Relationship discounts in corporate bond trading [*]

Simon Jurkatis
Bank of England
simon.jurkatis@bankofengland.co.uk

Andreas Schrimpf
BIS & CEPR
andreas.schrimpf@bis.org

Karamfil Todorov
BIS
karamfil.todorov@bis.org

Nicholas Vause
Bank of England
nicholas.vause@bankofengland.co.uk

This version: December 14, 2022

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 face a sizeable 67% drop in transaction costs relative to the median client—an effect which is particularly strong during the COVID-19 turmoil. We find clients' liquidity provision to be a key driver of relationship discounts: clients to whom dealers can turn to as a source of liquidity, are rewarded with relationship discounts. Another important motive for dealers to quote better prices 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.

JEL Classification: G12, G14, G23, G24

Keywords: Corporate bond, COVID-19, Dealer, Over-the-counter financial market, Trading relationship

[*]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, Jonathan Hall, Wenqian Huang, Ulf Lewrick, Jessica Li, and seminar participants at the Bank of England, the FSB conference on Systemic risks in NBFI and policies to address them, the 17th Central Bank Conference on the Microstructure of Financial Markets, and the JEF Conference.

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 and 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 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 data set 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. To the best of our knowledge, our paper is the first 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 are also the first to show that relationship clients generate the bulk of dealers' profits.

To measure the strength of dealer-client relationships, we rely on past bilateral trading volume between the two counterparties. Specifically, a client has a stronger 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. To do so, 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 tight panel regression specifications – with a rich set of fixed effects absorbing any bond, client, time, industry, and dealer-related variation – and control for other transaction-specific variables.

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

In the second part of our empirical analysis, we examine whether relationship clients obtain better prices. The results from our baseline panel regression show that one percentage point increase in our relationship measure ($Qrel$) is associated with a 0.28 basis points drop in transaction costs for clients. This implies that the top 1% of relationship clients face a sizeable *67% drop in transaction costs* relative to the median client. This drop amounts to total annual savings of at least £750,000 for top relationship clients. 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 baseline estimate is magnified more than threefold from 0.28 to 1 basis point 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.

In the third part of the empirical analysis, we dig deeper into the underlying mechanisms driving dealer-client relationships. To guide our empirical 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. To ensure continued access to this source of liquidity, dealers may reward such clients with relationship discounts. To test this hypothesis, we interact the main relationship metric with a dummy variable for liquidity-providing clients. We identify such clients as those to whom the dealer consistently off-loads bonds bought from other investors. If the hypothesis were true, we expect to see that dealers charge lower transaction costs for relationship clients who are consistently providing liquidity.

Our second hypothesis, *"profit maximisation"*, 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. 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 were true, we 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 top profit clients.

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 the order flow of these clients (e.g., informed hedge funds). 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 who consistently predict future price moves.

Our results show that the reduction in trading costs for clients with stronger relationships is consistent with the "liquidity provision" hypothesis. First, we document that dealers' balance sheet management considerations are related to relationship discounts. We define a dummy variable for transactions that are unmatched by an offsetting trade in the opposite direction executed at the same time to capture trades in which the dealer acts as a 'principal'. In normal times, discounts are higher for unmatched client purchases, consistent with dealers rewarding clients for reducing dealers' inventory. Discounts for relationship clients increased strongly during the COVID-19 crisis period when dealers were likely more constrained, and were higher for unmatched client sales.

Second, we find that liquidity providing clients consistently obtain larger relationship discounts. This extra discount is magnified six times and becomes highly significant during the COVID-19 episode. A striking fact is that the reduction in transaction costs for relationship clients during the crisis period is almost entirely driven by liquidity providing clients.

We also obtain some evidence in support of the "profit maximisation" hypothesis. We find that the average profit that dealers make on top relationship clients is *14 times larger* than the profit made on 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.

We find no evidence in support of the "private information" 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 with an eye towards information extraction. In addition, dealers may be also reluctant to trade with informed clients or even charge them higher costs to protect against adverse price changes. If dealers know that they face an informed trader who correctly predicts future price changes on average, dealers would face the risk of prices moving against them after the trade. Consistent with this conjecture, we find that informed relationship clients were actually charged larger costs, although this effect is not statistically significant.

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 turn to for "outsourcing" of liquidity provision. Dealers differentiate among clients especially for trades where they have to take additional inventory risk. As such, dealers' allocation of balance sheet space differs across clients, depending on how much dealers value the relationship with these clients. These relationship effects are particularly important during crisis times when relationship discounts almost exclusively accrue to liquidity-providing clients. Second, we find that dealers earn much higher profits from top relationship clients relative to others, creating a strong incentive for dealers to keep these types of counterparties as loyal customers.

The findings of our paper have several implications for investors and policy makers. First, we find evidence that dealers continued to provide liquidity 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 some of the largest clients who truly benefit from such dealer relationships. Many smaller market participants do not enjoy such benefits and may even see their costs of accessing dealers' balance sheets surge in stress episodes. Second, dealers appear to value especially the types of clients they can turn to as a source of liquidity provision. This could help dealers to operate with smaller inventories, which are cheaper to maintain, and increase dealers' overall liquidity provision. Third, the fact that dealers do not offer larger discounts to informed clients is consistent with dealers no longer implementing proprietary trades, which contributed to the buildup of risk in the run-up to the GFC and hence have been restricted by regulation. Albeit indirect, the evidence in this paper points to a cessation of proprietary trading by dealer-banks.

Related Literature. Our paper contributes to the growing literature on the value of relationships in over-the-counter markets. Two important papers in this literature which, similar to ours, focus on corporate bond markets are Di Maggio et al. (2017) and Hendershott et al. (2020). Di Maggio et al. (2017) show that dealers value relationships with other dealers in US corporate bond markets, particularly in stress times. Their study, however, focuses entirely on relationships in the inter-dealer market, which dealers use to manage inventory risk, whereas our paper covers the dealer-client segment (i.e., dealer trades with ultimate "end-users"). Also, the sample in Di Maggio et al. (2017) 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] Hendershott et al. (2020) in turn, like our paper, study the dealer-to-client segment. However, they do so for only a narrow subset of clients by studying primarily insurance companies. The authors document a non-monotonic impact of the number of trading relationships that insurance companies have with dealers 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). Insurance companies account for less than 9% of total trading volume in our sample.

Understanding the value of relationships from the client's perspective is particularly important against the backdrop of the evolving microstructure of the bond market. Key developments in past years have been the advent of electronic platforms, particularly based on request-for-quote (RFQ) protocols, which represent an electronic form of OTC trading. Despite this "electronification", all-to-all trading – a feature of equities or futures trading – has so far remained very limited in cash bond markets (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 centralised limit order books (Hendershott and Madhavan, 2015; Abudy and Wohl, 2018), the OTC structure has thus 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

[1] Di Maggio et al. (2017)'s sample ends in 2011. 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.

platform or bilaterally. The authors estimate that forgoing this benefit accounts for 40% of total costs for institutional investors. Our findings confirm the importance of including such loyalty components in structural models.

