Bank of England

Relationship discounts in corporate bond trading

Staff Working Paper No. 1,049
November 2023

Simon Jurkatis, Andreas Schrimpf, Karamfil Todorov and Nicholas Vause

Bank of England

Staff Working Paper No. 1,049

Relationship discounts in corporate bond trading

Simon Jurkatis,(1) Andreas Schrimpf,(2) Karamfil Todorov(3) and Nicholas Vause(4)

Abstract

We find that clients with stronger past trading relationships with a dealer receive consistently better prices in corporate bond trading. The top 1% of relationship clients enjoy transaction costs that are 51% lower than those of the median client – an effect which was particularly beneficial when transaction costs spiked during the Covid-19 turmoil. We find clients' liquidity provision to be a key motive why dealers grant relationship discounts: clients to whom balance-sheet constrained dealers can turn as a source of liquidity are rewarded with relationship discounts. Another important motive for dealers to give discounts to relationship clients is because these clients generate the bulk of dealers' profits. Finally, we find no evidence that extraction of information from clients' order flow is related to relationship discounts.

Key words: Corporate bonds, Covid-19, dealers, over-the-counter markets, trading relationships.

JEL classification: G12, G14, G23, G24.

---

(1) Bank of England. Email: simon.jurkatis@bankofengland.co.uk
(2) BIS and CEPR. Email: andreas.schrimpf@bis.org
(3) BIS. Email: karamfil.todorov@bis.org
(4) Bank of England. Email: nicholas.vause@bankofengland.co.uk

The views expressed in this paper are those of the authors and do not necessarily represent those of the Bank for International Settlements (BIS) or the Bank of England. For helpful comments and discussion we are grateful to Matteo Aquilina, Ben Charoenwong (discussant), Jaewon Choi, Francesco Franzoni, Thierry Foucault, Jonathan Hall, Wenqian Huang, Maureen O'Hara, Igor Kozhanov (discussant), Ulf Lewrick, Jane Li (discussant), Jessica Li (discussant), Gabor Pinter, Marco Rossi (discussant), Dimitri Vayanos, Adam Zawadowski (discussant), Alex Zhou, and seminar participants at WFA, EFA, Microstructure Exchange, 3rd CEU-ESSEC OTC Workshop, 10th SAFE Asset Pricing Workshop, Tenth Annual Conference on Financial Market Regulation at the SEC, Bank of England, Swiss Finance Institute in Lugano, 17th Central Bank Conference on the Microstructure of Financial Markets, 8th SDU Finance Workshop, Financial Stability Board Conference on Systemic risks in Non-Bank Financial Institutions and Policies to Address Them, and the Junior European Finance Conference.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1. Introduction

The over-the-counter (OTC) structure of the corporate bond market makes the value of bilateral interactions particularly important for investors. From a client's perspective, having an established relationship with a dealer can be valuable as it may allow the client to buy or sell bonds at greater ease or lower cost. For dealers, having a relationship with a client could be beneficial for managing inventory risk, for generating larger profits from loyal clients, or for extracting information from the client's order flow. The benefits of such trading relationships may be particularly pronounced during times of severe stress, like at the outbreak of the COVID-19 crisis in March 2020, when the corporate bond market experienced significant dislocations.

In this paper, we use a unique regulatory data set on corporate bond transactions to study bilateral trading relationships in the dealer-customer segment of the market. The dataset contains information about the identities of the traders, which allows us to dig deeper into the drivers of relationships than previous studies by exploring a rich cross-section of client types. Our primary goal is to quantify how the heterogeneity in prices faced by different clients – a common feature of corporate bond markets, similar to other OTC markets – can be traced to the strength of relationships with the dealer. To guide our empirical analysis, we formulate three hypotheses that could explain why relationships matter for dealers and test them in the data. Our contribution is to show that clients' liquidity provision and the attendant management of costly balance-sheet space are important factors that can explain why dealers value relationships with certain clients and quote them better prices. We also provide evidence that relationship clients generate the bulk of dealers' trading profits, suggesting a strong incentive for dealers to keep these high-value clients as customers.

To measure the strength of dealer-client relationships, we rely on past bilateral trading volume between the two counterparties. Specifically, we define a client as having a strong relationship with a dealer if the client accounts for a sizable share of the dealer's total trading volume ($Qrel$) in the past. To measure the transaction-cost benefits of relationships, we follow Hendershott and Madhavan (2015) by quantifying these costs as the log-difference between the transaction price and the closest reference price before the transaction. We use reference prices from MarketAxess at a quality and accuracy usually only available to sophisticated market participants. Our measure of relationship benefits captures the notion that a client who has a close relationship with a dealer obtains better prices relative to other clients of the same dealer. To precisely identify the reduction in transaction costs related to client-dealer-specific variation, we use panel regressions with a rich set of fixed effects (which absorb any bond, dealer, client, time, and industry-related variation) and controls for other transaction-specific variables.

In the first part of our analysis, we provide several new stylised facts on dealer-client relationships in corporate bond markets. These facts shed new light on price differentiation in OTC markets, which are often opaque and lack comprehensive data.

In the second part of our analysis, we examine whether relationship clients obtain better prices. The results from our baseline panel regression imply that the top 1% of relationship clients face a sizable *51% (4.6 basis points) reduction* in transaction costs relative to the median client. This relationship discount maps to total annual savings of around £1.3m in transaction costs for the average top relationship client. These results are robust to various alternative specifications and show that there are important client-dealer relationships in the corporate bond market.

Zooming in on the COVID-19 crisis stress episode, we find the decrease in transaction costs to be particularly important during stress times. The relationship discount more than doubled to above 10 basis points in this period. Having a relationship with a dealer is therefore particularly valuable during stress times, when top clients can trade bonds at much better prices compared to others.

We also show that relationship discounts are higher for trades that affect the dealer's balance sheet. This result is particularly pronounced during stress times. In the COVID-19 crisis period, when dealers were closer to inventory risk limits, relationship discounts rose sharply and became larger for client sales. This pattern is consistent with dealers steering their residual trading capacity towards favoured clients, especially for trades that would further increase their inventory risk.

In the third part of our analysis, we dig deeper into the underlying mechanisms driving dealer-client relationship discounts. To guide this analysis, we test three hypotheses in the data. Our first hypothesis, *"liquidity provision"*, is that dealers value relationships with clients to whom they can off-load bonds bought from other investors. Dealers may reward such clients with relationship discounts to ensure continued access to this source of liquidity. To test this hypothesis, we interact our relationship metric with a dummy variable for liquidity-providing clients. We identify such clients as counterparties to whom the dealer regularly off-loads bonds bought from other investors. If the hypothesis was true, we would expect to see that dealers charge lower transaction costs for relationship clients who regularly provide liquidity.

Our second hypothesis, *"retaining high-value clients"*, builds on the idea that an important motive for a dealer to offer better prices to certain clients is to maintain their loyalty and earn larger profits as a result of greater trading volumes (see e.g., Maskin and Riley (1984)). The effect is similar to shoppers continuing to go to the same shop due to discounts offered to them (e.g., via discount cards). To test this hypothesis, we compute dealers' total trading profits from top and non-top clients. If the hypothesis was true, we would expect to see larger profits from top clients relative to non-top clients on average. In addition, similar to the test for the liquidity provision hypothesis, we interact the main relationship metric with a dummy variable for high-value clients to test whether these clients pay lower transaction costs.

The third hypothesis, *"information extraction"*, is that dealers are willing to offer better prices to relationship clients from whom they can learn private information about the value of transacted bonds by observing their order flows (e.g., proprietary trading firms). If such considerations play a role, the relationship metric should have a more pronounced impact on transaction costs for such informed clients. To test this hypothesis, we interact the main relationship metric with a dummy variable for 'informed' clients, which we identify as those whose trades tend to predict future price moves.

Our results strongly support the "liquidity provision" hypothesis. In particular, we find evidence that dealers use relationship discounts to help control their inventories and that clients who regularly help dealers in this way receive larger discounts. We show that clients who consistently provide liquidity to dealers – i.e. by purchasing bonds shortly after dealers have absorbed sales from other clients – receive larger relationship discounts than other top clients. This extra discount was magnified six times during the COVID-19 episode. In fact, it was essentially only liquidity-providing clients who received relationship discounts during this period of stress.

We also find evidence in support of the "retaining high-value clients" hypothesis. We show that the average profit dealers make from top relationship clients is more than *18 times larger* than the profit from non-top clients. The much larger profit extracted from top clients gives dealers a strong incentive to retain trading volume from these clients by offering them more competitive prices. In addition, we find that dealers give high-value clients a larger discount than other relationship clients, after controlling for liquidity- and information-related motives.

However, we find no evidence in support of the "information extraction" hypothesis. Relationship clients receive no additional transaction-cost discounts when their order flow provides valuable signals about future bond returns. This finding could be explained by regulatory restrictions on proprietary trading by dealer-banks that were introduced in the wake of the global financial crisis (GFC). These regulations limited the extent to which dealers can trade on signals from customer flows, thus weakening the incentive to build relationships for the purpose of information extraction. In addition, dealers may also be reluctant to trade with informed clients who correctly anticipate future price changes on average, or could at least charge them higher transaction costs, due to the risk of the price moving against them after trading.

Taken together, our results point to two main economic mechanisms that could explain relationship discounts received by clients. First, dealers value clients to whom they can "outsource" some of their liquidity provision. In providing liquidity to the market, dealers often take on inventory risk. Shifting some of this risk onto liquidity-providing clients is valuable and rewarded with transaction-cost discounts. These discounts appeared particularly important during the COVID-19 crisis when inventory risk was magnified by asset price volatility, and liquidity-providing clients were the only client type to receive significant transaction-cost discounts. Second, we find that dealers earn much higher profits from a handful of high-value clients compared to the vast majority of clients. This creates a strong incentive for dealers to retain these high-value clients as loyal customers by offering them trading discounts.

The findings of our paper have several implications for investors and policy makers. First, we find evidence that dealers continued to provide liquidity at discounted prices to relationship clients during the COVID-19 shock. These results indicate that the OTC market structure centred around dealer intermediation proved relatively resilient for such clients, with relationships apparently incentivizing dealers to provide liquidity even during stress times. That said, our findings also suggest that it is only dealers' top clients who benefit from such relationships. Many smaller market participants do not enjoy such benefits and their cost of accessing dealers' balance sheets can surge in stress episodes. Second, dealers appear to particularly value clients to whom they can turn for liquidity provision. Such access to liquidity could help dealers to operate with smaller inventories, which are cheaper to maintain, and thus increase their own liquidity provision. Third, the fact that dealers do not offer larger discounts to informed clients is consistent with dealers scaling back proprietary trades, which contributed to the build-up of risk before the GFC and were subsequently discouraged by new regulations.

Related Literature. Our paper contributes to the growing literature on the value of relationships in OTC markets. Three important papers in this literature which, similar to ours, focus on corporate bond markets, are Di Maggio et al. (2017), Hendershott et al. (2020) and O'Hara et al. (2023). Di Maggio et al. (2017) show that dealers value relationships with other dealers in US corporate bond markets, particularly in stress times. That study, however, focuses entirely on relationships in the inter-dealer market whereas our paper covers the dealer-client segment (i.e., dealer trades with ultimate "end-users"). In our sample, the dealer-client market is roughly *four times larger* than the inter-dealer market as measured by the average daily trading volume. Also, the sample in Di Maggio et al. (2017) ends in 2011 and hence covers a period that predates important regulations implemented in the aftermath of the GFC, which had a material impact on dealer behavior (see, inter alia, Bao et al., 2018; Dick-Nielsen and Rossi, 2019).[1]

[1] Bao et al. (2018) study the effects of the Dodd–Frank Act over the period 21 July 2010 to 31 March 2014, and the effects of the Volcker Rule from 1 April 2014 to 31 March 2016. The authors show that dealers reduced their market-making activity in response to the Volcker Rule.

Hendershott et al. (2020) study the dealer-to-client segment, but they do so for only a narrow subset of clients, studying primarily insurance companies. The authors find that the number of trading relationships that insurance companies have with dealers has a non-monotonic impact on their transaction costs. Transaction costs decrease initially with a rise in the number of dealers due to increased competition, but eventually increase as bilateral relationships become more dispersed and weaker. Compared to Hendershott et al. (2020), we study transactions of dealers with a much broader set of clients, covering all clients with a legal entity identifier (LEI). This includes key players such as asset managers, non-dealer banks, principal trading firms and hedge funds. Insurance companies account for less than 9% of total trading volume in our sample. Consistent with the findings of our paper, O'Hara et al. (2023) show that dealers improve their liquidity provision in the corporate bond market when they have trading relationships with insurers.

