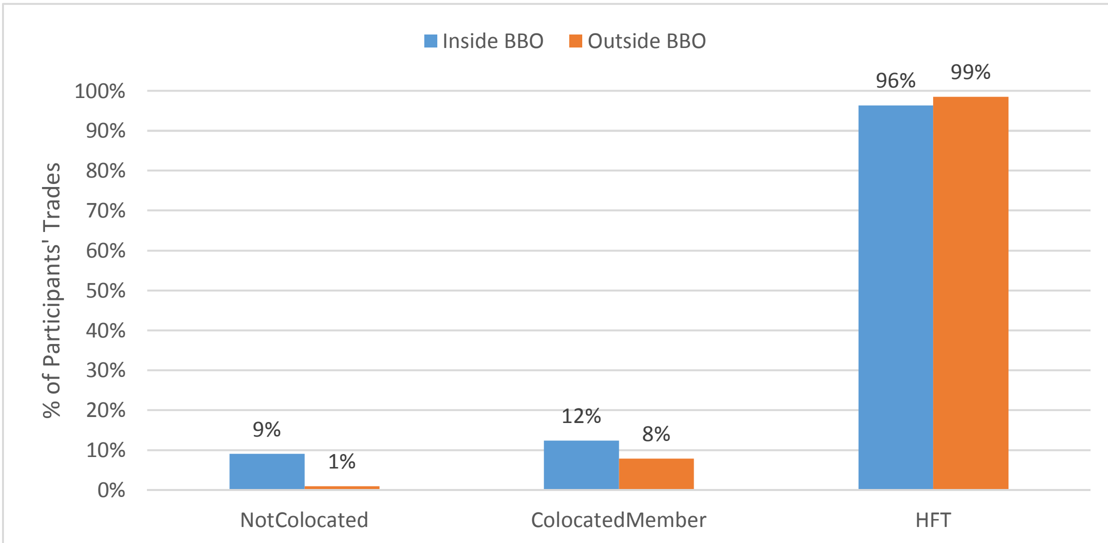

Figure 6: Proportion of Participants' Trades on Any Side (%)

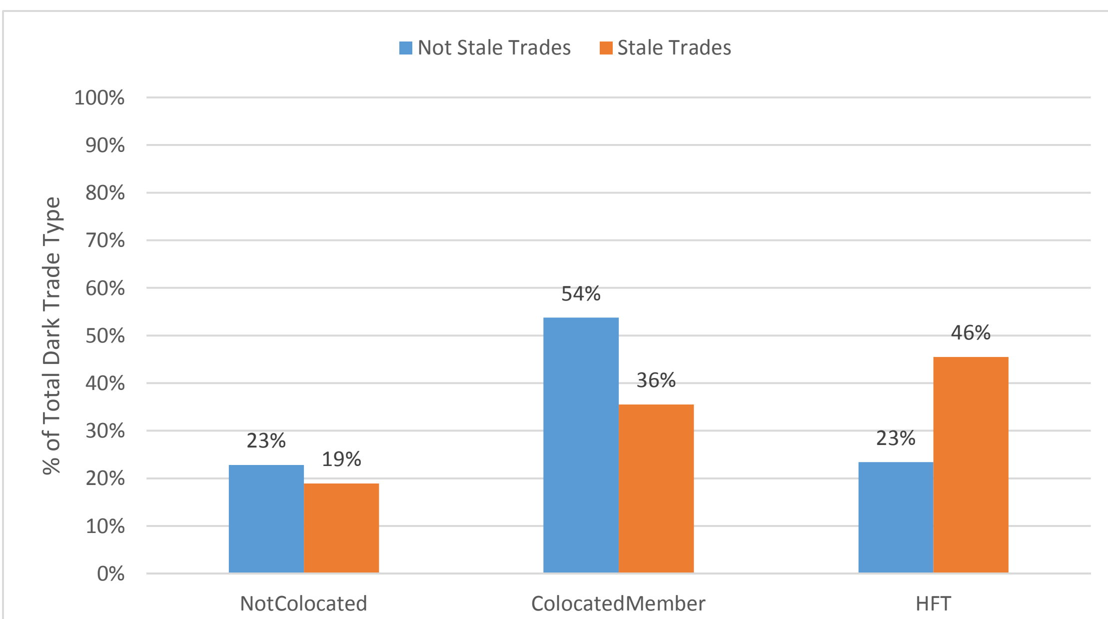

Figure 7: Proportion of Participants' Trades on Any Side - by BBO (%)