Our paper shows the importance of clients' heterogeneous trading relations with a dealer and the inherent benefits. 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 the clients enjoying relationship discounts. On the other hand, relationship benefits are typically given to selected clients and could be implicitly subsidized by other clients that 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 relationship clients incur larger transaction cost than non-relationship clients in FX derivatives. They attribute this finding to unsophisticated clients trading only with a particular dealer and, hence, being "captured". Relationships are measured as a dummy variable that is one for clients who also have a credit line with their dealer-bank. In contrast to their study, we measure the *strength* of bilateral relationships and focus on relationship benefits from the perspective of the dealer.

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 they mostly line up opposing client trades. 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 instead shows that asset managers and brokers are the majority of liquidity providing clients to dealers, accounting for more than 90% of the trading volume for such trades. Insurers, by contrast, typically account for less than 5% of the overall volume of liquidity providing clients. In addition, we show that dealers value relationships with liquidity providing clients by giving discounts to those clients, and show that these discounts are not only given for trades that help to reduce dealer inventories, but especially for those that *boost* inventories in stress times.

Finally, a number of papers study trading relations in OTC markets in the context of informed trading. Kondor and Pintér (2022) and Czech and Pintér (2020) show that clients that 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, perhaps because of post-GFC restrictions on proprietary trading or because of fear of being adversely selected.

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

This section provides an overview and basic descriptive statistics of the data on dealer-bond transactions and bond prices used in this paper.

2.1. Description of the regulatory data set

The main dataset for our analysis consists of transaction reports in corporate bonds submitted to regulatory authorities under MiFID 2, which took effect on 3 January 2018. Under the 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 MarketAxess, TradeWeb and Bloomberg) and organised trading facilities.[2]

[2] For precise definitions see point (20) to (23) of Article 4(1) of the Directive.

Trades have to be reported irrespective of whether they are actually carried out on a venue or not, 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. This sets our paper apart from other studies that used data under the MiFID I regime. Another benefit of our data is the granularity and depth of data coverage 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 executed by UK investment firms and on UK trading venues in reportable financial instruments. 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, even 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. In turn, 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 take on risk exposures.[3] Finally, we complement this data with information on bond characteristics, such as maturity and ratings, from S&P Capital IQ and ESMA's Financial Instruments Reference Files.

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,

[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 as much as possible. This is often achieved by trading in inter-dealer markets.

establishing a reference price is not an easy task.

For our reference prices, we use proprietary mid-quote data from MarketAxess Composite+ (CP+).[4] These data provide a level of pricing information that is usually only available to sophisticated market participants. CP+ is based on a proprietary machine-learning pricing algorithm developed by MarketAxess that generates pre-trade reference bid and ask 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. Basic descriptive statistics

Table 1 summarizes key descriptive statistics for the dealer-client transactions in our sample. Our sample runs from 3 January 2018 (the start of MiFID II reporting regime) to 18 March 2020 (the end of the 2020 dash-for-cash episode).[6] In much of our paper, we investigate relationship discounts and transaction costs separately during normal and crisis times. To do so, we split our sample into a pre-crisis (before March 2020) and crisis period (1 to 18 March 2020). We set the start of the turmoil period at the beginning of increased selling pressure in the bond market.[7] The end of the period, March 18, is when the "dash for cash" abated, following the

[4]Our approach of relying on actual market reference prices is superior to common practice in the literature where benchmark prices are inferred from inter-dealer transactions. The problem with the latter practice is that many bonds are not sufficiently liquid to have both a dealer-client and inter-dealer trade close enough to each other: e.g., sometimes the time period between such trades is several days. Thus, the inter-dealer price might not be an accurate representation of the fair value at the time of the dealer-client transaction. Moreover, inter-dealer prices could be influenced by the relationships among dealers.

[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 after the end of the COVID-19 turmoil period to avoid confounding our findings with the effects of policy measures introduced in response to the COVID crisis. See Table 1 on page 14 of the Bank of England's May 2020 Monetary Policy Report.

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

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 being active during the COVID-19 crisis as well (Panel A of Table 1). In total, dealers interact with almost 16,000 clients before the crisis.[8] We record a total of almost 5 million dealer-client transactions in our sample, worth a total nominal value of GBP 4.2 trillion. Overall, 35,380 (16,857) different bonds were traded at least once before the COVID-19 crisis (during the crisis period).

We next analyze how concentrated trading is, both across dealers and customers. Figure 1 (left panel) shows that dealers differ significantly in terms of size (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. Figure 1 (right panel) in turn displays the size distribution for the largest 500 clients. As can be seen in the graph, the concentration of client volumes is even more extreme than that of dealers. The trading volumes of the largest clients are equivalent to that of a medium-sized dealer.

Figure 1: Trading volume concentration across dealers and clients

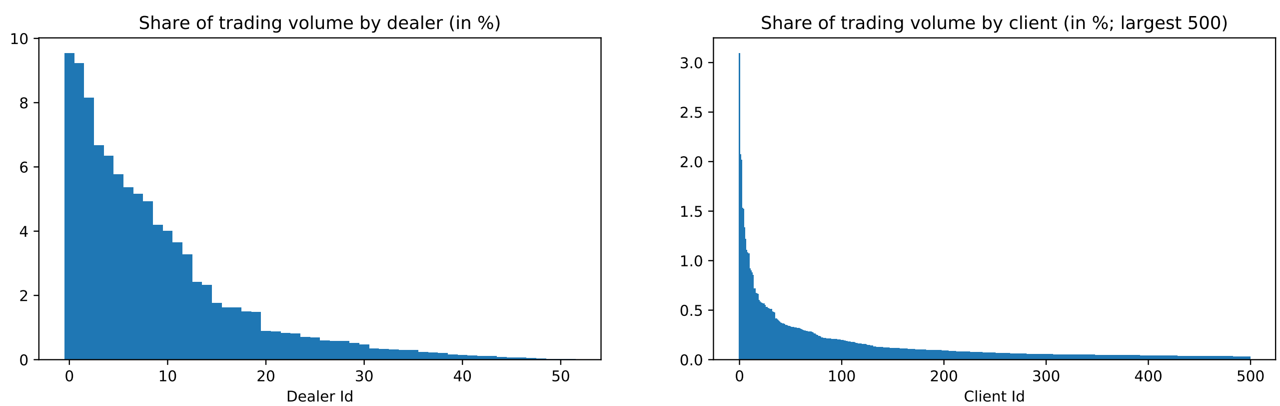

Notes: The Figure displays the share of trading volume (in %) for dealers (left) and clients (right) over the full sample period.

Some further insights into trading by dealers and customers can be gleaned from Panels B

[8] These are 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.