Understanding the value of relationships from the client's perspective is particularly important against the backdrop of an evolving microstructure of bond markets. Key developments in recent years have been the advent of electronic trading platforms, particularly based on request-for-quote (RFQ) protocols, which represent an electronic form of dealer-intermediated OTC trading. Despite this "electronification", all-to-all trading – a prominent feature of equities or futures markets – has so far remained very limited in cash bond markets (O' Hara et al., 2018; O'Hara and Zhou, 2021). Our finding that relationship clients trade on better terms than non-relationship clients, especially in stress times, could help explain why the market may be slow to adapt to a fully electronic and anonymous market structure. Despite the generally favourable liquidity of market structures built around centralized limit order books (Hendershott and Madhavan, 2015; Abudy and Wohl, 2018), the OTC structure has shown remarkable resilience in fixed income markets.

The preference of institutional investors for a bilateral OTC market structure is also discussed in Biais and Green (2019) and Wittwer and Allen (2021). In particular, Biais and Green (2019) highlight institutional traders' preference for low price impact and their relative bargaining power compared to retail traders as key factors in explaining the historical evolution of the US bond markets towards an OTC structure. Wittwer and Allen (2021) include a loyalty benefit in their structural model, where institutional investors choose between trading on an RFQ platform or bilaterally. Our findings confirm the importance of including such loyalty benefits in structural models.

Our paper highlights the heterogeneity of trading relationships between dealers and clients and shows that this leads to large differences in trading costs. These findings have implications for models analysing the welfare benefits of different market structures (Plante, 2018; Lee and Wang, 2018; Vogel, 2018; Wittwer and Allen, 2021). On the one hand, removing relationship benefits by mandating a centralized and anonymous market structure through all-to-all trading would directly impact the utility of clients enjoying relationship discounts. On the other hand, relationship benefits are typically given to selected clients and could be implicitly subsidized by other clients who are less capable of establishing a relationship (e.g., due to less-frequent trading activity).

Our paper is also related to the broader literature on dealer-client relationships in other OTC markets, notably Hau et al. (2021). In this paper, the authors show that unsophisticated clients, that have a relationship with only one dealer, incur larger transaction costs in FX derivatives.

More broadly, our paper also relates to the literature on dealer intermediation in the aftermath of the GFC. Several studies argue that dealers have become more constrained in their liquidity provision and are generally less willing to warehouse bonds in their inventory (Adrian et al., 2017; Bao et al., 2018; Dick-Nielsen and Rossi, 2019; Choi et al., 2021). This has led to a partial shift away from the principal model of market making, where dealers warehouse risk on their balance sheet, towards more "balance sheet-light" approaches where dealers mostly line-up opposing client trades in advance of simultaneous execution. Our paper shows that clients are asymmetrically impacted by these developments. Choi et al. (2021) document that clients complement the role of dealers by stepping in as liquidity providers in the post-GFC period. They show that dealers are more likely to turn to insurance companies with whom they have stronger relationships to offset trades with other clients. Our paper covers dealers' customer base comprehensively and shows that it is mostly asset managers and brokers that provide liquidity to dealers, accounting for almost 90% of the volume across liquidity-providing clients. Insurers, by contrast, account for less than 6% of this volume. In addition, we show that dealers value relationships with liquidity-providing clients. They do so not only by giving discounts to those clients in normal times, but also by being prepared to support them in times of stress by offering significant transaction-cost discounts to trade with them, and especially to absorb bond *sales* from them.

Finally, a number of papers study trading relationships in OTC markets in the context of informed trading. Kondor and Pintér (2022) and Czech and Pintér (2020) show that clients who spread their trades over a larger number of dealers outperform other clients, consistent with them hiding private information. We show, however, that dealers do not quote better prices to more informed clients, which may be due to post-GFC regulations that have discouraged proprietary trading or because of fears of being adversely selected.

2. Data on dealer-client transactions and corporate bond reference prices

This section introduces our two main data sets, which come from MiFID II regulatory reporting and MarketAxess. It then provides descriptive statistics on the corporate bonds and dealer-client transactions in our sample.

2.1. MiFID II data

The first main dataset for our analysis consists of transaction reports in corporate bonds submitted to regulatory authorities under MiFID II, which took effect on 3 January 2018. Under this regulation, investment firms and other trading institutions are mandated to submit reports for their trades in debt instruments that are permitted to trade on a venue. Venues include regulated markets (such as the London Stock Exchange), multilateral trading facilities (such as

RFQ platforms like TradeWeb, MarketAxess, or Bloomberg) and organised trading facilities.[2] Trades have to be reported irrespective of whether they are actually carried out on a venue, as long as the instrument is admitted to be traded on a venue. This amendment to the previous legal framework marks a significant improvement in terms of the data coverage of OTC markets. Under the preceding directive (MiFID I), trades had to be reported only when the instrument was permitted to trade on a regulated market—a requirement that many corporate bonds do not fulfil. Using these data distinguishes our paper from studies that used data collected under the MiFID I regime. Another benefit of our data is the level of granularity and detail compared to TRACE data for US corporate bonds. The key advantage of our MiFID II data is that they allow us to identify *both* counterparties of the trade, instead of only the dealer as in TRACE.

The data are made available to us by the Financial Conduct Authority (FCA), the UK's financial markets regulator. The FCA receives reports for all transactions in reportable financial instruments involving at least one UK investment firm or executed on a UK trading venues. Each report includes information on the ISIN of the instrument traded, the time of the transaction (time-stamped to at least the nearest second), the price and the quantity. As mentioned above, each report also identifies both counterparties of the trade, including for those trades where the counterparty is a client. This beneficial feature of our data allows us to study trading relationships between each dealer-client pair, going beyond the inter-dealer segment. As a result, we are able to shed light on the nature and importance of relationships from the perspective of clients, who are the ultimate end-users that shed or take on risk exposures.[3] Finally, we complement these data with information on bond characteristics, such as maturity and credit ratings, from S&P Capital IQ and ESMA's Financial Instruments Reference Database System (FIRDS).

[2]For precise definitions see point 20 to 23 of Article 4(1) of the MiFID II Directive. The trades in our data are executed on a UK venue, involve at least one UK counterparty, or executed on a EU venue in a bond regulated by the FCA. The vast majority of bonds in our sample are denominated in EUR and USD, as shown in columns 1 and 2 of Table A.1 in the Appendix.

[3]Inter-dealer transactions, by contrast, typically serve the purpose of inventory risk management. While taking bonds into inventory is commonplace for non-dealers, a key incentive of dealers is to minimize inventory risk. This is often achieved by trading in inter-dealer markets.

2.2. Reference prices from MarketAxess

For the computation of transaction costs, we require reference prices that reflect the fair value of the asset at the time of the transaction. However, in OTC markets, which are highly fragmented with only intermittent trading activity, establishing a reference price is not straightforward. A common practice in the literature is to use inter-dealer prices, but dealer-client trades and inter-dealer reference trades are not always proximate in the corporate bond market. Indeed, not infrequently they are days apart, by which time the reference price would likely have moved.

For our reference prices, we use proprietary mid-quote data from MarketAxess Composite+ (CP+). These data provide a level of pricing information that is usually only available to sophisticated market participants.[4] CP+ is based on a proprietary machine-learning pricing algorithm developed by MarketAxess that generates pre-trade reference prices for corporate bond investors. The pricing engine leverages data not only on the bond being priced, but also other related bonds. The sources of data include reported trade prices, RFQ responses sent by liquidity providers on the MarketAxess trading platform, and indications of trading interest streamed by dealers.[5] We received reference prices sampled at 8am London time, 8am New York time and 4pm London time each day. For the period 1 to 18 March 2020, which we will use as our crisis period, MarketAxess also provided us with reference prices on a tick-by-tick basis.

2.3. Descriptive statistics

Table 1 provides key descriptive statistics for the dealer-client transactions in our sample, which runs from 3 January 2018 (the start of MiFID II reporting) to 18 March 2020 (the end of the 2020 dash-for-cash episode).[6] In much of our paper, we investigate relationship discounts and

[4]The MarketAxess CP+ reference prices are observed much more frequently than inter-dealer quotes, which are typically used in the literature to date. Table A.3 in the Appendix shows that MarketAxess CP+ prices are observed almost four times more frequently, on average: every 2.2 hours compared to every 8 hours for inter-dealer trades.

[5]The input data is then fed into a tree-based machine learning algorithm called Gradient Boosting Method (GBM). For more information, see MarketAxess (2018).

[6]We do not use data shortly after the end of the COVID-19 turmoil period in order to avoid confounding our findings with the effects of policy measures introduced in response to the COVID-19 crisis. See Table 1 on page 14 of the Bank of England's May 2020 Monetary Policy Report.

transaction costs separately during normal and crisis times. Hence, we split our sample into a pre-crisis (before March 2020) and crisis period (1 to 18 March 2020). We set the beginning of the turmoil period to the start of increased selling pressure in bond markets.[7] We choose the end of the turmoil period, March 18, as the time when the "dash for cash" abated, due in large part to the actions of major central banks.

Overall, our dataset comprises transactions of more than 50 dealers in corporate bond markets pre-crisis, with almost all of them active during the COVID-19 crisis too (Panel A of Table 1). In total, dealers interact with more than 17,300 (4,500) clients before (during) the crisis.[8] We record a total of 6.7 million dealer-client transactions in our sample, worth a total nominal value of GBP 5.8 trillion. Overall, more than 39,200 (17,500) different bonds were traded at least once before the COVID-19 crisis (during the crisis period).

Next, we analyze the concentration of trading, both across dealers and clients. Figure 1 (left panel) shows that dealers differ significantly in terms of their market footprint (as measured by their share of aggregate trading volume). Trading is highly concentrated, with the largest 14 dealers accounting for 80% of trading volume of all dealer-client transactions.[9] The right panel of Figure 1, in turn, displays the size distribution for the largest 500 clients. The concentration of client volumes is even more extreme than that of dealers. The trading volumes of the largest clients are equivalent to those of medium-sized dealers.

The 'Averages' section of Panel A of Table 1 provides some further insights into trading by dealers and customers. On average, 48 dealers were active on any given day before the crisis and around 50 were active during the crisis. Client demand for dealers' intermediation services and the supply of those services both appeared to increase during the crisis: the number of clients active in the market increased from 1,400 to 1,600, and dealers' average daily trading volume increased by roughly 40%. Clients interact on average with about 3 dealers on a given day,

[7]See, for example, charts A.4 to A.6 in the Bank of England's August 2020 Financial Stability Report.

[8]This refers to all clients endowed with a Legal Entity Identifier (LEI). We exclude all transactions between dealers and clients where both belong to the same parent company, as identified by the Global Legal Entity Identifier Foundation (GLEIF) database.

[9]This finding is consistent with other studies of the corporate bond market, which also find that few dealers have a large share of total trading volume, e.g. O' Hara et al. (2018).