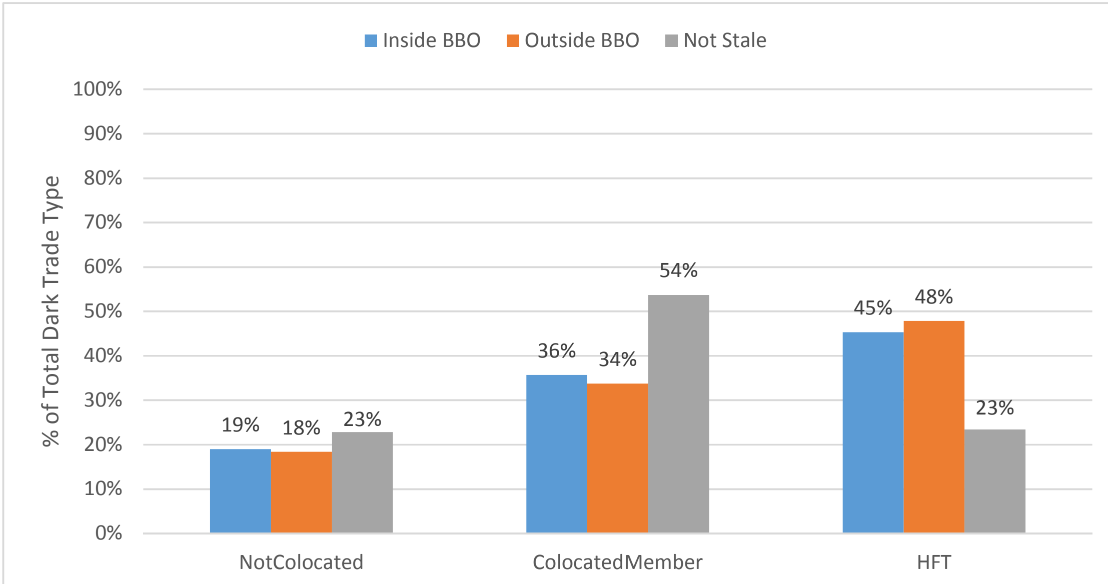

Figure 8: Participation on Aggressive Side of Dark Trades (% of all Dark Trades)

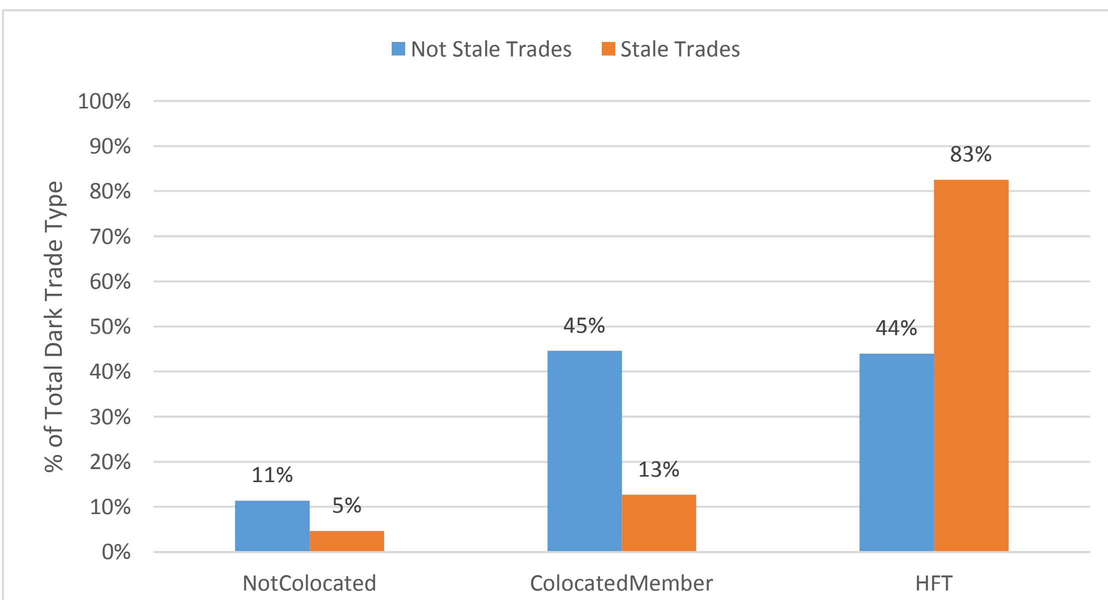

Figure 9: Participation on Passive Side of Dark Trades (% of all Dark Trades)

