FCA Public

Financial Conduct Authority

64

Occasional Paper

March 2024

OTC market frictions in stressed markets

Jef Van Cappellen, Fabian Garavito, Yordan Georgiev and Andrew Lepone

FCA occasional papers in financial regulation

The FCA occasional papers

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

The main factor in accepting papers is that they should make substantial contributions to knowledge and understanding of financial regulation. If you want to contribute to this series or comment on these papers, please contact Kieran Keohane (Kieran.Keohane@fca.org.uk)

Disclaimer

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

Authors

Jef Van Cappellen is with the FCA, Rozetta Institute, Antwerp University and Macquarie University.

Fabian Garavito was with the FCA at time of writing and is now at the Bank of International Settlements.

Yordan Georgiev is with the FCA.

Andrew Lepone is with Rozetta institute and Macquarie University.

Acknowledgements

We would like to thank Carla Ysusi and Hannes Ortner for help in preparing this study.

All our publications are available to download from www.fca.org.uk. If you would like to receive this paper in an alternative format, please call 020 7066 9644 or email publications_graphics @fca.org.uk or write to Editorial and Digital Department, Financial Conduct Authority, 12 Endeavour Square, London E20 1JN.

Executive Summary

This research explores the impact of over-the-counter (OTC) market frictions on liquidity within the UK gilt market, utilising the comprehensive MiFID II dataset spanning from September 2019 to November 2020. This period encompasses the critical phase of the COVID-19 pandemic, offering a unique context to analyse the dynamics of inventory, search, and bargaining frictions that are characteristic of OTC markets. Our findings reveal a substantial statistical correlation between the exacerbation of these prevalent OTC frictions and the observable decrease in market liquidity during the heightened stress of the COVID-19 crisis. The study provides a deeper understanding of the resilience of financial markets under stress and offers insights for both policymakers and market participants on market functioning and stability.

1) Introduction

Recently, the resiliency of global government bond markets has been repeatedly tested. The UK government bond (gilt) market experienced two "once in a generation" shocks over the span of less than three years. Both the March 2020 COVID-19 event and the September 2022 volatility led to the largest 1-day surges in 10-year gilt bond yields in decades of 24 BPS[1] and 30 BPS[2], respectively. These large spikes in volatility were accompanied by a significant deterioration of liquidity in government bond market that are usually considered 'safe-havens' for investors. In both events, investors faced difficulty in finding counterparties to trade with and transaction costs increased substantially. Understanding the drivers of this deterioration in liquidity is important for identifying any structural issues that might drive illiquidity in the gilt market. In addition, using a comprehensive regulatory dataset that captures most gilt trading activity, we provide summary statistics of gilt market trading. These statistics are crucial to understand gilt market functioning in the absence of a publicly available, consolidated feed of market data.

Government bond markets are crucial for the financial system, serving to finance government activities, act as a safe haven during stress, set benchmarks for pricing risky financial instruments, provide key collateral assets, and meet capital or liquidity regulations for financial institutions. A sound understanding of how government bond markets are operating is critical for both market participants and regulators alike. The liquidity, structure, and resilience of core government bond markets are analysed by national authorities and international bodies, including a recent FSB stocktake report on the topic[3].

For the FCA, market resilience is a strategic objective, ensuring that markets remain orderly in a variety of conditions. Moreover, market stability builds confidence in UK financial markets and provides a foundation for increasing investments in the UK, which supports productivity and growth. Therefore, insights in this paper have relevance for the FCA objectives linked to wholesale markets and broader policy thinking. FCA is already taking steps to improve the functioning and price formation in the gilt market, and fixed-income markets more broadly, through exploring enhancements in transparency requirements. There is also ongoing consultation on the framework for establishing a Consolidated Tape.[4]

Many government bond markets around the world are traded over-the-counter (OTC). Previous research has shown that OTC markets, segmented into distinct dealer-to-dealer (D2D) and dealer-to-client (D2C) markets and characterised by off-exchange, bilateral negotiations with dealers, lead to market frictions that can influence the price formation process and liquidity provision of these markets (see Bessembinder (2020) for a survey of fixed-income market microstructure). The most prominent frictions identified in the OTC markets literature that affect prices and liquidity are:

1. inventory risk of dealers who facilitate intermediation and

1 https://www.reuters.com/article/britain-bonds-idUSL9N29P01K

2 https://www.reuters.com/markets/europe/uk-bond-yields-set-biggest-daily-surges-decades-2022-09-23/

3 https://www.fsb.org/wp-content/uploads/P201022.pdf

4 https://www.fca.org.uk/publication/consultation/cp23-15.pdf

2. search-and-bargaining costs faced by clients when seeking counterparties to trade with.

Our analysis centres on the exogenous shock of the March 2020 COVID-19 event, which significantly disrupted global government bond markets. This provides an effective testing ground for analysing how the OTC market frictions impact prices and liquidity in normal times and in times of stress. The March 2020 COVID-19 shock showed that dealers struggled to facilitate the extraordinary selling pressure coming from clients. This compelled dealers to either curtail their liquidity offerings or temporarily exit the market. As a consequence, our analysis reveals that clients encountered a reduction in trading options which possibly weakened their bargaining power in bilateral negotiations (bargaining friction). As a result, the increased search-and-bargaining costs for clients seeking to trade could have caused trade prices to deviate from fundamental valuations.

To capture the impact of frictions on prices and liquidity, we use transaction cost as a proxy for liquidity (Aquilina et al. (2015), Hendershott & Madhavan (2015)) and we also compute measures aimed at proxying different market frictions. We follow Friewald and Nagler (2019) and use our non-anonymised regulatory dataset to compute frictions at the individual participant level. We document that OTC market frictions significantly worsen and are highly correlated with the deterioration in liquidity in periods of stress. Our results show that over the event period:

- Most of our proxies for inventory, search, and bargaining frictions increased significantly, with most frictions increasing by more than 1 standard deviation (spiking close to the Bank of England's intervention to stabilise markets on the 19th of March 2020[5]).
- There is a statistically significant correlation between our OTC friction proxies and the deterioration in liquidity, proxied by transaction cost. Across all frictions, we observe a 20% - 80% increase in transaction cost relative to the average transaction cost over the entire sample period.

Our findings suggest that the rapid deterioration of liquidity over the event period can be partially explained by frictions inherent in OTC markets.

This note is structured as follows. Section 2 describes the relevant OTC market literature, and the frictions present in these markets. We also summarise previous literature examining the impact of the COVID-19 shock on other fixed income markets. Section 3 provides a description of the institutional framework and section 4 describes the data used in this study. In section 5, we provide an overview of the gilt market, including descriptive statistics across transaction-level characteristics. Section 6 discusses how we measure liquidity and market frictions and provides summary statistics. In Section 7, we describe the event period by documenting the change in OTC market frictions over time together with a network analysis which illustrates how the market evolved when stressed. We also describe the change in trading behaviour of select client groups over the event period. In section 8, we empirically test whether the identified OTC market frictions impacted liquidity over the sample period. Section 9 concludes.

5 https://www.bankofengland.co.uk/-/media/boe/files/speech/2020/seven-moments-in-spring-covid-19-speech-by-andrew-auser.pdf?la=en&hash=43D022917D76095F1E79CBDD5D42FCD96497EA5F

2) Literature review

In this section, we briefly define market structure and its role in facilitating price formation and liquidity provision. We also identify the main frictions from the OTC market structure theory and explain how these frictions could amplify economic shocks.

Markets serve mainly two functions: i) incorporate new information into asset prices i.e., facilitate price formation and ii) provide a place where buyers and sellers are matched in a timely manner i.e. provide liquidity (O'Hara, 2003). Therefore, the structure of a market has an impact on price formation and liquidity provision. Financial markets can be broadly classified into two structures[6]. First, there are centralised exchanges, where all participants can trade with each other through central limit order books. Typically, equities are traded in centralised markets. Alternatively, there are decentralised, over-the-counter (OTC) markets, where buyers must search for sellers one by one in order to trade. This bilateral trading implies that when counterparties do meet for a trade their identities are known and bargaining over the terms of trade takes place. As prices are not widely available to all investors, OTC markets are regarded as opaquer with lower level of transparency. In addition, whenever quotes are available, they are indicative rather than firm commitments[7]. To limit investor search, intermediaries emerge to provide liquidity in these markets since investors seeking to buy and sell do not necessarily arrive simultaneously (Bagehot, 1971).

Gilts, like most other fixed-income assets, predominantly trade in an intermediated OTC market. To facilitate trades, intermediaries - or dealers in the case of gilts - commit their own capital and therefore take on inventory risk (Amihud & Mendelson, 1980; Garbade & Silber, 1976; Stoll, 1978). Consequently, the trading architecture of OTC markets results in search-and-bargaining frictions for clients as well as inventory frictions for dealers. These frictions could cause trade prices to deviate from fundamental frictionless prices (Feldhutter, 2012; Friewald & Nagler, 2019) and affect liquidity.

Focusing first on inventory frictions, dealers are the primary liquidity providers in decentralised OTC markets. By facilitating trades on a principal basis, dealers can build up substantial inventory positions and therefore accumulate risk on their balance sheet. In response, dealers set prices to mean-revert their inventory positions and thus manage their inventory risk. Ho (1978) and Ho and Stoll (1983) provide theoretical models that relate inventories to asset prices. The models predict that an increase in aggregate dealer inventories would lead to a widening of dealer spreads and vice versa. Although many empirical studies based on US equity markets (Camerton-Forde et al., 2010; Chordia et al., 2002) conclude that market makers are less willing to provide liquidity at favourable prices when inventory balances increase, studies investigating the impact of inventory levels on prices in fixed income markets (especially on non-US sovereign bonds) are lacking at present. In contrast to equity markets, trading in the gilt market is decentralised and bilaterally negotiated. Increased search effort by dealers in more fragmented markets could delay trading (Duffie, 2010). Moreover, the lack of anonymity also implies greater

6 Some markets are comprised of elements of both OTC and exchanges markets

7 Major reason why high frequency traders, whose business model relies on firm, executable quotations, are not present in many OTC markets like gilts (Bessembinder, 2020).

information revelation when managing inventories. This suggests that inventory management might be of greater concern to dealers in fixed income markets relative to equity markets. Friewald and Nagler (2019) show that in addition to search and bargaining frictions present, inventory frictions significantly determine prices in the US corporate bond market. They also find that inventory frictions increase when markets are stressed.

More complex inventory models incorporate dealer risk aversion and funding constraints in fixed income markets (Brunnermeier & Pedersen, 2009; Nagel, 2012). These models imply that elevated levels of risk and funding constraints lead to less liquidity supplied and worse prices for clients. Moreover, these models posit that inventory frictions increase in times of stress when holding cost and risk aversion increase due to tighter funding constraints and higher price volatility[8]. Lastly, Boyarchenko et al. (2017) study the effect of post GFC banking capital regulation on market liquidity. Although tighter capital regulation improves financial stability in normal times, the study showed that capital regulations may reduce banks' willingness to hold inventory. Furthermore, they argue that banks' reduced willingness to hold inventory due to either inventory risk concerns or capital regulations could reduce their ability to intermediate especially in times of stress, when inventory holding costs are at their highest.

Turning to search-and-bargaining frictions, investors who want to trade need to search for counterparties willing to trade. Contact might not be immediate and therefore intermediaries may need to be contacted sequentially which is costly in terms of time (need for immediacy) and information leakage. Once found, counterparties engage in bilateral bargaining regarding the terms of trade. Due to limited price transparency, this can result in substantial search-and-bargaining costs for the investor (Edwards, et al., 2007; Green, et al., 2007). The search-and-bargaining model of Duffie et al., (2005) relates this friction to prices. They contend that, in the absence of inventory risk and informational asymmetries between investors, prices are determined by an investor's ability to find counterparties more easily i.e., have more alternatives to trade. These models predict that increased investor access to intermediaries results in better prices since competition for investor order flow forces intermediaries to provide better prices. Empirically, (Friewald & Nagler, 2019; Jankowitsch et al., 2011) confirm that search-and-bargaining costs significantly impact prices and therefore liquidity in global OTC markets. Similarly, Pinter and Uslu (2022) document that search-and-bargaining frictions exist in the gilt market and ascribe most of the friction to the dealer's significant market power. Moreover, in stress periods it is harder to find a willing counterparty to sell asset to and bargaining positions for sellers deteriorate (Duffie, 2007). Therefore search-and-bargaining frictions could increase in times of stress, affecting prices and therefore the resiliency of OTC markets.

Liquidity can be defined as the extent to which asset prices deviate from their market value (Foucault, et al., 2013). In the absence of a clear benchmark for market value, past academic studies use average hourly or end of day prices (Aquilina et al, 2015; Pinter and Uslu, 2022). However, in our paper, we exploit the market structure of gilts and follow Hendershott and Madhavan (2015) by using the nearest D2D price of the same instrument as market value. The ability of dealers to manage their inventory risk in the D2D segment determines the cost of liquidity provision to clients in the D2C segment (Benos and Zikes,

8 In high volatility periods intermediaries refrain from making markets in riskier assets as these are more costly to fund (higher rates and margin requirements) and for risk management reasons (risk limits on inventory positions)

2016). Therefore, when a client trades at successively worse prices than the nearest D2D price, market frictions are impairing overall liquidity in the D2C segment.

In our analysis, we explore the impact of frictions on market liquidity in a period of stress. Global fixed income markets experienced a major deterioration of liquidity during the March COVID 19 shock. The event has been explored in the US context. The crisis was characterised by large and persistent selling pressures from clients (Ma, et al., 2022) in the so called 'dash for cash' period which overwhelmed dealer inventories (Kargar, et al., 2021). Meanwhile, regulatory capital requirements constrained dealers' ability to absorb the selling pressure (Duffie, 2020). As a consequence, intermediaries were unable to absorb client order flow which led to a 15 year high in US government bond volatility (Fleming & Ruela, 2020) and an 8-fold increase in US corporate bond transaction costs (Kargar, et al., 2021), resulting in intermediaries providing less depth at worse prices in both US government bond (He, et al., 2022) and US corporate bond (O'Hara and Zhou, 2021) markets. In response, central banks around the world intervened at a then never-before-seen speed and scale to restore liquidity (Haddad, et al., 2021). In what follows, we investigate how the UK government bond market responded to similar selling pressures in the face of OTC market frictions.

3) Institutional design

The UK Debt Management Office (DMO), an executive agency of the His Majesty's Treasury (HMT), is responsible for the issuance and management of UK government debt. To facilitate trading in this OTC market, the DMO designates some qualifying dealer-banks as Gilt-Edged Market Makers (GEMMs)[9]. These GEMMs (the sell-side) are obliged to provide continuous two-way spreads to clients (the buy side) in the secondary market in return for sole participation rights in the primary market for gilts[10]. In practice, 97%[11] of gilt trade volume is intermediated by GEMMs, which makes them the primary source of liquidity in the market. Although gilts are listed on the London Stock Exchange, 99% of dealer-to-client trade volume is done OTC.

