Specialization in Financial Markets*

Milena Wittwer⋆, and Andreas Uthemann+

May 20, 2025

Abstract

Intermediary asset pricing models typically assume integrated markets with homogeneous assets, akin to Walrasian auctions. In practice, however, markets are fragmented and financial products are diverse. Using a unique dataset linking intermediaries' trades across Canadian stock, bond, and derivative markets, we examine where, what, and at what prices intermediaries trade to assess whether market fragmentation, product diversity, and the resulting specialization are empirically relevant features that asset pricing models should incorporate. We document substantial specialization: intermediaries concentrate their trading activity unevenly across markets and products, with product specialization more pronounced than market specialization. Furthermore, more specialized intermediaries consistently obtain better prices. These findings highlight the need for asset pricing models that account for specialization—especially across products.

Keywords: Market segmentation, specialization, financial intermediaries, market design, asset prices

JEL: G00, G10, G12, G19, D40

*The views presented are those of the authors and not necessarily the Bank of Canada. We thank Mirela Sandulescu for an excellent discussion. We also thank Jason Allen, Vincent Vincent Bogousslavsky, Adi Sunderam, Jonathan Wallen, and the experts at the Bank of Canada and the Montreal Stock Exchange—in particular, Corey Garriott and the Triton project team. We also thank Connor Breck, Carel Chok, Costanza Didonna, Matthew Hagerty, Harrison Lynch, Brendan McLaughlin, Daniel Smith, and Joseph Wagner for excellent research assistance. Any errors are our own. Correspondence to: §Milena Wittwer (Boston College): wittwer@bc.edu, and +Andreas Uthemann (Bank of Canada): authemann@bankofcanada.ca

1 Introduction

Intermediary asset pricing suggests that frictions faced by financial intermediaries can constrain arbitrage and influence asset prices (e.g., Shleifer and Vishny (1997); Gromb and Vayanos (2002); Brunnermeier and Pedersen (2009); He and Krishnamurthy (2013); Brunnermeier and Sannikov (2014)). Most models assume frictionless markets, akin to a Walrasian auction. Consequently, the related empirical literature largely overlooks frictions arising from market fragmentation (e.g., Pasquariello (2014); Adrian et al. (2014); Du et al. (2018); He et al. (2017); Siriwardane et al. (2022)). This contrasts with the fragmented nature of financial markets, where different asset classes trade in distinct venues (Malamud and Rostek (2017); Weill (2020); Chen and Duffie (2021); Budish et al. (2024)). Moreover, since financial assets are treated as homogenous, the literature tends to overlook the diversity of financial products and, consequently, the role of product specialization (Babus and Hachem (2023); Babus et al. (2024); Mota and Siani (2024)).

We introduce a unique dataset to study the role of market and product 'segmentation'—or equivalently, 'specialization'—in financial markets to provide novel stylized facts that inform future asset pricing models. By linking trades across Canadian stock, bond, and derivative markets, we analyze cross-market and cross-product specialization by examining where, what and at what prices brokers and dealers trade, offering insights into the returns to specialization.[1] Cross-market specialization may arise from differences in market clearing rules or entry costs, while cross-product specialization within a market—where such frictions are absent—may instead reflect differences in trading expertise, relationships, or client preferences.

Our dataset covers all trades executed on Canada's fixed-income market and all exchanges owned by the Toronto Stock Exchange Group (TMX) from 2019 to 2022. TMX owns three stock exchanges, which account for roughly 60 percent of equity trade volume in Canada, and the country's only derivatives exchange. A key feature of this dataset is the ability to track dealers over time and across markets using legal entity identifiers (LEIs)—an attribute rarely available in trade-level datasets, particularly for stocks and derivatives. We link this information to public data in order to classify securities into products, for instance, corporate bonds, large-cap stocks, Exchange Traded Funds (ETFs), and Treasury futures. Additionally, we manually assign dealers to their parent institutions, i.e., the LEIs of their holding companies, and categorize them by

[1] An alternative label for the institutions analyzed in this paper is 'market-maker.' We avoid this term because the small set of financial institutions that dominate trading—the set we study in this paper—engage in activities beyond market-making, including executing trades for clients and responding to client needs. For instance, on exchanges, only a small set of firms are formally designated as market makers with an obligation to provide liquidity.

type, such as primary dealers or hedge funds.

Using these data, we establish three stylized facts about specialization. The first demonstrates its existence and quantifies its extent: dealers allocate their trading activity unevenly across markets and products—they specialize.

To quantify market specialization, we construct dealer-specific market specialization scores, which range from zero (no trades in a market) to one (exclusive trading within it). Banks tend to concentrate in bonds, high-frequency traders in derivatives, and primary dealers in government debt are the most active across markets. A similar pattern emerges within markets: trading is unevenly distributed across product segments, reflecting product specialization. We capture this using a dealer-specific product specialization score, defined as the ratio of a dealer's trade share in a product segment (relative to all dealers) to the sum of these trade shares across products in a market. Like the market score, it ranges from zero to one.

Product specialization appears shaped by both market structure—centralized versus decentralized—and product complexity. On centralized stock exchanges with standardized products, dealers typically trade broadly. In contrast, specialization is stronger in the decentralized fixed-income over-the-counter (OTC) market, where search and relationship frictions may push dealers to focus on specific bond types. Yet these frictions cannot fully explain specialization: it also arises in the centralized derivatives market, which, like stock exchanges, operates via an anonymous limit order book, but features less standarized products.

Across markets, product specialization is more pronounce than market specialization—this is our third fact. To derive it we introduce a specialization index that integrates market and product specialization scores, and allows us to decompose within-market from cross-market specialization for each dealer. For this, we adapt the Theil (1967) index to account for the fact that not all dealers participate in every market segment. While commonly used to measure inequality in socio-economic contexts (e.g., Anand and Segal (2015)), the Theil index has not, to our knowledge, been applied to trade settings.

The decomposition shows that, for most dealers, product specialization within a market is greater than market specialization. This finding suggests that, for large financial institutions, barriers to market entry are less restrictive than factors that limit trading across products within the same market. As a result, policies aimed at moderately changing entry costs or membership fees—such as the recently revised fee schedules for registered broker-dealers in the U.S. and Canada (FINRA (2024); CIRO (2024))—may have limited impact on market participation.

Next, we examine whether market and product specialization affect transaction prices, aiming to establish our third, and final, stylized fact that specialized dealers trade at better prices. We focus on relative prices across dealers, not on how specialization affects aggregate price levels. Dealer specialization could influence transaction prices by improving inventory management, or shaping beliefs about fundamentals. However, in a frictionless and competitive market, any price effects from specialization would be arbitraged away.

We, therefore, begin by showing that none of the markets is sufficiently frictionless to prevent some dealers to outperform others. We measure a trade's margin as its price advantage relative to the average price at which the same security trades on that day.[2] A margin of 1% indicates that the dealer pays 1% less than the daily average when buying (and sells at 1% more when selling). We show that dealers systematically differ in the prices they obtain, across all markets, even after controlling for trade size, security-time fixed effects, and other observables. High-frequency traders tend to outperform others on exchanges, while retail-facing brokers underperform in the bond and derivatives markets.

Having established that there is scope for price effects, we investigate whether the successful dealers trade across products or markets or whether they specialize. We show that some dealers who trade exclusively within a single market outperform those who trade across markets in the bond and derivatives market, but not the stock market. This is in line with the idea that a decentralized market structure and product complexity promote specialization. Across markets, dealers who consistently secure better prices for bonds do not achieve better prices in stocks or derivatives, and vice versa, suggesting limited trading synergies across markets and products, and reinforcing the role of specialization.[3]

Finally, we exploit cross-sectional variation in dealer specialization to show that more specialized dealers obtain better prices. To address concerns about reverse causality—where more successful dealers become more specialized—and omitted variables, such as dealer sophistication or efficiency, we implement two strategies. First, we relate lagged specialization scores to current trade margins. The idea is that last year's specialization is less likely to be influenced by, or directly affect, current prices. Second, we use an instrumental variables (IV) approach. While identifying exogenous variation in trading is notoriously difficult, our data offer a unique opportunity: we can distinguish whether a stock market trade is for the dealer's own account or for a client. Client orders serve as plausibly exogenous shocks for dealer own-account trades,

[2]Our approach follows the market microstructure literature, which commonly defines transaction costs as the trade price relative to a benchmark (e.g., Hendershott and Madhavan (2015); Hau et al. (2021); O'Hara and Zhou (2021); Pinter et al. (2024)). Ideally, one would compare the trade price to a mid-price or fundamental value, but data constraints prevent this. Instead, we use the average daily price, which we can construct consistently across markets. Our regressions control for security-week fixed effects to account for differences across securities and over time.

[3]This may reflect the tendency of individual traders or desks within institutions to focus on a narrow set of assets, optimizing trading within their domain (Lu and Wallen (2024)).

once observable trade characteristics and a rich set of fixed effects are controlled for.

While neither approach fully eliminates endogeneity concerns, they reveal a consistent pattern when taken together: more specialized dealers trade at better prices. The price effect is moderate at the trade level but becomes economically meaningful when aggregated over time. For example, on the stock exchange, moving from no market (product) specialization to full specialization increases margins by 4–39 (28–41) basis points per trade.

Taken together, our findings underscore the importance of market and product specialization in intermediary asset pricing—two dimensions largely overlooked in existing models. They call for asset pricing frameworks that incorporate specialization, particularly across asset classes, and point to several promising directions for future theoretical and empirical work.

One direction for future research is to analyze whether, and if so how, specialization shapes market-outcomes, including aggregate price levels. Doing so requires either large exogenous variation in specialization across several dealers or a structural framework—for example, by extending Vayanos and Vila (2021) to accommodate richer market structures and product complexity. Such analysis could lay the groundwork for a broader research agenda examining whether—and through what mechanisms—granular frictions observed in micro-level data aggregate into distortions in market equilibrium outcomes. It would also complement the empirical literature asset pricing literature (e.g., Pasquariello (2014); Adrian et al. (2014); Du et al. (2018); He et al. (2017); Siriwardane et al. (2022)), which typically relies on market-level data to capture broad effects but offers limited visibility into underlying mechanisms—aside from a few exceptions (e.g., Siriwardane (2019); Wittwer and Allen (2023)).

In line with the goal of unpacking mechanisms, another direction for future research is to examine why and how product complexity and market structure shape product specialization—as suggested by the contrast between product specialization on the stock market versus the OTC market and the derivatives exchange. This would require models that account for different market structures and product characteristics, shifting the focus beyond a single market or product, as seen in much of the existing literature. With few exceptions—such as Dougast et al. (2022) and studies on derivatives and their underlying assets dating back to Kumar and Seppi (1992)—existing models predominantly focus on a single market structure and asset class.

Another aspect for future research would be to explore how trading interconnectedness evolves during periods of distress. Our data shows that large banks dominate trading across markets, which raises concerns about financial stability. From a policy perspective, this implies that regulatory changes affecting dealer bank balance sheets—such as adjustments to the supplementary leverage ratio to accommodate increased government debt issuance—will impact all markets. Investigating which types of institutions amplify negative spillovers and which help mitigate them would contribute to the extensive literature on contagion, following Allen and Gale (2000).

Finally, our cross-market and multi-asset perspective highlights the need for empirical market microstructure studies (contributing to a large literature, including Hasbrouck and Sofianos (1993); O'Hara (2015); Menkveld (2016); Bessembinder et al. (2020)), and the growing literature on demand estimation (following Koijen and Yogo (2019)) to move beyond isolated markets or individual products within a market. Most existing studies in both literatures focus on a narrow set of assets, such as a single bond type or common equity, and largely constrain substitution across asset classes, limiting the potential for spillover effects. Allen et al. (2020), Chaudhary et al. (2022), Üslü and Pintér (2023), Allen and Wittwer (2024), and Dix and Wittwer (2025) take initial steps in this direction, but given the empirical patterns documented in this study, much remains to be explored.

Similarly, though more distantly related, the extensive asset pricing literature on factor structures, including common stochastic discount factors, has traditionally focused on asset-class-specific factors, but has begun shifting toward identifying joint factors that span multiple asset classes (e.g., Sandulescu (2020); Chen et al. (2024)). Sandulescu (2020), for example, documents significant integration between U.S. corporate bonds and equities, consistent with the empirical patterns we observe for Canada.

2 Institutional environment

Before detailing the construction of the dataset we use to examine dealer specialization, it is useful to review the key market features. The structure of Canadian financial markets closely mirrors that of other developed nations, including the United States. The three primary asset classes—bonds, stocks, and derivatives—each operate in separate markets.

Fixed-income market. Fixed-income instruments are issued in primary markets and traded in decentralized over-the-counter (OTC) markets. In traditional OTC markets, buyers must contact sellers individually to conduct bilateral trades. Consequently, these markets largely depend on large financial institutions—dealers—to intermediate between investors, such as firms, public entities, and individuals. Although not all trade occurs bilaterally today, the market remains fragmented.

Firms seeking to become fixed-income dealers must apply to the Canadian Investment Regulatory Organization (CIRO). CIRO membership is available to Canadian entities registered to

Table 1: Products in the fixed-income market

Product	Trade share
Government Bonds and Bills	63.44
Provincial, Municipal Bonds and Bills	9.33
Bankers' Acceptances	8.81
Bank, Agency Papers	7.96
Corporate Bonds	6.14
ABS, MBS, CMB	4.88
Strips	0.27

Notes: Table 1 shows the daily average share of total trade-volume, computed as the total amount of bonds (in terms of par value) traded on a day, in the bond market per product. ABS are Asset-Backed Securities, MBS are Mortgage-Backed Securities, and CMB are Canada Mortgage Bonds. Appendix Table A1 describes each product category.

operate as dealers or advisors in any province or jurisdiction in Canada. CIRO members must satisfy CIRO's financial and operations compliance, business conduct compliance and registration requirements, including minimal capital requirements (typically C$250,000), and pay annual membership fees (CIRO website).

Fixed-income securities range from long-term bonds to short-term money-market instruments. We classify all securities into product categories, as explained in Appendix Table A1. Government bills and bonds are traded the most, as shown in Table 1. Then we have provincial and municipal debt, and Bankers' Acceptance (which is a money market instrument that is issued by a business and guaranteed by a bank), bank or agency papers (money market instruments issued by banks or agencies), and corporate debt. Mortgage- or asset-backed securities (ABS, MBS, CMB), and strip bonds (which are debt instruments in which both the principal and regular coupon payments, that have been removed, are sold separately) are relatively small.

Equity market. Equity products are in most countries traded on centralized exchanges. Exchanges differ from OTC markets in that the market clears centrally on a limit order book.

In Canada there are nine exchanges.[4] Our focus lies on exchanges that are owned by the TMX group, given our data. The TMX group owns three exchanges: Toronto Stock Exchange (TSX), TSX Venture Exchange (TSXV), and TSX Alpha Exchange (Alpha). In our sample period, 2019-2022, about 58% of the total volume traded, and 63% of the total dollar value traded in an average month on any of the Canadian equity markets was traded on a TMX exchange.[5]

[4]The nine exchanges are: NEO Exchange Inc., Canadian Securities Exchange (CSE), Instinet Canada Cross Limited (ICX), Liquidnet Canada Inc. (Liquidnet), Nasdaq CXC Limited (Nasdaq Canada), Tradelogiq Markets Inc. (TMI), TMX Group (TSX, TSXV, TSX Alpha), TriAct Canada Marketplace (Match Now).

[5]These numbers are computed with data from CIRO, accessible here: https://

Table 2: Stock market products

Product	**Trade share**
Small Stock	53.58
Large Stock	32.33
Uncommon Shares	7.41
Exchange Traded Funds	5.98
Other or Missing	0.68

Notes: Table 2 shows the daily average share of total trade-volume on the stock market, computed as the total amount of stocks (i.e., the total number of shares) traded on a day, per product. Appendix Table A2 describes each product category.

Only exchange members can place orders for their own account, or on behalf of non-exchange members, i.e., their clients. To become a TMX exchange member, a firm must be a member of a self-regulatory organization (CIRO in Canada), have a CDS clearing agreement, and establish electronic access to TSX and/or TSX Venture Trading Engine.[6] In addition, a firm must pay an entry cost, which is relatively high for members who seek to be eligible to trade, roughly C$65,000. To keep the membership status, each exchange member must pay a monthly membership fee (in 2023 $1,500), in addition to trading fees, which are explained on TMX's website.

Exchange members can trade a variety of products, ranging from common stocks, and ETFs to more specialized products, such Exchange Traded Receipts (which let investors own gold bullion stored in the Royal Canadian Mint Gold Reserves). We group products in five categories, as explained in Appendix Table A2: large and small common company stocks, ETFs, non-common shares, and other/missing.

Derivatives market. Derivatives are traded over-the-counter, or on exchanges. We focus on exchange-traded derivatives. In Canada, there is a single derivative exchange, the Montreal Exchange (MX). It is owned by TMX group, and operates similarly to the other TMX exchanges.

To trade on the MX, a firm must become an MX exchange member. The requirements are similar to those for TMX. In particular, each MX member must be a CIRO member if the firm is Canadian and a member of the analogue regulatory entity of their nationality otherwise.[7]

`www.iiroc.ca/sections/markets/reports-statistics-and-other-information/reports-market-share-marketplace`, accessed on 08/10/2023.

[6]If the firm does not have a CDS clearing agreement, it must have a relationship with a clearing facilitator. For more details, see TMX's website.

[7]More specifically, requirements are different for Canadian and foreign firms. Canadian firms must be member of a Canadian self-regulatory organization (Investment Industry Regulatory Organization of Canada); must be a member of the Canadian Derivatives Clearing Corporation or conclude a clearing agreement with one of its members (MX). Foreign firms must be located in one of the following juris-

Table 3: Derivative products

Product	Trade share
Treasury Futures	37.44
Equity Options and Share Futures	27.77
Short Rate Derivatives	20.75
Bundles and Spreads	7.85
Index Options and Futures	6.51
Currency Options	0.04

Notes: Table 3 shows the daily average share of total trade-volume on the derivatives market, computed as the total amount of derivative contracts traded on a day, per product. Note that the amount of contracts does not reflect the value of the underlying assets. Appendix Table A3 describes each product category.

Members also have to pay MX-specific monthly membership fees and trading fees.

We group the derivative products into categories, closely following the MX website, as explain in Appendix Table A3. The largest category in terms of trade volume are Treasury futures, followed by equity options and share futures, and short rate derivatives (as shown Table 3). Trading activity in currency options is negligible.

Some derivative products, like Treasury futures or index futures, are highly standardized, while others are more complex. One example are 'user-defined-strategies' (UDS), which allow participants to create customized option strategies based on their individual risk management needs. We classify them under bundles and spreads since UDS tend to combine multiple derivative contracts. Even equity options and share futures are more complex than common stocks, because traders can specify the maturity and strike price, in addition to the underlying asset.

3 Data

We combine different data sources on five market segments, TSX, TSXV, Alpha, MX and the fixed-income market. These five market segments represent three markets: stock market (TSX, TSXV, Alpha), derivatives market, and the fixed-income market. The main data sources that allow us to observe trade information are proprietary to the TMX group and CIRO. We hand-collect publicly available information on CIRO and exchange members, financial products, and market conditions to enrich the data.

diction: United States, United Kingdom, Republic of Ireland, Israel, Jersey, the Netherlands and France; must be duly formed pursuant to the relevant laws of the country; must be registered with a securities or derivative instruments regulator, or a recognized self-regulatory organization, unless it is exempted from such registration in its jurisdiction and subject to all other applicable restriction; must have entered into a clearing agreement with a member of the Canadian Derivatives Clearing Corporation; must have a designated agent for service of process residing in Quebec (MX).

Fixed-income market. Our main data source for the fixed-income market is the Debt Securities Transaction Reporting System, MTRS2.0. This data base stores trades that involve at least one CIRO Member (who have an obligation to report all of their trades) since November 2015.[8] Our sample covers trades with all Canadian fixed-income products from 2019 until 2022. Trades between two institutions or individuals who are not CIRO Members are not reported. According to market experts, however, these trades are rare.

For each transaction we see which security is traded, and a series of security-characteristics which allows us to classify securities into product categories and assign industry sectors. We also observe the quantity and price of the trade, the time at which the trade is reported, and the side of the trade (buy/sell).

A rare feature of the MTRS2.0 data relative to most of the existing datasets that cover OTC markets is that most firms carry a unique identifier. In this study, we focus on CIRO dealers. Traders who act as dealers in the primary market have to report their own trades with their legal identifiers (LEIs). Other CIRO dealers are allowed (but not obligated) to mask their identity when they are reporting their own trades, but not when they are reported as counterparty (with LEIs). Given that most trades occur with at least one party acting as a primary dealer, masked identifiers are infrequent—roughly 5% of trades and 1% of trade volume.[9] Since we are interested in how much dealers buy and sell, we stack buy- and sell-side trades, and remove sales or purchases by non-dealers.

