BUNDLING TRADES IN OVER-THE-COUNTER MARKETS

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

Date: March 1, 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. An exception is Pintér and Üslü, 2022, who focus on a sub-sample of dealers and clients that actively participate in both the U.K. gilt and corporate bond markets. They document connections and substantial differences across these two markets, prompting a broader analysis of cross-sectional trading patterns in fixed-income markets.

In this paper, 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, oftentimes 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 between a dealer and their client 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 six key facts that offer new insights into dealer-client interactions in OTC markets over time and across asset classes. To preview: Facts 1 and 2 lay the foundation for exploring bundle trades as a possible source of dealer-client relationships (Fact 3). We then examine how these trades are executed—whether bilaterally or an electronic platform (Fact 4)—and analyze their pricing implications (Facts 5 and 6).

Our first fact is that clients are active in more than one asset class, and therefore might seek to engage in trading strategies that span across these classes. To quantify cross-asset market activities, we classify assets into eight categories (or classes), including Canadian government debt, foreign sovereign debt, and non-financial corporate bonds. We document that about 40 percent of clients are active in at least four asset classes and that the average client is active in two. Multi-category clients tend to be asset managers and investment banks, while hedge funds, mutual funds, and exchange-traded funds (ETF) often only trade assets within a single asset class.

Second, investors tend to trade with a small set of dealers—not only over time, but also across asset classes—highlighting the importance of dealer-client relationships. This is consistent with the existing literature on relationships in OTC markets (e.g., Di Maggio et al., 2017; Hendershott et al., 2020; Jurkatis et al., 2022; Pintér and Üslü, 2022; Allen and Wittwer, 2023; Pinter et al., 2024). We add to this literature by demonstrating that clients tend to rely on a single dealer, even when considering trades across all major fixed-income asset classes.

Third, bundling trades is common: investors often engage in multi-asset trading strategies with the same dealer—nearly 20 percent of all trades involve multiple assets. We identify three typical strategies: (*i*) an investor simultaneously buying one asset and selling another ('switch'), (*ii*) an investor simultaneously buying and selling multiple assets ('bi-directional bundle'), and (*iii*) an investor simultaneously buying (or selling) multiple assets ('unidirectional bundle'). Other permutations exist but are less frequent.

Bi-directional trades, which minimally impact the dealer's balance sheet as investors simultaneously buy and sell assets—often through switches—have gained popularity over unidirectional bundle trades, where investors exclusively buy or sell multiple assets. For certain asset classes, such as government debt, switches are more frequently conducted within the same asset class, whereas for other asset classes switches often span across asset classes. For example, during the COVID-19 pandemic, many clients switched corporate bonds for government bonds, underscoring the role of cross-asset switches as a mechanism for flight to safety.

Fourth, 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 disparity reflects the platform's ability to support trading strategies involving multiple assets, offering specialized services that streamline such complex transactions.[1] In other words, investors seem to have a taste for convenience, at least when bundling trades. As electronic platforms gain traction in OTC markets, we expect trade bundling to play an increasingly prominent role.

Fifth, 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

[1]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).

accrued to dealers in these trades are passed on to clients, in particular, their repeat clients.[2]

Finally, for the sixth fact, we compare trades made on electronic RFQ platforms with those made through direct negotiations to better understand what influences bundle pricing, beyond just 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 (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 that simplify multi-asset trades, which can increase transaction costs. We find cost-advantages for clients when bundling unidirectional 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.

[2]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.

Besides contributing to the literature on OTC markets and dealer-client relationships, our study intersects with other broader research areas. It connects to the industrial organization literature on product bundling, rooted in foundational work by Stigler, 1963, Whinston, 1990, and Salinger, 1995. Our contribution to this literature emphasizes the importance of bundling trades—across asset classes and over time—in decentralized OTC markets. This adds to a handful of papers that document trading of multiple bonds within the same asset class, namely corporate bonds, at one moment in time (Li et al., 2023; Meli and Todorova, 2023). The empirical evidence 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 contribute to the small but growing literature on trade models incorporating bundling in centralized markets (e.g., Wittwer, 2021; Rostek and Yoon, 2021; Chen and Duffie, 2021; Rostek and Yoon, 2025).

In addition, our study also 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). Our contribution to this literature highlights 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 cleaning and some sample selection choices, detailed in Appendix A, 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 Table Appendix D1) 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: (*i*) Government of Canada debt (GoC), (*ii*) Canadian provincial and municipal governments debt (Provi), (*iii*) non-financial corporate bonds, (*iv*) financial corporate bonds, (*v*) foreign sovereign debt, (*vi*) Banker Acceptances (BA), (*vii*) Canadian agency debt, and (*viii*) non-Canadian agency debt.[3] Detailed asset class definitions are provided in Appendix B.

Table 1 presents market shares by number of transactions and volume traded for the different asset classes, and some sub-categories. The most traded assets are Government of Canada bills and bonds, followed by foreign sovereign debt (mostly U.S. Treasuries), provincial debt, corporate debt, and BAs.

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.[4] 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 that have a LEI. This excludes less than 1 percent of trade volume executed with retail investors.

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

[3]We exclude two asset classes because they are so infrequently traded—supranational bonds and (non-mortgage) asset-backed securities. We also combine provincial bonds with 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.

[4]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.

TABLE 1. Secondary market trading shares by asset class

Shares	transactions (%)	volume (%)
Government of Canada (GoC)	34.52	55.73
Bills	3.83	6.42
Bonds	30.69	49.31
Canadian provincial governments (Provi)	13.94	9.32
Bills	1.12	1.01
Bonds	12.34	6.19
Corporate bonds (non-financial)	8.96	1.58
Corporate bonds (financial)	15.23	6.13
Foreign sovereign debt	17.53	15.94
U.S. Treasuries	16.34	16.99
Other	1.21	0.79
Bankers' Acceptances (BA)	4.51	7.42
Canadian agency debt	3.58	3.77
Non-Canadian Agency debt	1.73	1.04

Note: Table 1 presents the distribution of trades across the 8 asset classes in our sample, measured both by transaction count and trading volume percentages. Asset class definitions are provided in Appendix B.

FIGURE 1. Cumulative distribution function of trade activity

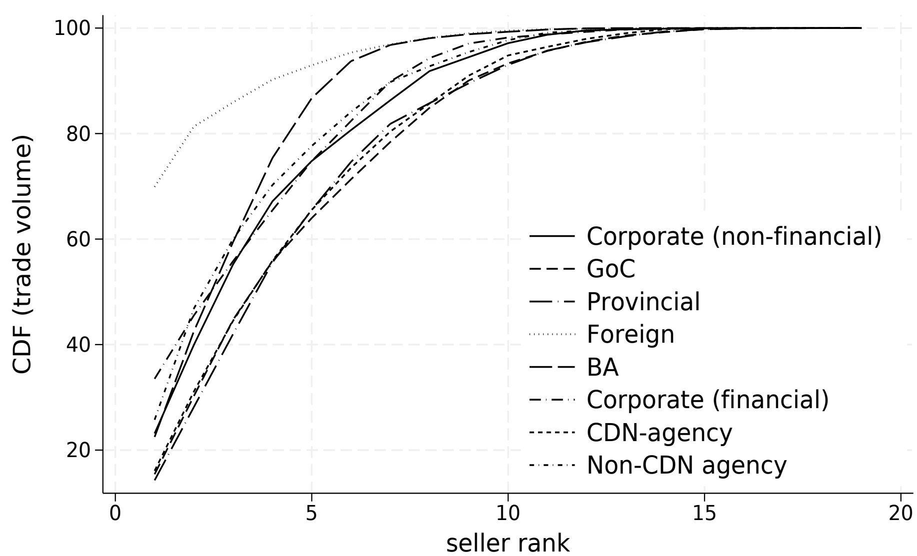

Note: Figure 1 is generated by sorting sellers by their trade volume from largest to smallest and graphing the cumulative distribution function (CDF) of trade volume. The CDF is similar if we sort by number of transactions instead of volume. Asset class definitions are provided in Appendix B.

and primary market for government debt. The fixed income market, however, is relatively concentrated. Figure 1 plots, for each asset class, the cumulative market share of dealers, ranked from largest to smallest share. Between 60 and 90 percent of trades are intermediated by the five largest dealers. The most concentrated market is the one for foreign sovereign bonds. The least concentrated markets are those for Government of Canada and corporate bonds.[5]

Clients. In addition to dealers, there are 7,982 clients with a unique LEI.[6] 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`.[7] 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 C. 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 C1.

Trading venues. For D2C trades, which are our focus, we observe whether a trade was conducted bilaterally with a dealer or electronically over a multilateral trading platform,

[5]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.