Table 1: Descriptive statistics

Panel A: Dealer-to-client trades

	Totals		Averages							
			Daily		Dealer-day		Client-day		Bond-day	
	Pre-crisis	Crisis	Pre-crisis	Crisis	Pre-crisis	Crisis	Pre-crisis	Crisis	Pre-crisis	Crisis
#Dealers	53	51	47.9	50.3			3.2	3.2	1.6	1.8
#Clients	17.3k	4.5k	1.4k	1.6k	94.9	104.9			1.9	2.1
#Bonds	39.2k	17.5k	5.2k	6.4k	177.2	221.9	6.9	8.1		
#Trades	6.5m	211k	11.8k	16.3k	245.8	323.1	8.3	9.9	2.3	2.6
Volume	5.6tn	181.9bn	10.2bn	14.0bn	212.3m	277.6m	7.1m	8.5m	1.9m	2.2m

Panel B: Bond characteristics

	Rating distribution (in %)				Maturity	Illiquidity
	AAA	A-AA	BBB	HY	(years)	%ZTD
pre-crisis	6.18	29.70	39.54	24.59	7.38	54.23
crisis	6.01	31.74	34.40	27.86	7.62	48.59

Panel C: Client sector

Sector	%
Asset Manager	58.98
Bank	15.39
PFLDI	8.42
Broker	8.16
Hedge Fund	7.85
PTF	1.19

Panel D: Transaction costs (bps)

		#Obs	mean	std	median
pre-crisis	buy	1.2m	9.2	36.7	5.2
	sell	1.2m	8.7	37.6	3.9
crisis	buy	41.6k	1.0	64.3	4.4
	sell	60.6k	50.2	81.3	21.7

Panel E: Qrel correlation

	(in %)
$(t, t-1)$	79.0
(i, j)	20.6
$(t, t-1)$	87.4
(i, j)	16.8

Panel F: Qrel distribution

	mean	std	1%	25%	50%	75%	99%
Qrel (in %)	1.30	3.53	0.00	0.06	0.28	1.01	16.78

Notes: This Table shows descriptive statistics for the pre-crisis (3 January 2018 to 29 February 2020) and crisis period (1 to 18 March 2020). Panel A displays the total number of dealers (#Dealers), clients (#Clients), bonds (#Bonds) and trades (#Trades), as well as total trading volume measured in GBP. Panel B presents bond characteristics. Maturity is the volume-weighted average residual maturity across transactions and is measured in years. %ZTD is a measure of illiquidity calculated as the share of days with no trade in a particular bond. The rating distribution is measured in terms of trading volume of all bonds that had a rating in February 2020. High-yield (HY) is all ratings below BBB. Panel C shows the share of trading volume by client sector: asset managers, banks, pension funds, insurers and liability-driven investors (PFLDI), brokers including executing and investing services firms, hedge funds, and proprietary trading firms (PTF). Panel D shows average transaction costs measured in basis points (bps) for client buys and sells. Panel E presents average relationship persistence (row $(t, t-1)$) and overlap (row (i, j)) as defined in Eqs. (4) – (5). The relationship metric is calculated as the client's trading volume relative to that of all other clients of the same dealer over a window of 180 days before the time of the transaction: see Eq. (3). Panel F presents the distribution of the relationship metric $Qrel$ measured in % across all transactions.

Figure 1: Trading volume concentration across dealers and clients

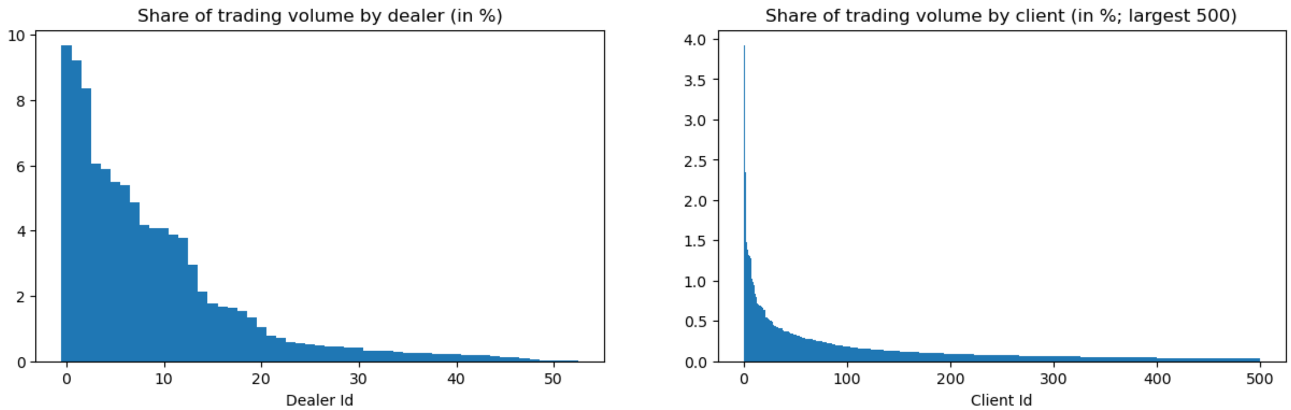

Notes: The Figure displays the share of trading volume (in %) for dealers (left) and clients (right) over the full sample period.

whereas dealers interact with 95 clients (105 during crisis), corroborating the earlier observation that clients' trading behavior is much more concentrated than that of dealers.

We next study the main characteristics of the bonds traded in our sample in terms of credit risk, duration, and liquidity. Specifically, we provide statistics for volume-weighted aggregate measures of the bonds' credit rating, time-to-maturity, and percentage of zero-trading-days (%ZTD). Panel B of Table 1 shows that the average bond has around 7.4 years to maturity and does not trade on 54% of trading days. The latter figure decreases to 49% in the crisis, suggesting that more liquid bonds were traded during the COVID turmoil. In terms of credit quality, around one (three) quarter(s) of trading volume is in high-yield (investment-grade) bonds.

Finally, we study the composition of investors active in corporate bond trading. Panel C of Table 1 reports the distribution of trading volume across different client sectors in the whole sample. The asset management sector is by far the largest, accounting for 59% of the trading volume. Banks that do not form part of the dealer community are the second largest group with more than 15% of trading volume. Pension funds, insurance companies and other liability-driven investors (PFLDI), as well as brokers and hedge funds each account for around 8% of trading volume. Principal trading firms (PTFs) have a small share of the total trading volume, which might be related to the low degree of electronification in corporate bond trading. We provide some additional summary statistics on currencies, countries of residence of issuers, and industries in the Online Appendix, Table A.1 and Table A.2.

3. Empirical strategy and measurement of key variables

The main goal of our paper is to study the extent to which dealers give transaction-cost discounts to clients, depending on the strength of their past trading relationships with those clients. In this section, we start by setting out our econometric approach—panel regressions with a rich set of fixed effects in order to tease out the effect of relationships on transaction costs. We then describe and present stylized facts on our main dependent and independent variables: transaction costs faced by different clients and our measure of dealer-client relationship strength ($Qrel$).

3.1. Econometric approach

To understand if the strength of relationships between dealers and their clients affects the transaction costs that clients incur, we estimate panel regressions with fixed effects of the form:

$$TC_{bdct} = \gamma\, Qrel_{dct} + \mathbf{X}'_{bdct}\,\beta + \mathbf{1}'\mu + \varepsilon_{bdct}, \tag{1}$$

where TC_{bdct} is the transaction cost for a trade between dealer d and client c in bond b at time t and $Qrel_{dct}$ is our relationship measure (which we describe below).

Fixed effects. The vector μ in Eq. (1) contains bond-month, dealer-month, client-month and industry-day fixed effects. This rich set of fixed effects allows us to control for many observable and unobservable variables that may influence transaction costs. For example, a bond's liquidity likely affects transaction costs for all clients. Including bond-month fixed effects thus ensures that our results are not affected by relationship clients systematically trading more liquid bonds. Similarly, the size, network centrality, or market power of clients and dealers could also affect transaction costs. However, including client-month and dealer-month fixed effects ensures that our results are not driven by such characteristics. Industry-day fixed effects absorb any variation in transaction costs related to systematic time-varying differences by day and across industries: e.g., if relationship clients systematically trade bonds from certain industries, which could affect the crisis results, for example, given the nature of the COVID shock. In several robustness tests, we also include dealer-day and bond-day fixed effects.

Additional controls. We control for additional variables in the vector $\mathbf{X}_{bdct}$. These include three dummy variables, $sell_{bdct}$, $match_{bdct}$ and MTF_{bdct}, which take the value 1 if the client sells, the dealer matches the trade with a transaction in the opposite direction, or the trade is executed on a regulated market or multilateral trading facility, respectively.[10] To compute $match_{bdct}$, a match is defined to be a transaction in which the dealer buys/sells the bond instantaneously and does not take on any balance-sheet risk. In practice, a dealer would line up these trades with different clients in advance and then execute them simultaneously, which is captured by our dummy.[11] In general, we would expect such transactions to be associated with lower costs for the client, as shown by Goldstein and Hotchkiss (2020). We also expect that having a strong relationship with a dealer could be more beneficial for non-matched trades, which would require the dealer to warehouse some risk. Since more volatile bonds could have larger transaction costs, we also control for bond's volatility by including the lagged squared intra-day return on the benchmark, r^2_{bt-1}.[12] Finally, we also control for the size of the transaction (in logs).

3.2. Measuring transaction costs

To capture price discounts that dealers offer to (certain) clients, we compute the transaction costs faced by clients trading with the dealer. We follow Hendershott and Madhavan (2015),

[10]Since trading on regulated markets, such as the London Stock Exchange, accounts for less than 1%, the MTF dummy effectively captures trading on multilateral trading facilities.

[11]All other trades in which the dealer trades the same bond in opposing directions - even if only a few seconds apart - are not pre-arranged and thus are risky for the dealer since she acts as a principal.

[12]We use lagged squared returns instead of contemporaneous ones to avoid bias in OLS estimates, which stems from using future benchmark price changes for transactions observed before the end of the day.

Hau et al. (2021) and others to measure transaction costs as

$$TC = \log\left(\frac{P}{P_b}\right) \times D, \tag{2}$$

where P denotes the transaction price, P_b is the closest CP+ mid-quote for the traded bond in the 24 hours prior to the transaction, and D is the trade direction of the client, taking the value +1 for a purchase and -1 for a sale. We multiply transaction costs in Eq. (2) by 10,000 to measure them in basis points (bps).

The transaction cost measure in Eq. (2) captures the extent to which the price paid by the client differs from the reference price prevailing in the market at the time of the transaction. This measure captures the transaction costs from the client's perspective.

Basic descriptive statistics on transaction costs. Panel D of Table 1 reports the volume-weighted average transaction costs for client purchases and sales (denoted by buy/sell), both before and during the March 2020 stress period.[13] We see that average transaction costs for client sales (dealer purchases) of 9.2 bps were similar to those for client purchases (8.7 bps) in the pre-crisis period. However, during the crisis period, transaction costs became highly asymmetric for client sales vs. client purchases. Transaction costs *dropped* to around 1 bp for client purchases but *jumped fivefold* to 50 bps for sales during the crisis. At this time, dealers faced significant net selling pressure from clients. Hence, they may have tried to discourage client sales (which would have further increased the size of their balance sheets) through higher transaction costs, and encouraged client purchases (which would provide balance sheet relief) through lower transaction costs.

3.3. Measuring dealer-client relationship strength

We now turn to the measurement of our key variable that seeks to capture the strength of a relationship between a dealer and a client.

[13]The number of observations in Panel D and in our regression sample is smaller relative to other panels in Table 1 because we use only transactions for which there is a benchmark price observed in the last 24 hours.

We measure the strength of dealer-client relationships ahead of each transaction. To do this, we begin by calculating the total trading volume of the dealer-client pair across all bonds over the previous 180 days, lagged by one week.[14] We then divide that measure by the total volume of trading between the dealer and all of its clients. This gives us the share of trading volume that dealer d obtains from trading with client c:

$$Qrel_{dct} = \frac{\sum_{\tau=t-187}^{t-7} Q_{dc\tau}}{\sum_{k\in C}\sum_{\tau=t-187}^{t-7} Q_{dk\tau}}, \tag{3}$$

where C is the set of clients of the dealer over the 180 days window. Intuitively, $Qrel$ captures, from a dealer's perspective, the importance of a particular client based on its contribution to the dealer's overall trading volume over the past 180 days.

Basic facts on dealer-client relationships. We now present some basic descriptive statistics on our dealer-client relationship measure. Panel F of Table 1 reports the distribution of the relationship metric over all dealer-client transactions in our sample. As the Table shows, the vast majority of clients only account for a small share of a dealer's trading volume: the median $Qrel$ client accounts for only 0.28% of dealer's trading volume. However, there is a significant heterogeneity among clients as indicated by the large standard deviation (relative to the median and mean). An important observation is that there is a small number of clients who account for a sizeable portion of a dealer's overall trading business: the top 1% of $Qrel$ clients accounts for about a sixth of a dealer's trading volume over the past 180 days.

To gain further insights into the nature of relationships in corporate bond trading, we examine the persistence and overlap of relationships. First, for each dealer, we measure the persistence of client relationships over time based on the correlation of client volume shares over two adjacent,

[14]We use the lagged measure to make sure that the information is readily available to a dealer and can serve as an input into the quote setting when responding to a client's trading requests. Our results are similar if we use one day or one week lag, or 90 days window (see Table A.4 in the Online Appendix).

non-overlapping periods:

$$corrQ_d(t, t-1) = \text{corr}(Qrel_{d,t}, Qrel_{d,t-1}), \tag{4}$$

where $Qrel_{d,t}$ denotes the vector of client volume shares in a dealer's overall trading volume at time t.[15] Panel E in Table 1 shows that relationship strengths are fairly persistent over time. The average $corrQ_d(t, t-1)$ is 79% in the pre-crisis period and the correlation measure increases to 87% during the COVID-19 turmoil.

