Specialization in Financial Markets *

Milena Wittwer[⋆], and Andreas Uthemann[+]

July 7, 2025

Abstract

Financial markets are fragmented, with different asset classes trading in separate markets. Yet most empirical asset pricing research pays little attention to this type of fragmentation. As a result, we know relatively little about how financial intermediaries operate across markets or how to design policies that account for cross-market dynamics. We construct a unique dataset that traces trading by registered broker-dealers and exchange members—dealers—across all major Canadian markets for bonds, stocks, and derivatives. We show that dealers concentrate their activity in specific markets and asset types, with stronger specialization across markets. We provide suggestive evidence that this specialization reflects a mix of firm organization, market frictions, and client demand, and show that more specialized dealers consistently obtain better prices, indicating a trading advantage.

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 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. In addition, we thank Jason Allen, Vincent Bogousslavsky, Mirela Sandulescu, Or Shachar, 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 are grateful for comments we received at the Stern Market Microstructure Conference, the Women in Market Microstructure Conference. Correspondence to: [§]Milena Wittwer (Columbia): mw3941@columbia.edu, and [+]Andreas Uthemann (Bank of Canada): authemann@bankofcanada.ca

1 Introduction

Financial markets are fragmented—different asset classes trade in distinct venues (Malamud and Rostek (2017); Weill (2020); Chen and Duffie (2021); Budish et al. (2024)). Yet the empirical asset pricing literature pays relatively little attention to the implications of fragmentation across asset classes, focusing instead on how intermediary frictions—such as balance sheet and funding constraints—affect asset prices (e.g., Pasquariello (2014); Adrian et al. (2014); Du et al. (2018); He et al. (2017); Siriwardane et al. (2022)). This is despite influential work in intermediary asset pricing showing that other frictions—distinct from market fragmentation—can constrain arbitrage and shape asset prices (e.g., Shleifer and Vishny (1997); Gromb and Vayanos (2002); Brunnermeier and Pedersen (2009); He and Krishnamurthy (2013); Brunnermeier and Sannikov (2014)). As a result, we lack a clear understanding of how intermediaries operate across markets, or how to design effective policies that account for cross-market dynamics.

We analyze the extent to which intermediaries specialize across financial markets rather than span them all, potential reasons for specialization, and how it impacts transaction prices. Using a unique trade-level dataset covering all major Canadian financial markets, we show that registered broker-dealers and exchange members (dealers) specialize in specific asset classes (markets) and specific types of assets within markets (products), with market specialization being more pronounced.[1] This pattern is not driven by any single factor but likely reflects an interplay of firm structure, market frictions, and clientele. Consistent with the idea that specialization is integral to dealers' business models, more specialized dealers tend to trade at better prices, suggesting a positive link between specialization and trading performance—and a monetary incentive to specialize.

Our dataset covers all trades executed on Canada's fixed-income market and all exchanges owned by the TMX Group 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 clas-

[1]We use a broad definition of dealers that includes both traditional institutions, such as primary dealers, and non-traditional ones, like high-frequency traders. While some might label these firms as 'market makers,' we avoid that term, as many also execute client trades and engage in other intermediation activities. Moreover, only a subset are formally designated as market makers with liquidity obligations. For readers interested in specific types of institutions, we show many of our results by dealer type.

sify securities into products, for instance, corporate bonds, common share stocks, funds, and Treasury futures. Additionally, we manually assign dealers to their parent institutions, i.e., the LEIs of their holding companies, and categorize them by type, such as primary dealers or high-frequency traders (which includes hedge funds).

Using these data, we document the existence and extent of dealer specialization—that is, the uneven allocation of trading activity across markets and products. This adds to a relatively small literature that emphasizes dealer specialization, typically within asset classes (e.g., Siriwardane (2019); Jotikasthira et al. (2024)), rather than across them.[2] In particular, our trade-level analysis complements Siriwardane et al. (2022), who analyze arbitrage spreads across bonds, stocks, and derivatives and infer specialization patterns.

To measure the degree of specialization, we construct market and product specialization scores for each dealer, ranging from zero (no participation) to one (exclusive trading in a single market or product segment), and use them to build a unified specialization index. This index combines information from both scores and decomposes overall specialization into cross-market and within-market components. Specifically, we adapt the Theil (1967) index —commonly used to study inequality in socio-economic contexts (e.g., Anand and Segal (2015)) but rarely applied in financial contexts—to account for the fact that not all dealers participate in every market segment.

For most dealers, cross-market specialization exceeds within-market product specialization, implying that segmentation is more pronounced across asset classes than within them. Standard portfolio theory does not predict such behavior, as organizing activity along asset class boundaries limits cross-asset hedging and arbitrage. Our unique cross-market data reveal that dealers nonetheless organize their trading and expertise around these market divisions. The finding supports theories that emphasize market segmentation at the level of intermediaries or arbitrageurs—such as Greenwood et al. (2018) and Gourinchas et al. (2025)—rather than at the investor level, which remains the more common assumption (e.g., Vayanos and Vila (2021)).

To understand why dealers specialize, we examine a non-exhaustive set of potential drivers, asking whether specialization in the cross-section reflects a single dominant factor or a combination. While much of the existing literature focuses on individual mechanisms within specific markets or institutions, we take a more aggregated perspective—closer in spirit to Siriwardane

[2]Several influential papers examine intermediary frictions across asset classes using return data in the cross-section (e.g., Adrian et al. (2014); He et al. (2017); Koijen et al. (2022)). While these studies provide evidence that intermediary balance sheet constraints help price a broad range of assets, they do not focus on institutional segmentation or fragmentation across trading venues, nor do they analyze how intermediaries allocate activity across markets.

et al. (2022), who emphasize funding and balance sheet segmentation. To complement their work, we investigate alternative mechanisms that can be inferred from trade-level data.

One possible driver is information advantages. Informed dealers may intermediate more trades and obtain better execution within a market (e.g., Hendershott and Madhavan (2015); Collin-Dufresne et al. (2020); Carlin et al. (2013)). However, the stability of specialization over time offers little support for time-varying information as a key driver.

A second potential driver of specialization is dealer size, which may matter due to economies of scale or network advantages. Using overall market share as a proxy for size, we find no systematic relationship between size and specialization. This contrasts with evidence from specific OTC markets—such as Jotikasthira et al. (2024), who document a positive link between market share and geographic concentration in U.S. corporate bond trading. Our findings suggest that this relationship breaks down outside illiquid OTC markets, where network connections that scale in size are more central. Instead, they are consistent with the view that smaller firms can adopt distinct business models to successfully specialize.

A third possible driver is in-house segmentation—that is, the division of trading activity across distinct legal entities within the same parent organization. This form of internal segmentation has been shown to affect risk-taking and trading behavior in financial institutions (e.g., Siriwardane (2019); Lu and Wallen (2024)). Using LEIs, we compare specialization at the entity and parent (holding company) level. While Theil indices decline when aggregating across affiliated entities, much of the heterogeneity remains. This suggests that in-house segmentation contributes to observed specialization, but does not fully account for it. Importantly, because we do not observe trading desks within entities, our estimates capture only formal legal boundaries and likely understates the degree of internal segmentation (when accounting for trading desk segmentation within the same legal entity).

Our next conjecture is that market frictions—such as search costs and entry barriers—explain specialization. If search frictions were the single driver, we would expect greater product specialization in the OTC bond market than in centralized markets. Yet, we observe the highest product specialization in derivatives, which trade on centralized exchanges with minimal search frictions. One possible explanation is that derivatives are more complex instruments, and existing research shows that asset complexity can influence both trading behavior and returns (e.g., Carlin et al. (2013); Eisfeldt et al. (2023)). Entry barriers may also contribute to segmentation; however, we find meaningful specialization even among some dealers active across all major markets—suggesting that, while market frictions might matter, they are not the sole explanation.

The last potential driver we consider is that dealers specialize in response to the preferred

habitats of their clients (e.g., Greenwood and Vissing-Jørgensen (2018); Khetan et al. (2023)). Supporting this view, we find that specialization varies systematically by dealer type, such as banks versus high-frequency trading firms, and, in the bond market, by the types of clients a dealer serves. For example, dealers trading with foreign central banks concentrate in Bankers' Acceptances—short-term instruments commonly used in monetary policy operations—while those interacting with hedge funds focus on liquid government bonds, consistent with their use as collateral or in hedging strategies.

Having established that specialization is widespread and driven by a combination of factors—including firm structure, market frictions, and clientele—, we turn to whether it is associated with systematic differences in transaction prices. We focus on relative prices across dealers, not on how specialization affects aggregate price levels. Theoretically, dealer specialization could influence transaction prices by improving inventory management, or shaping beliefs about fundamentals. However, in a frictionless and sufficiently competitive market, any systematic price effects from specialization at the individual-level would be arbitraged away.

We, therefore, begin by showing that none of the markets is sufficiently frictionless to prevent some dealers to consistently 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. We show that dealers systematically differ in the margins they obtain, across all markets, even after controlling for trade size, and security-time fixed effects.

We then investigate whether the successful dealers trade across products or markets or whether they specialize. 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 reinforcing the role of market specialization. Some dealers that concentrate their activity within a single market outperform those active across markets.

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 an 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, once we control for observable trade characteristics and a rich set of fixed effects.

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 specialization to full specialization increases margins by 6–27 basis points per trade. These price effects suggest that dealers face monetary incentives to specialize, supporting the view that specialization is an integral part of dealer business models.

Taken together, our findings underscore the importance of dealer specialization—both in terms of quantities and transaction prices. Specialization is not attributable to any single driver. Instead, the persistent heterogeneity we observe likely reflects a strategic response to a complex mix of institutional forces, including firm structure, market design, and client composition. This suggests that specialization is not merely a byproduct of frictions, but a defining feature of how intermediaries operate. Asset pricing models need not capture the full complexity of these forces, but incorporating some degree of segmentation across asset classes—-by intermediaries, not only investors—-may be essential for understanding cross-market dynamics, and how regulatory changes in one market may spill over into others.

From a policy perspective, our findings carry both cautionary and encouraging implications. On the one hand, since market specialization is not solely driven by market frictions, like market entry barriers, modest changes to market access or membership fees—such as the recently revised schedules for registered broker-dealers in the U.S. and Canada (FINRA (2024); CIRO (2024))—are unlikely to substantially alter dealer behavior unless they prompt broader institutional restructuring, which tends to occur only slowly. On the other hand, the absence of a systematic relationship between dealer size and specialization suggests that scale economies are not necessary for intermediation. This mitigates concerns about concentration or monopolistic dominance and points to a competitive landscape in which smaller, more specialized dealers can coexist with larger institutions.

There are several promising directions for future theoretical and empirical work. One direction 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. Such analysis would complement the empirical asset pricing literature (e.g., Gabaix et al. (2007); Pasquariello (2014); Adrian et al. (2014); Du et al. (2018); He et al. (2017); Siriwardane et al. (2022); Eisfeldt et al. (2023)), 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)).

A second direction for future research is to develop theoretical frameworks that more fully capture the drivers of dealer specialization. One idea is to extend market microstructure models to incorporate differences in market structure and product characteristics, moving beyond the typical focus on a single market or asset class (see Rostek and Yoon (2025)). With few exceptions—such as Dougast et al. (2022) and earlier work on derivatives and their underlying assets (e.g., Kumar and Seppi (1992))—existing models remain confined to narrow settings, limiting their ability to explain the cross-market patterns we document. Another path is to expand macroeconomic models that feature segmentation at the dealer-level (e.g., Greenwood et al. (2018, 2023); Gourinchas et al. (2025)), but currently emphasize exogenous limits to arbitrage—such as restricted market access—without explicitly modeling the origins of segmentation.

A third 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. With the exception of studies on options and their underlying securities (e.g., Figlewski (1989)), 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.[3]

[3]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 analyzing 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 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 (Provincial Bonds), and Bankers' Acceptance (which is a money market instrument that is issued by a business and guaranteed by a bank), bank or agency securities (money market instruments issued by banks or agencies), corporate bonds, and asset-backed securities (which includes mortgage-backed securities).

Equity market. Equities 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,

[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).

Table 1: Products in the fixed-income market

Product	Trade share
Government Bonds (and Bills)	63.71
Provincial Bonds	9.34
Bankers' Acceptance	8.82
Bank and Agency Securities	7.96
Corporate Bonds	6.15
Asset-Backed Securities	4.88

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. Appendix Table A1 describes each product category.

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]

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 securities, ranging from common stocks, and funds to more non-common shares. We group assets into six product categories, as detailed in Appendix Table A2: four types of common stocks based on company sector—natural resources; consumer and healthcare; technology and industry; and financials—along with two additional categories: non-common shares and funds (primarily ETFs).[7]

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

[5]These numbers are computed with data from CIRO, accessible here: https://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]Alternatively, we could group all common shares into a single category. However, as our analysis shows, common shares are more comparable to each other than to non-common shares or funds, so maintaining separate categories does not affect our main conclusions. We do so to ensure a consistent decomposition of market and product specialization—an approach that is cleanest when each market includes the same number of products, which we set to six.

Table 2: Stock market products

Product	Trade share
Stocks-Natural Resources	48.13
Stocks-Consumer & Healthcare	17.67
Stocks-Technology & Industry	13.30
Stocks-Financial Sector	7.16
Non-Common Shares	7.53
Funds	6.20

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.

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

We group the derivative products into six 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.

[8]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 jurisdiction: 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).

Table 3: Derivative products

Product	Trade share
Treasury Futures	37.48
Equity Options and Share Futures	27.91
Short Rate Derivatives	20.76
Bundles and Spreads	7.85
Index Options and Futures	6.52
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.

3 Data

We combine proprietary data on stock and derivative exchange trades from the TMX Group with fixed-income trading data from CIRO. To enrich these datasets, we hand-collect publicly available information on CIRO and exchange members, financial products, and market conditions.

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.[9] 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, but these trades are rare according to market experts.

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 (expressed in bond par-value) 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,

[9] 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.

masked identifiers are infrequent—roughly 5% of trades and 1% of trade volume.[10]

Since we are interested in how much CIRO dealers buy and sell, we stack buy- and sell-side trades, and remove sales or purchases by non-dealers. For example, in the raw data we observe CIRO dealer 1 buying from CIRO dealer 2, and CIRO dealer 1 selling to a client. The data we analyze lists CIRO dealer 1 and 2's transaction, separately, in addition to CIRO dealer 1's sale to their client.

Equity and derivatives market. We observe trade-level data for all exchanges that are owned by the TMX group between 2019 and 2022.[11] For each trade, we observe the time of the trade (up to milliseconds), the security (i.e., the TMX symbol), the amount, and the price. For stocks, the amount refers to the number of shares traded, while for derivatives, it reflects the number of contracts; the price corresponds to the share price in the case of stocks, and to the cost of entering the derivative position—for example, the option premium when the contract is an option—rather than the strike price. We also observe whether the trade is for the exchange member's own account, a client account, or whether the trade is conducted to fulfill market-making responsibilities. Although there are a few other types of accounts, these are negligible.

The most unique feature of the data is trading-firm IDs. In order to trace dealers across markets, we link the market-specific trading-firm IDs to the firm'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.

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, similar to the bond data, we stack buy- and sell-side trades to account for both sides of each transaction.

[10] 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.

[11] 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.

Table 4: Dealer types

Institution type	# of LEIs	# of Parents
Broker-Dealer	105	57
High-Frequency Trader	21	20
Primary Dealer	19	9
Retail-Facing Dealer	19	11
Asset Manager	12	16
Mutual Fund	6	1
Bank	4	28
Pension Fund and Insurance	3	2
Other	1	1

Notes: Table 4 shows the number of LEIs and parent-LEIs of each dealer type at the LEI and parent-level. Appendix Table A4 defines each type category we observe in our data. Note that the composition of trading firms changes slightly when moving from the LEI to the parent level, as some LEIs active in the bond market are not CIRO dealer members themselves, but their parent firms are.