OTC markets are characterised by having intermediaries (GEMMs in this case) that facilitate trading. They do this by either matching buyers and sellers if they meet contemporaneously or principally by taking on inventory positions on their balance sheet to match clients intertemporally. Dealers service their client's liquidity needs in the dealer-to-client (D2C) segment. In order to provide this service, dealers trade exclusively with each other in the dealer-to-dealer (D2D) segment to manage their inventory risk. These distinct segments of the overall market constitute the market structure of gilts (Diagram 1).

Trading in the gilt market is governed under MiFID/MiFIR, which organises trading within each segment onto specific venues. In the D2C segment, clients can engage with dealers in two ways: bilaterally or multilaterally on a trade venue. Bilateral trades either go through the dealer's own platform called Systematic Internaliser (SI) or are traded off-exchange (XOFF)[12]. On the other hand, Multilateral Trading Facilities (MTFs) provide the client access to multiple competing, non-anonymised, dealer quotes after the client submits a request-for-quote (RFQ). In the D2D segment, the majority of trading is facilitated by wholesale brokers (WSB) who run their own anonymised platforms called Organised Trading Facilities (OTFs) that facilitate trade between dealers on a matched principal basis[13].

9 https://www.dmo.gov.uk/responsibilities/gilt-market/market-participants/

The GEMMs over our sample period include:
Wholesale: J.P. Morgan Securities Plc, Banco Santander, S.A., Citigroup Global Markets Limited, Lloyds Bank Plc, HSBC Bank Plc, Merrill Lynch International, BNP Paribas, RBC Europe Limited, Goldman Sachs International, Natwest Markets Plc, Morgan Stanley & Co. International Plc, Barclays Bank Plc, Deutsche Bank Aktiengesellschaft, Nomura International Plc, UBS Ag, The Toronto Dominion Bank. Retail: Winterflood Securities Limited, Jefferies International Limited.

10 For more information see GEMM guidebook: https://www.dmo.gov.uk/media/22bbjndz/guidebook200921.pdf

11 Remaining 3% in our sample are so called client-to-client trades which mostly constitute intra-company or fund transfer trades we are not interested in for the purpose of this study.

12 These are classic, bilaterally negotiated, over-the counter type trades between counterparties on matched principle . XOFF trades are also not subject to the same level of transparency as SI trades. Due to reporting issues and inconsistencies some of these trades can be regarded as SI trades.

13 Over our sample period there are 5 wholesale brokers, which can be found on the DMO website. The three largest in terms are volume are TP ICAP Markets Limited, GFI Securities Limited and BGC Brokers L.P.

Diagram 1: Institutional Layout

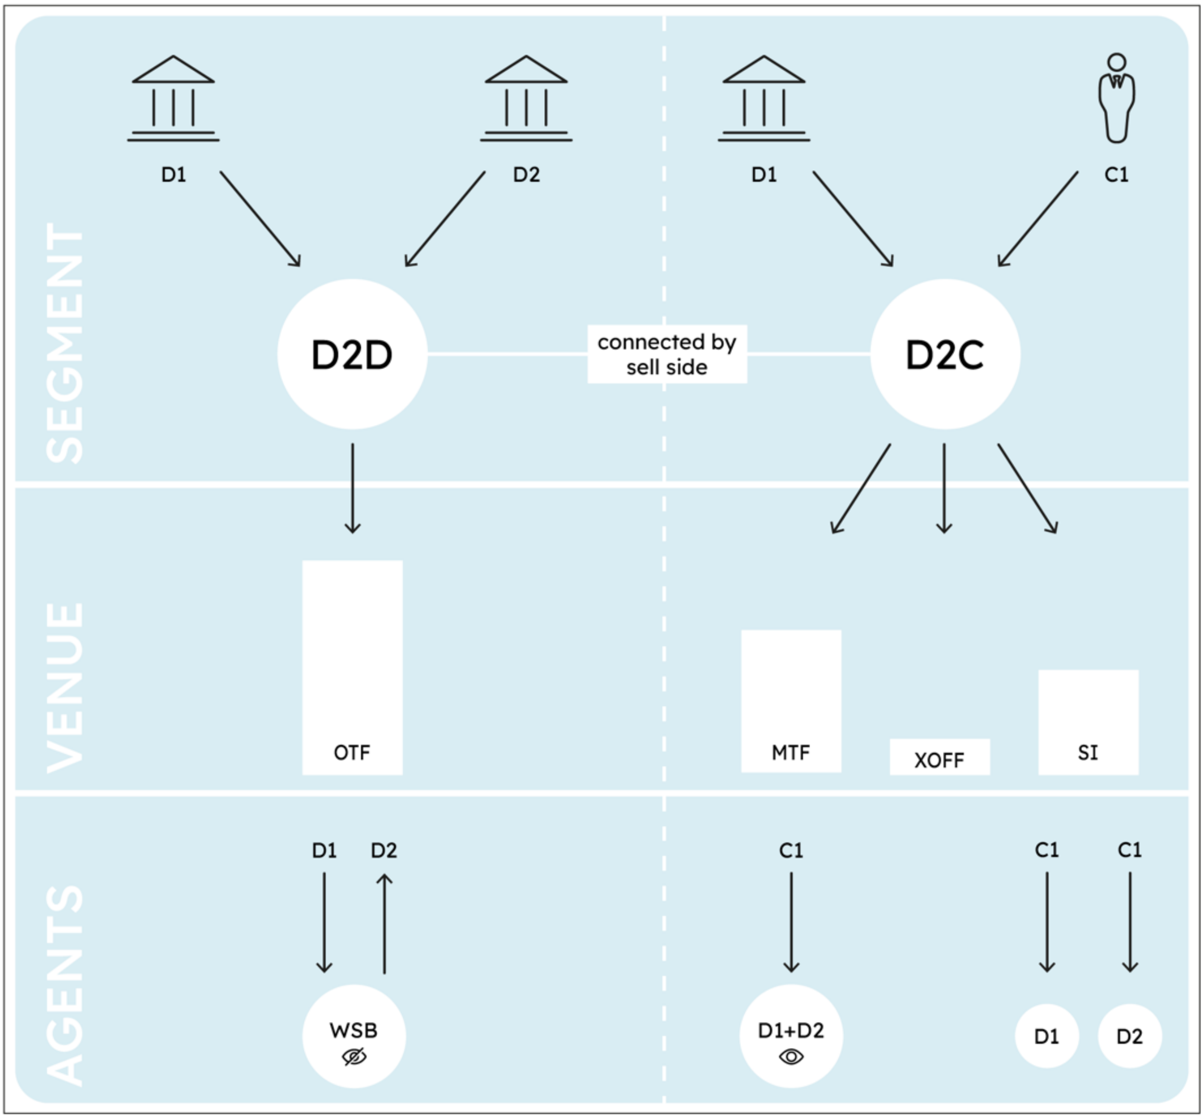

4) Data

The paper uses regulatory transaction-level data, submitted to the FCA, containing UK gilt trades – so-called MiFID II database[14]. Under UK MiFID/MiFIR, all FCA regulated parties to a trade must submit a transaction report to the FCA following strict UK MiFID/MiFIR transaction reporting guidelines[15]. Each transaction report contains information on the transaction execution date and time, instrument identifier (ISIN), execution price, transaction size, and in contrast to other datasets typically employed in the literature (e.g., TRACE), the legal identities of the buyer and seller. Since almost all trades are intermediated by GEMMs who are UK-domiciled, the MiFID II data set captures virtually all secondary market gilt trading activity. To capture secondary gilt market activity of clients only, we exclude primary market HMT activity and Bank of England trades. We use publicly available data from the Bank of England's website on their Asset Purchasing facility[16] when discussing its intervention over the event period. We supplement the transaction-level data with data from Bloomberg and Refinitiv Eikon databases to get (i) bond characteristics such as issuance date, maturity, benchmark status, amount issued and amount outstanding and (ii) min, max, open and close price data which we use to clean the data from transaction reports with implausible prices[17].

Due to the possibility of more than one transaction report for a given trade and to avoid overstating trade volumes, we employ a matching algorithm based on transaction information to clean the data from: (i) duplicate reports resulting from trades between two FCA-regulated firms; (ii) duplicate reports resulting from brokered trades where both broker and counterparties submit a transaction report. The latter trades typically occur in the D2D segment where most dealers trade on a wholesale broker platform. We also group counterparties using the GLEIF level 2 dataset[18] in combination with internal FCA records, to match entities to their respective ultimate parent entities. By doing this we increase matching efficiency across duplicate reports due to naming inconsistencies and identify intra-company trades which we discard from the sample. This ultimate parent grouping also makes classifying the sample into participant groups more efficient.

In total, our cleaned sample covering the period from September 2019 to November 2020 consists of 55 conventional gilts[19]. We end up with 657538 transactions of which around 30% are D2D and 70% are D2C. Over the period, we identify 16 wholesale dealers and two retail dealers using the GEMM list on the DMO website. The D2C segment consists of 1396 clients after grouping participants at the ultimate parent level and removing participants that trade less than £1 million over the entire sample period. This omits less than 0.001% of volume from the overall sample and ensures that these small participants do not skew client-level statistics.

14 MiFID II database is owned and maintained by the Financial Conduct Authority.

15For guidelines consult the ESMA guidelines https://www.esma.europa.eu/sites/default/files/library/2016-1452_guidelines_mifid_ii_transaction_reporting.pdf

16 https://www.bankofengland.co.uk/-/media/boe/files/markets/asset-purchase-facility/gilt-purchase-operational-results.xlsx

17 We identify prices that are more than 200 BPS outside of the (min, max) price band reported by Bloomberg for an ISIN on a given day. First, we repair these prices with the price of the counterparty report if that price is within the band. Otherwise, we discard the transaction report.

18 https://www.gleif.org/en/lei-data/access-and-use-lei-data/level-2-data-who-owns-whom

19 For the sake of simplicity and comparability of prices and volumes we focus our study only on conventional gilts i.e., plain vanilla bullet bonds.

5) Descriptive results

Descriptive statistics that depict the entire gilt market are hard to come by due to trade fragmentation across trade venues and the lack of a consolidated feed of trade data. It is therefore interesting to use this unique cross-venue regulatory dataset to illuminate previously, lesser-known aspects of the market and gain a better understanding of how the overall market operates. We start off with a high-level descriptive analysis of the two market segments before we analyse specific bond-level characteristics.

Table 1: Transaction level descriptive statistics

Characteristic	Panel A: D2D					Panel B: D2C				
	#ISIN	Mean (£ mln)	Volume %	#Trades %	Mean daily trades	#ISIN	Mean (£ mln)	Volume %	#Trades %	Mean daily trades
Overall	55	9.4	37.6	30.2	622	55	6.7	62.4	69.8	1400
Maturity										
0-2	9	13.0	4.9	3.6	22	9	5.6	10.5	12.6	181
3-5	12	10.7	18.3	16.1	101	12	9.5	19.3	13.7	197
6-10	13	11.9	46.6	36.7	230	13	11.8	38.2	21.7	307
11-20	10	6.1	8.5	13.2	82	10	4.6	10.3	15.2	218
21-50	19	6.7	21.6	30.4	190	19	4.0	21.6	36.7	524
Size										
Small (<1 £mln)	54	0.5	0.7	12.6	78	55	0.2	1.9	59.3	838
Medium (1 - 50 £mln)	55	9.1	82.9	85.3	534	55	9.7	54.4	37.6	535
Large (> 50 £mln)	51	75.3	16.4	2.0	13	55	96.6	43.7	3.0	44
Status										
Benchmark	8	11.0	53.8	45.9	286	8	10.2	38.5	25.4	361
Non Benchmark	47	8.0	46.2	54.1	338	47	5.6	61.5	74.6	1056
Venue										
MTF	55	2.7	6.4	22.3	139	55	4.5	51.3	77.5	1109
OTF	55	11.0	87.8	74.9	469	-				
SI	55	21.9	4.4	1.9	12	55	20.9	37.0	11.9	170
XOFF	52	14.6	1.4	0.9	7	55	7.5	11.7	10.5	150

This table reports the descriptive statistics for both market segments and per bond characteristic. #ISIN represents the number of distinct instruments trade. Mean represent the equally weighted mean trade size across all bonds. Volume represents the proportion of total notional traded. #Trades represent the proportion of the number of transactions. Mean daily trades is calculated by taking the average of the total amount of trades per day. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

Figure 1: Time series of weekly D2C market share

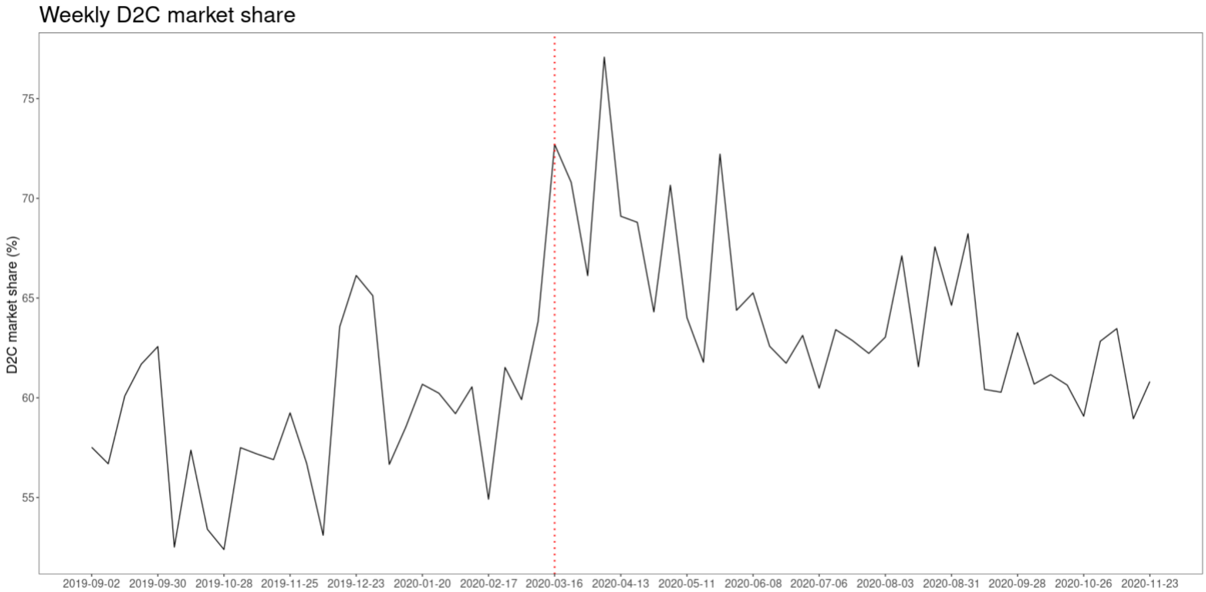

This graph plots the ratio of D2C trade volume divided by total trade volume on a weekly basis. Volumes represent notional traded. Dotted line represents the March 19th, 2020 Bank of England emergency COVID-19 intervention. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

5.1) Market overview

The top row of Table 1 reports summary statistics for the overall market. Relative to equities, gilts trade far less frequently, averaging around 622 trades per day in the D2D and 1400 in the D2C, compared to millions of trades in equity markets. However, the sizes of individual gilt trades are large. The average trade size in the D2D is £9.4 million which is larger than the average trade size in the D2C of £6.7 million, although the average trade size is highly dependent on bond-level characteristics discussed in section 5.2. In volume terms, the D2C accounts for nearly 2/3rd of total trade volume in gilts. However, this fluctuates considerably over our sample period depending on market conditions. Figure 1 reports the weekly time series of D2C market share calculated as the total D2C trade volume divided by total trade volume (D2D + D2C). There is substantial variation in D2C market share over time with values ranging from a low of 55% in late 2019 to nearly 75% in March 2020. The D2C market share is at its highest over the March 2020 COVID-19 event period. This could indicate that dealers in the D2D segment engaged in less risk sharing, relative to the increased D2C trade volume over the event period. It is however important to note that our sample excludes primary market transactions and Bank of England activity which might affect these statistics.