Second, to investigate whether relationship clients of one dealer (i) overlap with those of other dealers (j), we use a similar metric to $corrQ_d(t, t-1)$, but computed over the cross-section of dealers:

$$corrQ_t(i, j) = \text{Corr}(Qrel_{i,t}, Qrel_{j,t}). \tag{5}$$

Panel E in Table 1 shows that relationship clients are not especially common across dealers.[16] On average, the cross-sectional correlation measure $corrQ_t(i, j)$ is only 21% in the pre-crisis period. This finding suggests that when clients have a strong relationship with a dealer they tend to be quite loyal to that dealer and do not simultaneously show up as a strong relationship client for another dealer. The cross-sectional correlation between dealers' $Qrel$ measures decreased slightly during the crisis period. This pattern suggests that the trading of relationship clients became even more concentrated during that episode, with clients increasing their reliance on their relationship dealer.

[15] In unreported results we used the rank correlation and, as an alternative measure, the share of a dealer's top clients in one period that were also top clients in the previous 180-days period. Both measures also showed high persistence of relationships over time.

[16] Again, this finding is robust to using the rank correlation and an alternative measure based on the share of a dealer's top clients that are also top clients of another dealer during the same period. The results are excluded for brevity.

4. Relationship discounts in corporate bond trading

In this section, we report our baseline results on relationship discounts obtained from estimating Eq. (1). We begin with relationship discounts in normal times, before turning to discounts during the COVID-19 crisis period. We also report how transaction costs vary with several trade characteristics.

Table 2 (Panel A) presents results for the pre-crisis period. It shows that the relationship metric $Qrel$ has a negative and statistically significant effect on transaction costs, regardless of our fixed effects specification. The coefficient estimates in columns 1 and 2 imply that transaction costs are 0.28 basis points lower for every percentage point increase in $Qrel$. This suggests that top-percentile clients, with $Qrel = 16.8$ (as reported in Table 1), pay transaction costs that are 4.6 basis points lower than those of median clients with $Qrel = 0.3$ (calculated as $-0.28 \times (16.8 - 0.3) = -4.6$). This reduction for relationship clients represents a sizeable *51%* discount on the average transaction cost of 9 basis points that the typical clients face in the pre-crisis period. It amounts to annual savings of £1.3m based an average daily trading volume of £11m for top clients and 252 trading days (£11m $\times 4.6/10,000 \times 252 = 1.3m$).

Table 2 (Panel B) presents results for the COVID-19 crisis period. It shows coefficient estimates for $Qrel$ that are roughly double their pre-crisis values. For example, the medium estimate in column 5 suggests that transaction costs were 0.61 basis point lower for every percentage point increase in $Qrel$. This implies that top-percentile clients paid transaction costs that were 10.1 basis points lower than those of median clients during the crisis period ($-0.61 \times (16.8 - 0.3) = -10.1$), more than double the 4.6 basis point reduction in the pre-crisis period.

These pre-crisis and crisis results hold after controlling for other potential determinants of transaction costs. In particular, we control for several trade characteristics that may influence transaction costs, including through their effects on dealer balance sheets.

First, we allow for trade size to affect transaction costs. Consistent with previous studies (e.g. Pintér et al., 2022b), we find that larger trades are more expensive to execute. This may

Table 2: Baseline results on relationship discounts

	Panel A: Pre-crisis			Panel B: Crisis		
	(1)	(2)	(3)	(4)	(5)	(6)
Qrel	-27.93^{**}	-27.85^{**}	-31.93^{***}	-52.89^{**}	-60.89^{**}	-69.15^{**}
	(9.78)	(9.80)	(10.12)	(26.13)	(27.75)	(29.92)
ismatch	-3.45^{***}	-3.45^{***}	-3.64^{***}	-10.07^{***}	-7.63^{***}	-6.37^{***}
	(0.85)	(0.85)	(0.74)	(2.56)	(2.02)	(1.78)
sell	-0.77	-0.80	-0.02	38.49^{***}	36.19^{***}	39.60^{***}
	(1.44)	(1.43)	(1.63)	(3.96)	(3.79)	(4.56)
logQ	0.61^{***}	0.61^{***}	0.69^{***}	0.47^{*}	0.53^{*}	1.24^{***}
	(0.06)	(0.06)	(0.06)	(0.28)	(0.28)	(0.33)
MTF	2.34^{**}	2.31^{**}	2.59^{**}	-0.73	0.57	0.15
	(0.85)	(0.85)	(1.06)	(2.17)	(1.98)	(2.25)
r^2_{bt-1}	-0.07	-0.44		6.68^{***}	-0.95	
	(0.38)	(0.30)		(0.59)	(0.76)	
Nobs	2.3m	2.3m	1.9m	91.6k	91.6k	84.5k
R^2	0.20	0.21	0.35	0.31	0.36	0.48
Fixed effects						
dealer × month	Yes	Yes				
client × month	Yes	Yes	Yes			
bond × month	Yes	Yes				
industry × day		Yes			Yes	
dealer × day			Yes			Yes
bond × day			Yes			Yes
dealer				Yes	Yes	
client				Yes	Yes	Yes
bond				Yes	Yes	

Notes: This Table shows the results of our baseline regression (Eq. (1)) for the pre-crisis period from 3 Jan 2018 to 29 Feb 2020 (Panel A) and the crisis period from 1 to 18 March 2020 (Panel B). The dependent variable is the transaction cost of a corporate bond trade between a dealer and a client as described in Eq. (2) and measured in basis points. The independent variables are: *Qrel*, the share of the client's trading volume in total dealer's trading volume over a past window of 180 days as defined in Eq. (3); *match*, an indicator variable equal to one if the dealer offsets the trade with other trades executed at the same instant and in the opposite direction; *sell*, an indicator variable equal to one if the client is selling; MTF, an indicator variable equal to one if the trade is executed on a regulated market (e.g., London Stock Exchange) or a multilateral trading facility (e.g., MarketAxess); *logQ*, the natural logarithm of the trade size measured in GBP and r^2_{bt-1}, the MarketAxess benchmark's lagged squared daily return. Standard errors (shown in parentheses) are clustered at the dealer and month level for the pre-crisis period and at the dealer level for the crisis period. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%).

reflect compensation for dealers for taking on more inventory risk when absorbing larger trades.

Second, we include a control dummy variable that takes the value of one for trades executed electronically on a multilateral trading platform. We find that in non-crisis times, transaction costs for trades executed on such platforms are about 2.5 basis points *higher* than for other OTC trades.

Third, we control for matched trades. In a matched trade, a dealer simultaneously executes pre-arranged offsetting sales and purchases, such that the set of trades has no effect on its bond inventory. Transaction costs may be lower for matched trades as dealers do not require compensations for any additional inventory risk (Goldstein and Hotchkiss, 2020; Choi et al., 2021). Indeed, we find that in normal times transaction costs for matched trades are about 3.5 basis points lower than for non-matched trades. We also find this effect to be 2–3 times larger in the crisis period. This suggests that dealers charge more for balance-sheet-intensive trades, especially at times of crisis, when dealer balance sheets are likely to be more constrained.

Finally, we allow transaction costs to vary with the direction of the client's trade. In situations where dealers do not already have a pre-arranged matching trade, a client *sale* results in the dealer taking the bond into its inventory, whereas a client *purchase* draws from the dealer's inventory. While the former will tighten balance-sheet constraints, the latter will alleviate them. Despite this asymmetry, we find no significant difference between transaction costs for client purchases and sales in the pre-crisis period. In the crisis period (Panel B of Table 2), however, we find that client *sales* were much more expensive than client *purchases*: in fact, 36–40 basis points more expensive. A likely explanation is that higher volatility during the crisis boosted dealers' inventory risk and left them less able to absorb additional bond sales. This was likely compounded by expectations that bonds would need to be held in inventories for longer, given the imbalanced market with few buyers. Hence, dealers tried to deter trades that would *add* to their inventories by making them relatively expensive.

Relationship discounts for unmatched trades. We next study whether balance sheet motives interact with the size of the discount given to relationship clients. Since unmatched trades are more likely to affect dealer inventories, we first test whether the discount given on unmatched trades differs across relationship and non-relationship clients. We do so by estimating:

$$TC_{bdct} = \gamma \, Qrel_{dct} + \alpha_1 \, nomatch_{bdct} + \beta_1 \, nomatch_{bdct} \times Qrel_{dct} + \mathbf{X}'_{bdct}\beta + \mathbf{1}'\mu + \varepsilon_{bdct}. \qquad (6)$$

As shown in the first column of Panel A in Table 3, and consistent with earlier results, dealers charge roughly 4 basis points more to execute unmatched trades compared to matched trades in the pre-crisis period (α_1). More importantly, dealers offer a relationship discount of about 31 basis points per unit of $Qrel$ for unmatched trades ($\gamma + \beta_1$), compared to an insignificant 9 basis points discount for matched trades (γ). During the crisis period (Panel B), both the unconditional cost of unmatched trades (α_1) and the relationship discounts for such trades ($\gamma + \beta_1$) increased sharply. These results are consistent with dealers applying preferential pricing to relationship clients in principal trades that require dealers to warehouse risk on their balance sheet.

We next study whether dealers give larger discounts for unmatched client *sales* than unmatched client *purchases*. In the corporate bond market, a client sale usually increases the dealer's inventory, whereas a client purchase tends to be filled out of the dealer's inventory and hence reduces it.

Hence, we interact the dummy variables for unmatched trades and client sales and estimate the following regression:

$$\begin{aligned} TC_{bdct} = {} & \gamma \, Qrel_{dct} + \alpha_1 \, nomatch_{bdct} + \alpha_2 \, sell_{bdct} + \alpha_3 \, nomatch_{bdct} \times sell_{bdct} \\ & + \beta_1 \, Qrel_{dct} \times nomatch_{bdct} + \beta_2 \, Qrel_{dct} \times sell_{bdct} + \beta_3 \, Qrel_{dct} \times nomatch_{bdct} \times sell_{bdct} \\ & + \mathbf{X}'_{bdct}\beta + \mathbf{1}'\mu + \varepsilon_{bdct}. \qquad (7) \end{aligned}$$

The coefficient β_3 then captures the additional relationship discount for unmatched client sales, and the overall relationship discount for such trades is $\gamma + \beta_1 + \beta_2 + \beta_3$. The large negative estimate on β_3 in column 4 shows that dealers charge relationship clients significantly smaller costs for unmatched client sales during the crisis.[17] Thus, when dealers faced one-sided selling pressure that clogged their balance sheets, they gave preferential pricing to relationship clients.

The results from this section suggest two immediate conclusions about transaction costs in corporate bond markets. First, they differ markedly depending on the strength of the dealer-client trading relationship. Second, they depend on how the trade affects the dealer's inventory. In line with this, we observe that non-matched trades in general, and client sales in the crisis period, are more expensive. Taken together, these two inferences suggest that dealers ration their balance-sheet space through their pricing behavior, but that they do so in a differentiated way across clients depending on the strength of the trading relationships with these clients.

5. Why do dealers give relationship discounts?

In this section, we test three hypotheses that could explain why dealers give relationship discounts to clients: "liquidity provision", "retaining high-value clients" and "information extraction". We find that dealers give larger discounts to clients to whom they often turn for liquidity provision, which allows them to manage their bond inventories more efficiently. We also find that discounts appear related to dealers' profit motives, as "high value" clients, who generate the bulk of dealers' trading profits, receive lower transaction costs. However, we find no evidence that information extraction is a driver of relationship discounts.

To test the three hypotheses, we identify the top 1% of liquidity-providing, high-value, and informed clients, and estimate the following regression:

$$TC_{bdct} = \gamma\, Qrel_{dct} + \delta\, Qrel_{dct} \times \iota_{dct} + \alpha\, \iota_{dct} + \mathbf{X}_{bdct}\beta + \mathbf{1}'\mu + \varepsilon_{bdct}, \tag{8}$$

where ι_{dct} is a dummy for the respective category of clients. If $\delta < 0$, the category of clients receives a larger discount than other clients supplying the same share of the dealer's trade volume ($Qrel$).

[17] The total reduction in transaction costs during the crisis $\gamma + \beta_1 + \beta_2 + \beta_3 = -159.38$ is significant at the 12% level.