Equity and derivatives market. We observe trade-level data for all exchanges that are owned by the TMX group between 2019 and 2022. For each trade, we observe the time of the trade (up to milliseconds), the security (i.e., the TMX symbol), the amount, the price, and trading-firm IDs. For equities, we also see the best national bid and ask offer for each symbol that was valid right before each trade executes. Moreover, we know whether the trade is for the exchange member's own account or a client account. More specifically, for stock market trades, we can distinguish between an inventory account (IN), a client account (CL), and an account that members who are designated market makers use for their market making active (ST). Although there are a few other types of accounts, these are negligible. For trades in the derivatives

[8]A small group of Bank of Canada staff have access to the raw data, and this is anonymized before it can be shared, subject to a non-disclosure agreement, to external researchers.

[9]Whenever a masked dealer trades with a primary dealers or government distributor, it is possible to back out the identify of the masked dealer by relying on the fact that both dealers need to report the trade. The remaining trades by masked dealers are those between two masked dealers. In these cases, we cannot rule out the possibility that our data sample includes both sides of the trade due to double reporting. In all other cases, Bank of Canada staff has carefully removed one of the trade sides, so that each trade appears only once in our data set.

market, we observe analogous account-types.

To account for both sides of each transaction, we stack buy- and sell-side trades. Moreover, in order to link dealers across markets, we link the market-specific trading-IDs to each company's LEI. We achieve this by downloading exchange member lists containing trading IDs and company names for all instances where the TMX website was archived by the Wayback Machine during our sample period. We then identify each company's LEI using https://www.lei-lookup.com. Doing so, we account for mergers, acquisitions, and name changes over time.

Lastly, to categorize securities into products, we merge the data on stock market trades with publicly available listing information for each listed symbol in December of each year in our sample, relying on the wayback machine. Finally, to validate data quality, we verify that we observe the same daily trade volume and nominal value as CIRO and the MX exchange report publicly for the stock markets and the derivatives exchange. See Appendix A for details.

Dealers. A distinctive feature of our data is the ability to track firms registered as CIRO dealer members in the fixed-income market or as exchange members across these markets. Throughout the paper, we refer to these traders as 'dealers.' It is important to note that we do not classify firms as 'dealers' based on their trading or market-making activities. Instead, our definition relies solely on firms' membership status, which grants them the ability to place trade orders on their own behalf on exchanges and to trade with clients in OTC markets, regardless of their specific role in the market. We adopt this definition because it is more exogenous—at least conditional on market entry—than classifications based on endogenous trading behavior.

For each dealer LEI, we identify the LEI of its holding company parent using information from gleif.org. Doing so, we manually track mergers, acquisitions, and name changes found through Google searches to the best of our ability.

Further, we classify all LEIs and their parents into institution types: Table 4 shows that brokers constitute the largest category at the parent level, representing financial entities primarily engaged in brokerage services. Following them are asset managers and high-frequency traders, which include hedge funds, proprietary trading firms, and private equity firms. Investment banks come next, followed by other, typically smaller, banks and credit unions. At the LEI level, the dataset also includes some mutual funds and retail branches of larger institutions, such as banks, that focus on retail investing.

Summary statistics. Appendix Table A5 summarizes our trade data for stocks (TSX, TSXV, Alpha), bonds (MTRS), and derivatives (MX) to provide an overview of a typical trading day and trade in each market. The bond market is the largest in terms of trade volume. The number

Table 4: Dealer types

Institution type	# of LEIs	# of Parents
Asset Manager	15	21
Bank	7	17
Investment Bank	12	15
Broker	115	65
High-Frequency Trader	17	21
Mutual Fund	8	0
Pension Fund and Insurance	4	3
Retail	13	3
Other	3	2

Notes: Table 4 shows the number of LEIs and parent-LEIs of each dealer type at the LEI and parent-level. For the type classification we follow the methodology of the Bank of Canada used to classify institutions into types for the MTRS 2.0 data (explained in Appendix A). Appendix Table A4 defines each type category we observe in our data.

of dealers who actively trades on an average day is similar across markets, ranging from 48 in the derivatives market to 60 in the stock market.[10] For derivatives, trade size and volume reflect the number of contracts, not the underlying asset value. Similarly, the trade price reflects the price payed to exchange the derivative, i.e., the option fee in case the contract is an option, not the strike price.

4 Dealer specialization

We examine two types of specialization: market specialization, which may arise from differences in market-clearing rules and entry costs, and product specialization within a market, which could stem from variations in trading expertise or client relationships.

To preview, we will establish our first stylized fact in what follows:

Fact 1 (Specialization)**.** *Dealer trading is uneven both across markets and across products within a market—dealers specialize.*

To quantify specialization, we assign each dealer to their holding company. This uniformly removes any type of in-house segmentation for all dealers, which we know to play a role (as shown by Siriwardane (2019), and Lu and Wallen (2024), among others). As a result, our specialization measures are conservative and likely underestimates

[10]When the level of aggregation for institutions changes, the set of players adjusts slightly. This occurs because an entity identified by its LEI may not be classified as a dealer, whereas its parent institution might be.

Additionally, we distinguish between markets rather than market segments, treating TSX, TSXV, and Alpha as a single market. These exchanges are highly integrated, likely due to common ownership and structural similarities. Nearly all trade volume is executed by dealers active across all three exchanges (see Appendix Table A6).

Market specialization. Market specialization might arise because of different market structures. Bonds trade in a decentralized market that is not directly connected to stock markets, or the derivatives exchange. Moreover, each market is characterized by different entry costs.

Supporting the idea that frictions, or market-specific preferences, might hinder universal market participation, Figure 1a highlights that not all parent LEIs (the y-axis) participate in all markets (the x-axis). If a dealer trades at least once in a given market, we plot a black line for that market, otherwise, the line is white. Thus, if all dealers participated in every market, the graph would be entirely black. Instead, the presence of both black and white indicates that dealers do not participate in all markets.

Market activity, measured in trade volume, is also uneven across markets. This is evident for the 20 largest dealers in Figure 1b. If dealers maintained similar market shares—defined as their fraction of total trade volume in each market—the horizontal bars would have consistent color shading, with lighter shades indicating larger market shares.

Including all dealers, Figure 2 visualizes the pairwise correlation of market shares across markets.[11] We see that most dealers concentrate their trading in specific markets; if market shares were uniform across markets, the points would align along the 45-degree line. Furthermore, smaller dealers tend to specialize more strongly, as indicated by points near the axes, suggesting they transact almost exclusively in one market. Among larger dealers, a subset focuses more heavily on the bond market, primarily banks that act as primary dealers (see Appendix Figure A1a).

To control for the overall market size of a dealer, we introduce a market specialization score, which ranges between 0 (the dealer does not trade in the market under consideration) and 1 (the dealer only trades in the market). Formally, the score divides dealer j's market share in market m and year y, s_{yjm}, by the sum of the dealer's market shares in all markets:

$$\text{specialization}_{yjm} = \frac{s_{yjm}}{\sum_m s_{yjm}} \in [0,1]. \tag{1}$$

Figure 3 shows the market specialization scores for all dealers, averaged over the years. Many

[11] Appendix Figures A2 and A3 show similar correlation patterns when considering dealers at the LEI-level, and when excluding trades for client accounts.

Figure 1: Dealer presence and market shares per market

(a) Dealer presence

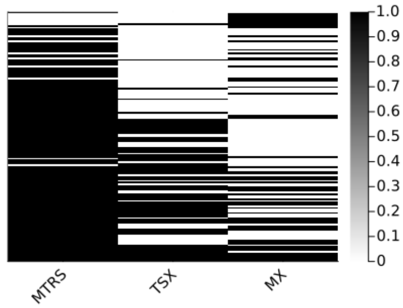

(b) Market shares

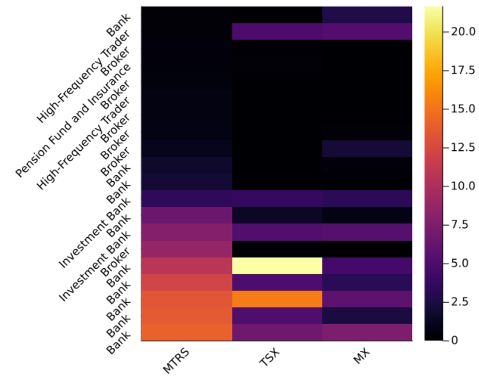

Notes: Figure 1a shows dealer presence in each market—bonds (MTRS), stocks (TSX), and derivatives (MX)—with black indicating that a dealer has traded at least once in the respective market. Figure 1b displays the average annual market shares of the 20 dealers with the highest average annual trade volume in any market. In both figures, each row represents a dealer, sorted by total trade volume across all markets. Dealers with the highest overall trade volume appear at the bottom, while those with the lowest appear at the top. Since the fixed-income market is the largest by trade volume, this sorting places dealers with relatively low participation in fixed income—but significant presence in other markets—toward the top.

dealers do not trade in all markets—their scores align along the x-or y-axes and on the diagonal that connects the 1 on the x-axis with the 1 on the y-axis. Some only trade in one market and therefore have a specialization score of 1.

Among the three markets, the derivatives market stands out as the most detached, with some dealers—particularly high-frequency traders—engaging almost exclusively in derivatives, even when their overall trading volume exceeds 5% (see Appendix Figure A1b). In total, approximately 35% of MX's trade volume is attributed to high-frequency traders when including client trades, rising to 72% when excluding them, as shown in Appendix Table A7. Around 25% of traders operate exclusively on MX, even at the parent-company level. This is largely driven by hedge funds, proprietary trading firms, and private equity firms that specialize in MX trading.

Banks acting as primary dealers are the most active dealers in linking markets (see Appendix Figures A4 and A5). In the bond market, nearly all dealers trading across all markets are primary dealers. On the stock market, primary dealers active in all markets account for approximately 68% of total trade volume. In the derivatives market, their share is lower at around 40%, yet no other dealer type plays a larger role in connecting derivatives with the other two markets. This underscores the significance of primary dealers beyond fixed income and highlights potential contagion risks during financial distress.

Figure 2: Dealer market shares, $s_{yjm} \in [0,1]$, in an average year

(a) All dealers (b) Largest dealers

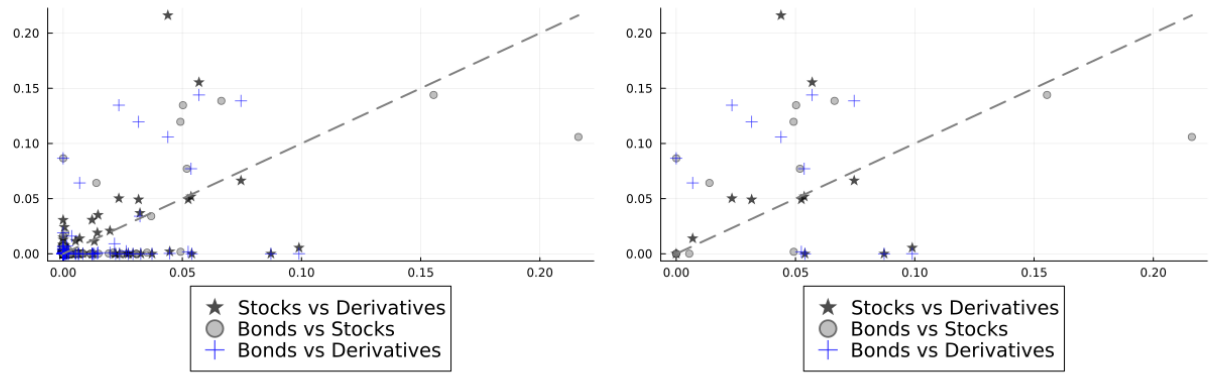

Notes: Figure 2a plots all dealer j's market shares, $s_{yjm} \in [0,1]$, for each market m, averaged across years. The stars show each dealer's stock market share on the y-axis and their derivatives market share on the x-axis; the circles show the stock market shares versus bond market shares, and the crosses the bond versus derivative market shares, on the y-axis and x-axis respectively. Figure 2b zooms in on dealers who trade at least 5% of the market share in one of the three markets.

Figure 3: Market specialization

(a) All dealers (b) Largest dealers

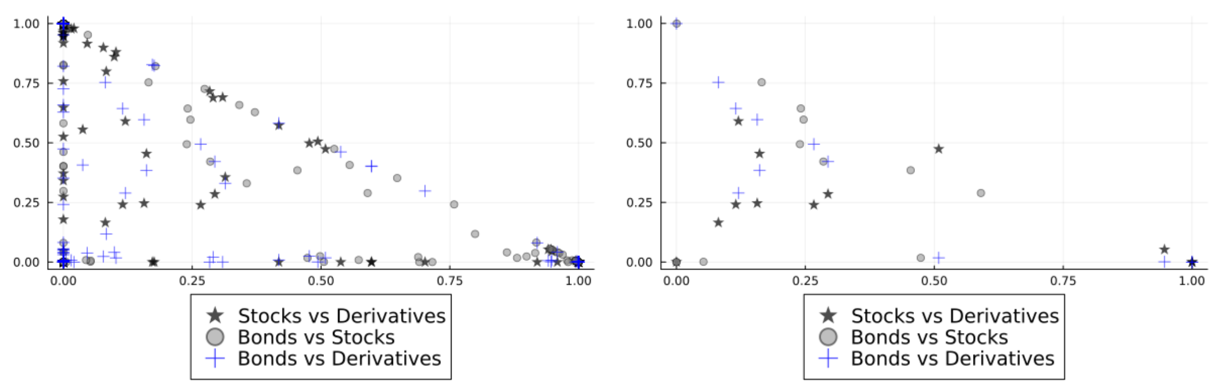

Notes: Figure 3a plots all dealer j's market specialization scores, $\text{specialization}_{yjm} = s_{yjm} / \sum_m s_{yjm} \in [0,1]$, for each market m, averaged across years. The stars show each dealer's stock market score on the y-axis and their derivatives market score on the x-axis; the circles show the stock versus bond market shares, and the crosses the bond versus derivative market score, on the y-axis and x-axis respectively. Figure zooms in on dealers who trade at least 5% of the market share in one of the three markets.

Product specialization. Product specialization may be driven by various factors, including differences in trading expertise, more effective inventory management, or relationships to clients with preferred habitat. On exchanges, some specialization also arises mechanically: firms designated as market makers are assigned subsets of securities and are obligated to trade them, leading to rule-based specialization. However, we show that this mechanical effect does not drive our results. When we exclude trades for market-making accounts (which designated market makers have to use when they are trading in their capacity as market maker), the main findings remain unchanged (see, for example, Appendix Figure 4 compared to Figure 4, which we explain below).

We assess the degree of product specialization within each market analogously to our assessment of market specialization.[12] First, Figure 4—analogous to Figure 1—visualizes dealer participation and product market shares—the fraction of a product's trade volume handled by each dealer—among the largest dealers across products. Comparing across markets, product specialization is lowest in the stock market. Not only do all dealers trade all products (as we see from the black box on the RHS of Figure 4a), dealers also distribute their trading more evenly across products. In the stock market colors on the LHS of Figure 4a are more consistent within dealers (horizontally) than across dealers (vertically), indicating a more even distribution of market shares, than in the bond and derivatives market.

Second, Figure 5—the analogue of Figure 3—presents the pairwise correlation of product specialization scores for a subset of products, which we define analogously to market specialization scores (1),

$$\text{specialization}_{yjmp} = \frac{s_{yjmp}}{\sum_p s_{yjmp}} \in [0,1], \tag{2}$$

where s_{yjmp} represents the fraction of dealer j's trade volume in product p within market m, relative to all other dealers. Since each market contains more than three products, the figure is less intuitive than, and not directly comparable to, Figure 3a. However, as with market specialization, dealers with scores near the x- or y-axis—or at the extreme value of 1 which means that the dealer only trades one product—demonstrate higher degrees of product specialization. Comparing across markets, we note that product specialization scores in the stock market (blue crosses) tend to be more moderate, whereas scores in the bond and derivatives markets are more frequently close to 1 or 0, reflecting stronger specialization.

Our interpretation of these empirical patterns is that product specialization is influenced

[12] In Appendix Figure A6 we further analyze the proportion of securities each dealer trades relative to the total number of securities in each market. For all dealers, this fraction is highest on the stock market, followed by the derivatives market, and then by the bond market.

Figure 4: Dealer presence and market shares across products in each market

(a) Stock market

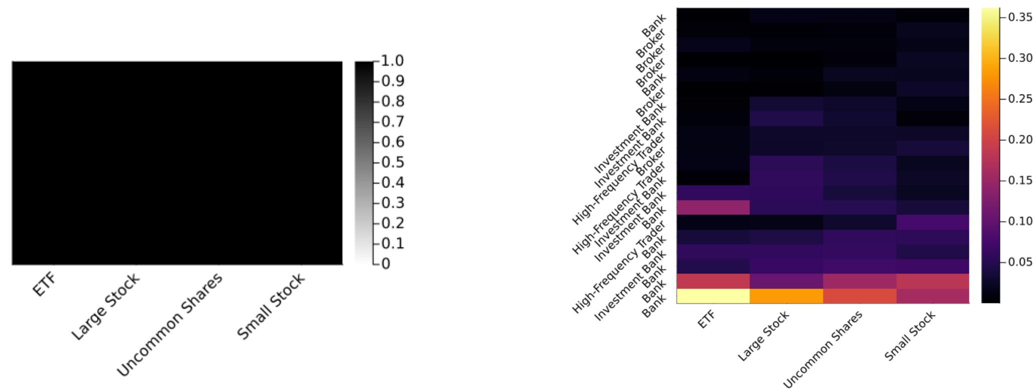

(b) Bond market

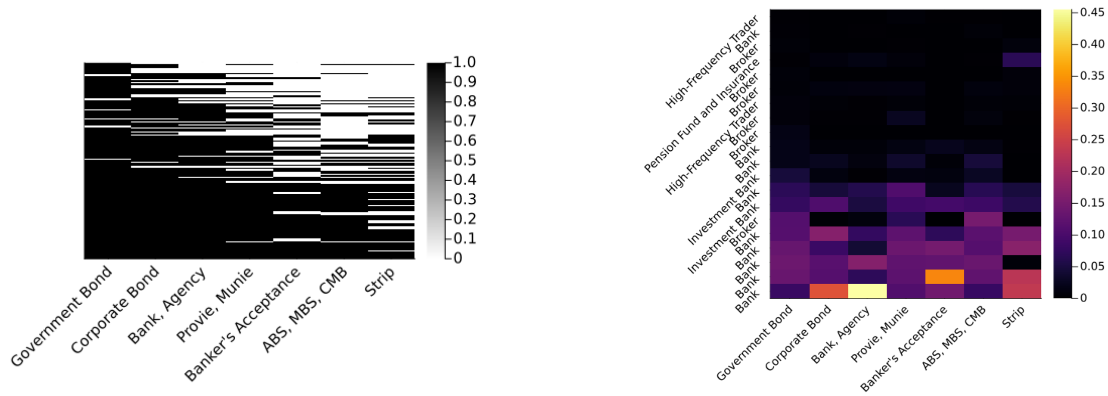

(c) Derivatives market

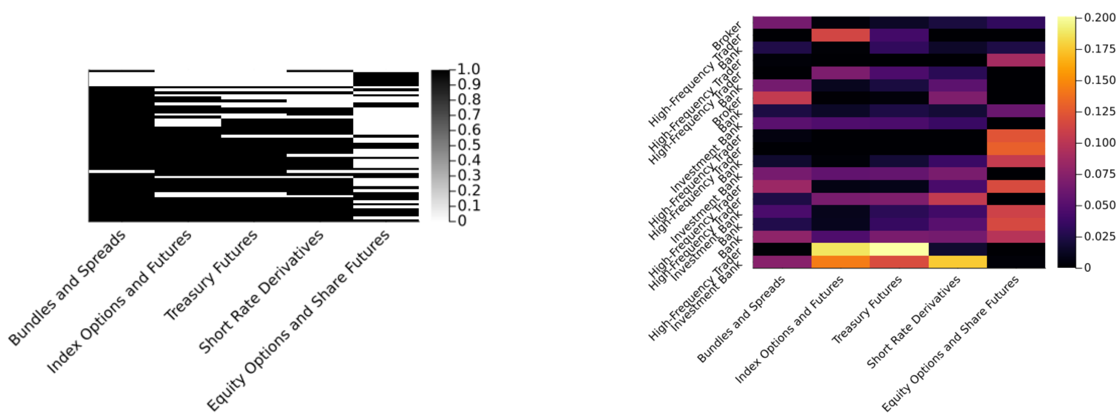

Notes: Figure 4 is similar to Figure 1 but replaces markets with products within each market, the stock market (a), bond market (b), and derivatives market (b), where we exclude currency options because they are so small. On the RHS we show whether dealers trade each product at least once in black. On the LHS we see the average annual product market shares of the largest dealers (on the LHS) for each market. In all figures, each row represents a dealer, sorted by total trade volume in the respective market. Dealers with the highest overall trade volume appear at the bottom, while those with the lowest appear at the top.