[6]We incorporate all institutional investors, regardless of how active they are across markets. There is clearly selection into who is and isn't active in multiple markets, and this (endogenous entry) is likely correlated with trading costs. We think that it is important, however, to document which types of firms trade across assets classes, how they trade, and implications for prices. In Appendix E we provide robustness results for our main contribution using a smaller set of 95 investors who trade the most volume across multiple asset classes.

[7]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'.

such as CanDeal. On the platform, an investor can send a request for quote to 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. The client selects the best offer, and the trade is executed shortly after.

Unlike bilateral negotiations, multilateral 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). 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 Figure 1).

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; 29 percent of Canadian government debt are traded electronically, which is consistent with Allen and Wittwer, 2023. Trade size also varies by asset class, as shown in Appendix Figure D1. The average D2C trade size is large—$10.88 million, but this is driven by the market for Canadian government debt and Banker's Acceptance. The median trade size is $1.3 million. For completeness, Table 2 also reports the share of transactions that are between related parties, which we drop from our analysis.

Table 2. Summary trading statistics (monthly)

Variable	GoC	Provi	CDN agency	Non-CDN agency	Corp. non-fin	Corp fin	BA	Foreign sov
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 (%)	29.0	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. 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.

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 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-favorite' to their 'least-favorite' dealer.

Multi-asset bundle trades. One possible reason for why we observe long-term relationships in OTC markets are synergies created from bundling trades. In our analysis, 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. Bundle trades can be initiated by clients and dealers alike. There are many types of strategies; we differentiate between two main types.

First, we define 'switches' as any trade between the same dealer and client within a five second window that involves exactly one buy and one sell. For example, an investor might have bought a 10-year Government of Canada bond when it was on-the-run (i.e., most-recently issued bond), and when that bond is no longer on-the-run, they wish to switch it for the most recent issuance.

Second, we define 'bundles' as all other types of multi-asset trading strategies. If the transaction involves all buy (sell) we call these 'buy (sell) direction bundles'. If they involve multiple buys and sells we call it a 'bi-directional bundle'.[8] There are countless types of trading strategies that involve bundles. For example, when investors need to re-balance their portfolio allocations, they have an incentive to bundle all required trades. Some of these 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').

Implementing all trades within a strategy simultaneously through a bundle trade is appealing to traders, as it eliminates the risk of partial execution—for instance, completing only one leg of a two-legged butterfly. Bundling also streamlines trade processing, settlement, and record-keeping by reducing the number of individual transactions, benefiting both clients and dealers. Additionally, as we will show, most switches and some bi-directional bundles have minimal impact on dealers' inventories, which is advantageous for managing balance sheet costs and capital requirements.

[8]Bundles are related, yet distinct from bond 'portfolio trading', which has been studied in the context of the U.S. corporate bond market by Meli and Todorova, 2023 and Li et al., 2023. Bond 'portfolio trades' in these papers are 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. To avoid confusion, we therefore do not use the term portfolio trade.

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

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

Here trade price is the transaction price for trade j; benchmark price is the average transaction price of the same bond on the same day; and trade sign is an indicator variable equal to 1 for an investor purchase and -1 for an investor sale.[9] 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 D3 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 and within client types.

Transaction costs for multi-asset 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 (CB) 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 CB 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 of all other investors on the same day. In the example, the average CB buy price is $(99.2 \times 500 + 99.3 \times 600)/1100 = 99.254$; the GoC average sell price is 100. Therefore, the investors saves (99.254-99)*1000+(100-100)*1100=$254 in transaction costs when bundling.

[9]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 average interdealer price and not the average across all trades. For robustness, in the Appendix we report results using the quantity-weighted average interdealer price. Appendix Figure D2 plots the time series of transaction costs based on different measures in the literature. Our four measures are highly correlated.

Table 3. Example of a multi-asset trading advantage

Investor	Asset class	Time	Side	Price	Quantity	Dealer
1	CB 123	9:05:01	buy	99	1000	A
1	GoC 456	9:05:01	sell	100	1100	A
2	CB 123	9:10:15	buy	99.2	500	A
3	GoC 456	9:11:22	sell	100	900	A
4	CB 123	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 corporate bond (CB) and selling 1,100 of Government of Canada bonds (GoC) simultaneously with the same dealer. In this example, investors 2,3, and 4 conduct single-asset trades.

This logic underpins the formal analysis we present in Section 4, where we establish our empirical findings on bundle trading. In this section, we follow Li et al., 2023 and estimate a fixed effects model, 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 six facts about bond trading, focusing on multi-asset strategies, relationships, and pricing. We begin with establishing that two facts about clients and dealers previously documented for a subset of OTC markets, generalizes when taking the entire market into account (Facts 1-2). Then we focus on our main contribution—analyzing bundle trading and its pricing implications with relationship dealers (Facts 3-6).

Fact 1. *Most investors trade multiple asset classes.*

Figure 2 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 D2.[10]

[10] If we aggregate at the holding company/parent level, the patterns are similar, with one exception: Appendix Table D3 shows that banks, as the ultimate parent of many funds and investment banks, are more similar to asset management companies.

FIGURE 2. Distribution of the number of asset classes that clients buy and sell

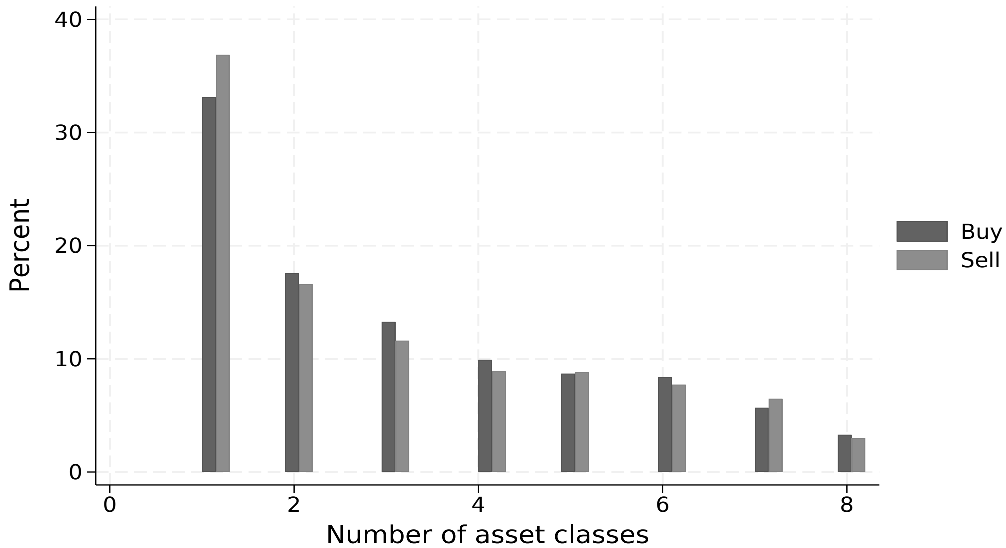

Note: Figure 2 shows the distribution of the number of asset classes across clients for when they buy and sell, separately.

Fact 2. *Most investors trade with a single dealer.*

This finding aligns with the existing literature, which has demonstrated similar patterns for subsets of fixed-income markets. 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-favorite dealers, averaged across clients. Conditional on having more than one relationship, 80.4% of trades are between a client and their most-favorite 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-favorite dealer and 14.7% for the second-favorite dealer. This reflects the fact that more active clients have more dealers; however, even the most active clients tend to have a large share of volume with a single dealer.[11]

[11]Appendix Table D4 shows that at the client level, banks, broker/dealers, asset managers, and pension fund & insurance companies all interact with a similar number of dealers, while hedge funds, mutual funds and ETFs are less likely to trade with multiple dealers. Aggregating clients at the parent level, however, only hedge funds and 'other' client-types have fewer dealers than asset managers.

TABLE 4. Mean trade-volume across most-favorite dealer-client relationships

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

Note: The first row of Table 4 shows the average trade share that clients execute with their most-favorite 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 aggregate trade shares with the two most favorite 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.

A natural question in light of this fact, in combination with the first fact—that most clients trade multiple asset classes—is to ask whether clients use the same dealer for all asset classes they are trading, or whether they segment dealers by asset classes, perhaps because of different dealer specializations. Nearly 40 percent of clients have the same number of dealers in an asset class as they have in total. However, in the median, some client-types have more dealers across than within assets classes. This is shown in Figure 3, which plots the distribution of number of dealers across asset classes in white and within asset class in gray. In conclusion, many clients are likely to have more dealers when they buy multiple asset classes than when trading in a single asset class.

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 these strategies (Facts 3-4) before examining their pricing implications (Facts 5-6).