Table 3: Relationship discounts and balance-sheet intensive trades

	Panel A: Pre-crisis		Panel B: Crisis	
	nomatch (1)	nomatch × sell (2)	nomatch (3)	nomatch × sell (4)
Qrel (γ)	-8.82	-11.86	-43.96	-39.63
	(7.57)	(7.03)	(46.61)	(63.86)
Qrel × nomatch (β_1)	-22.10^{***}	-22.58^{**}	-19.72	133.30
	(7.30)	(10.28)	(38.37)	(88.87)
Qrel × sell (β_2)		5.88		-9.51
		(8.16)		(66.43)
Qrel × nomatch × sell (β_3)		1.59		-243.50^{*}
		(13.45)		(142.60)
nomatch (α_1)	3.82^{***}	3.34^{***}	7.99^{***}	8.76
	(0.92)	(0.87)	(1.72)	(6.11)
sell (α_2)	-0.80	-1.77	36.20^{***}	40.98^{***}
	(1.43)	(1.67)	(3.79)	(9.11)
nomatch × sell (α_3)		0.95^{***}		-1.68
		(0.33)		(9.63)
$\gamma + \beta_1$	-30.91^{***}	-34.44^{**}	-63.67^{**}	93.62
$\gamma + \beta_2$		-5.98		-49.14
$\gamma + \beta_1 + \beta_2 + \beta_3$		-26.97^{*}		-159.38
Nobs	2.3m	2.3m	91.6k	91.6k
R^2	0.21	0.21	0.36	0.36

Notes: This Table shows the estimates from the regressions Eq. (6) (columns (1) and (3)) and Eq. (7) (columns (2) and (4)) for the pre-crisis and crisis period. Pre-crisis regressions (Panel A) include dealer-month, client-month, bond-month and industry-day fixed effects, and standard errors (shown in parentheses) are double clustered by month and dealer. For the crisis period (Panel B) regressions include dealer, client, bond and industry-day fixed effects, and standard errors are clustered by dealers. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%).

Table 4: Percent of trading volume by sector in different client categories

sector	top clients (1)	liquidity clients (2)	high-value clients (3)	informed clients (4)
Asset Manager	63.17	60.27	85.22	66.56
Bank	2.32	0.94	0.03	8.93
PFLDI	11.11	5.83	11.04	11.94
Broker	18.12	28.70	0.00	8.10
Hedge Fund	4.57	3.68	3.17	1.97
PTF	0.71	0.59	0.54	2.50

Notes: The table shows the percentage of trading volume by client sector (excluding other financials, non-financials and unclassified clients). 'Broker' includes brokers, executing and investing services, 'PFLDI' include pension funds, insurers and liability driven investors, 'PTF' are proprietary trading firms. The column 'top clients' shows the percentage by sector among the dealers' top-1% clients according to their $Qrel$ measure, the columns 'liquidity clients', 'high-value clients' and 'informed-clients' show the percentages among the top 1% liquidity-providing, profitable and informed clients (at a horizon of 5 days), respectively.

5.1. Hypothesis 1: "liquidity provision"

Relationships could be important for dealers to foster a convenient way to off-load bonds when they are faced with inventory imbalances or balance-sheet constraints. If so, we would expect the relationship discount to be more pronounced for such liquidity-providing clients. Given that balance-sheet intensive trades appear more costly for dealers, as shown in section 4, we now study whether clients who often supply liquidity to dealers receive larger discounts.

To identify liquidity provision by clients, we rely on a similar approach as in Choi et al. (2021). For each client sale absorbed by a dealer, we identify the clients who subsequently bought the bond from the dealer if there was such a trade *on the same day*. We measure the amount of liquidity provision by these clients as the total amount of purchases conducted in this way. We then sum the amount of liquidity provision over a 180-days window prior to the transaction. For each dealer at a given point in time, liquidity clients are then the top 1% clients according to this liquidity provision measure.

To provide some intuition on who the main liquidity-providing clients are, column 2 of Table 4 reports their sectoral composition. Asset managers are the most important liquidity-providing clients, followed by brokers. Together, these two groups account for roughly 90% of the trading volume of all liquidity-providing clients. The share of brokers in the population of liquidity-providing clients is more than 50% larger than their share in top relationship clients (column 1) and 3.5 times larger than their share in the overall population of client types (Panel C of Table 1). The outsized role of brokers as liquidity suppliers – despite their relatively small overall size – is intuitive.[18]

We now proceed by testing more formally whether dealers reward liquidity-providing clients by estimating Eq. (8). Panel A, column 1 of Table 5 shows that liquidity clients receive relationship discounts that are roughly 2.5 times as large (56.12 compared to 22.49) as for other clients. The discount for liquidity-providing relationship clients vs. other relationship clients is more than *six times larger* during the crisis period as seen from Panel B, column 4 (206.84 compared to 33.66). In fact, it was essentially only liquidity-providing clients receiving relationship discounts during this period of stress as the estimate on $Qrel$ is no longer significant. These results support the liquidity provision hypothesis. Clients who provide greater volumes of liquidity are a particularly valued set of relationship clients for dealers, especially during crisis times.

5.2. Hypothesis 2: "retaining high-value clients"

Another potentially important motive for dealers to offer a discount is to attract higher trading volume and, thereby, generate larger profits. By offering better prices to relationship clients, the dealer may be able to generate more business from those preferred customers. Thus, dealers may have an incentive to keep those clients loyal, benefiting from a larger trading volume over time. In addition, offering these clients a discount in corporate bond trading may also attract larger volumes from the same clients in other asset classes. Thus, similar to a shop on the high street offering discounts to loyal and profitable clients, dealers might offer a discount to certain groups of clients in order to make larger profits over the long run.

To investigate whether dealers treat high-value clients differently than other clients, we compute the total trading profits of each dealer from each of her clients. Dealer profits from each

[18] We did a robustness test excluding brokers from our sample and the main findings were unchanged. These results are excluded for brevity.

Table 5: Transaction costs for high-value clients, liquidity clients and informed clients

	Panel A: Pre-crisis			Panel B: Crisis		
	Liquidity client	High-value client	Informed client	Liquidity client	High-value client	Informed client
	(1)	(2)	(3)	(4)	(5)	(6)
Qrel (γ)	-22.49^{**}	-28.35^{**}	-27.81^{***}	-33.66	-73.61^{**}	-64.06^{**}
	(8.10)	(10.22)	(9.64)	(26.48)	(33.54)	(27.85)
Qrel × client type (δ)	-33.63^{**}	-13.59	-16.53	-173.20^{***}	9.44	31.52
	(14.69)	(11.77)	(46.29)	(40.17)	(46.90)	(92.31)
client type	-0.46	2.59^{**}	-0.77	3.09	6.78^{**}	1.49
	(0.77)	(1.22)	(1.16)	(2.04)	(3.20)	(7.13)
$\gamma + \delta$	-56.12^{***}	-41.94^{***}	-44.34	-206.84^{***}	-64.17	-32.54
Nobs	2.3m	2.3m	2.3m	91.6k	91.6k	91.6k
R^2	0.21	0.21	0.21	0.36	0.36	0.36

Notes: This table shows the results from fitting the regression model:

$$TC_{bdct} = \gamma\, Qrel_{dct} + \delta\, Qrel_{dct} \times \iota_{dct} + \alpha\, \iota_{dct} + \mathbf{X}'_{bdct}\beta + \mathbf{1}'\mu + \varepsilon_{bdct},$$

where TC_{bdct} are transaction costs, $Qrel_{dct}$ is our relationship measure, ι_{dct} is a dummy variable taking the value 1 if the client is in dealer's top 1% of liquidity-providing clients, high-value clients and informed clients, respectively, in the 180 days prior to the transaction. The controls $\mathbf{X}$ include dummies for matched trades, client sales, trades executed on a regulated market or multilateral trading facility, the log of the traded amount and the lagged squared return on the benchmark. The vector μ contains fixed effects. Pre-crisis regressions (Panel A) include dealer-month, client-month, bond-month and industry-day fixed effects, and standard errors (shown in parentheses) are double clustered by month and dealer. For the crisis period (Panel B) regressions include dealer, client, bond and industry-day fixed effects, and standard errors are clustered by dealers. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%).

Table 6: Dealers' profits from top and non-top clients

	Avg total profit (in £ m)	Avg total volume (in £ bn)	Avg number of clients	Avg profit per client (in £ k)	Avg volume per client (in £ m)
non-top	11.27	26.00	463	24.38	56.18
top	4.77	12.47	11	449.41	1174.91

Notes: The table shows profit and volume statistics across dealers for their top and non-top clients aggregated over the pre-crisis sample. The first three columns are averaged across dealers. The last two are averaged for all clients within each dealer, and then averaged across dealers.

client are calculated by summing the product of transaction cost and trade size over all trades between the client and the dealer in the 180 days window used to calculate $Qrel$. The intuition is that *cost* for a client is *profit* for a dealer. For consistency with the other hypotheses, we define 'high-value clients' as the top 1% of clients in total dealer's profits at a given point in time.

Again, it is useful to inspect basic descriptive statistics about high-value clients before proceeding with the formal tests. To this end, Table 6 shows the trading profit across all dealers for both top and non-top clients. Dealers have on average 11 top and 463 non-top clients and make on average around £5 million and £11 million profit over the pre-crisis sample period from those two groups, respectively. These facts show that only 11 clients (around 2% of all clients) account for around a third of dealers' total profits. Asset managers are the largest category of high-value clients and account for more than 85% of all such clients as seen from column 3 of Table 4. Importantly, the average profit made on a top client is more than 18 times larger than the average profit made on a non-top client. This fact suggests a strong incentive for dealers to focus on top clients and keep them as loyal customers, by offering them more competitive prices.

We next test whether high-value clients that contribute the most to dealers' profits receive larger discounts by defining a dummy for high-value clients when estimating Eq. (8). Table 5 shows that high-value clients receive relationship discounts, but these are only marginally greater than for other clients supplying the same trading volume to the dealer: δ is not statistically significant. In subsection 5.4 below, we test the "retaining high-value clients" hypothesis by taking into account also liquidity-providing and information motives and find that the estimate for high-value clients is significant in those tests.

5.3. Hypothesis 3: "information extraction"

Besides liquidity and profit maximization motives, dealers might also wish to build relationships with clients, whose order flow provides valuable trading signals (Pintér et al., 2022a). In that case, we should find that the relationship metric has a more pronounced impact on transaction costs for informed clients.

Following Kondor and Pintér (2022), we measure the informativeness of trades in terms of subsequent price returns, and classify clients as informed clients if their trades consistently predict future returns. Thus, for each trade, we compute the h-period-ahead directional return:

$$r_t(h) = [\log(p^*_{t+h}) - \log(p^*_t)] \times D_t, \tag{9}$$

where p^*_t is the benchmark price at time t, D_t is the direction of the client's trade (1 for a purchase and -1 for a sale), p^*_{t+h} is the end-of-day benchmark price h days after the trade. We then aggregate these directional returns into a performance metric $perf_{ct}$ that summarises the degree to which the client is informed, from the perspective of the dealer. Specifically, we compute the average return of all client trades with the dealer over a past 180-days window:

$$perf_{dct}(h) = \frac{1}{N_{dc}} \sum_{\tau \in \mathcal{T}_{dc}(t-h-180,\ t-h)} r_\tau(h), \tag{10}$$

where $\mathcal{T}_{dc}(t-h-180,\ t-h)$ is the set of all N_{dc} trades between the client and the dealer over the previous 180 days lagged by the return horizon h.[19] We scale $perf_{dct}(h)$ by its standard deviation in order to identify clients whose trades consistently (i.e., with little volatility) predict future returns, which gives the scaled measure $\widehat{perf_{dct}}(h)$. For each dealer we then define 'informed clients' as the top 1% clients according to their scaled performance, $\widehat{perf_{dct}}(h)$. We focus on results with $h = 5$ in the main text, but the results for other horizons $h \in \{1, 20, 30\}$ are

[19]We use equally-weighted rather than volume-weighted averages as informed traders may choose to trade in smaller sizes by splitting their trades (Kondor and Pintér, 2022).

generally similar as we show in the Online Appendix, Table A.6 and Table A.7.

Table 4, column 4 reports the sectoral composition of informed clients. It shows that asset managers account for the highest share of informed trading volume. Pension funds and other liability-driven investors (PFLDI) are the second largest group. The share of principal trading firms (PTFs) is more than 3.5 times larger than their share in top relationship clients (column 1, Table 4). Hedge funds account for a smaller share of informed-client trading volume.

We find no evidence in support of the information hypothesis: more informed clients do not receive larger relationship discounts. Although the estimate of δ from Eq. (8) is not statistically significant in Table 5, it is even *positive* during the crisis period. This indicates that dealers may actually charge informed clients *more* than other clients supplying a similar share of dealer's trading volume during crisis periods.