Figure 5: Product specialization

(a) Bond and stock market (b) Stock and derivatives market

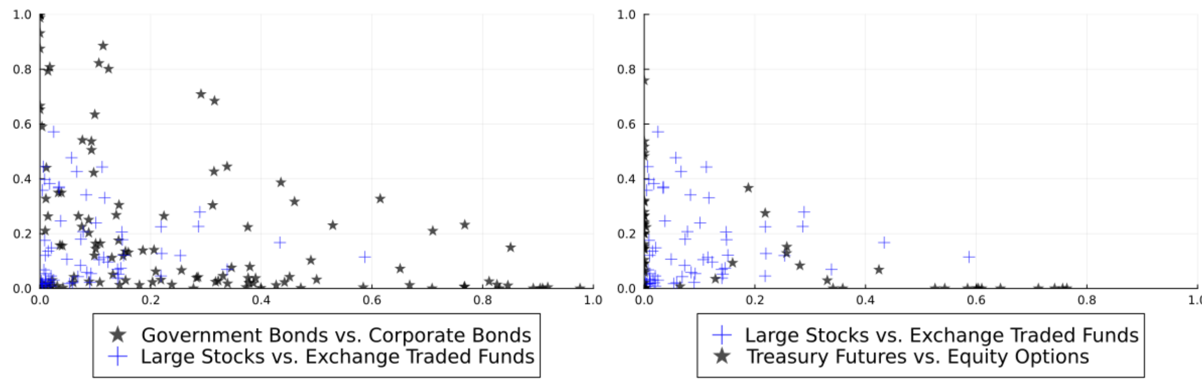

Notes: Figure 5 plots all dealer j's product specialization scores, $\text{specialization}_{yjmp} = s_{yjmp} / \sum_p s_{yjmp} \in [0, 1]$, for a subset of products p for each market m, averaged across years. In 5a we compare bonds versus stocks and in 5b we stocks versus derivatives. The crosses show the specialization scores for large stocks on the y-axis and ETFs on the x-axis. The stars in 5a show the scores for government bonds on the y-axis and for corporate bonds on the x-axis; they show the scores for Treasury futures versus equity options and share futures in 5b.

by both market structure—OTC versus exchange—and product complexity. On the centralized stock exchange, where products are largely standardized, specialization is minimal.[13] In contrast, it is more pronounced in the decentralized bond market and the derivatives exchange, which includes both standardized and complex products.

One reason for decentralized markets to feature higher product specialization is that trading is more strongly dictated by the network structure among dealers and clients. We know from existing studies that dealers form long lasting relationships with clients, and that clients tend to have tastes for specific bonds (Di Maggio et al. (2017); Hendershott et al. (2020); Jurkatis et al. (2023); Allen and Wittwer (2024)). This could attribute to stronger product specialization in OTC markets compared to stock markets where network structures and relationships are less relevant.

Yet, search frictions and relationships in OTC markets alone cannot explain product specialization; otherwise, we would not observe specialization on the derivative exchange, which operates similarly to stock exchanges. Unlike stock exchanges, the derivatives exchange accommodates both standardized products, like Treasury futures, and more complex derivative contracts. Standardized products are likely to attract a broader set of dealers, similar to the universal dealer presence in stock markets, as they require minimal customization—dealers simply

[13] To further support our idea, we examine different types of uncommon shares (such as preferred shares and debentures). We find that more complex products are not traded by all dealers (see Appendix Figure A7 and Appendix Table A8 for the complete list of suffices).

select from a predefined menu. Consistent with this, many more dealers trade Treasury futures, short-rate derivatives, and index futures—the most standardized derivatives—compared to options or share futures, despite the latter accounting for around 28% of total average daily trade volume. However, this pattern does not hold universally. Many dealers trade bundles and spreads, which are more complex, suggesting that product complexity alone does not determine specialization.

Comparing market and product specialization. Dealer specialization by market and product aligns with the growing literature on segmentation in financial markets and confirms our prior. However, it is less clear which type of specialization is more pronounced. This distinction matters for both policy and modeling, as it highlights the more relevant dimension of segmentation. We will now gather evidence to establish our second main fact:

Fact 2 (Market versus product specialization)**.** *Across-product specialization within a market is for most dealers larger than across-market specialization.*

To compare cross-market and cross-product specialization, we introduce a specialization index—based on Theil (1967)'s index—, which integrates the specialization scores from Figures 3 and 5 in a way that allows for direct comparison. The standard Theil T-Index, T_j, captures the distribution of dealer j's market share across product-market segments. If the dealer trades the same fraction of total volume in each segment, $T_j = 0$. A positive T_j indicates specialization, with higher values reflecting greater variation in the dealer's activity across segments.

For our purposes, the original index is not suitable because it fails to account for the fact that non-participation by dealers in a market or market-segment increases specialization. To address this, we introduce a non-participation cost, ξ, which applies when a dealer does not trade in a market segment. Since this punishment term is chosen arbitrarily, the magnitude of our index by itself is not informative. However, it is valuable for comparing specialization within a market across products to specialization across different markets, which is our primary objective. To see this, note that the index decomposes into two components: one measuring within-market specialization, T_j^w, and another measuring across-market specialization, T_j^a:

$$T_j = T_j^a + T_j^w, \text{ with } T_j^w = \frac{1}{M}\sum_m T_{jm}^w.$$

T_{jm}^w measures the how dealer j's market share in a market-product segment is distributed across products in market m, with more uneven distributions meaning higher specialization; and T_j^a

Table 5: Specialization decomposition

	Across-market (T_j^a)	Across-product (T_j^w)
Dealers who participate in all markets	0.02–1.10	0.09–5.22
Dealers who participate in two markets	2.09–2.77	1.87–5.87
Conditional on being active in only two markets	0.42-1.10	0.20–3.96
Dealer who participate in one market	–	3.48–6.71
Conditional on being active in only one market	–	0.15–1.95

Notes: Table 5 reports the range of across-market specialization indices (T_j^a) and within-market, cross-product specialization indices (T_j^w) for a punishment term of $\xi = 5$. Dealers are grouped by activity in all markets, two markets, or one market. For dealers who are active in $K \in \{1, 2\}$ markets, ξ increases both indices by $(3-K)/3\xi$. Indices are computed with and without accounting for non-participation penalties. For single-market dealers, only cross-product indices are calculated, since the cross-market index is uninformative.

measures how the average of product market shares in each market are distributed across markets. Consult Appendix B for mathematical details.

Table 5 provides the range of indices across dealers for three dealer groups, those active in all markets, those active in two out of three markets, and those active in only one market; Appendix Figure A11a shows all indices for each dealer separately. For dealers who are not active in all markets, we compute both indices subject to punishment for non-participation in the market they are not active in. In addition, we also compute the indices conditioning on market participation to compare cross-product specialization within a market across these dealer groups.

For most dealers, within-market specialization is larger than across-market specialization, implying larger cross-product specialization compared to cross-market specialization. Notably, this is not driven by the fact that there are more products in a market than entire markets, like it would when considering other measures, such as the variance. Instead, the difference is driven by unequal participation across submarkets.[14]

Greater product than market specialization suggests that, for large financial institutions, barriers to market entry are less restrictive than factors that limit arbitrage across products

[14]The magnitudes of the indices depend on the punishment ξ-term. However, the conclusion that within-market specialization is larger than across market specialization does not depend on the choice of ξ. To show this, we shut off non-participation punishments by setting ξ to zero. In that case, the average (median) within product index (across all dealers) is 0.83 (0.73), and the average across-product index is 0.91 (1.10). The indices are relatively similar in size, but this is driven by dealers who only participate in one market. For those dealers both indices are identical. When restricting attention to dealers who participate in at least two markets (for which the within-market and across market measures differ), the average (median) within product index (across all dealers) is 0.53 (0.45), which is significantly smaller than the average across-product index is 0.75 (0.83).

within a given market. Given that product specialization appears to depend on whether a market is centralized (e.g., exchanges) or decentralized (e.g., OTC markets) and on the complexity of the products traded, our findings underscore the need for theories that account for different market structures or different degrees product complexities, unlike most existing models that focus on a single market structure with standardized assets. While our analysis does not assess whether trading specialization improves welfare, it highlights the role of market design and product complexity in shaping market fragmentation.

5 Dealer specialization and transaction prices

Thus far, our analysis has focused on market shares—that is, quantities. The second part of the paper turns to prices. We ask whether market and product specialization affect transaction prices. This could occur if specialization improves inventory management, shifts beliefs about fundamentals, or allows some dealers to extract better prices in the presence of limited competition. We focus on price differences relative to market averages, leaving effects on aggregate equilibrium price levels for future research. Concretely, the remaining of our paper serves to establish our third stylized fact:

Fact 3 (Specialization and prices). *Dealers who are specialized trade at better prices relative to average market prices.*

To detect systematic differences in trade prices, we consider dealers at the LEI-level, but our findings are robust if we consider the dealers' parents instead. To reduce the sample size for our data from the exchanges, which is very large, we collapse the exchange data from the exchanges to the level of market segment (TSX/Alpha/TSXV/MX), day, security, dealer, trade direction (buy/sell), and trade type (active/passive). With slight abuse of terminology, we refer to each row in the collapsed dataset as a 'transaction' τ, and compute the total quantity traded, quantity_τ, and the average price, price_τ, for each τ.

Measuring price advantages. To detect systematic price differences across both sides of the trade and ensure comparability across securities with varying price levels, we follow the market microstructure literature and compare transaction prices relative to benchmark prices. Ideally, trading prices would be compared to fundamental values, but since these are rarely observable, equity studies use the prevailing mid-price, while bond market studies rely on inter-dealer prices, among other alternatives.

We seek a benchmark that is consistently available across markets and use each security's average daily price, which means our measure incorporates intra-day volatility. To account for variation across securities and over time, all regressions include security-week fixed effects. Moreover, to ensure the average price is meaningful, we focus on sufficiently liquid securities traded at least three times in a day, with results remaining qualitative robust when restricting to more liquid securities (e.g., those traded at least five times daily).[15] We also exclude approximately 2% of derivatives transactions executed at negative prices—common in certain spread types—as these would complicate the interpretation of our findings.

With these data, we want to measure the relative price advantage compared to the market. The easiest approach would be to consider the percentage difference of a transaction τ for security s relative to the average price for that security on that day t:

$$\text{Margin}_\tau = \frac{\text{average price}_{ts} - \text{trade price}_\tau}{\text{average price}_{ts}} \times 100 \times \text{ trade sign}_\tau$$

Trade sign is one when the trader buys and -1 when they sell. A 1% margin says that the trade price is 1% below the securities' daily avg. price when buying, and above when selling. However, since there are occasional outliers—which is common for trade-level data—we follow Hendershott and Madhavan (2015)'s measure, which is identical to Margin_τ for prices that are sufficiently close to the average price, and trims outliers:[16]

$$\text{margin}_\tau = -\ln(\text{trade price}_\tau / \text{average price}_{ts}) \times 100 \times \text{ trade sign}_\tau \approx \text{Margin}_\tau. \tag{3}$$

Margins vary significantly more in the derivatives market, where price volatility within a day is highest, followed by the stock market, and finally the bond market, where price volatility is more moderate (see Appendix Figures A13–A16). Due to differences in trade sizes and prices across markets, a 1% margin difference results in different total payment magnitudes for the median trade: approximately C$160 in the stock market, C$12,500 in the bond market, and 10 cents in the derivatives market.

Sufficient condition for price effects. Before analyzing how specialization affects transaction prices, we first verify that no market is frictionless enough to prevent some dealers to systematically outperform others. In fully frictionless markets, such differences would be ar-

[15] This restriction is particularly stringent in the derivatives market, as symbols often include detailed and flexibly specified contract information. Our restricted dataset covers over 99% of stock market trades, approximately 88% of bond market trades, and about 54% of derivatives trades.

[16] Appendix Figure A12 shows the relationship between our main margin measure (3) and its linear approximation. Relative to the linear percentage difference, the log-measure attenuates the poor trades (which, for buyers, are those executed at higher than average prices), and amplifies the successful trades.

bitraged away, and specialization would not impact prices—even if it influenced underlying inventory costs or beliefs.

We regress our margin measure on dealer-indicator variables separately for each market:

$$\text{margin}_\tau = \alpha + \sum_j \beta_j \mathbb{I}(\text{dealer} = j) + \gamma \text{control}_\tau + \zeta_t + \zeta_{ws} + \epsilon_\tau. \tag{4}$$

If there is no systematic difference of dealers across markets, all dealer coefficients should be statistically insignificant from zero. We include a day fixed effects, ζ_t, to absorb time-varying shocks that affect the entire market, and year-week-security fixed effects, ζ_{ws}, to account for the average weekly margin of a given security.[17] This addresses two potential biases. The first arises because different dealers trade different securities, which naturally have varying margins; the second comes from the feature that the set of traded securities varies over time.

Additionally, we include control variables, though they do not significantly affect the overall pattern of the dealer coefficients. First, given prior evidence that trade size influences outcomes, we control for trade size. Second, for exchange trades, we account for the account type associated with the trade. For bond trades, we distinguish the trade type—whether it occurs between dealers (i.e., CIRO dealer members), between a dealer and an inter-dealer broker, or between a dealer and a non-dealer. Across markets, we use the same large primary dealer as the baseline for consistency.

In this and all other margin regressions, we cluster standard errors at the daily level to account for arbitrary intra-day correlations across dealers, securities, and trades. This is crucial when traders split orders throughout the day or react to price shocks that impact multiple securities.[18] Because some days feature many more trades trade than others, the day-clusters are highly uneven in size. This can result in conventionally computed standard errors being underestimated (MacKinnon et al. (2023)). One common solution is to compute standard errors via wild (WCR) bootstrapping following Cameron et al. (2008), and Roodman et al. (2019). Unlike standard methods, this approach does not rely on asymptotic approximations to the

[17]We do not include day-security symbol fixed effects because some symbols are not traded frequently enough throughout the day. However, for robustness, we have estimated all regressions with symbol-day fixed effects, and our main conclusions remain unchanged.

[18]An alternative approach is to cluster by dealers, accounting for correlations in a dealer's trades across days while ignoring intra-day correlations across dealers. However, with fewer than 100 dealers in the stock and derivatives markets and uneven cluster sizes, we are not confident this would yield reliable standard errors, even when boostrapping standard errors (MacKinnon et al. (2023)). Another option is to cluster at the symbol level to capture correlated shocks affecting the same symbol over time. We do not adopt this approach because many price shocks are likely correlated across symbols, making symbol-level clustering insufficient for addressing cross-symbol dependencies.

test statistic's distribution, which can be inaccurate when clusters are uneven or few in number. Instead, it constructs confidence intervals using bootstrap resampling, and therefore yields more reliable test statistics when clusters are small or uneven. To ensure robustness, we report coefficients that are statistically significant under both wild-bootstrapped and conventionally computed standard errors.

Our findings (Figure 6 and Appendix Figure A19) show that some dealers consistently secure better prices across all markets, both at the LEI- and parent-level. In the stock market, the best dealer (a large broker) achieves margins 0.64% better than the baseline (a primary dealer), while the worst (a smaller broker) lags by 0.31%, translating into an annual benefit of approximately C$266 million for the best dealer and a C$6 million loss for the worst.[19] Dealer differences are more pronounced in the derivative market due to price volatility, with the best (a large hedge fund) outperforming by 1.38% and the worst (a proprietary trading firm) underperforming by 1.30%, though total payment differences remain modest given contract prices and trade volumes. In the bond market, most dealers earn lower margins than the dominant primary dealer baseline, yet the best (a large insurance company) outperforms by 0.08%, gaining C$17 million annually, while the worst (serving retail clients) underperforms by 0.48%, losing C$278 thousand.

Table 6 examines which dealer types achieve better margins using cross-dealer variation by estimating regression (4) with dealer-type indicators, setting asset managers as the baseline.[20] High-frequency traders outperform other types in derivatives and perform well in stocks, though mutual funds dominate. Pension funds and insurance companies, the weakest performers in stocks, achieve the highest bond margins. Dealers specializing in retail clients earn the lowest margins in bonds and derivatives but perform comparably to asset managers in stocks.

Specialization affects prices. Having established that no market is sufficiently frictionless to eliminate systematic price differences, we next ask whether and how specialization affects transaction prices. In theory, the relationship between dealer specialization and prices is ambiguous. Greater specialization might enable dealers to trade at lower prices. However, it could also be that specialized dealers are more efficient than their less specialized counterparts and,

[19] To translate margin percentages into annual monetary losses or gains, we assume that each trade is executed at the median price, using the total amount (e.g., number of shares in the stock market) that the dealer under consideration trades in an average year.

[20] Consistent with prior studies (Bernhardt et al. (2005)), larger trades receive worse prices on exchanges, and execution prices are worse for client trades than for inventory or market-making accounts. In the bond market, we find no trade-size discounts, aligning with mixed literature (Pinter et al. (2024); Allen and Wittwer (2024)). Dealers earn higher margins trading with clients than with other dealers or brokers.

Table 6: Margin regression with dealer-types (bonds, stocks, derivatives)

Bond market	Margin	Stock market	Margin	Derivative market	Margin
Trade size	+0.000	Trade size	−0.162***	Trade size	−24.255***
	(0.000)		(0.006)		(2.249)
	[0.000]		[0.006]		[1.909]
Counterparty is broker	+0.008**	Client account	−0.069***	Client account	−0.275***
	(0.003)		(0.003)		(0.066)
	[0.002]		[0.001]		[0.028]
Counterparty is client	+0.018***	Inventory account	−0.000	Inventory account	+1.332***
	(0.002)		(0.003)		(0.072)
	[0.002]		[0.002]		[0.030]
Bank	0.021*	Market-maker account	+0.070***	Bank	−0.304**
	(0.010)		(0.003)		(0.097)
	[0.004]		[0.002]		[0.078]
Broker	+0.003	Bank	+0.071***	Broker	+0.388***
	(0.007)		(0.002)		(0.034)
	[0.002]		[0.002]		[0.026]
High-Frequency Trader	0.027	Broker	+0.045***	High-Frequency Trader	+1.001***
	(0.032)		(0.001)		(0.041)
	[0.028]		[0.001]		[0.031]
Investment Bank	−0.049**	High-Frequency Trader	+0.115***	Investment Bank	+0.597***
	(0.015)		(0.004)		(0.043)
	[0.011]		[0.004]		[0.036]
Mutual Fund	−0.013	Investment Bank	+0.009***	Other	+0.404
	(0.008)		(0.002)		(0.275)
	[0.004]		[0.002]		[0.261]
Pension Fund and Insurance	+0.101***	Mutual Fund	+0.159***	Retail	−0.394***
	(0.008)		(0.019)		(0.067)
	[0.003]		[0.019]		[0.051]
Retail	−0.160***	Pension Fund and Insurance	−0.046***		
	(0.009)		(0.003)		
	[0.003]		[0.003]		
		Retail	+0.026***		
			(0.002)		
			[0.002]		
Date-& symbol-year-week fes	Yes	Date-& symbol-year-week fes	Yes	Date-& symbol-year-week fes	Yes
N	6,757,118	N	111,051,211	N	4,529,584
R^2	0.017	R^2	0.008	R^2	0.036
Within-R^2	0.000	Within-R^2	0.001	Within-R^2	0.006

Notes: Table 6 shows the estimation results from regressing margins (3) of stock market trades on trade-size, the account-type (client, inventory, market-market, or other—the baseline), dealer-types, and day- and security-year-week fixed effects on the LHS, for fixed-income trades in the middle, and for derivatives on the RHS. For bonds, we replace the account-type with the type of trade (dealer-broker, dealer-client, dealer-dealer—the baseline). Asset managers are the baseline. Margins are in percentage points, and quantities are in million C$. We cluster standard errors at the daily-level and report conventionally computed robust clustered standard errors in round brackets, and wild-bootstrapped standard errors in squared brackets. The stars reflect to the larger standard errors.