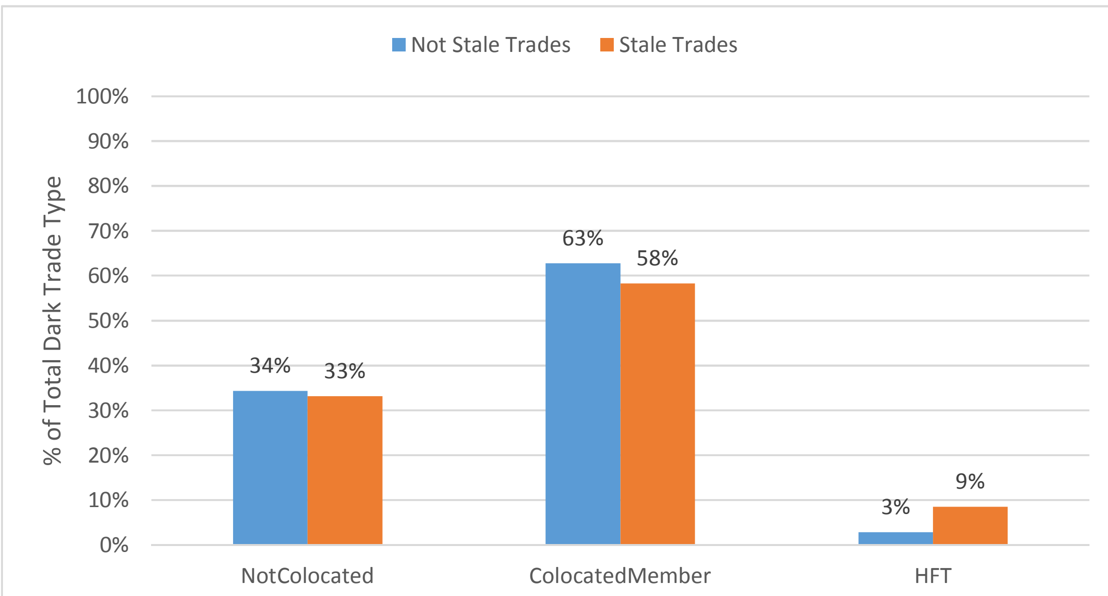

Figure 10: Stale Dark Trade Example (Aggressive Benefit)
This figure illustrates a lit market order book over time. The green shaded section represents liquidity at the Best Ask, the red shaded section represents liquidity at the Best Bid, and the unshaded section in the middle, represents 'the spread.' 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 resting bid prices at the time of the trade. This is because it is referencing the old midpoint towards the left of the figure.

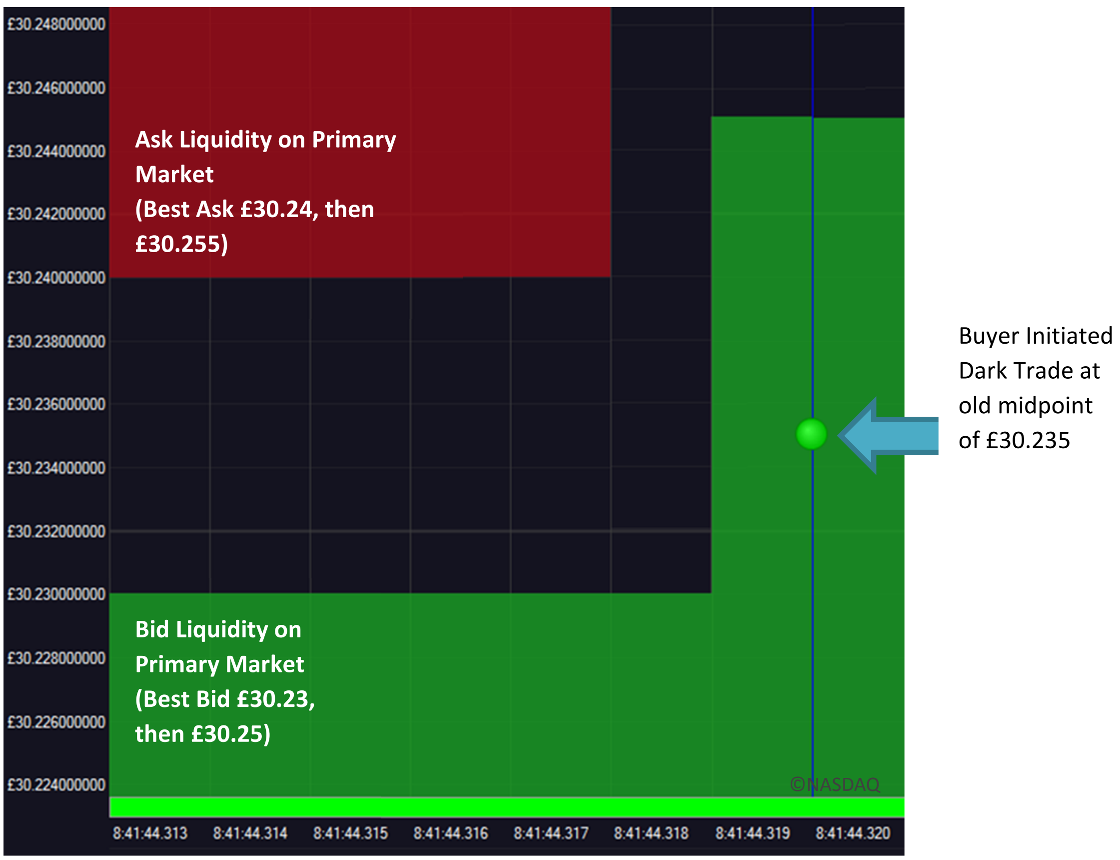

Figure 11: Stale Dark Trade Example (Passive Benefit)
This figure illustrates a lit market order book over time. The green shaded section represents liquidity at the Best Ask, the red shaded section represents liquidity at the Best Bid, and the unshaded section in the middle, represents 'the spread.' 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 resting ask prices at the time of the trade. This is because it is referencing the old midpoint towards the left of the figure.