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—which must also be CIRO members—in equity and derivatives markets. Throughout the paper, we refer to these firms as 'dealers.' This definition is intentionally broad and based on membership status rather than trading activity. We adopt it because membership is more exogenous than classifications based on observed trading behavior. Later in the analysis, we present our main results by dealer type, allowing readers to focus on specific dealer types—such as primary dealers—if desired.

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.

In addition, we classify all LEIs and their parents into institution types: Table 4 shows that broker-dealers constitute the largest category, representing financial entities primarily engaged in brokerage services. Following high-frequency traders, which are primarily proprietary trading firms, but also include few hedge funds and private equity firms. Then we have primary dealers, and retail-facing dealers. Those are firms that focus on serving retail investors; for example, Questrade Inc.—an online brokerage catering to retail investors.

Summary statistics. Appendix Table A5 provides 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 of dealers who actively trades on an average day is similar across markets, ranging from 54 in the derivatives market to 64 in the stock market.

4 Dealer specialization

We examine two types of specialization: market specialization—which may arise from differences in market structures or preferences for specific asset classes—and product specialization within a market—which could stem from variations in trading expertise or clients' preferred habitat for certain products.

4.1 Quantification

Our main approach to assessing the degree of market and product specialization, is to introduce specialization scores ranging from 0 to 1. We then aggregate information from these scores into a single index that can be decomposed into two components—specialization across markets and specialization across products within a market—to assess which one is more pronounced.

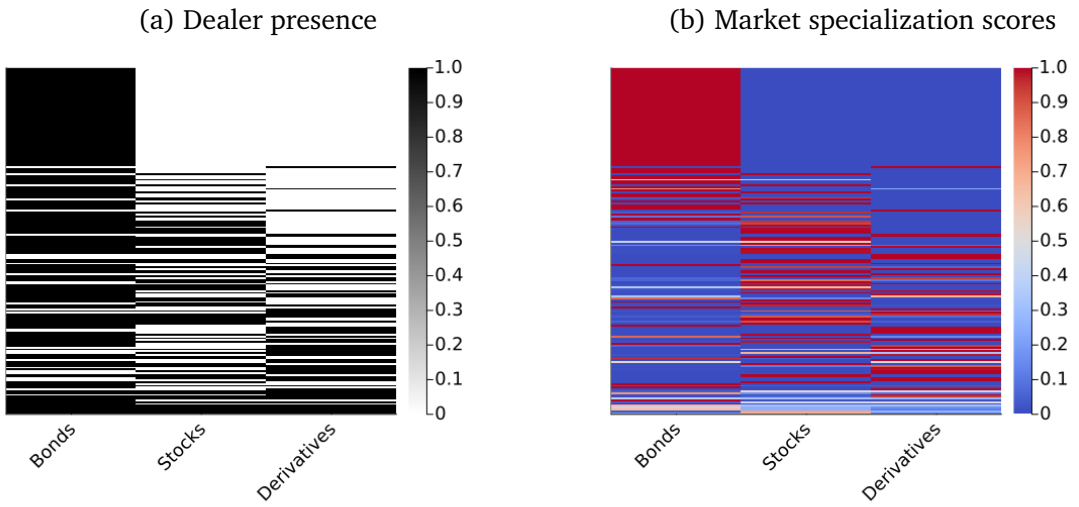
In Appendix B.1, we complement this approach by analyzing correlations in returns, dealer trade flows, and gross trade volumes across markets and products. This analysis provides additional insights—for instance, it helps identify whether dealers tend to go long in one product while short in another, suggesting potential hedging or strategic positioning across product pairs. In the broader picture, however, it reinforces the same qualitative takeaway: dealers specialize markets and, within markets, across products.

Throughout, we treat the three stock exchanges, TSX, TSXV, and Alpha, as one stock market, because 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. Supporting the idea that market entry barriers or market-specific preferences may limit universal market participation, Figure 1a highlights that not all dealer LEIs (y-axis) are active in all markets (x-axis). A black line indicates activity by a dealer in a given market. If all dealers were active in all markets, the chart would be fully black. Instead, the mix of black and white lines reveals gaps in participation.

Since our focus is on specialization rather than overall trading activity, we control for a dealer's total market size and introduce a market specialization score. This score ranges from 0 (the dealer does not trade in the market) to 1 (the dealer trades exclusively in that market). Formally, the score divides dealer j's market share—the fraction of trade volume the dealer

Figure 1: Market specialization



Notes: Figure 1a visualizes dealer presence. Each row represents a dealer, with black indicating that the dealer traded at least once in the respective market—bonds, stocks, and derivatives. Dealers are sorted by their average annual trade volume (computed as the sum across average annual market shares, $\sum_m s_{jm}$), with the most active dealers at the bottom. Figure 1b displays average annual market specialization scores (1) for each dealer. These scores ranges from 0 (no specialization) to 1 (full specialization). Appendix Figure A1 zooms in on the largest 20 dealers.

executes—in market m and year y, s_{yjm}, by the sum of the dealer's market shares in all markets:

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

Figure 1b shows market specialization scores for all dealers, averaged over the years. Many dealers specialize in specific markets.

Among the three markets, the derivatives market stands out as the most detached with 30% of trade volume being conducted by dealers who exclusively trade on the derivative exchange (see Appendix Table A7). This is because some dealers—particularly high-frequency traders—engage almost exclusively in derivatives. In total, approximately 30% of MX's trade volume is attributed to high-frequency traders when including trades for client accounts, rising to 67% when excluding them, as shown in Appendix Table A8.

Product specialization. Product specialization may be driven by various factors, including differences in trading expertise, or relationships to clients with preferred habitat.[12] To quantify

[12]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

product specialization, we define a product specialization score, analogously to our market specialization score:

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

where s_{yjmp} represents the fraction of dealer j's trade volume in product p within market m, relative to all other dealers.[13]

Figure 2 shows that many dealers only trade a subset of products. Conditional on market participation, Figure 3 tells us that product specialization scores are heterogeneous across dealers and products. In the bond market, government bond and corporate bonds feature most specialization, while equity options and share futures appear to have most dealer specialization in the derivative market. In the stock market, specialization is rather low, with most dealers trading all products and product specialization scores being more similar across dealers and products.

Market versus product specialization. To compare market and product specialization, we introduce an index based on Theil (1967) that integrates information from the specialization scores in equations (1) and (2) in a way that allows for direct comparison.

The Theil index, rooted in information theory, is commonly used to measure economic inequality. Its key advantage lies in its additive decomposability, which makes it particularly well-suited for comparing segmentation within and across groups (in our case, markets). This sets it apart from alternative measures or dispersion, such as variances, or market concentration, like the Herfindahl Index, which either depend on the number of observations (in a group) or lack this decomposability property.

To formally define the index, let there be M markets indexed by m (with $M = 3$ in our case), and P products in each market.[14] Let $s_{yjmp} \in [0,1]$ denote dealer j's market share in segment mp in year y, and define the market-level and overall averages in a year as $\bar{s}_{yjm} = \frac{1}{P}\sum_p s_{yjmp}$ and $\bar{s}_{yj} = \frac{1}{M}\sum_m \frac{1}{P}\sum_p s_{yjmp}$.

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 2 compared to Figure 2, which we explain below).

[13] In Appendix Figure A2 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.

[14] It is straightforward to generalize the framework to incorporate different numbers of products across markets. However, having markets with different numbers of products introduces a mechanical asymmetry in the cross-product index, which we explain in Appendix C .

Figure 2: Dealer presence across products in each market

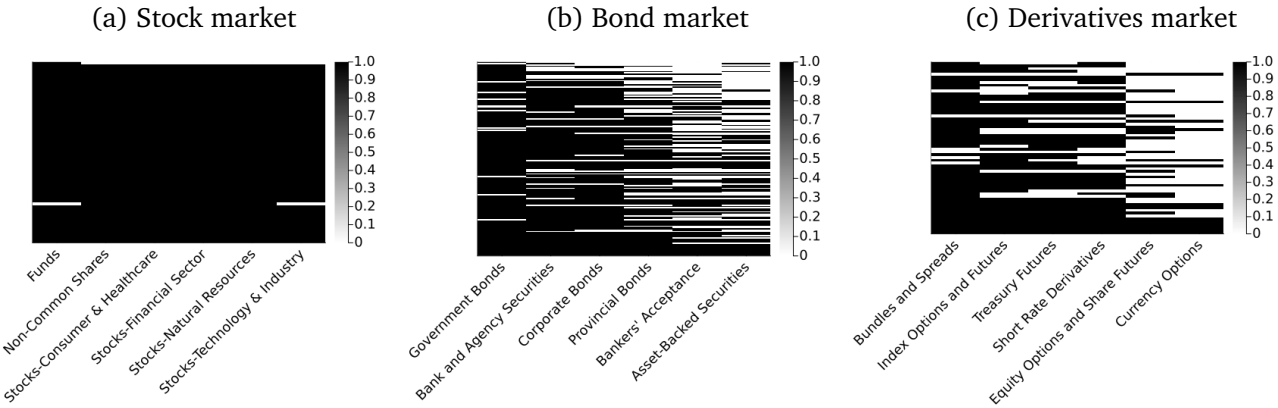
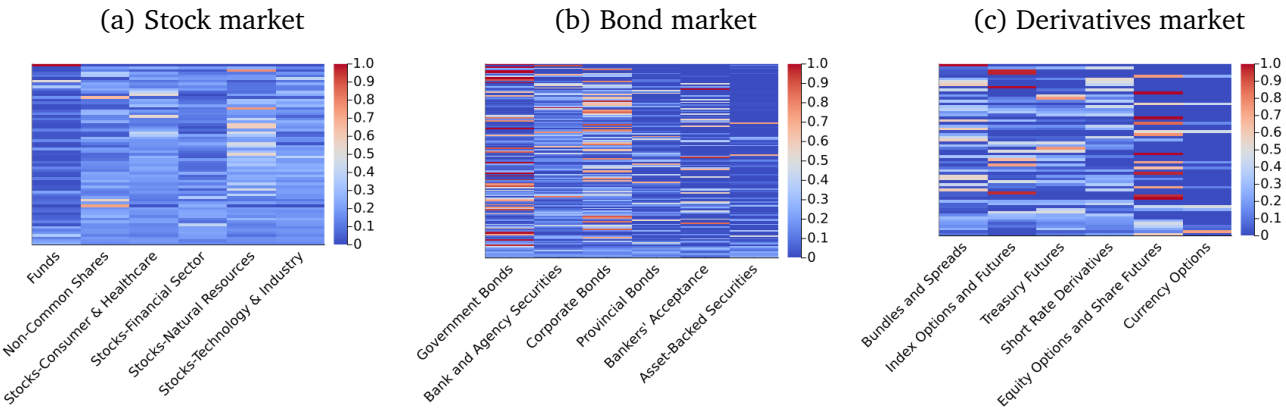


Figure 3: Dealer product specialization scores across products in each market



Notes: Figure 2 is similar to Figure 1a, but replaces markets with products within each market, the stock market (a), bond market (b), and derivatives market (c). We show whether dealers trade each product at least once in black. In all figures, each row represents a dealer. Dealers are sorted by their average annual trade volume (computed as the sum across average annual market shares, $\sum_m s_{jm}$, where $s_{jm} = \frac{1}{Y}\sum_y s_{yjm}$), with the most active dealers at the bottom. Figure 3 displays average annual product specialization scores (2) for each dealer. These scores ranges from 0 (no specialization) to 1 (full specialization). Appendix Figures A3 and A4 show the product market shares and product specialization scores for the top 20 largest dealers, respectively.

In this setting, the original Theil index measures the dispersion of dealer j's market shares in year y across product–market segments:

$$\mathcal{T}_{yj} = \frac{1}{M_{yj}} \sum_{m} \frac{1}{P_{yjm}} \sum_{p} \left(\frac{s_{yjmp}}{\bar{s}_{yj}} \right) \ln \left(\frac{s_{yjmp}}{\bar{s}_{yj}} \right), \tag{3}$$

where M_{yj} is the number of markets dealer j is active in and P_{yjm} is the number of products they trade in market m.

Intuitively, the index compares a dealer's market share in each segment to their average share across all segments. If a dealer allocates activity evenly across segments, each ratio $\frac{s_{yjmp}}{\bar{s}_{yj}}$ is close to one, and the logarithmic term is close to zero, resulting in a low index. Conversely, when activity is concentrated in a few segments, the index increases. As activity becomes more concentrated, the log-ratio terms grow, and the index increases. Hereby, the logarithmic form plays a central role by quantifying how much the dealer's distribution deviates from uniform allocation. It amplifies larger deviations while diminishing the effect of small ones, grows sublinearly to avoid overweighting extreme values, and ensures that the measure is additive across segments—in our case product-market segments— (Shannon (1948); Theil (1967)).

By construction, the original index does not penalize dealers for non-participation in certain product–market segments (i.e., when $s_{yjmp} = 0$) or even entire markets, since it conditions on observed participation. However, in our setting, such non-participation is common and should be reflected as higher specialization.

We therefore slightly modify the index:

$$T_{yj} = \frac{1}{MP \ln(MP)} \sum_{m} \sum_{p} \mathbb{I}(s_{yjmp} > 0) \left(\frac{s_{yjmp}}{\bar{s}_{yj}} \right) \ln \left(\frac{s_{yjmp}}{\bar{s}_{yj}} \right), \tag{4}$$

which equals zero when the dealer has equal market share across all product–market segments and approaches one when trading is fully concentrated in a single segment. Dividing by $\ln(MP)$, the maximum possible value, standardizes the index between 0 and 1.

Our measure preserves the useful mathematical property of the Theil index that it decomposes into two components—one capturing product specialization, T^{w}_{j}, and the other capturing market specialization, T^{a}_{yj}—allowing us to distinguish between these two dimensions of dealer specialization.

Proposition 1. *$T_{yj} = T^{w}_{yj} + T^{a}_{yj}$, where*

$$T^{w}_{yj} = \frac{1}{M \ln(MP)} \sum_{m} T^{w}_{yjm}, \text{ where } T^{w}_{yjm} = \frac{1}{P} \sum_{p} \mathbb{I}(s_{yjmp} > 0) \left(\frac{s_{yjmp}}{\bar{s}_{yj}} \right) \ln \left(\frac{s_{jmp}}{\bar{s}_{yjm}} \right) \tag{5}$$

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

$$T^a_{yj} = \frac{1}{M \ln(MP)} \sum_m \mathbb{I}(\bar{s}_{yjm} > 0) \left(\frac{\bar{s}_{yjm}}{\bar{s}_{yj}} \right) \ln \left(\frac{\bar{s}_{yjm}}{\bar{s}_{yj}} \right) \tag{6}$$

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

To formally see the connection to the product specialization scores (2), consider the cross-product index for market m, T^w_{yjm}. This index is mostly driven by the ratio that enters the logarithm,

$$\frac{s_{yjmp}}{\bar{s}_{yjm}} = P \times \text{product specialization}_{yjmp}, \tag{7}$$

which is re-weighted according to the importance of each product relative to all markets. Similarly, the cross-market index (6) depends on ratio

$$\frac{\bar{s}_{yjm}}{\bar{s}_{yj}} = M \times \text{market specialization}_{yjm} \tag{8}$$

if $\sum_p s_{yjmp} = s_{yjm}$, otherwise the equation holds approximately. However, Appendix Figure A1 shows that market specialization scores (1), are very similar to $\bar{s}_{yjm}/(M\bar{s}_{yj})$ in our data, implying that expression (8) holds closely.

Figure 4a presents the total and decomposed Theil indices, averaged over years, for dealers active in all markets. These dealers tend to be among the most active and may be of particular interest to those focused on intermediaries who represent the marginal investor. For comparison, Figure 4b shows the corresponding indices for all dealers (who are active in all years).