We also take a closer look at trade volumes over our sample period. Figure 2 depicts the monthly total volume traded. During our sample period of 14 months, the total notional value of gilts traded was about £4.95 trillion in both segments. Compared to the average notional amount outstanding of £1.4 trillion, that means that secondary market activity turned over 3.5 times the amount outstanding over the sample period. For the entire gilt market, monthly total trade volumes are between £220 billion and £420 billion. We see that over the sample period D2D trading remained relatively stable at an average level of around £112 billion traded per month. In contrast, D2C average trading is higher at a level of around £200 billion per month and increases by over 40% (month-on-month) over the

March 2020 COVID-19 event period. To illustrate the scale and timing of the Bank of England's intervention we plot their net asset purchases using data published on Bank of England's website (red bar). The spike in D2C trade volume is compounded to a 60% month-on-month increase, when one includes the Bank of England's asset-purchasing programme to the D2C volumes. In total, the Bank of England's intervention from March 2020 until the end of the sample period amounts to around £260 billion, with £163 billion released over the first three months after the intervention.

Figure 2: Time series of total trade volumes across segments

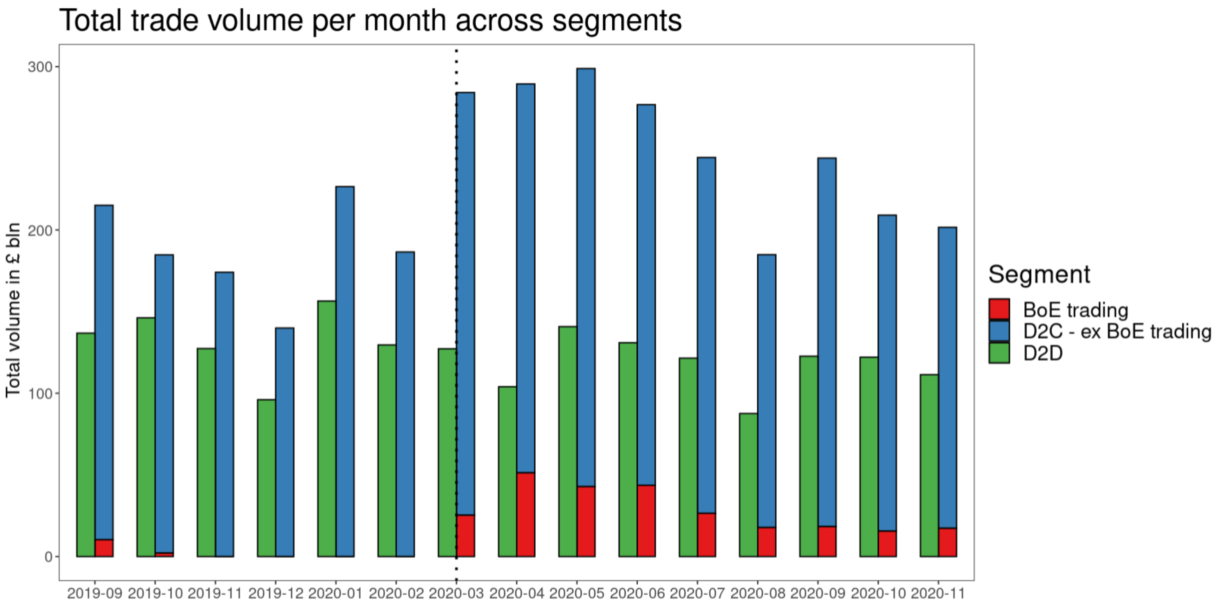

This graph plots the monthly total notional volume traded (£ billion) for each segment individually. We use data from the Bank of England website to calculate the monthly contribution of the Bank of England's Asset Purchasing Facility to the volume in the D2C segment. Dotted line represents the March 19th, 2020 Bank of England emergency COVID-19 intervention. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

5.2) Transaction level characteristics

Table 1 provides descriptive statistics on the most important bond-level characteristics across both segments. The most actively traded bonds, in terms of volume, are the 6-10-year gilts representing almost half of all D2D trading and 38% of D2C trading, followed by the longest maturity bonds (>20 years) representing approximately 22% of total trading in both segments, and then the shorter maturity bonds (<5 years) accounting for 20% of trading in both segments. The prevalence of trading in 10-year maturities can be explained by the fact that gilts with a 10-year duration represent a quarter of all existing gilts. Moreover, the 10-year gilt future is the sole liquid gilt future. Consequently, the 10-year gilt is the most liquid maturity.

Average trade size for D2D trading is approximately £10 million with a bimodal distribution around £5 million and £15 million pounds. This could indicate that dealers prefer to trade in 'standardised' sizes to minimise information leakage and to facilitate easier risk transfer among themselves. In contrast, D2C trading occurs in a variety of trade sizes. This could be indicative of the heterogeneity of clients, each trading for a diverse set of reasons, matching their individual needs and wealth levels. We see substantially more trade activity in the small size category in the D2C segment compared to the D2D (59% v 13%).

In the D2C trading in benchmark[20] gilts constitute 39% of volume traded while in the D2D benchmark gilts trading comprises 54% of trade volume. It should be noted that there are 47 non-benchmark gilts and only 8 benchmark gilts in our sample. Taking this into account, a given benchmark is twice as likely to trade compared to a non-benchmark gilt in the D2C and 5 times more likely in the D2D based on mean daily trades per ISIN. A possible explanation for this could be that most clients in the D2C have a buy-and-hold strategy coupled with longer investment horizons (e.g., pension funds) and can therefore take on more liquidity risk than dealers who, typically, do not keep gilts in inventory very long and therefore prefer the most liquid gilt per maturity i.e., the benchmark to manage their inventories and meet duration risk targets. In addition, non-benchmark trades are usually smaller than trades involving benchmarks in both D2D and D2C segments. Table 1 shows that the D2D mean trade size is £11 million for the benchmark and £7.9 million for the non-benchmark. More pronounced is the difference in the D2C segment where the average trade size of benchmark gilts are almost twice as large as non-benchmarks (£10.2 million v £5.6 million).

Trading across segments also vary depending on the MiFID II trade venue used. The vast majority, 88% of trade volume between dealers, occurs on OTFs. This is because trading on OTFs occurs anonymously which minimises a dealer's information leakage when they want to socialise their risk in the D2D. Trading in the D2C is more fragmented across bilateral and multilateral venues. Bilateral trading, such as SI and XOFF, account for 37% and 12% of total D2C trading volume respectively, whereas MTFs are responsible for 51% of D2C trading. As evidenced by the higher average trade size in Table 1, most D2C trading of large sized tickets are traded on bilateral (SI specifically) venues. This could be a consequence of the relatively lower level of transparency and direct competition which protects counterparties from information leakage. In contrast, the more transparent and competitive MTFs facilitate trading in smaller and medium trade sizes with a lower mean trade size of £4.5 million compared to £21 million on SIs.

20 Benchmarks are the on-the-run most liquid ISIN per standard maturity

6) Variable construction and summary statistics:

Table 2 presents summary statistics of the variables used later in the empirical analysis. All variables are computed on a daily frequency.

Following Hendershot and Madhavan (2015), we calculate per trade the log difference between the execution price and the nearest D2D price multiplied by the direction of the trade: +1 if client buys and -1 if the client sells. To get the full bid/offer spread, we multiply this spread by two and multiply by 10000 for conversion into basis points (BPS). Following standard naming convention in the literature, we use the term *transaction cost* for this measure. More details on the metric are presented in section 8. We compute the average daily transaction cost across all bonds. We do this by taking the equal-weighted average across all transactions on a given day. From Table 2 we see that the average daily transaction cost is about 1.2 BPS with a standard deviation of 4.4 BPS[21]. Since our transaction cost measure is directed, this means that on average a client pays 1.2 BPS to trade in the D2C. Or, alternatively, dealers charge an average two-way spread of 1.2 BPS for providing liquidity to clients in the gilt market. However, this average varies a lot with market conditions and bond-level characteristics that are explored further in section 7. Consistent with Jurkatis et al. (2022), who observed that transaction costs are predominantly positive, our findings reveal that transaction cost can sometimes be negative, suggesting dealers occasionally offer price concessions to clients with whom they have established relationships.

Table 2: Summary statistics of regression variables

Variable	Mean	Median	sd	q10	q90
Transactio cost (BPS)	1.18	0.47	4.40	-2.36	5.16
Market volume	8.84	8.14	3.95	4.71	14.21
Dealer volume	0.52	0.38	0.64	0.01	1.12
D2D Market share (%)	30.68	31.18	5.62	23.26	36.94
Dealer connections	11.72	13.00	3.10	7.00	14.00
HHI D2C	0.10	0.10	0.03	0.08	0.13
Outside options	1.60	1.00	1.23	1.00	3.00

This table reports the summary statistics of variables used in the empirical analysis. All variable are calculated on a daily basis. Transaction cost represents the average daily 2-way spread in BPS. Transaction cost is calculated as the log difference between executed trade price and the nearest D2D price per trade. We only use observations for which the nearest interdealer price is within 3 hours – which 92% of the transactions are. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

21 This is much lower than the average daily absolute transaction cost of 6.4 BPS with standard deviation of 2.5 BPS reported by Pinter and Uslu (2022), for gilts over a sample period ranging from 2011 to 2017. This discrepancy is partly driven by their use of hourly average price as benchmark price and by using the absolute (not directed) difference between transacted and benchmark prices. When we replicate their method, we find a more comparable average daily transaction cost of 4.2 BPS with standard deviation of 2.6 BPS. Using an hourly average price, rather than the nearest D2D price induces an upward bias due to the intra-hour volatility of prices due to the arrival of news in the market. Additionally, we prefer to use the D2D price since, for some off the run gilts, D2C trading is limited and it does not include costs dealers charge to other clients to trade.

Dealers' balance sheets are the primary source of liquidity in OTC markets. Therefore, any pressure on dealer inventories would impact prices and liquidity in the gilt market. To determine the extent of inventory frictions, we calculate market wide and dealer specific volumes traded using this as a proxy for dealer balance sheet constraints (Benos & Zikes, 2016; Friewald & Nagler, 2019; Duffie, et al., 2023). *Market volume* is calculated as the sum of (unsigned) volume traded in the D2C on a daily basis across all dealers. An increase in this measure suggests that dealers' balance sheets, collectively, face greater pressure since they need to intermediate increased order flow from clients. Dealers intermediate £8.84 billion of volume on an average day over the sample period. *Market volume* does fluctuate significantly on a day-to-day basis indicating varying demand for liquidity by clients. *Dealer volume* is calculated similarly but at the individual dealer level i.e., is the gross volume traded with clients per dealer per day. Table 2 reports that the average dealer intermediates £500 million per day. We observe a lot of variation in this measure with some dealers trading over £1billion on certain days[22].

Dealers offload their inventory risk accumulated when providing liquidity to clients in the D2C segment by trading with other dealers in the D2D segment. A more active D2D market where dealers are better connected and therefore do not have to search extensively to find other willing dealers to trade with, would improve the dealer's ability to provide liquidity to clients in the D2C segment (Di Maggio, 2016). Following Benos and Zikes (2016), a market wide measure reflecting the ease with which dealers can search for each other is *D2D market share*. We calculate this as the ratio of number of D2D trades divided by total number of D2D plus D2C trades per day. Lower values might suggest a less active D2D segment which would make it harder for dealers to offload their inventories to other dealers. We see that the D2D segment is always smaller than the D2C segment and accounts for 31% of total gilt trades (33% in terms of volume), although this fluctuates over the sample period. An alternative measure of search, calculated at the dealer level, is to calculate connectivity by calculating the number of *dealer connections* per dealer in the D2D segment (Soramaki et al., 2007). A link forms when dealers trade a gilt with each other. This measure resembles how well dealers are connected with each other. A higher value implies easier risk sharing in the D2D. As a result, dealers might be more willing to service the needs of clients at a more favourable price in the D2C. As expected, the dealers in the D2D are well connected with each other with the average dealer being connected to 11.72 (or 73%) other dealers in the D2D segment. The high connectivity among dealers is primarily due to the limited number of primary dealers participating in the market. However, there is a noticeable variation in *dealer connection*s throughout the sample period.

After searching for dealers to trade with, clients need to bargain the terms of trade. The search and bargaining framework of Duffie (2005) states that clients have less bargaining power vis a vis dealers if they have fewer *outside options* (fewer dealers to trade with) or if client order flow becomes more concentrated. We compute two widely used measures to capture the bargaining power of both dealers and clients in the D2C market. Following

22 We also measure net order flow as the aggregate (signed) volume across all dealers. On average dealers are net sellers per day to the tune of 4% of daily market volume (significantly different from 0). Something we would expect since: i) dealers act as conduit between the primary market and clients in the secondary market and ii) the largest investors in this market, such as pension fund managers and insurance companies, most often buy and hold gilts in their portfolios for long periods of time. However, there is substantial variation in net order flow over the sample period depending on wider market conditions and across dealers.

