BUNDLING TRADES IN OVER-THE-COUNTER MARKETS

JASON ALLEN[a], MILENA WITTWER[b]

Abstract. This paper documents trade bundling as a source of relationships between dealers and their clients in over-the-counter markets. Using trade-level data on the near-universe of the Canadian fixed income market, we show that investors bundle a sizable fraction of their trades—within and across asset classes—with the same dealer. This includes combining a group of buy or sell transactions, as well as bi-directional trading that involves swapping one asset for another. Data sets that focus on only one type of asset miss this aspect of trading, which we show impacts trading costs, providing monetary incentives for trading relationships.

Keywords: OTC markets, dealers, bonds, portfolio trading, financial innovation
JEL: G10. G12. G20

Acknowledgments: The views in this paper do not reflect those of the Bank of Canada. We are grateful to Mustapha Al Bachir Faye for phenomenal research assistance. We thank Jean-Francois Houde, Sebastien Plante, Jun Yang, Andreas Uthemann, Adrian Walton, Chaojun Wang (discussant), and participants at the 2025 NBER session Big Data and High-Performance Computing for Financial Economics. We thank the Bank of Canada trading desk for excellent discussions and providing real-time observations of multi-asset trading. We thank Matthew Arnot at CanDeal for his valuable insights on electronic RFQ platforms. [A]Jason Allen - Wisconsin School of Business and Bank of Canada, jallen39@wisc.edu; [B]Milena Wittwer - Columbia Business School, mw3941@columbia.edu.

Date: August 14, 2025.

1. Introduction

It is well documented that in over-the-counter (OTC) markets investors form long-term relationships with a small set of financial intermediaries ('dealers'), and never trade with most other dealers (Ashcraft and Duffie, 2007; Di Maggio et al., 2017; Hollifield et al., 2017; Li and Schurhoff, 2019; Hendershott et al., 2020; Allen and Wittwer, 2023). One explanation for this is that investors ('clients') trade off execution speed, i.e., immediacy, with transaction costs (e.g., Vayanos and Wang, 2007 and Uslü, 2019); or that they value access to liquidity during market turmoil (e.g., Carlin et al., 2007). The literature has primarily examined relationships over time within a single asset class, such as repeated interactions in corporate or municipal bond markets, with limited attention to relationships spanning across asset classes. Furthermore, due to the absence of persistent client identifiers, most studies have focused on inter-dealer relationships rather than client-dealer interactions.

We leverage trade-level data on the near-universe of the Canadian fixed income market with unique dealer and client identifiers to analyze trading over time and across assets classes. Our contribution is to highlight that investors commonly bundle trades, that is, they simultaneously buy or sell, or buy and sell multiple assets with the same dealer, within and across asset classes. This can both incentivize and strengthen dealer-client relationships. We show how trade bundling affects transaction costs, and highlight two possible offsetting pricing channels. On the one hand, improved inventory risk management may lower transaction costs; on the other hand, the convenience of bundling trades with one dealer may increase transaction costs.

For our analysis we create a novel dataset that covers essentially all fixed income trades involving a Canadian dealer, including those for non-Canadian products, such as U.S. Treasuries. The dataset is unique in that it includes identifiers for market participants and securities, so that we can trace both through the entire market. We observe the time, price, and size of trades. We also know if clients are institutional investors—our focus—and their type: asset manager, bank, hedge fund, etc. Moreover, we observe how trades between dealers and their clients are executed—either bilaterally through private negotiations, or on an electronic request-for-quote (RFQ) platform, where clients can request quotes from multiple dealers simultaneously.

With these data, we establish five facts that offer new insights into dealer-client interactions in OTC markets. Our first fact is that clients often trade in multiple asset classes with the same dealer. To show this, we group assets into eight classes, which include Canadian government debt, foreign sovereign debt, and non-financial corporate bonds. Roughly 40 percent of clients are active in at least four asset classes, with the average trading in two, and they tend to use the same set of dealers across them.

Documenting the tendency of investors to rely on the same dealer across asset classes broadens the OTC market literature, which has so far documented persistent relationships within asset classes (e.g., Di Maggio et al., 2017; Hendershott et al., 2020; Jurkatis et al., 2022; Allen and Wittwer, 2023; Pinter et al., 2024). This complements Pintér et al., 2025, who develop a search and bargaining model of the UK gilt and corporate bond market.

Our second finding is that trade bundling—simultaneously trading multiple bonds with the same dealer within a short time frame—is common, particularly with a client's preferred dealer. Nearly 20 percent of all trades involve multiple assets and bundling happens approximately 10 percent more often with the client's preferred dealer. One rational for bundling is that it minimizes execution risk when investors wish to trade several assets as part of a strategy—for example, ensuring both legs of a yield-curve trade are executed together.[1]

This fact complements a small set of studies documenting the simultaneous trading of multiple bonds within the same asset class, specifically U.S. corporate bonds (Li et al., 2023; Meli and Todorova, 2023). Lacking client identifiers, these studies rely on matching algorithms to infer take-it-or-leave-it offers involving long lists of corporate bonds. In contrast, our data identify counterparties directly, allowing us to track who trades with whom—even across asset classes. This is crucial for studying dealer-client relationships in the cross-section.

To analyze trade bundling more systematically, we distinguish between two main types of multi-asset trades: 'switches', which involve exactly one buy and one sell order, and 'bundles' which include all other multi-asset trades. We show that in recent years bi-directional trades—combinations of buy and sell orders in comparable amounts, often executed as switches—have gained popularity over uni-directional bundles, in which investors only buy or only sell multiple assets. This trend may be driven by stricter regulations that have increased the cost for dealers of carrying inventory (e.g. Basel III), coupled with the fact that bi-directional trades have minimal impact on dealers' balance sheets.

[1]A yield-curve trade is meant to be sensitive to relative movements between different points on the yield curve. An example is buying a long-term bond and selling a short-term bond of the same issuer. The trade is profitable if long rates rise relative to short rates.

Our third fact is that trade bundling, particularly in the form of switches, is significantly more common on electronic RFQ platforms than in bilateral negotiations between dealers and their clients. On an average day, roughly 26 percent of volume traded electronically involve switches, compared to only about 6 percent of bilateral trades. During the COVID-19 pandemic, the share of platform switches surged to 45 percent. This pattern may reflect the platform's ability to support trading strategies involving multiple assets, offering specialized services that streamline such complex transactions.[2] In other words, investors seem to value convenience when bundling trades, especially for switches. As electronic platforms gain traction in OTC markets, we expect trade bundling to play an increasingly prominent role.

Our fourth fact is that bundling trades tends to save transaction costs, especially when intermediated by a dealer that the client has a relationship with, providing a monetary incentive to maintain relationships. Specifically, conducting a multi-asset trade with a relationship dealer is cheaper than separately trading the bonds in the bundle on the same day and summing the individual costs, either with different dealers, or even with the same dealer. Switches, which typically result in small changes in a dealer's net inventory position because they combine a buy and sell-side transaction, are the cheapest type of multi-asset trade. This suggests that at least some of the balance sheet cost-savings accrued to dealers in these trades are passed on to clients, in particular, their repeat clients.[3]

Finally, for our fifth fact, we compare trades made on electronic RFQ platforms with those made through direct negotiations to better understand what influences bundle pricing, beyond balance sheet costs. Theoretically, it is not clear whether bundle trading should be cheaper or more expensive on an electronic platform. On the one hand, platforms can foster competition between dealers, since it is easier for clients to ask multiple dealers for price quotes (O'Hara and Zhou, 2021, Allen and Wittwer, 2023, and Kargar et al., 2024), which can reduce transaction costs. On the other hand, clients might be willing to pay a convenience premium for services provided by the platform

[2]For example, the platform allows the client to request a switch to trade bonds at par-value or make adjustments for changes in credit risk or duration risk. The platform calculates different net settlement amounts depending on the adjustments prior to sending out the RFQ so that the client has an idea about the net settlement, given indicative quotes (which are also provided by the platform).

[3]Note that switches are different from what are frequently called agency, matched, or paired roundtrip trades. See, for example, Goldstein and Hotchkiss, 2020 and Jurkatis et al., 2022. These are trades where a dealer buys from one client and (almost) immediately sells to a different client.

that simplify multi-asset trades, which can increase transaction costs. We find cost-advantages for clients when bundling uni-directional trades (such as buying multiple assets) but not when conducting switches on the electronic platform relative to bilateral trading. One possible explanation for this result is that dealers are less concerned about being adversely selected by a client with private information for a specific bond (or its issuer) when that client places an order across multiple bonds. As a result, the platform competitive effect dominates the convenience-premium charged for bundling.

Our analysis excludes several understudied factors that may also be a source of relationships in OTC markets, leaving room for future research. For instance, our data do not cover the repo market, where dealers and large institutional investors borrow and lend bonds against collateral to finance bond purchases. Examining dealer-client relationships across the repo and cash market could provide valuable insights into how funding needs influence these connections in fixed-income markets. Moreover, our study focuses on institutional investors, as retail investors in our data are not consistently assigned unique identifiers across dealers and market segments. It would be interesting to analyze relationships to retail investors, who tend to pay higher transaction costs. Finally, we abstract from relationships that connect fixed-income markets to equity and derivatives markets. For these, see Wittwer and Uthemann, 2024, who are the first to document dealer connections across these markets.

Besides contributing to the literature on OTC markets and dealer-client relationships, our study intersects with other research areas. It connects to the industrial organization literature on product bundling, rooted in seminal work by Stigler, 1963, Whinston, 1990, and Salinger, 1995. In contrast to this literature's focus on bundling products, our contribution emphasizes the importance of bundling trades—across asset classes and over time—in decentralized OTC markets. Unlike standard (seller-driven) product bundling, trade bundling involves both buyers and sellers, creating a different strategic interplay. This identifies opportunities to develop theories of trade bundling suited to decentralized markets, where incentives are driven by bilateral negotiations and relationships rather than centralized mechanisms. This would extend the small but growing literature on trade models with bundling in centralized markets (e.g. Chen and Duffie, 2021; Rostek and Yoon, 2021; Wittwer, 2021; Rostek and Yoon, 2025).

Finally, our study relates to the extensive literature on convenience yields for financial assets, which dates back to Kaldor, 1939 for commodities and has recently regained prominence in the context of safe assets (Krishnamurthy and Vissing-Jørgensen, 2012).

We highlight that convenience is not solely determined by the asset type (such as U.S. Treasury bills), but also by how an asset is traded—independent of liquidity differences. In our context, the electronic platform offers convenience services to investors, resulting in convenience premia for certain bundled trades. With the increasing adoption of electronic trading and the growing use of Artificial Intelligence to streamline trading processes, we anticipate that such convenience premia might become relevant in a wide range of settings. We hope our findings encourage further research into when trading mechanisms influence convenience yields.

2. Data

To analyze bundle trades we create a unique data set spanning across all main fixed-income instruments in Canada.

Data sources. The primary data source for this paper is the Market Trade Reporting System 2.0 (MTRS2.0). The Canadian Investment Regulatory Organization (CIRO) collects data on every fixed income trade in Canada, and this is provided to the Bank of Canada on a confidential basis. Our sample spans 2,068 trading days between January 4th, 2016 and December 31, 2023. After removing holidays, canceled orders, trades inputted more than once, trades where the trading entity identifier is missing, as well as Government of Canada bonds not denominated in Canadian dollars, we observe 12,928,044 trades for 176,331 unique securities.

We match the transactions data to data feeds from Refinitiv and Markit. These sources provide information (summarized in Appendix Table F1) about bonds characteristics, such as coupon, time of issuance and maturity, par amount outstanding, duration, convexity, 5-digit NAICS industry code, country of issuance, currency, etc.

Asset classes. We classify assets into the following categories: Government of Canada debt (GoC), Canadian provincial and municipal governments debt (Provi), non-financial corporate bonds, financial corporate bonds, foreign sovereign debt, Bankers' Acceptances (BA), Canadian agency debt, and non-Canadian agency debt.[4] Detailed asset class definitions are provided in Appendix A.

[4]We exclude two infrequently asset classes—supranational bonds and (non-mortgage) asset-backed securities. We also combine provincial and municipal bonds into a single category because the delineation in the data is often unclear when comparing the bond classification and the bond description. In our regression analysis we include bond (ISIN) fixed effects to control for the individual bond properties.

Table 1. Secondary market trade shares by asset class

Shares	Transactions (%)	Volume (%)
Government of Canada (GoC)	34.52	55.73
Canadian provincial governments (Provi)	13.94	9.32
Corporate bonds (non-financial)	8.96	1.58
Corporate bonds (financial)	15.23	6.13
Foreign sovereign debt	17.53	15.94
Bankers' Acceptances (BA)	4.51	7.42
Canadian agency debt (CDN)	3.58	3.77
Non-Canadian Agency debt (non-CDN)	1.73	1.04

Note: Table 1 presents the distribution of trades across the eight asset classes in our sample, measured both by transaction count and trading volume percentages. Asset class definitions are provided in Appendix A.

Table 1 presents market shares by number of transactions and volume traded for the different asset classes. The most traded assets are Government of Canada debt, followed by foreign sovereign debt (mostly U.S. Treasuries), provincial debt, corporate debt, and Bankers' Acceptances.

Trade data fields. We observe security identifiers (International Securities Identification Number, ISIN), the transaction time (rounded to the second), the side (buy/sell), the price, and the quantity of each trade.[5] A unique feature of the data is that dealers, brokers, and institutional investors carry a Legal Entity Identifier (LEI). Other investors, such as retail investors, have an anonymous identifier, which may be dealer-specific. Since we are interested in relationships, we only include investors with a LEI.

Dealers. There are 19 primary dealers in Canada—financial institutions that actively trade in the secondary market but who also have direct access to the Bank of Canada and primary market for government debt. The fixed income market, however, is relatively concentrated. Between 60 and 90 percent of trades are intermediated by the five largest dealers (see Appendix Figure F1).[6]

[5]We use date-time-ISIN-dealer-client-side-price-quantity to aggregate block trades. For example, a trade for $100 million might be reported as 10 different trades of $10 million. We report these types of trades as a single $100 million trade. We do not consider block trades to be bundles.

[6]The concentration of trades among the largest dealers in the Canadian market is not substantially different from other countries. O'Hara and Zhou, 2021, for example, report that over 70 percent of U.S. corporate bonds are intermediated by the top 10 dealers.

Clients. In addition to dealers, there are 7,982 clients with a unique LEI.[7] Many of these clients are subsidiaries of larger institutions—something that we also observe by leveraging the family structure reported by the Global Legal Entity Identifier Foundation, `https://www.gleif.org/en`.[8] When calculating transaction costs of dealer-to-client (D2C) trades, we separate out related party trades, i.e., trades between subsidiaries within a holding company or trades between a subsidiary and the holding company. There are likely to be non-cost reasons for a firm to transfer assets within its organizational structure. Although interesting, we leave this for future work, and focus instead on transactions between dealers and their unaffiliated clients.

We classify clients and their parent companies into types, as explained in Appendix B. The largest share of market participants are asset managers, followed by broker/dealers and banks. The parent, or holding companies, tend to be banks or asset managers. See Appendix Figure B1.

Trading venues. For D2C trades, which are our focus, we observe whether a trade was conducted bilaterally with a dealer or electronically over a RFQ trading platform. Most electronic D2C trades in our data take place on CanDeal, a platform owned by the largest dealers. Investors pay membership fees but no per-trade fees.[9] On the platform, an investor can request quotes from multiple dealers, specifying the securities, quantities, and settlement dates of the desired trades. Dealers, aware of the number of competing participants, but not their identities, respond with prices (within 10 minutes). The client selects the best offer, and the trade is executed shortly after.