For most dealers, market specialization exceeds product specialization. This is consistent with the visual patterns in Figures 1b and 3, which suggest greater heterogeneity in specialization across markets than across products. Importantly, this pattern is not mechanically driven by the greater number of products relative to markets—as would be the case with other dispersion measures such as the variance—but instead reflects asymmetric participation across markets. In fact, when there are more products than markets ($P > M$), the potential for product specialization exceeds that for market specialization, as shown in Corollary 1, which derives upper and lower bounds for all indices.

This finding is consistent with the view—central to standard portfolio theory—that asset classes differ more in their risk and return characteristics than individual securities. Standard theory does not predict segmentation by asset class since this limits hedging and arbitrage

possibilities. Our cross-market data show that dealers nonetheless structure their trading and specialization along asset class boundaries. The finding lends support to theories that emphasize segmentation at the level of intermediaries or arbitrageurs—such as Greenwood et al. (2018) and Gourinchas et al. (2025)—rather than investors, as assumed in much of the literature (e.g., Vayanos and Vila (2021)). It calls for more models that endogenize segmentation among intermediaries and explain how it arises across asset classes.

4.2 Drivers of dealer specialization

Before analyzing pricing implications of specialization, we aim to shed some light on why dealers specialize, while recognizing that a full explanation likely requires data beyond trading activity. Our goal is to motivate future empirical research that—with access to richer institutional information—can more precisely unpack the underlying mechanisms, and to help guide the development of new theories on dealer specialization.

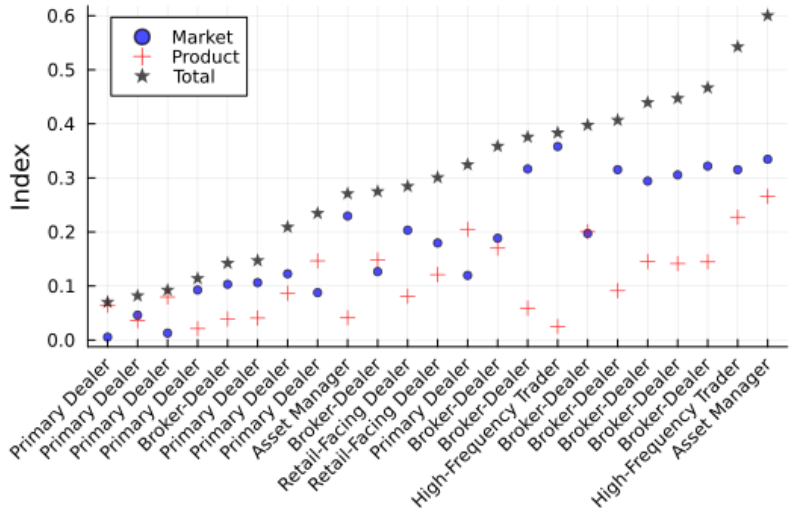
We relate our analysis to existing studies, many of which focus on a single market segment to document how specific frictions or incentives shape trading behavior or prices. In contrast, our approach examines whether these mechanisms remain relevant when aggregating to the product and market level and analyzing variation in the cross-section. We view this as a complement to micro-level studies, and we expect multiple drivers to matter—particularly when examined at finer levels of granularity.

Information-driven specialization. We begin with the conjecture that specialization is information -driven: dealers may focus on securities where they possess an informational edge over other market participants. Micro-level evidence shows that informed dealers achieve better execution outcomes and, in some cases, intermediate a larger share of trades (e.g., Easley et al. (1998) for options and their underlying securities; Hendershott and Madhavan (2015) for U.S. corporate bonds; Collin-Dufresne et al. (2020) for index CDS). Experimental studies also link information asymmetries to trading strategies (Carlin et al. (2013)). However, whether such advantages scale across products and markets remains an open question.

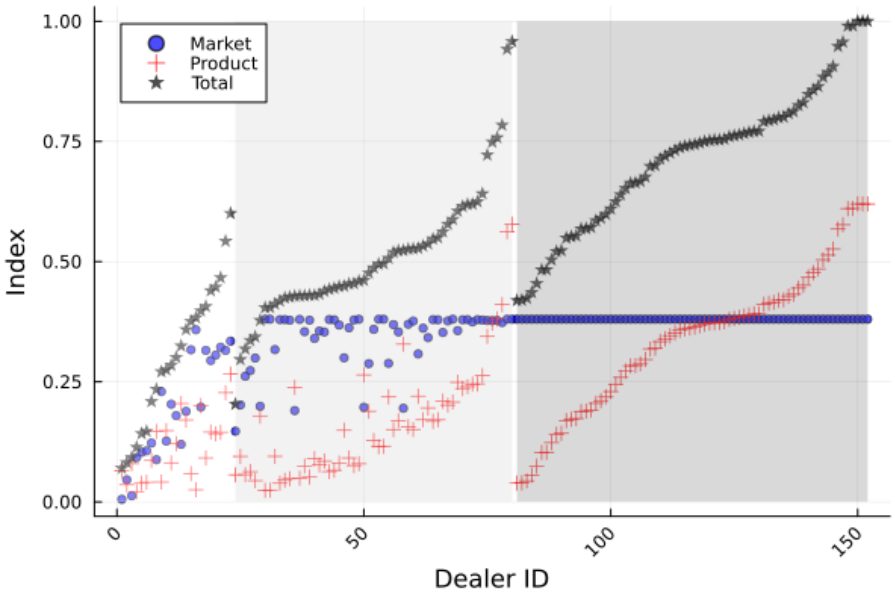
If information advantages are a key of specialization, we would expect specialization to vary over time, reflecting the changing nature of information. To test this, we compute dealer-level Theil indices (4) at different frequencies—monthly, quarterly, and annually—denoted by T_{tj}, where t indexes the period and j the dealer. We then decompose the variation in T_{tj} into two components: one capturing variation across time for the average dealer (ζ_t), and the other capturing variation across dealers for the average time period (ζ_j). This is done by regressing

Figure 4: Within- and cross-market specialization indices

(a) Dealers who are active in all markets

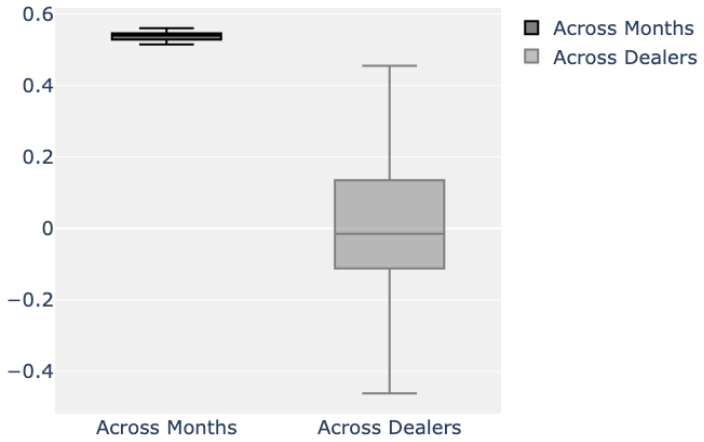


(b) All dealers



Notes: Figure 4a shows the within-market across-product index (Product), the cross-market index (Market), and the total index (Total) for each dealer ID (at the parent level), for an average year. It focuses on the subset of dealers who are active in all three markets. Figure 4b displays the indices for all dealers who are active in all years. Dealers are grouped by market coverage: those active in all three markets appear in the white area, those active in two markets in the light gray area, and those active in only one market in the dark gray area.

Figure 5: Theil index over time vs. across dealers



Notes Figure 5 displays the distribution of dealer and month fixed effects from a regression of the monthly Theil index (4) on dealer and month indicators. We only include dealers active in every month to avoid biases stemming from non-participation, which implies a missing Theil index for that month.

the Theil index on period and dealer fixed effects:

$$T_{tj} = \zeta_t + \zeta_j + \epsilon_{tj}. \tag{9}$$

Figure 5 shows the distributions of ζ_t and ζ_j for the monthly specification. The patterns are similar at quarterly and annual frequencies.

The results show that variation over time for a given dealer is negligible compared to variation across dealers in a given period. This temporal stability suggests that specialization is not primarily driven by shifting information advantages—unless such advantages are persistent and heterogeneous across dealers—but rather by stable factors such as dealer preferences or regulatory and structural frictions.

Size-driven specialization. Another potential driver of specialization is a dealer's overall level of market activity. For instance, Jotikasthira et al. (2024) show that dealers concentrating on specific geographic segments of the U.S. corporate bond market tend to capture larger market shares. This may reflect dealers' reliance on trading networks that expand with market share, reinforcing specialization in segments where their networks are strongest—a mechanism that may be absent outside of OTC markets.

To determine whether dealers with larger market shares are more specialized when going across markets, we regress market specialization scores on aggregated market shares, control-

ling for year-market fixed effects:

$$\text{market specialization}_{yjm} = \alpha + \tilde{s}_{yj} + \zeta_{ym} + \epsilon_{yjm}, \tag{10}$$

where $\tilde{s}_{yj} = \sum_m s_{yjm}$ denotes dealer j's aggregated market share in year y. Similarly, we relate product specialization to product-level market shares, $\tilde{s}_{yjm} = \sum_p s_{yjmp}$, controlling for year–product–market fixed effects.[15]

We find no statistically significant relationship between market specialization and summed market shares, nor between product specialization and summed product-market shares (see Appendix Table A9 and Appendix Figure A6, which presents the corresponding scatter plots). Both small and large dealers can be specialized, indicating that specialization reflects a distinct aspect of dealer behavior, separate from overall market presence. From a welfare perspective, this is encouraging: it suggests, for instance, that scale economies enjoyed by large institutions do not prevent smaller, more specialized dealers from competing effectively—and in some cases, outperforming their larger peers (as we show below).

Legal in-house specialization. In light of empirical studies, including Siriwardane (2019) and Lu and Wallen (2024), we expect some degree of specialization to be driven by in-house segmentation—that is the division of tasks within a firm. Ideally, to analyze the extent of such segmentation, one would observe trading desks within each institution. Our data does not provide this level of granularity.

However, we observe the dealer's LEI and the LEI of their parent institution, and can therefore assess the role of legal in-house segmentation. Some dealer types maintain more LEIs per parent than others (Table 5). For example, it is common for banks to operate multiple subsidiaries, and for fund managers (such as pension funds or asset managers) to manage various funds, each with its own LEI.

We compute Theil indices and their decompositions at both the LEI level and the parent level and compare their magnitudes. If specialization were primarily driven by segmentation across legal entities within a parent firm, we would expect a marked decline in Theil indices when aggregating to the parent level. Figure 6 presents the distribution of Theil indices across LEIs and across parent entities (averaged across years). While the median index does decline, the reduction is modest. This pattern holds even when we focus exclusively on the bond market—where LEI data is most accurate—and on parent holding companies that are banks or

[15] As an alternative size proxy, one could use total trade volume or the fraction of total volume a dealer trades. We avoid this measure because trade volume varies substantially across markets and products (see Appendix Table A5, which creates a mechanical bias.

Table 5: Average number of LEIs per parent holding company

Parent type	Avg # of LEIs
Bank	8.57143
Pension Fund and Insurance	8.5
Asset Manager	4.75
Primary Dealer	4.66667
Broker-Dealer	1.49123
High-Frequency Trader	1.35
Retail-Facing Dealer	1.09091
Other	1.0
Mutual Fund	1.0

Table 5 displays the average number of LEIs per holding company for each dealer type of the holding company. Types are defined in Appendix Table A4.

primary dealers, which tend to have multiple LEIs.

These results suggest that legal in-house segmentation plays a role in shaping specialization, but it is not the primary driver. Our analysis does not capture internal segmentation at the trading desk level within a single legal entity, as we lack desk-level identifiers. As a result, we likely understate the overall degree of internal segmentation. However, given the limited impact of legal segmentation, it seems unlikely that unobserved desk-level segmentation alone could account for the full extent of specialization we observe.

Specialization due to market frictions. Another possible source of specialization is market frictions—impediments that arise because different securities trade in distinct market environments. We discuss two examples of these frictions, while acknowledging that others may also play a role.

One potential market friction is search frictions, particularly in OTC markets (e.g., Duffie (2010)). Dealers known for trading specific products—such as corporate bonds—may find it easier to attract clients seeking those products, which in turn facilitates matching, reduces search costs, and lowers inventory risk. If search frictions were the main source of specialization, we would expect them to be most pronounced in the bond market, and less so in exchange-traded markets like stocks and derivatives. However, Figure 7, which shows the distribution of cross-product Theil indices by market, indicates that the median index is actually highest in derivatives—despite it being exchange-traded. While this does not rule out a role for search frictions, it suggests they are unlikely to be the sole explanation for product specialization.

Figure 6: Assessing legal-in house specialization

(a) Theil index (b) Bond product specialization of banks or PDs

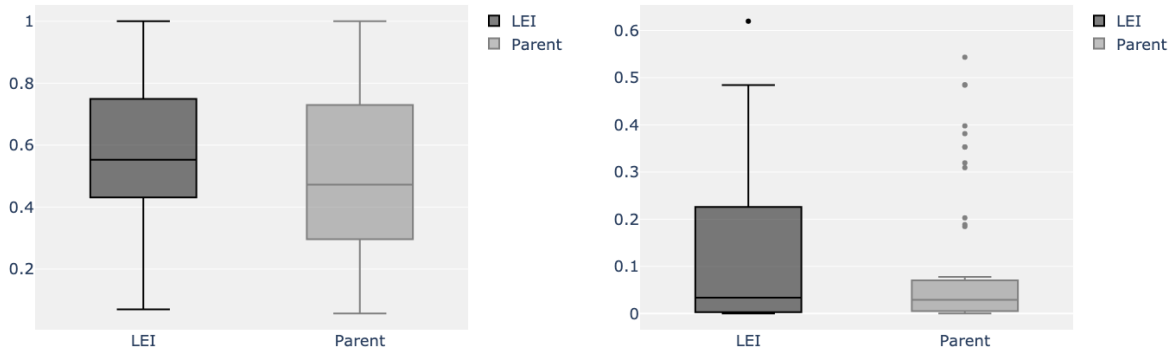


Figure 6a shows the distribution of average annual Theil indices (4) across dealers, when a dealer is defined as one LEI (in black), and when we aggregate LEIs to the holding company parent (in gray). Figure 6b shows annual average cross-produce Theil indices, $\frac{1}{Y}\sum_y T^w_{yjm}$ defined in expression (5), for the bond market, conditional on participation in the bond market, across LEIs and their parents.

Another type of market friction stems from market entry costs, which may deter some trading firms from becoming CIRO-dealer members. While it is difficult to directly assess this mechanism—since market entry costs and market structures remain largely unchanged during our sample period—we can examine dealers that are active across all markets and thus have paid entry fees everywhere. Some of these dealers still exhibit high levels of specialization, with Theil indices exceeding 0.5 (as shown in Figure 4b). While this does not imply that market entry barriers are irrelevant, they alone cannot explain specialization.

In later analysis of transaction prices, we provide suggestive evidence that none of the markets in our sample are entirely frictionless. However, the overall evidence suggest that market frictions do not fully account for specialization.

Clientele. The last potential driver we consider is that dealers specialize in response to the preferred habitats of their clients (e.g., Greenwood and Vissing-Jørgensen (2018); Khetan et al. (2023)).

As a starting point, we compare average market specialization scores and average product specialization scores—conditional on market participation—across dealer types (Appendix Figures A7 and A8). The idea being that different types of dealers cater to different clientele, or, more broadly, follow distinct business models.