Duffie (2005), we measure a client's bargaining power with their *outside options.* This measure calculates how many dealers a client trades with on a given day. A higher value would indicate that clients on average have more bargaining power since they have the ability to trade with more dealers. Table 2 reports that clients are on average connected to less than 2 dealers (1.6). Moreover, by looking at the narrow range of values between the top and bottom 10% (only 2) and low standard deviation (1.23) client's *outside options* seem very persistent across the sample. Possible explanations for this lack of variation could be: i) the cost of setting up a trading relationship with a dealer is non negligible and ii) once a trading relationship is established, they tend to be persistent due to the benefit of repeat business (Hendershott et al., 2020; Jurkatis, et al., 2022). Dealer bargaining power in the D2C could also be proxied by the Herfindahl-Hirschmann index or *HHI D2C (Bicu et al*., 2017). Daily *HHI D2C* is calculated as the squared sum of each dealer's share of total trade volume. HHI can take on a value between 0 and 1, where 1 indicates a perfect monopoly and near-zero full competition. In our context, the index captures the concentration of market-making activity. Higher values indicate higher concentration and potentially less competition among dealers for client order flow. We see that the average daily *HHI D2C* is 0.1 which indicates that D2C trading activity is not very concentrated. This can be expected since the DMO expects that GEMMs provide liquidity to clients in the D2C in return for sole participation rights in the primary market. We do however see some variation over the sample period indicating that some dealers intermediate more client order flow than others in certain periods.

7) Gilt market under stress

In this section we analyse how the gilt market evolved over the March 2020, COVID-19 event period. First, we describe what happened over the event period by looking at what happened to prices and trade volumes. Next, we describe the evolution of frictions when markets are under stress. Specifically, we analyse how inventory, search and bargaining frictions evolved over the event period. Lastly, we conduct a network analysis to see how the market and its participants changed over time and analyse the behaviour of certain client groups.

7.1) Prices and volumes over event period

The event can be subdivided into 3 subperiods[23] each with their own characteristics and implications for prices (Figure 3) and volumes traded (Figure 4):

Period 1: 'Flight-to-safety' - ranging from March 1-9, 2020 (represented by the yellow shaded bar in Figure 3). In response to the growing risk of a global pandemic investors substituted, as is often the case in stress periods, risky assets for less risky assets which lead to an increase in demand for gilts. This led to an increase in prices for most gilts, especially for longer maturity gilts and non-benchmarks. Whereas prices for short (er) term gilts (0-5 years) remained relatively unchanged over this period, longer term gilts (10-50 years) increased substantially, with the longest maturity bracket increasing over 10% during the week of March 1-9, 2020. Prices for non-benchmark gilts (6%) rose twice as fast as for benchmark gilts (3%) over this period.

Period 2: "Dash-for-cash" - on the 11th of March 2020, the WHO officially declared a global pandemic. In response, many countries around the world announced strict lockdown measures. These actions significantly worsened the global macro-economic outlook and increased market uncertainty. For various reasons[24] investors responded by broad-based selling of financial assets and even assets normally considered as safe, such as gilts, were not immune to the selling frenzy[25]. This period is referred to as the 'dash-for-cash' (represented by the red shaded bar in Figure 3) and lasted from March 11-18, 2020. Prices plummeted by nearly 5% over this period, on average, across all bonds. For context, this is twice as much as any other weekly aggregate price change over the sample period. Similar to "Flight-to-safety" period longer maturity gilts were more affected than shorter maturity gilts, with the prices of the longest maturities declining by nearly 10% over the week. Prices of non-benchmarks were more reactive and fell twice as fast as benchmark gilts (6% vs 3%).

Period 3: "Recovery phase" - on the 19th of March, 2020 the Bank of England took substantial remedial action to calm the gilt market by announcing it would roll out its largest and fastest asset purchasing scheme at the time and cut the Bank rate to 0.1% (Hauser, 2020). This phase is the so called 'recovery phase' (represented by the green shaded bar in Figure 3). Figure 4 illustrates the market dynamics that necessitated the Bank of England's intervention and its scale in order to restore market functioning. In

23For comparability and accuracy, we use the same dates used in other Bank of England studies covering the COVID-19 event (see Czech et al., (2021)).

24 Reasons for broad based selling include redemptions from investment funds, fulfilling margin requirements, deleveraging and general fears concerning macro-economic outlook.

25 For more details see The FSB's Holistic review of the March market turmoil: https://www.fsb.org/wp-content/uploads/P171120-2.pdf

intermediated markets, like the gilt market, intermediaries aim to maintain fairly neutral inventory positions[26] over longer time periods like a month, as shown by the neutral *net dealer balances* in Figure 4 for the months prior to March 2020 (dashed vertical line). However, during and for some months after March 2020, the significant increase in one-sided net selling volume i.e., *client selling pressure*[27] led dealer inventories to substantially increase (before taking into account the Bank of England's asset-purchasing programme - blue bars in Figure 4). Large client gilt selling coupled with post GFC regulatory constraints[28], severely impacted dealers' ability to intermediate in this market. To help dealers cope with the increased *client selling pressure*, the Bank of England intervened by buying gilts directly from dealers (red bars in Figure 4)[29]. This freed up necessary dealer balance sheet space which restored their ability to intermediate in this market resulting in prices (Figure 3) and transaction cost (see Figure 5) returning quickly to pre-event levels.

Figure 3: Price evolution over event date

Time series of prices for the 10-year benchmark gilt (ISIN = GB00B24FF097). We use the MiFID II database to plot daily transaction prices. From Bloomberg, we retrieve daily closing prices. The yellow bar indicates the 'flight-to-safety' period ranging from March 1-9, 2020. The red bar resembles a 'dash-for-cash' episode from March 10-18, 2020. The green bar portrays the first two weeks after the Bank of England and government intervention on March 19, 2020 – the 'recovery phase'. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

26 We approximate dealer inventories by calculating net dealer trade volumes over a given month. The resulting net dealer balance does not account for accumulated inventories over past months.

27 For the purpose of this chart only, we include all trading activity that affect dealers' balance sheets. We therefore include both secondary market client trades and primary market, Her Majesty's Treasury trades, retrieved from the MiFID II database.

28 FSB papers discuss 'regulatory constraints' like leverage and liquidity which contributed to intermediaries' inability to absorb excessive client selling. See https://www.fsb.org/wp-content/uploads/P171120-2.pdf

29 The bank also stepped in to mitigate inventory balance sheet constraints by exempting central bank reserves in the calculation. See https://www.bis.org/bcbs/publ/d521.pdf

Figure 4: Trade volumes over event period

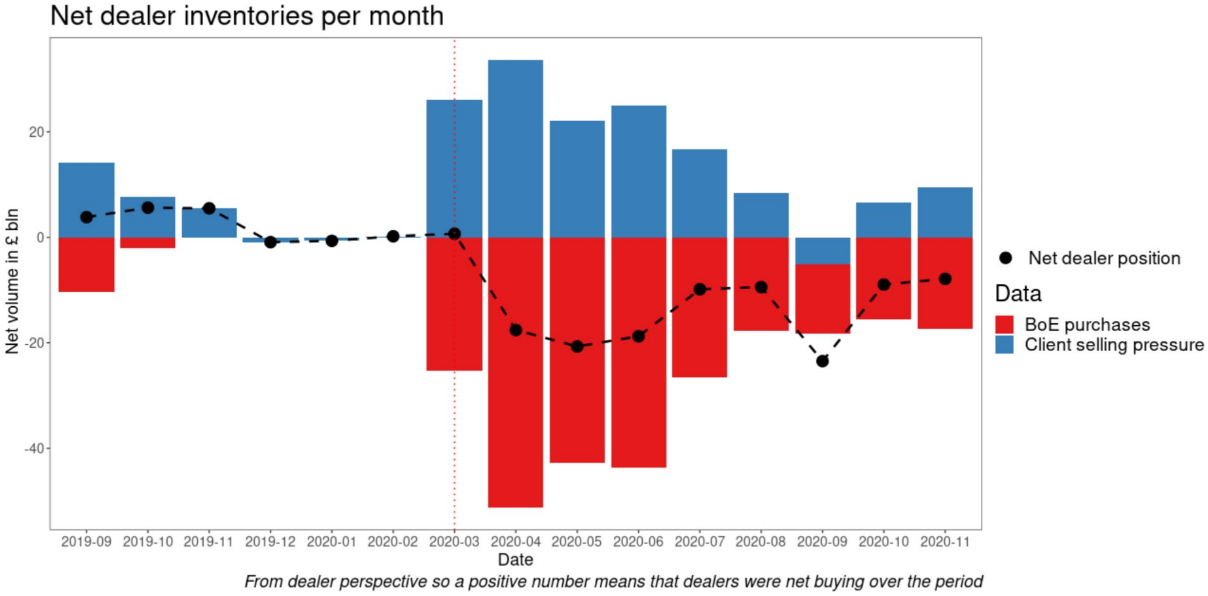

Breakdown of monthly trading balances over time. Bank of England purchases (red bar) is calculated using data from the Bank of England website. Client selling pressure reflects the net selling volume dealers intermediated based on MiFID II database. This includes both client and His Majesty's Treasury trading activity. Net dealer balances are calculated as the difference between client selling pressure and Bank of England purchases. Vertical dotted line represents the March 19th, 2020 Bank of England emergency COVID-19 intervention. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

7.2) Frictions over event period

To get a sense of the time-series trends of our OTC market frictions proxies we plot them in Figures 5-7. In general, we find that most of the proxies exhibit substantial variation over time and they materially worsen over the event date suggesting that OTC market frictions worsen in times of market stress - in line with theoretical predictions and past empirical results.

First, we trace the evolution of *market volume* that captures how much order flow dealers have to intermediate. The metric is calculated as the total trade volume (in £billions) in the D2C. Figure 5 plots the average weekly *market volume* calculated on a daily basis. *Market volume* increased during the flight to safety period and even more so during the dash for cash. On the 19th of March, the day Bank of England intervened, the measure peaked at £12 billion gilts traded by clients[30] during a single day. Our second, related measure, d*ealer volume* (Figure 5) captures the extent to which individual dealer balance sheets were pressured. During the week leading up to March 19th dealers, on average, intermediated £0.7 billion client volume per day. However, the top 10% of dealers, intermediated double the volume (£1.5 billion). Confirmed by the network analysis in section 7.3, the data suggests that some dealers intermediate substantially more client order flow in times of stress than others. Greater *dealer volume* increases the cost of intermediation since dealers struggle to match incoming trades, especially when order flow is increasingly one-sided as we see over the event period (Figure 4), which results in dealers increasing their inventory position. To manage inventory risk dealers could react by exerting pressure on prices – we will discuss this in section 8.

30 As mentioned in the data description section - all gilt market volumes, unless otherwise reported, only include secondary market activity and non-Bank of England trades i.e., client activity only.

Figure 5: Inventory frictions over time

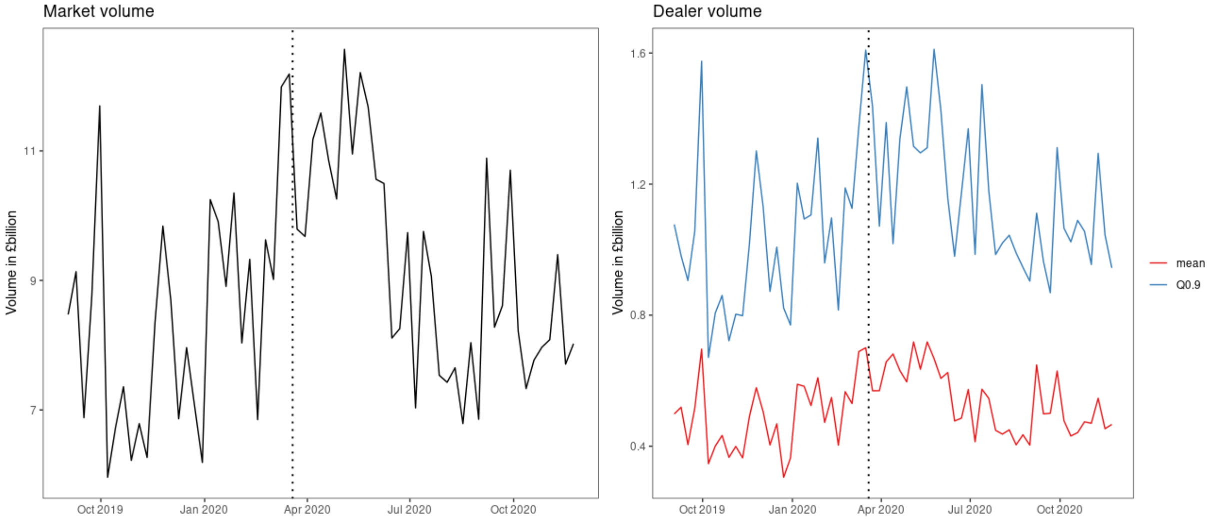

Evolution of inventory frictions over the sample period. For the purpose of this graph, we plot the weekly average of both measures calculated on a daily basis. Market volume is calculated as the sum of all D2C trading volume across dealers and is reported in £ billions. Dealer volume is calculated as the sum of all D2C trading volume per dealer/day and is reported in £ billions. Red line represents the mean of the distributions and the blue line the top 10% dealer/day observations. Dotted line represents the Bank of England's intervention on the 19th of March 2020. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

Search and bargaining frictions are seen as major impediments to efficient trade in OTC markets. Not only because participants need to search for eligible counterparties, which might be costly if they are pressured to trade, but the fragmentation of trading and lack of anonymity in this market could lead to substantial bargaining power for those who provide liquidity. We proxy search frictions by assessing the ease with which dealers can offload their inventory risk in the D2D segment. By extension if dealers are able offload their inventories more easily in the D2D segment they could provide liquidity more readily to clients in the D2C which could lead to reduced search for clients in the D2C. Figures 6 plots two proxies aimed at capturing search frictions. Our first proxy is *D2D market share* (Figure 6) which captures D2D activity by calculating the ratio of D2D trade volumes over total trade volumes. We see a sudden drop in *D2D market share* over the event period to a low of 21%, whereafter it slowly recovers to pre-event levels over the following 6 months. This indicates that the D2D became relatively less active which might affect the provision of liquidity to clients in the D2C. Next, we assess how well-connected dealers are in the D2D segment i.e., *dealer connections* (Figure 6). This measure calculates the number of dealers each dealer trades with on a given day. As in Di Maggio et al., (2016) and Friewald and Nagler (2019), we assume that if dealers in the D2D are less connected, dealers find it harder or more costly to establish a connection. We see a sharp decrease in *dealer connections* over the event period with the average dealer connected to 4 fewer dealers 2 weeks after March 19th relative to 2 weeks before. Notably this is not the case for the largest 10% of dealers who remain connected to all other dealers over the event period. Collectively, these measures suggest that search frictions increased over the event date which might affect prices in the gilt market.