Table 1: Descriptive statistics

Panel A: Totals			Panel B: Per day avg across dealers			Panel C: Per day avg across clients		
	pre-crisis	crisis		pre-crisis	crisis		pre-crisis	crisis
#Dealers	52	50		34	36		3	3
#Clients	15,580	4,013		78	86		910	1,043
#Trades	4,824 k	154 k		183	237		7	8
Volume	4,069 bn	130 bn		154 m	200 m		6 m	7 m
#Bonds	35,380	16,857		151	189		6	7

Panel D: Bond stats			Panel E: Sector distribution			Panel F: $Qrel$ and tc		
	pre-crisis	crisis		pre-crisis	crisis		pre-crisis	crisis
T2M	7.29	7.55	AM	53.12	53.07	$top_{t,t-1}$	0.62	0.68
%ZTD	55.53	50.07	Bank	17.67	17.58	$corrQ_{t,t-1}$	0.81	0.86
AAA	6.51	6.34	Broker	12.34	8.15	$top_{i,j}$	0.14	0.12
A-AA	30.66	33.2	PFLDI	8.70	10.60	$corrQ_{i,j}$	0.16	0.13
BBB	39.64	34.48	HF	6.95	8.04	tc buy	7.06	9.84
HY	23.19	25.98	PTF	1.22	2.57	tc sell	8.11	34.18

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), trades (#Trades) and bonds (#Bonds), as well as total trading volume measured in GBP. Panels B and C show corresponding daily averages across dealers (Panel B) and clients (Panel C). Panel D presents bond characteristics. Time-to-maturity (T2M) is the volume-weighted average 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 in Panel D is measured in terms of trading volume among all bonds that had a rating in February 2020. High-yield (HY) is below BBB. Panel E shows the share of trading volume by client sector: asset managers (AM), banks, hedge-funds (HF), pension funds, insurers and liability-driven investors (PFLDI), proprietary trading firms (PTF), and brokers (Broker) including executing and investing services firms. Panel F presents average relationship persistence and overlap as defined in Eq. (5) – (8), as well as average transaction costs measured in bps for client buys and sells, respectively.

and C of Table 1. Panel B shows that on average, 34 dealers were active on any given day before the crisis and 36 were active during the crisis. The need for intermediation services on behalf of clients and the supply of those services by dealers picked up notably during the crisis: daily dealer volume rose by roughly 1/3 and and the number of clients went up 10% in that period. Panel C of Table 1 shows the equivalent statistics for clients. Roughly 900 clients were active on a typical day before the crisis compared to 1,000 during the crisis. Clients interact on average with 3 dealers on a given day, corroborating the earlier observation that clients' trading behavior is much more concentrated than that of dealers.

We next study the main characteristics of bonds traded in terms of duration, liquidity and credit risk. Specifically, we provide statistics for volume-weighted aggregate measures of the bonds' time-to-maturity, the percentage of zero-trading-days (%ZTD), and credit rating. Panel D of Table 1 shows that the average bond has around 7-8 years to maturity and does not trade on 55% of trading days. The latter number decreases to 50% in the crisis, suggesting that slightly more liquid bonds were traded during the COVID turmoil. In terms of credit quality, around one (three) quarter of trading volume is in high-yield (investment-grade) bonds.

Finally, it is important to understand the composition of investors active in corporate bond trading. To this end, Panel E of Table 1 reports the distribution of trading volume across different client segments. The asset management (AM) sector is by far the largest and accounts for more than the half of the trading volume. Banks are the second largest group with 18%, followed by brokers with 12% of total volume (pre-crisis). Hedge funds (HF), as well as pension funds, insurance companies and other liability-driven investors (PFLDI) account for around 7-11% of trading volume. Principal trading firms (PTFs) are only a small share of the total, which might be related to the low degree of electronification in corporate bond trading.

3. Empirical strategy and measurement of key variables

The main goal of our paper is to study the extent to which dealers give price discounts to clients, depending on the degree of their past trading relationships with those clients. In this subsection, we start by laying out our main econometric approach—panel regressions with a rich set of fixed effects to tease out the effect of relationships on prices. We then proceed by describing the measurement of our main variables: (i) relationship discounts in transaction costs faced by different clients and (ii) dealer-client relationships. We then present some basic stylised facts about these two main variables.

3.1. Empirical strategy

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:

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

where tc_{dcbt} 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.

Fixed effects. The vector μ in Eq. (1) contains bond-month, dealer-month and client-month fixed-effects. This rich set of fixed-effects allows us to control for many observable and unobservable variables that would influence transaction costs. For example, bond's liquidity likely affects transaction costs for all clients. Including bond-month fixed effects ensures that our results are not driven 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. Including client-month and dealer-month fixed-effects ensures that our results are not driven by such characteristics. In several robustness tests, we also include week, industry $\times$ week, day, and industry $\times$ day fixed effects.

Additional controls. We control for additional variables in the vector $\mathbf{X}_{dcbt}$. These include three dummy variables, $sell_{dcbt}$, $match_{dcbt}$ and $elec_{dcbt}$, 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. Furthermore, we control for intra-day price moves, $r_{bt} \times dt_{dcbt}$ (explained below) and for size of the transaction (in logs).

To compute $match_{dcbt}$, 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 or does not need to incur any additional search costs to locate the bond. In practice, a dealer would line up these trades with different clients in advance and then execute them simultaneously. In general, we would expect such transactions to be associated with lower costs for the client, as shown in Goldstein and Hotchkiss (2020). That said, we would also expect that having a good relationship with a dealer is beneficial for a client precisely in transactions that involve risk warehousing and are thus more balance sheet intensive for the dealer.

One might worry that in instances of extreme intra-day price moves, the benchmark price may not represent the actual fair value of the bond, if the market is moving fast and the benchmark was observed long before the trade. To ensure that our findings are not driven by such intra-day price drifts, we include the intra-day return, r_{bt}, multiplied by the distance between the time of the trade and the time of the benchmark price.

3.2. Measuring differences in transaction costs faced by clients

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), 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 faced by the client and is different from the measure used by Di Maggio et al. (2017) in their study of inter-dealer relationships.[9]

Basic descriptive statistics on transaction costs. Panel F of Table 1 reports the volume-weighted average transaction costs for client purchases and sales, both before and during the March 2020 stress period.[10] We see that the average transaction cost for client sales (dealer purchases) of 8 bps were similar to those for client purchases (7 bps) in the pre-crisis period. However, during the turmoil period, transaction costs surged markedly, and became highly asymmetric for client sales vs. client purchases. Costs increased to 10 bps for a client purchase and jumped fourfold to 34 bps for a sale during the COVID-crisis. At this time, dealers faced significant net selling pressure from clients, which could help explain why transaction costs were markedly higher for client sales than client purchases.