We discover substantial heterogeneity across dealer types. To illustrate, we highlight two

Figure 7: Product specialization in each market

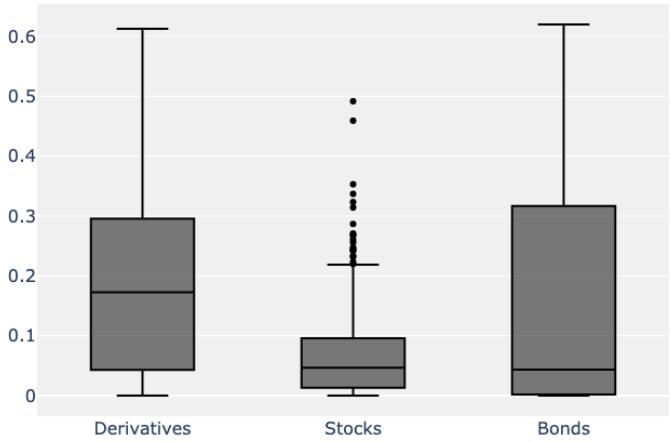


Figure 7 shows annual average cross-produce Theil indices, $\frac{1}{Y}\sum_y T^w_{yjm}$ defined in expression (5), for each market, conditional on participation in the respective market across dealers, i.e., LEIs.

examples: high-frequency traders and retail-facing dealers. High-frequency traders concentrate on derivatives—particularly Treasury futures, index options, and share futures—likely due to the leverage and intraday volatility these instruments offer. In the bond market, they focus on government bonds, consistent with Treasury basis trading strategies. Retail-facing dealers, by contrast, specialize in corporate bonds, likely reflecting retail demand for yield and stable income. Moreover, their focus on equity funds and equity derivatives aligns with the prominence of these products in standardized, diversified investment strategies commonly marketed to retail investors.

In the bond market, we go one step further by relying on the fact that we observe which institutional clients dealers trade with. We observe 4,060 such clients. Although client types share labels with dealer types, they represent separate entities that are not CIRO-registered dealer members. Figure 8 shows average product specialization scores by client type, conditional on dealers trading a given bond product with that client type. For each product–client-type pair, we compute the average score across all dealers involved in those trades, pooling across years. As a robustness check, we also run a saturated regression with product–client-type indicators, which recovers the same group means but allows us to suppress statistically imprecise averages (Appendix Figure A9).

We find significant variation in product specialization by client type—consistent with the idea that dealers tailor their activity to the needs of their specific clientele. Again, we pro-

Figure 8: Bond product specialization scores by client-types



Note: Figure 8 displays average product specialization scores in the bond market, defined in equation (2), by client type. For each client type r and bond product p, we compute the average specialization product specialization$_{yjmp}$ across all dealers j who trade product p with client type r, pooling across years. As an alternative, we can estimate a saturated regression with product–client-type indicators, which yields equivalent group means while allowing us to suppress imprecise estimates.

vide two examples. Dealers trading with foreign central banks tend to specialize in Bankers' Acceptances, which are short-term, high-quality instruments often used in monetary policy operations. Those interacting with high-frequency trading clients—primarily hedge funds within the client base—also concentrate in Bankers' Acceptances, in addition to government bonds. This pattern suggests that hedge funds focus on safe and liquid fixed-income instruments, likely because they can be used as collateral in the repo market or to hedge exposures to riskier assets—such as equities or derivatives—traded through (exchange-member) dealers.

In summary, the evidence suggests that specialization is not driven by a single dominant factor, but rather by a combination of influences. It does not appear to be merely a byproduct of exogenous frictions. Instead, specialization seems to be an integral part of a dealer's business model.

5 Returns to specialization

Thus far, our analysis has focused on market shares—that is, quantities. We now examine whether more specialized dealers (LEIs) trade at different prices than less specialized ones.[16] Such price differences could arise if specialization improves inventory management, influences beliefs about fundamentals, or enables dealers to implement more effective trading strategies—such as timing the market. We focus on relative price differences compared to market averages, leaving potential effects on aggregate price levels to future research.

To analyze systematic price effects in the cross-section—and to reduce the size of our exchange-trade sample—we collapse the exchange data 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.

Measuring price advantages. To detect price differences across both sides of the market and enable comparability across securities with different price levels, we benchmark each transaction against the security's average daily transaction price. This allows us to measure a dealer's relative price advantage compared to the market using a simple, interpretable metric: the percentage difference between a transaction price and the security's average price that day. Formally, for trade τ in security s on day t, we define:

$$\text{margin}_\tau = \frac{\text{average price}_{ts} - \text{trade price}_\tau}{\text{average price}_{ts}} \times 100 \times \text{trade sign}_\tau \tag{11}$$

The trade sign is $+1$ for purchases and -1 for sales. A 1% margin thus means the trade price is 1% below the daily average when buying, or 1% above when selling. To ensure the benchmark price is meaningful, we restrict attention to sufficiently liquid securities, defined as those with at least three trades per day, and confirm that results are robust to stricter thresholds (e.g., five trades per day). We also exclude approximately 2% of derivatives trades executed at negative prices—common in certain spreads—as these would complicate interpretation.

Using an average price as a benchmark differs from market microstructure studies of individual asset classes—such as equities—which often rely on the prevailing mid-quote prior to execution. We depart from that approach for two main reasons. First, our goal is not to unpack fine-grained trading strategies, such as timing or strategic liquidity provision in specific securities, but to assess whether dealers consistently achieve better prices in the cross-section—that

[16] All results are at the LEI level, but they are qualitatively robust to aggregation at the parent level.

is, at a more aggregated level. Second, mid-quotes are not available for many of the securities in our sample, particularly those traded off-exchange, making them infeasible to use.

By construction our margin measure increases in intra-day price volatility, and therefore is largest in the derivative market, followed by equities, and smallest in bonds. The margin mean is approximately 0.14% for derivatives, 0.05% for stocks and close to 0% for bonds; the standard deviation is 15.37 for derivatives, 1.98 for stocks, 0.66 for bonds. Appendix Figure A10 displays the margin distributions per market. In terms of monetary value, a 1% margin translates into roughly C$160 for the median stock trade, C$12,500 for bonds, and 10 cents for derivatives; with the difference being driven from differences in price levels and trade sizes.

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

To find out, we regress our margin measure on dealer-indicator variables separately for each market,

$$\text{margin}_\tau = \alpha + \sum_j \beta_j \mathbb{I}(\text{dealer} = j) + \gamma \text{control}_\tau + \zeta_t + \zeta_{ws} + \epsilon_\tau, \tag{12}$$

and test if any of the dealer coefficient is different 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. All results are relative to the performance of that dealer, and

[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.

therefore not informative in absolute terms. We cluster standard errors at the daily level to account for arbitrary intra-day correlations across dealers, securities, and trades In this and all other margin regressions. This is crucial when traders split orders across exchanges or react to price shocks that impact multiple securities.[18]

Our findings (Figure 9) show that in all markets, certain dealers consistently achieve better prices, both at the LEI and parent level. In stocks, a large broker earns 0.65% more than the baseline (a primary dealer), gaining about C$266 million annually, while a smaller broker loses 0.32%, or C$6 million.[19] Differences are larger in derivatives due to volatility: a hedge fund outperforms by 1.45%, a broker-dealer underperforms by 1.14%, though payment gaps are smaller. In bonds, most underperform the primary dealer, but an insurance company gains 0.11% (C$22 million), while a retail-focused dealer loses 0.48% (C$278k).

In Appendix B.2, we examine whether certain dealer types—such as high-frequency traders or retail-focused dealers—systematically earn better margins using cross-dealer comparisons. This shows, for example, that retail-facing dealers underperform compared to primary dealers in all markets, while high-frequency traders are particularly successful in the derivatives market.

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.

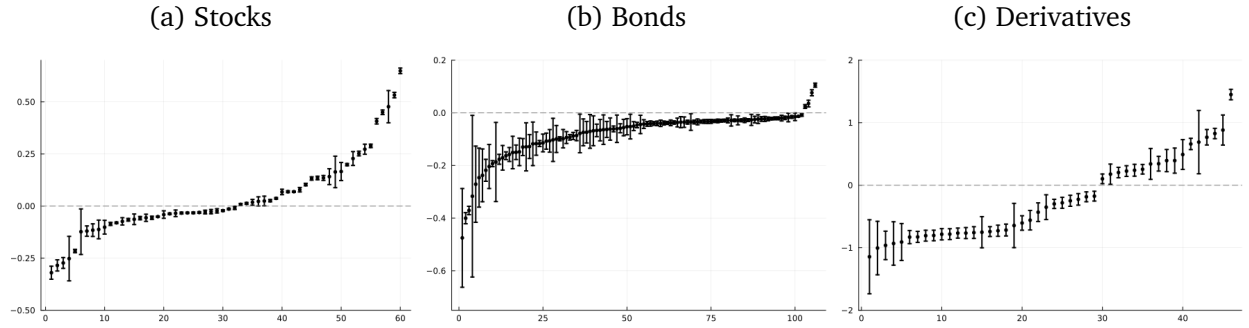
In theory, the relationship between dealer specialization and prices is ambiguous. Greater specialization might enable dealers to buy at lower prices and sell at higher ones. However, it could also be that specialized dealers are more efficient than their less specialized counterparts and, as a result, 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 comparing the prices obtained by extremely specialized dealers—those active in only one market segment—to those of a baseline dealer active across all major markets.

[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.

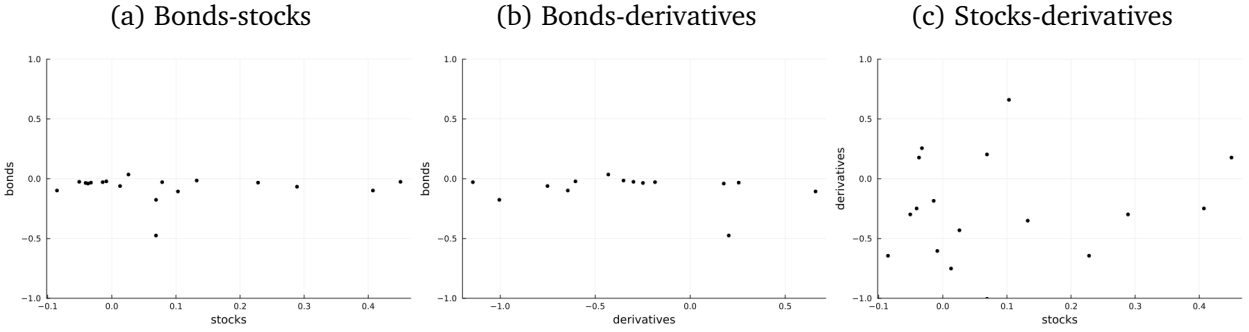
[19] To translate margin percentages into annual monetary losses or gains, we assume that each trade is executed at the median price, using the total amount (e.g., number of shares in the stock market) that the dealer under consideration trades in an average year.

Figure 9: Dealers obtain systematically different margins



Notes: Figure 9a shows the dealer coefficients, and 95% confidence intervals (clustered at the daily-level), when regressing margins (11) 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 9b and 9c 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%. Since we sort coefficients from small to large, the x-axis are not comparable across markets, as they do not reflect dealer identifiers.

Figure 10: Cross-market correlation between dealer coefficients of dealers active in all markets



Notes: Figure 10a shows the within-dealer correlation of coefficients in the bond (y-axis) versus stock market (x-axis), 10b and 10c 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%.

Specifically, we examine dealer coefficients from regression (12) for single-market dealers. Appendix Figure A11 shows that most earn below-average prices relative to the multi-market baseline, a large primary dealer—a result that is unsurprising given the baseline dealer's prominent role across markets. More striking is that some of the extremely specialized dealers outperform the baseline, particularly in the derivatives market, possibly reflecting the complexity and niche expertise required in these products.

We then analyze how dealer performance correlates across markets. Supporting the notion of market specialization, we find little evidence that dealers who achieve better prices in one market do so in others. Figure 10 shows no significant cross-market correlations in dealer coefficients from regression (12), focusing on dealers active in multiple markets. A similar pattern emerges when we run analogous regressions across products within a given market—dealers rarely outperform consistently across multiple products.

Subsequently, we leverage cross-sectional variation across dealers in their product and market specialization scores (1) and (2). A first approach is to regress margins on specialization scores, in addition to the same control variables and fixed effects we have used above, for regression (12).

$$\text{margin}_{\tau} = \alpha + \beta \text{market specialization}_{yjm} + \gamma \text{controls}_{\tau} + \zeta_{t} + \zeta_{ws} + \epsilon_{\tau}, \tag{13}$$

and similarly with product specialization, which varies by year, dealer, market and product: $\text{product specialization}_{yjmp}$. Appendix Table A10 shows that market and product specialization scores correlate positively with margins.

To mitigate concerns about reverse causality—where more successful dealers may become more specialized—and omitted variables—such as dealer sophistication or efficiency—we implement two complementary strategies. While neither approach offers definitive causal identification, each provides suggestive evidence, and together with the simple correlations, they paint a consistent and coherent picture that specialization is associated with better prices.

Our first approach is to lag specialization scores to examine whether dealers who were more specialized last year obtain better prices today. Specifically, we estimate regression (13), but replace $\text{market specialization}_{yjm}$ with $\text{market specialization}_{y-1,jm}$, and similarly for the product specialization scores. This specification would recover the causal effect of specialization on margins if past specialization were exogenous to current pricing—that is, if dealers did not choose last year's specialization in anticipation of future margin opportunities, and if there were no time-persistent, unobserved dealer-level characteristics that influence both specialization and pricing.

Although unobserved factors like business models or expertise could, in principle, confound

the relationship, two features of our setting reduce this concern. First, because specialization is stable over time, lagged specialization is unlikely to reflect short-term responses to anticipated pricing gains, helping to mitigate reverse causality. Second, while business models may shape where dealers operate, they are less likely to drive time-varying differences in within-market specialization or transaction-level pricing, conditional on market participation.

Table 6 presents the regression results, showing that dealers with higher specialization scores in the previous year earn larger margins today. The magnitudes of these (per-trade) effects are relatively modest, yet they aggregate to significant amounts over time in the stock and bond market. For example, on the stock exchange, moving from no market specialization (score of 0) to full specialization (score of 1) increases margins by 5.9 basis points—equivalent to 0.03 standard deviations; a full shift in product specialization raises margins by 46.4 basis points, or 0.23 standard deviations. In the bond market, margins increase by 0.017 and 0.047 standard deviations for market and product specialization, respectively. In the derivatives market, the corresponding increases are 0.13 and 0.14 standard deviations.

Product coefficients are larger in the stock and bond markets and similar in size for derivatives. Combined with the fact that market-level specialization is more pronounced when measured in quantities, this pattern is consistent with the idea that part of the observed specialization reflects market frictions or internal organizational boundaries. If such frictions are more severe across markets than within them—which is plausible, particularly for market frictions—then some of the observed market specialization may not arise from strategic choices. In that case, it would not necessarily translate into trading advantages.

Our second approach is to instrument for specialization scores in regression (13)—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).[20]

We use daily client orders on the exchanges as an instrument for dealer specialization when dealers trade for their own accounts. The idea is that dealers have limited discretion over client orders, which must be executed promptly. For example, a retail investor placing a stock

[20]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)).

Table 6: Correlation between margins and last year's specialization scores

	Stocks		Bonds		Derivatives	
Lagged market specialization	0.059***		0.011**		2.016***	
	(0.002)		(0.003)		(0.041)	
Lagged product specialization		0.464***		0.031***		2.077***
		(0.006)		(0.005)		(0.051)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	87,299,519	87,299,519	5,322,564	5,322,184	4,186,707	4,186,707
R^2	0.008	0.008	0.041	0.041	0.158	0.158
Within-R^2	0.000	0.000	0.000	0.000	0.002	0.002

Table 6 shows the estimation results from regressing margins (11) 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 (A11): trade size, date fixed effects and security-year-week fixed effects. We cluster standard errors at the daily-level.

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 would allow us to identify the causal effect of specialization on stock market margin.