Unlike bilateral negotiations, RFQ trading platforms can enhance competition among dealers, because they run auctions with multiple dealers (see O'Hara and Zhou, 2021 and Allen and Wittwer, 2023 for more details). However, because clients can query at most five dealers in an auction and not all of them may respond, the competitive

[7]We include all institutional investors, regardless of their activity across markets. While participation in multiple markets is likely correlated with trading costs, we seek to document which firms trade across asset classes, how they trade, and the price implications. Appendix G reports robustness results using a subset of 95 investors who trade the most volume across multiple asset classes.

[8]Consider the following example taken from the GLEIF database. Royal Bank of Canada has LEI code=ES7IP3U3RHIGC71XBU11. This uniquely identifies Royal Bank of Canada. Using the GLEIF database, we observe that this bank has 22 direct children (subsidiaries). One example is RBC Dominion Securities (LEI=549300QJJX6CVVUXLE15). It also has 72 ultimate children, that is, subsidiaries that report to one of the 22 direct subsidiaries. One example is the RBC Canadian bond index ETF (LEI=549300I23D6OJU4GU489). We consider any transaction between these parties as 'related'.

[9]Since the platform is dealer-owned, it matters little whether charges are explicit or embedded in higher prices for certain trades.

advantage can be limited.[10] Platforms also simplify the trading process by facilitating complex trades such as switches and bundles. For example, the platform provides indicative quotes for a wide range of bonds and calculates expected net settlement (cash transfer) for bundled trades that might have a credit or interest rate risk component.

Summary statistics. To get a sense of a typical trade, we report in Table 2 monthly trading averages for some of the variables in our data. The number of dealers trading each asset class is large, highlighting that most dealers are multi-class intermediaries, and yet all markets are rather concentrated (in line with Appendix Figure F1).

On average, there are between 291 and 684 unique clients trading per month, depending on the asset class. Eighty-three percent of the trades are D2C trades—our focus—but this varies substantially across asset classes. The share of trades that are executed on an electronic platform also varies by asset class. On the one end, 29 percent of Canadian government debt are traded electronically, which is consistent with Allen and Wittwer, 2023. On the other end, the platform trade share of foreign sovereign bonds is zero—this is because foreign sovereign bonds are the only asset class that cannot be traded there. For completeness, Table 2 also reports the share of transactions that are between related parties, which we drop from our analysis.

Trade size varies by asset class, as shown in Appendix Figure F2. The average D2C trade size is large—\$10.88 million, but this is driven by the market for Canadian government debt and Bankers' Acceptances. The median trade size is \$1.3 million.

3. Measurements for relationships and bundle trades

Before examining trade bundling and its association with relationships, we define the way we measure relationships, bundle trades, and transaction costs.

Relationships. We define relationships as repeat interactions between clients and the same dealer(s). Following Di Maggio et al., 2017 and Pinter et al., 2024, we measure the strength of client j's relationship with dealer i at date t as the fraction of trade (quantity) that client j intermediates with dealer i relative to the total amount the

[10] In contrast to other bonds, clients can request quotes from up to eleven dealers for corporate bonds. In a theoretical model, Wang, 2023 also shows that competition on RFQ platforms can be limited by dealer's strategic avoidance of competition.

Table 2. Summary trading statistics (monthly)

Variable	GoC	Provi	CDN	Non-CDN	Corp-NF	Corp-F	BA	Foreign
Volume (B$)	772	101	52	14.4	22	85.1	103	236
# trades	43,408	17,476	4,506	2,170	11,267	19,156	5,675	22,105
# dealers	18.3	18.1	16.4	14.8	16.6	17.6	15.1	13.6
# investors	684	450	234	291	488	681	325	453
HHI	1,270	1,328	1,459	1,383	2,095	1,798	2,445	3,921
D2C share (%)	51.1	68.6	66.4	86.7	88.2	89.5	96	97.3
Platform/D2C (%)	29	29	28.5	7.5	3.7	3.0	14.2	0
Related party (%)	14	14.3	14.5	21.7	23	27.4	30.5	41.6

Note: Table 2 provides monthly trading averages for some of the key variables in our data based on trading activity between January 4, 2016 and December 31, 2023 for each asset class: Government of Canada bonds (GoC), Canadian provincial and municipal debt (Provi), Canadian agency debt (CDN), non-Canadian agency debt (Non-CDN), non-financial corporate bonds (Corp-NF), financial corporate bonds (Corp-F), Bankers' Acceptances (BA), and foreign sovereign debt (Foreign). Currency is Canadian dollars and reported in billions. All foreign dollar trades are converted to Canadian dollars using the foreign exchange rate at the end of day of the trade. HHI is the Herfindahl–Hirschman Index and measures market concentration of dealers in each asset class. D2C share is the share of total trades that are dealer-to-client. Platform/D2C is the fraction of D2C trades conducted on one of the electronic trading platforms, such as CanDeal. Related party share is the percentage of trades that are conducted between a dealer and an affiliated counterparty. We drop these trades in our analysis.

client trades with any dealer, over a 180 day rolling window:

$$\Gamma^c_{i,j,t} = \frac{\sum_l^{N_{i,j,t}} \text{quantity}_l}{\sum_k^{N_{i,t}} \text{quantity}_k}. \tag{1}$$

For each client, we then use this measure to rank dealers by how much volume they trade from the client's 'most-favored' to their 'least-favored' dealer.

We define relationships at the LEI level. One could alternatively take a more granular view at the individual trader level or a more aggregated view at the parent-firm level. Our choice reflects both institutional realities and data limitations: we lack detailed information on individual traders, and dealer interviews suggest that relationships are maintained at the institutional LEI-level. For example, one dealer described a client calling the corporate bond desk to bundle a corporate bond with a Government of Canada bond; the dealer then contacted the government bond desk, and the quoted price depended on the client's identity or tier.[11] This pattern is consistent with Bak-Hansen and Sloth, 2024, who document that dealer–client relationships in Europe are not person-specific but are managed by sales teams to ensure that relationship pricing applies across desks.

Multi-asset bundle trades. One possible reason for why we observe long-term relationships in OTC markets are synergies created from bundling trades. To find out whether this is indeed the case, we define any transaction involving more than one asset that occurs between the same dealer and client within a five second window as a multi-asset (or bundle) trade.[12]

We differentiate between two main types of bundle trades, which can be initiated by clients and dealers alike. First, 'switches' are any trade between the same dealer and client within a five second window that involves exactly one buy and one sell. Second, 'bundles' are all other types of multi-asset trading strategies. If the transaction involves all buy (sell) we call these 'uni-directional', or, more specifically, 'buy (sell) bundles'. If they involve multiple buys and sells we call it a 'bi-directional bundle'.

Bundles are related yet broader than 'portfolio trades', which have been studied in the context of the U.S. corporate bond market by Meli and Todorova, 2023 and Li et al., 2023. Portfolio trades in these papers are assumed to be take-it-or-leave-it offers, often related to mutual fund redemptions. They are defined as trading protocols according to which an investor buys or sells a portfolio of individual corporate bonds with a single dealer. Bundle trades also relate to, but are broader than, 'list trades.' List-trades are a protocol used on electronic trading platforms where clients send a RFQ for multiple bonds and multiple dealers can win. As with TRACE, we cannot observe whether a given bundle was executed as part of a list trade.

Trade bundling may arise from a variety of mechanisms, including take-it-or-leave-it offers, bilateral negotiations over multiple assets, or list trades. Moreover, countless trading strategies may give rise to switches or bundles. In Appendix Table C we provide a high-level breakdown of the main strategies. Common switches are 'rate' trades,

[11]Another example, although still small, is the emergence of auto-filling on electronic platforms. Rather than respond to individual requests, dealers have pre-arranged programs that automatically respond with a quote for certain clients and assets.

[12]Given the precision of the time stamps, the time-window for defining bundles is not crucial, so long as it's in seconds and not hours or days. Extending the definition of a bundle to a day would substantially increase the size of a bundle for some client-dealer interactions.

which are bets on yield curve shape movements, or 'credit' trades, which involve swapping bonds of different credit. The most common bundles are uni-directional, with a median of four bonds per trade. Some strategies even have catchy names, such as a 'butterfly'. Executing a positive butterfly means buying an asset with a time-to-maturity somewhere in the middle part of the yield curve (the 'belly') and simultaneously selling assets with short and long time-to-maturities (the 'wings').

Independent of the specific trading strategy, executing all trades in a strategy at once through a bundle trade can be attractive to investors for several reasons. First and foremost, it eliminates the risk of partial execution—for instance, completing only one leg of a two-legged butterfly.[13] Moreover, most switches and some bi-directional bundles have minimal impact on dealers' inventories (as we will show). This is advantageous for managing balance sheet costs and capital requirements. Bundling also streamlines trade processing, settlement, and record-keeping by reducing the number of individual transactions, benefiting both clients and dealers. In addition, cash-constrained investors might prefer to sell a bond to finance the purchase of a different bond—sometimes called self-financed trades—and documented by Barardehi et al., 2024 in equities markets. Further, some types of investors, such as asset managers, prefer to swap out riskier bonds for Treasuries, instead of cash, since holding cash is a drag on fund performance. Lastly, investors may bundle a large set of assets to signal that the trade is 'uninformed' and not driven by adverse selection.[14]

Conversely, investors may avoid bundling trades with the same dealer to conceal their trading strategy. For example, an investor with private information about the economy might wish to buy a longer-dated bond while shorting a short-dated bond without risking that the dealer front-runs the trade. Splitting the order helps mask

[13]Execution failure is a real possibility in some markets and in some periods (e.g., Hendershott and Madhavan, 2015, Kargar et al., 2024, and Hendershott et al., 2024). In general, there are a number of risks related to execution. One example is timing risk. With separate trades, an investor is exposed to market moves between when the first leg and the second leg are executed. The dealer quoting a bundle is pricing the trade simultaneously, which protects the investor from adverse moves during execution. Another example is settlement risk. Having a single dealer means one settlement process rather than coordinating two separate settlements that might not align perfectly.

[14]In Appendix Table D1, we present a simple test for adverse selection. We compare the transaction costs of bundles of 10 or more assets to those with fewer than 10 assets. We choose 10 since FINRA defines portfolio trades as those with 10 or more assets and these are typically assumed to be uninformed. In all our specifications, trades with at least 10 assets are cheaper than smaller bundles. This suggests that uninformed trades receive better pricing than those that may signal the investor has private information, which can lead dealers to charge more to protect against adverse selection.

the strategy. However, because investors typically trade with relationship dealers they appear to trust, we expect this incentive to be small relative to execution risk.

Transaction costs for single-asset trades. To estimate transaction costs of a single trade j, we follow Hendershott and Madhavan, 2015:

$$\text{cost}_j = \log(\text{trade price}_j/\text{benchmark price}_j) \times \text{trade sign}_j. \tag{2}$$

Here trade price is the transaction price for trade j; benchmark price is the average transaction price of the same bond (ISIN) on the same day; and trade sign is an indicator variable equal to 1 for an investor purchase and -1 for an investor sale.[15] Costs are multiplied by 10,000 to put units in basis points of the benchmark price. Following the literature, we winsorize the top and the bottom 1 percent of the transaction cost estimates to eliminate outliers. Appendix Figure F4 plots transaction costs by asset class as well as client-type. There is substantial variation across and within asset class in daily transaction costs as well as variation across clients. Over the entire sample, the mean transaction cost is 1.10 basis points with a standard deviation of 17 basis points. This is slightly higher than those reported in the UK gilt and corporate bond market by Pinter et al., 2024.

Transaction costs for bundle trades. In order to determine whether multi-asset trading is more or less expensive than single-asset trades, we need to define transaction costs for bundle trades and compare them to the cost of hypothetical individual transactions of the components in the bundle trade.

To provide an intuition for this, it helps to provide an example of what we observe in the data—see Table 3. The first two trades represent a switch by investor 1. They bought a corporate bond (Corp) for 99 and sold a government of Canada bond (GoC) for 100. To determine the switch (dis)advantage, we need the counterfactual price this investor would have paid if they had conducted the same transaction in two pieces.

Our approach is to compare investor i's corporate buy price to the average buy price of all other investors on the same day and their GoC sell price to the average sell price

[15] There are many alternatives proposed in the literature to quantify transaction costs and what to use as benchmark prices. For example, one benchmark in Hendershott and Madhavan, 2015 is the last inter-dealer price and not the daily average across all trades. For robustness, in the Appendix we report results using the inter-dealer price. Rather than using the last inter-dealer price, we take the quantity-weighted average to account for the substantial variation in trade sizes across asset classes and the documented sensitivity of prices to trade size. In addition, Appendix Figure F3 plots the time series of transaction costs based on four different measures in the literature: they are highly correlated.

TABLE 3. Example of a multi-asset trading advantage

Investor	Asset	Time	Side	Price	Quantity	Dealer
1	Corp k	9:05:01	buy	99	1,000	A
1	GoC j	9:05:01	sell	100	1,100	A
2	Corp k	9:10:15	buy	99.2	500	A
3	GoC j	9:11:22	sell	100	900	A
4	Corp k	10:12:01	buy	99.3	600	B

Note: Table 3 provides a hypothetical example of what we observe in the data. We label the investor 1 trade as a switch since they are buying 1,000 of a specific corporate bond (Corp),k, and selling 1,100 of a specific Government of Canada bonds (GoC), j, simultaneously with the same dealer. In this example, investors 2,3, and 4 conduct single-asset trades.

of all other investors on the same day. In the example, the average corporate buy price is $(99.2 \times 500 + 99.3 \times 600)/1{,}100 = 99.254$; the GoC average sell price is 100. Therefore, the investors saves $(99.254 - 99) \times 1{,}000 + (100 - 100) \times 1{,}100 = \$254$ in transaction costs when bundling.

This logic underpins the formal analysis we present in Section 4, where we establish our empirical findings on bundle trading. In this section, we estimate fixed effects models using the transaction costs calculated in equation (2) as the dependent variable. The underlying idea of the regressions aligns with the example.

4. Empirical facts about relationships and bundle trades

We develop five facts about bond trading, focusing on multi-asset strategies, relationships, and pricing. We begin with establishing that most investors trade multiple asset classes and do so with a single dealer. Then we focus on our main contribution—analyzing bundle trading and its pricing implications with relationship dealers.

Fact 1. *Most investors trade multiple asset classes and do so with a single dealer.*

Figure 1 shows the percentage of clients who trade multiple asset classes over our sample period, separated by buy and sell side: approximately two-thirds of clients trade multiple asset classes. Asset management companies and investment banks trade the most number of asset classes while non-financial entities and hedge funds trade the least, as shown in Appendix Table F2.[16]

[16] If we aggregate at the holding company/parent level, the patterns are similar, with one exception: Appendix Table F3 shows that banks, as the ultimate parent of many funds and investment banks, are more similar to asset management companies.