3.3. Measuring dealer-client relationships in corporate bond trading

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

For each transaction, we calculate the total trading volume of the dealer-client pair across all bonds over the previous 180 days, lagged by one week. We then divide that measure by the total volume of trading between the dealer and all its clients. This in turn, gives us the share

[9] Di Maggio et al. (2017) focus on the profits from a dealer's perspective, as measured via a round-trip of a buy followed by a sell transaction.

[10] For the computation of average transaction costs, we adjust the measure in Eq. (2) to account for intra-day price moves that could lead to a discrepancy between the available benchmark price and the fair value at the precise moment when a transaction occurs. Namely, we compute

$$tc_{adj} = tc - r\, dt/T, \tag{3}$$

where r is a bond's intra-day return derived from MarketAxess' opening to closing prices, dt is the time-distance of a transaction to its benchmark measured in seconds, and T is the total time between the open to close, also measured in seconds.

of trading volume that dealer d obtains from trading with the respective client c:

$$Qrel_{dct} = \frac{\sum_{\tau \in \mathcal{T}_d(t-187,t-7)} \mathbb{1}_{\{c_\tau = c_t\}} Q_\tau}{\sum_{\tau \in \mathcal{T}_d(t-187,t-7)} Q_\tau}, \tag{4}$$

where $\mathcal{T}_d(t-187, t-7)$ is the set of all transactions of dealer d with all its clients over the 180 days prior to the transaction at t lagged by 7 days to ensure that the relationship metric is not influenced by the most recent trades. $\mathbb{1}_{\{c_\tau = c_t\}}$ is an indicator function taking the value 1 if the client in transaction τ is the same as in the current transaction, and Q is the nominal volume of a transaction measured in GBP. Intuitively, $Qrel$ captures, from a dealer's perspective, how important a particular client is in terms of 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 main dealer-client relationship measure. Table 2 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 0.29% 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 almost a fifth 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 look at how persistent such relationships are. For each dealer, we measure the persistence of client relationships by:

$$top_d(t, t-1) = \#\{istop_{d,t} \cap istop_{d,t-1}\} / \#\{istop_{d,t}\} \tag{5}$$

$$\text{and} \quad corrQ_d(t, t-1) = \text{corr}(Qrel_{d,t}, Qrel_{d,t-1}), \tag{6}$$

where $istop_{d,t}$ is the set of top-1% clients according to $Qrel_{dct}$ for dealer d at period t, and

$Qrel_{d,t}$ is the vector of clients' share in a dealer's overall trading volume. Hence, $top_d(t, t-1)$ is the share of a dealer's top clients that were already top clients in the previous period, and $corrQ_d(t, t-1)$ is the correlation of a dealer's client share over two adjacent 180-days periods.

Panel F in Table 1 shows that relationships are fairly persistent over time. The average correlation across dealers and periods is above 80%. On average, more than 60% of dealers' top clients at a given point in time were also top clients in the previous period. Interestingly, both measures increase during the COVID-19 turmoil, suggesting that trading by dealers with their prior top clients intensified over that period.

It is also informative to study the degree of overlap in client-relationships across dealers. To understand how relationship-clients of one dealer overlap with those of other dealers, we use similar metrics to the ones above, but compute them over the cross-section of dealers:

$$top_t(i,j) = \#\{istop_{i,t} \cap istop_{j,t}\}/\#\{istop_{i,t}\} \tag{7}$$

$$\text{and} \quad corrQ_t(i,j) = \text{Corr}(Qrel_{i,t}, Qrel_{j,t}). \tag{8}$$

Panel F in Table 1 shows that relationship-clients are not shared across dealers. On average, only 14% of a dealer's top clients are also top clients of another dealer. Together, these findings suggest that, when clients have a strong relation with a dealer, they tend to be quite loyal and do not simultaneously show up among the top clients of another dealer. Both the cross-sectional overlap of top clients $top_{i,j}$, as well as the cross-sectional correlation between dealers' $Qrel$ measures $corrQ_t(i,j)$, decrease slightly during the crisis period. These patterns suggest that top clients' trading became even more concentrated during that episode, thereby increasing the reliance on the relationship dealer.

4. Relationship discounts in corporate bond trading

In this section, we present our baseline results from estimating the main regression Eq. (1). The goal is to quantify the extent to which dealers give transaction-cost discounts to relationship

Table 2: Distribution of the relationship measure

	mean	std	percentile				
			1%	25%	50%	75%	99%
Qrel	1.39	3.53	0.00	0.06	0.29	1.14	18.32

Notes: Th Table shows the distribution of the relationship metric $Qrel$ measured in % across all transactions. 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. (4).

clients. We start by analyzing the relationship discounts that clients receive in normal times, before turning to the magnitude of such discounts during the COVID-19 crisis and the role of various trade characteristics.

Table 3 reports the results of our baseline regression Eq. (1) for the pre-crisis period (Panel A). In all regressions, we cluster standard errors at the dealer and month level, allowing for arbitrary correlation between transactions of the same dealer and arbitrary correlation between transactions within the same month.

We find that our main dealer-client relationship variable $Qrel$ has negative and statistically significant effect on transaction costs, irrespective of the fixed effects specification. The estimates consistently suggest that transaction costs are 0.28 basis points lower for every percentage point increase in $Qrel$. Quantitatively, these estimates imply that a top-percentile client (with $Qrel = 18.3$, as reported in Table 2) pays about five basis points lower transaction costs than the median client ($-0.28 \times (18.3 - 0.3) \approx -5$). This reduction is a sizeable *two-thirds* discount on the average transaction cost of 7.5 basis points in the pre-crisis period and amounts to annual savings of £750,000 for a top client assuming daily trading volume of £6 m (see Table 1) and 250 trading days ($6 \times 5/10000 \times 250 = 0.75$).

Note that in the regressions reported in Table 3, we allow for client, dealer, bond and time fixed effects to separately affect transaction costs. As explained before, bond-time fixed effects allow bonds with different time-varying characteristics, such as their liquidity, to have systematically different transaction costs. Similarly, dealer-time and client-time fixed effects