Figure 6: Dealer coefficients that are statistically different from zero at 5% significance level

(a) Stocks (b) Bonds (c) Derivatives

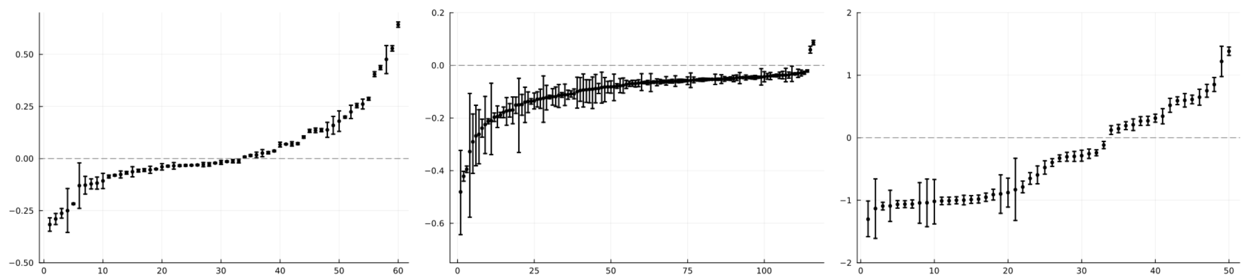

Notes: Figure 6a shows the dealer coefficients, and 95% wild-bootstrapped confidence intervals (clustered at the daily-level), when regressing margins (3) of a stock market trade on dealer indicator variables and control variables (trade-size, the account-type, security-week-year and day fixed effects). Figures 6b and 6c show the analogue for the bond and derivatives market, respectively. For the bond market, we replace the account-type with a variable that indicates the type of trade (dealer-dealer, dealer-client, dealer-broker). In all graphs we exclude dealer coefficients that are not significantly different from zero at a significance level of 5% according bootstrapped and conventional inference to be conservative. Since we sort coefficients from small to large, the x-axis are not comparable across markets, as they do not reflect dealer identifiers. Appendix Figures A18 shows the analogous figures with conventionally computed confidence intervals; Appendix Figures A19 aggregates dealers to the parent-level.

Figure 7: Cross-market correlation between dealer coefficients of dealers active in all markets

(a) Bonds-stocks (b) Bonds-derivatives (c) Stocks-derivatives

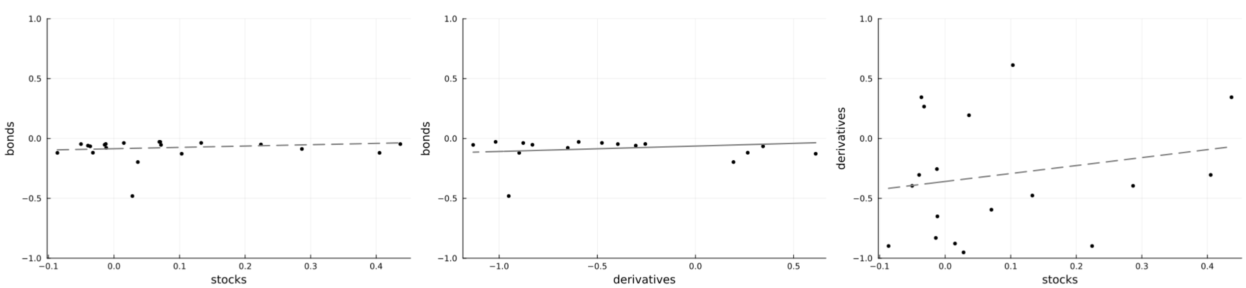

Notes: Figure 7a shows the within-dealer correlation of coefficients in the bond (y-axis) versus stock market (x-axis), 7b and 7c show the correlation for the other two market pairs. We exclude dealer coefficients that are not significantly different from zero at a significance level of 5% according to bootstrapped and conventional inference to be conservative. Appendix Figure A21 shows the cross-market correlation between dealer coefficients when aggregating dealers to the parent-level.

as a result, are willing to trade at prices that less specialized dealers avoid. To find out we use two complementary strategies: examining dealers individually and analyzing patterns in the cross-section.

We begin by zooming in on the dealer level to assess whether dealers who trade across markets and products achieve better prices, or whether specialized dealers outperform. Specifically, we examine dealer fixed effects from regression (4) for dealers exclusively active in a single market—an extreme form of specialization. Appendix Figure A17 shows that many of these dealers earn worse-than-average prices relative to the baseline dealer, potentially due to lower trading volume. However, outcomes vary by market. In the bond market, the second-best dealer is a bond-only trader, suggesting successful specialization. In the derivatives market, specialization appears even more advantageous: a substantial share of high-performing dealers trade only derivatives, possibly reflecting the relative complexity of these products compared to equities and bonds.

Supporting the notion of market specialization, we find little evidence that dealers who perform well in one market also perform well in others. Figure 7 shows no significant cross-market correlations in dealer fixed effects from regression (4) for dealers active in multiple markets. Appendix C provides analogous evidence in favor of strong product specialization, showing that few dealers outperform across multiple products.

Next, we leverage cross-sectional variation across dealers in their product and market specialization scores (1) and (2). Ideally, we would want to know if specialization causes specialized dealers to trade at different prices compared to non-specialized ones.

A first naive approach would be to regress margins on specialization scores, in addition to the same control variables and fixed effects we have used above, for regression (4).

$$\text{margin}_{\tau} = \alpha + \beta \text{market specialization}_{yjm} + \gamma \text{controls}_{\tau} + \zeta_t + \zeta_{ws} + \epsilon_{\tau}, \tag{5}$$

and similarly with product specialization, which varies by year, dealer, market and product: product specialization$_{yjmp}$.[21]

A natural and pressing concern in interpreting these regressions is reverse causality: specialization may help dealers obtain better prices, but better prices may also lead dealers to specialize. As a result, the estimates may be biased, reflecting the broader endogeneity of prices and quantities in equilibrium.

We adopt two complementary approaches to mitigate endogeneity concerns. Our first approach is to lag specialization scores to examine whether dealers who were more specialized

[21]For completeness, we report estimation outputs from regression (5) in Appendix Table A9

last year obtain better prices today. Specifically, we estimation regression (5), but replace market specialization$_{yjm}$ by market specialization$_{y-1,jm}$ and similarly for the product specialization scores. This would provide the causal effect of specialization on margins if past specialization is exogenous to current pricing—that is, if dealers did not choose last year's specialization based on anticipated future margin opportunities, and if there are no unobserved time-varying dealer-level factors affecting both specialization and pricing.

Table 7 presents the regression results, showing that dealers with higher specialization scores in the previous year earn larger margins today. The relationship is stronger for product specialization than for market specialization. For example, on the stock exchange, moving from no market specialization (score of 0) to full specialization (score of 1) increases margins by 3.8 basis points; a full shift in product specialization raises margins by 41 basis points. Given the large trading volumes dealers manage over the year, even these seemingly modest per-trade gains translate into substantial monetary value.

Our second approach is to instrument for specialization scores in regression (5)—that is, to find an observable variable that affects margins only through its influence on specialization. Identifying valid instruments that allow us to disentangle quantity and price effects is notoriously difficult, as emphasized in the growing literature on demand estimation following Koijen and Yogo (2019).[22]

We use client orders on the stock exchanges—where we observe a sufficient volume of such orders—as an instrument for dealer specialization when dealers trade for their own accounts.[23] The idea is that dealers have limited discretion over client orders, which must be executed promptly. For example, a retail investor placing a stock order through a Fidelity brokerage account will have that order executed automatically by Fidelity. These client orders generate variation in dealer specialization that is plausibly unrelated to the margins dealers earn on their own-account trades. If this exclusion restriction holds—conditional on our standard control variables and fixed effects—the instrument allows us to identify the causal effect of specialization on stock market margin.

[22]Common instruments include stock index inclusions (e.g., Shleifer (1986); Chang et al. (2015); Pavlova and Sikorskaya (2023)); capital flows (e.g., Coval and Stafford (2007); Ben-David et al. (2022)); announcements of quantitative easing (e.g., Krishnamurthy and Vissing-Jørgensen (2011)); COVID-19 stimulus programs (e.g., Greenwood et al. (2022)); variation in government bond supply (e.g., Krishnamurthy and Vissing-Jørgensen (2012)); unexpected inventory shocks to dealers (e.g., Allen and Wittwer (2023)); and regulatory constraints such as investment mandates (e.g., Koijen and Yogo (2019)).

[23]We cannot apply the same strategy to the derivatives or bond markets. For derivatives, too few active traders receive enough client orders to construct a meaningful instrument. For bonds, client orders are not observed.

Table 7: Correlation between margins and last year's specialization scores

	Stocks		Bonds		Derivatives	
Lagged market specialization	0.038***		0.008		0.218***	
	(0.002)		(0.006)		(0.049)	
	[0.001]		[0.001]		[0.011]	
Lagged product specialization		0.413***		0.065***		0.290***
		(0.004)		(0.014)		(0.052)
		[0.001]		[0.004]		[0.009]
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	87,299,519	87,299,519	5,358,555	5,358,555	3,350,296	3,350,296
R^2	0.008	0.009	0.017	0.017	0.038	0.038
Within-R^2	0.000	0.001	0.000	0.000	0.000	0.000

Table 7 shows the estimation results from regressing margins (3) on our specialization measures (1) and (2) from last year, respectively, for each market separately, using all trades. In all regressions we include the same control variables and fixed effects as in regression (6): trade size, account-types for the exchange, trade-type for the bond market, dealer-types, date fixed effects and security-year-week fixed effects. We cluster standard errors at the daily-level and report conventionally computed robust clustered standard errors in round brackets, and wild-bootstrapped standard errors in squared brackets. The stars reflect to the larger standard errors.

The IV-estimates, reported in Tables 8, suggest that more specialized dealers obtain prices that are roughly 30 basis points better than average.[24] The OLS estimate for market specialization is smaller than the IV estimate, while the OLS estimate for product specialization is larger, though the difference is not statistically significant at the 5% level. A downward bias in the OLS estimate could arise from omitted variables that induce a negative correlation between the error term and specialization. For example, more specialized dealers may be more efficient and thus willing to accept lower margins. Conversely, an upward bias could result from positive correlations—for instance, if dealers choose to specialize in products or markets where they already enjoy favorable pricing conditions due to strong client relationships or reputational advantages.

Neither of our strategies to address endogeneity is without limitations. For example, if dealers are forward-looking and build specialization in anticipation of future pricing advantages, lagged specialization scores are not exogenous. Similarly, the IV approach hinges on an exclusion restriction that could be violated if clients systematically direct orders to dealers who secure better prices on their own-account trades. However, we view this risk as limited: clients do not observe transaction prices, which are private to exchange members and the exchange

[24]Appendix Table A10 shows the analogous results for the derivative exchange, where we include margins from all trades, not just trades for dealer-accounts to obtain sufficient power. The exclusion restriction is therefore more restrictive.

Table 8: IV regressions of margins on specialization scores for stocks

	(First Stage)	(OLS)	(IV)
s^c_{yjm}	-0.442***		
	(0.003)		
	[0.002]		
Market specialization		0.134***	0.385***
		(0.004)	(0.024)
		[0.001]	[0.000]
Controls	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes
N	27,300,881	30,356,466	27,300,881
R^2	0.349	0.124	0.126
Within-R^2	0.012	0.000	0.000

	(First Stage)	(OLS)	(IV)
s^c_{yjmp}	-0.206***		
	(0.002)		
	[0.001]		
Product specialization		0.335***	0.284***
		(0.008)	(0.037)
		[0.002]	[0.001]
Controls	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes
N	27,300,881	30,356,466	27,300,881
R^2	0.349	0.124	0.126
Within-R^2	0.012	0.000	0.000

Table 8 shows the IV estimation results for own-account stock market trades. Consider the LHS of 8. In column (First Stage), we show the first stage of the two stage least square estimator—regressing the market specialization score (1) on the fraction of all client-orders dealer j executes in market m in year y relative to other dealers, $s^c_{yjm} \in [0, 1]$. In column (OLS) we present the OLS coefficient from regressing margins on market specialization, using trades for the dealer's own account for the stock market, and all trades for the derivatives market. In column (IV) we depict the corresponding IV estimate. The table on the RHS shows the analogous for product specialization, where the instrument is the fraction of all client-orders for product p dealer j executes in market m in year y relative to other dealers, $s^c_{yjmp} \in [0, 1]$. In all regressions we include the same control variables and fixed effects as in regression (6): trade size, account-types for the exchange, dealer-types, date fixed effects and security-year-week fixed effects. We cluster standard errors at the daily-level and report conventionally computed robust clustered standard errors in round brackets, and wild-bootstrapped standard errors in squared brackets. The stars reflect to the larger standard errors.

operator, making such selection unlikely to be a first-order concern. More importantly, both approaches—despite their individual limitations—yield a consistent pattern: more specialized dealers obtain better prices.

6 Conclusion

We analyze dealer specialization across bond, stock, and derivative markets using a unique dataset that tracks trading activity across all major Canadian financial markets. Our findings show that product specialization within a market is stronger than market specialization, though not all dealers participate in every market. While no market is frictionless enough to prevent some dealers from consistently securing better prices, we find no evidence of cross-market or cross-product trading synergies. These results challenge the traditional view that financial intermediaries operate seamlessly across markets and products, and underscore the importance of market structure and product complexity in driving market fragmentation.

References

Adrian, T., E. Etula, and T. Muir (2014). Financial intermediaries and the cross-section of asset returns. *The Journal of Finance 69*(6), 2557–2596.

Allen, F. and D. Gale (2000). Financial contagion. *Journal of Political Economy 108*(1), 1–33.

Allen, J., J. Kastl, and M. Wittwer (2020). Estimating demand systems with bidding data. Working paper.

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

Allen, J. and M. Wittwer (2024). Bunlidng trades in over-the-counter markets. Working paper.

Anand, S. and P. Segal (2015). Chapter 11: The global distribution of income. In: Handbook of Income Distribution, A. B. Atkinson, Bourguignon, F. (eds.), 2: 937-979.

Babus, A. and K. Hachem (2023). Markets for financial innovation. *Journal of Economic Theory 208*, 1–39.

Babus, A., S. Moreira, and M. Marzani (2024). The rise of specialized financial products. Working paper.

Ben-David, I., J. Li, A. Rossi, and Y. Song (2022). Ratings-driven demand and systematic price fluctuations. *The Review of Financial Studies 35*, 2790–2838.

Bernhardt, D., V. Dvoracek, E. Hughson, and I. M. Werner (2005). Why do larger orders receive discounts on the london stock exchange? *The Review of Financial Studies 18*(4), 1343–1368.

Bessembinder, H., C. Spatt, and K. Venkataraman (2020). A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis 55*(1), 1–45.

Brunnermeier, M. K. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *The Review of Financial Studies 22*(6), 2201–2238.

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review 104*(2), 379–421.

Budish, E., R. S. Lee, and J. J. Shim (2024). A theory of stock exchange competition and innovation: Will the market fix the market? *Journal of Political Economy 132*(4), 1209–1246.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economic Statistics 90*, 414–427.

Chang, Y.-C., H. Hong, and I. Liskovich (2015). Regression discontinuity and the price effects of stock market indexing. *The Review of Financial Studies 28*(1), 212–246.

Chaudhary, M., Z. Fu, and J. Li (2022). Corporate bond multipliers: Substitutes matter. Working paper.

Chen, D. and D. Duffie (2021). Market fragmentation. *American Economic Review 111*(7), 2247–2274.

Chen, Z., N. Roussanov, X. Wang, and D. Zou (2024). Common risk factors in the returns on stocks, bonds (and options), redux. Working paper.

CIRO (2024). Proposed integrated fee model. Public announcement, available at: https://www.ciro.ca/newsroom/publications/proposed-integrated-fee-model?utm_source=chatgpt.com.

Coval, J. and E. Stafford (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics 86*(2), 479–512.

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

Dix, R. and M. Wittwer (2025). Estimating demand systems with trade data. Working Paper.

Dougast, J., S. Üslü, and P.-O. Weill (2022). The theory of participation in otc and centralized markets. *The Review of Economic Studies 89*(6), 3223–3266.

Du, W., A. Tepper, and A. Verdelhan (2018). Deviations from covered interest rate parity. *The Journal of Finance 73*(3), 915–957.

FINRA (2024). Self-regulatory organizations; financial industry regulatory authority, inc.; notice of filing and immediate effectiveness of a proposed rule change to adjust FINRA fees to provide sustainable funding for FINRA's regulatory mission. Public announcement, available at: https://www.federalregister.gov/documents/2024/11/27/2024-27764/self-regulatory-organizations-financial-industry-regulatory-authority-inc-notice utm_source=chatgpt.com.

Greenwood, R., T. Laarits, and J. Wurgler (2022). Stock market stimulus. Working paper.

Gromb, D. and D. Vayanos (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics 66*, 361–407.

Hasbrouck, J. and G. Sofianos (1993). The trades of market makers: An empirical analysis of NYSE specialists. *The Journal of Finance 48*(5), 1565–1593.

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

He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics 126*(1), 1–35.

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *American Economic Review 103*(2), 732–770.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020). Relationship trading in over-the-counter markets. *The Journal of Finance 75*(2), 683–743.

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

Jurkatis, S., A. Schrimpf, K. Todorov, and N. Vause (2023). Relationship discounts in corporate bond trading. *BIS Working Papers, No. 1140*, 1–46.

Koijen, R. S. J. and M. Yogo (2019). A demand system approach to asset pricing. *Journal of Political Economy 127*(4), 1475–1515.

Krishnamurthy, A. and A. Vissing-Jørgensen (2011, Fall). The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity*.

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

Kumar, P. and D. J. Seppi (1992). Futures manipulation with "cash settlement". *Journal of Finance 47*(4), 1485–1502.

Lu, L. and J. Wallen (2024). What do bank trading desks do? Working paper.

MacKinnon, J. G., M. O. Nielsen, and M. D. Webb (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics 232*, 272–299.

Malamud, S. and M. Rostek (2017). Decentralized exchange. *American Economic Review 107*(11), 3320–3362.

Menkveld, A. J. (2016). The economics of high-frequency trading: Taking stock. *Annual Reviewof Financial Economics 8*, 1–24.

Mota, L. and K. Siani (2024). Financially sophisticated firms. Working paper.

O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics 116*, 257–270.

O'Hara, M. and X. A. Zhou (2021). Anatomy of a liquidity crisis: Corporate bonds in the COVID-19 crisis. *Journal of Financial Economics 142*(1), 46–68.

Pasquariello, P. (2014). Financial market dislocations. *The Review of Financial Studies 27*(10), 1868–1914.

Pavlova, A. and T. Sikorskaya (2023). Benchmaring intensity. *The Review of Financial Studies 36*(3), 859–903.

Pinter, G., C. Wang, and J. Zou (2024). Size discount and size penalty: Trading costs in bond markets. *Review of Financial Studies 37*(7), 2156–2190.

Roodman, D., M. O. Nielsen, J. G. MacKinnon, and M. Webb (2019). Fast and wild: Bootstrap inference in stata using boottest. *The STATA Journal 19*(1), 4–60.

Sandulescu, M. (2020). How integrated are corporate bond and stock markets? Working paper.

Shleifer, A. (1986). Do demand curves for stocks slope down? *The Journal of Finance 41*(3), 579–590.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance 52*(1), 35–55.

Siriwardane, E., A. Sunderam, and J. L. Wallen (2022). Segmented arbitrage. Working paper.

Siriwardane, E. N. (2019). Limited investment capital and credit spreads. *The Journal of Finance 74*(5), 2303–2347.

Theil, H. (1967). *Economics and Information Theory*. North-Holland Publishing Company.

Üslü, S. and G. Pintér (2023). Comparing search and intermediation frictions across markets. Working paper.

Vayanos, D. and J.-L. Vila (2021). A preferred-habitat model of the term structure of interest rates. *Econometrica 89*(1), 77–112.

Weill, P.-O. (2020). The search theory of over-the-counter markets. *Annual Review of Economics 12*, 747–773.

Wittwer, M. and J. Allen (2023). Market power and capital constraints. Working paper.

ONLINE APPENDIX

Market and Product Specialization in Financial Markets

by Milena Wittwer, and Andreas Uthemann

Section A provides details regarding data cleaning.
Section B provides mathematical details for our specialization indices.

A Data cleaning

Data Restrictions. Our bond data includes all bond trades that are reported by CIRO-dealers to MTRS.20, excluding foreign sovereign bonds. We exclude primary market trades. In rare cases, trades are reported on a weekend. We treat those cases as trades that occur on the Monday following the weekend.

We include all stock-market trades, including those executed during the opening and closing auctions. In rare cases, trades are associated with negative trade amounts. We exclude those trades.