Figure 1. Distribution of the number of asset classes that clients buy and sell

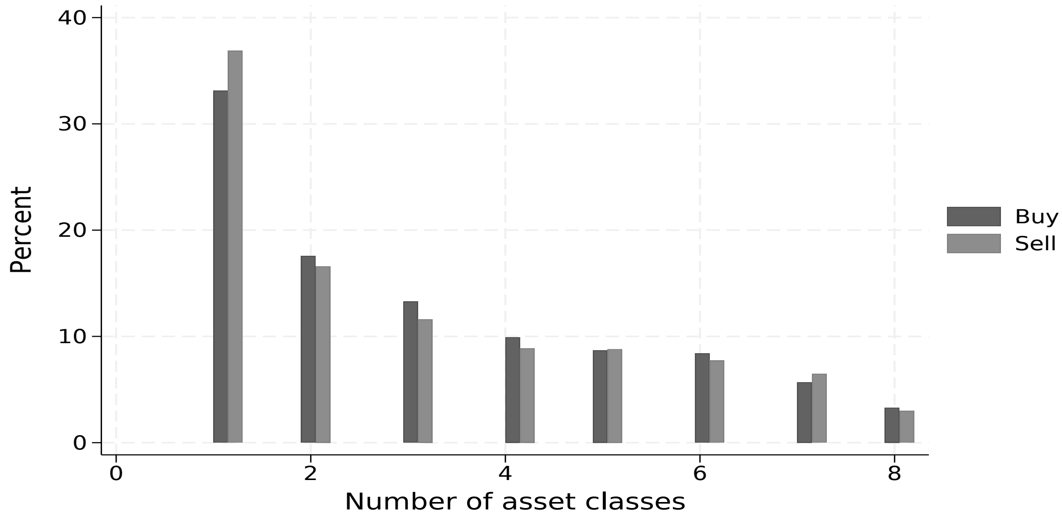

Note: Figure 1 shows the distribution of the number of asset classes across clients for when they buy and sell, separately.

The existing literature has shown that within an asset class most investors trade with very few dealers. We confirm, using our relationship measure (1), that this remains true when considering an entire fixed-income market. Table 4 shows the average trade share with the most-favored dealers, averaged across clients. Conditional on having more than one relationship, 80.4% of trades are between a client and their most-favored dealer, while 92.7% of trades are between a client and their top two dealers. If we weight each client-favorite-dealer pair by each clients' total trade volume rather than treating all pairs equally, the trade shares are lower but still substantial—51% for the most-favored dealer and 14.7% for the second-favorite dealer. This reflects the fact that more active clients have more dealers; however, even these clients tend to have a large share of volume with a single dealer.

Given the fact that clients tend to trade with few dealers and trade multiple asset classes, the question is whether they use the same dealer for all asset classes or segment dealers, perhaps because of dealer specialization. Nearly 40 percent of clients have the same number of dealers in an asset class as they have in total. In the median, some client-types, such as asset managers, have more dealers across than within assets classes—see Figure 2, which plots the distribution of number of dealers across asset classes in white and within asset class in gray. Despite this, the total number of dealers

Table 4. Mean trade volume across most-favored dealers

	EW share (%)	VW share (%)
Most-favored dealer	88.1	53.0
Most-favored dealer | # dealers > 1	80.4	50.8
Second-favored dealer	12.3	14.7

Note: The first row of Table 4 shows the average trade share that clients execute with their most-favored dealer, based on relationship measure (1). In the second row, we condition on clients who have more than one dealer. In the third row, we show the average trade shares with the two most favored dealers. In the columns, we report both equal-weighted (EW) and volume-weighted (VW) averages. To calculate the volume-weighted shares, we weight each client-dealer pair by the clients total trade volume. The value-weighted shares are substantially lower than the equal-weighted shares because active clients have more dealers than less active clients.

remains low. Furthermore, clients are more likely to trade with their most-favored dealer across asset classes than within an asset class—see Appendix Table F4.

Having established that dealer-client relationships are prevalent not only over time but also across asset classes, we now turn to investigating the role of bundling trades. The underlying idea is that the advantages of these trades may help facilitate and strengthen relationships. Given the limited knowledge about multi-asset trading in fixed-income markets, we first conduct a detailed analysis of bundle trades (Facts 2-3) before examining their pricing implications (Facts 4-5).

Fact 2. *Investors engage in multi-asset trading strategies, and this is often with their most-favored dealer. In recent years, bi-directional trades are particularly common.*

We start by establishing that multi-asset trading is common. Figure 3 panel (A) plots, in the red dashed line, the share of total volume traded that we classify as switches. Over our sample period the share of switches increases from 6 percent to over 12 percent. In the blue solid line, we plot the share of all other bundle trades (which we refer to as bundles) over time—they also increased. In total, multi-asset trading represents over 20 percent of trades.[17] This is a substantial proportion, especially considering our conservative definition of bundle trades (as transactions within 5 seconds between the same dealer-client pair), and underscores the importance of analyzing trading across asset classes.

[17] In Appendix Table F5 we show that asset managers are the client-type most engaged in bundles. This is not surprising given that asset managers (and investment banks) trade the most asset classes (recall Appendix Table F2). In contrast, hedge funds and pension funds & insurance, who tend to trade fewer asset classes, are more actively switching assets—often within an asset class.

Figure 2. Number of dealers per client-type

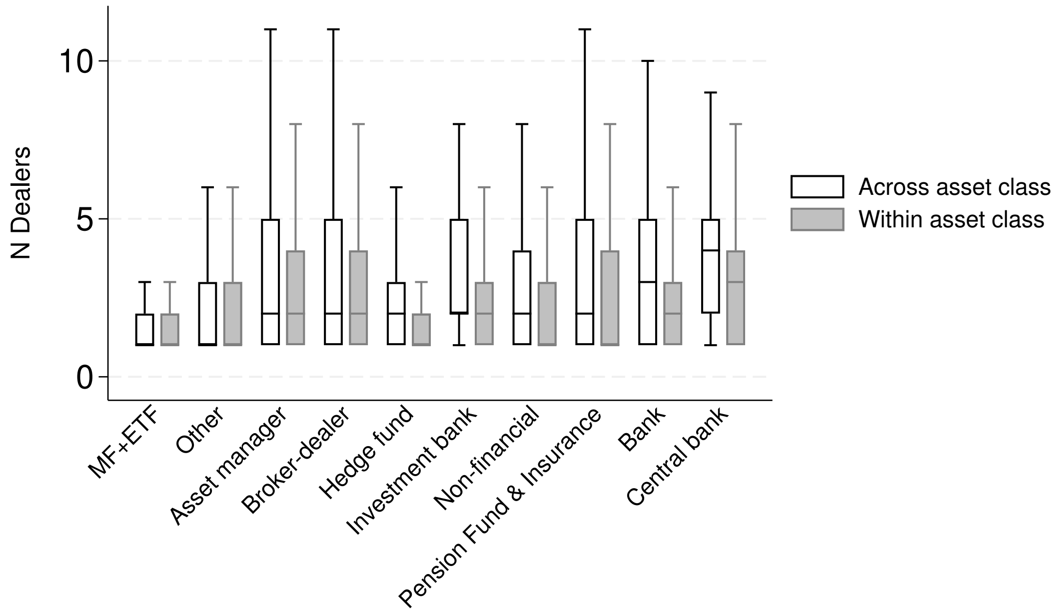

Note: Figure 2 plots the distribution of the number of dealers per client-type across asset classes in white and within asset class in gray, excluding outliers. MF stands for mutual fund while ETF stands for Exchange-Traded Fund. The definition of client-types are provided in Appendix B.

Table 5 shows that investors are more likely to bundle trades with their most-favored dealer than with other dealers (we distinguish between bundles and switches in Appendix Table F7). To show this, we regress an indicator variable for whether the trade is with the most-favorite dealer on an indicator variable for whether the trade is a bundle trade plus different sets of fixed effects. Results are relatively similar. Consider, for example, column (2) which includes client-day fixed effects. The constant says that 46.2% of trades are with the most-favorite dealer. A bundle trade is 6.2 pp (or 13.4%) more likely to be with the most-favorite dealer. This suggests that investors are intentionally bundling trades at a one-stop-shop.

In recent years, switches and bi-directional bundles have become more popular than buy or sell bundles (see the dotted red lines in Figure 3). One possible explanation for this trend is that dealers are increasingly constrained by limited balance sheet capacity, leading them to favor transactions that involve both buying and selling—particularly switches. This behavior parallels the growing preference for agency trading (or riskless principal trading) over principal trading in the U.S. corporate bond market as a result of the Volcker Rule requiring dealers to meet minimum capital requirements in stress

FIGURE 3. Share of volume that is bundled

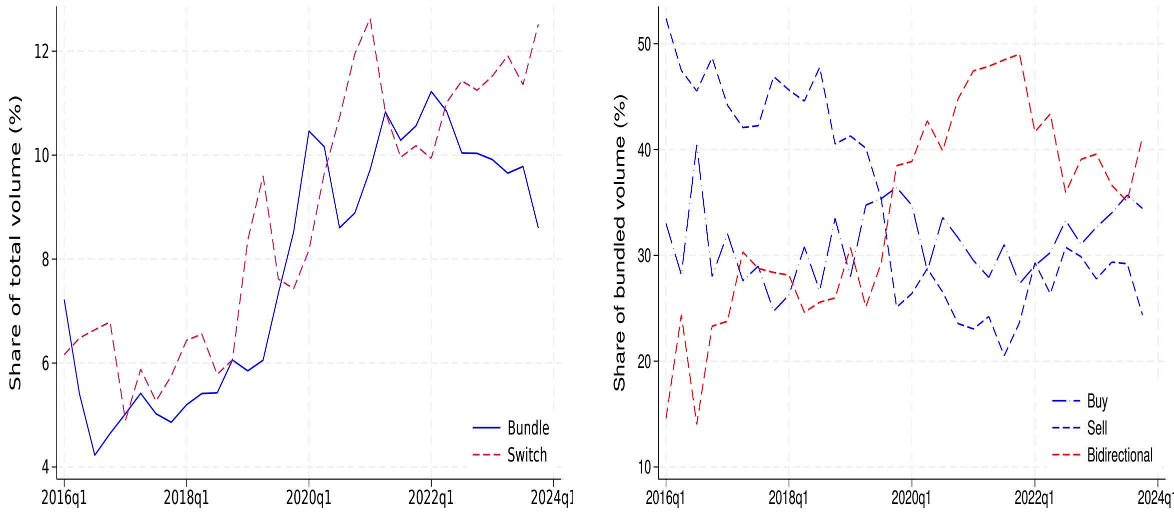

(A) Bundles vs. switches

(B) Types of bundles

Note: Figure 3 panel (A) plots the share of total volume that are bundles (blue solid line) and switches (red dotted line). Bundles are defined as any combination of buying and selling between the same dealer and client in a 5 second window, excluding switches. Switches are trades between the same dealer and client in a 5 second window that involves exactly one buy and one sell. Panel (B) decomposes the bundles into the three types: buy, sell, and bi-directional. Bi-directional trades involve at least 3 assets and not all can be on the same side.

TABLE 5. Trading with most-favored dealer

	(1)	(2)	(3)
I(bundle trade)	0.045***	0.062***	0.042***
	(0.002)	(0.001)	(0.001)
Constant	0.482***	0.462***	0.544***
	(0.001)	(0.000)	(0.000)
R2	0.379	0.571	0.834
Obs.	6,646,752	6,333,414	2,661,910
Client FE	Y	N	N
Client-day FE	N	Y	N
Client-day-bond FE	N	N	Y

Note: Table 5 shows the linear regression coefficients from regressing an indicator variable equal to 1 if a client trades with their most-favored dealer and 0 otherwise on an indicator variable equal to 1 if the trade is bundled and 0 if it's a single-asset trade. The columns allow for different fixed effects. Standard errors are clustered at the date level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

FIGURE 4. Changes in net quantity in switches

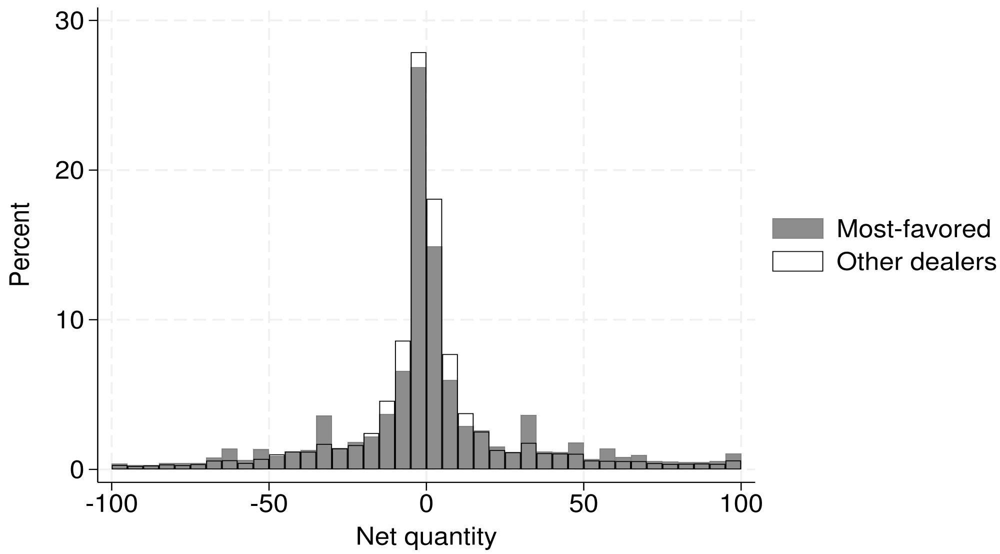

Note: Figure 4 plots net quantity exchanged in a switch with the most-favorite dealer (in gray) and dealers other than the most-favorite (in white). Net quantity is defined as the buy quantity minus the sell quantity as a fraction of total quantity exchanged. Each bin is 5%.

scenarios (e.g., Bao et al., 2018). Unlike agency and riskless principal trades, bundle trades can impact dealer balance sheets. Even if a dealer is able to offload a position within a day, they face intraday position limits and greater execution risk from bundled trades than agency trades—markets can move between when they acquire (or short-sell) a bond and sell (buy) it on the inter-dealer market or find a different client.

To support the idea that bi-directional bundle trades—in particular switches—are popular thanks to their limited impact on the dealers' balance sheet, we show in Figure 4 that the net change in the dealer's inventory (as a percentage of the amount of the trade) is close to zero for most switches.[18] Moreover, Figure 5 panel (A) shows that most switches are within the same asset class—for instance, switching Government of Canada bonds.[19] This type of trade leaves the characteristics of the dealer's balance sheet largely unaffected, and hence there is no change in the balance sheet constraints a dealer faces (risk-weighted capital and leverage requirements). These facts together imply that dealers pay negligible inventory costs, if any, to intermediate switches—an appealing property that holds regardless of their relationship with the client. When a

[18]Appendix Figure F5 displays a similar, slightly weaker pattern for bi-directional bundles.

[19]Appendix Table F6 provides insight into the specific types of assets involved in switches. In half of the eight asset classes, we find that within-asset class switches occur the most frequently, and within the other half, across-asset class switches are more predominant.

client conducts a switch (or bi-directional bundle) with their most-favored dealer the amount of balance space used is larger; consistent with dealers being constrained and favoring certain clients (see Figure 4).

To shed additional light on how dealer balance sheets relate to client relationships, we conduct a difference-in-differences analysis in Appendix E. The idea is to exploit a temporary change in balance sheet requirements that affected some dealers (treated) but not others (control). The findings, with some important caveats, support the notion that balance sheet constraints influence dealer–client interactions, with dealers prioritizing clients with whom they have close relationships.

Dealer balance sheets, however, are certainly not the only consideration for trade bundling; client preferences also play a role. For example, during the COVID-19 pandemic, we observe a preference for safety as clients purchased Government of Canada bonds and sold riskier asset classes.[20] Moreover, client preferences, particularly a preferred habitat for specific asset classes, could explain why the majority of uni-directional bundles are within an asset class (shown in Figure 5 panel (B)), even though they have a sizable impact on the dealer's balance sheet. In line with this idea, over 37 percent of clients who bundle assets within a single class trade exclusively in that class, indicating a preferred habitat for it.