Figure 6: Search frictions over time

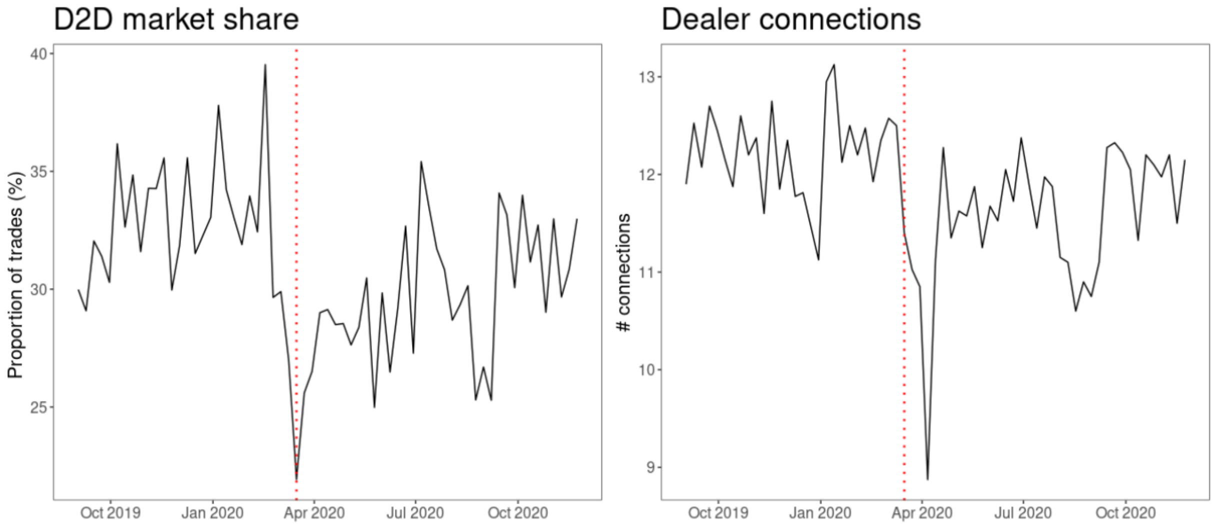

Evolution of search frictions over the sample period. For the purpose of this graph, we plot the weekly average of both measures calculated on a daily basis. D2D market share is calculated as the ratio of D2D trade count divided by total trade count. Dealer connections is calculated as the number of actual dealer links in the D2D segment on a given day. Dotted line represents the Bank of England's intervention on the 19th of March, 2020. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

Lastly, we look at client bargaining frictions over the event period in Figure 7. Bilateral trading, where counterparty identities are known, usually involves bargaining over terms of trade. In times of stress, pressure on dealers' balance sheets increases (seen in Figure 5) and dealers refrain from risk sharing among themselves (seen in Figure 6). The increase in these frictions when markets are stressed implies that it becomes harder for clients to find willing dealers to provide liquidity. This leads to a deterioration in clients' bargaining power since dealer activity become more concentrated and clients no longer have alternatives (many dealers) to trade with (Duffie, 2007). First, we assess the concentration of dealer activity. To assess how dealer concentration in the D2C evolves over time, we calculate the Herfindahl-Hirschman index (*HHI*). The *HHI* is a standard proxy for market concentration used to measure bargaining power (Bicu et al., 2017; Friewald and Nagler, 2019). It is calculated by aggregating each dealers' market share based on their daily D2C traded volume. *HHI* takes on a value closer to 1 if the market becomes more concentrated. Trading activity concentrates in the hands of fewer dealers over the event period suggesting that dealers have greater bargaining power in stressed periods (Figure 7). In general, dealers are not heavily concentrated in the D2C market with *HHI* levels below 0.13. In relative terms however, dealer concentration in the D2C (*HHI D2C*) increased by more than 40% over the event period. Notably, the peak came two weeks after the Bank of England's intervention. Next, we asses clients' alternatives to trade by calculating *outside options i.e., t*he number of distinct dealers a client trades with on a daily basis. *Outside options* (Figure 7), decreases over the event period, indicating that clients have fewer alternatives to trade when markets are stressed. However, it should be accounted for that the average number of dealers a client trades with on a daily basis is very low over the sample we use. This is because 50% of clients trade with only 1 dealer within a day – more over this in the network analysis section.

Figure 7: Bargaining frictions over time

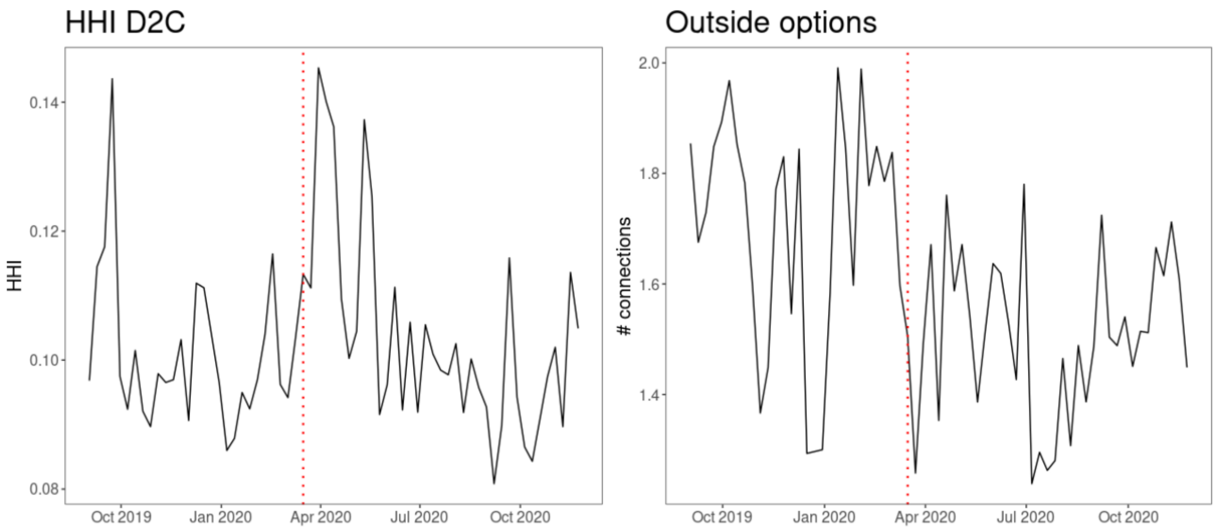

Evolution of bargaining frictions over the sample period. For the purpose of this graph, we plot the weekly average of both measures calculated on a daily basis. HHI D2C is calculated as the squared sum of each dealer's share of total market volume in the D2C on a daily basis. Outside options is calculated as the daily number of dealers a client trades with in the D2C. Dotted line represents the Bank of England's intervention on the 19th of March, 2020. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

7.3) Network analysis

To better understand the gilt market structure, how participants interact and how the market evolved in the stress episode, we conduct a network analysis.

Graphs 1 and 2 are network representations of both D2D (Panel A)[31] and D2C (Panel B) segments. For illustration purposes, we plot a one-day snapshot, however the market structure remains fairly stable across other non-event days and is therefore a fair representation of how the market is structured in normal times. Graph 1 plots the gilt market on a random non-event day of January 10th, 2020. An edge in the network represents a connection between two participants i.e., they have traded with each other. Nodes depict the type of participant distinguished by colour in each segment. The size of the node represents the number of connections per participant also referred to as the participant's network (degree) centrality (Hasbrouck & Levich, 2021). Graph 1 illustrates that the D2D and D2C segments are structured very differently. Panel A shows that trading in the D2D segment is all-to-all and not concentrated. This means that an individual dealer can potentially trade with all other dealers and that there is no single dealer that dominates trading volume. On the other hand, trading in the D2C segment (Panel B) has a core-periphery structure. This means that most trading is done by a few participants, mostly dealers and large institutional clients (the core) and the rest of the trading is done by many smaller clients (the periphery) who are less connected.

We already know that dealers are well connected to each other in the D2D with an average daily *dealer connection* of around 11.71 (see Table 2) – confirming the graphical evidence in Graph 1 Panel A. Due to the D2C's core-periphery structure where dealers intermediate almost all trades, most dealers are well connected to clients, but most clients are only connected to a limited number of dealers. Heterogeneity in dealer-client connections is illustrated by the different sizes of nodes. Panel B shows that although most dealers are well connected there is some heterogeneity in dealer-client connections. For example, over

31 For graphical purposes we only plot non-retail GEMMs. Retail GEMMs do not contribute to risk sharing in D2D

the entire sample period most dealers are connected to around 400 clients. The top 3 dealers, however, each trade with approximately 650 clients and the bottom 2 trade around 50 clients each. This suggests that although most dealers provide liquidity to a similar number of clients (~400), there are 3 players in the market that intermediate substantially more clients. The bottom 2 dealers are retail GEMMs.

There is substantial heterogeneity among clients in the D2C. Panel B shows that there are a small number of clients who are connected to many dealers (red dots) whereas most other clients are connected to only 1 dealer on a given day. This heterogeneity is also observed when comparing *dealer connections* per client over the entire sample period. Figure 8 plots the density and cumulative distribution of *dealer connections* per client over the entire sample period. The density plots illustrate just how skewed this distribution is, with approximately 50% of clients trading with less than 4 dealers over the entire sample period. The average client is connected to about 5 dealers and there is substantial variation in dealer-client connections with a difference between the 90th and 10th percentile of 13. The positively skewed distribution is even more pronounced when looking at trade volumes - where 5% of clients are connected to 15 or more dealers and they account for 42% of total D2C trade volume. This indicates that most gilt market participants trade low volumes with only a few dealers (think small asset managers and individual investors) and that only a few participants (institutional asset managers, insurance companies and hedge funds) dominate the market in terms of volume and they trade with most dealers. We cover the individual participant behaviour in the next subsection.

Graph 1: Network graph for normal period

Panel A: D2D segment Panel B: D2C segment

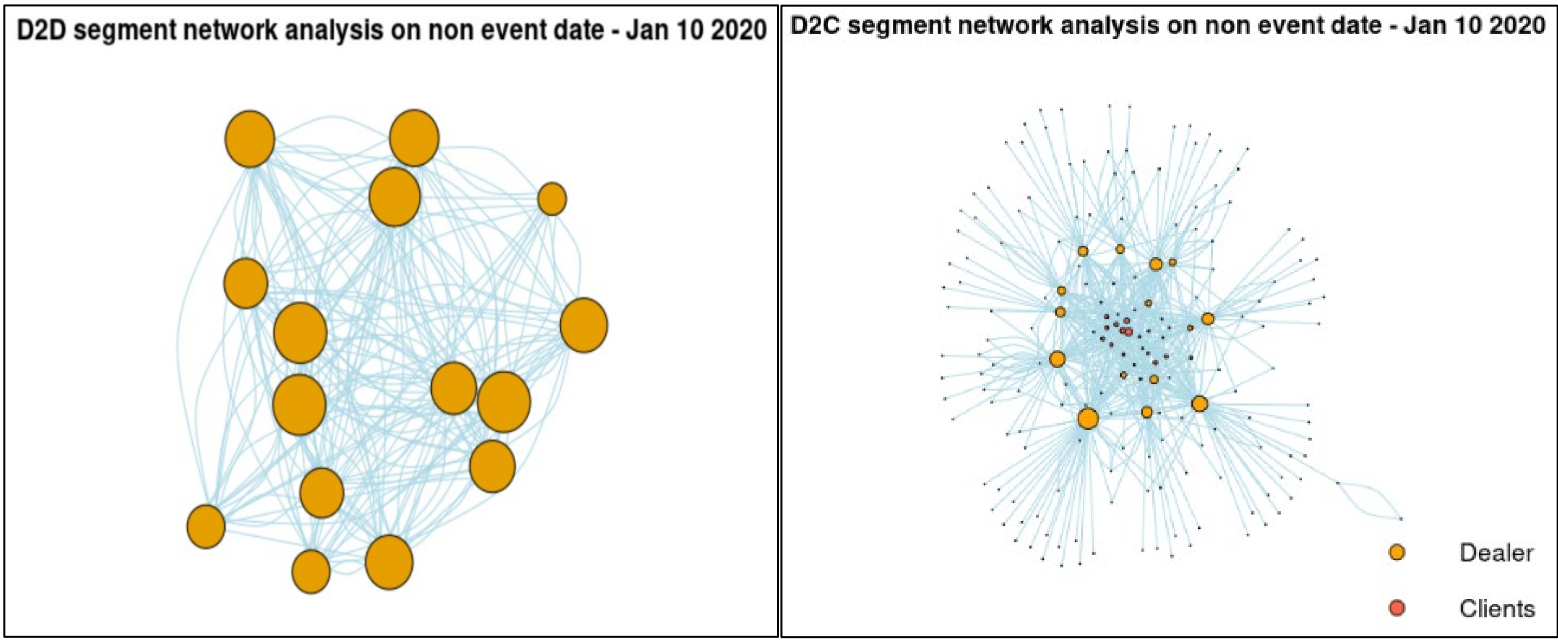

Network graph of the gilt market segments. An edge in the network represents a connection between two participants i.e., they have executed a trade with each other. Nodes depict participants and type of participants are distinguished by colour in each segment. The size of the node represents the number of connections per participant – larger nodes have executed trades with more counterparties. To make the graph legible we plot the network for a random non-event day of 10 January 2020.

Figure 8: Distribution of dealer connections per client

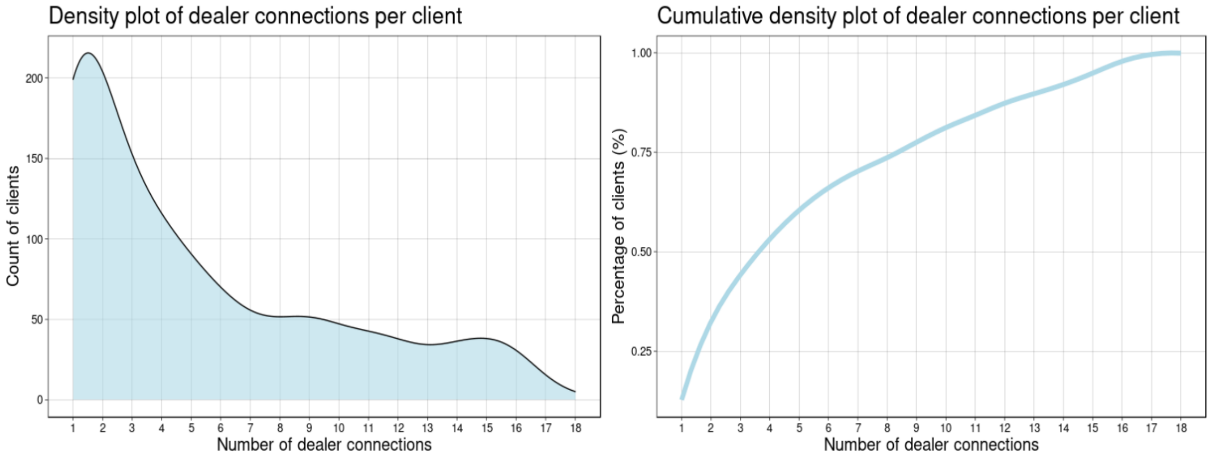

We plot the density and cumulative density of dealer connections per client. Number of dealer connections counts the number of distinct dealer connections per client over the entire sample period. The sample period ranges from 2 September 2019 to 30 November 2020.