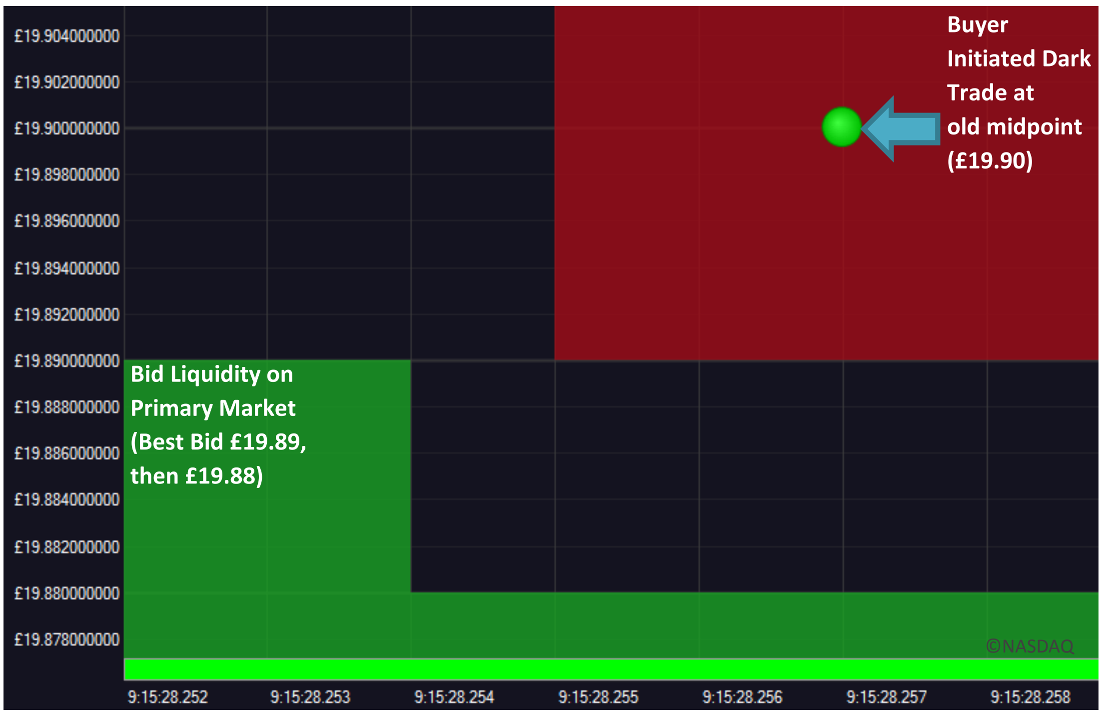

Figure 12: Histograms of Marketable Orders in Midpoint Dark Pools by Participant