Fact 3. *Trade bundling, particularly in the form of switches, is common on electronic platforms.*

Figure 6 plots the share of bilateral trade volume that are bundles and switches versus the share of bundles and switches on a platform. We see that a strikingly large fraction of trade volume on the electronic platform, about 26 percent, comes from switches.

To determine why switches are over-represented on the platform, we collect information about the services the platform provides to investors for switches. Appendix Figure F7 provides a wireframe of a hypothetical request for a switch. The platform, for example, allows an investor to request a switch to be 'par-for-par', that is, an exchange of securities with the same par value regardless of their market price. This is typically

[20]Appendix Figure F6 plots the share of switches over time that are Government of Canada for a corporate bond and vice versa. The share of switches that exchange a corporate bond for government bond nearly doubles with the COVID-19 pandemic. In unreported results, we show this is mainly switches with an investors most-favorite dealer. This result is consistent with Carlin et al., 2007 and the role of relationships during flight to safety episodes.

FIGURE 5. Share of within-asset class switches and bundles over time

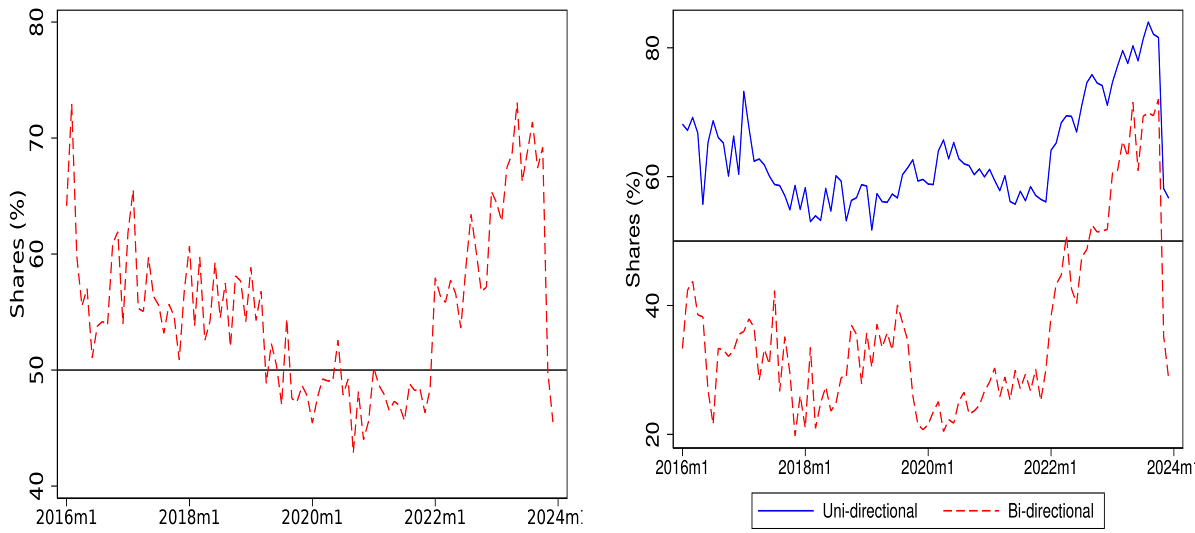

(A) Switches (B) Bundles

Note: Figure 5 panel (A) plots the share of switches that are within an asset class, as opposed to across asset classes. Panel (B) shows the share of uni-directional bundles (all buy or all sell bundles) versus bi-directional bundles (mix of buy and sell) that are made up entirely of a single asset class as opposed to a mix of assets from different asset classes. We combine all-buys and all-sells since their share of trades that are within-asset class are nearly identical. The large increase in within-asset class trades in 2023 comes almost entirely from an increase in trades within our foreign sovereign bond asset class, specifically U.S. Treasuries.

used when swapping similar securities with the same maturity and credit rating. Alternatively, if a switch leads to a change in credit risk or duration risk ('risk-adjusted switch'), the platform automatically provides an adjustment to how much cash needs to exchange hands (in addition to the bonds)—what is called 'net settlement' (based on indicative quotes). When the transaction is completed, the platform also provides an audit email providing all of the details of the trade. These conveniences can be one reason why we observe so many switches on the platform.

In contrast to switches, bundles are not that popular on the platform. One reason is that some of the platform services just described are tailored to switches. Another reason is that foreign sovereign bonds—which are often included in bundles—cannot be traded there. As a result, an investor wishing to buy a bundle of U.S. Treasuries and Government of Canada bonds does so bilaterally. In addition, there is endogenous selection by asset class: less liquid bonds, as shown in the prior literature, tend to trade off-platform. In our data (recall Table 2), corporate bonds and non-Canadian

FIGURE 6. Where do clients trade bundles?

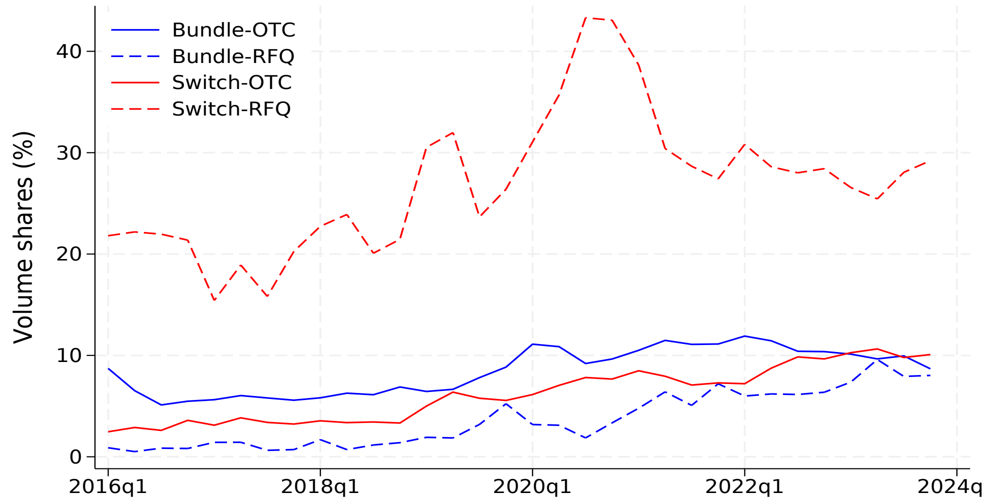

Note: Figure 6 plots the share, based on volume, of all bilateral trades that are bundles and switches (solid lines) and the share, based on volume, of all electronic platform trades that are bundles and switches (dashed line).

agency bonds—which are also often included in bundles—are only rarely traded on the platform. While bundling could, in principle, mitigate liquidity-driven selection, our evidence indicates it does not, at least not substantially. This holds even though, as we show later, bundles are cheaper on the platform—suggesting that selection effects remain strong, even when moving beyond single-asset trades.

Having documented the rising popularity and analyzed the nature of bundle trades, we examine their pricing implications to shed some light on the value of bundle trades and, by extension, the value of relationships. Theoretically, it is ambiguous whether bundle trades are priced higher or lower than single-asset trades, as pricing depends on several factors. We focus on three factors: the dealer's balance sheet costs, their market power relative to the investor, and the investor's preference for the convenience of bundling trades over executing them separately. The relative importance of these factors likely varies based on the type of bundle trade and the method of execution.

Fact 4. *Investors face lower transaction costs when bundling trades with their most-favored dealer compared to single-asset transactions.*

To establish this fact, we separately estimate regression (3) for trade bundles and switches. For example, for bundles:

$$\text{cost}_k = \beta \text{I(asset } k \text{ is part of a bundle)}_k + \gamma_k \text{log(quantity)} + \text{fixed-effects} + \epsilon_k. \quad (3)$$

We are interested in whether or not the transaction costs for trades that are part of a bundle (switch) have lower costs than trades that are not part of a bundle (switch). Throughout, we control for log-trade quantity, and confirm Pinter et al., 2024's finding for UK corporate bonds and gilts that trade-size discounts become markups when identified from within-client rather than cross-client variation.

To highlight the impact of dealer-client relationships on costs, we also estimate the regression with an additional indicator variable, I(most-favored), which equals 1 if the trade is with the client's most favored dealer and 0 otherwise, based on our relationship measure (1). Our focus is on its interaction with the bundle-trade dummy. For average trades, our findings are consistent with Issa and Jarnecic, 2019 and Allen and Wittwer, 2023, who document relationship markups.[21]

Table 6 reports estimation results. We include bond-day fixed effects in columns (1), (2), and (4), and add dealer-client fixed effects in columns (3) and (5). This is our preferred specification, since it captures the within-relationship effect of bundles and switches on transaction costs, purging out day-bond variation.[22]

We find that switches are consistently cheaper than single-asset trades across all specifications. A coefficient of -0.249 bps, for example, represents a 22 percent reduction in average trading costs (1.10 bps). Since switches are less balance sheet-intensive than most other trades, this suggests that dealers pass some of the resulting balance sheet cost savings on to clients. This is especially true when the switch is conducted with the most-favorite dealer (as shown from the interaction term in column (2)).

In contrast, bundles—typically uni-directional—are not consistently cheaper than single-asset trades. On average, they are more expensive, as indicated by the positive bundle coefficient in columns (1) and (2). This is intuitive: unlike most switches, uni-directional bundles do not have a limited balance sheet impact. Instead, they are predominantly sell-side, with clients selling more than they buy, leaving dealers to either absorb the bonds onto their balance sheet or find another buyer before settlement.

[21]Evidence on relationship-pricing is mixed, and depends on the market. For instance, Jurkatis et al., 2022 find relationship discounts between dealers and clients in the European corporate bond market.

[22]In unreported results, we found that our estimates when only controlling for dealer-client fixed effects are similar to those reported here that has both day-bond and dealer-client fixed effects.

However, bundles with the most-favorite dealer are cheaper. We see this from the negative interaction terms for trading a bundle with the most-favored dealer in column (2). Moreover, when controlling for the relationship status in form of dealer-client fixed effects, bundles are also cheaper.

Next, we try to unpack some of the other factors that influence bundle pricing, beyond dealer balance sheet costs, with a focus on the investors' taste for convenience and dealer market power. The idea for this is to leverage differences in the way bonds are traded—bilaterally via negotiations, or on an electronic platform.

Fact 5. *On the platform, switches cost more, whereas bundles are cheaper than in bilateral trading.*

From Allen and Wittwer, 2023 we know that government bond trades are, on average, cheaper on than off the platform—a result we confirm holds, on average, across asset classes. Given that private information plays a smaller role for the pricing of government debt than for other bonds, Allen and Wittwer, 2023 interpret this finding as evidence that dealers more strongly compete for clients on the platform. However, since clients tend to trade with their relationship dealers(s), the competition benefit is limited—a result that generalizes across asset classes.

In this paper, we analyze bundle trades. Here, there is a trade-off that does not exist for single-asset trades. Some bundle trades are attractive for dealers, because they are associated with lower balance sheet costs, so that dealers have incentives to offer these trades at a discount. Greater competition among dealers could amplify the discount. However, conditional on a client wanting to trade with a dealer, for example, because of their existing relationship, the dealer might (still) charge a markup, especially if it is convenient for the client to conduct the bundle trade all at once. Which effect dominates should depend on whether the trade is executed bilaterally via private negotiation, or on an electronic platform, since both the nature of competition and the services provided differ across these trading venues.

To find out whether pricing of bundle trades differs across venues, and provide evidence for our hypotheses, we estimate a similar regression to equation (3), but now we include indicator variables for whether the trade is executed on the platform (or not) and the interaction between the platform-indicator and the bundle-trade-indicator.

Table 7, which is similar to Table 6, shows the estimation results separately for switches (in columns (3) and (4)), and for bundles (in columns (5) and (6)). Columns (1) and (2) present the average platform discount across all trades, first with day-bond fixed effects and second adding dealer-client fixed effects. In columns (3)-(6) we include day-bond-platform fixed effects to control for selection onto the platform. Since we have already addressed the overall pricing differences between bundles, switches, and

TABLE 6. Transaction costs of bundle trades

	(1)	(2)	(3)
	Panel (A) - switches		
I(switch)	-0.275***	-0.204***	-0.249***
	(0.052)	(0.062)	(0.076)
I(most-favored)		0.186***	
		(0.042)	
I(most-favored) $\times$ I(switch)		-0.154**	
		(0.073)	
log(quantity)	-0.013	-0.013	0.088***
	(0.011)	(0.011)	(0.013)
Constant	0.787***	0.687***	-0.667***
	(0.161)	(0.170)	(0.190)
R2	0.165	0.165	0.174
Obs.	3,854,865	3,854,865	3,853,449
	Panel (B) - bundles		
I(bundle)	0.155**	0.418***	-0.150**
	(0.076)	(0.159)	(0.059)
I(most-favored)		0.270***	
		(0.038)	
I(most-favored) $\times$ I(bundle)		-0.483***	
		(0.176)	
log(quantity)	-0.138***	-0.134***	0.043***
	(0.026)	(0.025)	(0.014)
Constant	2.678***	2.480***	0.268
	(0.352)	(0.339)	(0.197)
R2	0.186	0.186	0.205
Obs.	5,035,640	5,035,640	5,034,277
Day-Bond FE	Y	Y	Y
Dealer-Client FE	N	N	Y

Note: Table 6 reports results from estimating regression (3) for bundles and switches. The dependent variable is trading costs. Panel (A) compares switches to single-asset trades. Panel (B) compares bundles to single-asset trades. The independent variables depend on the column: I(bundle) is an indicator equal to 1 if the trade is a bundle trade and 0 otherwise, while I(switch) marks a switch. I(most-favored) indicates a trade with a client's most-favored dealer according to our relationship measure (1). In all regressions we control for trade size using log(quantity). Standard errors are double clustered at the date and bond level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

TABLE 7. Transaction costs of bundle trades on- versus off-platform

	(1)	(2)	(3)	(4)	(5)	(6)
I(platform)	-0.988***	-0.648***				
	(0.110)	(0.107)				
I(switch)			-0.591***	-0.530***		
			(0.082)	(0.119)		
I(platform) × I(switch)			0.812***	0.541***		
			(0.138)	(0.170)		
I(bundle)					0.269**	-0.120
					(0.123)	(0.086)
I(platform) × I(bundle)					-0.559***	-0.030
					(0.172)	(0.123)
log(quantity)	-0.193***	0.043***	-0.030**	0.098***	-0.194***	0.044**
	(0.034)	(0.016)	(0.015)	(0.017)	(0.033)	(0.017)
Constant	3.921***	0.588***	1.188***	-0.664***	3.729***	0.550**
	(0.494)	(0.223)	(0.220)	(0.242)	(0.434)	(0.235)
R2	0.203	0.226	0.224	0.234	0.251	0.275
Obs.	4,087,034	4,085,632	2,721,846	2,720,345	3,574,801	3,573,369
Day-bond FE	Y	Y	N	N	N	N
Day-bond-platform FE	N	N	Y	Y	Y	Y
Dealer-client FE	N	Y	N	Y	N	Y

Note: Table 7 presents regression results where the dependent variable is trading costs. I(platform) is an indicator variable equal to 1 if the trade is on the electronic platform and 0 otherwise. I(bundle) is an indicator variable equal to 1 if the trade is part of a bundle and 0 otherwise. I(switch) is an indicator variable equal to 1 if the trade is part of a switch and 0 if part of a single-asset trade. We control for trade size using log(quantity). Standard errors are clustered at the date and bond level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

single-asset trades, as well as the transaction cost advantages for an average platform trade, our discussion focuses on the interaction terms.