Table 3: Baseline results on relationship discounts

	Dependent variable: transaction costs					
	Panel A: Pre-crisis			Panel B: Crisis		
	(1)	(2)	(3)	(4)	(5)	(6)
Qrel	-27.83^{**}	-27.74^{**}	-27.86^{**}	-77.28^{**}	-97.64^{***}	-99.69^{***}
	(11.79)	(11.77)	(11.76)	(37.68)	(34.69)	(35.05)
match	-2.15^{***}	-2.15^{***}	-2.15^{***}	-8.66^{***}	-6.57^{***}	-6.53^{***}
	(0.52)	(0.52)	(0.51)	(2.19)	(1.62)	(1.57)
client sell	-1.24	-1.26	-1.26	31.35^{***}	28.73^{***}	28.76^{***}
	(1.27)	(1.26)	(1.26)	(3.13)	(2.88)	(2.89)
elec	1.87^{**}	1.84^{**}	1.84^{**}	-3.93	-1.28	-1.26
	(0.71)	(0.71)	(0.71)	(2.63)	(2.21)	(2.20)
logQ	0.55^{***}	0.56^{***}	0.56^{***}	0.21	0.30	0.32
	(0.06)	(0.06)	(0.06)	(0.23)	(0.22)	(0.22)
dt*r	0.01	0.02	0.02	-0.40^{***}	-0.06	-0.06
	(0.03)	(0.03)	(0.03)	(0.14)	(0.12)	(0.12)
R^2	0.21	0.22	0.22	0.30	0.35	0.36
nobs	1,933,568	1,933,560	1,932,525	74,551	74,551	74,549
Fixed Effects						
dealer × month	Yes	Yes	Yes			
client × month	Yes	Yes	Yes			
bond × month	Yes	Yes	Yes			
day		Yes			Yes	
industry × day			Yes			Yes
dealer				Yes	Yes	Yes
client				Yes	Yes	Yes
bond				Yes	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. (4); *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; *elec*, 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; $dt \times r$, the bond's intra-day return (estimated as the difference between the closing and opening benchmark price) times the seconds between the time of the trade and the time of the benchmark. 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%).

allow for heterogeneity in transaction costs of different market participants, perhaps related to their bargaining power or centrality in the trading network. For example, client-time fixed effects prevent confounding our estimates of the effect of a client being important to a *particular dealer* with the effect of the client being important for the *overall market*. Finally, we control also for day and industry-day fixed effects to absorb any variation in transaction costs related to systematic time-varying differences by day and across industries.[11]

Relationship discounts during the Covid-crisis. Panel B of Table 3 reports the results of our baseline regression Eq. (1) for the March 2020 crisis period. In the strictest fixed effects specification (column 6), the coefficient estimate for $Qrel$ is nearly quadruple its pre-crisis value, suggesting that transaction costs were 1 basis point lower for every percentage point increase in $Qrel$. Thus, a top-percentile client paid transaction costs about 18 basis points lower than the median client ($-1 \times (18.30 - 0.29) \approx -18$) during the crisis period. As average transaction costs roughly tripled during the crisis period to 22 basis points (unconditional on the trade direction), this estimate implies an economically sizable discount for top-percentile clients of around 80% of transaction costs faced by the average client.

Controlling for trade characteristics. The pre-crisis and crisis results hold after controlling for other potential determinants of transaction costs. We control for several trade characteristics that are likely to influence transaction costs due to dealers' balance sheet management considerations.

First, we control for matched trades. In a matched trade, a dealer simultaneously executes pre-arranged offsetting sales and purchases, such that the package 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.,

[11] Another control we include is the intra-day benchmark-price returns between the last observed benchmark price and the time of the transaction ($dt \times r$). The control is intended to correct for any staleness of benchmark prices, which could bias our estimates of transaction costs. However, the estimated coefficient for this variable is not statistically significant in most specifications as seen from Table 3. This is not surprising as the benchmark price is observed every few hours and intra-day price changes during non-stress times are typically small for corporate bonds.

2021). Indeed, we find that in normal times transaction costs for matched trades are about 2 basis points lower than for unmatched trades (Panel A of Table 3). In the crisis period, the discount for matched trades is more pronounced and 3 times larger than the pre-crisis level. This spike in transaction costs for unmatched trades shows that dealers charged clients more for balance-sheet-intensive or search-intensive trades during the crisis when dealers were likely more constrained.

Second, we analyse to what extent transaction costs depend on the direction of the client's trade. Where dealers do not already have a pre-arranged match, a client *sale* means that the dealer needs to take the bond into its inventory. By contrast, a client *purchase* needs to be filled by drawing on the dealer's inventory, or by sourcing the bond from another client (or dealer). While the former may induce some balance sheet constraints, the latter would alleviate such constraints. Despite this asymmetry, we find no significant difference between transaction costs for client purchases and sales in the pre-crisis period.

However, in the crisis period (Panel B of Table 3), we find that client *sales* are much more expensive than client *purchases*. A likely reason is that dealers incorporated balance sheet considerations in quotes offered to clients during the crisis period. Absorbing client sales raised inventory risk by adding to dealers' long positions, and dealers may have expected to bear this risk for longer than usual because of a highly imbalanced market with few buyers.

Other controls in our main regression include a dummy that takes the value of one for trades executed electronically on a platform or multilateral trading facility, and trade size. We find that during pre-crisis times, transaction costs for electronic trades are nearly 2 basis points higher compared to other OTC trades. Regarding trade size, we find larger trades to incur higher transaction costs, consistent with previous studies (e.g. Pinter et al., 2022). These higher transaction costs may reflect a compensation for dealers for taking on greater amounts of inventory risk when absorbing large trades. The estimates on both controls are however, insignificant during the crisis period as indicated in panel B of Table 3.

Interim conclusion. Taken together, the findings presented in this section provide strong evidence of a significant degree of price differentiation across clients, depending on their relationship with dealers. Relationship discounts are magnified during the COVID-19 crisis period. Our results further indicate that dealers have limited capacity to manage inventory risk arising from corporate bond trades. As such, dealers charge different costs to different clients for trades that use this capacity, while at the same time differentiating according to the strength of trading relationships. Thus, when transaction costs increased substantially during the crisis period, they went up particularly sharply for *sales* to dealers (which likely *increased* inventories and, hence, inventory risk) and were higher for unmatched trades (which boost inventory risk whenever they do not offset an existing inventory position). Even as transaction costs soared, however, dealers charged their top relationship clients substantially less than other clients, offering them 80% discounts compared to the median client.

5. Why dealers give relationship discounts?

In this section, we test three hypotheses that could explain why dealers give relationship discounts to clients: "liquidity provision", "profit maximization" and "information extraction". We find that dealers give larger discounts to clients who they can turn to for liquidity provision, thereby allowing dealers to manage balance sheet costs more efficiently. In addition, we find discounts to be also related to dealers' profit maximisation motives since top relationship clients generate the bulk of dealers' profits. However, we do not find evidence that information extraction is a main driver of relationship discounts.

5.1. Hypothesis 1 - "liquidity provision": dealers value clients that step in as liquidity suppliers

Relationships could be important for dealers to facilitate a convenient way to off-load bonds when the dealer is faced with an inventory imbalance or is balance-sheet constrained. The relationship discount should then be more pronounced for such liquidity-providing clients. Given that balance-sheet intensive trades appear more costly for dealers as shown in the previous section, we now study whether balance sheet motives interact with the size of the discount given to relationship clients.