Table 7 reports the IV estimates separately for the stock and derivative exchanges. The instrument—the fraction of volume a dealer executes for clients within a given market segment—is negatively correlated with both market and product specialization scores. This suggests that specialization is not primarily driven by client order flow, which is intuitive given that dealers with large client businesses tend to be banks with broad investment mandates. The negative correlation is particularly pronounced on the derivative exchange, consistent with the presence of high-frequency traders who are less engaged in client trading.

Across both exchanges, the IV estimates confirm that more specialized dealers obtain better prices. The results are qualitatively consistent with the correlation patterns, with effect sizes in the same general range—though not identical. The estimated IV effects range from 27 to 208 basis points, corresponding to 0.13 to 1.05 standard deviations. The IV estimates are consistently larger than their OLS counterparts, suggesting that OLS may be downward biased. One explanation is that more specialized dealers operate more efficiently and are therefore willing to accept lower margins, which would attenuate the OLS relationship. The OLS coefficients in this table differ from those in Appendix Table A10 because the current analysis is restricted to trades for dealers' own accounts, whereas Appendix Table A10 and Table 6 include all trades.

Neither of our strategies to address endogeneity is without limitations. For example, if

Table 7: IV regressions of margins on specialization scores for stocks

	(First Stage)	(OLS)	(IV)		(First Stage)	(OLS)	(IV)
s^c_{tjm}	-0.887***			s^c_{tjmp}	-0.091***		
	(0.010)				(0.001)		
Market specialization		0.163***	0.271***	Product specialization		0.570***	2.085***
		(0.003)	(0.014)			(0.014)	(0.110)
Controls	Yes	Yes	Yes	Controls	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Fixed effects	Yes	Yes	Yes
N	30,356,466	30,356,466	30,356,466	N	30,356,466	30,356,466	30,356,466
R^2	0.217	0.128	0.128	R^2	0.298	0.128	0.128
Within-R^2	0.031	0.001	0.000	Within-R^2	0.009	0.001	0.000

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

	(First Stage)	(OLS)	(IV)		(First Stage)	(OLS)	(IV)
s^c_{tjm}	-10.702***			s^c_{tjmp}	-4.775***		
	(0.290)				(0.198)		
Market specialization		0.706***	0.791***	Product specialization		0.557***	0.640***
		(0.034)	(0.066)			(0.039)	(0.087)
Controls	Yes	Yes	Yes	Controls	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Fixed effects	Yes	Yes	Yes
N	2,835,371	2,835,371	2,835,371	N	2,835,371	2,835,371	2,835,371
R^2	0.497	0.194	0.194	R^2	0.607	0.194	0.193
Within-R^2	0.344	0.000	0.000	Within-R^2	0.320	0.000	0.000

Table 7 shows the IV estimation results for own-account stock market trades. Consider the LHS of 7. 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 trade-volume dealer j executes for clients in market m on day t, $s^c_{tjm} \in [0,1]$. 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 trade-volume dealer j executes of product p for clients in market m on day t, $s^c_{tjmp} \in [0,1]$. In all regressions we include the same control variables and fixed effects as in regression (A11): trade size, date fixed effects and security-year-week fixed effects. We cluster standard errors at the daily-level.

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. We view this risk as limited given that clients do not observe transaction prices, which are private to exchange members and the exchange operator, making such selection unlikely to be a first-order concern. More importantly, all approaches—despite their individual limitations—reveal 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 captures trading activity across all major Canadian financial markets. Dealers tend to concentrate their activity not only within specific markets but also within particular types of securities, with market-level specialization being more pronounced. This specialization does not appear to be driven by any single factor, but likely reflects a combination of firm boundaries, market frictions, and client composition. More specialized dealers consistently obtain better prices, indicating that specialization is not merely a byproduct of constraints but a strategic choice embedded in the dealer's business model. In a nutshell, our findings highlight that specialization matters—both in terms of quantities traded and prices obtained—and is likely the result of strategic choice. It therefore warrants greater attention in both theoretical and empirical research on financial intermediation and market structure.

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.

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.

Carlin, B. I., S. Kogan, and R. Lowery (2013). Trading complex assets. *Journal of Finance 68*(5), 1937–1960.

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.

Collin-Dufresne, P., B. Junge, and A. Trolle (2020). Market structure and transaction costs of index cdss. *Review of Financial Studies 33*(5), 1969–2003.

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

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.

Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. *The Journal of Finance 65*(4), 1237–1267.

Easley, D., M. O'Hara, and P. S. Srinivas (1998). Option volume and stock prices: Evidence on where informed traders trade. *The Journal of Finance 53*(2), 431–465.

Eisfeldt, A. L., B. Herskovic, and S. Liu (2023). Interdealer price dispersion. Working Paper.

Eisfeldt, A. L., H. Lustig, and L. Zhang (2023). Complex asset markets. *Journal of Finance 78*(5), 2519–2562.

Figlewski, S. (1989). Options arbitrage in imperfect markets. *The journal of Finance 44*(5), 1289–1311.

FINRA (2024). 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.

Gabaix, X., A. Krishnamurthy, and O. Vigneron (2007). Limits of arbitrage: Theory and evidence from the mortgage-backed securities market. *The Journal of Finance 62*(2), 557–595.

Gourinchas, P.-O., W. Ray, and D. Vayanos (2025). A preferred-habitat model of term premia, exchange rates, and monetary policy spillovers. *American Economic Review forthcoming*.

Greenwood, R., S. Hanson, J. C. Stein, and A. Sunderam (2023). A quantity-driven theory of term premia and exchange rates. *The Quarterly Journal of Economics 138*(4), 2327–2389.

Greenwood, R., S. G. Hanson, and G. Y. Liao (2018). Asset price dynamics in partially segmented markets. *The Review of Financial Studies 31*(9), 3307–3343.

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

Greenwood, R. and A. Vissing-Jørgensen (2018). The impact of pensions and insurance on global yield curves. 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.

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

Jotikasthira, C., C. T. Lundblad, and J. Xue. (2024). Dealer specialization and market segmentation. Working paper.

Khetan, U., J. Li, I. Neamtu, and I. Sen (2023). The market for sharing interest rate risk: Quantities and asset prices. Working paper.

Koijen, R. S. J., T. J. Moskowitz, and L. H. Pedersen (2022). Intermediary asset pricing: New evidence from many asset classes. Working Paper 21920, National Bureau of Economic Research. Previously NBER Working Paper No. 21920; empirical evidence across equities, bonds, FX, commodities.

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). The effects of quantitative easing on interest rates: Channels and implications for policy. *Brookings Papers on Economic Activity*, 215–287.

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 Review of Financial Economics 8*, 1–24.

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

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.

Rostek, M. J. and J. H. Yoon (2025). Imperfect competition in financial markets: Recent developments. *Journal of Economic Literature forthcoming*.

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

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal 27*(3), 379–423.

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

Specialization in Financial Markets

by Milena Wittwer, and Andreas Uthemann

Section A provides details regarding data cleaning.
Section B gives additional empirical evidence.
Section C provides proofs.

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 Additional empirical evidence

B.1 Dealer specialization — Correlation patterns

In this appendix, we complement our analysis of dealer specialization by examining correlations in monthly returns, trade volumes, and dealer trade flows across products and markets.[21] These correlation patterns provide an alternative lens through which to assess the degree of interdependence or segmentation across markets and product categories, reinforcing the findings based on specialization scores in the main text.

Returns and aggregate trade volumes. As a starting point, we compute the pairwise correlation of average monthly returns and total trade volumes across all products and markets in Appendix Figure A12. To construct these returns, we first compute the average monthly return per symbol (conditional on being traded), as the percentage change in the average monthly symbol-price. Then we average across all symbols within a given product category. We see that correlations tend to be stronger within markets than across markets, suggesting that market segmentation is reflected not only in dealer behavior but also in asset price and volume co-movements.

There is a relatively high degree of interdependence between the stock and bond markets, particularly compared to their correlations with the derivatives market, which appears more isolated. This is in line with the fact that there are some trading firms that only trade derivatives, leading to higher segmentation of that market compared to the others.

At the same time, there is substantial heterogeneity within markets, supporting the idea

[21] When we shorten the time horizon—from months to weeks or trading days—the correlations weaken, but the overall structure remains similar. This likely reflects the increasing influence of idiosyncratic shocks at higher frequencies, which dominate over systematic cross-product or cross-market trading patterns.

product specialization, especially in the bond market, but also in the stock market. For example, returns on funds are weakly negatively correlated with returns on other equity products.

Dealer net demand and gross trade volume. Next, we examine trade flows at the dealer level to assess whether dealers take systematically aligned or offsetting positions across products. Specifically, we analyze how a dealer's net demand in one product, p, defined as the total quantity bought minus the total quantity sold in a given month, relates to their net demand in another product, $q \neq p$.

To do this, we estimate a series of pairwise regressions of the form:

$$\tilde{dem}_{ymjp} = \alpha + \beta_{pq} \times \tilde{dem}_{ymjq} + \text{fixed effects} + \epsilon_{ymjp},$$

where $\tilde{dem}_{ymjp}$ denotes the standardized net demand of dealer j for product p in month m of year y. Standardization ensures that products with different overall trading volumes are comparable:

$$\tilde{dem}_{ymjp} = \frac{dem_{ymjp} - \mu_p}{\sigma_p},$$

where μ_p and σ_p are the mean and standard deviation of dem_{ymjp} for product p, computed across all dealers and time periods.

We first include year-month fixed effects, ζ_{ym}, to control for common shocks that affect all dealers in a given period. In this specification, a positive coefficient β_{pq} indicates that dealers who tend to buy product q more than others in a given month also tend to buy more of product p. In other words, β_{pq} captures whether dealers co-move across products within a month, using cross-sectional variation across dealers.

Next, we add dealer fixed effects, ζ_j, to absorb persistent differences in overall trading levels across dealers. In this second specification, β_{pq} is identified purely from within-dealer variation over time. A positive coefficient now implies that when a given dealer increases their net demand for product q relative to their own average, they also tend to increase their net demand for product p in the same month. If the coefficient shrinks toward zero compared to the first specification, this suggests that the earlier pattern was driven by differences in specialization across dealers, rather than coordinated variation in behavior across products within a dealer.

Appendix Figure A13 show the estimated coefficients β_{pq} from these regressions. When we include only year-month fixed effects, we observe substantial positive and negative correlations across some of the products. The graph looks similar, even though the point estimates

differ marginally when we do not include year-month fixed effects, which implies that the estimated co-movement in demand across products reflects dealer-level variation, rather than being mechanically induced by shared time shocks. The findings show that, in a given month, some dealers tend to take offsetting positions across products—for example, going long in one product while going short in another.

However, once we include both month and dealer fixed effects, most of these correlations attenuate or disappear. This suggests that the cross-product patterns in the first specification are not driven by consistent multi-product strategies within dealers. Rather, they reflect differences in specialization across dealers—some focus more on certain products than others—reinforcing the idea of segmented dealer activity across product.

At the aggregate level—comparing across dealers—some products are more closely connected than others, as reflected in stronger correlations between their net demands. Within the stock market, funds exhibit the weakest correlation with other equity products, which are otherwise similarly correlated with each other. In the bond market, asset-backed securities are weakly negatively correlated with other bond types, whereas the remaining fixed-income products show positive correlations.

Comparing across asset classes, correlations among stock products are more homogeneous than among bond products, consistent with our finding that product specialization is more pronounced in the bond market than in equities. Within the derivative market, most coefficients close to zero or statistically insignificant, although some derivative products show notable correlations with bond or equity instruments. Currency options stand out: they are negatively correlated with most equity products and positively correlated with most fixed-income products. Somewhat surprisingly, Treasury futures show no strong connection with government bonds—the estimated coefficient is statistically insignificant.

Appendix Figure A14 displays the estimated coefficients from regressions where net demand is replaced with standardized gross trade volume (i.e., total purchases or sales). As before, the inclusion of dealer fixed effects causes the coefficients to shrink toward zero, indicating that much of the correlation is driven by dealer-level heterogeneity, i.e., their specialization. Among the three markets, the derivative market appears the most segmented, with the weakest correlations across products—also supporting previous findings.

Taken together, the correlation patterns in net- and gross demands complement our analysis with specializations cores. They are in line with the idea that dealers specialize on specific products and markets.

B.2 Margins and dealer types

In this appendix, we examine whether different dealer types systematically earn different margins. The regression average across substantial within-type heterogeneity (we know exist from Figure 9) and uneven type category sizes (recall Table 4). For instance, findings for pension funds and insurance companies are based on just three institutions and may not be representative of the broader sector. In contrast, the high-frequency trader category includes 21 institutions with varied business models and trading strategies.

We estimate regression (12) using dealer-type indicators, with primary dealers as the reference category. Results are reported in Appendix Table A11.

Consistent with prior work (e.g., Bernhardt et al. (2005)), larger trades receive worse prices on exchanges. Moreover, prices are also worse for client trades than for inventory or market-making accounts. In the bond market, we find no evidence of trade-size discounts, consistent with the mixed findings in the literature (Pinter et al. (2024); Allen and Wittwer (2024)). Dealers earn higher margins when trading with clients rather than with other dealers or brokers—again in line with existing evidence.

Performance varies across dealer types and markets. In all markets, retail-focused dealers earn significantly worse margins than primary dealers. High-frequency traders outperform in the derivative market but show no advantage in stocks or bonds. Asset managers consistently underperform relative to primary dealers across all three markets. Pension funds and insurance companies do poorly on the stock exchange but well in the bond market. Mutual funds, by contrast, are the best-performing traders on the stock exchange—the only category outperforming primary dealers (the baseline).

C Mathematical details

Proof of Proposition 1. We derive the decomposition for the general framework with P_m number of products in market m (while we assume $P_m = P$ for all m in the main text to simplify notation). Moreover, we drop the year subscript, y, for notational convenience.

