DEUTSCHE BUNDESBANK EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 36/2026

Anatomy of a safe asset: Bund market structure and reform

Mark Kerssenfischer
(Deutsche Bundesbank)

Konrad Lucke
(Goethe University)

Willy Scherrieble
(Deutsche Bundesbank)

Discussion Papers represent the authors' personal opinions and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

Editorial Board: Daniel Foos
Stephan Jank
Thomas Kick
Martin Kliem
Malte Knüppel
Christoph Memmel
Hannah Paule-Paludkiewicz

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank, Press and Public Relations Division, at the above address or via email: www.bundesbank.de/contact

Internet http://www.bundesbank.de

Reproduction permitted only if source is stated.

DOI https://doi.org/10.71734/DP-2026-36
ISBN 978-3-98848-095-8
ISSN 2941-7503

Anatomy of a Safe Asset: Bund Market Structure and Reform

Mark Kerssenfischer
Deutsche Bundesbank

Konrad Lucke
Goethe University

Willy Scherrieble
Deutsche Bundesbank

October 2026

Abstract

This paper constructs a non-anonymous transaction-level dataset for German Bunds, the euro area's benchmark safe asset, from MiFID I and MiFID II transaction reports since 2008. We show that a concentrated dealer core intermediates this safe asset through persistent dealer–client relationships. Trading is segmented across investor sectors, with order flow fragmented across bilateral OTC and electronic venues. Identified stress flows show foreign official accounts supplying Bunds whether prices fall in a dash for cash or rise in a flight to quality. Hedge funds absorb Bunds in both episodes, while primary dealers absorb Covid selling but distribute Bunds after fiscal supply news. Hedge funds and official sector accounts carry the most permanent price-relevant flow yet pay middling effective half-spreads, while retail households pay by far the highest. Execution costs do not increase monotonically with the price content of flow. We evaluate two reforms targeting this structure: mandatory central clearing would cut dealers' net settlement obligations in Bunds by 35%, less than comparable estimates for U.S. Treasuries and U.K. Gilts. In the benchmark structural calibration, platform centralization lowers welfare because matching gains from broader dealer access do not offset lost relationship surplus under imperfect dealer competition.

Keywords: safe assets, sovereign bonds, German Bunds, market microstructure, OTC markets, dealer intermediation, electronic trading platforms, central clearing, MiFID II transaction reports.

JEL classification: G12, G14, G18, G23

Contact address: mark.kerssenfischer@bundesbank.de and willy.scherrieble@bundesbank.de. We thank Jean-Charles Wijnandts (discussant), Katia Vozian (discussant), Thomas Perreten (discussant), Tomy Lee, Loriana Pelizzon, Marti G. Subrahmanyam, an anonymous referee, and participants at the 14th Bundesbank Term Structure Workshop, the 2026 ESMA Research Conference, the 32nd Annual Meeting of the German Finance Association (DGF), and CREDIT 2026, for valuable feedback. The views expressed in this paper are those of the authors and do not necessarily coincide with the views of the Deutsche Bundesbank or the Eurosystem. We declare that we have no relevant or material financial interests that relate to the research described in this paper. A supplementary online appendix is available at https://www.bundesbank.de/dp-en/2026-36-appendix. Appendices A–D, with the figures and tables numbered A1, A2, and so on, are in the online appendix.

1 Introduction

Repeated stress episodes have put the design, regulation, and reform of sovereign bond markets back at the center of policy debate.[1] Yet despite the central role of sovereign bonds as safe assets (Gorton, 2017; Caballero, Farhi, and Gourinchas, 2017), their market microstructure remains opaque. To diagnose stress and prescribe reform reliably, policymakers need an anatomy of the market. We build that anatomy from what is, to our knowledge, the most granular and comprehensive transaction-level record of any major sovereign bond market. Covering nearly two decades of German Bund transactions, the record shows who meets whom, through which venues, and at what prices. We use this microdata to document market structure and evaluate two major reforms as counterfactuals.