We find that switches are more expensive on the platform. This is notable, given that switches are generally attractive to dealers due to their low balance sheet costs. Combining this fact with the evidence in Figure 6 showing that switches are frequently executed on the electronic platform suggests that investors are willing to pay a premium for the convenience of executing switches on the platform. This premium is not driven by liquidity differences between bonds traded on and off the platform, as we control for bond-day-platform fixed effects. Thus, it is distinct from the classical liquidity premium.

The exact size of the convenience premium is difficult to determine, as transaction costs are influenced by multiple factors and represent equilibrium outcomes. However, the coefficient (on the platform–switch interaction term) is large, representing about half of the average transaction cost. This means that a client pays about 50 percent more when conducting switches on the platform than off it—a sizable relative difference, though moderate in dollar terms (C$ 416) given the small overall transaction costs.

In contrast to switches, bundles are cheaper on than off the platform.[23] Even though clients might potentially be willing to pay a convenience premium for platform services that facilitate bundling, the competition effect dominates, while the reverse holds for switches. We interpret this contrasting findings between bundles and switches as reflecting both the platform's convenient services that are specific to switches and dealers' willingness to compete on the platform.

The dealers' willingness to compete hinges on their concerns about adverse selection—a factor for which we provide suggestive evidence in Appendix Table D1. When investors seek to buy or sell multiple bonds across asset classes or from different issuers, it is less likely that their trades are motivated by private information about the miss-pricing of a specific bond. In such cases, dealers perceive lower adverse selection risk and are more willing to compete for the client. As a result, the competition effect dominates. This dynamic differs for switches, which are more likely to be influenced by private information about specific bonds.[24] Therefore, dealers are more cautious, and competition is weaker. Additionally, some switches require dealers to precisely or near-perfectly offset the investor's buy or sell order (recall Figure 4). This complexity can reduces the number of dealers participating in the RFQ process compared to single-asset trades, further weakening competition. Consequently, the convenience premium outweighs the competition effect, leading to higher transaction costs for switches.

Looking ahead, our findings have implications for the design of trading platforms and the evolution of bond trading. In traditional OTC markets, trade bundling is facilitated by close dealer-investor relationships that reduce search frictions, such as swapping assets with the same dealer over time. As electronic platforms and algorithms evolve, these mechanisms could be replaced by centralized services that replicate or enhance these benefits. Our results highlight investors' preference for convenience, suggesting that bilateral trading may give way to centralized trading, but only if platforms offer

[23]In unreported results, we also find that this is true for all three bundles types, but the largest cost-savings are for the client buys (dealer sell).

[24]Clients switching out of risky bonds and into safe bonds might signal that a client knows something the dealer does not about the market. Even within some types of safe assets, private information can play a role—see Kondor and Pinter, 2022 for the UK gilt market.

features like streamlined execution, and efficient transaction processing. A potential concern of trade bundling is that it can act as a barrier to dealer entry: new entrants must account for client preferences to bundle across asset classes and for how incumbent dealers adjust prices depending on their competitors. Policymakers should be mindful of whether incumbents use bundled pricing strategically to discourage entry.

5. Conclusion

Using a novel data set on the near-universe of fixed-income trades in Canada, we document that institutional investors trade multiple asset classes, and that they have long-term relationships with few dealers across these asset classes. We show that a large fraction of trades involve multiple assets, and argue that bundling trades, in particular, the associated balance sheet cost savings, might reinforce dealer-client relationships. We also document a convenience premium for implementing multi-asset trades on electronic RFQ platforms, which speaks to the broad literature on convenience-yields of safe assets. This literature highlights asset characteristics as main drivers for the convenience yield. Our evidence suggests that, in some cases, it might be less about the asset itself, and more about the way it is traded.

References

Allen, J. and Wittwer, M. (2023). Centralizing over-the-counter markets? *Journal of Political Economy*, 131(12):3310–3351.

Ashcraft, A. B. and Duffie, D. (2007). Systemic illiquidity in the Federal Funds market. *American Economic Review*, 97(2):221–225.

Bak-Hansen, M. and Sloth, D. (2024). Customers, dealers and salespeople: Managing relationships in over-the-counter markets. mimeo.

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

Barardehi, Y. H., Da, Z., and Warachka, M. (2024). The information in industry-neutral self-financed trades. *Journal of Financial and Quantitative Analysis*, 59(2):796–829.

Carlin, B., Lobo, M., and Viswanathan, S. (2007). Episodic liquidity crises: Cooperative and predatory trading. *The Journal of Finance*, 62(5):2235–2274.

Chen, D. and Duffie, D. (2021). Market fragmentation. *American Economic Review*, 111(7):247–2274.

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

Eisfeldt, A. L., Herskovic, B., and Liu, S. (2024). Interdealer price dispersion and intermediary capacity. NBER working paper 32998.

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

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

Hendershott, T., Li, D., Livdan, D., and Schürhoff, N. (2024). When failure is an option: Fragile liquidity in over-the-counter markets. *Journal of Financial Economics*, 157:103859.

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

Hollifield, B., Neklyudov, A., and Spatt, C. (2017). Bid-Ask spreads, trading networks, and the pricing of securitizations. *The Review of Financial Studies*, 30(9):3048–3085.

Issa, G. and Jarnecic, E. (2019). Effect of trading relationships on execution costs in low-information-asymmetry over-the-counter markets. *The Journal of Financial and Quantitative Analysis*, 54(6):2605–2634.

Jurkatis, S., Schrimpf, A., Todorov, K., and Vause, N. (2022). Relationship discounts in corporate bond trading. Working paper.

Kaldor, N. (1939). Speculation and economic stability. *The Review of Economic Studies*, 7:1–27.

Kargar, M., Lester, B., Plante, S., and Weill, P.-O. (2024). Sequential search for corporate bonds. Forthcoming, *The Journal of Finance*.

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

Krishnamurthy, A. (2002). The bond/old-bond spread. *Journal of Financial Economics*, 66(2):463–506.

Krishnamurthy, A. and Vissing-Jørgensen, A. (2012). The aggregate demand for Treasury debt. *Journal of Political Economy*, 120(2):233–267.

Li, D. and Schurhoff, N. (2019). Dealer networks. *The Journal of Finance*, 74(1):91–144.

Li, J., O'Hara, M., Rapp, A., and Zhou, X. (2023). Bond market illiquidity: Is portfolio trading the solution? Working paper.

Meli, J. and Todorova, Z. (2023). Portfolio trading in corporate bond markets. Working paper.

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

Pintér, G., Üslü, S., and Wijnandts, J.-C. (2025). Comparing search and intermediation frictions across markets. Working paper.

Pinter, G., Wang, C., and Zou, J. (2024). Size Discount and Size Penalty: Trading Costs in Bond Markets. *The Review of Financial Studies*, 37(7):2156–2190.

Rostek, M. and Yoon, J. H. (2021). Exchange design and efficiency. *Econometrica*, 89:2887–2928.

Rostek, M. and Yoon, J. H. (2025). Financial product design in decentralized markets. Forthcoming, *Journal of Political Economy*.

Salinger, M. (1995). A graphical analysis of bundling. *Journal of Business*, 68:85–98.

Stigler, G. (1963). United States v. Loew's inc.: A note on block-booking. *Supreme Court Review*, pages 152–157.

Uslü, S. (2019). Pricing and liquidity in decentralized asset markets. *Econometrica*, 87:2079–2140.

Vayanos, D. and Wang, T. (2007). Search and endogenous concentration of liquidity in asset markets. *Journal of Economic Theory*, 136:66–104.

Wang, C. (2023). The limits of multi-dealer platforms. *Journal of Financial Economics*, 149(3):434–450.

Whinston, M. D. (1990). Tying, foreclosure, and exclusion. *American Economic Review*, 80(4):837–859.

Wittwer, M. (2021). Connecting disconnected financial markets? *American Economic Journal: Microeconomics*, 13(1):252–82.

Wittwer, M. and Allen, J. (2025). Market power and capital constraints. Working paper.

Wittwer, M. and Uthemann, A. (2024). Market segmentation across stocks, bonds, and derivatives. Working paper.

Appendix A. Asset class definitions

Here we provide definitions for our asset class classification:

Asset class	**Description**
Government of Canada (GoC)	Interest-bearing bills and bonds issued by Canada and redeemed at maturity for par value. Maturity ranges from 90-days to 50 years.
Provincial debt (Provi)	Provincial bills and bonds (median tenor is 10 years) and debt securities issued by cities, counties, and other governmental entities to fund day-to-day obligations and finance capital projects.
Corporate, non-financial (Corp-NF)	Short-term debt and bonds from non-financial Canadian and foreign firms. Short-term median tenor < 1 year; bonds 1-99 years with median 10 years.
Corporate, financial (Corp-F)	Financial institution debt including bills, bonds, commercial papers, and bearer bonds/notes. Median tenor < 7 years.
Bankers' Acceptances (BA)	Bank commitment for future payment, substituting bank's creditworthiness for borrower's. Tradeable in secondary market. Median tenor < 1 year.
Foreign sovereign debt (Foreign)	Non-Canadian debt issued in Canada or non-Canadian dollar debt issued by Canadian firms. Includes U.S. Treasuries, Euro bonds, etc. Median tenor 5 years.
Canadian agency (CDN)	Includes National Housing Act-Mortgage-backed securities, i.e. MBS backed by insured residential mortgages) and CMB (Canada Mortgage Bonds; packaged MBS). MBS carries prepayment risk but no credit risk. CMB converts monthly amortizing cash flows into typical bond-like payments without prepayment risk.
Non-Canadian agency (Non-CDN)	Miscellaneous financing category; securities issued by Government-Sponsored Enterprises.

Appendix B. client-types

Our data do not come with client-types. We generate these by first searching each firms' official website and internal documents, as well as the regulatory organization to which it reports. Second, we use Bloomberg, Yahoo Finance, and CapEdge. If these do not yield any results, we search on FT.com, WSJ.com, Preqin.com, Investorx.ca, WhaleWisdom.com, Fundlibrary.com, Morningstar, aum13f.com, and privatefunddata.com.

Client-type	**Description**
Asset manager	Financial entity whose purpose is to manage assets (or investments) and offer investment advising services. Includes entities that manage multiple types of funds such as hedge fund, mutual fund.
Bank	A deposit-taking/accepting institution.
Broker-dealer	Financial entity whose purpose is to offer brokerage services.
Central bank	An institution that manages the currency and monetary policy of a country.
Hedge fund	Financial entity that is a hedge fund or a hedge fund manager. They have greater flexibility than other asset managers have take on leverage, short sell securities and use derivatives.
Investment bank	A bank that provides financial services for corporate and institutional clients.
MF+ETF	An investment fund that pools money from many investors to purchase securities. Includes mutual funds, mutual fund managers, exchange-traded funds, and exchange-traded fund managers.
Non-financial	Non-financial entity, including retail and real estate, among other groups.
Pension fund & insurance	Financial entity whose purpose is to manage investments (and provide services) related to pension, retirement, insurance, reinsurance, benefits, and superannuation funds.
Other	Financial entity that does not fall in any of the aforementioned classifiers (e.g., private equity, financial planner).

Appendix Figure B1. Fraction of trades by client-types and their parents

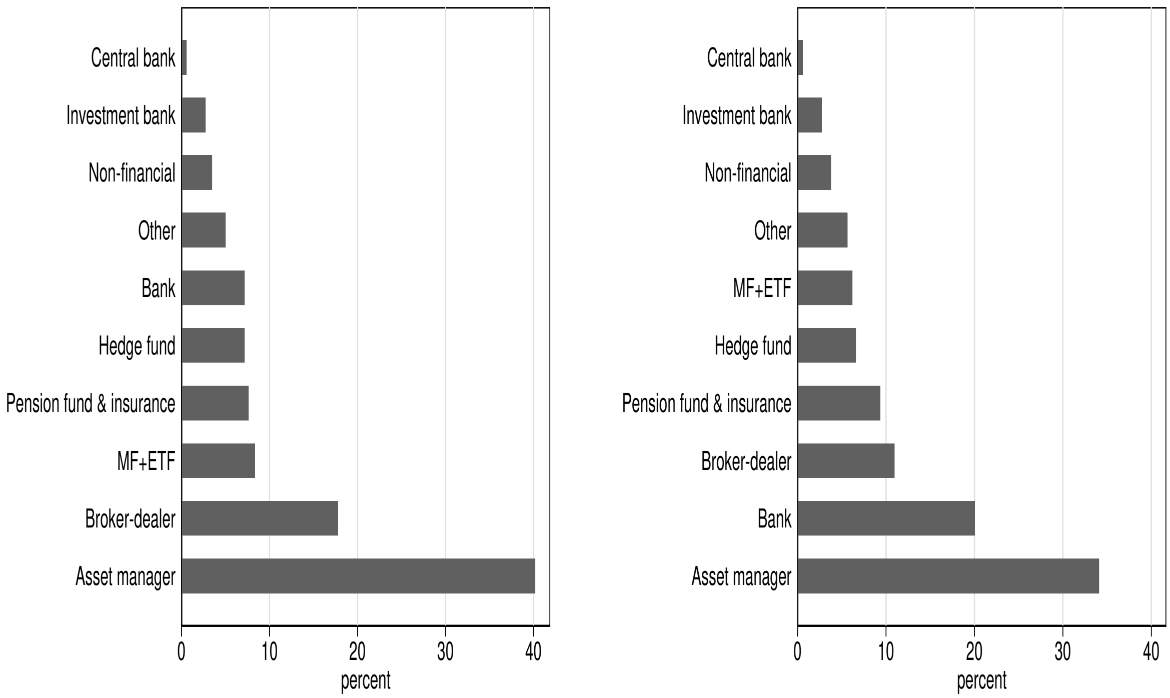

(A) Client-type (B) Client-parent type

Note: Appendix Figure B1 shows the fraction of dealer-to-client trades in our data by client-type in panel (A) and by the client's parent company type in panel (B). MF is an abbreviation for Mutual Fund while ETF is an abbreviation for Exchange-Traded Fund.

Appendix C. Trading strategies

Appendix Table C1 groups bundle trades into different trading strategies. We identify these based on characteristics of the trade, so they should be thought of as 'potential reasons for observing a bundled trade.' We report trade strategy frequencies by transaction count within bundles and switches (in columns (2) and (3), respectively) and as the share of trade volume within bundles and switches (in columns (4) and (5), respectively). Each column sums to 100 percent.

The most common bundles are uni-directional buys or sells. Investors may bundle purchases or sales for both strategic and non-strategic reasons. Strategic bundles can reflect directional bets—for example, buying multiple government bonds of different maturities to profit from an expected fall in interest rates. Non-strategic bundles can stem from portfolio rebalancing, tax loss harvesting, i.e., selling to realize losses for tax purposes, liquidity events (e.g., redemptions triggering broad sales), or cash deployment across several bonds. The median directional bundle contains four securities—well below FINRA's 10-bond threshold for a portfolio trade. This suggests that other mechanisms beyond take-it-or-leave-it offers of long bond lists play a role in the fixed income market.

Bi-directional trading strategies can be market 'neutral' or 'directional'. Specifically, we define a trade to be 'neutral' if the net notional position is close to 0 and 'directional' if the net notional amount is positive or negative. Our cut-off for 'close' is if the net quantity is less than 10 percent of the total trade size.

Common bi-directional strategies are 'rate', 'credit', and 'rich/cheap' trades. A 'rate' trade is a bet on yield curve shape movements. Examples include bets on the yield curve steeping, flattening, or twisting, as well as butterflies. An investor might sell $10 million of a 2-year Government of Canada bond with a coupon of 3% and buy $10 million of a 10-year Government of Canada bond with a coupon of 3.5%. The current spread is 50 bps and the bet is that the spread will widen by more—which would happen if the central bank cuts its policy rate but the long rate remains high for fears of inflation.