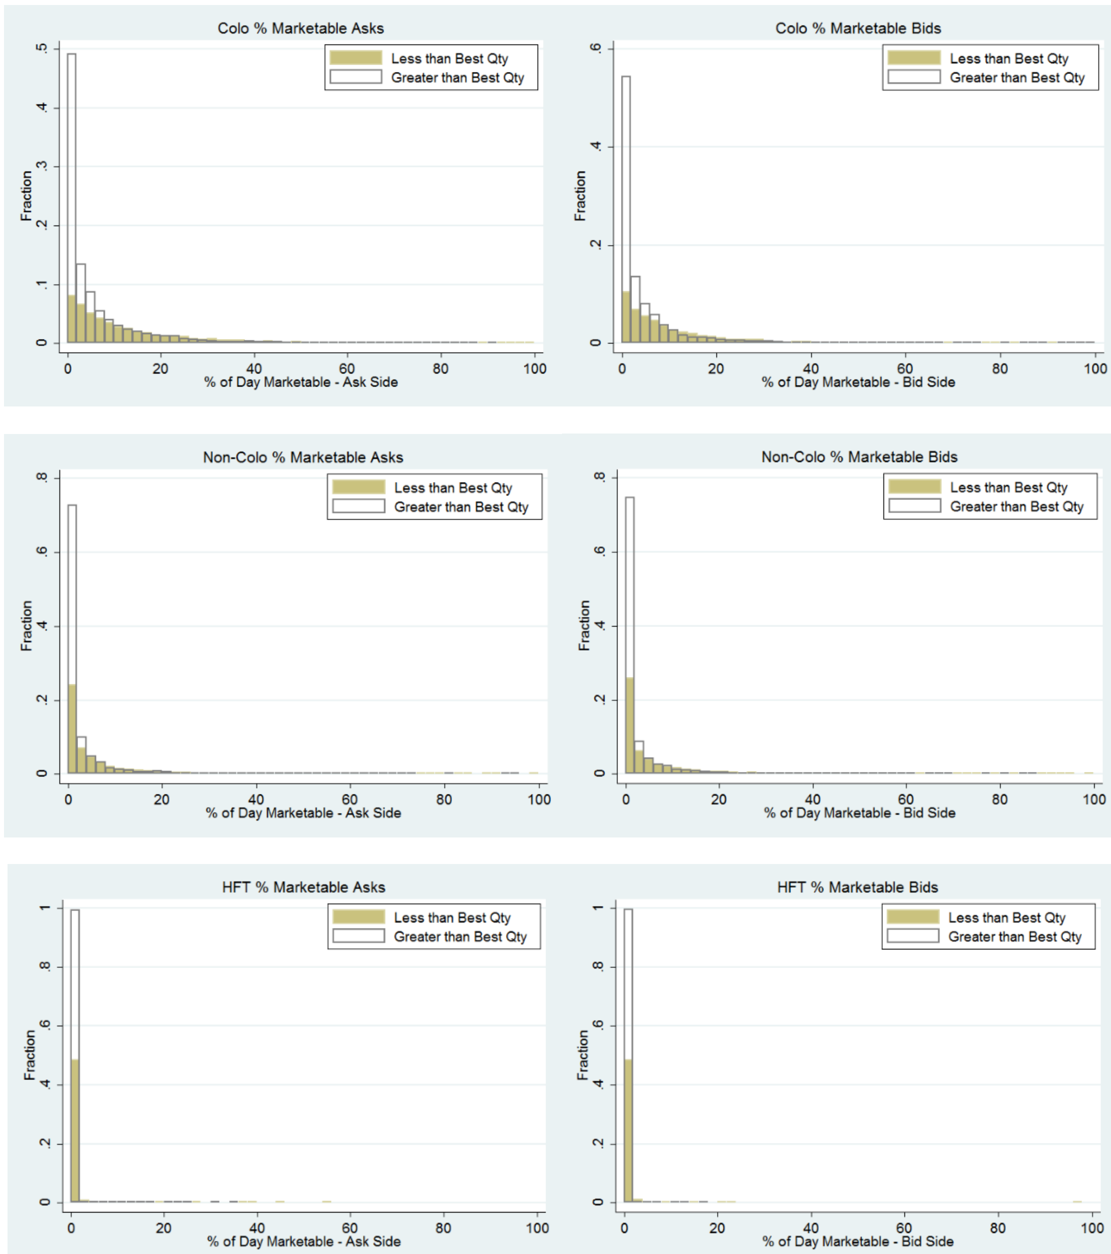

Figure 13: Histograms of Non-Marketable Orders in Midpoint Dark Pools by Participant