Finally, we discuss how the market evolved over the stress period. Graph 2 plots a one-day snapshot of how the gilt market changed when it came under stress over the event period on the day the Bank of England started its intervention. From Panel A we see some graphical evidence to suggest that the network structure of the D2D becomes more concentrated relative to normal times. Although all 16 dealers still participate in the market, some dealers become less connected to other dealers than they would when markets are not under stress (see Graph 1 Panel A). This was also confirmed in Figure 6 where we see *dealer connections and D2D market share,* separate but related measures, both decrease over the event period. Taken together, these findings suggest that dealers become more concentrated and overall D2D market activity decreases when markets are stressed. This could imply that dealers find it harder to socialise their inventory risk. The change in dealer concentration over the event date seems more pronounced in the D2C (Graph 2 Panel B). We see a clear increase in the size of some nodes, indicating that some dealers provide liquidity to more clients in times of stress. Moreover, Figure 7 shows that fewer dealers intermediate a greater proportion of the total trade volume over the event period i.e., *HHI D2C* increased.

Graph 2: Network graph in stressed period

Panel A: D2D segment | Panel B: D2C segment

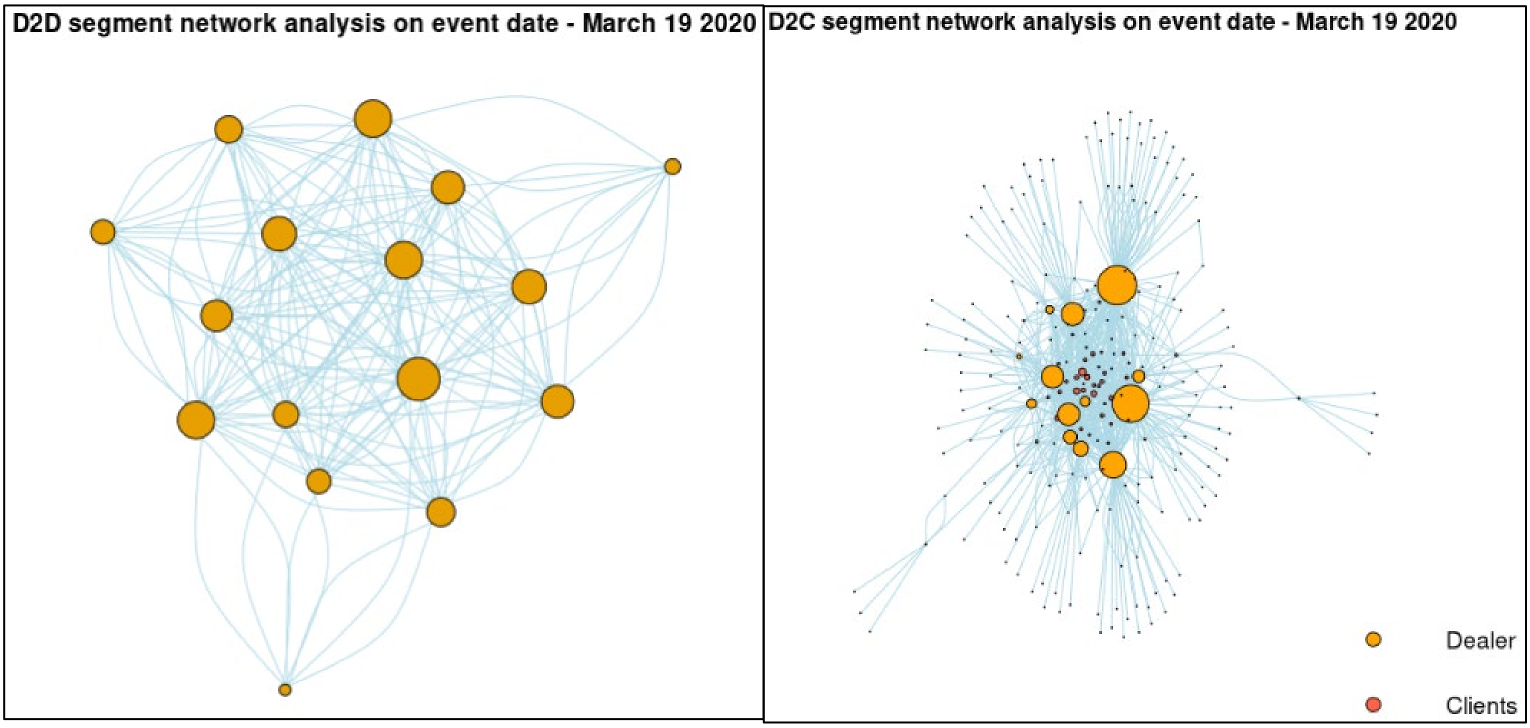

Network graph of the gilt market segments. An edge in the network represents a connection between two participants i.e., they have traded with each other. Nodes depict the type of participant distinguished by colour in each segment. The size of the node represents the number of connections per participant. To make the graph legible we plot the network only for the day the event peaked – the day the Bank of England intervened - 19 March 2020.

7.4) Client behaviour over event period:

By having access to non-anonymised counterparty identities, our dataset uniquely allows us to breakdown each client's behaviour to see how they reacted to the stress event. Using our own internal records and a best endeavour approach, we classify counterparties into client groups. To avoid subjective decisions regarding the allocation of pension funds and insurers who are also involved in asset management, we take a more conservative approach and group them all as asset managers. The 4 client groups we use are: asset managers (AM), hedge funds (HF), official institutions (e.g. local authorities, foreign central banks, sovereign wealth funds and international organisations like the IMF) and other (e.g. mainly commercial banks, retail brokers and retail participants). These four groups serve different clienteles with different objectives and face different regulations. For example, AMs typically have a longer investment horizon and have stricter fiduciary responsibilities (best execution) than HFs. Official institutions are motivated by financial stability concerns unlike commercial banks and individual investors (Other category) who invest for their own benefit. It is therefore interesting to examine each group's behaviour in the gilt market according to their own motivation.

AMs account for 50% of total trade volume and 80% of all transactions indicating that some AMs also trade smaller tickets – typically smaller private wealth managers and for portfolio rebalancing purposes. Conversely, HFs account for 34% of volume but only 11% of trades in the market. This suggests that HFs typically trade larger transactions compared to the average AM. Similarly, official institutions contribute 8% of total trade volume but only account for 2% of trades. Other clients account for 8% of trade volume and trades. These proportions remain relatively stable over time. Furthermore, as is the case across the entire sample, constituents within each client type are also heavily skewed. Across all groups the top 25% of clients account for nearly 95% of trade volume in each group.

From Figure 9 Panel A, we see that all client groups increased their monthly trading activity over the event period and the increase was most pronounced for AMs. AMs increased their monthly trading activity by around 65% (from £85bln in February to ~£141bln in March). From Panel B we also see that AMs were the biggest net sellers over the event period and therefore mainly responsible for the spike in client selling pressure we observed in Figure 4, previously. The Other client category remains neutral in net volume terms over the event period. Interestingly, the HFs and Official institutions in fact remained net providers of liquidity (net buyers) over the same period. Contrary to findings based on the US market (BIS, 2020 and FSB, 2021), UK hedge funds in fact provided much needed liquidity to dealers during and for the months after the event date by becoming net buyers of gilts over the event period until the end of the sample period.

Figure 9: Client trading activity over sample period

Panel A: Total trade volume per client group

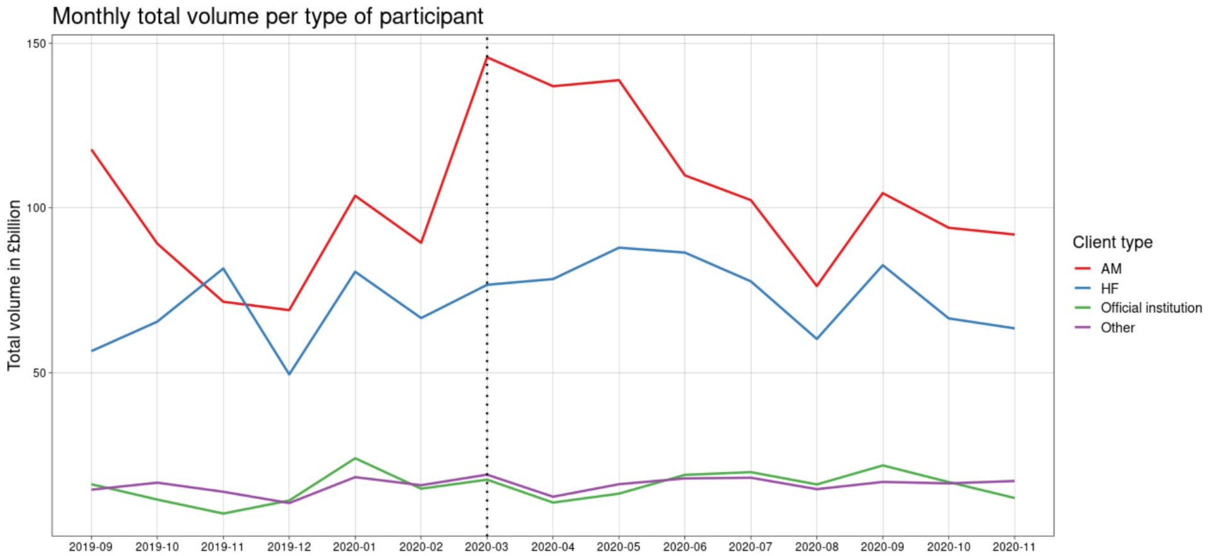

Panel B: Net trading volume per client group

Times series of client trading behaviour. Panel A plots the monthly total trade volume per client group in £ billions. Panel B plots the monthly net volume per client group in £ billions. A positive net volume value indicates that the client was a net buyer (bought more gilts over the period than it sold). Dotted line represents the Bank of England's intervention on the 19th of March, 2020. The sample period ranges from September 2nd, 2019 to November 30th, 2020.

8) What drove market liquidity to deteriorate

In this section, we aim to determine whether there is a statistical relationship between market liquidity and our market friction proxies. Market liquidity is the degree to which a security can be traded at a price close to its fundamental value (Foucault, et al., 2013). For equities, market liquidity is measured using observed bid-ask spreads. In the absence of reliable bid-ask spreads to measure market liquidity in the gilt market (and other OTC markets), we construct a measure of liquidity directly from our transaction data. Following Hendershott and Madhavan (2015) and Aquillina et al., (2015), we calculate transaction cost for each transaction k as:

$$Transaction\ cost_k = \ln\left(\frac{TradePrice_k}{BenchmarkPrice_k}\right) * 2\mathrm{D}$$

where $TradePrice_k$ is the transaction price for trade k. $BenchmarkPrice_k$ is the transaction price of the nearest D2D trade for that bond. D is equal to +1 (-1) when a client buys from (sells to) a dealer. Similar to Hendershot and Madhavan (2015), we use the nearest D2D price as benchmark price for the assets fundamental value, although, our results are not affected by our choice of benchmark price[32]. The D2D price is widely used as reference price in OTC market studies since it resembles a frictionless market (Duffie et al., 2005). For the purpose of this study, transaction cost can be interpreted as the deviation in trade price from fundamental value in the presence of OTC market frictions in the D2C segment. The less liquid a bond or market is, the more likely it is to observe a significant deviation between trade price and reference price (Jankowitsch, et al., 2011).

To assess how market liquidity evolved over time, we plot the average daily transaction cost by computing the simple average across all trades each day. Figure 10 indicates that transaction cost is mean reverting to a level around 1.2 BPS, equal to the average transaction cost reported in Table 2, which reflects the compensation dealers charge for their intermediary services. This means that in normal times prices in the D2C segment do not consistently differ much from prices in the D2D segment, implying that clients face minimal trade frictions since dealers can effectively socialise their risk in the D2D. This was not the case, when the market came under stress over the event period. Transaction cost spiked to 30 BPS on March 19th and remained elevated in the weeks before and after the Bank of England's intervention. To contextualise this number: this is nearly twice the size of the spike in bid-ask spreads in the UK equities market on the same day (Mittendorf, et al., 2021) and comparable to the deterioration in the U.S. Treasury market during COVID-19 (He, et al., 2022). This suggests that dealers find it more costly to intermediate in times of stress.

32 We also use the end of day Tradeweb closing price and the daily average price for the bond as robustness checks. We find that our results do not change when we use alternative reference prices. Results are available on request.