These additional transaction charges could reflect dealers' aversion to trading against informed clients since dealers could suffer a loss by taking the opposite position to an informed client. Another possible explanation for why dealers do not offer discounts to informed clients could be that post-GFC regulations limited the scope to profit from trading on such information by deterring proprietary trading.

5.4. Horse race between the three hypotheses

The same client can be important to a dealer for more than one reason as shown in the Venn diagram of clients in Figure 2. In particular, there is a notable intersection between liquidity and high-value clients. Informed clients are a rather separate group and have little intersection with the other two groups of clients.

We next run a horse race between the three main hypotheses by including dummies for all three client types in a single regression, and accounting for the large intersection between high-value and liquidity-providing clients.[20] Table 7 shows that the results are qualitatively the same as in the individual regressions. The only major difference is that after accounting for the

[20]We do not put dummies for all possible intersections in the regression since there are very few observations in some groups as seen from Figure 2.

Figure 2: Intersection between liquidity, high-value, and informed clients.

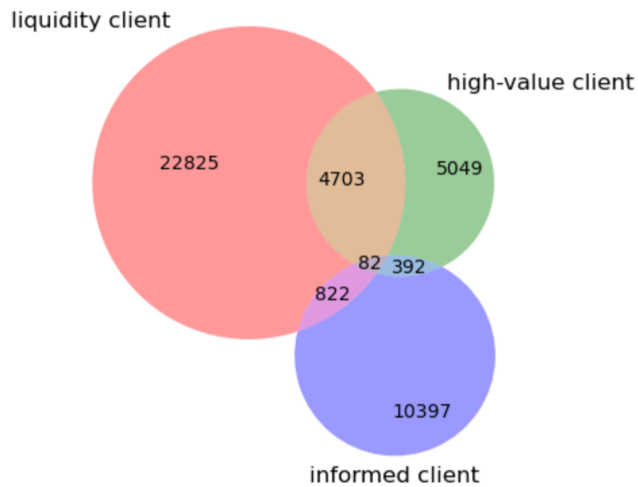

Notes: The figure displays the intersection between liquidity, high-value, and informed clients. The numbers illustrate the number of observations in each category over the regression sample.

overlap with other types of clients, the estimate on the high-value clients dummy becomes also statistically significant in the pre-crisis period. This result shows that after isolating liquidity- and information-related motives, high-value clients also get a relationship discount, which lends additional support to our second hypothesis.

5.5. Additional results and robustness

Finally, we confirm our main results in a series of additional robustness tests in the Online Appendix. Besides providing additional descriptive statistics, we investigate the impact of sensible modifications to our modeling choices.

We find that the main results are qualitatively similar for different horizons used in the calculation of the relationship metric $Qrel$ and if we focus on transactions for which the benchmark price is observed more frequently than three times a day (Table A.4). The results are also similar if we include only large dealers or only large clients, or exclude the largest clients (Table A.5). The major results for the three hypotheses are robust to using different cutoffs for top clients (instead of 1%): Table A.6 and Table A.7. The results are also generally similar if we use

Table 7: Horse race between the three hypotheses

	Panel A: Pre-crisis		Panel B: Crisis	
	(1)	(2)	(3)	(4)
$Qrel\,(\gamma)$	-22.10^{**}	-21.98^{**}	-34.25	-36.16
	(8.06)	(8.04)	(33.81)	(30.74)
$Qrel \times lc\,(\beta_{lc})$	-36.68^{**}	-37.68^{**}	-172.20^{***}	-176.20^{***}
	(15.87)	(15.85)	(54.93)	(49.36)
$Qrel \times hv\,(\beta_{hv})$	-33.39^{*}	-33.10^{*}	-4.75	-51.51
	(19.09)	(18.99)	(55.21)	(54.94)
$Qrel \times ic\,(\beta_{ic})$	-14.15	-16.86	-5.60	31.88
	(46.33)	(46.99)	(97.44)	(105.30)
$Qrel \times hv \times lc\,(\beta_{hvlc})$	37.74	40.41	105.20	168.50
	(24.86)	(24.66)	(236.30)	(167.10)
lc	-0.41	-0.35	4.53	3.42
	(0.79)	(0.77)	(2.78)	(2.49)
hv	4.03^{*}	3.99^{*}	4.05	7.68^{*}
	(1.99)	(2.00)	(4.96)	(4.39)
ic	-0.86	-0.87	6.50	0.98
	(1.15)	(1.16)	(7.25)	(7.26)
$hv \times lc$	-2.37	-2.39	-0.57	-4.00
	(1.79)	(1.77)	(8.48)	(6.46)
Nobs	2.3m	2.3m	91.6k	91.6k
R^2	0.20	0.21	0.31	0.36
$\gamma + \beta_{lc}$	-58.78^{***}	-59.66^{***}	-206.50^{***}	-212.41^{***}
$\gamma + \beta_{hv}$	-55.49^{***}	-55.09^{***}	-39.00	-87.67^{**}
$\gamma + \beta_{ic}$	-36.26	-38.85	-39.85	-4.28
$\gamma + \beta_{lc} + \beta_{hv} + \beta_{hvlc}$	-54.43^{***}	-52.36^{***}	-106.00	-95.46
Fixed effects				
dealer × month	Yes	Yes		
client × month	Yes	Yes		
bond × month	Yes	Yes		
day × industry		Yes		Yes
dealer			Yes	Yes
client			Yes	Yes
bond			Yes	Yes

Notes: This table shows the result from fitting the regression model:

$$TC_{bdct} = \gamma\, Qrel_{dct} + \beta_{lc}\, Qrel_{dct} \times lc_{dct} + \beta_{hv}\, Qrel_{dct} \times hv_{dct} + \beta_{ic}\, Qrel_{dct} \times ic_{dct} + \beta_{hvlc}\, Qrel_{dct} \times hv_{dct} \times lc_{dct} + \alpha_{lc}\, lc_{dct} + \alpha_{hv}\, hv_{dct} + \alpha_{ic}\, ic_{dct} + \alpha_{hvlc}\, hv_{dct} \times lc_{dct} + \mathbf{X}'_{bdct}\beta + \mathbf{1}'\mu + \varepsilon_{bdct},$$

where lc_{dct}, hv_{dct} and ic_{dct} are dummy variables taking the value 1 if the client is in dealer's top 1% of liquidity-providing clients, high-value clients and informed clients, respectively, in the 180 days prior to the transaction. The controls $\mathbf{X}$ include dummies for matched trades, client sales, trades executed on a regulated market or multilateral trading facility,, the log of the traded amount and the lagged squared return on the benchmark. The vector μ contains fixed effects. Standard errors (shown in parentheses) in the pre-crisis regressions (Panel A) are double clustered by month and dealer, and those in the crisis period (Panel B) are clustered by dealers. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%)

other time horizons for the information hypothesis. Finally, in unreported results, we find that "captured" clients who rely on only one dealer, do not get significant relationship discounts.

6. Conclusion

Drawing on regulatory and proprietary data sets, we document several new findings on dealer-client relationships in the corporate bond market. Our results show that clients with a stronger relationship with a dealer receive better prices. Top relationship clients pay approximately half the transaction costs faced by the median client, which amounts to annual cost savings of around £1.3m per client on average. These relationship benefits were particularly important during the dash-for-cash episode in March 2020, when the absolute reduction in trading costs for relationship clients more than doubled.

Our results point to two major economic mechanisms that could explain the significant reduction in transaction costs for relationship clients. First, dealers value clients to whom they can turn for liquidity provision. Second, dealers earn much higher profits from relationship clients than other clients, which creates a strong incentive for dealers to keep these counterparties as loyal customers.

Our findings show that the OTC market structure centred around dealer intermediation in corporate bonds proved largely resilient for relationship clients during the COVID-19 shock. The results also suggest that the OTC market structure might be more sustainable in the presence of relationship benefits, as they could help dealers to operate with smaller inventories, which are cheaper to maintain. On the other hand, relationship benefits are by their nature reserved for particular clients. Clients not able to build meaningful relationships with dealers pay significantly larger transaction costs, especially during stress times. These findings have implications for the debate about alternative market structures (dealer-centred OTC vs all-to-all), which has intensified in the aftermath of the COVID-19 crisis.

References

Abudy, M. M. and A. Wohl (2018). Corporate bond trading on a limit order book exchange. *Review of Finance 22*(4), 1413–1440.

Adrian, T., N. Boyarchenko, and O. Shachar (2017). Dealer balance sheets and bond liquidity provision. *Journal of Monetary Economics 89*, 92–109.

Bao, J., M. O'Hara, and A. Zhou (2018). The Volcker rule and corporate bond market making in times of stress. *Journal of Financial Economics 130*(1), 95–113.

Biais, B. and R. Green (2019). The microstructure of the bond market in the 20th century. *Review of Economic Dynamics 33*, 250–271.

Choi, J., Y. Huh, and S. S. Shin (2021). Customer liquidity provision: Implications for corporate bond transaction costs. *Management Science forthcoming*.

Czech, R. and G. Pintér (2020). Informed trading and the dynamics of client-dealer connections in corporate bond markets. *Bank of England Working Paper No. 895*.

Di Maggio, M., A. Kermani, and Z. Song (2017). The value of trading relations in turbulent times. *Journal of Financial Economics 124*(2), 266–284.

Dick-Nielsen, J. and M. Rossi (2019). The cost of immediacy for corporate bonds. *The Review of Financial Studies 32*(1), 1–41.

Goldstein, M. A. and E. S. Hotchkiss (2020). Providing liquidity in an illiquid market: Dealer behavior in US corporate bonds. *Journal of Financial Economics 135*(1), 16–40.

Hau, H., P. Hoffmann, S. Langfield, and Y. Timmer (2021). Discriminatory pricing of over-the-counter derivatives. *Management Science 67*(11), 6660–6677.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020). Relationship trading in over-the-counter markets. *The Journal of Finance 75*(2), 683–734.

Hendershott, T. and A. Madhavan (2015). Click or call? Auction versus search in the over-the-counter market. *The Journal of Finance 70*(1), 419–447.

Kondor, P. and G. Pintér (2022). Clients' connections: Measuring the role of private information in decentralized markets. *The Journal of Finance 77*(1), 505–544.

Lee, T. and C. Wang (2018). Why trade over-the-counter? when investors want price discrimination. *Working Paper*.

Maskin, E. and J. Riley (1984). Monopoly with incomplete information. *The RAND Journal of Economics 15*(2), 171–196.

O'Hara, M. and A. Zhou (2021). The electronic evolution of corporate bond dealers. *Journal of Financial Economics 140*(2), 368–390.

O' Hara, M., Y. Wang, and A. Zhou (2018). The execution quality of corporate bonds. *Journal of Financial Economics 130*(2), 308–326.

O'Hara, M., A. C. Rapp, and A. Zhou (2023). The value of value investors. *Working paper*.

Pintér, G., C. Wang, and J. Zou (2022a). Information chasing versus adverse selection. *Bank of England Working Paper No. 971*.

Pintér, G., C. Wang, and J. Zou (2022b). Size discount and size penalty: Trading costs in bond markets. *Bank of England Working Paper*.

Plante, S. (2018). *Should Corporate Bond Trading be Centralized? Theory and Evidence*. Doctoral, University of Pennsylvania.

Vogel, S. (2018). When to introduce electronic trading platforms in over-the-counter markets? *Working Paper*.

Wittwer, M. and J. Allen (2021). Centralizing over-the-counter markets? *SSRN Discussion Paper*.

Online Appendix

This supplementary Online Appendix provides some additional summary statistics and performs several robustness tests.

Summary statistics on currency, country, and industry. The vast majority of bonds in our sample are denominated in EUR and USD, and only a small share is denominated in GBP. EUR-denominated bonds account for more than a half of all trades, whereas GBP-denominated bonds are only around 8% as seen from columns 1 and 2 of Table A.1.

Around a third of bonds is issued by European Union (EU) entities and a similar fraction is issued by US entities as seen from column 3. Trading, however is more concentrated in EU-issued bonds as seen in column 4. Less than 12% of bonds are issued by UK firms. In terms of industries, more than half of the bonds are issued by financial entities, whereas all other industries account for less than 7% of total volume each, as seen from Table A.2.

Alternative time horizons. Our main results are robust to other horizons used to calculate the relationship metric $Qrel$ (90 days window with a lag of 1 day) as shown in columns 1–2 and 5–6 of Table A.4.