We keep regular derivative trades, and excludes rare cases of trades involving 'test futures'. In rare cases, trades are reported on a weekend. We treat those cases as trades that occur on the Monday following the weekend. We exclude a handful of dates where only Buy-ins are trading.

Quality Check. We compare the average monthly trade volume on the stock markets with the publically available information that is provided in CIRO's website to confirm that we observe close to 100% of the trades we should observe.

We also compare the derivative trade volume with information provided on MX's website. After restricting the raw data, as explained above, we observe roughly 90% of trade volume on average.

Our bond data is provided directly by the regulator and serves as the source for publicly available information on bond market trading volumes. We therefore do not cross-validate it against reported figures.

Type Classification for Dealers. The Bank of Canada classifies traders based on their LEIs into types following their in-house methodology. We replicate their approach to classify dealers on the stock exchanges and the derivative exchange, and to classify the parent-holding company of each LEI (see Appendix Table A4). Here we briefly describe their approach.

We use two types of information to classify entities – "Direct" and "Indirect" sources of information. "Direct" information refers to any information provided by the entity itself - either through its official website, internal documents, spokespersons, a regulatory organization to which it reports, etc. "Indirect" information refers to any information which is not direct information. The latter is further broken down into two sub-types: "Reliable" or "Weakly reliable". For the dealers in this project, all information comes from reliable sources, such as Bloomberg, Yahoo Finance, CapEdge, etc.

B Theil Index

To define the dealer-specific measure, consider a fixed dealer j. Let there be M markets, indexed by m, and P_m products within each market m, and $P = \sum_m P_m$ products overall, indexed by p. Denote dealer j's share of total volume traded by dealers in product-market segment mp by $s_{jmp} \in [0,1]$. The cross-product average for a dealer within a market is given by $\bar{s}_{jm} = \frac{1}{P}\sum_p s_{jmp}$, and $\bar{s}_j = \frac{1}{M \times P}\sum_m \sum_p s_{jmp}$ the overall average.

The standard Theil T index in this setting is defined as:

$$T_j = \frac{1}{M \times P}\sum_m \sum_p \left(\frac{s_{jmp}}{\bar{s}_j}\right) \ln\left(\frac{s_{jmp}}{\bar{s}_j}\right). \tag{6}$$

This index captures the distribution of dealer j's market share across product-market segments. If the dealer trades the same fraction of total volume in each segment, $T_j = 0$. A positive T_j indicates specialization, with higher values reflecting greater variation in the dealer's activity across segments.

For our purposes, the original index is not suitable because it fails to account for the fact that non-participation by dealers in a market or market-segment increases specialization. The original formulation only sums the trade volume of dealers who are active in a market, ignoring those who are inactive. To address this, we introduce a non-participation cost, ξ, which applies when a dealer does not trade in market mm (i.e., when $s_{jm} = 0$):

$$T_j = \underbrace{\frac{1}{M \times P}\sum_m \sum_p \mathbb{I}(s_{jmp} > 0)\left(\frac{s_{jmp}}{\bar{s}_j}\right) \ln\left(\frac{s_{jmp}}{\bar{s}_j}\right)}_{\text{Standard Theil index conditional on participation}} + \underbrace{\frac{1}{M \times P}\sum_m \sum_p \mathbb{I}(s_{jmp} = 0)\xi}_{\text{Non-participation}}.$$

The magnitude of the index depends on the size of the penalty, ξ, which can be chosen arbitrarily. As a result, the index by itself is not informative in absolute terms. However, it is valuable for comparing specialization within a market across products to specialization across different markets, which is our primary objective.

To see this, note that the index decomposes into two components: one measuring within-market specialization, T_j^w, and another measuring across-market specialization, T_j^a:

$$T_j = T_j^a + T_j^w, \text{ with } T_j^w = \frac{1}{M}\sum_m T_{jm}^w, \text{ where}$$

$$T_{jm}^w = \frac{1}{P}\sum_p \mathbb{I}(s_{jmp} > 0 \cup \bar{s}_{jm} > 0)\left(\frac{s_{jmp}}{\bar{s}_j}\right)\ln\left(\frac{s_{jmp}}{\bar{s}_{jm}}\right) + \xi\mathbb{I}(s_{jmp} = 0 \cup \bar{s}_{jm} > 0)$$

measures the how dealer j's market share in a market-product segment is distributed across products in market m, with more uneven distributions meaning higher specialization; and

$$T_j^a = \frac{1}{M}\sum_m \mathbb{I}(\bar{s}_{jm} > 0)\left(\frac{\bar{s}_{jm}}{\bar{s}_j}\right)\ln\left(\frac{\bar{s}_{jm}}{\bar{s}_j}\right) + \xi\mathbb{I}(\bar{s}_{jm} = 0)$$

measures how the average of this market share across products is distributed across markets.

As for the standard Theil index, the minimum value for both measures is 0, which is the case when a dealer distributes their trading activity evenly across products in a market for T_{jm}^w, or on average across markets for T_j^a. The maximum value is given by M, P_m and ξ, namely, $\overline{T}_{jm}^w = \frac{1}{P_m}[MP_m\ln(P_m) + (P_m - 1)\xi]$, and $\overline{T}_j^a = \frac{1}{M}[M\ln(M) + (M-1)\xi]$.

C Additional evidence: product specialization and prices

To detect cross-product price effects, we add product indicators to regression (4), and estimate the following regression for each market separately:

$$\text{margin}_\tau = \alpha + \sum_p\sum_j \beta_{jp}\mathbb{I}(\text{dealer} = j \text{ and product} = p) + \gamma\cdot\text{control}_\tau + \zeta_t + \zeta_{ws} + \epsilon_\tau. \quad (7)$$

If the same dealer obtains similar margins across products within a market compared to the baseline, all β_{jp} coefficients would be similar in size. If the dealer is more successful when trading some products relative to others, these coefficients would differ. As before, we include day and security-week fixed effects to avoid potential biases that arise from time-variation in the traded securities.[25] For bonds, the baseline is a large primary dealer trading government bonds, for stocks it is that bank trading large stocks, and for derivative it is that bank trading

[25]As robustness, we also estimate a specification with only include day-fixed effect to exploit variation of margins across securities within the same product category. While the size of the coefficients differs, the main take away (Fact **??** is robust).

Treasury futures.[26]

We visualize the estimation outcome through heatmaps, one for each market, in Appendix Figure A22. Since estimating regression (7) is computationally intensive, especially for the stock market, we estimate it for each of the years in our sample separately, and report results for 2022. A row in the heatmap correspond to a dealer j. A column corresponds to a product p. When dealer j obtains systematically worse margins for product p, we color the corresponding jp cell red, meaning that the β_{jp} coefficient is negative and statistically different from zero. The cell is black if the dealer outperforms the other dealers, and empty if they either do not trade product p or the coefficient is not statistically significantly different from zero.

If the one dealer were to outperform (underperformed) the baseline across products, we would observe a black (red) line for that dealer. This is not the case for any dealer in any market—a take away that is robust across years, while the β estimates vary. Crucially, since the margin measure reflects price volatility over a day, this analysis does not imply that some products are inherently more profitable—we do not account for underlying trading costs. Rather, the key takeaway is that no dealer consistently outperforms all others across products, pointing towards product specialization in all markets.

[26]We clustered at the daily-level, and compute standard errors via wild-bootstrapping. This is useful not only because it circumvents issues that arise from uneven cluster sizes, but also because many indicator variables in regression (7) are zero, since dealers tend to specialize in specific products. This implies that the standard cluster-robust covariance matrix is close to singular (non-invertible) due to high correlation within some clusters with many zeros. Since wild-bootstrapping resamples residuals with cluster-dependent perturbations, and does not directly rely on inverting the covariance matrix, bootstrapping circumventing the issue.

Appendix Table A1: Fixed-income products

Product	Description
Government Bond	Government of Canada Bond, Government of Canada Real Return Bond, Government of Canada T-bill
Corporate Bond	Corporate Bond
Provie, Munie	Provincial Bill, Provincial Bond, Provincial Commercial Paper, Municipal Bond
Bank, Agency Paper	Bank Commercial Paper and Bank Security - Note/Bond/Debenture. Agency Bond and Agency Commercial Paper
Bankers' Acceptance	Bankers' Acceptance
ABS, MBS, CMB	Mortgage-Backed Security, Asset-Backed Security, Canada Mortgage Bond.
Strip	Agency Strip Bond, Bank Strip Bond, Corporate Strip Bond, Finance company Strip Bond, Government of Canada Strip Bond, Municipal Strip Bond, Provincial Strip Bond

Appendix Table A2: Equity products

Product	Description
Large Stock	Symbols without suffices (i.e., common shares) that are listed with missing sp-type with more than 2 billion of quoted market value
Small Stock	Symbols without suffices (i.e., common shares) that are listed with missing sp-type with less than 2 billion of quoted market value
Exchange Traded Funds	Symbols that are listed with sp-type being Exchange Traded Funds
Uncommon Shares	Symbols which suffices that aren't listed as Exchange Traded Funds, which include the the following types: preferred stocks, class A-C, notes, debentures, equity dividends, when-issued capital pool companies, warrants, redeemable common stocks, U.S. funds, units, subscr. receipts, and stocks that trade on the NEX market
Others or Missing	Symbols without suffices that have a non-missing sp-type, which include the following sp-types: Income Trust, Fund of Equities, Commodity Funds, Exchange Traded Receipt, Split Shares, Fund of Mortgages/MBS, Fund of Debt

Appendix Table A3: Derivative products

Product	Description/Symbols if available
Treasury futures	Government bond futures and future options; CGZ, CGF, CGB, LGB, OGZ, OGF, OGB
Short-term derivatives	BAX futures and future options, CORRA futures; BAX, OBW, OBX, OBY, or OBZ, CRA
Equity options and share futures	Equity option, weekly option, option on ETFs, share futures
Currency options	Options on USD; USX
Index options and futures	Index futures and options; SXF, SCF, SXB, SXY, SXK, SXJ, SEG, SXM, SXA, SXH, SXO, SXU, SXV, SCG, SDV
Bundles and spreads	User-defined strategy, inter-group strategies, spreads

Notes: Appendix Tables A1–A3 describe our product classification for the bond, stock, and derivative market, respectively.

Appendix Table A4: Dealer types classification

Broker	Financial entity whose purpose is to offer brokerage services
Investment Bank	Investment bank
Bank	Bank, retail bank or credit union, and any entity that is deposit taking
Asset Manger	Financial entity whose purpose is to manage assets (or investments) and/or offer investment advising services. Entities that manage multiple types of funds such as HF, MF or ETFs are also classified as such
Mutual Fund	Financial entity that is a mutual fund or a mutual fund manager
High-Frequency Trader	Financial entity that is a hedge fund or a hedge fund manager; Private Equity, or Proprietary Trader
Pension Fund and Insurance	Financial entity whose purpose is to manage investments (and/or provide services) related to pension, retirement, insurance, re-insurance, benefits, and superannuation funds
Retail	Financial entity whose purpose is to offer financial services to retail (non-institutional) investors
Other	This category includes all other types which we observe in our traded data. From the Bank of Canada classification we pool the following types under this category "Real Estate (a financial or non-financial entity that is involved in the construction, financing, management, or sale of commercial, industrial, or residential real estate), "Other" (Financial entity that does not fall in any of the aforementioned classifiers (e.g., Financial Planner, Financial Research Services, Execution Platform)), Uncategorized (entity that can neither be classified as a financial nor a non-financial entity due to lack of information), "Non-financial entity". We also include Buy-Ins that execute some trades on the exchanges here.

Notes: Appendix Table A4 explains the classification of trader types we adopt following the methodology of Bank of Canada staff.

Appendix Table A5: Daily trade volume, number of active dealers, trade-sizes, and prices

Variable	**Mean**	**Median**	**Min**	**Max**	**Std**
Daily Trade Volume					
Stocks (in mil)	661.053	613.824	155.859	1746.480	209.427
Bonds (in bn C$)	71.723	68.410	0.196	1,504.660	54.771
Derivatives (in k)	340.366	323.155	1.487	1,017.150	126.632
Number of Active Dealers (LEIs)					
Stocks	64.157	64.0	61.0	69.0	1.630
Bonds	61.041	62.0	5.0	84.0	7.794
Derivatives	53.617	54.0	21.0	57.0	3.019
Number of Active Dealers (Parents)					
Stocks	60.705	61.0	58.0	64.0	1.180
Bonds	57.394	58.0	11.0	71.0	6.636
Derivatives	47.737	48.0	21.0	51.0	2.668
Trade size					
Stocks	11875.7	1300.0	0.1	7.361×10^7	74,694.1
Bonds (in mil C$)	10.934	1.250	$1.0/10^6$	7.189×10^5	436.177
Derivatives	78.143	10.0	1.0	170,478.0	667.365
Trade price					
Stocks (in C$)	27.280	12.240	0.005	2,392.4	86.936
Bonds (in C$)	102.344	100.24	1.0	980.0	12.398
Derivatives (in C$)	20.809	1.060	-142.050	21,800.0	123.298

Notes: Appendix Table A5 summarizes trade data for stocks (TSX, TSXV, Alpha), bonds (MTRS), and derivatives (MX) from 2019 to 2022. It provides the mean, median, minimum, maximum, and standard deviation of daily trade volume, the number of active dealers ("Number of Dealers (LEI)") and parent institutions ("Parents"), trade size, and trade price. For derivatives, trade size and volume reflect the number of contracts, not the underlying asset value. There are 1,004 active trading days for stocks, 994 for derivatives, and 1,035 for bonds. Some bond trades occur on Canadian holidays, when the Investment Industry Association of Canada (IIAC) recommends pausing trading. These days, typically involving minimal activity, are excluded from the table but included in the analysis with either lower-frequency aggregation or day-fixed effects to account for special cases. The stock market features 6,449 symbols traded by 72 dealers and 66 parent institutions. The derivatives market, where symbols often include contract details like expiration dates, has 503,056 symbols traded by 64 dealers and 56 parent institutions. In the bond market, 107,516 CUSIPs are traded by 163 dealers (CIRO dealer members in the raw data) and 131 parent institutions. Only a subset of these dealers is active daily.

Appendix Table A6: Member intersection: RHS: LEI-level of members; LHS—Parent level

Intersection	**TSX**	**ALPH**	**TSXV**	**Intersection**	**TSX**	**ALPH**	**TSXV**
All	99.97	100.0	99.68	All	99.97	100.0	99.68
TSX and ALPH	0.0	0.0	0.0	TSX and ALPH	0.0	0.0	0.0
TSX and TSXV	0.03	0.0	0.32	TSX and TSXV	0.03	0.0	0.32
ALPH and TSXV	0.0	0.0	0.0	ALPH and TSXV	0.0	0.0	0.0
TSX only	0.0	0.0	0.0	TSX only	0.0	0.0	0.0
ALPH only	0.0	0.0	0.0	ALPH only	0.0	0.0	0.0
TSXV only	0.0	0.0	0.0	TSXV only	0.0	0.0	0.0

Notes: Appendix Table A6 shows the percentage of total volume traded in each of the three stock exchanges (TSX, TSXV, and Alpha) that is traded by brokers who trade on all segments (All), on TSX and Alpha, etc. Each column sums to 100%. Total volume traded is computed by summing all quantities of all brokers including both sides of the trade.

Appendix Table A7: Avg. weekly share traded by dealers per type (parent-level)

Dealer type	**MTRS**	**TSX**	**TSX-IN**	**MX**	**MX-IN**
Asset Manager	0.27	1.70	0.98	0.68	0.00
Bank	74.13	56.05	35.77	27.12	19.05
Broker	12.32	13.47	16.46	9.12	2.02
High-Frequency Trader	0.86	10.70	24.24	34.94	72.07
Investment Bank	11.88	17.21	22.49	27.98	6.75
Other	0.00	0.00	0.00	0.11	1.11
Pension Fund and Insurance	0.50	0.23	0.03	0.00	0.00
Retail	0.02	0.60	0.00	0.07	0.00

Notes: Appendix Table A7 shows the fraction of trade volume by dealer type (at the parent-level) per market in columns MTRS, TSX, and MX, respectively. In columns TSX-IN, and MX-IN we show the analogue but excluding trades for client accounts for TSX and MX.

Appendix Table A8: Symbol Suffixes on TSX/TSXV/Alpha

Symbol suffix	Description
None	Common shares
A, B, C	Class A, B, C of shares is typically related to voting rights, access to dividend
DB	Debenture, stock type that makes fixed payments at scheduled intervals of time, operates similar to preferred stock
E	Equity dividend
F	
G	
H	NEX market provides a trading forum for listed companies that no longer meet the TSX Venture's ongoing listing standards; designed for companies that have low levels of business activity or have ceased to carry on active business. It benefits such companies by giving their stocks a degree of liquidity and providing visibility that may attract potential acquirers or investors.
K	NEX market
IR	Installment receipts, is an equity issuance in which the purchaser does not pay the full value of the issue up front. In the purchase of an installment receipt, an initial payment is made to the issuer at the time the issue closes; the remaining balance must be paid in installments, usually within a two-year period. Although the purchaser has not paid the full value of the issue, he or she is still entitled to full voting rights and dividends.
J	
L	Legended shares. A legend is a statement on a stock certificate noting restrictions on the transfer of the stock. A stock legend is typically put in place due to the requirements established by the Securities and Exchange Commission (SEC) for unregistered securities. A stock legend may or may not be legally required on the certificate itself, depending on state laws.
M	Booms
N	Subscription receipts (second issue trading)
O	Subscription receipts (third issue trading)
NO, NS, NT	Exchange traded note (ETN) are unsecured debt securities that tracks an underlying index of securities.
P	Capital pool company (CPC) is an alternative way for private companies in Canada to raise capital and go public. The capital pool company system was created and is currently regulated by the TMX Group, and the resulting companies trade on the TSX Venture Exchange in Toronto, Canada.
Q	
PR, PF, PS	Preferred shares; similar to common shares, no maturity date, ownership, fixed distribution rate, no voting rights
PR.CLASS, PS.CLASS, PF.CLASS	Preferred class
R	Subscription receipts are defined as those limited term securities issued via prospectus, which are convertible into another security class of the issuer (predominantly common shares) at a set conversion rate based on the successful completion of a planned reorganization or transaction. Where completion is not successful, security proceeds are either returned to the subscriber or a more generous conversion rate is made available to the subscriber.
RT	Rights are instruments issued by companies to provide current shareholders with the opportunity to preserve their fraction of corporate ownership. Rights are short-term instruments that expire quickly, usually within 30-60 days of issuance. The exercise price of rights is always set below the current market price, and no commission is charged for their redemption.

Appendix Table A9: Correlation between margins and specialization scores

	Stocks		Bonds		Derivatives	
market specialization	0.022***		0.007		0.440***	
	(0.002)		(0.005)		(0.039)	
	[0.000]		[0.001]		[0.008]	
product specialization		0.388***		0.062***		0.237***
		(0.004)		(0.009)		(0.043)
		[0.001]		[0.003]		[0.008]
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	111,051,211	111,051,211	6,757,118	6,757,118	4,529,585	4,529,585
R^2	0.008	0.009	0.017	0.017	0.036	0.036
Within-R^2	0.000	0.001	0.000	0.000	0.000	0.000

Appendix Table A9 shows the estimation results from regressing margins (3) on our specialization measures (1) and (2),respectively, for each market separately, using all trades. In all regressions we include the same control variables and fixed effects as in regression (6): trade size, account-types for the exchange, trade-type for the bond market, dealer-types, date fixed effects and security-year-week fixed effects. We cluster standard errors at the daily-level and report conventionally computed robust clustered standard errors in round brackets, and wild-bootstrapped standard errors in squared brackets. The stars reflect to the larger standard errors.

S	Special U.S. terms
T	Special US trading terms (second issue trading)
U	U.S. dollar
V	U.S. dollar (second issue trading)
UN	Units are a securities that is made up of one common share and half a warrant. Units are commonly offered by special-purpose acquisition companies, or SPACs that are seeking to raise money in a public stock offering and trade on a stock exchange with the primary goal of merging with a private business and taking it public.
W	When issued
WB	
WR	
I	When issued (second issue trading)
WT	Warrants give the holder the right to purchase a company's stock at a specific price and at a specific date.
X	
Y	Redeemable common. Redeemable shares are shares that a company has agreed it will, or may, redeem (in other words buy back) at some future date. The shareholder will still have the right to sell or transfer the shares subject to the articles of association or any shareholders' agreement.

Notes: Appendix Table A8 describes the meaning of all suffixes of symbols trades on TSX, TSXV, or Alpha. An empty cell means that we were not able to find the description of a symbol that we observe in the raw data.