A 'credit' trade is one involving bonds of different credit. For example, an investor might sell $10 million of Government of Canada bonds and buy $10 million of a corporate bond when the spread is 200 bps. The trade is profitable if the corporate spread tightens, perhaps as the economy improves. Another example is a trade where an investor buys a technology firms' bond and sells a manufacturing firms' bond on the bet that technology firms will improve relative to manufacturing firms.

A 'rich/cheap' trade is a bet on the price relationship between the expensive bond and cheap bond converging to the fair value, but that is not classified as a rate trade. A classic example is on-the-run versus off-the-run government bond trades. If an investor believes that the off-the-run 5-year bond is relatively cheap, perhaps due to a liquidity premium for on-the-run bonds, they should buy off-the-run and sell on-the-run. Spreads will converge when the government issues a new 5-year bond and the current on-the-run bond becomes off-the-run. The profitability of such a trade, of course, depends on the cost of carry, which we do not observe and can be non-negligible (Krishnamurthy, 2002). The cost of carry might also depend on relationships, something we aim to explore in future work.

Trades involving foreign bonds or foreign currency ('international mixes') are less common and typically involve bets on sovereign spread changes. An investor, for example, expecting the Canada-U.S. government bond spread to narrow would buy Canadian bonds and sell U.S. Treasuries.

Appendix Table C1. Trading strategies for all bundled trades

	Transactions (%)		Volume(%)	
Trade type	Bundles	Switches	Bundles	Switches
Uni-directional-buy	43.0		31.76	
Uni-directional-sell	42.0		29.28	
Rates-neutral	1.52	16.88	5.24	19.13
Rates-directional	5.62	28.87	8.87	28.86
Credit-neutral	1.54	28.37	6.52	17.6
Credit-directional	5.29	9.99	14.07	3.81
Rich/Cheap-neutral	0.07	10.26	0.40	26.89
Rich/Cheap-directional	0.06	1.05	0.11	1.36
International mix	0.81	4.05	3.48	2.01
Other	0.13	0.52	0.27	0.36

Note: Appendix Table C1 lists the main strategies we classify in the data. Strategies are based on characteristics of the trade and not on any information provided by traders. Frequencies are presented in terms of transaction count and trade volume share. Uni-directional trades are those that result in all buy or all sell. Neutral trades are those that result in approximately the same value of sell and buy. Directional trades are bi-directional bundles that take a particular side of a bet—either more buy than sell or more sell than buy. The cut-off for neutral versus directional is 10 percent of net quantity (the amount the client buys minus sells, expressed in percentage of the total amount traded). Rate trades are bets on yield curve shape movements. Credit trades involve bonds of different credit quality. Rich/cheap trades is a bet on the relative value of an expensive bond and cheap bond converging to it's fair value over time. International mix trades is any trade involving a non-Canadian asset.

Appendix D. Adverse selection

Appendix Table D1 reports results of our adverse selection test. The main idea is to compare transaction costs of bundles that might be considered uninformed to bundles that might be driven by private information. If dealers are concerned about adverse selection when trading informed bundles, they should charge a higher price as compensation.

To distinguish uninformed from potentially informed bundles, we follow Li et al., 2023, who argue that portfolio trades are typically motivated by liquidity needs rather than private information. Using FINRA's definition—a portfolio trade involves at least ten bonds—we classify bundles exceeding this threshold as uninformed and those with fewer than ten bonds as potentially informed.

Specifically, in Appendix Table D1 we regress transaction costs on an indicator equal to one if the bundle contains at least ten unique bonds, and zero otherwise. Columns 'All' include all bundles, while columns 'Bi-d' restrict to bi-directional bundles. All specifications control for trade size and various fixed effects. Our preferred specification adds date-client-platform fixed effects to account for unobserved client characteristics and platform selection.

The main finding is that the coefficient on the indicator for bundles with at least ten bonds is negative, consistent with dealers being less concerned about adverse selection when a trade spans many bonds, and thus charging lower transaction costs. Adding date-client-platform fixed effects reduces the magnitude of this effect, likely reflecting both client heterogeneity and platform selection.

Appendix Table D1. Adverse selection

	All	Bi-d	All	Bi-d
I(≥ 10 bonds)	-1.785***	-1.157***	-0.500***	-0.443***
	(0.148)	(0.124)	(0.173)	(0.165)
log(quantity)	-0.566***	-0.453***	-0.240***	-0.021
	(0.022)	(0.029)	(0.021)	(0.029)
Constant	-3.039***	-1.809***	-0.614***	1.185***
	(0.177)	(0.217)	(0.168)	(0.227)
R2	0.018	0.018	0.230	0.175
Obs.	1417921	339240	1408443	338264
Date FE	Y	Y	N	N
Date-client-platform FE	N	N	Y	Y

Note: Appendix Table D1 presents our test for adverse selection. The columns represent different sample restrictions. 'All' includes all bundle trades and 'Bi-d' includes only bi-directional bundles. The dependent variable is transaction costs. I(≥10 bonds) is an indicator variable equal to 1 if the bundle has at least 10 unique assets and 0 otherwise. We control for trade size through log(quantity) and consider two different levels of fixed effects: date and date-client-platform. Standard errors are clustered at the date level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix E. Balance sheet space and relationships

In this appendix, we examine whether balance sheet constraints interact with trade bundling and client relationships. For this, we leverage the temporary exemption of Government of Canada bonds from the Basel III leverage ratio (LR) requirement between April 9, 2020, and December 21, 2021. This allowed affected dealers to hold these bonds without being constrained by the LR, and thus impacted their overall balance sheet space. Examining this exemption, Wittwer and Allen, 2025 show that the LR constraint, though infrequently binding, can have sizable effects in the primary market for government debt. In this paper, we use our secondary-market trading data from January 4, 2016, to December 31, 2023, to study whether balance sheet constraints influence dealers' willingness to trade and how this varies with the strength of their client relationships.

The main idea is to compare trading behavior of dealers who face the constraint and are affected by the temporary exemption of the LR constraint, with dealers who do not face this constraint, depending on the relationship status with the client. To implement this idea we employ a difference-in-difference estimation strategy; with the treated dealers being the eight dealers subject to the LR constraint—the largest dealers—, and the control group being the remaining dealers who do not have a LR requirement. We include trades across all asset classes, even though the policy change directly affected only government debt, because balance sheet constraints depend on the total assets held rather than solely on the directly affected bonds.

Appendix Table E1 reports difference-in-difference estimation results using different sample restrictions. Across all the dependent variable is net quantity—defined as 100 times the net amount traded (buy minus sell) divided by the total amount traded. For uni-directional trades this is always 100 (buy) or -100 (sell). For bi-directional trades it is somewhere in between (see Figure 4 and Appendix Figure F5). The independent variables are indicator variables and their interaction terms: I(most-favorite) indicates the most-favorite dealer; I(T) indicates dealers subjected to the leverage ratio constraint; and I(relief) indicates the exemption period (April 9, 2020 and December 21, 2021). The coefficients of interest are on the triple-interaction terms (I(T) $\times$ I(LR relief) $\times$ I(most-favorite)). We include client-day fixed effects to ensure that our results are not driven by variation across time or across clients, but instead is coming from variation across dealers who face different regulation.

The five columns represent different sample restrictions. In columns (1) and (2) we include all trades; columns (3) and (4) uses only bundles; column (5) is only switches. Furthermore, in column (1) and (3) we restrict the sample to trades where net quantity is greater than or equal to zero. That is, the net trade is a buy from the client's perspective (a sell from the dealer's perspective). In column (2) and (4) we restrict the sample to trades where the net quantity is negative, i.e. the net trade is a sell for the client (and buy for the dealer). Column (5) looks only at switches.

When a client is a net-seller (columns 2 and 4), the coefficient of interest—the interaction of I(T), I(relief), and I(most-favored) is positive, indicating that clients sold more to their most-favored dealer—relative to other dealers—when that dealer received a temporary exemption from the LR regulation. This implies that dealers who are buying expand their balance sheet for clients who have them as their most-favored dealer when the LR exemption came into effect.[25]

In contrast, when a client is a net-buyer (column 1 and 3), the interaction term is negative. This says that when a client bought from a dealer who was given a temporary exemption from the LR regulation, they bought less with their most-favored dealer than other dealers; possibly because these dealers prioritized buying from other clients.

Finally, when the client switches (column 5) the balance sheet impact can be net zero or directional—more buy than sell or vice versa. The negative coefficient says that dealers who received an LR exemption increased the amount they bought and decreased the amount they sold to clients who have them as their most-favored dealer—in line with the direction of the effects in the other columns.

Taken together, the results indicate that balance sheet constraints shape how dealers interact with clients and support the view that dealers prioritize those with whom they have close relationships. That said, there are several confounding factors, that warrant mention. First, dealers subject to the Basel III leverage constraint differ from those exempt not only in regulatory treatment but also in size, business model, and capital allocation strategies—the constrained group consists of the eight largest Canadian dealers—so size differences could drive the observed relationships. Second, the leverage-ratio exemption coincided with an unprecedented period of market stress and policy intervention during COVID-19, making it difficult to disentangle the effects of LR relief

[25]Note that the coefficient on I(most-favorite) is positive when the client is a net-seller and negative when the client is a net-buyer. This means that clients use more of their favorite dealer's balance sheet than other dealers. This is somewhat muted for the dealers who face the LR constraint—the interaction of I(T) and I(most-favorite)—suggesting that the LR constraint might be binding.

from those of contemporaneous shocks, even with client-day fixed effects. Third, our analysis focuses on traded quantities rather than prices, limiting our ability to capture the full impact of balance sheet constraints. Lastly, potential spillovers between constrained and unconstrained dealers—such as strategic responses by the latter to their competitors' regulatory relief—could bias our estimates.

Appendix Table E1. Balance sheet space and relationships

	Net q≥0	Net q<0	Bundles Net q≥0	Bundles Net q<0	Switches
I(most-favorite)	-3.350***	3.812***	-12.551***	4.854***	0.664
	(0.454)	(0.407)	(2.154)	(1.546)	(0.921)
I(T)	-3.102***	3.073***	-1.501***	2.304***	1.137***
	(0.164)	(0.153)	(0.529)	(0.427)	(0.356)
I(T) × I(most-favorite)	1.838***	-1.853***	8.627***	-0.137	0.508
	(0.478)	(0.444)	(2.197)	(1.618)	(1.021)
I(relief) × I(most-favorite)	4.501***	-1.446	13.663***	-7.725	3.499*
	(0.791)	(0.901)	(4.838)	(4.957)	(1.872)
I(T) × I(relief)	0.130	0.182	0.859	-1.513**	1.554**
	(0.268)	(0.242)	(0.967)	(0.743)	(0.663)
I(T) × I(relief) × I(most-favorite)	-6.118***	3.881***	-12.990***	11.365**	-4.103**
	(0.847)	(0.955)	(4.906)	(5.006)	(2.051)
Constant	94.055***	-95.205***	89.297***	-92.041***	-0.939***
	(0.121)	(0.108)	(0.413)	(0.319)	(0.273)
R2	0.435	0.432	0.658	0.639	0.521
Obs.	3340929	2805344	869043	810713	377862

Note: Appendix Table E1 presents results from our difference-in-difference specification. The dependent variable is net quantity (the amount the client buys minus sells), expressed in percentage of the total amount traded. In columns (1) and (2) we include all trades, columns (3) and (4) focus on bundles, and column (5) is switches. Further, in columns (1) and (3) we restrict the sample to trades where the net quantity is at least zero, and in columns (2) and (4) net quantities are negative. The indicator I(most-favorite) equals 1 if the trade is made with the client's most-favored dealer and 0 otherwise. The indicator I(T) is equal to 1 for the banks subjected to the leverage ratio constraint. The indicator I(relief) equals 1 when government of Canada bonds were exempt from the leverage ratio calculation and 0 otherwise. We include client-day fixed effects. Standard errors are clustered at the date level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix F. Additional tables and figures

This appendix displays all appendix tables and figures we referred to in the main text.

Appendix Table F1. Asset characteristics

Variable	GoC	Provi	CDN	Non-CDN	Corp-NF	Corp-F	BA	Foreign
Unique # of ISIN	474	16,528	113	6,584	22,289	69,821	13,436	4,745
Tenure (yrs)	3.82	3.77	6.9	3.78	5.33	2.38	0.32	7.36
CAD (%)	96.8	98.2	100	51.5	59.6	57	79.1	0
I(domicile=Canada)	100	100	100	61	66.3	80.4	96.7	0
I(callable) (%)	0	0	0	0.01	25.5	6.6	0	0.03
Coupon type (%)								
Plain vanilla	28	65.2	69	43.1	43.2	16.8	0.4	55.8
Fixed zero coupon	0	22.6	0	53.8	55.6	64.2	99.6	2.7
Fixed discount	69.8	11.3	0	0.2	0	0	0	39.4
Other	1.9	0.01	31	3.0	1.2	17.7	0	2.1

Note: Appendix Table F1 reports key bond statistics based on characteristics at issuance. Tenor is the time to maturity at issuance. CAD is the share of bonds that are in Canadian dollars. I(domicile=Canada) is an indicator for whether the domicile (the country where the issuing entity has a legal address) is Canadian or not. I(callable) is an indicator equal to 1 if the bond is callable and 0 otherwise. The last four rows are the percentage of coupon types. Plain vanilla bonds offer fixed coupon payments determined at issuance. A zero-coupon bond is an investment in debt that does not pay interest. Fixed discount means the bond trades on a discount basis rather than on a yield basis—bills and commercial paper. The only asset class where 'other' is non-negligible is corporate (financial). These are almost entirely fixed-then-floating coupons.