Figure 10: Time series of transaction cost

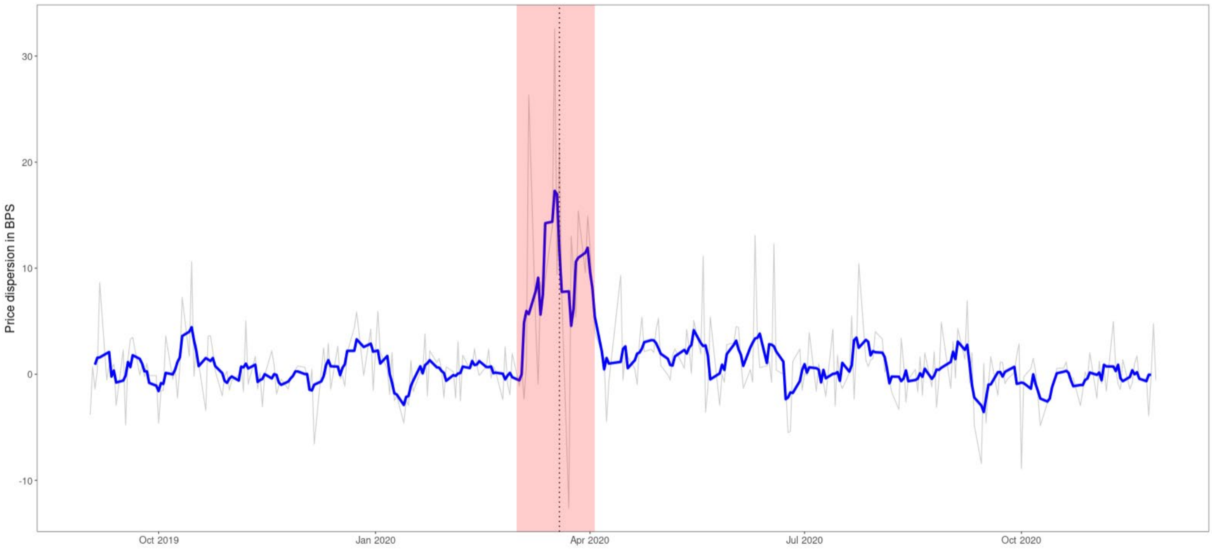

Evolution of daily average transaction cost over the sample period. Average daily transaction cost is calculated by taking the equal-weighted average transaction cost across all transactions within a day. The raw transaction cost is plotted in grey. The blue line represents a 5-day (business week) moving average transaction cost. Dotted line represents the Bank of England's intervention on the 19th of March, 2020. Red shaded ranges from the onset of the flight-to-safety period on the first of March and ends on the last day of March. The sample period ranges from September 2nd, 2019 to November 30th, 2020

8.1) Liquidity and frictions

We closely follow Friewald and Nagler (2019) to determine whether there is a statistical relationship between transaction cost (a proxy for market liquidity) and our measures of OTC market frictions. For this purpose, we run separate regression models for each OTC market friction and two multivariate models that include all market-wide and participant-level frictions separately, using the following model specification:

$$Transaction\ cost_k = a + \beta_1.Friction_{t(k)} + \beta_2.X_k + Event_t + FE_i + FE_d + \varepsilon_k \qquad (1)$$

Where *Transaction cost* is calculated per transaction k. $Friction_{t(k)}$ represents market wide frictions - *market volume, D2D market share and HHI D2C* - which are calculated on a daily frequency across all dealers and participant level frictions - *dealer volume, dealer connections and outside options* - which are calculated per day per dealer d or client c (for detailed calculation see section 6). The subscript $t(k)$ indicates the day t when transaction k takes place. X_k represents a set of time-unvarying transaction-level controls[33] and controls for client heterogeneity[34]. FE_i, FE_d represent bond (ISIN) fixed effects and dealer fixed effects. This way we compare prices that the same dealer charges for the same bond. Standard errors are clustered at ISIN, day and dealer levels. $Event_t$ is a dummy that takes the value of one if trade execution occurs between the start of the flight-to-safety period (March 1st) and the day the Bank of England intervenes (March 19th). To control for confounding effects not controlled for in the single-variable models (1-6) we perform a

33 These include log trade size, a dummy that is equal to one if a dealer is the buyer in the transaction, time to maturity dummy (bucketed as in section 5), benchmark status dummy and venue traded dummy.

34 These include log of total trade volume per client and a client category dummy using the same categorisation as in section 7.4.

multivariate regression analysis combining all market-level frictions in model 7 and participant-level frictions in model 8.

Table 3: Regression results

	Dependent variable: Transaction cost							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Market volume	0.099***						0.009	
	(0.020)						(0.020)	
Dealer volume		0.422***						0.578***
		(0.112)						(0.144)
D2D Market share			-0.184***				-0.176***	
			(0.015)				(0.015)	
Dealer connections				-0.120**				-0.073
				(0.053)				(0.078)
HHI D2C					17.176***		13.722***	
					(3.332)		(3.263)	
Outside options						-0.202***		-0.204***
						(0.065)		(0.065)
Log(trade size)	0.333***	0.325***	0.356***	0.345***	0.337***	0.632***	0.351***	0.599***
	(0.038)	(0.038)	(0.038)	(0.038)	(0.037)	(0.061)	(0.038)	(0.061)
Log(client size)	-0.473***	-0.479***	-0.470***	-0.464***	-0.472***	-1.296***	-0.471***	-1.310***
	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)	(0.091)	(0.038)	(0.092)
Official institution dummy	-4.023***	-3.999***	-4.056***	-4.038***	-3.997***	-6.358***	-4.041***	-6.437***
	(0.518)	(0.518)	(0.518)	(0.519)	(0.518)	(1.513)	(0.518)	(1.515)
HF dummy	-3.951***	-3.932***	-3.780***	-3.970***	-3.967***	-7.776***	-3.808***	-7.906***
	(0.337)	(0.337)	(0.336)	(0.337)	(0.337)	(0.763)	(0.337)	(0.763)
AM dummy	-2.733***	-2.769***	-2.673***	-2.715***	-2.734***	-4.439***	-2.692***	-4.571***
	(0.276)	(0.275)	(0.276)	(0.276)	(0.276)	(0.498)	(0.275)	(0.497)
Event dummy	8.218***	8.329***	7.691***	8.434***	8.422***	10.930***	7.765***	11.026***
	(0.722)	(0.723)	(0.713)	(0.731)	(0.723)	(1.076)	(0.713)	(1.085)
Dealer BUY dummy	4.401***	4.419***	4.422***	4.437***	4.403***	4.987***	4.400***	4.952***
	(0.167)	(0.167)	(0.167)	(0.168)	(0.167)	(0.262)	(0.167)	(0.263)
Maturity 3-5	1.320***	1.191***	1.510***	1.321***	1.190***	2.334***	1.528***	2.340***
	(0.155)	(0.151)	(0.156)	(0.148)	(0.151)	(0.259)	(0.159)	(0.262)
Maturity 6-10	2.449***	2.205***	2.772***	2.219***	2.136***	4.021***	2.777***	3.952***
	(0.289)	(0.280)	(0.289)	(0.279)	(0.279)	(0.442)	(0.294)	(0.447)
Maturity 11-20	2.477***	2.126***	2.920***	2.032***	1.972***	4.863***	2.952***	4.999***
	(0.564)	(0.554)	(0.563)	(0.554)	(0.553)	(0.765)	(0.569)	(0.774)
Maturity 21-50	1.589	1.137	2.230*	1.032	0.943	4.856***	2.263*	5.085***
	(1.176)	(1.169)	(1.176)	(1.170)	(1.168)	(1.443)	(1.181)	(1.449)
Benchmark dummy	2.978***	3.079***	2.794***	3.190***	3.145***	1.013	2.768***	0.864
	(0.827)	(0.826)	(0.827)	(0.827)	(0.826)	(1.047)	(0.827)	(1.049)
SI dummy	0.332	0.285	0.293	0.455	0.348	-2.867***	0.278	-3.111***
	(0.325)	(0.321)	(0.325)	(0.327)	(0.325)	(0.699)	(0.324)	(0.684)
Off exchange dummy	4.484***	4.357***	4.501***	4.491***	4.492***	5.450***	4.457***	5.041***
	(0.332)	(0.334)	(0.332)	(0.336)	(0.332)	(0.591)	(0.332)	(0.599)
Dealer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ISIN FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	375857	375857	375857	375857	375857	375857	375857	375857
Note:								*p<0.1; **p<0.05; ***p<0.01

Regression results using OLS. Transaction cost is measured at the transaction level. Independent variables are measured on a daily frequency. Controls include log of individual trade size, log of total client volume traded across the entire sample period, a dummy variable of the client category with reference category Other, a dummy for when a dealer is buying in the transaction, a dummy for which maturity bucket the ISIN is in with reference category 0-3 years, a benchmark status dummy with reference category non-benchmarks and a venue dummy with reference category MTF trades. Event dummy refers to the period 1st of March until the Bank of England's intervention on the 19th of March. Numbers between parentheses reflect robust standard errors clustered at day, ISIN, dealer level. Dealer fixed effects account for difference in size of individual dealers. Bond fixed effects account for differences across ISIN. The sample period ranges from September 2nd, 2019 to November 30th, 2020 and included trades only if the nearest D2D price is less than 3 hours away – which captures 92% of traded volume.

Across all models in Table 3, we observe that the OTC friction coefficients are in line with theory predictions. First, we assess the relationship between inventory frictions and liquidity. Both measures have their predicted signs and are statistically significant. That is, an increase in both *market volume* and *dealer volume* increases transaction costs. Model 1 implies that liquidity deteriorates when markets become more active, and dealers need to intermediate larger client order flow. That is a one standard deviation increase of *market volume* leads to an increase in transaction cost of 0.4 BPS (3.95 x 0.099). Compared to the mean daily transaction cost in Table 2 of 1.18BPS, this is a 34% increase in transaction cost. From the start of the flight to safety period (March 1st) till the end of the dash for cash (the day the Bank of England intervened – 19th March) henceforth referred to as the 'event period', the average daily *market volume* rose by nearly 1.5 standard deviations from approximately £7 billion to £12.5 billion. In response to the increase in *market volume,* transaction cost increased by 0.6 BPS (1.5 x 0.4) over the event period. Therefore, the impact of *market volume* is economically significant. Model 2 suggests that individual dealers find it more costly to intermediate increased client order flow. A one standard deviation increase in *dealer volume* results in a 0.47 BPS (0.64 x 0.422) or a 28% increase in transaction cost. The coefficient of *dealer volumes* become even larger when we control for confounding effects of the other frictions in model 8. Thus, inventory frictions significantly impact transaction cost even after controlling for other frictions. In both models 1 and 2, the underlying mechanism is such that dealers charge a higher intermediation premium to clients when they are faced with higher inventory costs due to constrained inventories. This finding suggests that inventory frictions impact liquidity and that the deterioration in liquidity over the event date can be partly ascribed to the increase in inventory frictions.

Next, we analyse the impact of our search proxies on liquidity. Model 3 and 4 indicate that there is a significant inverse relationship between transaction cost, *dealer connections and D2D market share.* That is, transaction cost increases (liquidity deteriorates) when dealers are not well connected, and their market activity reduces relative to overall market activity. However, only *D2D market share* remains statistically significant in the mixed regressions in model 7. A possible reason why we do not find a statistically significant relationship between *dealer connections* and liquidity is because dealers are generally well connected in the D2D segment which may be explained by exclusive dealer participation in D2D and limited number of GEMMs. We also notice that *dealer connections* only decrease 2 weeks after the Bank of England's intervention suggesting that dealers did not reduce their *dealer connections* over the event period. Model 3 shows that a one standard deviation decrease in *D2D market share* among dealers leads to 1 BPS increase in transaction cost (5.62 x -0.186). Even after controlling for other market frictions in model 7, *D2D market share* remains statistically significant and economically meaningful. When dealers trade 1.5 standard deviation less in the D2D (as they did during the event period) relative to overall trade volume, transaction cost doubled relative to the average transaction cost over the sample period. This result is consistent with models of dealer competition like Li and Schurhoff (2019) where greater D2D activity leads to better execution. This result suggests that when search frictions increase, as was the case during the COVID-19 event, liquidity deteriorates for clients in the D2C.

To examine the impact of bargaining frictions on transaction cost, we rerun our baseline regression with proxies for dealer bargaining power and report the results in models 5 and 6. Both proxies exhibit their predicted sign, and both are highly statistically significant. Model 5 implies that dealer concentration in the D2C (*HHI* D2C) significantly impacts transaction cost. This relationship remains significant even after controlling for other frictions in model 7 and predicts that a one standard deviation increase, experienced over the event period, leads to a nearly 30% increase in transaction cost. When clients have less *outside options,* dealers have more bargaining power and spreads deteriorate (Duffie, 2005). Over the event period, clients on average become 1 standard deviation less connected to dealers, transaction cost across all clients increase by 0.25 BPS (1.23 x -0.204). This result holds after controlling for other frictions in model 8. That means that over the event period transaction cost increased by 20% in response to the decrease in client o*utside options.* Collectively, these results imply that bargaining power impacts liquidity.

These results hold true even after controlling for increase in transaction cost over the event period. Across all models the *event dummy* coefficient is statistically significant around 8-12 BPS. That is, relative to the rest of the sample period transaction cost was 8-11 BPS higher over the event period, on average. Overall, we find that frictions in the gilt market partly explain the deterioration in liquidity over the event period as predicted, after controlling for the increase in transaction cost. An increase in inventory, search and bargaining frictions increase the cost of both dealers and clients to trade in this market. This is represented by the significant increase in transaction cost when these frictions spiked over the event period.

8.2) Liquidity and transaction level characteristics

In addition to controlling for transaction level characteristics in the baseline regression model in section 8.1, we discuss here how transaction level characteristics relate to market liquidity. Figure 11 plots the daily time series of average transaction cost for a brief period before, during and after the event period for different bond characteristic. Benchmark bonds, a designation given to the on-the-run bond per standard maturity, seem to be less volatile than non-benchmarks over the event period. This could be explained by the fact that non-benchmark bonds have higher liquidity risk since they are seldom traded. This makes them more sensitive in periods of stress. However, Table 3 shows they have a higher average transaction cost of around 2-3 BPS on average. Although this result may seem counterintuitive since benchmarks should be more liquid, we should note that the average trade size for benchmarks is double that of non-benchmarks and a larger proportion of smaller clients trade benchmarks.

From Table 3, we see that the coefficient of *log (trade size)* is statistically significant and positive across all models with similar magnitudes to Jurkatis et al., 2022. Thus, larger trades result in a larger transaction cost. Trades concerning small transaction sizes also seem to be more sensitive to periods of stress (figure 11). A possible explanation for this could be that smaller transactions are disproportionately traded by less active clients who may not enjoy the benefits of regular clients who bring repeat business (Hendershott et al., 2020).

Similar to the literature on client-dealer relationships (most notably Jurkatis et al., 2022 who look at the UK corporate bond market) we also find that larger clients receive prices that are closer to the fundamental value of the asset traded i.e., lower transaction cost. The result is statistically significant and robust across all model specifications in Table 3. We also regress client category dummies to see if some client groups receive different prices on average. Relative to the Other client category, which mainly consist of commercial banks, retail brokers and retail participants, Official institutions get the best prices, HFs get marginally worse prices followed by AMs. Many factors could motivate why Official institutions, HFs and AMs get better prices than the other category but the most plausible based on the distribution of client sizes per category seems to be that Other category clients are on average smaller in total trade volume and we therefore capture a similar effect to that of *log(client size).*

Figure 11 shows that bonds with a longer time to maturity have a higher transaction cost in periods of stress. This is because longer maturity bonds are more sensitive to changes in interest rates due to their higher duration risk. Similarly, the maturity bucket coefficients in Table 3 exhibit the same trend - longer maturity gilts have a higher average transaction cost relative to shorter maturity gilts. Finally, transaction cost seems to vary across trade venue types with the multilateral venues (MTFs) having lower levels of average transaction cost than bilateral venues (SI and especially XOFF) (Hendershot and Madhavan, 2015). While dealer competition might promote better prices in some cases, there are also arguments that direct one-to-one interaction with SI can bring the benefit of lower information leakage. Further analysis is needed to clearly determine the effect of trading venue type on transaction cost and resiliency in times of stress.