Appendix Table A10: IV regressions of margins on specialization scores for derivatives

	(First Stage)	(OLS)	(IV)		(First Stage)	(OLS)	(IV)
s^c_{yjm}	-1.291***			s^c_{yjmp}	-1.353***		
	(0.061)				(0.039)		
	[0.014]				[0.007]		
Market specialization		0.440***	1.691***	Product specialization		0.237***	0.389*
		(0.039)	(0.258)			(0.043)	(0.155)
		[0.008]	[0.007]			[0.008]	[0.008]
Controls	Yes	Yes	Yes	Controls	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Fixed effects	Yes	Yes	Yes
N	2,911,210	4,529,585	2,911,210	N	2,911,210	4,529,585	2,911,210
R^2	0.495	0.036	0.125	R^2	0.534	0.036	0.125
Within-R^2	0.032	0.000	0.000	Within-R^2	0.164	0.000	0.000

Appendix Table A10 show the IV estimation results for all trades on the derivatives market. Consider the LHS of 8. In column (First Stage), we show the first stage of the two stage least square estimator—regressing the market specialization score (1) on the fraction of all client-orders dealer j executes in market m in year y relative to other dealers, $s^c_{yjm} \in [0,1]$. In column (OLS) we present the OLS coefficient from regressing margins on market specialization, using trades for the dealer's own account for the stock market, and all trades for the derivatives market. In column (IV) we depict the corresponding IV estimate. The table on the RHS shows the analogous for product specialization, where the instrument is the fraction of all client-orders for product p dealer j executes in market m in year y relative to other dealers, $s^c_{yjmp} \in [0,1]$. In all regressions we include the same control variables and fixed effects as in regression (6): trade size, account-types for the exchange, dealer-types, date fixed effects and security-year-week fixed effects. We cluster standard errors at the daily-level and report conventionally computed robust clustered standard errors in round brackets, and wild-bootstrapped standard errors in squared brackets. The stars reflect to the larger standard errors.

Appendix Figure A1: Dealer market shares, $s_{yjm} \in [0,1]$, in an average year (parent-level)

(a) Banks

(b) High-frequency traders

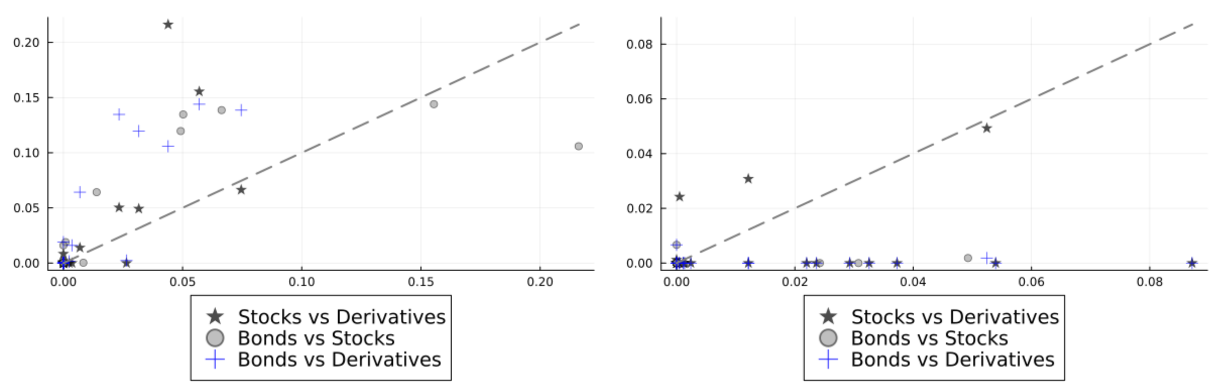

Notes: Appendix Figures A1a and A1b are analogous to Figure 2a but only includes banks, and high-frequency traders, respectively. It plots all bank/high-frequency dealer j's market shares for each market m, averaged across years. The stars show each dealer's stock market share on the y-axis and their derivatives market share on the x-axis; the circles show the stock versus bond market shares and the crosses the bond versus derivative market shares, on the y-axis and x-axis respectively.

Appendix Figure A2: Dealer market shares, excluding client-accounts, in an average year

(a) All dealers

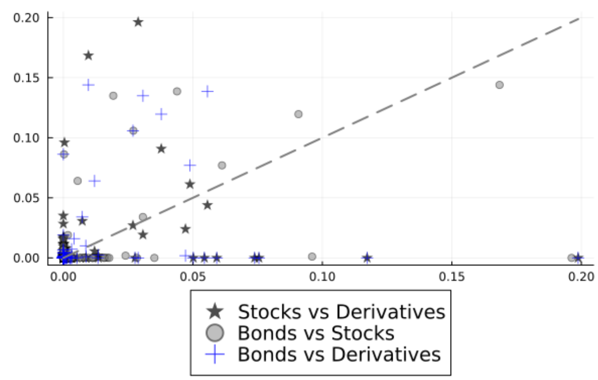

(b) Largest dealers

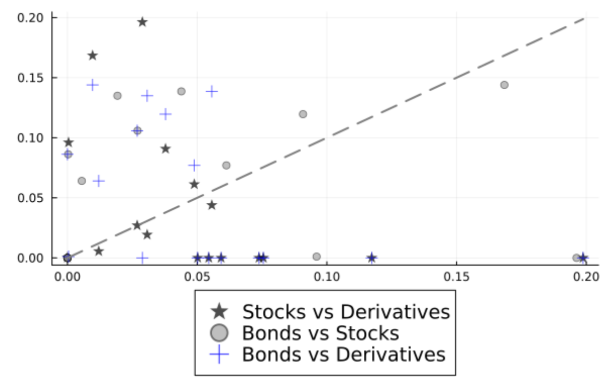

(c) Banks

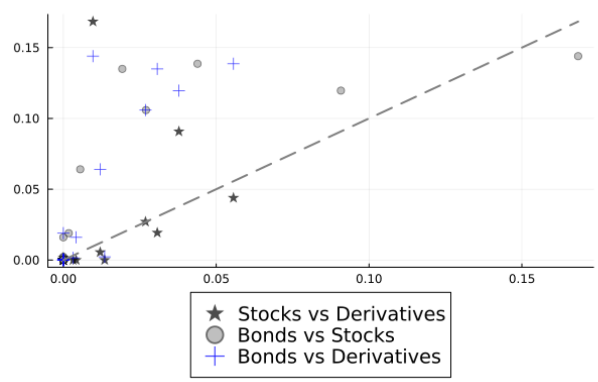

(d) High-frequency traders

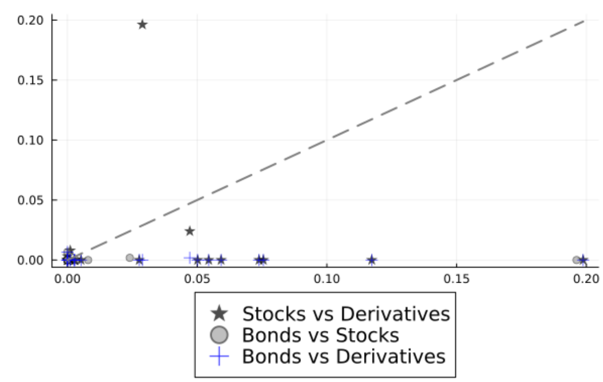

Notes: Appendix Figure A2a is analogous to Figure 2a but excludes trades for client accounts on the exchanges. It plots all dealer *j*'s market shares (of trades for non-client accounts) for each market *m*, averaged across years. The stars show each dealer's stock market share on the y-axis and their derivatives market share on the x-axis; the circles show the stock versus bond market shares and the crosses the bond versus derivative market shares, on the y-axis and x-axis respectively. Figure zooms in on dealers who trade at least 5% of the non-client market share in one of the three markets. In Figures A2c and A2d we only consider banks, and high-frequency traders, respectively.

Appendix Figure A3: Dealer market shares, excluding client-accounts, in an average year (LEI-level)

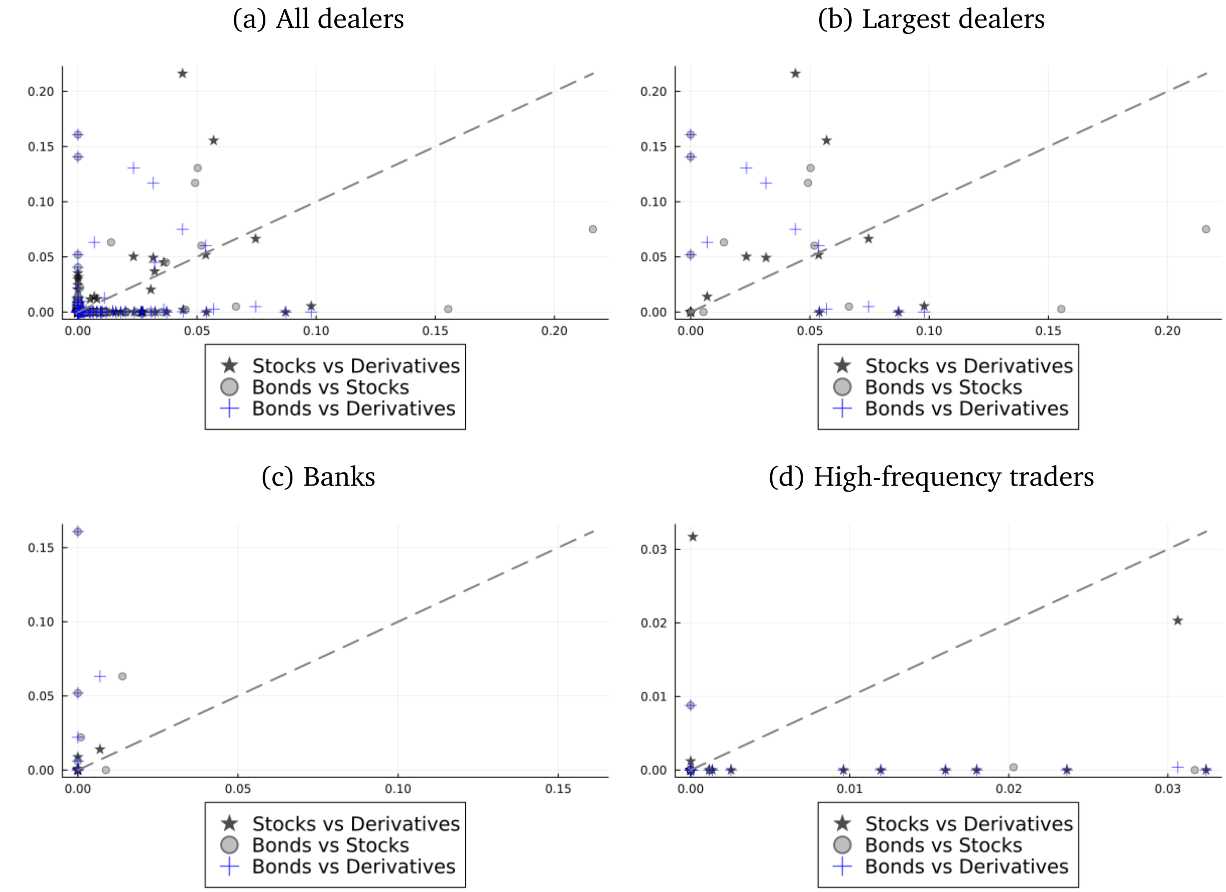

Notes: Appendix Figure A3 is analogous to Figure 2a but for dealers at the LEI-level rather than the parent-level. It plots all dealer *j*'s market shares for each market *m*, averaged across years. The stars show each dealer's stock market share on the y-axis and their derivatives market share on the x-axis; the circles show the stock versus bond market shares and the crosses the bond versus derivative market shares, on the y-axis and x-axis respectively. Figure zooms in on dealers who trade at least 5% of the non-client market share in one of the three markets. In Figures A3c and A3d we only consider banks, and high-frequency traders, respectively.

Appendix Figure A4: Fraction of monthly trade volume by dealers active in all markets

(a) Including all trades

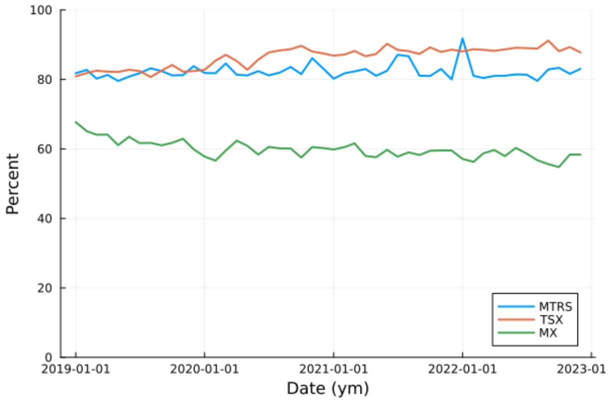

(b) Excluding client trades (for TSX/MX)

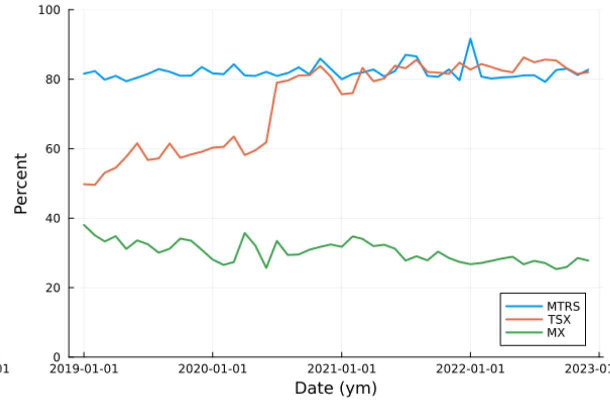

Notes: Appendix Figure A4a shows the fraction of monthly trade volume by those active in all markets, for each market (MTRS, TSX, MX). Appendix Figure A4b excludes trades for client accounts.

Appendix Figure A5: Fraction of monthly trade volume by primary dealers active in all markets

(a) Including all trades

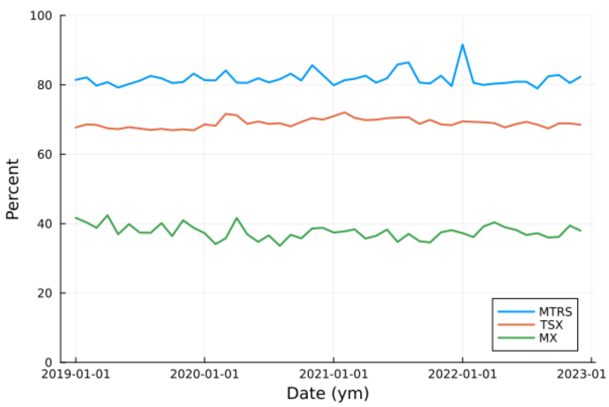

(b) Excluding client trades (for TSX/MX)

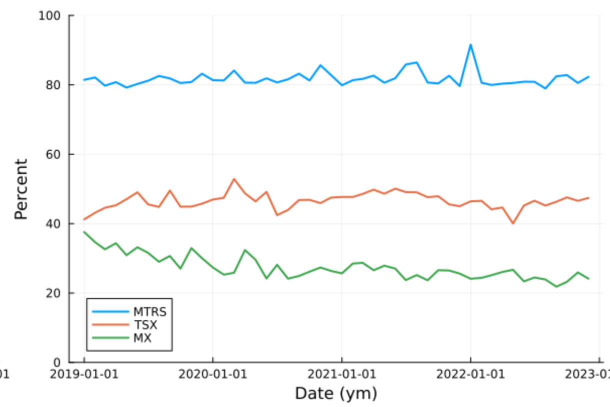

Notes: Appendix Figure A5a shows the fraction of monthly trade volume by primary dealers who are active in all markets. The graph implies that, in the bond market, essentially all dealers who are active in all markets are primary dealers. Roughly 68% of trade volume on TSX is executed by primary dealers who are active in all markets, which means that that 80%-68%=12% of trade volume is executed by dealers who are active across markets but are not primary dealers. Figure A5b excludes trades for client accounts.

Appendix Figure A6: Fraction of securities traded by each dealer out of all securities on TSX, MX, MTRS

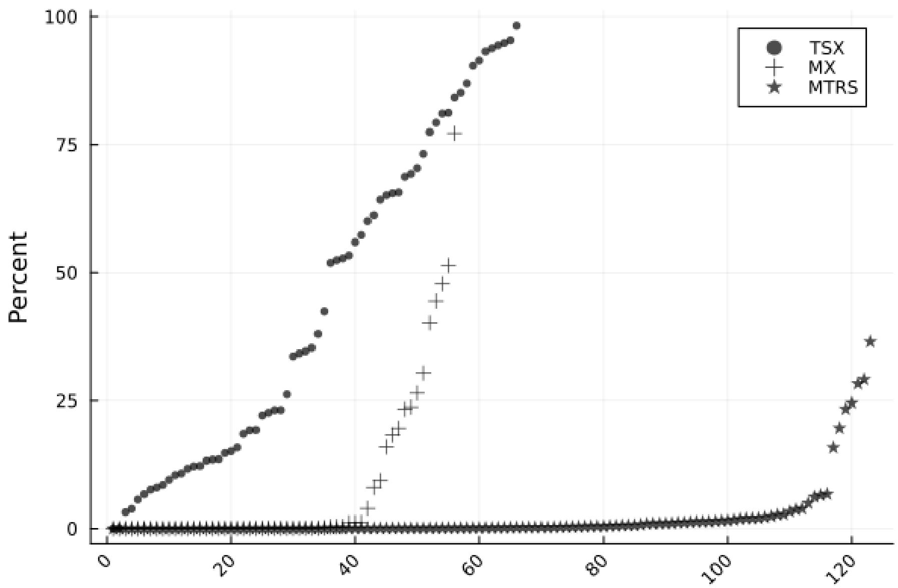

Notes: Appendix Figure A6 shows the fraction of symbols (x-axis) that each dealer trades (at the parent-level, on the y-axis) out of all traded symbols for each market. Since we sort by size within each market, the y-axis doesn't represent dealer-IDs that are common across markets. This is an alternative way of showing product differentiation: it is highest in the fixed-income market, and lowest on the stock exchange. The derivative exchange is in between.

Appendix Figure A7: Dealer presence across asset-types (defined by the symbol suffix)

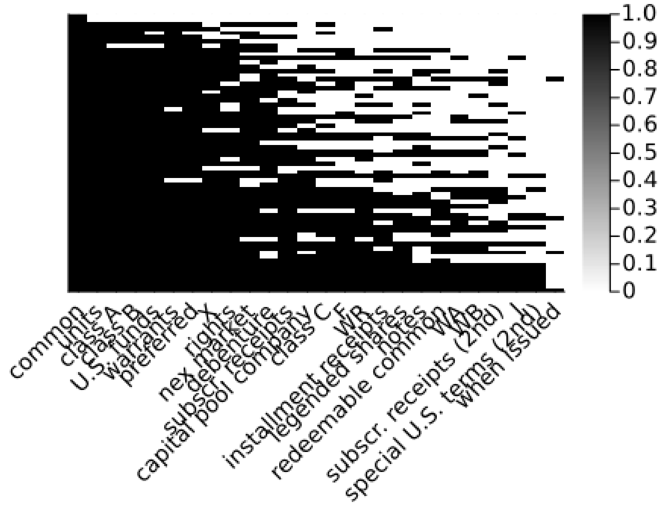

Notes: Appendix Figure A7 shows whether each of the dealers is active (i.e., trades at least ones) in white, versus in-active in white for each asset-class, defined according to the symbol-suffix within the stock markets at the parent-level. Suffices are explained in Appendix Table A8.

Appendix Figure A8: Dealer presence and market shares across all derivative products

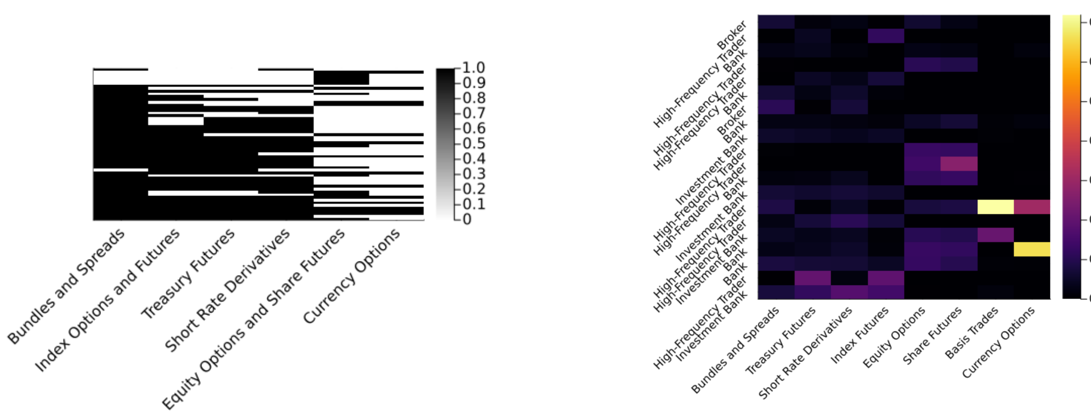

Notes: The RHS of Appendix Figure A8 is the analogue to Figure 4c but includes the small product "Currency Options". On the RHS we show whether dealers trade each product at least once in black. On the LHS we see the average annual product market shares of the largest dealers (on the LHS) for each market. In all figures, each row represents a dealer, sorted by total trade volume in the respective market. Dealers with the highest overall trade volume appear at the bottom, while those with the lowest appear at the top.