5.1.1. Relationship discounts for unmatched trades

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 by estimating:

$$tc_{dcbt} = \gamma\, Qrel_{dct} + \alpha_1\, nomatch_{dcbt} + \beta_1\, nomatch_{dcbt} \times Qrel_{dct} + \mathbf{X}'_{dcbt}\beta + \mathbf{1}'\mu + \varepsilon_{dcbt}. \qquad (9)$$

As shown in the first column of Panel A in Table 4, and consistent with earlier results, dealers charge a little over 2 basis points more to execute unmatched trades compared to matched trades in the pre-crisis period (α_1). More importantly, dealers give a relationship discount of about 37.5 basis points per unit of $Qrel$ for unmatched trades ($\gamma + \beta_1$), compared to 16 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.

We next study whether dealers give larger discounts for unmatched client *sales* than unmatched client *purchases*. 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:

$$
\begin{aligned}
tc_{dcbt} &= \gamma\, Qrel_{dct} \\
&\quad + \beta_1\, Qrel_{dct} \times nomatch_{dcbt} + \beta_2\, Qrel_{dct} \times sell_{dcbt} + \beta_3\, Qrel_{dct} \times nomatch_{dcbt} \times sell_{dcbt} \\
&\quad + \alpha_1\, nomatch_{dcbt} + \alpha_2\, sell_{dcbt} + \alpha_3\, nomatch_{dcbt} \times sell_{dcbt} \\
&\quad + \mathbf{X}'_{dcbt}\beta + \mathbf{1}'\mu + \varepsilon_{dcbt}. \qquad (10)
\end{aligned}
$$

In this regression, β_1 captures any additional relationship discount for *unmatched* client purchases, while β_2 captures any additional relationship discount for matched client *sales*. Moreover, β_3 captures relationship discount for trades that are *both* unmatched and client sales. Thus, the overall relationship discount for unmatched client sales is given by $\gamma + \beta_1 + \beta_2 + \beta_3$.

As shown in the 'nomatch × sell' column of Table 4, dealers offer an additional relationship discount over the baseline for unmatched *client purchases* outside of crisis times, but no significant discount for unmatched *client sales*. This result suggests that dealers reward relationship clients when they help to reduce dealers' inventories, but not when they increase them. In the crisis period (panel B), the unconditional cost of unmatched client sales ($\alpha_1 + \alpha_2 + \alpha_3$) rose strongly, but dealers gave large relationship discounts on such trades ($\beta_1 + \beta_2 + \beta_3$ is significantly negative). Thus, when dealers were flooded with orders that swelled their inventories, they gave preferential pricing to relationship clients.

Table 4: Trade characteristics and relationship discounts

	Panel A: Pre-crisis					Panel B: Crisis				
			Conditional on nomatch-trades					Conditional on nomatch-trades		
	nomatch	nomatch × sell	high-yield × sell	long-maturity × sell	illiquid × sell	nomatch	nomatch × sell	high-yield × sell	long-maturity × sell	illiquid × sell
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Qrel	-15.94^{*}	-20.81^{**}	-41.93^{**}	-57.65^{**}	-56.54^{**}	-43.79	-72.49	29.21	1.27	16.01
	(8.31)	(8.93)	(18.54)	(22.73)	(23.67)	(43.94)	(58.91)	(51.33)	(61.69)	(58.20)
Qrel × ι	-21.65^{**}	-31.17^{**}	-27.76	-25.33	3.66	-78.61^{**}	73.88	-32.05	-93.36	-51.46
	(9.81)	(13.90)	(16.72)	(16.23)	(11.65)	(36.63)	(71.38)	(64.76)	(62.54)	(94.68)
Qrel × sell		10.61^{**}	32.41	32.96^{*}	30.61^{*}		42.72	-250.90^{**}	-190.22	-196.05^{*}
		(4.03)	(20.75)	(16.98)	(16.76)		(43.31)	(101.33)	(124.65)	(115.06)
Qrel ×ι× sell		20.49	-7.87	-23.43	-7.53		-241.26^{**}	241.59^{**}	163.33^{*}	47.41
		(17.31)	(27.45)	(36.83)	(15.97)		(117.83)	(101.14)	(84.48)	(168.92)
ι	2.47^{***}	1.66^{**}				7.90^{***}	5.54			
	(0.59)	(0.62)				(1.42)	(3.53)			
sell	-1.27	-3.08^{**}	-1.46	-1.51	-1.37	28.79^{***}	27.84^{***}	28.49^{***}	28.77^{***}	30.84^{***}
	(1.26)	(1.39)	(1.32)	(1.32)	(1.36)	(2.88)	(4.61)	(2.19)	(2.23)	(2.58)
ι× sell		1.65^{**}					3.86			
		(0.73)					(5.47)			
R^2	0.22	0.22	0.21	0.23	0.23	0.36	0.36	0.36	0.37	0.37
nobs	1,932,525	1,932,525	1,373,460	1,523,402	1,644,036	74,549	74,549	53,844	55,518	60,226

Notes: This Table shows the estimates from the regression

$$tc_{dcbt} = \gamma Qrel_{dct} + \beta_1 Qrel_{dct} \times \iota_{dcbt} + \beta_2 Qrel_{dct} \times sell_{dcbt} + \beta_3 Qrel_{dct} \times \iota_{dcbt} \times sell_{dcbt}$$
$$+\alpha_1 \iota_{dcbt} + \alpha_2 sell + \alpha_3 \iota_{dcbt} \times sell_{dcbt} + \mathbf{X}'_{dcbt}\beta + \mathbf{1}'\mu + \varepsilon_{dcbt},$$

where $Qrel$ is the relationship metric defined in Eq. (4) and ι_{dct} is one of the risk-dummies: nomatch, high-yield, long-maturity, illiquid. The controls $\mathbf{X}$ include a dummy for electronic trades, the log of the traded amount and the intra-day return on the benchmark. μ is a vector of 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. Note that there are no estimates on ι and $\iota\times$ sell in columns 3–5 and 8–10 as these are spanned by bond and bond-month fixed effects. Asterisks indicate significance levels (***= 1%, **= 5%, *= 10%).

Next, we define a set of dummies ι_{dcbt} that capture the riskiness of the traded bond to test if riskier bonds, which are likely costlier to hold on balance sheet, are associated with larger discounts than safer bonds. We consider three dimensions of risk: credit risk, market risk and liquidity risk. Thus, in turn, we define ι_{dcbt} to equal one if the bond is high-yield, if the remaining maturity of the bond is in the upper 10th percentile (long-term bonds are typically more sensitive to interest rates) and if the proportion of zero trading days (ZTDs) for the bond is in the upper 10th percentile. The results, shown in columns 3-5 of Table 4, suggest that discounts are generally larger (transaction costs are smaller) for riskier bonds sold by relationship clients, but the estimates are not statistically significant. This result shows that dealers do not seem to discriminate bonds on their risk when giving discounts on trades that affect inventories.