Fact 3. *Investors engage in multi-asset trading strategies; especially switches, where investors simultaneously exchange one asset for another. For certain asset classes, such as government debt, switches typically occur within the same asset class, whereas for others, like corporate debt, they tend to span across asset classes.*

We start by establishing that multi-asset trading is common. Figure 4 panel (A) plots, in the red dashed line, the share of total volume traded that we classify as

Figure 3. Number of dealers per client type

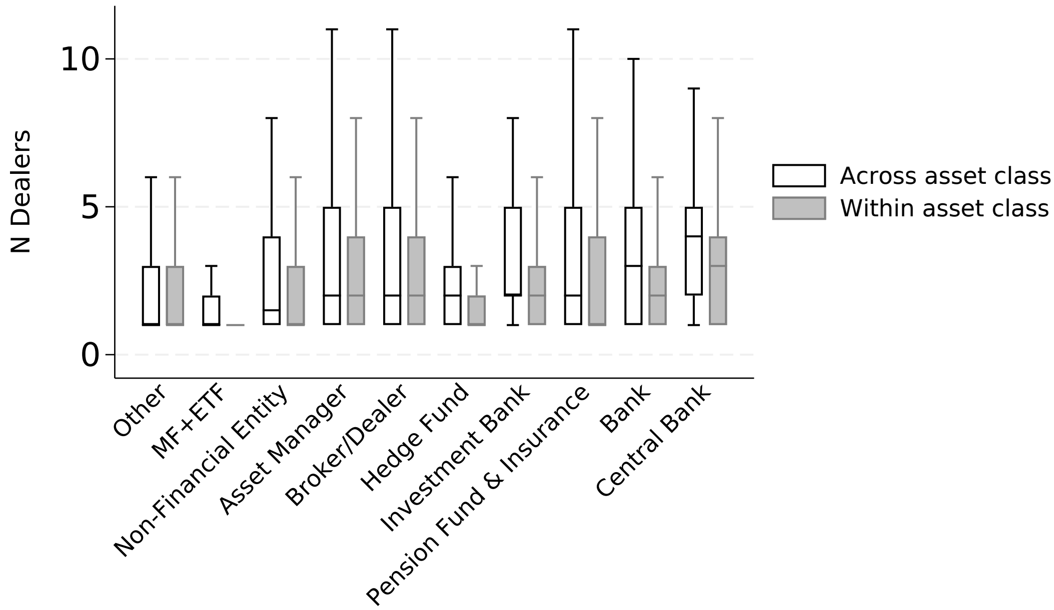

Note: Figure 3 plots the distribution of the number of dealers per client-type across asset classes in white and within asset class in gray. We exclude outliers. MF stands for mutual fund while ETF stands for Exchange-Traded Fund. The definition of client-types are provided in Appendix Table C1.

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.[12] This is a substantial proportion, especially considering our conservative definition of bundle trades, which requires transactions to occur within a strict 5-second window between the same dealer and client.

Over time, bi-directional bundle trades have become more popular than uni-directional trades. To show this, Figure 4 decomposes the bundles of panel (A) into: (i) directional (buy), directional (sell), and bi-directional. We see substantial time variation over our pooled sample of asset classes; however, since 2020, the largest share of bundle trades are bi-directional.

[12]In Appendix Table D5 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 Table D2). 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 4. Share of volume that is bundled

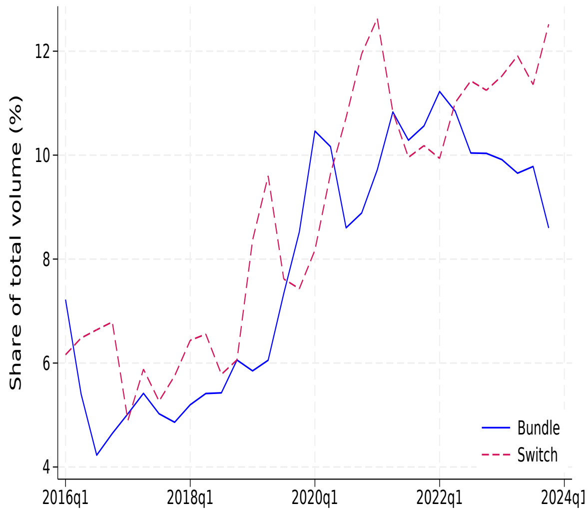

(A) Bundles vs. switches

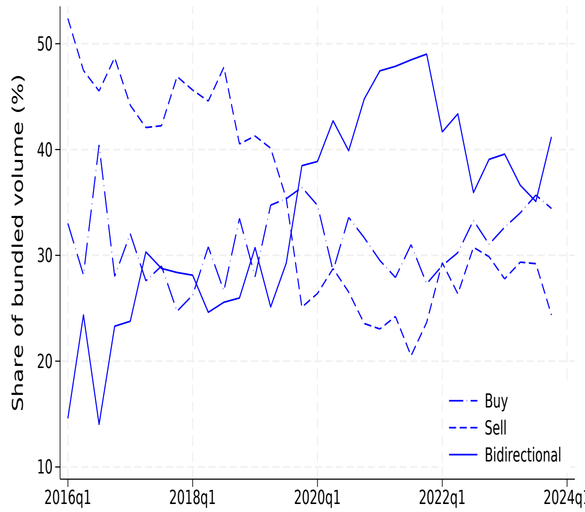

(B) Types of bundles

Note: Figure 4 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: (i) buy, (ii) sell, and (iii) bi-directional. Bi-directional trades involve at least 3 assets and not all can be on the same side.

FIGURE 5. Changes in net quantity in switches

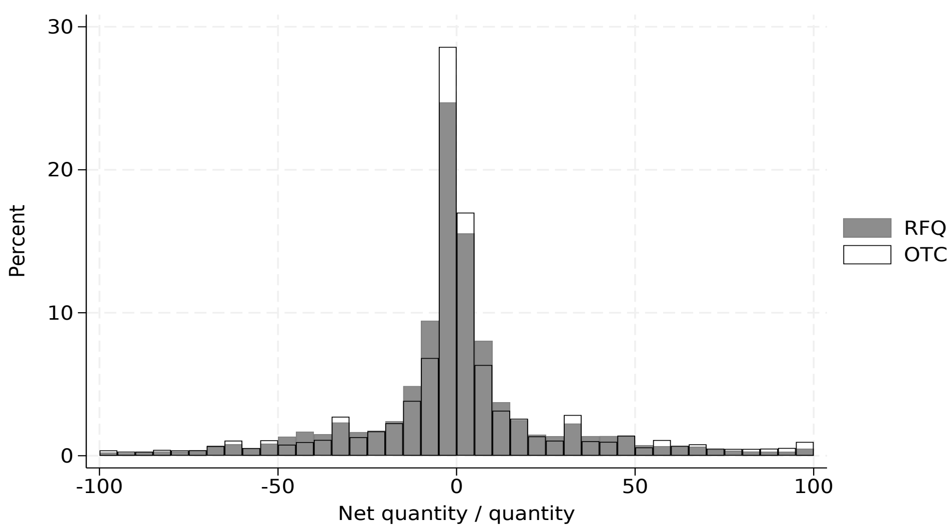

Note: Figure 5 plots net quantity exchanged in a switch, where net quantity is defined as the buy quantity minus the sell quantity as a fraction of total quantity exchanged. Each bin is 5%. In white are bilaterally negotiated switches (OTC), and in gray are switches on CanDeal.

One possible explanation for these trends 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 over principal trading in the U.S. corporate bond market (e.g., Bao et al., 2018).

To support the idea that switches are popular thanks to their limited impact on the dealers' balance sheet, we show in Figure 5 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. An implication of switches leading to only small changes in a dealer's balance sheet is that dealers pay negligible inventory costs, if any, to intermediate these trades. This is true both where negotiations are conducted bilaterally or if a switch is conducted on an electronic RFQ platform.[13]

If balance sheet cost-savings are the main driver of trade bundling, switches should predominantly occur within the same asset class rather than across different asset classes. Figure 6 panel (B), which plots the fraction of switches that consists entirely of assets within the same asset class, confirms that this is the case. Over 50 percent of switches are within an asset class. A typical switch is a client selling one Government of Canada bond and buying a different Government of Canada bond.[14] Not only does a dealer's inventory not change with these types of switches, but there is also no change in the credit risk characteristics of their inventory. Bond duration might change slightly, but in terms of the balance sheet constraints a dealer faces, there is effectively no change.

Dealer balance sheets, however, are certainly not the only consideration for trade bundling; client preferences also play a role. For example, during the COVID pandemic, we observe a preference for safety as clients purchased Government of Canada bonds and sold riskier asset classes. Generally, trade bundling eliminates the risk that a client can only execute part of a multi-asset trading strategy, and might therefore be preferred.