Appendix Figure A9: Dealer presence and market shares across products in the stock and derivative market (excluding market-maker accounts)

(a) Stock market

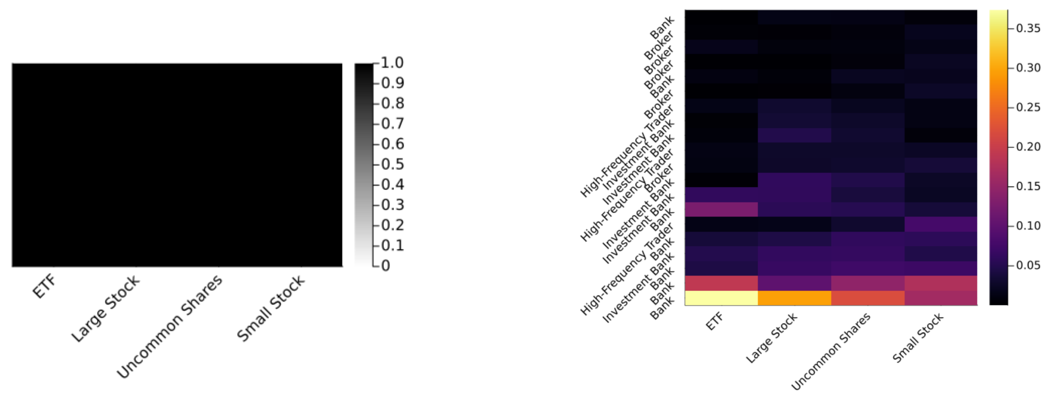

(b) Derivatives market

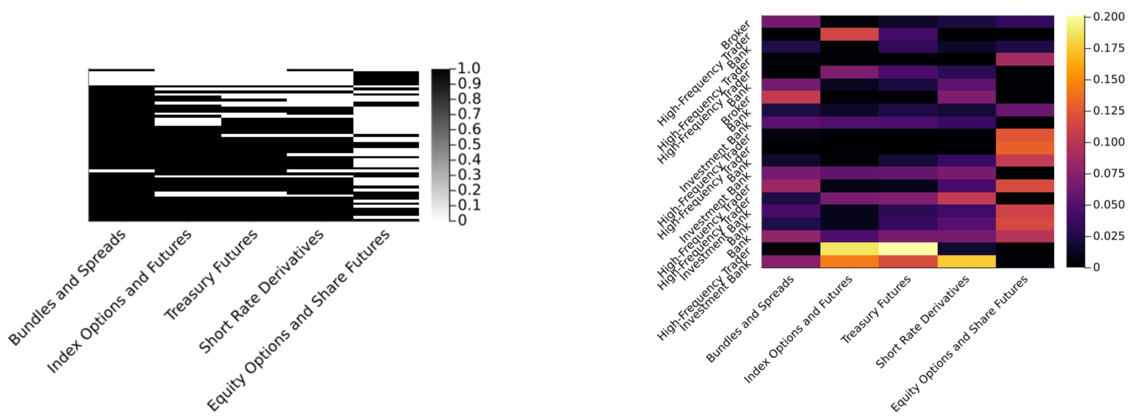

Notes: Appendix Figure A9 is similar to Figure 4 but excludes trades for market-marker accounts. On the RHS we show whether dealers trade each product at least once in black. On the LHS we see the average annual product market shares of the largest dealers (on the LHS) for each market. In all figures, each row represents a dealer, sorted by total trade volume in the respective market. Dealers with the highest overall trade volume appear at the bottom, while those with the lowest appear at the top.

Appendix Figure A10: Dealer presence across products of dealers present in all markets

(a) MTRS

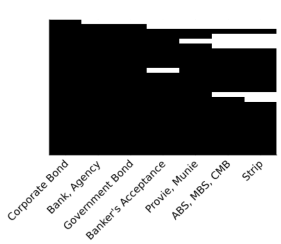

(b) TSX

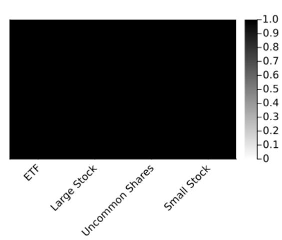

(c) MX

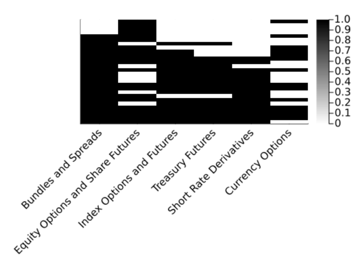

Notes: Appendix Figure A10 is the analogue to Figure 4, but includes the small product “Currency Options" for MX. It shows whether dealers who trade in all markets trade a product at least once in black, versus not in white within the bond (MTRS), stock (TSX), and derivatives market (MX).

Appendix Figure A11: Within-and cross-market segmentation indices

(a) Indices of all dealers

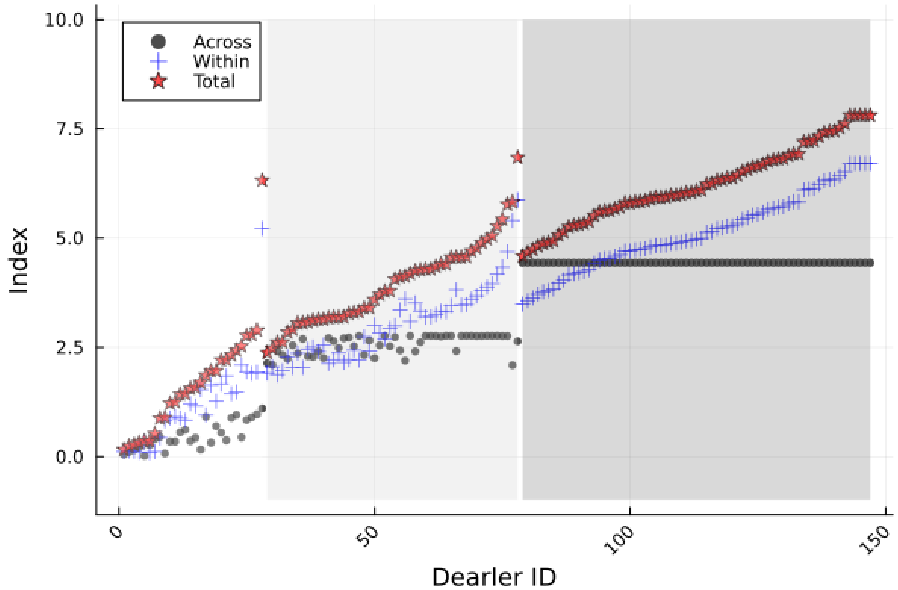

(b) Indices of largest dealers

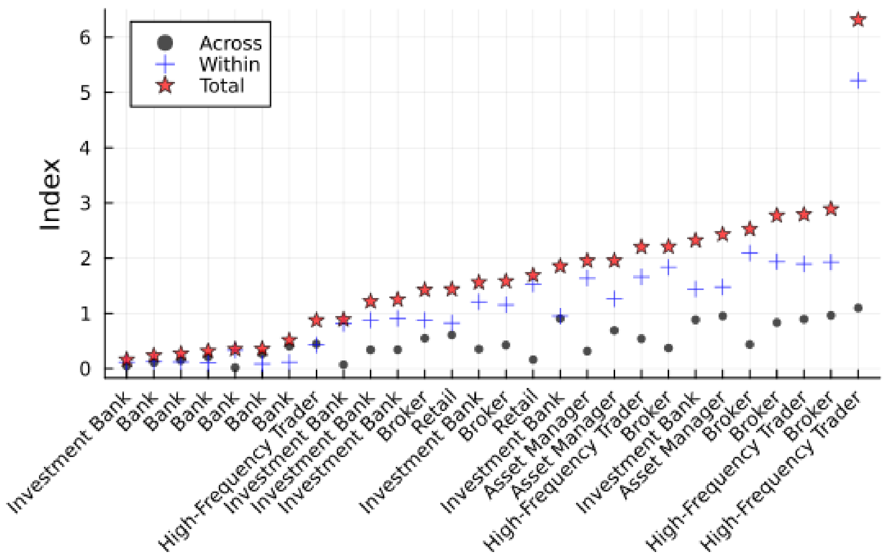

Notes: Appendix Figure A11a shows the within-market, across-market, and total adjusted Theil index for each dealer ID (at the parent-level). In the white area are dealer's who are active in all markets, in the light gray area are dealers who are active in only two markets, and in the darker gray shaded area are dealers who are active in only one market. The punishment term is 5; the maximal across market index is 4.431, and the maximal within market index is 8.82 for the stock market, 10.1234 for the bond market, and 10.6133 for the derivatives exchange. Figure A11b zooms in on the dealers who are active in all three markets.

Appendix Figure A12: Margins

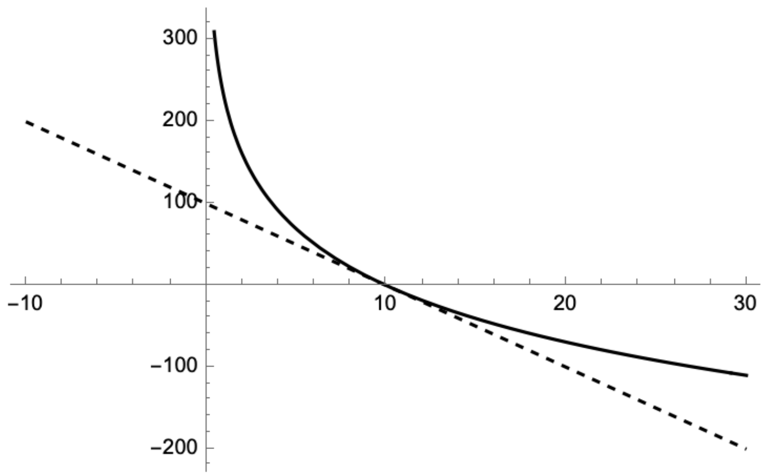

Appendix Figure A12 shows the margin (3) for an average price of 10 in black, and the linear approximation in dashed lines.

Appendix Figure A13: Margin distribution in each market

(a) Bonds (b) Stocks (c) Derivatives

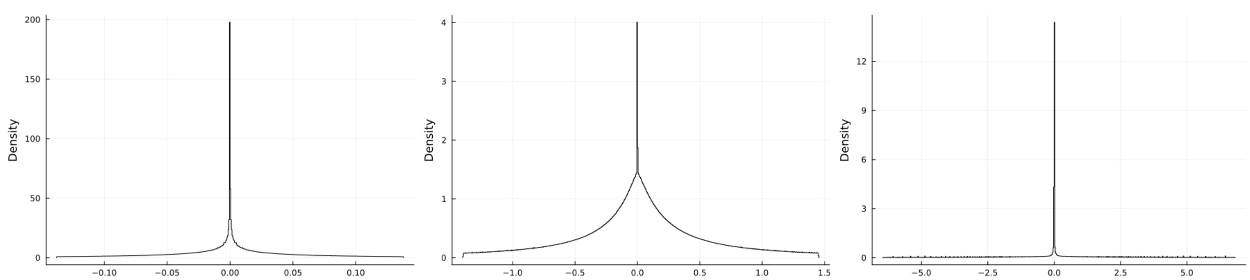

Appendix Figure A13 shows the distribution of our margin measure (3), which approximates how much less (more) a trader paid compared to the average price for a security in a day when buying (selling) for for each market. We exclude outliers, which are outside of the interquartile range. The median (average) margin is 0% for bonds, 0.005% for stocks, and 0% (0.18%) for derivatives. The standard deviation of margins is 1.59 for bonds, 2.03 for stocks, and 12.50 for derivatives. In comparison, the median trade price is C$100.29 for bonds, C$12.27 for stocks and C$1.1 for derivatives. The standard derivation in prices is 12.11 for bonds, 87.18 for stocks, and 157.35 for derivatives (where most of the variation is coming from the cross section of derivative contracts).

Appendix Figure A14: Price and margin distribution for equity products

(a) Small Stock

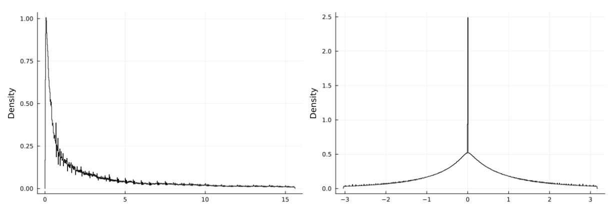

(b) Large Stocks

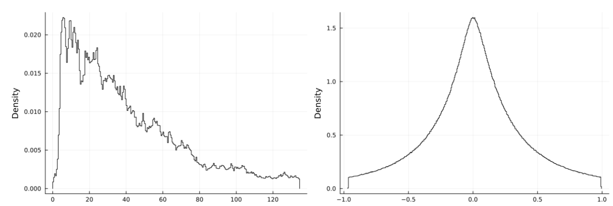

(c) Uncommon shares

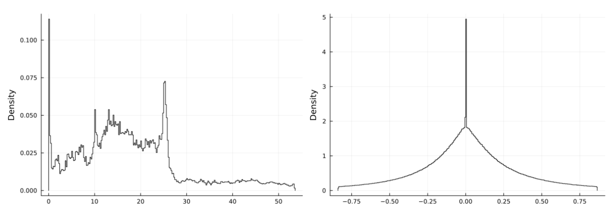

(d) ETFs

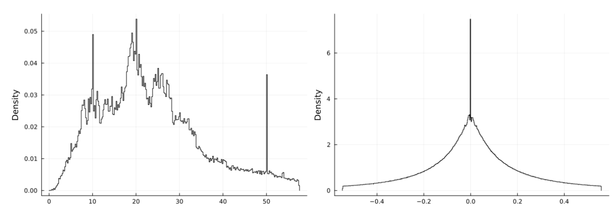

(e) Other or Missing

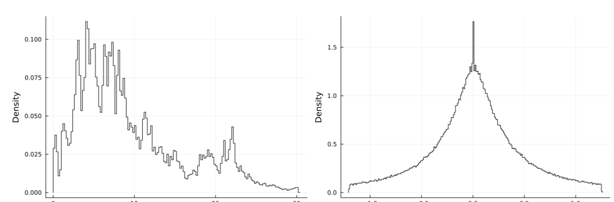

Notes: Appendix Figure A14 shows density histograms of prices for each product on TMX, excluding observations outside of the inter-quartile range.

Appendix Figure A15: Price and margin distribution for bonds

(a) Government Debt

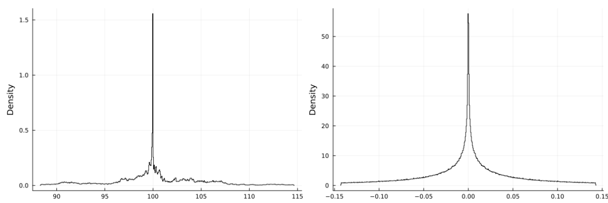

(b) Provincial and Municipal

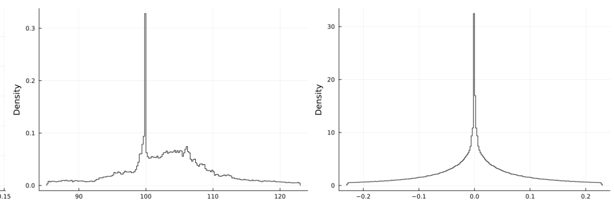

(c) Bankers' Acceptance

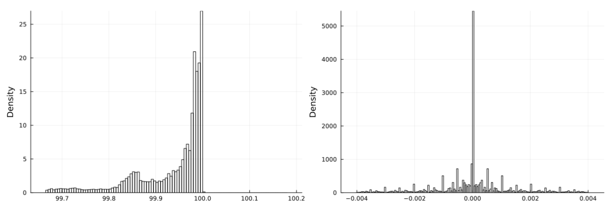

(d) Bank/Agency bonds

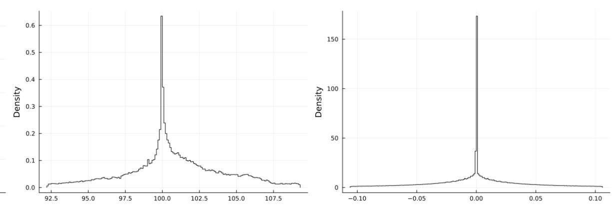

(e) Corporate Debt

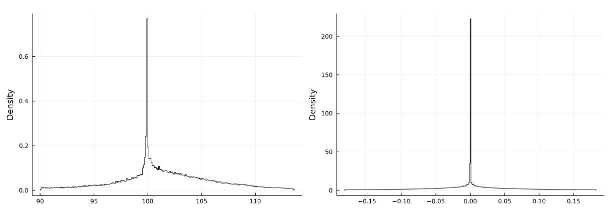

(f) ABS/MBS/CMB

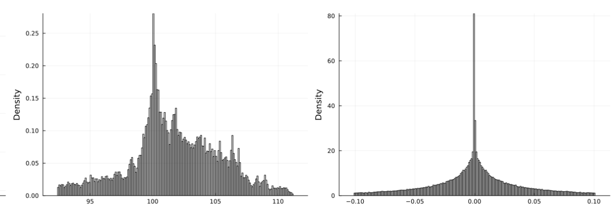

(g) Strips

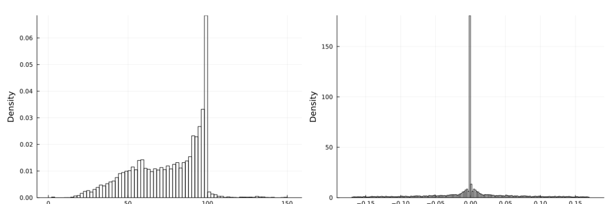

Notes: Appendix Figure A15 shows density histograms of prices for each product on the fixed-income market, excluding observations outside of the inter-quartile range.

Appendix Figure A16: Price and margin distribution for derivatives

(a) Treasury futures

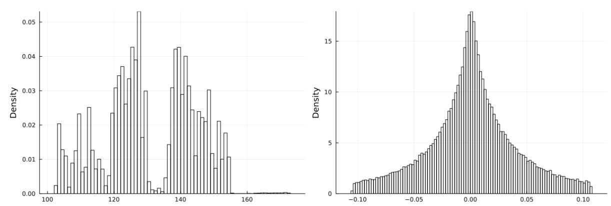

(b) Short Rate Derivatives

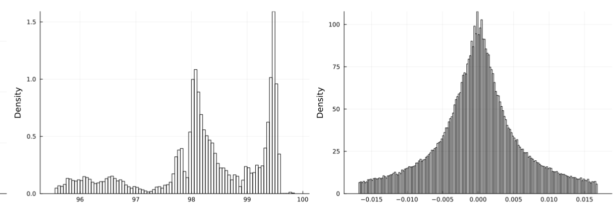

(c) Index Options and Futures

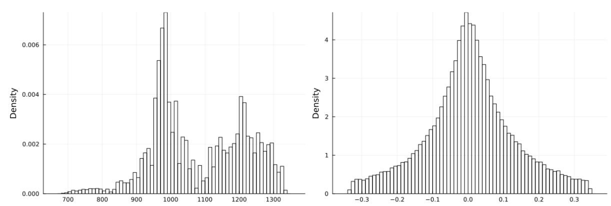

(d) Bundles and Spreads

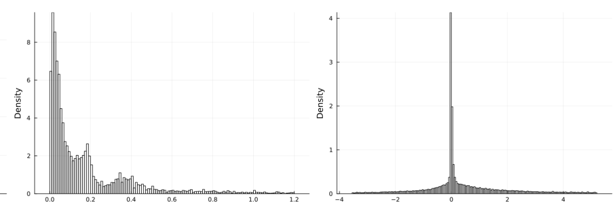

(e) Equity Options and Share Futures

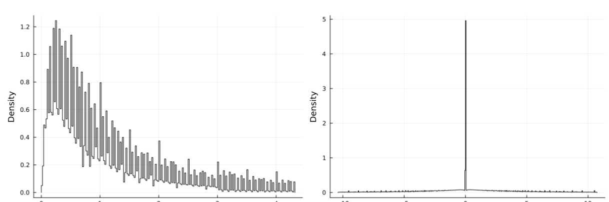

(f) Currency Options

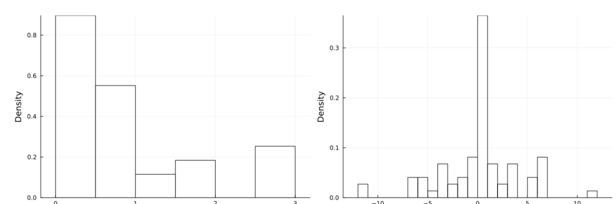

Notes: Appendix Figure A16 shows density histograms of prices for each product on MX, excluding observations outside of the inter-quartile range.