To start, note that:

$$\bar{s}_{jm} = \frac{1}{P_m} \sum_p s_{jmp},$$
$$\bar{s}_j = \frac{1}{M} \sum_m \frac{1}{P_m} \sum_p s_{jmp} = \frac{1}{M} \sum_m \bar{s}_{jm}.$$

To derive the decomposition, we consider the non-standardized adjusted Theil index: $\tilde{T}_j =$

$\ln(MP)T_j$. The non-standardized adjusted Theil index is:

$$\begin{aligned}
\tilde{T}_j &= \frac{1}{M}\sum_m \frac{1}{P_m}\sum_p \mathbb{I}(s_{jmp}>0)\left(\frac{s_{jmp}}{\bar{s}_j}\right)\ln\left(\frac{s_{jmp}}{\bar{s}_j}\right) \\
&= \frac{1}{M}\sum_m \frac{1}{P_m}\sum_p \mathbb{I}(s_{jmp}>0)\left(\frac{s_{jmp}}{\bar{s}_j}\right)\left[\ln\left(\frac{s_{jmp}}{\bar{s}_{jm}}\right)+\ln\left(\frac{\bar{s}_{jm}}{\bar{s}_j}\right)\right] \\
&= \frac{1}{M}\sum_m \frac{1}{P_m}\sum_p \mathbb{I}(s_{jmp}>0)\left(\frac{s_{jmp}}{\bar{s}_j}\right)\ln\left(\frac{s_{jmp}}{\bar{s}_{jm}}\right) \\
&+\frac{1}{M}\sum_m \mathbb{I}(\bar{s}_{jm}>0)\ln\left(\frac{\bar{s}_{jm}}{\bar{s}_j}\right)\frac{1}{P_m}\sum_p \mathbb{I}(s_{jmp}>0)\left(\frac{s_{jmp}}{\bar{s}_j}\right).
\end{aligned}$$

Using

$$\sum_p \mathbb{I}(s_{jmp}>0)\left(\frac{s_{jmp}}{\bar{s}_j}\right)=\frac{\sum_p \mathbb{I}(s_{jmp}>0)s_{jmp}}{\bar{s}_j}=\frac{P_m\bar{s}_{jm}}{\bar{s}_j},$$

we simplify:

$$\begin{aligned}
\tilde{T}_j &= \frac{1}{M}\sum_m \frac{1}{P_m}\sum_p \mathbb{I}(s_{jmp}>0)\left(\frac{s_{jmp}}{\bar{s}_j}\right)\ln\left(\frac{s_{jmp}}{\bar{s}_{jm}}\right)+\frac{1}{M}\sum_m \mathbb{I}(\bar{s}_{jm}>0)\left(\frac{\bar{s}_{jm}}{\bar{s}_j}\right)\ln\left(\frac{\bar{s}_{jm}}{\bar{s}_j}\right). \\
&= \tilde{T}_j^w+\tilde{T}_j^a, \quad \text{with } \tilde{T}_j^w=\frac{1}{M}\sum_m \tilde{T}_{jm}^w.
\end{aligned}$$

$\tilde{T}_j$ assumes its smallest value, 0, if dealer j trades the same amount in all market product segments, i.e., $s_{jmp}=s$ for all m,p. It assumes the highest value when the dealer trades in one single market product segment, in the market with the most number of products. Refer to this market by k and let s be the dealer's only non-zero product market share in that market. In that case:

$$\begin{aligned}
&\bar{s}_{jk}=\frac{s}{P_k} \quad \text{and } \bar{s}_{jm}=0 \quad \text{for all } m\neq k. \\
&\bar{s}_j=\frac{1}{M}\sum_m \bar{s}_{jm}=\frac{s}{MP_k}. \\
&\tilde{T}_j=\frac{1}{M}\sum_m \frac{1}{P_m}\sum_p \mathbb{I}(s_{jmp}>0)\left(\frac{s_{jmp}}{\bar{s}_j}\right)\ln\left(\frac{s_{jmp}}{\bar{s}_j}\right)=\frac{1}{M}\frac{1}{P_k}(MP_k)\ln(MP_k)=\ln(MP_k)
\end{aligned}$$

Therefore, the standardized Theil index T_j lies between 0 and 1 where $T_j=T_j^w+T_j^a=\tilde{T}_j/\ln(MP_k)$ with $T_j^w=\tilde{T}_j^w/\ln(MP_k)$ and $T_j^a=\tilde{T}_j^a/\ln(MP_k)$. □

Corollary 1. *Let $P_m = P$, and T_j denote the standardized adjusted Theil index for dealer j. Then:*

(i) For dealers who are active in all markets (i.e., $\bar{s}_{jm} > 0$ for all m),

$$T_j \in [0,1), \quad T_j^w \in \left[0, \frac{\ln(P)}{\ln(MP)}\right], \quad T_j^a = \left[0, \frac{\ln M}{\ln(MP)}\right).$$

(ii) For dealers who are active in two markets only,

$$T_j \in \left[\frac{\ln(M/2)}{\ln(MP)}, 1\right), \quad T_j^w \in \left[0, \frac{\ln(P)}{\ln(MP)}\right], \quad T_j^a = \left[\frac{\ln(M/2)}{\ln(MP)}, \frac{\ln M}{\ln(MP)}\right).$$

(iii) For dealers who are active in one market only,

$$T_j \in \left[\frac{\ln(M)}{\ln(MP)}, 1\right], \quad T_j^w \in \left[0, \frac{\ln(P)}{\ln(MP)}\right], \quad T_j^a = \frac{\ln(M)}{\ln(MP)}.$$

Proof of Corollary 1. *Bounds for T_j.* First, we show that the upper bound of the Theil index is 1 for all three cases, but that this upper bound is only attained when a dealer is only active in one market. To show this, we assume that dealer j is only active in one product segment per market. In all other product segments the market share is 0. Let k refer to the market with the largest number of products P_k. In market k, the dealer's market share is $s > 0$ in the single product segments they are active in. For all other markets $m \neq k$, the product share is αs in the active product segment of that market. We have

$$\bar{s}_{jk} = \frac{s}{P_k} \quad \text{and } \bar{s}_{jm} = \frac{\alpha s}{P_m} \text{for all } m \neq k.$$

Together, this implies

$$\bar{s}_j = \frac{1}{M}\left[\frac{s}{P_k} + \sum_{m \neq k} \frac{\alpha s}{P_m}\right] \equiv h(\alpha)s \quad \text{with } \lim_{\alpha \to 0} h(\alpha) = \frac{1}{MP_k}.$$

It follows that

$$\tilde{T}_j = \frac{1}{M}\left[\frac{1}{P_k}\left(\frac{1}{h(\alpha)}\right)\ln\left(\frac{1}{h(\alpha)}\right) + \sum_{m \neq k} \frac{1}{P_m}\left(\frac{\alpha}{h(\alpha)}\right)\ln\left(\frac{\alpha}{h(\alpha)}\right)\right].$$

Note that, as $\lim_{x \to 0} x \ln x = 0$, we have $\lim_{\alpha \to 0} \tilde{T}_j = \ln(MP_k)$. If the dealer is only active in market k, we can set $\alpha = 0$ and obtain $\tilde{T}_j = \ln(MP_k)$. If the dealer is only active in $N < (M-1)$ of the other markets, we can set $\alpha = 0$ for the inactive markets and the above argument holds

when restricted to the active markets m.

We next consider the lower bound for the Theil index for all three case. We assume an equal number of product segments in all markets, that is $P_m = P$ for all m.

First, suppose that dealer j is active in all markets and allocates the same volume across all market-product segments. That is, $s_{jmp} = s$ for all m, p. Then the average shares satisfy $\bar{s}_{jm} = s$ and $\bar{s}_j = s$, and it follows immediately that $T_j = 0$.

Second, consider a dealer who is active in exactly two markets. If the dealer splits volume evenly across the two markets and allocates volume equally across all products within each market, then within-market variation is zero, $T_j^w = 0$, and the across-market component becomes $\tilde{T}_j^a = \ln(M/2)$. The total index is then $T_j = \ln(M/2)/\ln(MP)$.

Finally, consider a dealer who is active in a single market. In this case, the across-market component is maximized, with $\tilde{T}_j^a = \ln(M)$, since all volume is allocated to one market only. If the dealer spreads volume equally across all products in that market, then $T_j^w = 0$. Hence, the minimum value for the total index is $T_j = \ln(M)/\ln(MP)$.

Bounds for T_j^w. To derive the upper and lower bounds for T_j^w we again assume equal number of product segments in all markets, $P_m = P$. Setting market shares in all product segments of all (active) markets m to a constant $s > 0$, we immediately obtain a lower bound of 0 as $s_{jmp}/\bar{s}_{jm} = 1$ in these markets.

The upper bound is obtained by assuming a market share of $s > 0$ in a single product segment in market k and a zero market share in all other product segments and markets. In that case, we have $\bar{s}_{jk} = s/P$, $\bar{s}_{jm} = 0$ for all $m \neq k$ and $\bar{s}_j = s/(MP)$. Together, this implies $\tilde{T}_j^w = \ln P$ and therefore $T_j^w = \ln P/\ln(MP)$.

Bounds for T_j^a. To obtain an upper bound for T_j^a, we again assume that the dealer has a market share of $s > 0$ in a single product segment of market k, and a market share of αs in a single product segment of all other markets $m \neq k$ where the dealer is also active. Using the same type of argument as in the derivation of the upper bound for T_j, we find that the upper bound for $\tilde{T}_j^a$ is $\ln(M)$ and is only attained when the dealer is only active in a single market but it can be reached arbitrarily closely by letting α go to zero when the dealer is active in several markets.

To obtain the lower bound for T_j^a, assume that the dealer has equal average market shares $\bar{s}_{jm} = s > 0$ in all active markets. Suppose the dealer is active in $N \leq M$ markets. Then $\bar{s}_j = (N/M)s$. It follows that $\tilde{T}_j^a = \ln(M/N)$ and therefore $T_j^a = \ln(M/N)/\ln(MP)$. □

Appendix Table A1: Fixed-income products

Product	Description
Government Bonds	Government of Canada Bonds, Government of Canada Real Return Bonds, Government of Canada T-bills
Provincial Bonds	Provincial Bills, Provincial Bonds, Provincial Commercial Papers, Municipal Bonds
Bank/Agency Securities	Bank Commercial Papers and Bank Security - Notes/Bonds/Debentures. Agency Bonds and Agency Commercial Papers
Bankers' Acceptance	Bankers' Acceptance
Corporate Bonds	Corporate Bonds
Asset-Backed Securities	Mortgage-Backed Securities, Asset-Backed Securities, Canada Mortgage Bonds.

Appendix Table A2: Equity products

Product	Description
Stocks	Common shares (symbols without suffixes), grouped by sector:
– Natural Resources	Companies in Mining or Oil, Gas
– Financial Sector	Companies in Financial Services or Real Estate
– Technology & Industry	Companies in Technology, Utilities and Pipelines, and Renewable Energy
– Consumer & Healthcare	Companies in Media and Life Science or Products and Services (including Consumer and Industrial Products and Services)
Non-Common Shares	Securities with suffixes that are not classified as fund, including preferred shares, warrants, notes, debentures, split shares, income trusts, exchange-traded receipts, CPC/SPAC listings, and other structured equity instruments
Funds	Symbols which are marked as Exchange Traded Funds, Fund of Equities, Fund of Debt, Fixed Income, Fund of Mortgages/MBS, Commodity Funds, Fund of Multi-Asset, Fund of Multi-Asset/Other, and Fund of Funds

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

Asset Manager	Financial firm that manages investment portfolios and makes investment decisions on behalf of clients, typically charging fees based on assets under management. A key focus of the firm is long-term investment strategy and portfolio construction rather than securities transactions. Entities can manage multiple types of funds such as hedge funds, mutual funds or ETFs.
Bank	Bank, investment bank, retail bank, or credit union, and any other entity that is deposit taking, which is not a primary dealer.
Broker-Dealer	Financial firm that buys and sells securities either for their own account (as a dealer) or on behalf of clients (as a broker). A key focus of the firm is facilitating securities transactions.
High-Frequency Trader	Proprietary trading firm: Financial firm that trades securities, commodities, or other financial instruments using its own capital rather than client money distinguishing it from broker-dealers. A key focus of the firm is to use sophisticated strategies and technology to generate returns from their trading activities. Hedge funds: Financial firm that uses alternative investment strategies, including leverage, derivatives, and short selling, and is typically organized as a private partnership. It charges both management and performance fees, limits the number of investors and imposes restrictions on investor withdrawals. Private Equity Fund: Financial firm that invests in private companies or public companies with the intent to take them private, typically through leveraged buyouts. Private Equity funds have longer investment horizons than hedge funds, require larger minimum commitments, typically offer no liquidity during the investment period, and focus on improving operational performance of portfolio companies rather than trading securities.
Mutual Fund	Financial firm that is a mutual fund or a mutual fund manager. A mutual fund pools money from many investors to purchase a diversified portfolio of stocks, bonds, or other securities.
Pension Fund and Insurance	Financial firm whose purpose is to manage investments (and/or provide services) related to pension, retirement, insurance, re-insurance, benefits, and superannuation funds
Primary Dealer	Banks and broker-dealers that are authorized to trade government securities directly with the Bank of Canada and are required to participate in Canadian government debt auctions. A primary dealer is designated by the Department of Finance or Bank of Canada based on having significant market share in Government of Canada bonds and treasury bills, plus sufficient resources and commitment to actively make markets in these securities. For a current list of primary dealers, see https://www.bankofcanada.ca/markets/government-securities-auctions/.
Retail-Facing Dealer	Financial firm whose purpose is to offer financial services to retail (non-institutional) investors. This includes retail brokerage firms, financial planners, and other entities that provide investment advice or services to individual 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".

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)	660.997	613.824	155.859	1746.480	209.429
Bonds (in bn C$)	71.724	68.410	0.196	1,504.660	54.771
Derivatives (in k)	340.366	323.155	1.487	1,017.150	126.646
Number of Active Dealers (LEIs)					
Stocks	64	64	60	68	1.479
Bonds	59	60	5	82	7.583
Derivatives	54	54	20	56	3.068
Number of Active Dealers (Parents)					
Stocks	60	60	57	63	1.009
Bonds	56	57	11	70	6.519
Derivatives	48	48	20	51	2.715
Trade size					
Stocks	11,875.3	1,300	0.1	7.361×10^{7}	74,693.4
Bonds (in mil C$)	10.935	1.250	1×10^{-6}	7.189×10^{5}	436.178
Derivatives	78.146	10	1	170,478	667.407
Trade price					
Stocks (in C$)	27.281	12.240	0.005	2,392.4	86.937
Bonds (in C$)	102.344	100.24	0.001	980	12.398
Derivatives (in C$)	20.812	1.063	-142.050	21,800	123.305

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 (LEIs) and parent institutions (Parents), the trade size, and trade price. Note that the composition of trading firms changes when we move from the LEI to the parent-level, because an entity identified by its LEI may not be classified as a dealer, whereas its parent institution might be. For bonds, trade volume and size is measured in par-value, for stocks it's the number of shares, and for derivatives the number of contracts. 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: Stock market 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: Market intersection: : RHS: LEI-level of members; LHS—Parent level

Intersection	TSX	MTRS	MX	Intersection	TSX	MTRS	MX
All	78.57	50.13	53.31	All	89.66	83.16	64.06
TSX and MTRS	19.52	2.95	0.0	TSX and MTRS	9.59	3.03	0.0
TSX and MX	1.15	0.0	4.63	TSX and MX	0.0	0.0	0.11
MTRS and MX	0.0	2.73	12.32	MTRS and MX	0.0	3.86	10.73
MTRS only	0.0	44.2	0.0	MTRS only	0.0	9.96	0.0
TSX only	0.76	0.0	0.0	TSX only	0.75	0.0	0.0
MX only	0.0	0.0	29.74	MX only	0.0	0.0	25.09

Notes: Appendix Table A7 reports the share of trade volume in each of the three main market segments (TSX, MX, and MTRS) accounted for by dealers with different patterns of cross-market participation—for example, those active in all three segments (All), in TSX and MTRS only, and so on. Each column sums to 100%. The left-hand side (LHS) of the table shows results at the LEI level, while the right-hand side (RHS) aggregates to the parent-firm level.

Appendix Table A8: Avg. weekly share traded by dealers per type (parent-level)

Dealer type	**MTRS**	**TSX**	**TSX-IN**	**MX**	**MX-IN**
Bank	55.09	34.54	31.83	39.45	17.57
Primary Dealer	32.94	47.32	27.39	21.03	13.24
Broker-Dealer	11.24	9.23	15.07	8.40	1.69
High-Frequency Trader	0.00	5.78	23.36	29.76	67.41
Pension Fund and Insurance	0.51	0.24	0.03	0.00	0.00
Retail-Facing Dealer	0.02	1.71	1.34	0.62	0.00
Asset Manager	0.21	1.18	0.98	0.69	0.00
Mutual Fund	0.00	0.00	0.01	0.00	0.00
Other	0.00	0.00	0.00	0.68	1.12

Notes: Appendix Table A8 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 A9: Specialization and size

	market specialization		product specialization	
	(OLS)	(FE)	(OLS)	(FE)
(Intercept)	0.333***		0.167***	
	(0.011)		(0.003)	
Summed market share, $\tilde{s}_{yj}$	0.000	-0.000		
	(0.648)	(0.636)		
Summed product-market share, $\tilde{s}_{yjm}$			0.000	-0.000
			(0.097)	(0.094)
N	2,013	2,013	6,294	6,294
R^2	-0.000	0.042	-0.000	0.102
Within-R^2		-0.000		-0.000

Notes Table A9 displays the estimates from regressing dealer-specific annual market specialization scores (1) on summed market shares $\tilde{s}_{yj} = \sum_m s_{yjm}$ in column (OLS) in addition to market-year fixed effects in column (FE). The last two columns show the analogue but for product specialization scores (2) and summed product market shares, $\tilde{s}_{yjm} = \sum_p s_{yjmp}$.