[13]We observe a similar pattern in Appendix Figure D4, which plots the net change in quantity for bi-directional trades. The mass is more evenly distributed than in the case of switches; however, about one-third of trades have a net-zero impact on a dealers' balance sheet. This is in contrast to uni-directional bundles, which by definition increase or decrease a dealers' inventory by the total size of the trade.

[14]Appendix Table D6 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.

Switching bonds with the same dealer, for instance, means that the client does not have to worry about having sold a bond with one dealer and failing to execute the buy trade with a different dealer—a real possibility in some markets and in some time periods (e.g., Hendershott and Madhavan, 2015, Kargar et al., 2024, and Hendershott et al., 2024). Moreover, 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.[15]

Which of these effects dominates—balance sheet considerations, client needs or preferences—varies over time and across asset classes, and type of bundle trade. For example, 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 6 panel (A)), even though they have a sizable impact on the dealer's balance sheet. This is indeed what we find—over 37 percent of clients who bundle assets within the same class only ever trade that asset class. Furthermore, the high share of trades executed across asset class in bi-directional bundles, in some periods, underscores the complexity of the factors driving trade bundling. These types of trades typically involve changes in credit risk, duration risk, etc., and so client preferences can increase the complexity of balance sheet management.[16]

Fact 4. *Bundling trades is more common on electronic platforms than when trading bilaterally.*

Figure 7 plots the share of bilateral trade volume that are bundles and switches versus those shares for trades on a platform. We see that a strikingly large fraction of trade volume on the electronic platform, 40 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. Appendix Figure D6 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 without the need for a premium or discount. This is typically

[15]Asset managers have a legal obligation to invest assets according to guidelines set by their clients.
[16]The large increase in within-asset class bi-directional trades in 2023 comes entirely from an increase in trades within our foreign sovereign bond asset class, specifically U.S. Treasuries.

FIGURE 6. Within-asset class bundles and switches over time

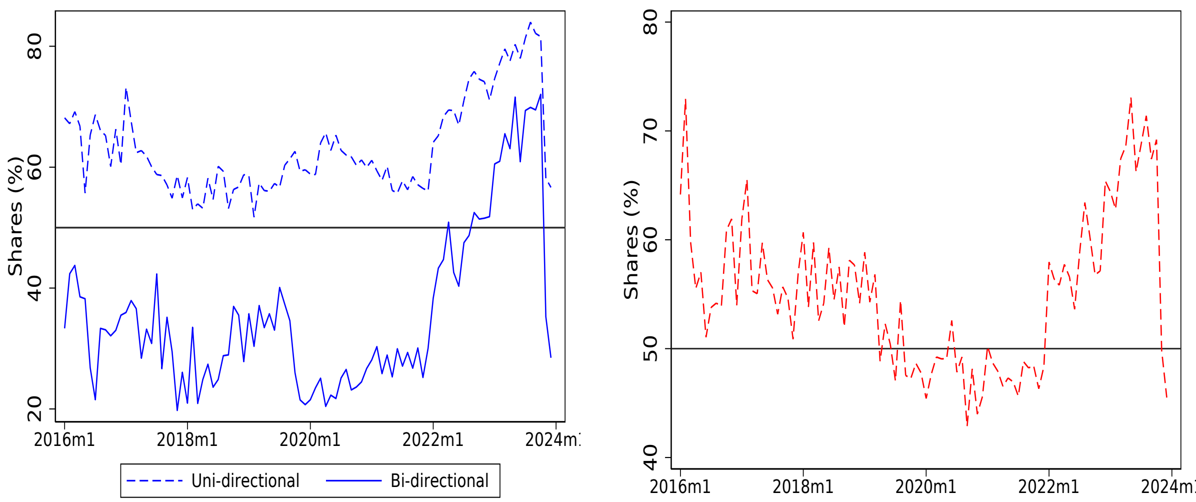

(A) Bundles (B) Switches

Note: Figure 6 plots, in panel (A), 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. Panel (B) plots the share of switches that are within an asset class, as opposed to across asset classes.

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 the net settlement estimate (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.

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.

FIGURE 7. Where do clients trade bundles?

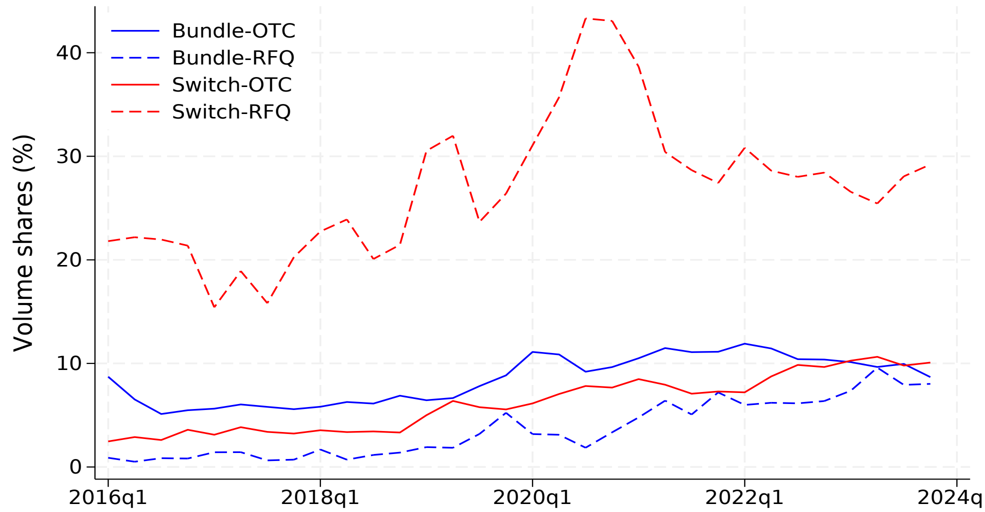

Note: Figure 7 plots the share, based on volume, of all bilateral trades that are bundles and switches (solid lines) and the share, based on volume, of all electronic platform trades that are bundles and switches (dashed line).

Fact 5. *Switches are cheaper than single-asset trades, while the transaction cost for bundles depends on dealer-client relationships.*

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

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

We control for log-trade quantity (Pinter et al., 2024) and different fixed effects. 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). Across specifications, 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).

Table 5 reports estimation results; Appendix Table D7 reports the results if we pool both types of bundle trades into the same regression. 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.[17]

We find that switches, which involve minimal balance sheet costs, are cheaper than single-asset trades. In contrast, the average bundle is not consistently cheaper than a single-asset trade. The differences arise from the direction of the bundle and the dealer-client relationship. For example, from column (4), we see that sell-side bundles, where clients sell multiple assets, are more costly than single-asset trades. Bi-directional bundles are also more expensive, likely because, on average, dealers add more bonds to their balance sheets than they remove in these types of trades, as illustrated in Appendix Figure D5, though there is some heterogeneity in this trend.

Importantly, dealers appear to pass on balance sheet cost savings, particularly to clients with whom they have a relationship. We see this from the negative interaction terms for trading a bundle with the most-favorite dealer in column (2). Even though clients pay a markup to trade with their most-favored dealer, they receive a discount when bundling trades.[18] Furthermore, bundles and switches are cheaper than single-asset trades within a dealer-client relationship (as shown in column (3); furthermore, column (5) shows that the savings from bundling come from the client buy-side). Consequently, bundling trades may serve as a means to foster and strengthen relationships both over time and across asset classes.

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 6. *The transaction cost of bundle trades vary depending on whether the trade is conducted on an electronic platform or bilaterally.*

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

[17]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.

[18]Similarly, Issa and Jarnecic, 2019 and Allen and Wittwer, 2023 document relationship markups between dealers and clients. However, evidence on relationship-pricing is mixed; for instance, Jurkatis et al., 2022 find relationship discounts between dealers and clients in the European corporate bond market.