Appendix Figure A17: Dealer coefficients for dealers who exclusively trade in one market

(a) Stock market (b) Bond market (c) Derivative market

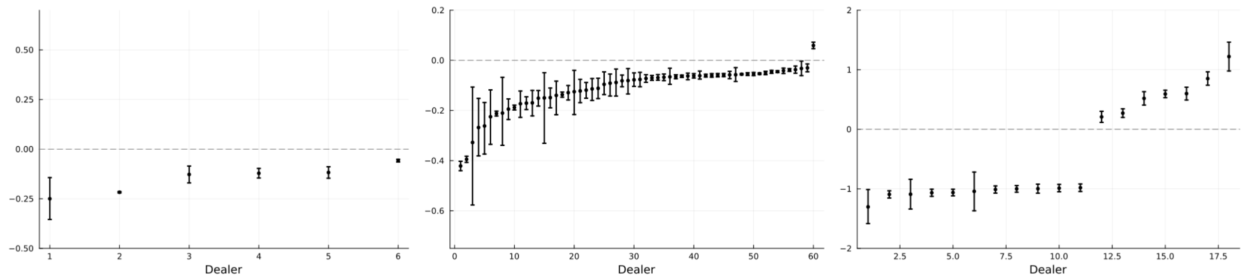

Notes: Appendix Figures A17a shows the dealer coefficients and 95% confidence intervals, which are obtained via WCR bootstrapping, when regressing margins (3) of a trade on indicator variables for each dealer that is only active on stock exchanges, and not on the bond market, (at the LEI-level) in addition to control variables (trade-size, the account-type, security-week and day fixed effects). Figure A17b shows the analogue for the fixed-income market, where we replace the account-type with a variable that indicates the type of trade. Figure A17c shows the analogue for the derivatives market. In all graphs, we exclude dealer coefficients that aren't significantly different from zero at a significance level of 5% according to bootstrapped and conventional inference to be conservative. We sort coefficients from small to large. Therefore, the x-axis are not comparable across markets, since they don't reflect the dealer's IDs.

Appendix Figure A18: Robustness—Dealer coefficients that are statistically different from zero at 5% significance level (conventional inference)

(a) Stocks (b) Bonds (c) Bonds

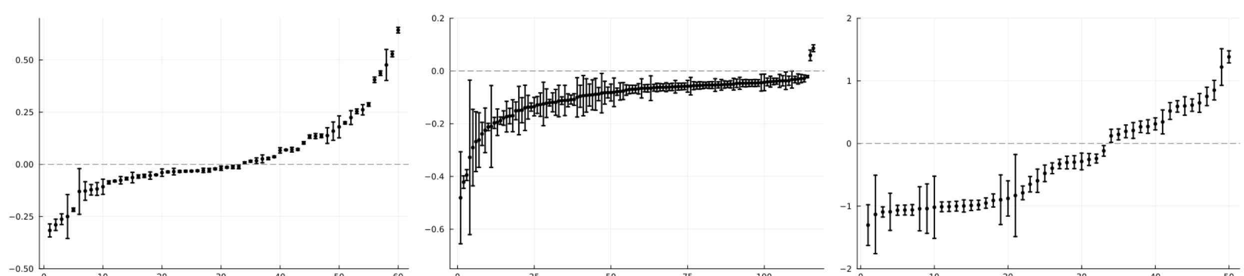

Notes: Appendix Figure A18 is the analogue to Figure 6 but with confidence intervals that are computed in the conventional way (without bootstrapping). For bonds, where clusters are most uneven in size, conventionally computed standard errors and confidence intervals differs slightly from bootstrapped confidence intervals—confirming expectations. Figure A18a shows the dealer coefficients when regressing margins (3) of a trade on indicator variables for each dealer active on the stock exchanges in addition to control variables (trade-size, the account-type, security-week and day fixed effects). Figure A18c shows the analogue for the derivatives exchange. Figure A18b shows the analogue for the fixed-income market, where we replace the account-type with a variable that indicates the type of trade (dealer-dealer, dealer-client, dealer-broker). In both graphs we exclude dealer coefficients that aren't significantly different from zero at a significance level of 5% according to bootstrapped and conventional inference. We sort coefficients from small to large. Therefore, the x-axis are not comparable across markets, since they don't reflect the dealer's IDs.

Appendix Figure A19: Robustness—Dealer coefficients that are statistically different from zero at 5% significance level (parent-level)

(a) Stocks (b) Bonds (c) Bonds

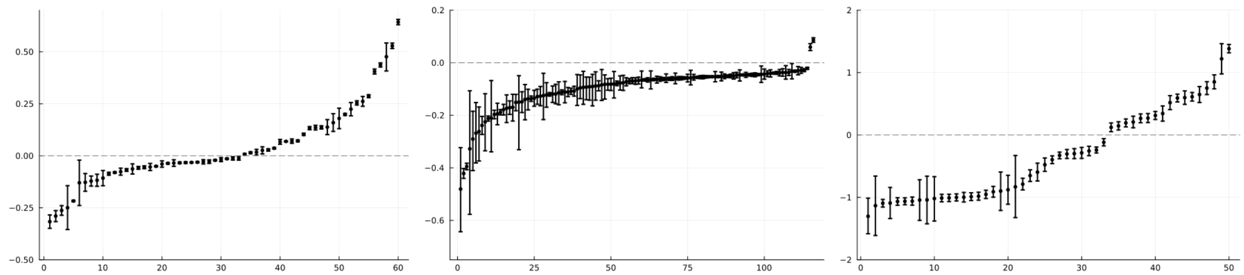

Notes: Appendix Figure A19 is the analogue to Figure 6 but aggregating dealer LEIs to the parent-level. Figure A19a shows the dealer coefficients when regressing margins (3) of a trade on indicator variables for each dealer active on the stock exchanges in addition to control variables (trade-size, the account-type, security-week and day fixed effects). Figure A19c shows the analogue for the derivatives exchange. Figure A19b shows the analogue for the fixed-income market, where we replace the account-type with a variable that indicates the type of trade (primary dealer/broker with non-primary dealer/non-broker, on-primary dealer/non-broker with non-primary dealer/non-broker, or primary dealer/broker with primary dealer/broker). In both graphs we exclude dealer coefficients that aren't significantly different from zero at a significance level of 5%. We sort coefficients from small to large. Therefore, the x-axis are not comparable across markets, since they don't reflect the dealer's IDs.

Appendix Figure A20: Robustness—Cross-market correlation between dealer coefficients of dealers active in all markets (conventional inference)

(a) Bonds-stocks (b) Bonds-derivatives (c) Stocks-derivatives

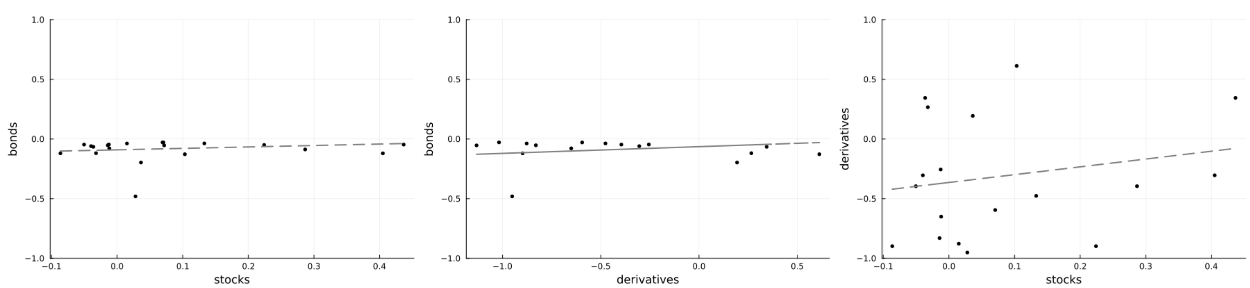

Notes: Appendix Figure A20a is analogous to Figure 7a but with confidence intervals computed in the conventional way without bootstrapping. It shows the within-dealer correlation of coefficients in the bond (y-axis) versus stock market (x-axis), (b) and (c) show the correlation for the other two market pairs. We exclude dealer coefficients that aren't significantly different

Appendix Figure A21: Robustness—Cross-market correlation between dealer coefficients of dealers active in all markets (parent-level)

(a) Bonds-stocks (b) Bonds-derivatives (c) Stocks-derivatives

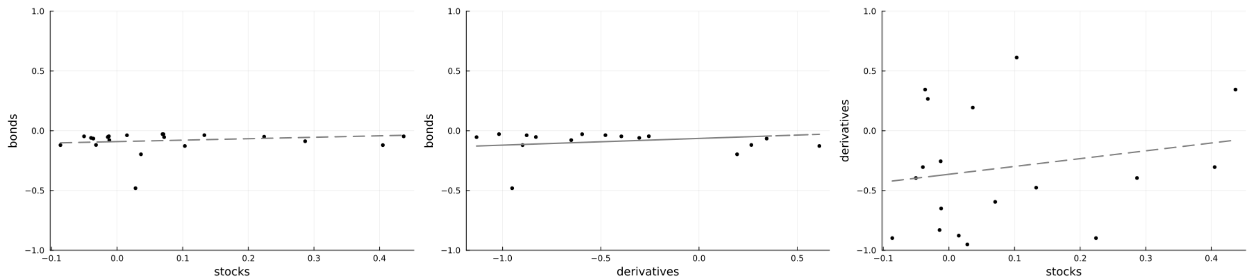

Notes: Appendix Figure A21 is analogous to Figure 7a but uses dealer LEIs at the parent-level. It shows the within-dealer correlation of coefficients in the bond (y-axis) versus stock market (x-axis), (b) and (c) show the correlation for the other two market pairs. We exclude dealer coefficients that aren't significantly different from zero at a significance level of 5%.

Appendix Figure A22: Dealer specialization across products within a market — 2022

(a) Stocks (b) Bonds (c) Derivatives

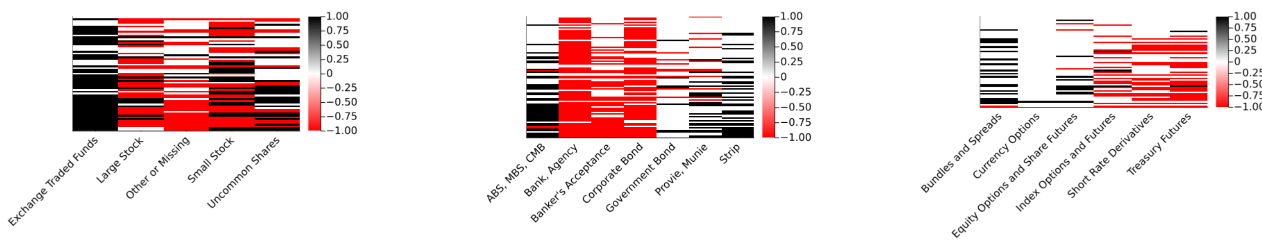

Notes: Appendix Figure A22a visualizes the dealer-product coefficients, β_{jp}, from regression (7)—which regresses trade margins (3) on indicator variables for each dealer-product combination plus control variables (trade-size, the account-type, security-year-week and day fixed effects)—using data from 2022. A row in each heatmap correspond to a dealer j, a column corresponds to a product p. When the corresponding β_{jp} is positive and statistically significant from zero at a 5% significance level, a p-j cell is black; if it is negative it is red; and empty if the coefficient is not statistically different from zero. Figure A22b and A22c show the analogue for bonds and derivatives. For bonds, the baseline β_{pj} coefficient is a large primary dealer trading government bonds; for stocks it's that bank trading large stocks, and for derivative it's that bank trading Treasury futures. In all graphs, dealers are sorted according to their trade-volume, with the dealer trading the most in the given market being at the bottom, and the dealer trading the least at the top. Standard errors are clustered at the daily-level.

Appendix Figure A23: Robustness: Dealer specialization across products within a market (parent-level) — 2022

(a) Stocks (b) Bonds (c) Derivatives

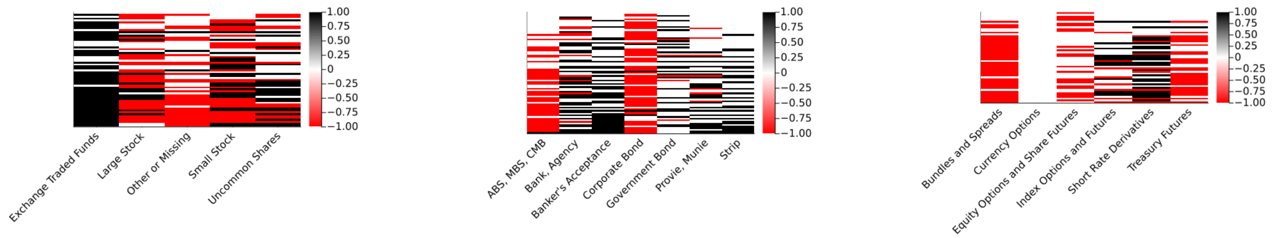

Notes: Appendix Figure A23 is analogous to Appendix Figure A22 but when aggregating dealers to the parent-level. Appendix Figure A23a visualizes the dealer-product coefficients, β_{jp}, from regression (7)—which regresses trade margins (3) on indicator variables for each dealer-product combination plus control variables (trade-size, the account-type, security-week and day fixed effects)—using data from 2022. Each row in a heatmap correspond to a dealer j. A column corresponds to a product p. When the corresponding β_{jp} is positive and statistically significant from zero at a 5% significance level, a p-j cell is black; if it is negative it is red; and empty if the coefficient is not statistically different from zero. Appendix Figures A23b and A23c show the analogue for bonds and derivatives. For bonds, the baseline β_{pj} coefficient is a large primary dealer trading government bonds; for stocks it's that bank trading large stocks, and for derivative it's that bank trading Treasury futures. In all graphs, dealers are sorted according to their trade-volume, with the dealer trading the most in the given market being at the bottom, and the dealer trading the least at the top. Standard errors are clustered at the daily-level.

Appendix Figure A24: Dealer specialization across products within a market — 2019, 2020, 2021, 2022

(a) Stocks (b) Bonds (c) Derivatives

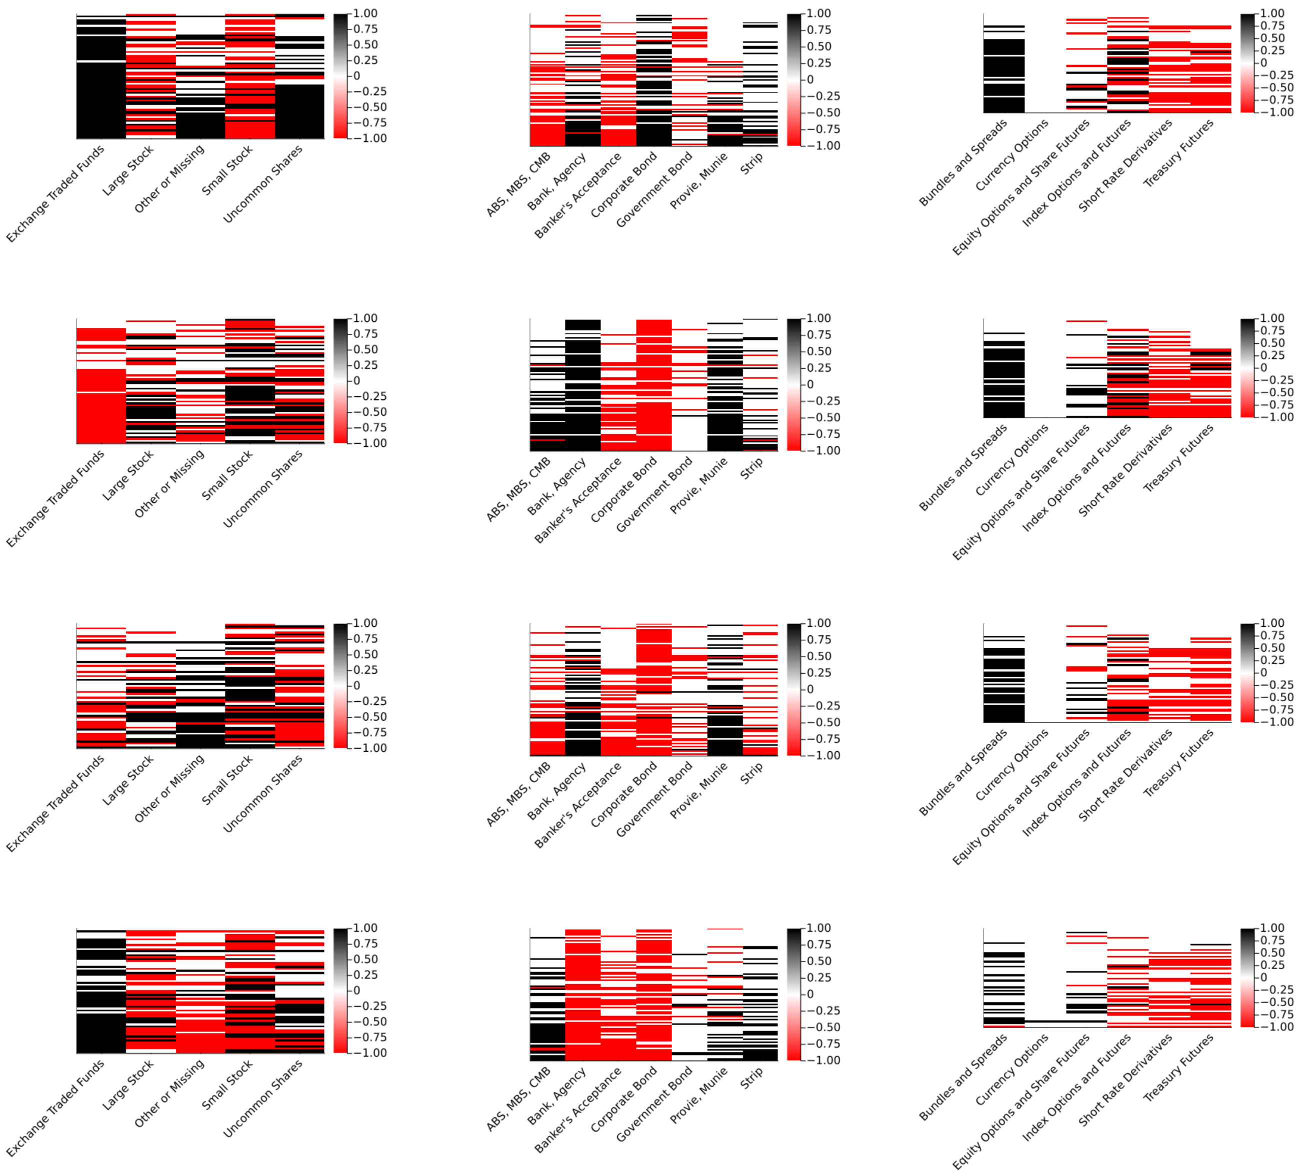

Notes: Appendix Figure A24 is analogous to Figure A22, but for the other years in our sample. We note that dealer coefficients vary across years. However, the main takeaway that no dealer outperforms across products is robust for all years. Appendix Figures A24a visualizes the dealer-product coefficients, β_{jp}, from regression (7)—which regresses trade margins (3) on indicator variables for each dealer-product combination plus control variables (trade-size, the account-type, security-week and day fixed effects) using data from 2019, 2020, and 2021, respectively. Each row in a heatmap correspond to a dealer j. A column corresponds to a product p. When the corresponding β_{jp} is positive and statistically significant from zero, a p-j cell is black; if it is negative it is red; and empty if the coefficient is not statistically different from zero. Appendix Figures A24b and A24c show the analogue for bonds and derivatives. For bonds, the baseline β_{pj} coefficient is a large bank trading government bonds; for stocks it's that bank trading large stocks, and for derivative it's that bank trading Treasury futures. In all graphs, dealers are sorted according to their trade-volume, with the dealer trading the most in the given market being at the bottom, and the dealer trading the least at the top.