In contrast, during the dash-for-cash episode, transaction costs for riskier bonds sold by relationship clients were significantly *higher* as seen from the estimates on high-yield and long-maturity bonds (columns 8 and 9). These results indicate that bond's riskiness becomes important for dealers in times of stress, when balance sheet space is particularly costly, but is less important in normal times. Another reason for why safer investment-grade bonds received higher discounts during the crisis period could be that investors such as asset managers (who constitute a large part of relationship clients) sold mostly higher-rated bonds to bolster their cash balances (see Haddad et al. (2021)).

Taken together, our findings indicate that dealers' balance sheet management considerations are related to relationship discounts. In normal times, discounts are higher for unmatched client purchases, consistent with dealers rewarding clients for reducing their inventory. Discounts increase strongly during the crisis period for relationship clients when dealers were likely more constrained, and were higher for unmatched client sales. The latter result is consistent with dealers applying preferential pricing on balance-sheet intensive trades for relationship clients.

5.1.2. *Testing the liquidity provision hypothesis*

We next test more formally whether clients who provide liquidity to dealers, get larger discounts. To identify liquidity provision by clients, we rely on a similar approach as in Choi et al. (2021). For each client sale to a dealer, we find the clients who subsequently bought the bond from the dealer on the same day, and measure the amount of their liquidity provision as the value of bond purchases. 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, Table 5 reports their sectoral composition. The Table shows that 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 outsized role of brokers as liquidity suppliers – despite their relatively small overall size – is intuitive as some of the agents in that category include financial trading houses engaged in liquidity provision.[12]

We now proceed by testing more formally whether dealers reward liquidty providing clients. To do so, we add a dummy ϕ_{dct} for the top 1% of liquidity providing clients and estimate:

$$tc_{dcbt} = \gamma\, Qrel_{dct} + \delta\, Qrel_{dct} \times \phi_{dct} + \alpha\, \phi_{dct} + \mathbf{X}_{dcbt}\beta + \mathbf{1}'\mu + \varepsilon_{dcbt}. \tag{11}$$

If $\delta < 0$, liquidity-providing clients receive a larger discount than other clients supplying the same share of the dealer's trade volume ($Qrel$). Indeed, Table 6 shows that liquidity clients receive relationship discounts that are roughly twice as large (40.00 compared to 22.46) as for other clients. This result supports the liquidity providing hypothesis. Clients who provide greater volumes of liquidity are a particularly valued set of relationship clients for dealers.

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

Table 5: Dealers' client types and total profits

Panel A: Percent trading volume by sector

sector	all	top clients	liquidity clients	profit clients	informed-clients 1 day	informed-clients 30 days
Asset Manager	53.12	54.71	50.10	87.58	73.02	76.08
Bank	17.67	4.23	2.20	0.02	3.33	5.59
Hedge Fund	6.99	3.34	2.35	1.62	4.85	2.04
Broker	12.21	27.83	40.02	0.34	2.46	0.41
PFLDI	8.76	9.52	4.64	9.96	7.85	9.74
PTF	1.26	0.37	0.69	0.48	8.48	6.14

Panel B: Dealer profits for top and non-top clients

Client group	Total Profit (in £ m)	Total Volume (in £ bn)	Avg number of clients	Avg profit per client (in £ k)	Avg Vol per client (in £ m)
non-top	1043.29	1190.25	640	31.36	35.77
top	297.67	534.88	14	439.04	788.91

Notes: Panel A 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', 'profit clients' and 'informed-clients' show the percentages among the top 1% liquidity providing, profitable and informed clients (at horizons of 1 and 30 days), respectively. Panel B shows profit and volume statistics across dealers for their top and non-top clients aggregated over the sample (excluding 2020).

Table 6: Transaction costs for profit clients, liquidity clients and informed clients

	liquidity clients	profit clients	informed clients $h = 1$	informed clients $h = 30$
Panel A: Pre-crisis				
Qrel (γ)	-22.46^{*}	-28.69^{**}	-28.11^{**}	-27.92^{**}
	(11.85)	(12.44)	(11.87)	(11.83)
Qrel $\times$ client-type (δ)	-17.54^{*}	-9.51	17.29	4.95
	(9.06)	(14.54)	(15.46)	(7.69)
$\gamma + \delta$	-40.00^{***}	-38.20^{***}	-10.82	-22.97^{**}
Panel B: Crisis				
Qrel	-62.49	-72.38^{*}	-98.41^{***}	-100.84^{***}
	(37.89)	(36.33)	(36.00)	(34.95)
Qrel $\times$ client-type	-115.06^{***}	-84.40	2.44	87.34
	(36.82)	(97.08)	(67.02)	(53.00)
$\gamma + \delta$	-177.54^{***}	-156.78^{*}	-95.97	-13.51

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

$$tc_{dcbt} = \gamma\, Qrel_{dct} + \delta\, Qrel_{dct} \times \phi_{dct} + \alpha\, \phi_{dct} + \mathbf{X}'_{dcbt}\beta + \mathbf{1}'\mu + \varepsilon_{dcbt}, \tag{12}$$

where ϕ_{dct} is a dummy variable taking the value 1 if the client is in dealer's top 1% of liquidity providing clients, profit-generating clients and informed clients, respectively, in the 180 days prior to the transaction. The controls $\mathbf{X}$ include dummies for matched trades, client sales, electronic trades, the log of the traded amount and the intra-day return on the benchmark. μ is a vector of 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%).

5.2. Hypothesis 2 - "profit maximization": dealers maximize long-term profits

Another important motive for dealers to offer a discount is to attract a 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 as loyal customers, benefiting from their larger overall volumes. In addition, offering these clients a discount in corporate bond trading may also attract larger volumes from the same clients in other asset classes, i.e. there might be ”relationship spillovers”. 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 profit considerations play a role for dealers, we compute total trading profits of each dealer generated from top 1% relationship clients and from other (non-top) clients. Dealer profits from each 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 *revenue* for a dealer. For consistency with the other hypotheses, we define 'profit 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 descriptives about profit clients before proceeding with the formal tests. To this end, Table 5, Panel B, shows the trading profit across all dealers for both top and non-top clients. Dealers have on average 14 top and 640 non-top clients, and make in total around £0.3 billion and £1 billion profit over the sample period from those two groups, respectively. These facts show that only 14 clients account for more than 20% of total dealers' profits. Asset managers are the largest category of top profit clients and account for almost 90% of all such clients as seen from Panel A in Table 5. Importantly, the average profit made on a top client is 14 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 top profit clients receive larger discounts by defining a dummy for top profit clients similar to Eq. (11). Table 6 shows that profit clients receive relationship discounts, but these are no greater than for other clients supplying the same trading volume to the dealer: δ is not statistically significant. A possible explanation for the insignificant estimate on δ is that top profit clients are also typically top $Qrel$ clients since they also have the largest trading volume with the dealer (see Panel B in Table 5). Thus, most of the reduction in transaction costs for profit clients seems already captured in the estimate on $Qrel$.