Potential benchmark price staleness. In our main analysis, we use benchmark prices from MarketAxess observed three times a day and even more frequently during the stress episode (1 to 18 March). To address the concern that benchmark prices might be stale, i.e. not reflecting the market price close to the trade time, we did a robustness test by restricting the sample to trades for which there was a recently observed benchmark price (less than an hour before the trade timestamp). Columns 3–4 and 7–8 of Table A.4 show that the relationship discount is still strongly significant in that sample, and is even larger in size compared to our baseline estimate from Table 2.

Excluding small and large dealers and clients. One concern might be that relationship discounts are driven by a particular set of dealers or clients. For example, the right panel of Figure 1 shows that 2 clients are particularly large and account for just over 5% of all dealer-to-client trading volume. The left panel of the figure illustrates that our sample includes a tail of smaller dealers who may be more inclined to provide discounts than larger dealers. Alternatively, relationship discounts may be driven by the largest clients in our sample, if they have a higher bargaining power. Client-time and dealer-time fixed effects in our main specification already address some of these concerns, but we also performed a robustness test by dropping the largest clients, the smallest dealers and the smallest clients from the sample.

In columns 1–2 and 7–8 of Table A.5, we run our main regression for a subsample of trades that excludes the largest 2 clients. In columns 3–4 and 9–10 of the table, we keep only the largest 500 clients, who account for about 80% of total dealer-to-client trading volume. Similarly, in columns 5–6 and 11–12 of Table A.5 we keep only the largest 15 dealers, who also account for about 80% of total dealer-to-client trading volume. Our relationship discount estimates remain large and significant across all of these settings, which illustrates that the main results are not driven by the 2 largest clients, by small dealers, or by small clients.

Alternative cutoffs to identify client types. Our results are also robust to using other cutoffs than the top 1% to identify liquidity, high-value, and informed clients as shown in Table A.6 and Table A.7. Namely, the results are robust to using the top 10% or the top 5%.

For the information hypothesis, the results are generally similar if we use other horizons than 5 days: 1 day, 20 days, or 30 days as shown in Table A.6 and Table A.7. In some specifications (20 days or 30 days), the estimate for informed clients β_{ic} is positive and significant, indicating that such clients are charged larger costs.

Clients that rely on one dealer only. We also analyzed clients with only one dealer vs. those with more than one dealer, since dealers could potentially extract "monopoly rents" from the former group of clients. Clients with one dealer might find it harder to switch to another dealer compared to other clients, and dealers might thus not give relationship discounts to such "captured" clients.

In line with this conjecture, we find that the estimate on $Qrel$ in our main regression Eq. (1) is insignificant for clients who trade with one dealer only (these are only 2% of all trades). These results are excluded for brevity but are available on request.

Table A.1: Share of bonds and trading volume by currency and issuer country

	Share (in %) by			Share (in %) by	
Currency	#Bonds (1)	Volume (2)	Issuer country	#Bonds (3)	Volume (4)
EUR	29.18	51.02	EU	34.66	46.36
USD	62.97	40.85	United States	32.12	13.72
GBP	7.85	8.13	United Kingdom	11.55	12.39
			RoW	21.67	27.53

Notes: This table shows the share of bonds issued in different currency (columns 1 and 2) and by issuer's country (columns 3 and 4). The numbers are based on the dealer-client transaction sample from 3 Jan 2018 to 18 March 2020. The information on issued currency and issuer country comes from S&P Capital IQ. "RoW" is the rest of the world, excluding the EU, US, and the UK.

Table A.2: Share of bonds and trading volume by issuer industry

	Share (in %) by	
Sector	#Bonds	Volume
Financials	54.53	58.11
Communication Services	4.42	6.48
Energy	5.38	6.18
Industrials	6.96	5.14
Utilities	5.52	4.32
Consumer Discretionary	3.97	3.80
Materials	3.69	3.41
Consumer Staples	3.34	3.16
Real Estate	1.86	3.03
Health Care	3.32	2.27
Information Technology	2.31	1.59

Notes: This table shows the share of bonds by issuer's industry sector classification provided by S&P Capital IQ. The numbers are based on the dealer-client transaction sample from 3 Jan 2018 to 18 March 2020.

Table A.3: MarketAxess CP+ and inter-dealer price comparison

	Distribution of time-distance (in hours)								tc (in bps)	
	#Obs	mean	std	1%	25%	50%	75%	99%	mean	median
Panel A: MarketAxess CP+										
pre-crisis	3.2m	2.31	2.51	0.03	0.99	1.90	2.84	16.00	8.43	4.37
crisis	106k	0.29	1.47	0.00	0.00	0.01	0.02	6.76	41.28	11.78
Panel B: Inter-dealer price										
pre-crisis	2.6m	8.01	8.65	0.00	0.81	3.31	17.86	23.80	7.79	2.71
crisis	106k	7.16	8.34	0.00	0.67	2.66	16.47	23.74	30.73	7.70

Notes: This table shows the distribution of time differences (in hours) between the time of the transaction between a dealer and client, and the time when the benchmark price was observed prior to that transaction. The last two columns on the right show the average and median transaction costs based on the respective benchmark. Panel A displays MarketAxess CP+ benchmark prices, Panel B uses inter-dealer prices. The numbers are based on our transaction sample running from 3 Jan 2018 to 18 March 2020, split into pre-crisis (before 1 March) and crisis sub-periods, but restricted to observations where the benchmark price is less than 24 hours old.

Table A.4: Relationship metric window size and lag, and benchmark price staleness

	Panel A: Pre-crisis				Panel B: Crisis			
	Qrel ($w = 90, l = 1$)		$\Delta t \leq 1h$		Qrel ($w = 90, l = 1$)		$\Delta t \leq 1h$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Qrel$	-26.72^{***}	-29.99^{***}	-31.80^{***}	-35.74^{**}	-56.51^{**}	-64.12^{**}	-55.67^{*}	-72.20^{**}
	(8.83)	(9.42)	(10.81)	(14.16)	(26.89)	(26.25)	(28.22)	(29.03)
$ismatch$	-3.45^{***}	-3.65^{***}	-3.26^{***}	-3.40^{***}	-7.67^{***}	-6.43^{***}	-7.84^{***}	-6.57^{***}
	(0.86)	(0.74)	(0.88)	(0.95)	(2.01)	(1.76)	(2.09)	(1.80)
$sell$	-0.46	0.37	0.14	1.05	36.20^{***}	39.60^{***}	30.49^{***}	27.47^{***}
	(1.35)	(1.55)	(1.32)	(1.70)	(3.79)	(4.56)	(3.46)	(3.81)
$logQ$	0.62^{***}	0.71^{***}	0.81^{***}	1.00^{***}	0.53^{*}	1.23^{***}	0.57^{**}	1.14^{***}
	(0.06)	(0.06)	(0.09)	(0.13)	(0.29)	(0.33)	(0.26)	(0.27)
MTF	2.31^{**}	2.63^{**}	3.13^{***}	5.38^{***}	0.50	0.08	0.74	0.15
	(0.84)	(1.08)	(0.90)	(1.75)	(2.00)	(2.27)	(2.04)	(2.38)
$r^2_{b,t-1}$	-0.35		0.47		-0.95		-0.96	
	(0.28)		(0.41)		(0.76)		(0.80)	
Nobs	2.6m	2.1m	573.5k	326.3k	91.6k	84.5k	88.0k	78.5k
R^2	0.21	0.35	0.33	0.54	0.36	0.48	0.37	0.50
Fixed effects								
dealer × month	Yes		Yes					
client × month	Yes	Yes	Yes	Yes				
bond × month	Yes		Yes					
day × industry	Yes		Yes		Yes		Yes	
dealer × day		Yes		Yes		Yes		Yes
bond × day		Yes		Yes		Yes		Yes
dealer					Yes		Yes	
client					Yes	Yes	Yes	Yes
bond					Yes		Yes	

Notes: This table shows the results of our baseline regression (Eq. (1)) based on alternative time horizons used to calculate our relationship metric, $Qrel$, and for a sample of trades with recently observed benchmark prices. In columns 1, 2, 5 and 6 we measure $Qrel$ as the share of a client's trading volume in a dealer's total trading volume over a past window of 90 days, lagged by one day. In columns 3, 4, 7 and 8 we measure the dependent variable, transaction costs, based on trades where the benchmark price is observed less than 1 hour prior to the trade. Panel A presents the results for the pre-crisis period from 3 Jan 2018 to 29 Feb 2020, and Panel B the results for the crisis period from 1 to 18 March 2020. The rest of the variables remain as defined in the main text: *match*, an indicator variable equal to one if the dealer offsets the trade with other trades executed at the same instant and in the opposite direction; *client sell*, an indicator variable equal to one if the client is selling; MTF, an indicator variable equal to one if the trade is executed on a regulated market (e.g., London Stock Exchange) or a multilateral trading facility (e.g., MarketAxess); $logQ$, the natural logarithm of the trade size measured in GBP and r^2_{bt-1}, the MarketAxess benchmark's lagged squared daily return. Standard errors (shown in parentheses) are clustered at the dealer and month level for the pre-crisis period and at the dealer level for the crisis period. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%).

Table A.5: Excluding large and small clients and dealers

	Panel A: Pre-crisis						Panel B: Crisis					
	excl. largest 2 clients		keep largest 500 clients		keep largest 15 dealers		excl. largest 2 clients		keep largest 500 clients		keep largest 15 dealers	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Qrel	−33.03**	−39.09***	−20.46**	−24.09**	−76.58**	−87.04**	−80.03***	−92.10***	−50.96*	−61.22**	−114.50***	−116.00**
	(12.40)	(13.20)	(7.79)	(8.48)	(25.74)	(30.81)	(29.29)	(33.41)	(27.93)	(28.39)	(37.80)	(46.92)
ismatch	−3.53***	−3.78***	−2.96***	−3.16***	−4.32***	−4.99***	−7.75***	−6.18***	−7.79***	−5.64***	−6.43***	−8.13**
	(0.89)	(0.77)	(0.75)	(0.71)	(1.42)	(1.26)	(2.07)	(1.81)	(2.15)	(2.10)	(1.92)	(3.11)
sell	−0.90	−0.04	−0.66	0.19	−0.53	0.52	36.34***	40.00***	32.32***	33.90***	31.60***	33.83***
	(1.45)	(1.67)	(1.39)	(1.56)	(1.47)	(1.66)	(4.01)	(4.82)	(3.51)	(3.94)	(4.36)	(5.47)
logQ	0.58***	0.69***	0.69***	0.77***	0.76***	0.85***	0.52	1.23***	0.52*	1.09***	0.94**	1.86***
	(0.06)	(0.06)	(0.07)	(0.07)	(0.06)	(0.07)	(0.31)	(0.38)	(0.29)	(0.35)	(0.33)	(0.40)
MTF	1.97**	2.44**	2.04**	2.12*	1.51	1.45	0.33	0.14	0.42	−0.06	0.30	−0.26
	(0.91)	(1.13)	(0.89)	(1.11)	(0.94)	(1.24)	(2.08)	(2.39)	(2.30)	(2.31)	(2.44)	(2.91)
$r^2_{b,t-1}$	−0.41		−0.38		−0.49		−0.72		−0.93		−1.22	
	(0.30)		(0.42)		(0.29)		(0.76)		(0.98)		(1.00)	
Nobs	2.2m	1.7m	1.8m	1.4m	1.7m	1.3m	85.9k	78.8k	70.8k	61.2k	66.3k	57.2k
R^2	0.21	0.36	0.17	0.33	0.21	0.37	0.36	0.49	0.36	0.50	0.37	0.50
Fixed effects												
dealer × month	Yes		Yes		Yes							
client × month	Yes	Yes	Yes	Yes	Yes	Yes						
bond × month	Yes		Yes		Yes							
day × industry	Yes		Yes		Yes		Yes		Yes		Yes	
dealer × day		Yes		Yes		Yes		Yes		Yes		Yes
bond × day		Yes		Yes		Yes		Yes		Yes		Yes
dealer							Yes		Yes		Yes	
client							Yes	Yes	Yes	Yes	Yes	Yes
bond							Yes		Yes		Yes	

Notes: This table shows the results of our baseline regression (Eq. (1)) based on a restricted sample by excluding certain sets of dealers or clients. In columns 1, 2, 7 and 8 we exclude the largest two clients based on their share of trading volume over the sample period as shown in right Panel of Figure 1. In columns 3, 4, 9 and 10, we restrict our sample to trades by the largest 500 clients. In columns 5, 6, 11 and 12, we restrict the sample to trades by the top 15 dealers based on their share of trading volume as shown in the left Panel of Figure 1. Panel A shows the results for the pre-crisis period from 3 Jan 2018 to 29 Feb 2020, Panel B the results for the crisis period from 1 to 18 March 2020. Standard errors (shown in parentheses) are clustered at the dealer and month level for the pre-crisis period and at the dealer level for the crisis period. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%).