Table 5. Transaction costs and multi-asset trading strategies

	(1)	(2)	(3)	(4)	(5)
		Panel (A) - 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)xI(bundle)		-0.483***			
		(0.176)			
Bidirectional				0.176**	-0.076
				(0.070)	(0.068)
Buy				0.028	-0.218*
				(0.133)	(0.114)
Sell				0.265**	-0.140
				(0.131)	(0.120)
log(quantity)	-0.138***	-0.134***	0.043***	-0.138***	0.043***
	(0.026)	(0.025)	(0.014)	(0.026)	(0.014)
Constant	2.678***	2.480***	0.268	2.680***	0.278
	(0.352)	(0.339)	(0.197)	(0.345)	(0.197)
R2	0.186	0.186	0.205	0.186	0.205
Obs.	5,035,640	5,035,640	5,034,277	5,035,640	5,034,277
		Panel (B) - 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)xI(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		
Day-Bond FE	Y	Y	Y	Y	Y
Dealer-Client FE	N	N	Y	N	Y

Note: Table 5 reports regression results from estimating transaction cost equations for bundles and switches. The dependent variable is trading costs with the benchmark price being the average (signed) trading price per day-ISIN. Panel (A) compares bundles to single-asset trades. Panel (B) compares switches 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, I(most-favored) indicates a trade with a client's most-favorite dealer according to our relationship measure (1), I(Bidirectional) indicates a bi-directional bundle, I(Buy) and I(sell) mark all buy, or all sell bundles, respectively. 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 6. Correlation between transaction cost, bundling and switching

	(1)	(2)	(3)	(4)	(5)
I(platform)	-0.988***	-0.810***	-0.587***	-0.617***	-0.523***
	(0.110)	(0.088)	(0.115)	(0.075)	(0.132)
I(bundle)		0.407***	-0.015		
		(0.121)	(0.080)		
I(platform)× I(bundle)		-1.196***	-0.496***		
		(0.184)	(0.116)		
I(switch)				-0.562***	-0.517***
				(0.074)	(0.109)
I(platform) × I(switch)				0.716***	0.487***
				(0.120)	(0.155)
log(quantity)	-0.193***	-0.194***	0.041**	-0.031**	0.099***
	(0.034)	(0.033)	(0.017)	(0.015)	(0.016)
Constant	3.921***	3.844***	0.659***	1.343***	-0.560**
	(0.494)	(0.446)	(0.230)	(0.220)	(0.230)
R2	0.203	0.217	0.242	0.190	0.201
Obs.	4,087,034	3,799,402	3,797,992	2,900,303	2,898,838
Day-bond FE	Y	Y	Y	Y	Y
Dealer-client FE	N	N	Y	N	Y

Note: Table 6 presents regression results where the dependent variable is trading costs with the benchmark price being the average (signed) trading price per day-ISIN. We control for trade size using log(quantity). 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. Standard errors are clustered at the date and bond level.

asset classes.[19] 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

[19]The only asset class not traded on multilateral electronic platforms are foreign sovereign bonds—recall Table 2. We therefore exclude them from our analysis of trading costs on and off the platform.

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 6, which is similar to Table 5, shows the estimation results separately for bundles (in columns (2) and (3)), and for switches (in columns (4) and (5)). Column (1) presents the average platform discount across all trades—about 25 percent. Since we have already addressed the overall pricing differences between bundles, switches, and 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 bundles are cheaper on than off the platform.[20] Even though clients might potentially be willing to pay a convenience-premium for platform services that facilitate bundling, the competition effect dominates.[21] This is not true for switches: within day and bond switches are more expensive, and if we zoom in on within dealer-client, switches cost the same on versus off the platform. This is notable, given that switches are generally attractive to dealers due to their low balance sheet costs. Despite this, higher competition on the platform does not reduce transaction costs for switches relative to bilateral trading.

Combining the facts that switches are not cheaper on the platform than off with the evidence in Figure 7 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

[20] 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).

[21] A potential opposing factor is front-running (Baldauf and Mollner, 2024). Dealers might be concerned that the losing bidders in the RFQ front-run them. For example, a dealer might have bought three bonds without having them in inventory. The losers in the auction could then drive up the winning dealers costs by selling and subsequently buying these bonds on the market. If front-running was the primary concern then we should see higher prices for bundles on the platform than in bilateral trades.

between bonds traded on and off the platform, as we control for bond-day fixed effects. Thus, it is distinct from the classical liquidity premium.

The size of the convenience premium is difficult to determine, as transaction costs are influenced by multiple factors and represent equilibrium outcomes. Our estimate, derived from comparing the cost of a switch on and off the platform in column (4), likely represents a lower bound for the convenience premium, which is about 0.10 basis points (or 18 percent of the total cost-savings from conducting a switch in bilateral negotiations). This is because our comparison of switches across trading venues (of the same security on the same day) controls for factors that influence pricing but are unrelated to the execution method, such as dealer balance sheet costs. One remaining key factor—competition—is expected to increase when moving from bilateral trading to the platform, which would reduce the observed price difference and, consequently, our estimate of the premium.

We interpret the contrasting findings between bundles and switches as reflecting dealers' willingness to compete on the platform, which hinges on their concerns about adverse selection. 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.[22] Therefore, dealers are more cautious, and competition is weaker. Additionally, switches—particularly swaps—require dealers to precisely or near-perfectly offset the investor's buy or sell order (recall Figure 5). This complexity 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

[22]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.

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 features like streamlined execution, liquidity pooling, and efficient transaction processing.

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.

Baldauf, M. and Mollner, J. (2024). Competition and information leakage. *Journal of Political Economy*, 132(5):1603–1641.

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.

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.

Feldhutter, P. (2012). The same bond at different prices: Identifying search frictions and selling pressures. *Review of Financial Studies*, 25:1155–1206.

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.

Harris, L. and Piwowar, M. (2006). Municipal bond liquidity. *Journal of Finance*, 61(3):1330–1366.

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.

Hong, G. and Warge, A. (2000). An empirical study of bond market transactions. *Financial Analysts Journal*, 56:32–46.

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, *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. 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. and Üslü, S. (2022). 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. (2021). Exchange design and efficiency. *Econometrica*, 89:2887–2928.

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

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.

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

Appendix A. Data sample selection and cleaning

The raw sample includes 33,714,669 observations, of which 32,197,123 are secondary market trades. The largest sample reduction is dropping all counterparties that do not have a globally unique legal entity identifier (LEI, `https://www.gleif.org/en`). The majority of these trades are retail investors. Prior to August 2019, dealers were not required to report counterparty LEI, and dealers are still not required to report an LEI for retail clients (since they would not have one). Dropping these observations leaves us with 13,682,819 observations. Next, we drop trades that were on a Canadian holiday, canceled or inputted more than once, or having missing quantities—dropping these observations leaves us with 15,493,314 trades. Next, we drop trades involving reporting dealers who do not have to report trades to the Bank of Canada. We are left with 13,393,580 trades. We then drop Government of Canada bonds that are not denominated in Canadian dollars. This is 29,968 transactions. Our final data set has 12,928,044 trades over 2,068 days, which is approximately 6,251 secondary market trades per day.

Appendix B. Asset class definitions

Table B1 provides definitions for our asset class classification.

Appendix Table B1. Types of Fixed Income Securities

Issuer Type	**Description**
Government of Canada	Interest-bearing bonds redeemed at maturity for par value. Maturity ranges from 90-days to 50 years. Includes Government of Canada bonds (GoC) and Treasury bills (T-bills). Bank of Canada acts as auctioneer.
Provincial governments	Includes Provincial bills and bonds ("provis"). Median tenor is 10 years. Mostly issued through syndication.
Municipalities	Debt securities ("munies") issued by cities, counties, and other governmental entities to fund day-to-day obligations and finance capital projects. Included with provincial government debt for the purpose of our analysis.
Corporate (non-financial)	Short-term debt and bonds from non-financial firms. Short-term median tenor < 1 year; bonds 1-99 years with median 10 years. Includes Canadian and foreign firms.
Corporate (financial)	Financial institution debt including bills, bonds, commercial papers, and bearer bonds/notes. Median tenor < 7 years.
Bankers' Acceptances	Bank commitment for future payment, substituting bank's creditworthiness for borrower's. Tradeable in secondary market. Median tenor < 1 year.
Foreign sovereign debt	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	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	Miscellaneous financing category; securities issued by Government-Sponsored Enterprises.

Appendix C. 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. Table C1 provides definitions.

Appendix Table C1. Client Types

Type	**Description**
Central bank	An institution that manages the currency and monetary policy of a country.
Bank	A deposit-taking/accepting institution.
Investment bank	A bank that provides financial services for corporate and institutional clients.
Pension fund & insurance	Financial entity whose purpose is to manage investments (and provide services) related to pension, retirement, insurance, reinsurance, benefits, and superannuation funds.
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.
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.
Hedge fund	Greater flexibility than other asset managers have take on leverage, short sell securities and use derivatives. Financial entity that is a hedge fund or a hedge fund manager.
Broker-dealer	Financial entity whose purpose is to offer brokerage services.
Non-financial	Non-financial entity, including retail and real estate, among other groups.
Other	Financial entity that does not fall in any of the aforementioned classifiers (e.g., private equity, financial planner).