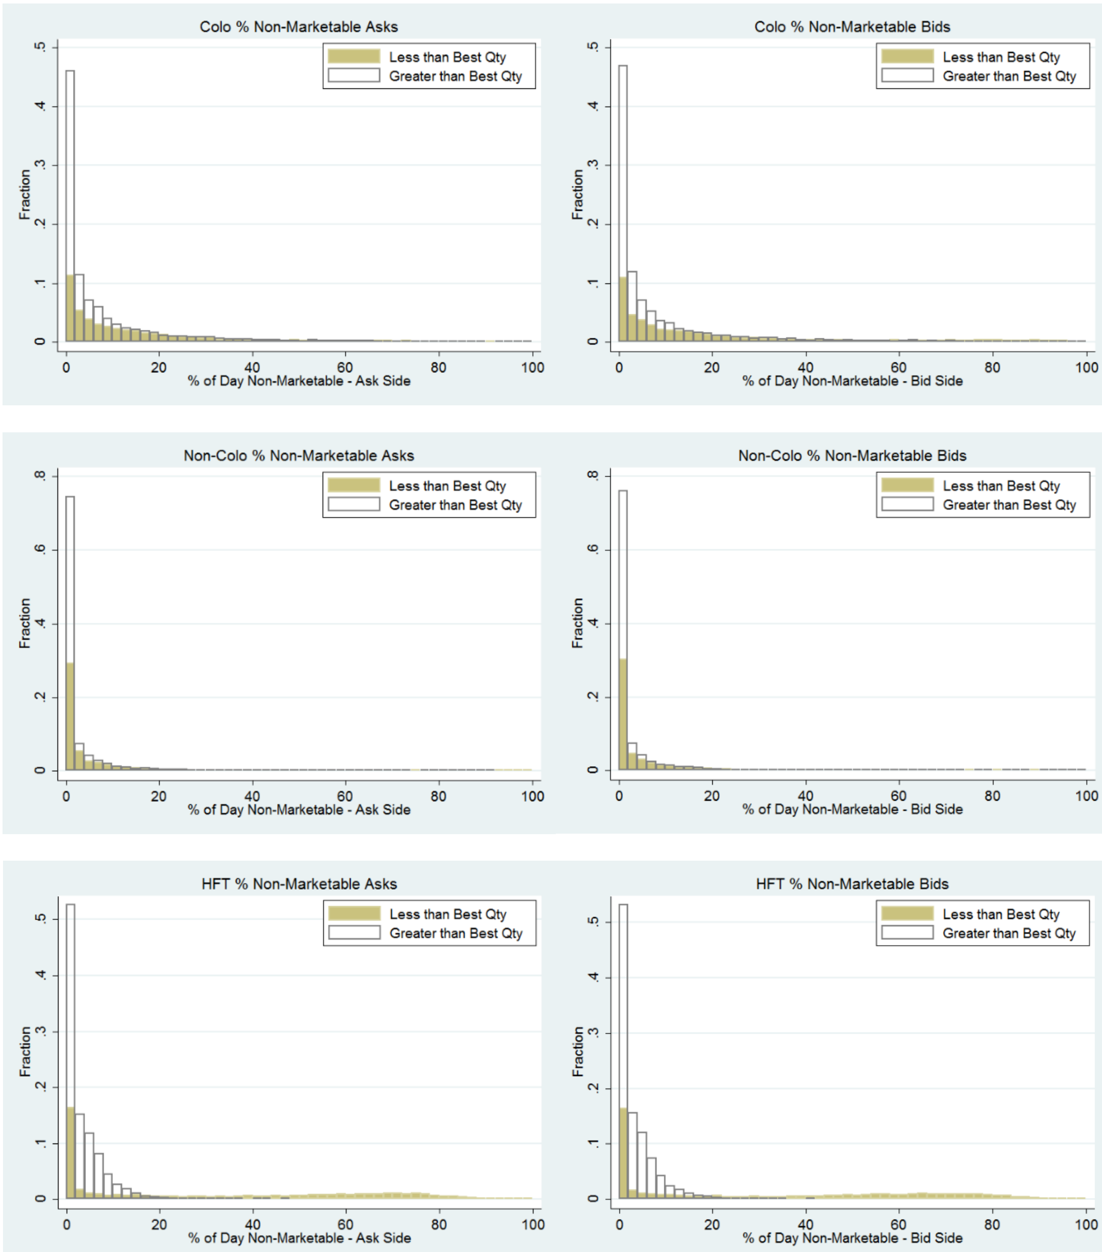

Table 1: Regression of Stale Trades over Time
This table reports coefficient estimates for the percentage of stale trades across all venues for each stock date in our sample using the following specifications:

(1) $Percstale_{s,d} = Trend + FE_{s,v} + \epsilon_{s,d}$
(2) $Percstale_{s,d} = Trend + Spread_{s,d} + PriceVol_{s,d} + VFTSE_d + Messages_d + FE_{s,v} + \epsilon_{s,d}$
(3) $Percstale_{s,d} = Trend + Spread_{s,d} + PriceVol_{s,d} + Messages_d + FE_{s,v} + \epsilon_{s,d}$
(4) $Percstale_{s,d} = Trend + Spread_{s,d} + PriceVol_{s,d} + VFTSE_d + FE_{s,v} + \epsilon_{s,d}$

Where $Trend$ is our time trend variable which increments one for each calendar date in our sample. $Spread_{s,d}$ is the time weighted quoted spread for stock s and day d in pence. $PriceVol_{s,d}$ is the % change in a stock's opening and closing price. $VFTSE_d$ is the natural log of the average FTSE Volatility Index across 15 second time intervals on day d. $Messages_d$ is the average number of messages across two millisecond time intervals on date d across the top 400 stocks on all markets and dark pools in our full orderbook sample. $FE_{s,v}$ represents fixed effects for stock and dark venues v in our sample. Our sample includes five distinct weeks in 2014; the 13th of Jan, 31st of March, 16th of June and 1st of September and the 22nd of June in 2015. Standard errors are clustered at the stock level.

VARIABLES	(1) PercStale	(2) PercStale	(3) PercStale	(4) PercStale
Trend	0.000**	-0.000	-0.000	0.000
	(2.613)	(-0.211)	(-0.104)	(0.632)
Spread		-0.581***	-0.605***	-0.598***
		(-4.856)	(-5.237)	(-4.778)
PriceVol		0.000	0.000	0.000
		(0.270)	(0.375)	(0.110)
VFTSE		0.030***		0.045***
		(2.772)		(4.979)
Messages		0.002**	0.004***	
		(2.086)	(4.095)	
Constant	0.033***	-0.049**	0.017***	-0.074***
	(19.849)	(-2.067)	(2.950)	(-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

Table 2: Probit Regression of Stale Trades Against Market Messages
This table reports coefficient estimates for a probit regression on individual dark trades in our sample with the dependent variable $StaleTrade$ with value one for a dark trade which is stale, and zero for a trade which is not stale. $TotMktMessages$ is the change in total number of messages for a 2 millisecond period, relative to the previous, around a dark trade across the top 400 stocks across all exchanges in our full orderbook sample, excluding messages in the stock for which the trade occurs.