We first map the market structure. A tight primary dealer core intermediates most activity and connects a broad client periphery; direct client-to-client trading is rare. Clients reach that core through narrow, persistent relationships: most bilateral clients use a single dealer, and electronic platforms broaden access without delivering all-to-all trading. The periphery is segmented across venue types: investor sectors reach the market in systematically different ways. This segmentation lines up with information. Hedge funds and the official sector have the most price-relevant flow and trade mostly on electronic platforms. The venue gradient is not just composition: the same hedge fund flow carries more permanent price impact on multilateral platforms than bilaterally. Trading is fragmented across venues and concentrated in a small set of bonds, especially on-the-run and futures-deliverable issues, while most outstanding bonds trade only sporadically on a given venue. The market is therefore hierarchical, relationship-based, and segmented.

Identifying both counterparties shows directly which sectors supply and absorb Bunds in stress. In the 2020 dash for cash and the 2023 flight to quality, despite opposite price moves, foreign official accounts are the dominant net sellers and hedge funds buy, while primary dealers absorb material flow only in the cash shock. After the 2025 fiscal supply news, primary dealers distribute Bunds rather than warehousing them. In normal times, hedge funds and official sector accounts carry the most permanent price-relevant flow, but they do not pay the highest effective half-spreads. Retail households pay by far the most, and, within a given investor, larger trades pay more. The effective half-spread is therefore an execution cost that does not map monotonically to adverse selection.

Both reform counterfactuals target the dealer core, but at different stages. Central clearing keeps observed trades fixed and changes how dealer obligations net after execution. Netting each dealer's own trades at a clearing house would cut dealers' net settlement obligations by 35%, less than comparable estimates for U.S. Treasuries and U.K. Gilts (Fleming and Keane, 2021; Baranova, Holbrook, MacDonald, Rawstorne, Vause, and Waddington, 2023). Platform cen-

[1] Duffie (2025) provides a recent overview. Current reform proposals in the United States include mandatory central clearing (SEC) and capital regulation affecting dealer intermediation capacity (Federal Reserve). Reforms in the EU include faster settlement from $T+2$ to $T+1$ (ESMA), a consolidated tape (ESMA), and venue consolidation (Reuters). The Savings and Investments Union adds market integration to the EU agenda (European Commission, COM(2025) 124, 19 March 2025).

tralization changes the matching problem before execution: investors may reach more dealers, dealers may compete more aggressively, and bilateral relationship value may be lost. Applying the dealer–investor framework of Allen and Wittwer (2023) to Bunds, we find that forced platform centralization lowers welfare by 7.73% relative to the current access regime. Matching gains from broader dealer access do not offset lost relationship surplus under imperfect dealer competition. The sign matches the Canadian estimates (a loss of 1.75%), and in the benchmark calibration the losses are larger. We import only the demand elasticity, which the Bund data cannot identify (Appendix A.6). The loss relative to a fully bilateral market persists across its entire confidence interval.

2 Literature

The paper relates to four aspects of decentralized fixed income markets: dealer cores, relationships, venue segmentation, and market design. The facts show where existing models fit the Bund market and where they do not; no single structural model matches all of them. The first aspect is the dealer core. Bond markets remain more decentralized and relationship-based than equity markets (Bessembinder, Spatt, and Venkataraman, 2020; Biais and Green, 2019), and sovereign bond price formation depends on liquidity and intermediation (Fleming and Remolona, 1999; Pasquariello and Vega, 2007; Cheung, Rindi, and De Jong, 2005; Mizrach and Neely, 2008). A concentrated intermediary layer can improve inventory offsetting and execution capacity, but it can also concentrate market power and balance sheet obligations, as in core–periphery and dealer network theories (Wang, 2016; Li and Schürhoff, 2019; Colliard, Foucault, and Hoffmann, 2021) and related work on dealer specialization, primary market participation, and balance sheet provision (Wittwer and Uthemann, 2025; Allen, Hortaçsu, Richert, and Wittwer, 2026; Huh and Infante, 2021; Duffie, 2023). Our facts bear on both sides: the tight hierarchy and the rarity of direct client-to-client trading measure how concentrated intermediation has become, while the stress episodes show the core absorbing cash shock selling in 2020, taking little net flow in the 2023 safety shock, and distributing Bunds after the 2025 fiscal news; relative to trade-level safe asset stress work on U.K. Gilts (Pintér, 2023; Pintér, Siriwardane, and Walker, 2026), the identified Bund flows span episodes of opposite price sign and a domestic supply contrast (Duffie, 2025). They do not identify whether the Bund core is too thin or too broad. The safe asset literature asks why investors demand safety and which assets supply it (Gorton, 2017; Caballero et al., 2017; He, Krishnamurthy, and Milbradt, 2019). We take the safe asset status of Bunds as given and document the market through which that demand is met: who buys and sells in safety shocks, and what they pay to trade.