Appendix Table A10: Correlation between margins and specialization scores

	Stocks		Bonds		Derivatives	
market specialization	0.051***		0.009***		2.000***	
	(0.001)		(0.003)		(0.035)	
product specialization		0.460***		0.032***		1.981***
		(0.006)		(0.004)		(0.040)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N	111,051,211	111,051,211	6,713,819	6,757,118	5,646,884	5,646,884
R^2	0.008	0.008	0.041	0.017	0.150	0.150
Within-R^2	0.000	0.000	0.000	0.000	0.002	0.002

Appendix Table A10 shows the estimation results from regressing margins (11) 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 (A11): trade size, date fixed effects and security-year-week fixed effects. We cluster standard errors at the daily-level.

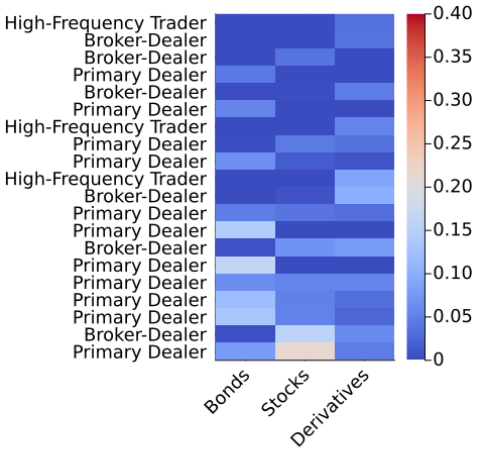
Appendix Table A11: Margin regression with dealer-types (bonds, stocks, derivatives)

Bond market	Margin	Stock market	Margin	Derivative market	Margin
Trade size	−0.000	Trade size	−0.158***	Trade size	−54.411***
	(0.000)		(0.006)		(4.924)
Counterparty is broker	+0.005***	Client account	−0.070***	Client account	−0.155**
	(0.001)		(0.002)		(0.058)
Counterparty is client	+0.014***	Inventory account	+0.006*	Inventory account	+1.179***
	(0.001)		(0.002)		(0.063)
Asset Manager	−0.074***	Market-maker account	+0.089***	Asset Manager	−0.130**
	(0.003)		(0.003)		(0.039)
Bank	−0.008	Asset Manager	−0.127***	Bank	+0.127**
	(0.018)		(0.004)		(0.047)
Broker-Dealer	+0.015***	Bank	−0.023***	Broker-Dealer	−0.012
	(0.003)		(0.004)		(0.022)
High-Frequency Trader	0.010	Broker	−0.015***	High-Frequency Trader	+0.567***
	(0.053)		(0.001)		(0.023)
Mutual Fund	+0.034	High-Frequency Trader	−0.013***	Other	+0.302
	(0.032)		(0.003)		(0.485)
Pension Fund and Insurance	+0.093***	Mutual Fund	+0.094***	Retail-Facing Dealer	−0.805***
	(0.004)		(0.019)		(0.056)
Retail-Facing Dealer	−0.223***	Pension Fund and Insurance	−0.093***		
	(0.005)		(0.003)		
		Retail-Facing Dealer	−0.025***		
			(0.002)		
Date-& symbol-year-week fes	Yes	Date-& symbol-year-week fes	Yes	Date-& symbol-year-week fes	Yes
N	6,713,819	N	111,051,211	N	5,646,884
R^2	0.044	R^2	0.008	R^2	0.151
Within-R^2	0.003	Within-R^2	0.001	Within-R^2	0.003

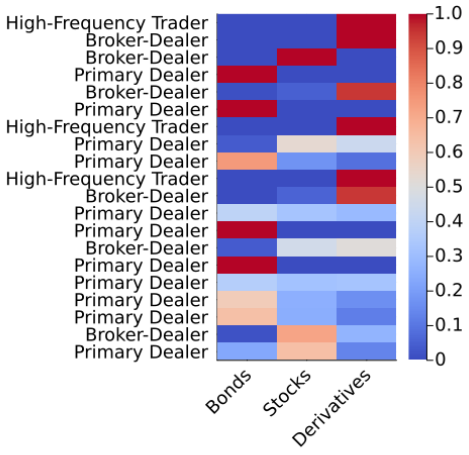
Notes: Table A11 shows the estimation results from regressing margins (11) of stock market trades on trade-size, the account-type (client, inventory, market-market, or other—the baseline), dealer-types, and day- and security-year-week fixed effects on the LHS, for fixed-income trades in the middle, and 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.

Appendix Figure A1: Dealer market shares and market specialization scores (Top 20)

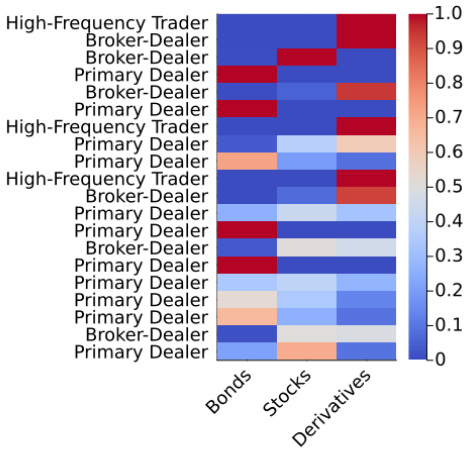
(a) Market shares



(b) Market specialization scores

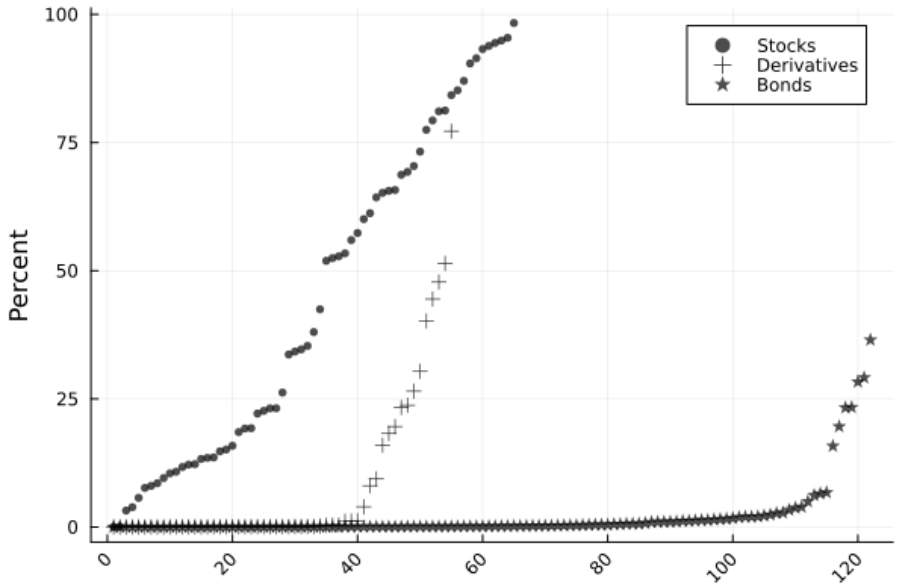


(c) Alternative market specialization scores



Notes: Appendix Figure A1a displays the average annual market, $\frac{1}{Y}\sum_y s_{yjm} = s_{jm}$, shares of the 20 dealers with the highest average annual trade volume (computed as the sum across average annual market shares, $\sum_m s_{jm}$) in each market, sorted with the largest dealer at the bottom and the smallest one at the top. Appendix Figure A1b displays the corresponding average annual market specialization scores (1). Figure A1c shows the average across years of the ratio $\bar{s}_{yjm}/\bar{s}_{yj}$, which enters the cross-market Theil index (6), normalized by M. While this ratio differs slightly from the market specialization scores, the patterns across dealers are similar, as the figure illustrates.

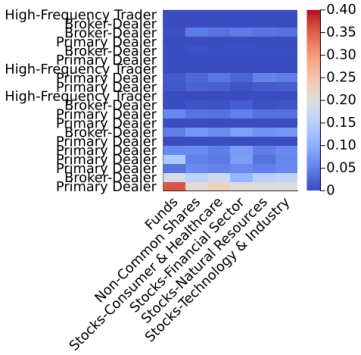
Appendix Figure A2: Fraction of securities traded by each dealer out of all securities on TSX, MX, MTRS



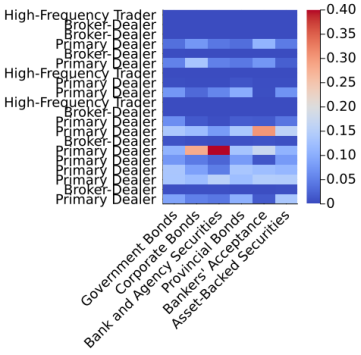
Notes: Appendix Figure A2 shows the fraction of symbols (x-axis) that each dealer trades (at the parent-level, on the y-axis) out of all traded symbols for each market. Since we sort by size within each market, the y-axis doesn't represent dealer-IDs that are common across markets. This is an alternative 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 A3: Product market shares (Top 20)

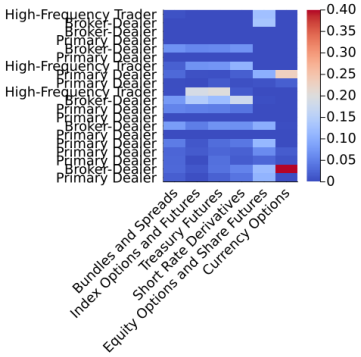
(a) Stock market



(b) Bond market

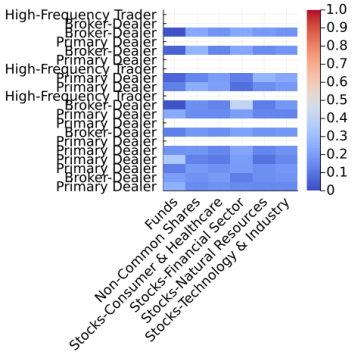


(c) Derivatives market

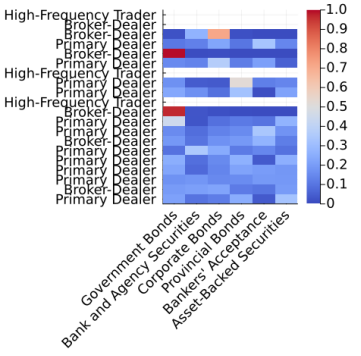


Appendix Figure A4: Product specialization scores (Top 20)

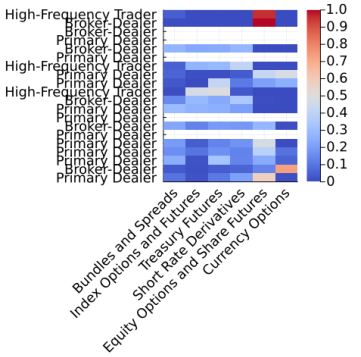
(a) Stock market



(b) Bond market



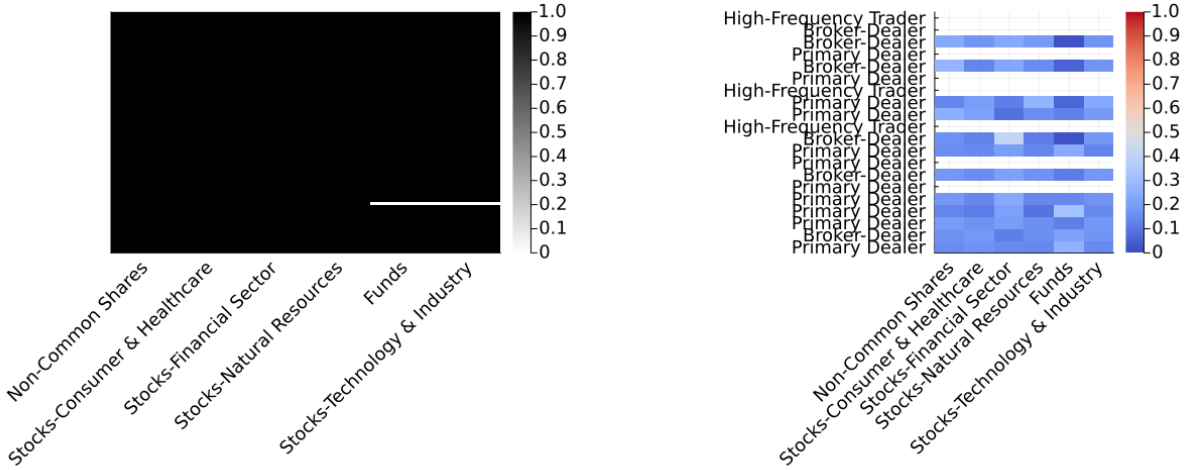
(c) Derivatives market



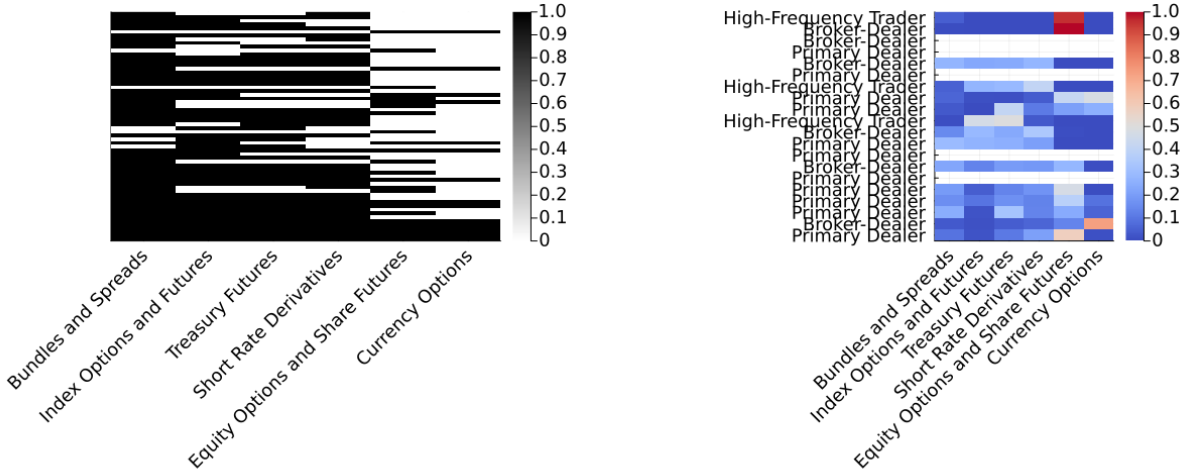
Notes: Appendix Figure A3 plots product market shares, $s_{yjmp} \in [0,1]$, averaged across years, for the 20 dealers with the highest average annual trade volume (computed as the sum across average annual market shares, $\sum_m s_{jm}$) in each market. In all graphs, a row represents a dealer; they are sorted by size with larger dealers at the bottom and smaller ones at the top. Appendix Figure A4 plots product specialization scores, $\text{specialization}_{yjmp} = s_{yjmp} / \sum_p s_{yjmp} \in [0,1]$, averaged across years, for the 20 dealers with the highest average annual trade volume (computed as the sum across average annual market shares, $\sum_m s_{jm}$) in each market. In all graphs, a row represents a dealer; they are sorted by size with larger dealers at the bottom and smaller ones at the top. When a dealer is not active in the market, the row is white.