VARIABLES	StaleTrade
TotMktMessages - 14	0.005***
	(2.640)
TotMktMessages - 13	0.004*
	(1.911)
TotMktMessages - 12	0.006***
	(2.610)
TotMktMessages - 11	0.009***
	(4.175)
TotMktMessages - 10	0.006***
	(3.947)
TotMktMessages - 9	0.010***
	(4.636)
TotMktMessages - 8	0.006***
	(3.436)
TotMktMessages - 7	0.008***
	(4.295)
TotMktMessages - 6	0.006***
	(3.202)
TotMktMessages - 5	0.008***
	(3.599)
TotMktMessages - 4	0.005**
	(2.569)
TotMktMessages - 3	0.004**
	(2.568)
TotMktMessages - 2	0.005***
	(4.611)
TotMktMessages - 1	0.003***
	(6.382)
TotMktMessages - 0	0.002***
	(5.934)
TotMktMessages + 1	0.000
	(0.439)
TotMktMessages + 2	-0.000
	(-0.039)
TotMktMessages + 3	-0.001
	(-0.824)
TotMktMessages + 4	0.001
	(0.634)
TotMktMessages + 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

Table 3: Costs of Stale Reference Prices
The following table presents estimates of the cost of dark pool latency arbitrage to users of dark pools in the UK. This is presented separately for two subsets of stale trades. Stale reference prices that are inside the Primary BBO, and those that are outside, which sum to the total column. Outside BBO costs represents real opportunity costs, whilst inside BBO trades represent assumed opportunity costs. The first row expresses the average cost of a dark pool trade at a stale price in basis points (bps) = basis points. This is calculated as, for trade i, $Cost_i/Consideration_i * 10000$. We assume 253 trading days in the year when scaling the 25 days in our sample. We assume constant proportionality of rates across stocks when scaling our 114 to the FTSE 350. We also assume constant proportionality when scaling our sample of dark venues to the population. We estimate our sample as 32.41% of total dark trading using Rosenblatt's Dark Liquidity report for 2014.

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

Table 4: Regression of Price Impact and Stale Trades
This table reports coefficient estimates of trade by trade regressions of the price impact of dark trades in our sample using the following specification:

$$PriceImpact_{istd+m} = \alpha + \beta_1 Stale_{istd} + \beta_2 consideration_{istd} + \beta_3 momentum_{istd} + \beta_4 aggrHFT_{istd} + \beta_5 aggrColo_{istd} + \beta_6 passHFT_{istd} + \beta_7 passHFT_{istd} + \beta_8 passColo_{istd} + \beta_9 takebook_{istd} + \beta_{10} VFTSE_{istd} + \beta_{11} spread_{istd} + \beta_{12} FE_{sd} + \epsilon_{istd}$$

This model is estimated for three samples, the full sample of dark trades (non-stale and stale), only stale trades which benefit the aggressor counterparty and non-stale and only stale dark trades that benefit the passive counterparty and non-stale trades. Where $PriceImpact_{istd+m}$ is the price impact in basis points for trade i in stock s at time t on date d over the 100 millisecond period m. $Stale_{istd}$ is a dummy variable which is one if the dark trade is deemed to be stale. $consideration_{istd}$ 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 which take the value of one 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 one if the passive side of the trade is an HFT or a co-located participant respectively. $takebook_{istd}$ is a dummy variable with a value of one if the aggressor to the trade also consumes all LSE liquidity in the same trade direction. $VFTSE_i$ is the natural log of the value of the FTSE 100 volatility index 15 seconds prior to the trade. $spread_{istd}$ is the NBBO quoted spread in basis points at the time of the trade. FE_{sd} represents fixed effects for stock and date. Our sample includes five distinct weeks in 2014; the 13th of Jan, 31st of March, 16th of June and 1st of September and the 22nd of June in 2015. Standard errors are clustered at the stock level.

	Price Impact (100 Millisecond)		
VARIABLES	Full Sample	Aggressive Benefit	Passive Benefit
Stale	0.649***	0.883***	-1.095***
	(24.806)	(30.042)	(-15.603)
Consideration	-0.039***	-0.039***	-0.039***
	(-9.150)	(-9.068)	(-8.920)
Spread	0.016***	0.016***	0.015***
	(7.754)	(7.766)	(7.347)
VFTSE	-0.783***	-0.775***	-0.755**
	(-2.668)	(-2.645)	(-2.570)
Momentum	0.479	0.509	0.526
	(1.278)	(1.349)	(1.413)
AggressiveHFT	1.511***	1.492***	1.499***
	(53.381)	(53.101)	(53.563)
AggressiveColo	0.466***	0.466***	0.462***
	(24.989)	(24.993)	(25.362)
PassiveHFT	-1.444***	-1.264***	-1.266***
	(-31.638)	(-27.414)	(-28.727)
PassiveColo	0.058***	0.057***	0.063***
	(4.132)	(4.059)	(4.460)
Takebook	0.503***	0.515***	0.484***
	(5.631)	(5.801)	(5.361)
Constant	2.425***	2.402***	2.357***
	(3.288)	(3.257)	(3.189)
Observations	723,979	720,570	698,862
R-squared	0.135	0.134	0.126