Appendix Figure C1. Fraction of trades by client and parent type

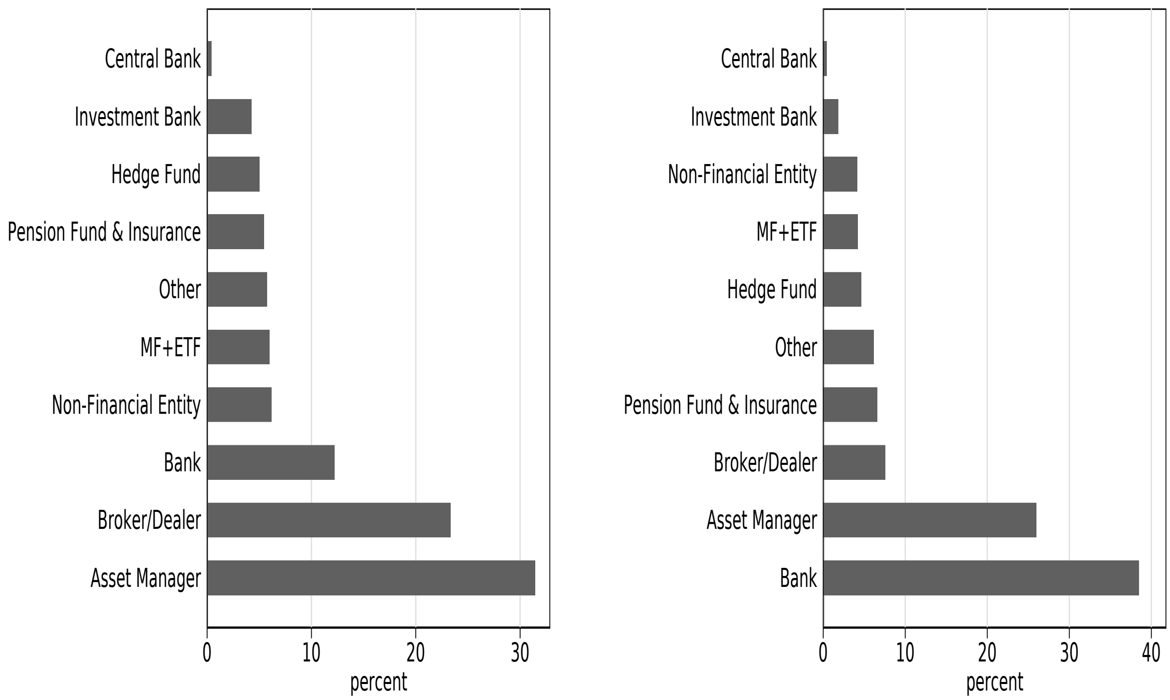

(A) client type (B) Type of the client's parent

Note: Appendix Figure C1 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 D. Additional tables and figures

In this section we provide additional information about the assets' traded in our sample and details on the types of market participants.

Appendix Table D1. Asset characteristics

Variable	GoC	Provi	CDN agency	non-CDN agency	Corp. non-fin.	Corp fin	BA	Foreign sov.
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
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 D1 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. A country of domicile is the country where the issuing entity has a legal address. 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, which remain constant until the predetermined maturity date. 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 Figure D1. Average trade size by asset class

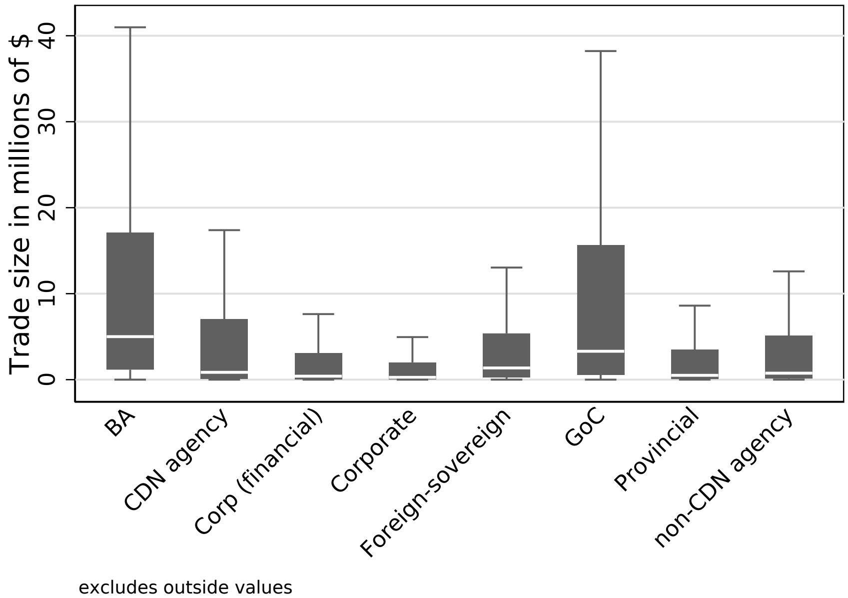

Note: Appendix Figure D1 plots the interquartile range of trade sizes for each asset class in our data set. The average trade size is large—$10.88 million, but this is driven by the Government of Canada (GoC) and Bankers' Acceptances (BA) market. The median trade size is $1.3 million.

Appendix Table D2. Number of asset classes by client type

	Probit		Ordered Probit	
	coeff.	std. error	coeff.	std. error
Categories				
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 Entity	-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		

Note: Appendix Table D2 reports Probit and ordered Probit results for the number of asset classes by client type. In order to correlate multi-asset trading with client characteristics, 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 client i. The sample period is January 4, 2016 to December 31, 2023. There are 135,439 observations. 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 D3. Number of asset classes by parent client type

	Probit		Ordered Probit	
	coeff.	std. error	coeff.	std. error
Categories				
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 Entity	-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		

Note: Appendix Table D3 reports Probit and ordered Probit results for the number of asset classes by parent-client type. In order to correlate multi-asset trading with parent characteristics, 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 parent i. The sample period is January 4, 2016 to December 31, 2023. There are 99,919 observations. 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 D4. Poisson regression for number of dealers

	Client		Parent	
	coefficient	std. error	coefficient	std. error
Asset class				
Agency	0.861***	0.014	0.832***	0.016
Agency-mtg	1.168***	0.025	1.150***	0.029
BA	0.756***	0.017	0.688***	0.019
Corp (financial)	0.895***	0.012	0.844***	0.014
Corporate	0.878***	0.013	0.832***	0.015
Foreign-sovereign	0.711***	0.014	0.659***	0.014
Provincial	1.035**	0.014	1.018	0.016
Client type				
Bank	0.999	0.061	1.097	0.084
Broker/Dealer	0.961	0.065	0.893	0.067
Central Bank	1.121	0.081	1.021	0.079
Hedge Fund	0.766***	0.054	0.842*	0.075
Investment Bank	0.844**	0.073	0.909	0.132
Non-Financial Entity	0.964	0.076	0.956	0.082
Other	0.794***	0.052	0.737***	0.055
Pension Fund & Insurance	1.013	0.081	1.059	0.094
MF+ETF	0.657***	0.044	1.130	0.116
Constant	2.577***	0.114	2.932***	0.156

Note: Appendix Figure D4 shows Poisson regression results where we pool data from January 4, 2016 to December 31, 2023. The first two columns are at the client-level and the last two columns are at the parent-level. Estimates are incidence rates. The omitted asset class are GoCs and the omitted client type is asset manager. MF+ETF is mutual funds and exchange-traded funds. There are 18,344 observations at the client level and 12,625 observations at the parent level. Standard errors are clustered at the client level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix Figure D2. Transaction costs