Appendix Figure A5: Dealer presence and product specialization scores across products in the stock and derivative market excluding market-maker accounts (Top 20)

(a) Stock market



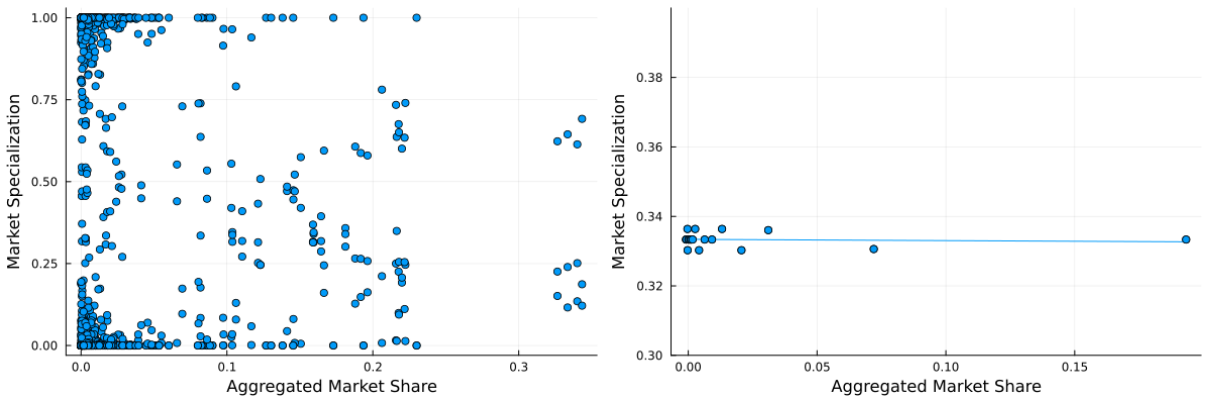
(b) Derivatives market



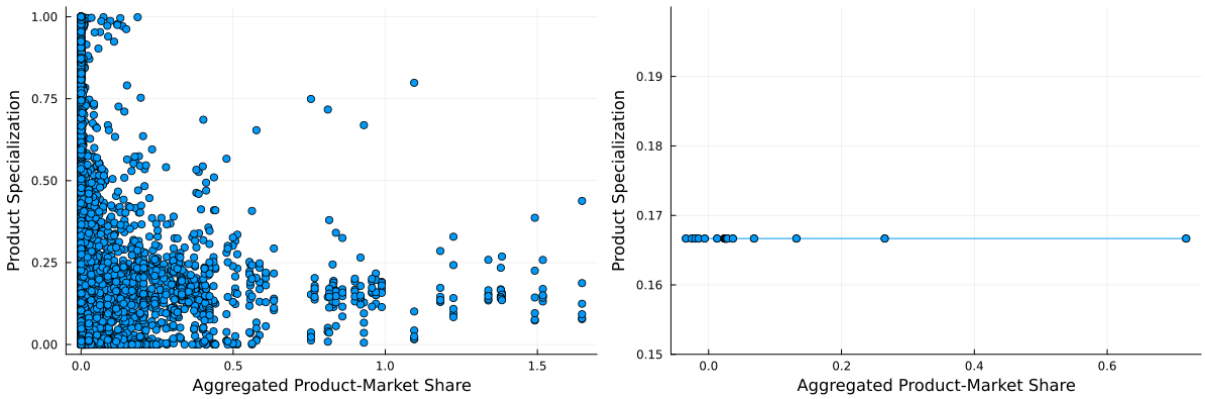
Notes: Appendix Figure A5 shows dealer presence and product specialization scores for the largest 20 dealers in the stock and derivative market, excluding market-marker accounts. On the RHS we show whether dealers trade each product at least once in black. On the LHS we see the average annual product specialization scores. 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 A6: Specialization scores and market presence (LHS scatter, RHS binned scatter)

(a) Market specialization and average market share

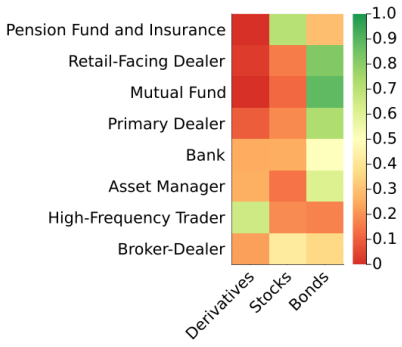


(b) Product specialization and average product-market share



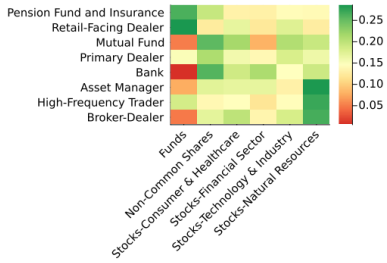
Notes The LHS of Appendix Figure A6 shows a scatter plot of annual dealer-specific market specialization scores on the y-axis versus average summed annual market shares, $\sum_m s_{yjm}$, on the x-axis, in panel (a). In panel (b) we see the correlation between product specialization cores and average summed product-market shares, $\sum_m s_{yjmp}$. On the RHS, the figure shows the respective binned scatter plot, where we control for year-market, and year-market-product fixed effects.

Appendix Figure A7: Market specialization scores across dealer types

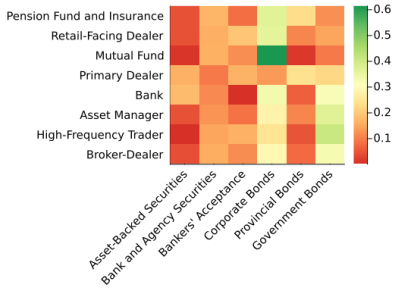


Appendix Figure A8: Average product specialization scores across dealer types

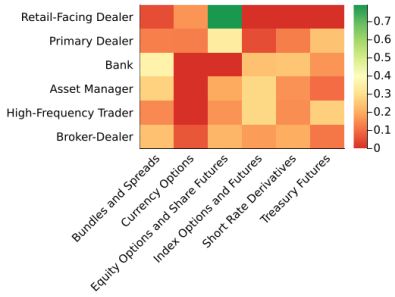
(a) Stocks



(b) Bonds



(c) Derivatives



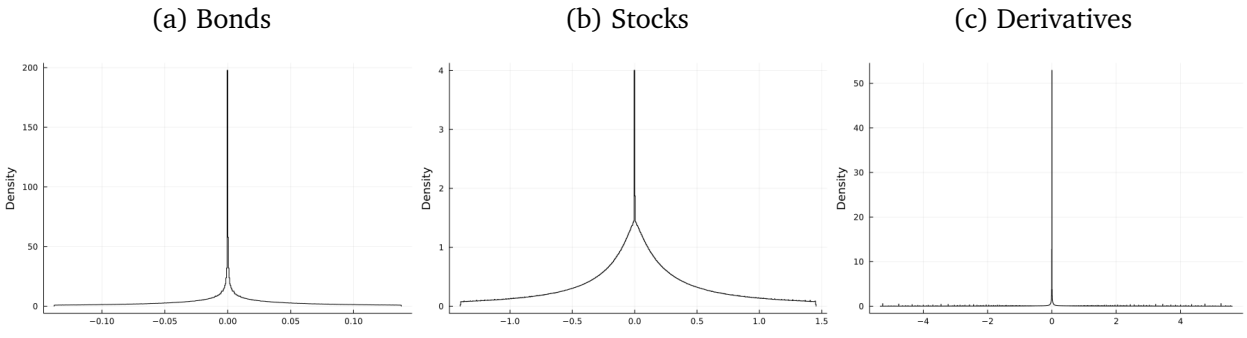
Notes: Appendix Figure A7 plots average annual market specialization scores (1) for each dealer types. Appendix Figure A8 visualizes average annual product specialization scores (2) for each dealer type and each market—stocks in (a), bonds in (b), and derivatives in (c). Note that the color scheme ranges from 0 to 1 for market specialization, but not for product specialization to allow for comparison across products within a market (subject to the overall degree of specialization in the respective market, for example, we have already established that product specialization in the stock market is low).

Appendix Figure A9: Bond product specialization scores by client-types (excluding noisy means)



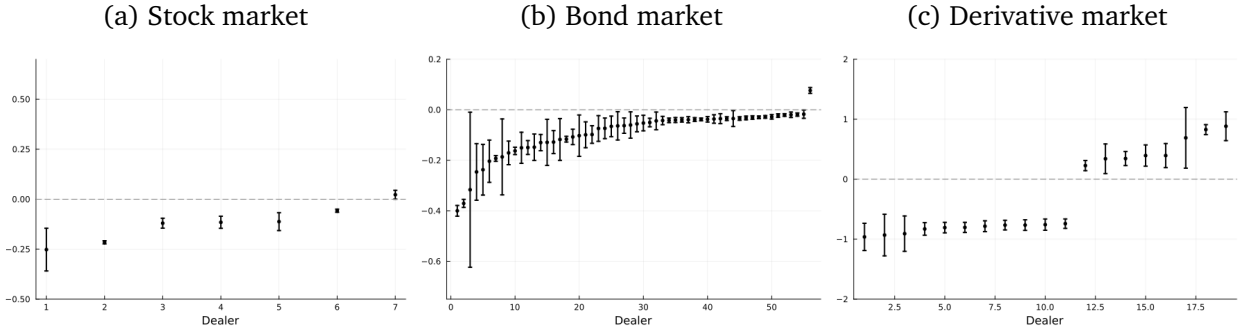
Notes: Appendix Figure A9 complements Figure 8 by excluding noisy estimates. To construct this graph, we assemble a panel of product specialization scores, product specialization$_{yjmp}$, for all dealers j, years y, and bond products p, restricted to cases where the dealer traded with client type r in the bond market (m = bonds). For each observation, we include a set of indicator variables for product–client-type pairs, marking whether dealer j traded product p with client type r in year y. This panel is unbalanced, as it conditions on observed trades. We then estimate the following regression: product specialization$_{yjmp} = \sum_{p_{\in\mathcal{P}}} \sum_{r\in\mathcal{R}} \beta_{pr} \mathbb{I}$(product $= p$ and client type $= r$) $+ \epsilon_{yjmp}$, where $\mathcal{P}$ denotes the set of bond products and $\mathcal{R}$ the set of client types. Figure 8 displays all β_{pr} estimates, while Figure A9 shows only those that are statistically significant at the 10% level.

Appendix Figure A10: Margin distribution in each market



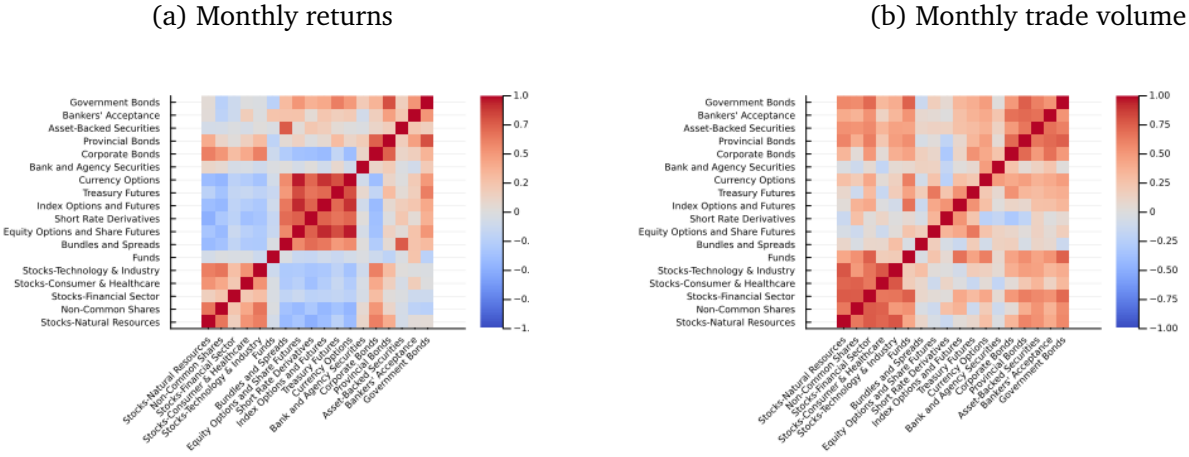
Notes: Appendix Figure A10 shows the distribution of our margin measure (11), which approximates how much less (more) a trader paid compared to the average price for a security in a day when buying (selling) for for each market. We exclude outliers, which are outside of the interquartile range.

Appendix Figure A11: Dealer coefficients for dealers who exclusively trade in one market



Notes: Appendix Figures A11a shows the dealer coefficients and 95% confidence intervals when regressing margins (11) 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 A11b shows the analogue for the fixed-income market, where we replace the account-type with a variable that indicates the type of trade. Figure A11c 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%. 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.

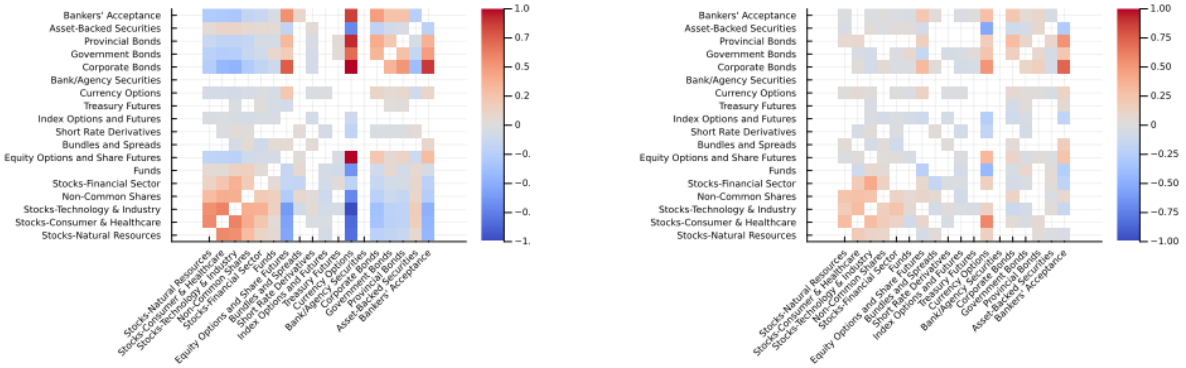
Appendix Figure A12: Correlations of monthly returns and trade volumes across products



Notes: Appendix Figure A12a shows the pairwise correlation of average monthly returns across all products in from the stock over the derivative to the bond market. Figure A12b shows analogous correlations for total monthly trade volume.

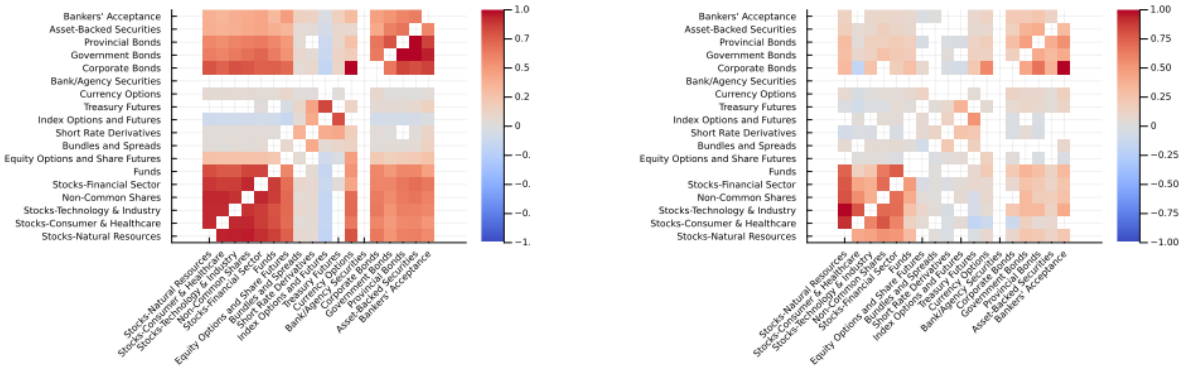
Appendix Figure A13: Cross-product correlations in standardized dealer net demand

(a) Across dealers (b) Within dealers



Appendix Figure A14: Cross-product correlations in standardized dealer gross trading volume

(a) Across dealers (b) Within dealers



Notes: Appendix Figure A13 shows the estimated coefficients β_{pq} from regressions of standardized net demand in product p on net demand in product $q \neq p$, controlling for day fixed effects (top panel) and both day and dealer fixed effects (bottom panel). Appendix Figure A14 shows the analogous graph but with (gross) trade volume instead of net demand. Both, net demand and trade volume, is standardized at the dealer-day level to allow comparability across products. We exclude coefficients that are not statistically significant at a 10% significance level.