Robust t-statistics in parentheses, clustered by date and stock.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 5: Event Studies - Impact on Adverse Selection
This table reports coefficient estimates of our adverse selection proxies for three different event studies using the following difference in differences specification:

$$AdverseSelection_{sdv+m} = \alpha + \beta_1 event_d + \beta_2 treated_{sv} + \beta_3 eventTreated_{sdv} + \beta_4 Spread_{sd} + \beta_5 VIX_d + \beta_6 consideration_{sdv} + \beta_7 FE_{sd} + \epsilon_{sdv}$$

Where $AdverseSelection_{sdv+m}$ is (a) the trade value weighted price impact in basis points for stock s on date d for venue v over the time interval m which is the 100 millisecond period before to after the trade and (b) the % of dark trades for which there is a change in the midpoint over the 100 millisecond period before and after the trade, for stock s on date d for venue v over the time interval m. $Event_d$ is a dummy variable which takes the value of one on and after the event in question, and zero otherwise. $Treated_{sv}$ is a dummy variable which takes the value of one for stocks traded on venues which are subject to the market design change, and zero otherwise. $EventTreated_{sdv}$ is the interaction term for venues which are subject to the treatment, after the event occurs. $Spread_{sd}$ is the stock's time weighted average quoted spread. $Consideration_{sdv}$ is the natural log of the sum of trading consideration in British Pounds or US dollars. We also use stock and date fixed effects. VIX_d is the natural log of the mean value of the VIX volatility index, or the VFTSE (for UK event studies), measured at 15 seconds intervals throughout the day. We also use stock and date fixed effects.

	Uncross		NASDAQ SIP	
VARIABLES	(a)	(b)	(a)	(b)
Event	1.105	0.103	-0.743	-5.507
	(1.628)	(1.289)	(-1.249)	(-0.440)
Treated	-0.248***	-0.091***	0.113*	3.337***
	(-14.471)	(-31.660)	(1.846)	(2.668)
Treated_event	-0.126***	-0.041***	0.012	1.695***
	(-5.631)	(-10.440)	(1.044)	(6.895)
Consideration	-0.053***	-0.022***	-0.015***	0.952***
	(-6.252)	(-15.553)	(-2.577)	(5.118)
Spread	0.000*	-0.000**	-0.000***	0.000
	(1.820)	(-2.224)	(-3.473)	(1.291)
VIX	5.181*	0.584*	4.758	46.288
	(1.754)	(1.930)	(1.160)	(0.537)
Constant	-12.510	-0.593	-16.995	-183.971
	(-1.529)	(-0.712)	(-0.942)	(-0.485)
Observations	37,921	39,684	196,689	196,689
Adj. R-squared	0.143	0.246	0.463	0.644

Robust t-statistics in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 6: BATS Batch Auctions - Impact on Adverse Selection
This table reports coefficient estimates of our adverse selection proxies for the following specification:

$$AdverseSelection_{sdv+m} = \alpha + \beta_1 batch_{sdv} + \beta_2 Spread_{sd} + \beta_3 VIX_d + \beta_4 consideration_{sdv} + \beta_5 FE_{sd} + \epsilon_{sdv}$$

Where $AdverseSelection_{sdv+m}$ is (a) the trade value weighted price impact in basis points for stock s on date d for venue v over the time interval m which is the 100 millisecond period before to after the trade and (b) the % of dark trades for which there is a change in the midpoint over the 100 millisecond period before and after the trade, for stock s on date d for venue v over the time interval m. $batch_{sv}$ is a dummy variable which takes the value of one for trades that are flagged as BATS periodic batch auction trades for stock s date d and venue v. $Spread_{sd}$ is the stock's time weighted average quoted spread. $Consideration_{sdv}$ is the natural log of the sum of trading consideration in British Pounds or US dollars. We also use stock and date fixed effects. VIX_d is the natural log of the mean value of the VIX volatility index, or the VFTSE (for UK event studies), measured at 15 seconds intervals throughout the day. We also use stock and date fixed effects.

VARIABLES	(a)	(b)
Batch	-0.373***	-0.166***
	(-8.211)	(-15.905)
Consideration	-0.053***	-0.018***
	(-7.432)	(-15.032)
Spread	0.001***	-0.000
	(4.345)	(-0.249)
VFTSE	-0.348*	-0.073*
	(-1.900)	(-1.784)
Constant	2.445***	1.167***
	(4.742)	(10.225)
Observations	40,119	41,306
Adjusted R-squared	0.359	0.263

Robust t-statistics in parentheses
Standard errors adjusted for clustering in stock and date
*** p<0.01, ** p<0.05, * p<0.1