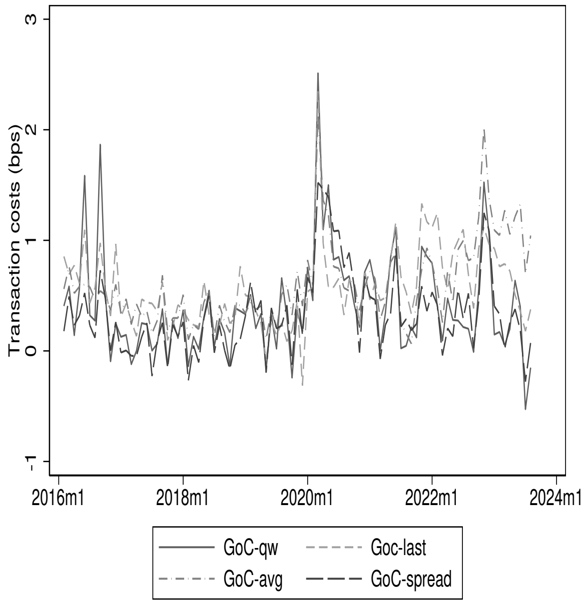

(A) Government of Canada

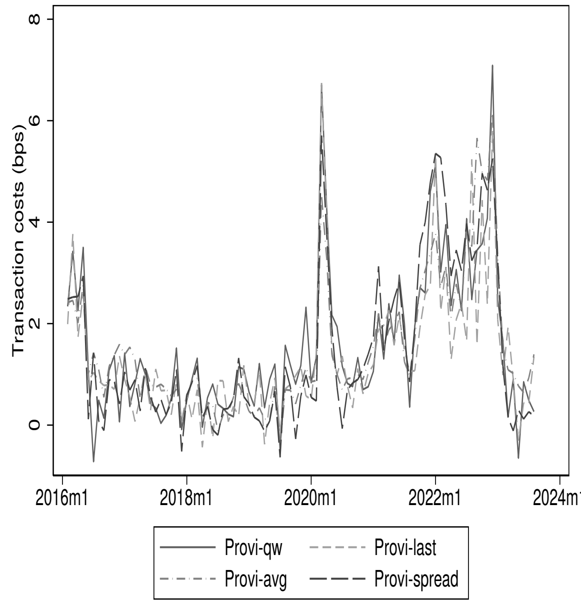

(B) Provincial

Note: Appendix Figure D2 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 average of D2D trades as the benchmark price in equation (2); (ii) '-avg' uses a simple average as the benchmark price (our baseline); (iii) '-last' uses the last observed transaction price as the benchmark; and (iv) 'spread' uses the buy-sell spread (Hong and Warge, 2000, Harris and Piwowar, 2006, and Feldhutter, 2012). Transaction costs using the quantity-weighted D2D daily price as the benchmark is less volatile than other measures, however, the patterns are similar.

Appendix Figure D3. Transaction costs per asset class and client type

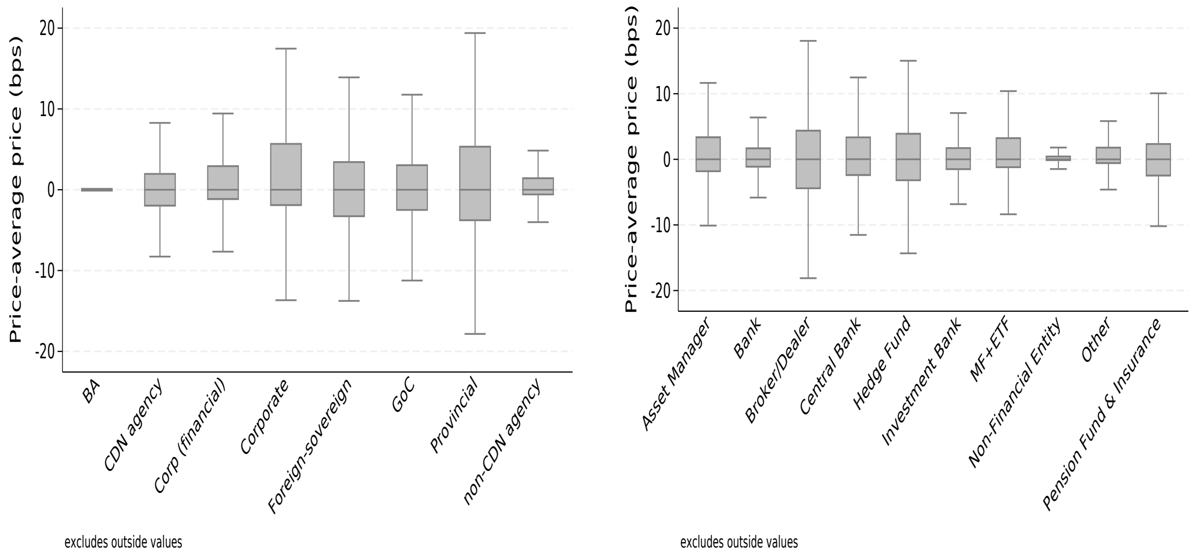

(A) Asset class (B) Client

Note: Figure D3 plots the interquartile range of transaction costs for each asset class (panel A) and client type (panel B).

Appendix Figure D4. Changes in net quantity in bidirectional trades

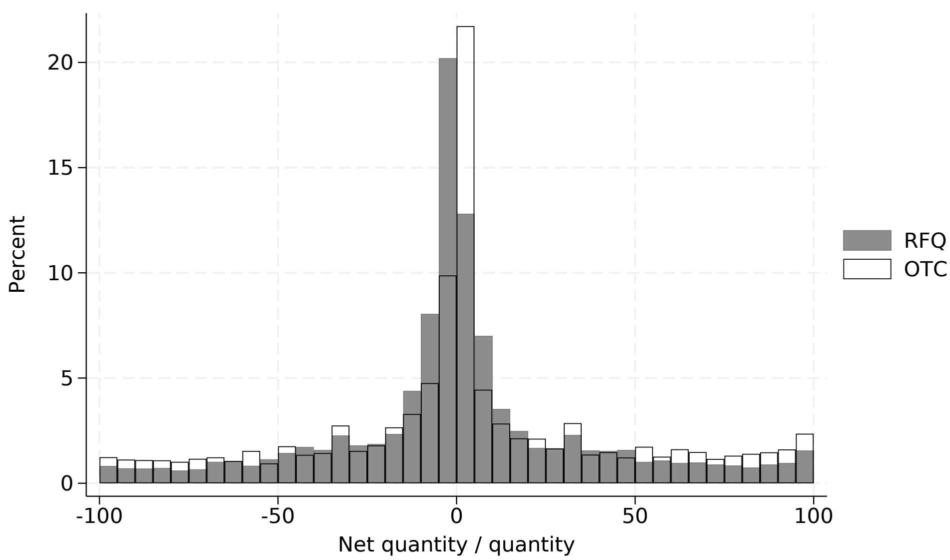

Note: Appendix Figure D4 plots the change in net quantity in bidirectional trades. Each bin is 5%. In white are bilaterally negotiated bidirectional trades (OTC), and in gray are bidirectional trades on Candeal. Net quantity is defined as the buy quantity minus the sell quantity as a fraction of total quantity exchanged.

Appendix Table D5. Who conducts multi-asset trades?

	I(bundle)		I(switch)	
Categories	coefficient	std. error	coefficient	std. error
Bank	-0.744***	0.080	-0.139	0.158
Broker/Dealer	-0.180	0.116	0.075	0.094
Central Bank	-1.375***	0.118	0.062	0.187
Hedge Fund	-0.445***	0.079	0.606***	0.082
Investment Bank	-0.843***	0.175	-0.389**	0.142
Non-Financial Entity	-1.027***	0.109	-0.393***	0.098
Other	-0.549***	0.124	-0.125	0.095
Pension Fund & Insurance	-0.616***	0.116	0.322**	0.115
MF+ETF	-0.155*	0.074	-0.081	0.131
Constant	-1.704***	0.083	-2.018***	0.057
N	5,013,896		4,769,086	

Note: Appendix Table D5 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 D6. Buy and sell in a switch

sale/purchase	GoC	Provi	Corp. non-fin.	Corp. fin.	For. sov.	CDN agency	BA	non-CDN agency
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 (non-fin.)	70.4	2.2	17.2	6.0	2.2	0.5	0.13	0.4
Corp.(fin)	68.6	2.2	5.0	16.3	6.5	0.6	0.5	0.4
For. sov.	0.6	1.1	1.1	2.9	90.8	0.0	0.01	3.5
CDN-agency	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 agency	19.3	9.3	1.5	1.9	52.2	0.4	0.1	15.3

Note: Appendix Table D6 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 Figure D5. Net sales in bidirectional trades

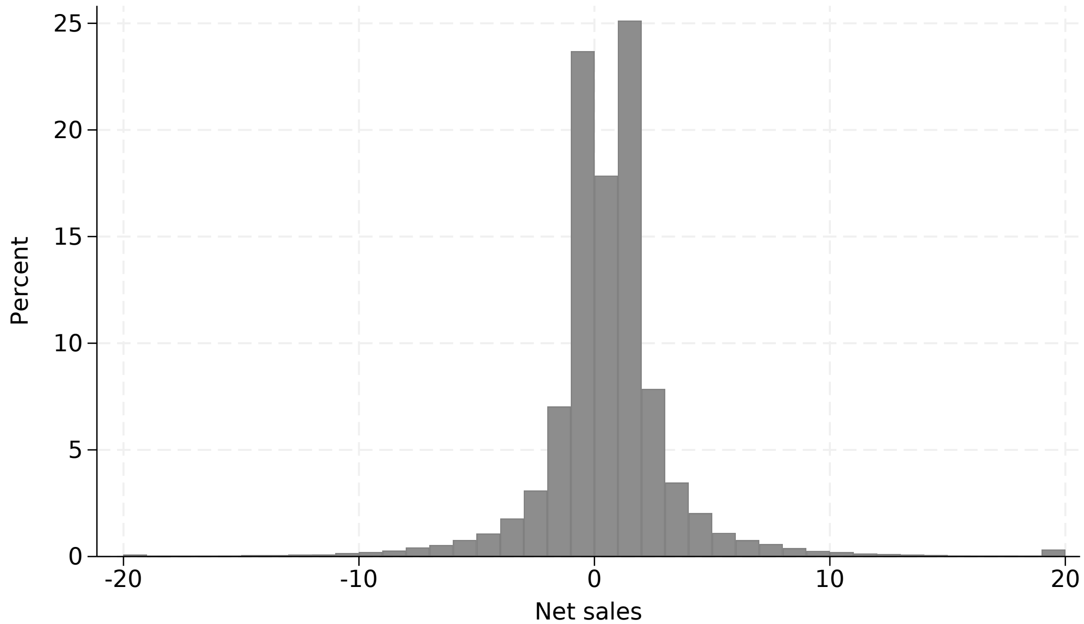

Note: Appendix Figure D5 plots a histogram of the net sales across all bidirectional bundles. The graph is winsorized at -20 and +20. Each bin corresponds to 1 bond.