Appendix Table F2. Number of asset classes by client-type

	Probit		Ordered Probit	
Client-type	coeff.	std. error	coeff.	std. error
Bank	-0.172***	0.058	-0.200***	0.069
Broker-dealer	-0.255***	0.063	-0.302***	0.072
Central bank	-0.302***	0.092	-0.467***	0.093
Hedge fund	-0.488***	0.063	-0.553***	0.072
Investment bank	-0.160	0.101	-0.253**	0.114
Non-financial	-0.666***	0.076	-0.661***	0.089
Other	-0.424***	0.063	-0.482***	0.069
Pension fund & insurance	-0.182***	0.068	-0.127	0.084
MF+ETF	-0.397***	0.057	-0.416***	0.070
Constant	0.421***	0.048		
R2 (pseudo)	0.0193		0.0125	
Obs.	135,438		135,438	

Note: Appendix Table F2 reports Probit and ordered Probit results for the number of asset classes by client-type. Specifically, in columns (2) and (3), we estimate a Probit model where the dependent variable is equal to 1 if a client trades more than one asset class in a month and 0 otherwise; and an ordered Probit model where the count variable is the number of asset classes traded in a month by a client in columns (4) and (5). The sample period is January 4, 2016 to December 31, 2023. The omitted category is asset manager. MF+ETF is mutual funds and exchange-traded funds. Regressions include year fixed effects. Standard errors are clustered at the client level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Table F3. Number of asset classes by parent client-type

	Probit		Ordered Probit	
Parent client-type	coeff.	std. error	coeff.	std. error
Bank	-0.014	0.067	-0.028	0.084
Broker-dealer	-0.231***	0.080	-0.306***	0.085
Central bank	-0.281***	0.095	-0.472***	0.096
Hedge fund	-0.446***	0.078	-0.526***	0.090
Investment bank	-0.071	0.146	-0.086	0.190
MF+ETF	0.082	0.095	0.047	0.107
Non-financial	-0.592***	0.082	-0.614***	0.096
Other	-0.431***	0.071	-0.516***	0.078
Pension fund & insurance	-0.073	0.073	-0.049	0.088
Constant	0.452***	0.055		
R2 (pseudo)	0.0254		0.0167	
Obs.	98,919		98,919	

Note: Appendix Table F3 reports Probit and ordered Probit results for the number of asset classes by parent client-type. Specifically, we estimate a Probit model where the dependent variable is equal to 1 if a parent trades more than one asset class in a month and 0 otherwise; and an ordered Probit model where the count variable is the number of asset classes traded in a month by a client-parent. The sample period is January 4, 2016 to December 31, 2023. The omitted category is asset manager. MF+ETF is mutual funds and exchange-traded funds. Regressions include year fixed effects. Standard errors are clustered at the parent level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Table F4. Trading with the most-favored dealer across versus within asset classes

	(1)	(2)
I(across asset class)	0.054***	0.021***
	(0.002)	(0.002)
Constant	0.462***	0.438***
	(0.000)	(0.001)
R2	0.572	0.596
Obs.	5,224,089	480,187

Note: Appendix Table F4 shows the linear regression coefficients from regressing an indicator variable equal to 1 if a client trades with their most-favored dealer and 0 otherwise on an indicator variable equal to 1 if the trade is a mix of asset classes and 0 if it's a trade within an asset class. Column (1) is for all trades and column (2) are bundle trades. By definition, all single-asset trades are within an asset class. Each observation is one trade, i.e., a single-asset trade or a bundle trade if it involves multiple assets. Regressions include client-day fixed effects. Standard errors are clustered at the date level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Table F5. Who conducts bundle trades?

	I(bundle)		I(switch)	
Client-type	coeff.	std. error	coeff.	std. error
Bank	-0.718***	0.078	-0.137	0.157
Broker-dealer	-0.334***	0.097	0.077	0.091
Central bank	-1.402***	0.123	0.065	0.187
Hedge fund	-0.450***	0.081	0.607***	0.082
Investment bank	-0.822***	0.187	-0.387**	0.142
Non-financial	-1.017***	0.118	-0.389***	0.098
Other	-0.580***	0.126	-0.123	0.095
Pension fund & insurance	-0.611***	0.126	0.324**	0.115
MF+ETF	-0.112	0.075	-0.074	0.130
Constant	-1.742***	0.074	-2.014***	0.057
R2 (psuedo)	0.0517		0.0406	
Obs.	5,011,237		4,841,504	

Note: Appendix Table F5 reports Probit regression results for the probability of trading a bundle and the probability of conducting a switch. I(bundle) is an indicator variable equal to 1 if the trade is a bundle and 0 otherwise. I(switch) is an indicator variable equal to 1 if the trade is a switch and 0 if it is a single-asset trade. The sample period is January 4, 2016 to December 31, 2023. The omitted category is asset manager. MF+ETF is mutual funds and exchange-traded funds. Regressions include year fixed effects. Standard errors are clustered at the client level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Table F6. Buy and sell in a switch

Sale/purchase	GoC	Provi	Corp-NF	Corp-F	Foreign	CDN	BA	Non-CDN
GoC	59.5	15.6	8.8	11.2	0.24	3.6	0.2	0.8
Provi	45.1	46.2	1.0	11.2	2.1	2.8	0.23	0.8
Corp-NF	70.4	2.2	17.2	6.0	2.2	0.5	0.13	0.4
Corp-F	68.6	2.2	5.0	16.3	6.5	0.6	0.5	0.4
Foreign	0.6	1.1	1.1	2.9	90.8	0.0	0.01	3.5
CDN	48.1	12.1	1.0	1.7	0.02	36.7	0.1	0.3
BA	14.1	9.5	14.4	18.1	0.4	0.4	42.0	1.2
Non-CDN	19.3	9.3	1.5	1.9	52.2	0.4	0.1	15.3

Note: Appendix Table F6 shows what is bought and sold in a switch. The sale of an asset is in the rows and the purchase of assets in the columns. The rows add up to 100%. Therefore, when a client sells a GoC bond, they buy a GoC bond in 59.5% of cases, buy a provincial bond in 15.6% of cases, etc. Thus, for GoC bonds we conclude that clients are more likely to switch within the same asset class than across asset classes.

Appendix Table F7. Bundle trades with the most-favored dealer

(A) Switches			
I(switch)	0.016***	0.026***	0.018***
	(0.001)	(0.001)	(0.002)
Constant	0.491***	0.472***	0.546***
	(0.001)	(0.000)	(0.000)
R2	0.385	0.580	0.832
Obs.	4,966,988	4,644,101	1,805,455
Client FE	Y	N	N
Client-day FE	N	Y	N
Client-day-bond FE	N	N	Y

(B) Bundles			
I(bundle)	0.055***	0.074***	0.052***
	(0.002)	(0.001)	(0.002)
Constant	0.482***	0.462***	0.553***
	(0.001)	(0.000)	(0.001)
R2	0.381	0.577	0.848
Obs.	6,268,888	5,949,605	2,421,739
Client FE	Y	N	N
Client-day FE	N	Y	N
Client-day-bond FE	N	N	Y

Note: Appendix Table 5 shows the coefficients from regressing a dummy variable equal to 1 if a client trades with their most-favored dealer and 0 otherwise on a dummy variable equal to 1 if the trade is bundled and 0 if it's a single-asset trades. The columns allow for different fixed effects. The top panel bundle compares switches to single-asset trades and the bottom panel compares bundles to single-asset trades. The coefficient on I(switch) is relatively small because many switches are on CanDeal, where the most-favored dealer is chosen less frequently. In unreported results, we show that focusing only on bilateral trades the coefficient on I(switch) is basically the same as I(bundle). Standard errors are clustered at the date level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Figure F1. Cumulative distribution function of trade activity

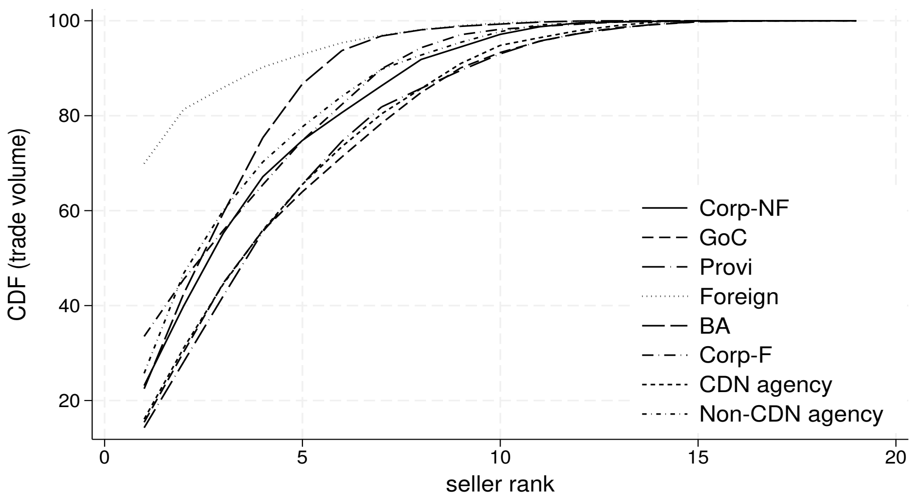

Note: Appendix Figure F1 is generated by sorting sellers by their trade volume from largest to smallest and graphing the cumulative distribution function (CDF) of trade volume. The CDF is similar if we sort by number of transactions instead of volume. Asset class definitions are provided in Appendix A. The most concentrated market is the one for foreign sovereign bonds. The least concentrated markets are those for Government of Canada and corporate bonds.

Appendix Figure F2. Average trade size by asset class

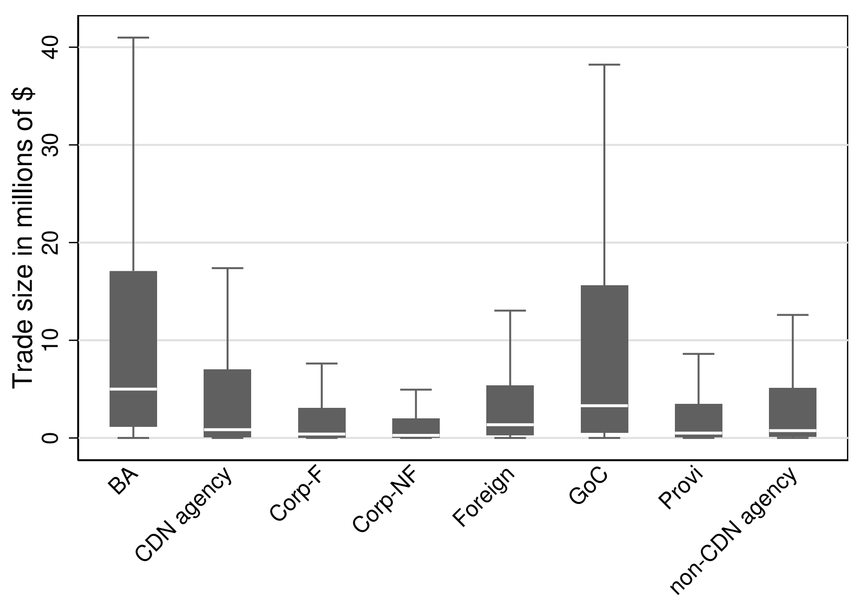

Note: Appendix Figure F2 plots the inter-quartile range of trade sizes for each asset class in our data set, excluding outliers. The average trade size is large—$10.88 million, but this is driven by the Government of Canada (GoC) and Bankers' Acceptances (BA) market. The median trade size is $1.3 million.

Appendix Figure F3. Transaction costs for different benchmark prices

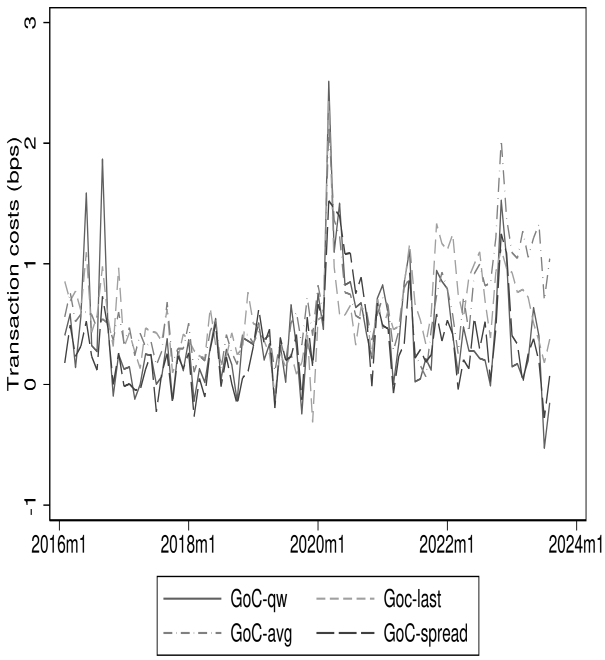

(A) Government of Canada

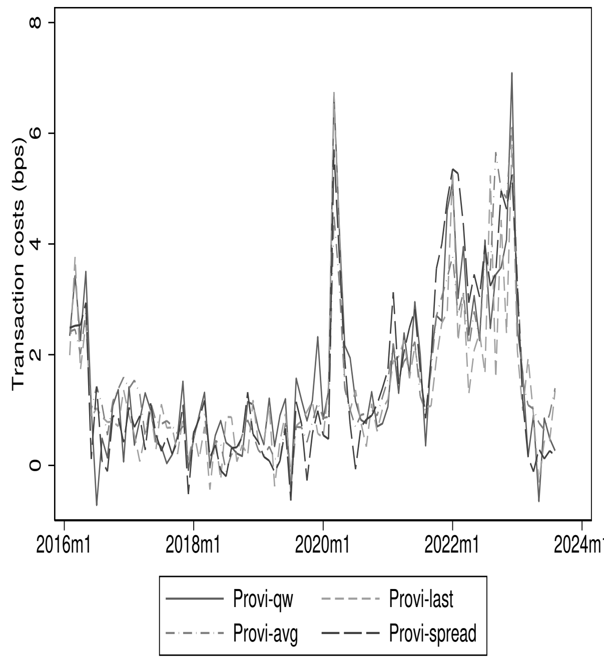

(B) Provincial

Note: Appendix Figure F3 Panel (A) plots the monthly average transaction costs for Government of Canada debt and Panel (B) does the same for provincial debt. We plot four different measures of transaction costs: (i) '-qw' is using the quantity-weighted daily average of D2D trades (per ISIN) as the benchmark price in equation (2); (ii) '-avg' uses a simple daily average as the benchmark price per ISIN (our baseline); (iii) '-last' uses the last observed transaction price of an ISIN as the benchmark; and (iv) 'spread' uses the daily buy-sell spread per ISIN. Transaction costs using the quantity-weighted D2D daily price as the benchmark are less volatile than other measures, however, the general patterns are similar.

Appendix Figure F4. Transaction costs per asset class and client-type

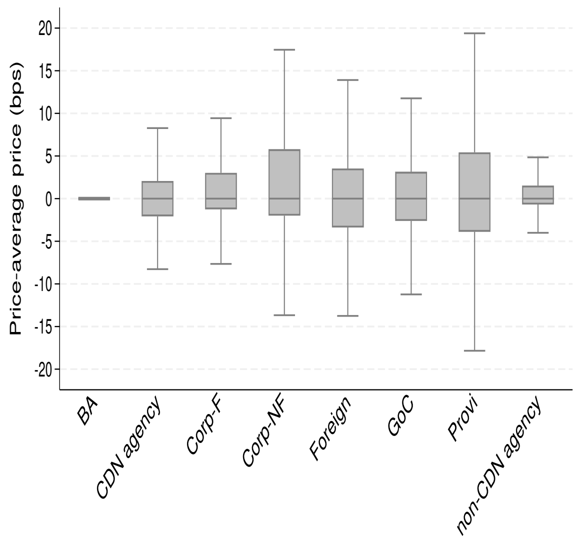

(A) Asset class

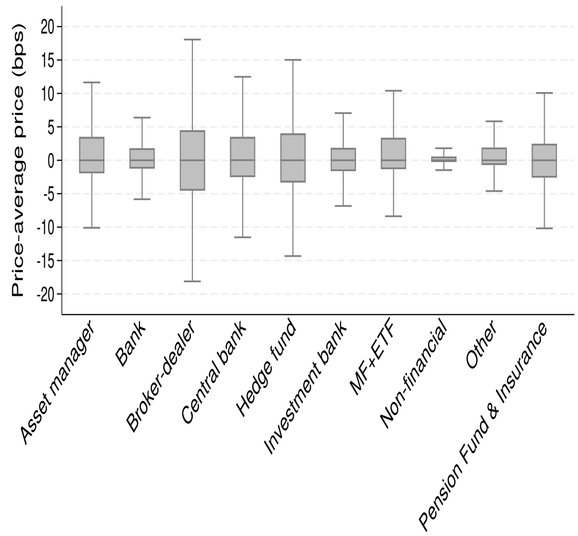

(B) Client

Note: Appendix Figure F4 plots the inter-quartile range of transaction costs for each asset class (panel A) and client-type (panel B) excluding outliers.