The second aspect is relationship trading. Search and relationship models show why clients use a small set of dealers, why platforms need not create all-to-all competition, and why a dealer connection can carry non-price value (Allen and Wittwer, 2023; Wang, 2023, 2022; Allen and Wittwer, 2025; Di Maggio, Kermani, and Song, 2017; Kondor and Pintér, 2022). Our relationship facts show that Bund clients broaden access on platforms but remain far below random matching, which is the market structure input for the centralization exercise.

The third aspect is venue segmentation and information. Trading protocols allocate order flow across venues, and the identity of that flow matters for price discovery (Hasbrouck, 1995; Menkveld, Yueshen, and Zhu, 2017; Dugast, Üslü, and Weill, 2022; Lee and Wang, 2025; Lee, Wang, and Zawadowski, 2025; Hagströmer and Menkveld, 2019; Hagströmer, 2022; Hagströmer and Menkveld, 2023). Outages provide causal evidence in the same market: when the Bund futures market halts, cash market liquidity evaporates, dealers withdraw, and transaction prices deviate from fundamentals (Kerssenfischer and Helmus, 2025)—time-series evidence that market functioning depends on specific venues, for which our venue-localized price discovery facts are the cross-sectional counterpart. Identified dealer-to-client data also let us separate whose flow has permanent price content from who pays the highest effective half-spread (Czech, Huang, Lou, and Wang, 2021; Kondor and Pintér, 2022; Pasquariello and Vega, 2007; Pintér, Wang, and Zou, 2024, 2022; Lee, Nathan, and Wang, 2024). That separation is central here: the most price-relevant flow is not the flow charged the highest execution cost.

The fourth aspect is market design. Central clearing, platform centralization, and all-to-all trading are distinct reforms. Clearing changes post-trade netting (Fleming and Keane, 2021; Baranova et al., 2023; Duffie and Wang, 2017; Duffie, Singh, and Wang, 2026). Platform centralization changes matching, competition, and relationship value; in that structural model, markups and spreads are transfers, so welfare turns on gains from trade, not on price concessions alone (Allen and Wittwer, 2023; Kutai, Nathan, and Wittwer, 2025; Chaboud, Correia Golay, Cox, Fleming, Huh, Keane, Lee, Schwarz, Vega, and Windover, 2025; Lee and Wang, 2025; Duffie, 2023). We therefore treat clearing as a mechanical settlement counterfactual and platform centralization as a structural welfare counterfactual.[2]

3 Data and Measurement

We assemble a new, non-anonymous transaction-level dataset for the German sovereign bond market. As Figure 1 summarizes, we distill roughly 19 million raw regulatory reports from two reporting regimes into approximately 16.4 million Bund trade records spanning 2008–2025. We then match each transaction—on ISIN, LEI, MIC, and timestamp—to bond, investor, venue, and price benchmarks from more than a dozen external reference datasets. Section D provides details.

The underlying transaction reports are confidential regulatory microdata and cannot be posted publicly. We will maintain the final research dataset and its documentation as a reusable resource for Bundesbank researchers whose projects are authorized to use the underlying data.

3.1 Transaction-Level

Our transaction-level dataset for German sovereign bonds ("Bunds") combines two regulatory reporting regimes: MiFID I (2008–2017) and MiFID II/MiFIR transaction reports (2018–2025).

[2]The same transaction data support a companion paper tracing an on-the-run discount in Bunds to the Finanzagentur's tapping policy (Kerssenfischer, 2026). We keep that pricing question separate and focus here on market structure and trading costs.