5.3. Hypothesis 3 - "information extraction": dealers extract private information from relationship clients' trades

Besides profit maximization and liquidity motives, dealers might also wish to build relationships with clients, whose order flow provides valuable trading signals (Pintér et al., 2022). 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:

$$r_t(h) = [\log(p^*_{t+h}) - \log(p^*_t)] \times D_t, \tag{13}$$

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, $h \in \{1, 30\}$. 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 volume-weighted average directional return over all client trades with the dealer over a past 180-days window:

$$perf_{ct}(h) = \frac{\sum_{\tau \in \mathcal{T}_c(t-h-180, t-h)} r_\tau(h) Q_\tau}{\sum_{\tau \in \mathcal{T}_c(t-h-180, t-h)} Q_\tau}, \tag{14}$$

where $\mathcal{T}_c(t-h-180, t-h)$ is the set of all trades between the client and the dealer over the previous 180 days lagged by the return horizon h, and Q is the trade size. Thus, the metric gives higher weight to correctly predicted returns on larger trades. We scale $perf_{ct}(h)$ 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 $\widehat{perf_{ct}(h)}$. For each dealer we then define 'informed clients' as the top 1% performing clients according to their scaled performance, $\widehat{perf_{ct}(h)}$.[13]

Table 5, Panel A reports the sectoral composition of informed clients. It illustrates that asset managers account for the highest share of informed trading volume, particularly at longer horizons. Principal trading firms (PTFs) account for the second largest share at shorter horizons. This share is more than 20 times larger than the share of PTFs in top-relationship clients. Hedge funds account for a larger share of informed-client trading volume over the short-run versus the long-run, whereas the opposite is true for asset managers. These patterns are consistent with the idea that asset managers' trades have more information about long-term returns, whereas hedge funds' trades are more informative about short-term price changes (Czech et al., 2021).

We estimate Eq. (11) for top information clients but find no evidence in support of the information hypothesis: more informed clients do not receive larger relationship discounts. Although the δ coefficients are not statistically significant, they are often *positive*. This indicates that dealers may actually charge informed clients *more* than other clients supplying a similar share of dealer's trading volume. 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. To compensate for that risk, they could charge informed clients worse prices as reflected in higher transaction costs. Another reason why dealers do not offer discount to informed clients is that they can no longer profit from trading on information inferred from such client trades, given the clampdown on proprietary trading following the 2008 crisis.

[13] We require that a client must have at least 100 trades with the dealer in the 180-day window. Similarly, we require that the dealer must have at least 10 such clients.

In addition to studying the information effect for clients who consistently correctly predict future price changes, we also study whether clients who consistently "get it wrong", face different prices. The idea behind this test is that dealers might learn and profit not only from trades of informed clients, but also from those of uninformed ones (e.g., by taking the opposite position). We re-run regression Eq. (11) with a dummy for clients in the bottom 1% of informed clients, instead of the top 1% as before. The estimates in this regression were negative but still statistically insignificant. We excluded these results for brevity. In summary, we do not find evidence in support of the information extraction hypothesis.

5.4. Intersection between top profit, liquidity, and information clients

Figure 2 illustrates the degree of overlap between profit clients, liquidity clients and informed clients. The diagram shows that the same client can be important to a dealer for more than one reason. In particular, the majority of profit clients are also liquidity clients, whereas informed clients are less often important for profit or liquidity reasons.

Given the intersections between the different types of clients, we also tested whether the significant estimates on liquidity providing clients are partly driven by profit or information motives. We run a horse race between the three hypotheses by including dummies for all three client types in a single regression and found that the results are qualitatively the same as in the individual regressions. This regression is excluded from the paper for brevity.

Figure 2: Overlap between profit clients, liquidity clients and informed clients

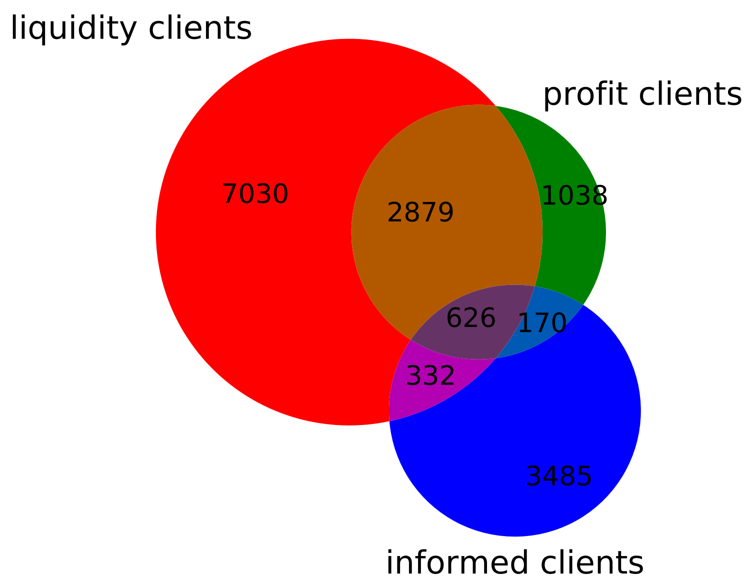

Notes: This figure shows the number of client-days in which a client is identified as a profit client, liquidity client or informed client.

6. Conclusion

We study the effects of dealer-client relationships in corporate bond trading. Our results show that clients with a stronger relationship with a dealer (top clients), as measured by past trading volume, receive better prices. The top 1% of relationship clients face a sizable 67% drop in transaction costs relative to the median client, which amounts to total annual cost savings of more than £750,000. The relationship benefits were particularly important during the dash-for-cash episode in March 2020 when the reduction in trading costs increased threefold.

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 to for outsourcing liquidity provision. Dealers differentiate among clients especially for trades where they have to take additional inventory risk. Second, dealers earn much higher profits from relationship clients relative to others, which creates a strong incentive for dealers to keep these types of counterparties as loyal customers.

Our findings show that the OTC structure centred around dealer intermediation in corporate bonds largely proved 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 it could help dealers to operate with smaller inventories, which are cheaper to maintain. On the other hand, relationship benefits are by nature reserved for particular clients. Clients not able to build significant relationships with dealers pay significantly larger transaction costs, especially during stress times. These findings have implications for the debate on the costs and benefits of different market structure designs – OTC markets vs All-to-All – that have 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 X. 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., S. Huang, D. Lou, and T. Wang (2021). Informed trading in government bond markets. *Journal of Financial Economics 142*(3), 1253–1274.

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.

Haddad, V., A. Moreira, and T. Muir (2021). When selling becomes viral: Disruptions in debt markets in the COVID-19 crisis and the Fed's response. *The Review of Financial Studies 34*(11), 5309—-5351.

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*.

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

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

Pinter, G., C. Wang, and J. Zou (2022). 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*.