Appendix Figure D6. Switch trade on CanDeal

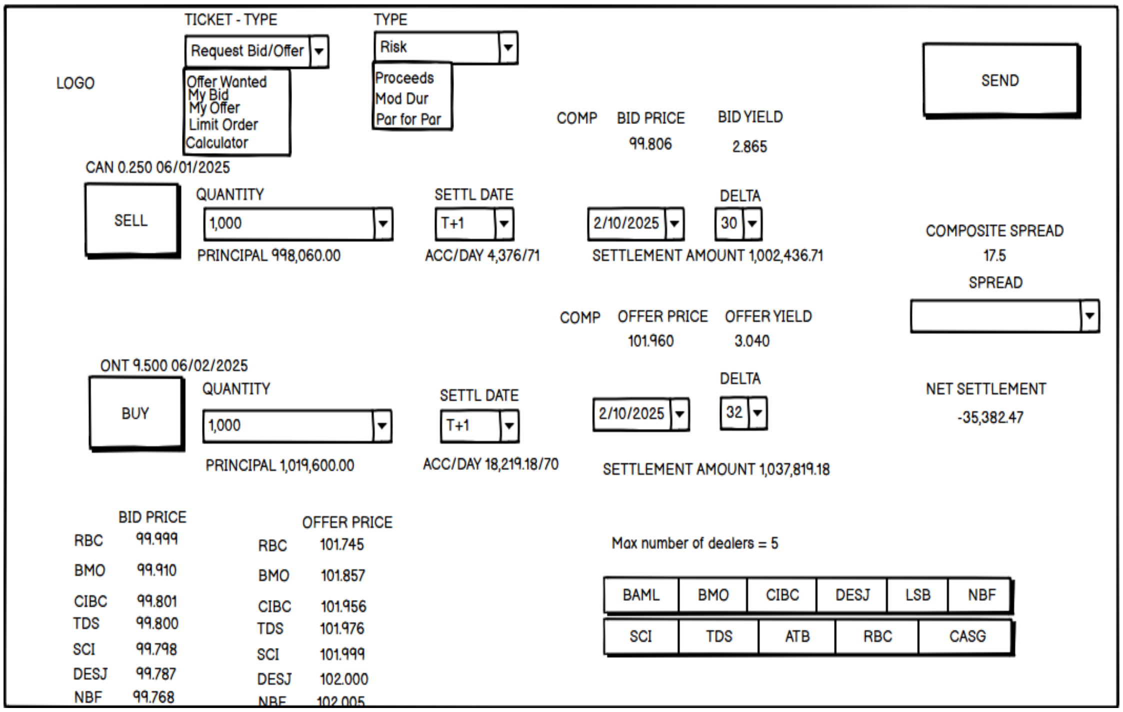

Note: Appendix Figure D6 is a depiction of what clients see on their CanDeal platform for a hypothetical switch trade. Although a screenshot was not possible, this graphic captures everything a client sees and was generously provided to us by CanDeal. In the SELL section is 1,000 units of a Government 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. The maximum number of dealers in the government bond market is 5; it is 11 for coporate bonds.

Appendix Table D7. Transaction cost and multi-asset trading strategies

	(1)	(2)	(3)
I(multi-asset)	0.175**	0.435***	-0.126**
	(0.078)	(0.162)	(0.060)
I(most-favored)		0.226***	
		(0.036)	
I(most-favored)xI(multi-asset)		-0.473***	
		(0.178)	
log(quantity)	-0.135***	-0.131***	0.041***
	(0.026)	(0.024)	(0.013)
Constant	2.624***	2.447***	0.267
	(0.345)	(0.332)	(0.188)
R2	0.174	0.174	0.192
Obs.	5,410,480	5,410,480	5,409,135
Day-Bond FE	Y	Y	Y
Dealer-client FE	N	N	Y

Note: Appendix Table D7 presents regression results for equation (2), where the dependent variable is trading costs and the benchmark price being is the average (signed) trading price per day-ISIN. We control for trade size using $log(quantity)$. I(multi-asset) is an indicator variable equal to 1 if the trade involves more than one asset and 0 otherwise. I(most-favored) is an indicator variable equal to 1 if the client trades with their highest-volume dealer and 0 otherwise. Standard errors are double clustered at the date and bond level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Appendix E. Robustness

We first present robustness results for our main results using the dealer-to-dealer benchmark price as in Hendershott and Madhavan, 2015. Appendix Table E1 replicates Table 5. The results are similar, even though the number of transactions drops substantially when using the interdealer price. The reason is that some asset classes, for example, BAs, are infrequently traded on the interdealer market, and therefore we do not have a benchmark price. See Table 2.

Second, 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 in the Canadian fixed income markets. They account for approximately 40 percent of all trade. Appendix Table E2 reports results for this sample of most active investors. The premium for bundles and the discount for switches are both somewhat larger in this sample of investors than in the full sample.

Appendix Table E1. Transaction costs and multi-asset trading strategies (benchmark = quantity-weighted average)

	(1)	(2)	(3)	(4)	(5)
			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)xI(bundle)		-0.614**			
		(0.250)			
Bidirectional				0.162	0.013
				(0.124)	(0.129)
Buy				-0.543	-0.815**
				(0.342)	(0.319)
Sell				1.282***	0.739**
				(0.310)	(0.309)
log(quantity)	-0.143***	-0.137***	0.097***	-0.140***	0.097***
	(0.037)	(0.035)	(0.024)	(0.036)	(0.024)
Constant	2.693***	2.499***	-0.635*	2.650***	-0.635*
	(0.511)	(0.491)	(0.342)	(0.489)	(0.345)
R2	0.151	0.151	0.168	0.151	0.168
Obs.	2,126,988	2,126,988	2,125,798	2,126,988	2,125,798
			Panel (C) - 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)xI(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		
Day-Bond FE	Y	Y	Y	Y	Y
Dealer-client FE	N	N	Y	N	Y

Note: Appendix Table E1 reports regression results from estimating transaction cost equations for bundles and switches. The dependent variable is trading costs with the benchmark price being the quantity-weighted average (signed) trading price per day-ISIN. Panel (A) compares bundles to single-asset trades. Panel (B) compares switches 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 E2. Transaction costs and multi-asset trading strategies for a sample of the 95 most active investors

	(1)	(2)	(3)	(4)	(5)
			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)xI(bundle)		-0.903***			
		(0.229)			
Bidirectional				0.033	-0.214**
				(0.107)	(0.095)
Buy				0.278	-0.055
				(0.181)	(0.147)
Sell				0.270*	-0.325**
				(0.159)	(0.142)
log(quantity)	-0.106**	-0.097**	0.083***	-0.103**	0.084***
	(0.048)	(0.046)	(0.019)	(0.047)	(0.019)
Constant	2.230***	2.088***	-0.178	2.191***	-0.185
	(0.635)	(0.607)	(0.259)	(0.620)	(0.262)
R2	0.247	0.247	0.269	0.247	0.269
Obs.	2,106,686	2,106,686	2,106,667	2,106,686	2,106,667
			Panel (C) - switches		
I(switch)	-0.488***	-0.437***	-0.507***		
	(0.122)	(0.116)	(0.140)		
I(most-favored)		0.025			
		(0.057)			
I(most-favored)xI(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		
Day-Bond FE	Y	Y	Y	Y	Y
Dealer-Client FE	N	N	Y	N	Y

Note: Appendix Table E2 reports regression results from estimating transaction cost equations 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 bundles to single-asset trades. Panel (B) compares switches 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.