Appendix Figure F5. Changes in net quantity in bi-directional bundles

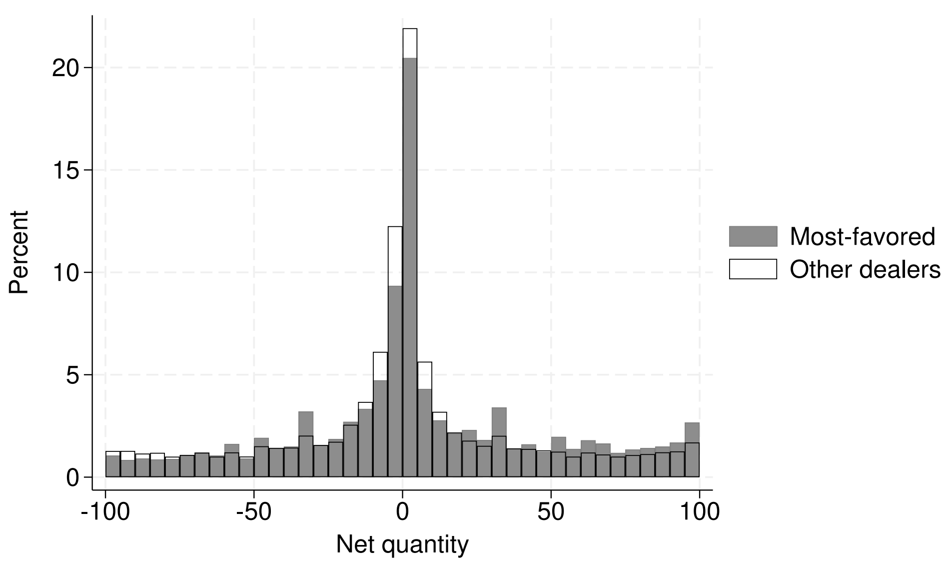

Note: Appendix Figure F5 plots the change in net quantity in bi-directional bundle with the most-favorite dealer (in gray) and all other dealers (in white). Net quantity is defined as the buy quantity minus the sell quantity as a fraction of total quantity exchanged. Each bin is 5%.

Appendix Figure F6. Government of Canada—Corporate switch

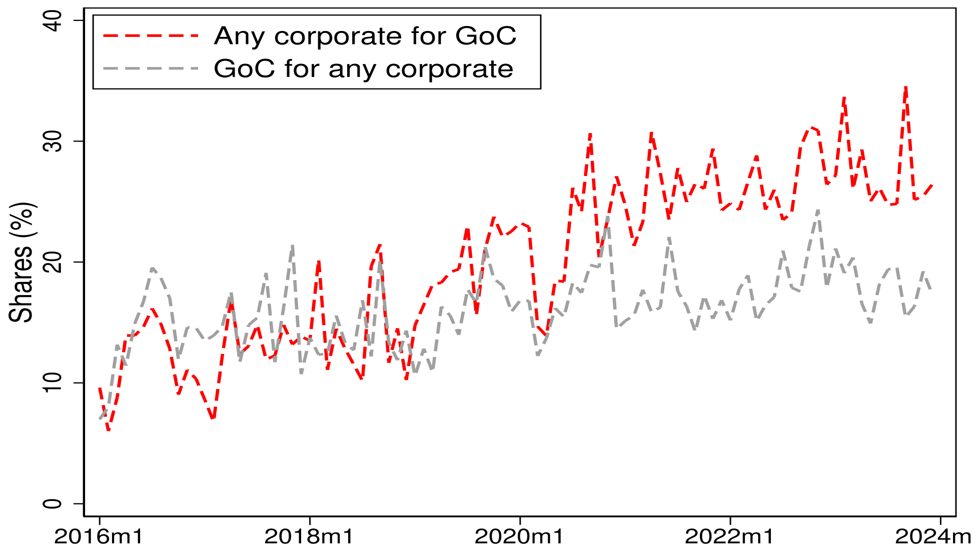

Note: Appendix Figure F6 plots the time series of the share of switches that exchange Government of Canada bonds to any corporate bond, and vice-versa.

Appendix Figure F7. How to conduct a switch on CanDeal

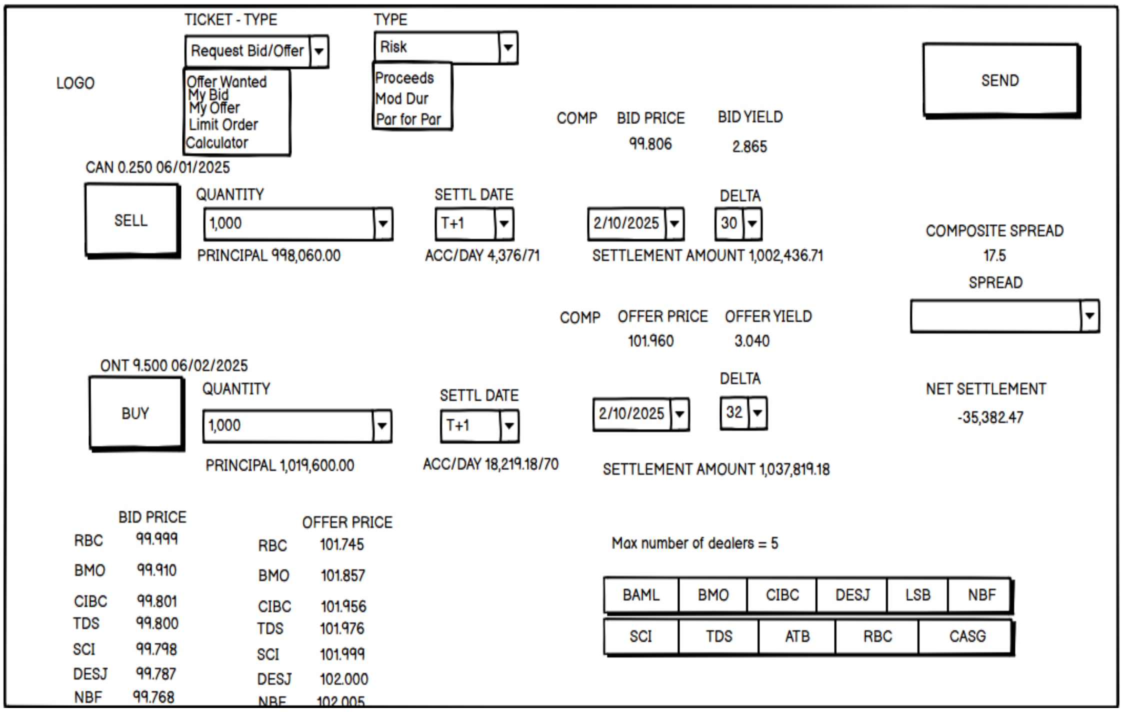

Note: Appendix Figure F7 is a depiction of what clients see on their CanDeal platform for a hypothetical switch trade. Although a screenshot was not possible, this graphic captures everything a client sees and was generously provided to us by CanDeal. In the SELL section is 1,000 units of a Government of Canada bond. In the BUY section is 1,000 units of a province of Ontario bond. The typical ticket type on the top left corer is 'Request Bid/Offer'. The different types of trades are: (i) Risk, (ii) Proceeds, (iii) Mod Dur, and (iv) Par for Par. Risk means that the quantities of the securities the client wants to trade are calculated in order to obtain an equal 'delta' (price sensitivity) for each bond. Proceeds means that the quantity of one leg is calculated in order to obtain a settlement amount for each bond that is almost equal. Mod Dur means that the quantity of one leg is calculated in order to obtain a duration weighted amount for each bond that is almost equal. Par-for-Par means that the buy quantity is equal to the sell quantity. This graphic also shows that the maximum number of dealers that can be invited is 5 and provides a list of dealers as well as their indicative bid and offer prices.

Appendix G. Robustness

We first present robustness results for our main results on transaction costs using the dealer-to-dealer benchmark price. This is similar to Hendershott and Madhavan, 2015, however, we use the quantity-weighted dealer-to-dealer price to account for the substantial heterogeneity in trade sizes and prices in the inter-dealer market (c.f., Eisfeldt et al., 2024). Appendix Table G1 replicates Table 6 and Appendix Table G2 replicates Table 7. The results are similar, even though the number of transactions drops substantially when using the inter-dealer price as our benchmark price. The reason is that some asset classes, for example, BAs, are infrequently traded on the inter-dealer market, and therefore we do not have a benchmark price. See Table 2.

Second, we conduct a robustness analysis to address the concern that the choice to buy and sell different asset classes is endogenous to the characteristics of the market. Our comparison of transaction costs across asset classes, therefore, suffers from selective entry. To partially offset this, we present results for a set of 95 clients who are the most active across the Canadian fixed income market. They account for approximately 40 percent of all trade. Appendix Tables G3 and G4 reports results for this sample of most active investors. The discount and premium estimates are somewhat larger in this sample of active investors than in the full sample.

Appendix Table G1. Transaction costs of bundle trades—Robustness with inter-dealer price as benchmark

	(1)	(2)	(3)
		Panel (A) - switches	
I(switch)	-0.475***	-0.341***	-0.522***
	(0.088)	(0.085)	(0.122)
I(most-favored)		0.198***	
		(0.073)	
I(most-favored) × I(switch)		-0.318**	
		(0.126)	
log(quantity)	-0.003	-0.003	0.140***
	(0.016)	(0.016)	(0.024)
Constant	0.552**	0.466*	-1.562***
	(0.244)	(0.257)	(0.349)
R2	0.123	0.123	0.132
Obs.	1,812,288	1,8122,88	1,811,046
		Panel (B) - bundles	
I(bundle)	0.309**	0.616***	-0.008
	(0.124)	(0.218)	(0.109)
I(most-favored)		0.254***	
		(0.066)	
I(most-favored) × I(bundle)		-0.614**	
		(0.250)	
log(quantity)	-0.143***	-0.137***	0.097***
	(0.037)	(0.035)	(0.024)
Constant	2.693***	2.499***	-0.635*
	(0.511)	(0.491)	(0.342)
R2	0.151	0.151	0.168
Obs.	2,126,988	2,126,988	2,125,798
Day-Bond FE	Y	Y	Y
Dealer-client FE	N	N	Y

Note: Appendix Table G1 reports regression results from estimating equation 2 for bundles and switches. The dependent variable is trading costs with the benchmark price being the quantity-weighted average (signed) trading price in the interdealer market, per day-ISIN. Panel (A) compares switches to single-asset trades. Panel (B) compares bundles to single-asset trades. Standard errors are double clustered at the date and bond level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Table G2. Transaction cost of bundle trades on- versus off-platform—Robustness with inter-dealer price as benchmark

	(1)	(2)	(3)	(4)	(5)	(6)
I(platform)	-0.969***	-0.789***				
	(0.135)	(0.173)				
I(switch)			-0.770***	-0.807***		
			(0.131)	(0.185)		
I(platform) × I(switch)			0.965***	0.755***		
			(0.172)	(0.218)		
I(bundle)					0.481**	0.187
					(0.188)	(0.136)
I(platform) × I(bundle)					-0.688***	-0.213
					(0.254)	(0.196)
log(quantity)	-0.179***	0.079***	-0.022	0.135***	-0.181***	0.088***
	(0.042)	(0.023)	(0.019)	(0.025)	(0.041)	(0.025)
Constant	3.605***	-0.119	0.865***	-1.457***	3.350***	-0.415
	(0.631)	(0.325)	(0.281)	(0.360)	(0.554)	(0.352)
R2	0.171	0.189	0.192	0.202	0.223	0.242
Obs.	2,031,864	2,030,741	1,492,822	1,491,578	1,753,755	1,752,573
Day-bond FE	Y	Y	N	N	N	N
Day-bond-platform FE	N	N	Y	Y	Y	Y
Dealer-client FE	N	Y	N	Y	N	Y

Note: Appendix Table G2 presents regression results where the dependent variable is trading costs with the benchmark price being defined as the quantity-weighted average (signed) trading price in the interdealer market, per day-ISIN. I(platform) is an indicator variable equal to 1 if the trade is on the electronic platform and 0 otherwise. I(bundle) is an indicator variable equal to 1 if the trade is part of a bundle and 0 otherwise. I(switch) is an indicator variable equal to 1 if the trade is part of a switch and 0 if part of a single-asset trade. We control for trade size using log(quantity). Standard errors are clustered at the date and bond level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Table G3. Transaction costs of bundle trades—Robustness using a sample of the 95 most active investors

	(1)	(2)	(3)
		Panel (A) - switch	
I(switch)	-0.488***	-0.437***	-0.507***
	(0.122)	(0.116)	(0.140)
I(most-favored)		0.025	
		(0.057)	
I(most-favored) × I(switch)		-0.203**	
		(0.099)	
log(quantity)	0.071***	0.071***	0.129***
	(0.021)	(0.021)	(0.020)
Constant	-0.404	-0.410	-1.236***
	(0.297)	(0.298)	(0.278)
R2	0.220	0.220	0.224
Obs.	1,445,036	1,445,036	1,445,016
		Panel (B) - bundles	
I(bundle)	0.200**	0.522***	-0.204***
	(0.089)	(0.153)	(0.066)
I(most-favored)		0.061	
		(0.053)	
I(most-favored) × I(bundle)		-0.903***	
		(0.229)	
log(quantity)	-0.106**	-0.097**	0.083***
	(0.048)	(0.046)	(0.019)
Constant	2.230***	2.088***	-0.178
	(0.635)	(0.607)	(0.259)
R2	0.247	0.247	0.269
Obs.	2,106,686	2,106,686	2,106,667
Day-Bond FE	Y	Y	Y
Dealer-Client FE	N	N	Y

Note: Appendix Table G3 reports regression results from estimating equation 2 for bundles and switches. The sample is the 95 most active clients and all dealers. The dependent variable is trading costs with the benchmark price being the average (signed) trading price per day-ISIN. Panel (A) compares switches to single-asset trades. Panel (B) compares bundles to single-asset trades. Standard errors are double clustered at the date and bond level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Table G4. Transaction costs for bundle trades on- versus off-platform—Robustness using a sample of the 95 most active investors

	(1)	(2)	(3)	(4)	(5)	(6)
I(platform)	-0.985***	-0.653***				
	(0.148)	(0.108)				
I(switch)			-0.765***	-0.802***		
			(0.170)	(0.199)		
I(platform) ×I(switch)			0.830***	0.794***		
			(0.193)	(0.219)		
I(bundle)					0.256**	-0.246***
					(0.119)	(0.084)
I(platform) × I(bundle)					-0.518***	0.033
					(0.189)	(0.141)
log(quantity)	-0.124**	0.079***	0.055***	0.112***	-0.121**	0.075***
	(0.051)	(0.019)	(0.021)	(0.019)	(0.049)	(0.019)
Constant	2.836***	-0.016	-0.153	-0.970***	2.533***	0.024
	(0.734)	(0.256)	(0.297)	(0.275)	(0.652)	(0.258)
R2	0.230	0.251	0.259	0.263	0.288	0.308
Obs.	2242209	2242193	1293965	1293940	1899400	1899379
Day-bond FE	Y	Y	N	N	N	N
Day-bond-platform FE	N	N	Y	Y	Y	Y
Dealer-client FE	N	Y	N	Y	N	Y

Note: Appendix Table G4 presents regression results where the dependent variable is trading costs with the benchmark price being the average (signed) trading price per day-ISIN. I(platform) is an indicator variable equal to 1 if the trade is on the electronic platform and 0 otherwise. I(bundle) is an indicator variable equal to 1 if the trade is part of a bundle and 0 otherwise. I(switch) is an indicator variable equal to 1 if the trade is part of a switch and 0 if part of a single-asset trade. We control for trade size using log(quantity). Standard errors are clustered at the date and bond level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.