Figure 11: Transaction level characteristics and transaction cost

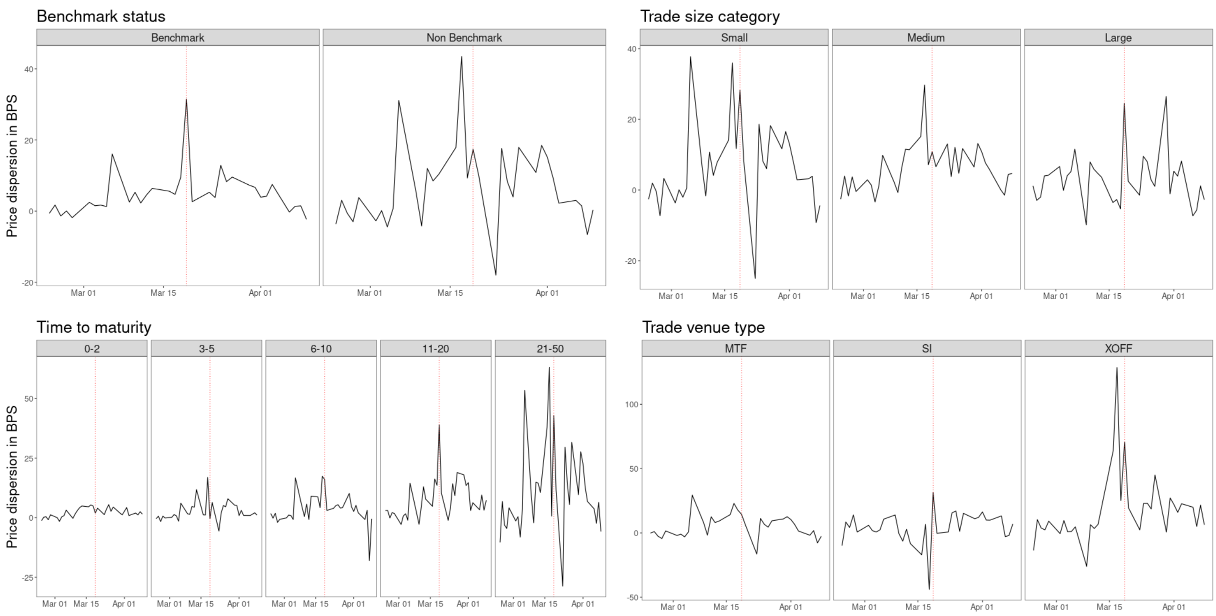

Transaction cost over event period broken down per transaction-level characteristic. Average daily transaction cost is calculated by first taking the equal-weighted average transaction cost across all transactions within a day for a given bond ISIN. Then we average this equally across all bonds with the same bond characteristic to attain a transaction characteristic-wide average daily transaction cost. Dotted line represents the Bank of England's intervention on the 19th of March, 2020. The sample period ranges from February 22 nd, 2020 to April 12th, 2020.

9) Conclusion

Due to the lack of a representative[35] consolidated feed of data that is publicly available to market participants, statistics describing the entire gilt market are hard to come by. Our unique regulatory dataset, which captures almost all trading in gilts, provide useful statistics that describe the gilt market. We document that gilt market trading is far less frequent than equities trading with an average number of daily trades of around 500 in the D2D and 1500 in the D2C segments of the market. The average size of an individual trade is however far larger than in equities with average trade size being £9.4 million in the D2D and £6.7 million in the D2C. D2C trading constitutes nearly 60% of overall market volume over our sample period. Within the D2C segment, trading is split roughly equally between multi-dealer platforms (MTFs) and single-dealer platforms (SI and XOFF) that respectively constitute 51%, 37% and 12% of total trade volume. In the D2C, more than 70% of small to medium sized trades (below £50 million) are traded on MTFs, whereas over 60% of large trades go through SI. Large trades are also more likely to execute on benchmark bonds due to their lower liquidity risk and higher amount outstanding. 6-10-year maturity gilts command the highest proportion of total trade volume in both D2D (47%) and D2C (38%) segments.

Due to high levels of sovereign and corporate debt, the resilience of government bond markets is key. The March Covid-19 event and more recently the September 2022 fiscal event showed that sovereign bond markets around the world including the UK gilt market are not immune to stress events. We analyse how features of OTC market trading can impact prices and liquidity in sovereign debt markets, specifically the gilt market. Moreover, we look at frictions that could impact the provision of liquidity during times of stress and hence can affect the market resiliency.

Using proxies informed by the academic literature, we document the evolution of OTC market frictions in the gilt market. We find that dealers faced a significant increase in the number of trades they needed to intermediate both on a market wide level (*market volume*) and dealer level (*dealer volume*) over the event period. These measures suggest that dealers' inventories were put under severe pressure during the event. Additionally, risk sharing activity in the D2D segment decreased, with *D2D market share* of total gilt trading contracting to 21% during the event period from the average level of 31%. Lastly, we find that clients' bargaining power reduced over the event since dealers became more concentrated in the D2C (*HHI increased*) and each client traded with fewer dealers (*outside options decreased*).

In our empirical analysis we examine whether these OTC market frictions help explain the deterioration in liquidity seen over the event period. Our analysis shows that there is a statistically and economically meaningful relationship between the deterioration in liquidity and the increase in frictions over the event period. During the event period, frictions typically rose by over one standard deviation, leading to a reduction in liquidity (proxied by transaction cost) of approximately 0.25 to 1.5 basis points. This outcome varied for

35 Post trade data is available through FITRS, but the data is not consolidated across venues/APAs and not representative since nearly 60% of trade volume is indefinitely aggregated (ICMA & Propellant report, 2022)

each individual friction, resulting in a 20% to 80% decrease in liquidity relative to the average transaction cost across the entire sample. The analysis in the paper confirms that frictions present in OTC fixed income markets have bearing on the liquidity present in those markets with different frictions having a varying degree of impact. Our results also demonstrate that liquidity in stressed periods is significantly affected by transaction level characteristics like trade size and bond maturity bucket. In addition, the choice of trading venue affects the observed transaction costs.

In conclusion, this paper’s comprehensive examination of OTC fixed income markets underscores the importance of accounting for frictions and considering transaction-level characteristics when assessing market liquidity. The implications of our research extend beyond the specific context of the UK, fostering a broader understanding of fixed income market resilience in global settings.

References

Amihud, Y. & Mendelson, H., 1980. Dealership market: Market-making with inventory. *Journal of Financial Economics,* Volume 8, pp. 31-53.

Aquilina, M. et al., 2015. Transparency in the UK Bond Markets: An overview. *FCA Occasional Paper No.6.*

Bagehot, W., 1971. The Only Game in Town. *Financial Analyst Journal ,* p. Vol 27 issue 2.

Bao, J., O'Hara, M. & Zhou, A., 2018. The Volcker Rule and Market-Making in Times of Stress. *Journal of Financial Economics,* pp. 130, 95–113.

Benos, E. & F, Z., 2016. Liquidity Determinants in the UK Gilt Market. *Bank of England working paper,* p. 600.

Benos, E. & Zikes, F., 2016. Liquidity Determinants in the UK Gilt Market. *Bank of England Working Paper ,* p. No. 600.

Bessembinder, H., Spatt, C. & Venkataraman, K., 2020. A Survey of the Microstructure of Fixed-Income Markets. *Journal of Financial and Quantitative Analysis,* 55(1), pp. 1-45.

Biais, B. & Green, R., 2019. The Microstructure of the Bond Market in the 20th Century. *Review of Economic Dynamics,* pp. 33 , 250–271..

Bicu, A., Chen, L. & Elliott, D., 2017. The leverage ratio and liquidity in the gilt. *Staff Working Paper No. 690,* p. Bank of England .

Boyarchenko, N., Shachar, O. & Tobias, A., 2017. Dealer balance sheets and bond liquidity provision. *Journal of Monetary Economics,* pp. 89, 92-109.

Brunnermeier, M. K. & Pedersen, L. H., 2009. Market liquidity and funding liquidity. *The Review of Financial Studies,* pp. 22, 2201-2238.

Chordia, T., Roll, R. & Subrahmanyam, A., 2002. Order imbalance, liquidity, and market returns. *Journal of Financial Economics,* p. 111–130.

Comerton-Forde, C. et al., 2010. Time variation in liquidity: The role of market-maker inventories and revenues. *The Journal of Finance,* pp. 65, 295-331.

Czech, R., Huang, S. & Lou, D., 2020. Informed Trading in Government Bond Markets. *Bank of England Working Paper No. 871.*

Czech, R., Huang, S., Lou, D. & Wang, T., 2021. An unintended consequence of holding dollar assets. *Staff Working Paper No. 953, Bank of England.*

Di Maggio, M., Kermani, A. & and Song, Z., 2017. The Value of Trading Relationships in Turbulent Times. *Journal of Financial Economics,* pp. 124, 266–284.

Duffie, D., 2010. Presidential Address: Asset Price Dynamics. *The Journal of Finance,* p. Vol(4).

Duffie, D., 2020. Still the World's Safe Haven? -- Redesigning the U.S. Treasury Market After the COVID-19 Crisis. *Hutchins Center Working Paper,* p. 62.

Duffie, D. et al., 2023. Dealer Capacity and. *Federal Reserve Bank of New York Staff Reports,* p. https://doi.org/10.59576/sr.1070.

Duffie, D., Garleanu, N. & Pedersen, L. H., 2005. Over-the-Counter Markets. *Econometrica,* pp. 73, 1815–1847.

Duffie, D., Garleanu, N. & Pedersen, L. H., 2007. Valuation in Over-the-Counter Markets. *Review of Financial Studies ,* pp. 20, 1865–1900..

Edwards, A., Harris, L. E. & Piwowar, M. S., 2007. Corporate Bond Market Transaction Costs and Transparency. *The Journal of Finance,* pp. 3, 1421-1451 .

Feldhutter, P., 2012. The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures. *Review of Financial Studies,* pp. vol. 25, issue 4, 1155-1206.

Fleming, M. & Ruela, F., 2020. *Treasury Market Liquidity during the COVID-19 Crisis,* s.l.: Liberty Street Economics.

Foucault, T., Pagano, M. & Roell, A., 2013. *Market Liquidity: Theory, Evidence, and Policy.* s.l.:Oxford University Press.

Friewald, N. & Nagler, F., 2019. Over-the-Counter Market Frictions and Yield Spread Changes. *Journal of Finance.*

Garbade, K. & Silber, W., 1976. Price Dispersion in the Government Securities Market. *Journal of Political Economy,* pp. 721-740.

Green, R., Hollifield, B. & Schurhoff, N., 2007. Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies,* Volume 20, pp. 275-314.

Grossman, S. J. & Miller, M. H., 1988. Liquidity and Market Structure. *The Journal of Finance,* pp. 617-633.

Haddad, V., Moreira, A. & Muir, T., 2021. When Selling Becomes Viral: Disruptions in Debt Markets in the COVID-19 Crisis and the Fed's Response. *The Review of Financial Studies,* Volume 34, p. 5309–5351.

Harris, L. E. & Piwowar, M. S., 2006. Secondary Trading Costs in the Municipal Bond Market. *Journal of Finance,* pp. 61, 1361–1397.

Hasbrouck, J. & Levich, R., 2021. Network structure and pricing in the FX market. *Journal of Financial Economics,* Volume 141, pp. 705-729.

Hauser, A., 2020. *Seven Moments in Spring: Covid-19, financial markets and the.* London, Bloomberg: Bank of England.

Hendershott, T., Li, D. & Schurhoff, N., 2020. Relationship Trading in OTC Markets. *The Journal of Finance,* 75(2), pp. 683-734.

Hendershott, T. & Madhavan, A., 2015. Click or Call? Auction versus Search in the Over-the-Counter Market. *The Journal of Finance,* pp. 70, 419–447.

He, Z., Nagel, S. & Song, Z., 2022. Treasury inconvenience yields during the covid-19 crisis. *Journal of Financial Economics,* Volume 143, p. 57–79.

Ho, T. & Stoll, H., 1983. The Dynamics of Dealer Markets under Competition. *Journal of Finance,* 38(4), pp. 1053-1074.

Jankowitsch, R., Nashikkar, A. & Subrahmanyam, M., 2011. Price Dispersion in OTC Markets: A New Measure of Liquidity. *Journal of Banking and Finance,* 35(2), p. 343–357.

Joseph, A. & Vasios, M., 2022. OTC Microstructure in a period of stress: A Multi-layered network approach. *Journal of Banking and Finance,* pp. 106400,138.

Jurkatis, S., Schrimpf, A., Todorov, K. & Vause, N., 2022. Relationship Discounts in Corporate Bond Trading. *Working paper.*

Kargar, M. et al., 2021. Corporate Bond Liquidity during the COVID-19 Crisis. *The Review of Financial Studies,,* Volume 34, p. 5352–5401.

Li, D. & Schurhoff, N., 2019 (74). Dealer Networks. *Journal of Finance,* pp. 91-144.

Market, L. D. i. t. U. G., 2016. Benos, E.; Zikes, F.. *Bank of England Working Paper No. 600.*

Ma, Y., Xiao, K. & Zeng, Y., 2022. Mutual Fund Liquidity Transformation and Reverse Flight to Liquidity. *The Review of Financial Studies,* 35(10), p. 4674–4711.

Mittendorf, D., Neumeier, C., O'Neill, P. & Rahimi, K., 2021. Capital market liquidity in the 2020 coronavirus crisis. *FCA Research note .*

Nagel, S., 2012. Evaporating Liquidity. *The Review of Financial Studies,* 25(7), p. 2005–2039.

O'Hara, M., Wang, Y. & Zhou, X., 2018. The Execution Quality of Corporate Bonds. *Journal of Financial Economics,* pp. 130, 308–326.

O'Hara, M., 2003. Presidential Address: Liquidity and Price Discovery. *The Journal of Finance,* pp. 1335-1354.

O'Hara, M. & Zhou, A., 2021. Anatomy of a liquidity crisis: Corporate bonds in the COVID-19 crisis. *Journal of Financial Economics,* pp. 1, 46-68.

Pinter, G. & Uslu, S., 2022. Comparing Search and Intermediation Frictions across Markets. *Bank of England working paper No. 974.*

Soromaki, K. et al., 2007. The topology of interbank payment flows. *Physica A: Statistical Mechanics and its Applications,* 379(1), pp. 317-333.

Stoll, H., 1978. The Supply of Dealer Services in Securities Markets. *Journal of Finance,* 33(4), pp. 1133-1151.

12 Endeavour Square, London E20 1JN
Telephone: +44 (0)20 7066 1000 Website: www.fca.org.uk