Table A.6: Hypotheses horse-race – pre-crisis

	Panel A: top 10%				Panel B: top 5%				Panel C: top 1%			
horizon	1	5	20	30	1	5	20	30	1	5	20	30
$Qrel\,(\gamma)$	-17.50^{***}	-17.50^{***}	-17.53^{***}	-17.54^{***}	-17.97^{***}	-17.97^{***}	-17.98^{***}	-17.99^{***}	-21.99^{**}	-21.98^{**}	-22.00^{**}	-22.07^{**}
	(4.89)	(4.89)	(4.90)	(4.90)	(4.86)	(4.86)	(4.86)	(4.87)	(8.06)	(8.04)	(8.08)	(8.16)
$Qrel \times ic\,(\beta_{ic})$	−6.01	−3.84	−9.76	−11.71	−10.25	−10.55	−19.99	−12.30	−14.75	−16.86	−24.34	−8.63
	(16.02)	(15.96)	(17.61)	(17.37)	(30.69)	(31.55)	(30.51)	(25.54)	(41.10)	(46.99)	(44.81)	(33.94)
$Qrel \times pc\,(\beta_{pc})$	-27.75^{*}	-27.87^{*}	-27.38^{*}	-27.28^{*}	-33.41^{*}	-33.50^{*}	-33.15^{*}	-33.31^{*}	-33.07^{*}	-33.10^{*}	−32.95	-32.92^{*}
	(14.32)	(14.29)	(14.42)	(14.39)	(16.62)	(16.63)	(16.66)	(16.62)	(18.93)	(18.99)	(19.07)	(19.02)
$Qrel \times lc\,(\beta_{lc})$	-82.27^{*}	-82.30^{*}	-81.99^{*}	-82.26^{*}	-75.29^{*}	-75.32^{*}	-75.01^{*}	-75.28^{*}	-37.57^{**}	-37.68^{**}	-37.60^{**}	-37.32^{**}
	(45.81)	(45.77)	(45.72)	(45.89)	(41.00)	(40.85)	(40.94)	(41.25)	(15.88)	(15.85)	(15.88)	(15.81)
$Qrel \times pc \times lc\,(\beta_{pclc})$	68.14	68.11	67.59	68.03	60.10	60.06	59.51	60.02	40.27	40.41	40.25	40.01
	(49.95)	(49.91)	(49.77)	(50.11)	(40.18)	(40.02)	(40.04)	(40.41)	(24.71)	(24.66)	(24.61)	(24.59)
ic	−0.18	−0.37	−0.25	−0.33	−0.28	−0.50	−0.41	−0.62	−0.71	−0.87	−0.83	−0.79
	(0.45)	(0.50)	(0.54)	(0.57)	(0.79)	(0.77)	(0.88)	(0.85)	(1.19)	(1.16)	(1.31)	(0.97)
pc	2.48^{*}	2.49^{*}	2.48^{*}	2.49^{*}	3.38^{**}	3.40^{**}	3.38^{**}	3.39^{**}	3.98^{*}	3.99^{*}	3.99^{*}	3.96^{*}
	(1.23)	(1.23)	(1.23)	(1.23)	(1.60)	(1.59)	(1.60)	(1.60)	(1.99)	(2.00)	(1.99)	(2.00)
lc	-0.88^{**}	-0.87^{**}	-0.88^{**}	-0.87^{**}	−0.76	−0.75	−0.75	−0.75	−0.35	−0.35	−0.35	−0.37
	(0.36)	(0.35)	(0.35)	(0.36)	(0.48)	(0.47)	(0.48)	(0.48)	(0.78)	(0.77)	(0.77)	(0.77)
$pc \times lc$	−0.63	−0.64	−0.63	−0.63	−1.41	−1.42	−1.39	−1.41	−2.37	−2.39	−2.39	−2.33
	(1.12)	(1.12)	(1.12)	(1.12)	(1.39)	(1.38)	(1.40)	(1.40)	(1.77)	(1.77)	(1.76)	(1.76)
$\gamma + \beta_{ic}$	−23.51	−21.35	−27.29	−29.25	−28.22	−28.51	−37.97	−30.29	−36.73	−38.85	−46.35	−30.70
$\gamma + \beta_{pc}$	-45.25^{***}	-45.38^{***}	-44.91^{***}	-44.82^{***}	-51.38^{***}	-51.47^{***}	-51.13^{***}	-51.31^{***}	-55.05^{***}	-55.09^{***}	-54.95^{***}	-54.99^{***}
$\gamma + \beta_{lc}$	-99.77^{**}	-99.80^{**}	-99.52^{**}	-99.80^{**}	-93.26^{**}	-93.29^{**}	-92.99^{**}	-93.28^{**}	-59.56^{***}	-59.66^{***}	-59.60^{***}	-59.39^{***}
$\gamma + \beta_{lc} + \beta_{pc} + \beta_{pclc}$	-59.38^{***}	-59.56^{***}	-59.30^{***}	-59.05^{***}	-66.58^{***}	-66.74^{***}	-66.63^{***}	-66.57^{***}	-52.35^{***}	-52.36^{***}	-52.30^{***}	-52.30^{***}

Notes: This table shows the result from fitting the regression model:

$$TC_{bdct} = \gamma\, Qrel_{dct} + \beta_{lc}\, Qrel_{dct} \times lc_{dct} + \beta_{hv}\, Qrel_{dct} \times hv_{dct} + \beta_{ic}\, Qrel_{dct} \times ic_{dct} + \beta_{hvlc}\, Qrel_{dct} \times hv_{dct} \times lc_{dct} +$$
$$\alpha_{lc}\, lc_{dct} + \alpha_{hv}\, hv_{dct} + \alpha_{ic}\, ic_{dct} + \alpha_{hvlc}\, hv_{dct} \times lc_{dct} + \mathbf{X}'_{bdct}\beta + \mathbf{1}'\mu + \varepsilon_{bdct},$$

for the pre-crisis period from 3 Jan 2018 to 29 Feb 2020 where lc_{dct}, hv_{dct} and ic_{dct} are dummy variables taking the value 1 if the client is in dealer's top 10% (Panel A), 5% (Panel B) or 1% (Panel C) of liquidity providing clients, high-value clients and informed clients, respectively, in the 180 days prior to the transaction. Informed clients are identified based on their performance to anticipate price moves at the horizon of $h = 1,\ 5,\ 20,\ 30$ days ahead. The controls $\mathbf{X}$ include dummies for matched trades, client sales, trades executed on a regulated market or multilateral trading facility, the log of the traded amount and the lagged squared return on the benchmark, while μ is a vector of dealer-month, client-month, bond-month and industry-day fixed effects. Standard errors (shown in parentheses) are double clustered by month and dealer. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%).

Table A.7: Hypotheses horse-race – crisis

	Panel A: top 10%				Panel B: top 5%				Panel C: top 1%			
horizon	1	5	20	30	1	5	20	30	1	5	20	30
$Qrel\,(\gamma)$	−54.54	−54.55	−56.28	−54.50	−42.28	−42.58	−44.31	−42.58	−34.51	−36.16	−37.86	−36.68
	(35.48)	(35.28)	(35.48)	(35.43)	(34.27)	(33.95)	(34.24)	(34.43)	(30.95)	(30.74)	(31.02)	(31.46)
$Qrel \times ic\,(\beta_{ic})$	−96.48	45.05	140.80*	24.80	−75.63	76.99	272.60**	172.40**	−127.80	31.88	171.30	94.90
	(76.39)	(87.50)	(70.11)	(54.82)	(73.10)	(106.80)	(127.90)	(66.22)	(85.55)	(105.30)	(157.50)	(90.31)
$Qrel \times pc\,(\beta_{pc})$	27.17	−1.42	−31.71	7.45	36.55	11.21	5.97	5.19	−16.72	−51.51	−53.23	−54.03
	(47.25)	(45.65)	(42.35)	(46.54)	(53.46)	(52.82)	(47.78)	(46.35)	(57.72)	(54.94)	(45.36)	(46.49)
$Qrel \times lc\,(\beta_{lc})$	−79.25*	−93.53**	−108.10**	−88.13*	−133.30***	−146.10***	−153.60***	−149.40***	−167.30***	−176.20***	−175.00***	−175.80***
	(44.66)	(43.64)	(46.45)	(44.87)	(43.30)	(43.33)	(42.95)	(42.45)	(49.67)	(49.36)	(46.63)	(46.84)
$Qrel \times pc \times lc\,(\beta_{pclc})$	212.30***	243.50***	288.60***	229.50***	194.10**	222.70***	233.50***	233.60***	80.43	168.50	184.90	186.40
	(74.93)	(75.97)	(71.71)	(77.74)	(81.74)	(79.38)	(75.20)	(74.47)	(148.80)	(167.10)	(132.10)	(128.30)
ic	7.29**	−3.15	−3.59	−4.10**	6.79*	−4.98	−7.10	−5.38	12.39**	0.98	−3.78	−0.70
	(3.03)	(5.12)	(3.26)	(1.95)	(3.68)	(7.00)	(5.77)	(3.88)	(5.16)	(7.26)	(5.80)	(5.81)
pc	2.48	3.32	4.24	3.21	−1.16	−0.47	−0.22	−0.25	6.26	7.68*	7.90*	7.91*
	(3.22)	(3.31)	(3.12)	(3.42)	(4.72)	(5.05)	(5.09)	(4.98)	(4.41)	(4.39)	(4.24)	(4.33)
lc	−1.60	−1.03	−1.05	−0.95	3.12	3.24	3.16	3.29	3.58	3.42	3.26	3.32
	(1.68)	(1.55)	(1.60)	(1.65)	(2.21)	(2.11)	(2.08)	(2.14)	(2.55)	(2.49)	(2.52)	(2.52)
$pc \times lc$	−0.87	−1.50	−2.54	−1.38	1.83	1.62	1.21	1.14	−1.83	−4.00	−4.51	−4.50
	(3.13)	(3.45)	(3.21)	(3.58)	(5.72)	(5.80)	(5.92)	(5.81)	(6.20)	(6.46)	(5.52)	(5.44)
$\gamma + \beta_{ic}$	−151.03*	−9.50	84.55	−29.70	−117.91	34.41	228.32	129.86*	−162.29	−4.28	133.39	58.22
$\gamma + \beta_{pc}$	−27.38	−55.97	−87.99**	−47.04	−5.73	−31.37	−38.33	−37.39	−51.24	−87.67**	−91.10**	−90.71**
$\gamma + \beta_{lc}$	−133.79***	−148.08***	−164.39***	−142.63***	−175.57***	−188.64***	−197.87***	−191.99***	−201.83***	−212.41***	−212.83***	−212.44***
$\gamma + \beta_{pc} + \beta_{lc} + \beta_{pclc}$	105.65	93.99	92.48	94.29	55.13	45.27	41.61	46.81	−138.12	−95.46	−81.15	−80.09

Notes: This table shows the result from fitting the regression model:

$$TC_{bdct} = \gamma\, Qrel_{dct} + \beta_{lc}\, Qrel_{dct} \times lc_{dct} + \beta_{hv}\, Qrel_{dct} \times hv_{dct} + \beta_{ic}\, Qrel_{dct} \times ic_{dct} + \beta_{hvlc}\, Qrel_{dct} \times hv_{dct} \times lc_{dct} + \alpha_{lc}\, lc_{dct} + \alpha_{hv}\, hv_{dct} + \alpha_{ic}\, ic_{dct} + \alpha_{hvlc}\, hv_{dct} \times lc_{dct} + \mathbf{X}'_{bdct}\beta + \mathbf{1}'\mu + \varepsilon_{bdct},$$

for the crisis period from 1 March to 18 March 2020 where lc_{dct}, hv_{dct} and ic_{dct} are dummy variables taking the value 1 if the client is in dealer's top 10% (Panel A), 5% (Panel B) or 1% (Panel C) of liquidity providing clients, high-value clients and informed clients, respectively, in the 180 days prior to the transaction. Informed clients are identified based on their performance to anticipate price moves at the horizon of $h = 1,\ 5,\ 20,\ 30$ days ahead. The controls $\mathbf{X}$ include dummies for matched trades, client sales, trades executed on a regulated market or multilateral trading facility, the log of the traded amount and the lagged squared return on the benchmark, while μ is a vector of dealer, client, bond and industry-day fixed effects. Standard errors are clustered by dealers. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%).