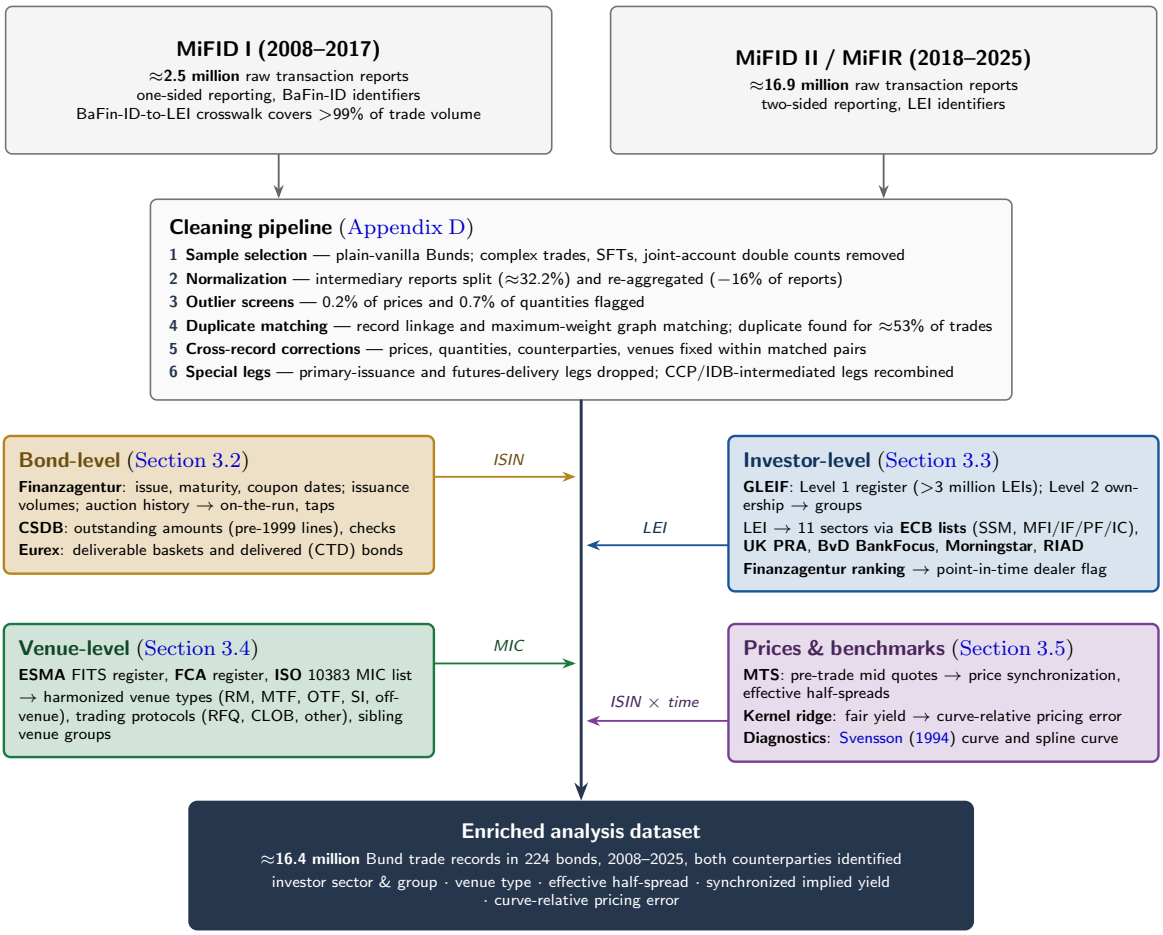


Figure 1: Construction of the analysis dataset. Regulatory transaction reports from the two reporting regimes (top) pass through the cleaning pipeline (middle). The cleaned records are enriched with bond-, investor-, venue-, and price-level reference data (side panels), each merged on the indicated key (bottom). Percentages give the share of reports or trades affected by each cleaning step; references point to the relevant sections and Appendix material.

Together they cover approximately 16.4 million trade records; the unit of the cleaned file is the trade record as reported, and each record has two legs, its buyer side and its seller side. One execution can appear as several records. An agency or matched principal report, in which the reporting firm is neither buyer nor seller, becomes two records, one against each side, and on an anonymous platform each member files its own record against the venue. We link the records of each execution (Section D.8). Statistics of market size—volume, trade counts and sizes, venue shares, and bond and venue concentration—count each execution once; summed over records, average daily MiFID II volume would be €24 bn instead of €20 bn. Statistics that need a buyer or a seller—sectors, dealers, counterparties, net flows, and the price regressions—use records, because the records of one execution name different buyers and sellers, such as an agent and its client. In those analyses, a trade is a record. The reports are regulatory records, not research data, so we normalize reporting formats, aggregate mechanically related reports, screen implausible prices and quantities, and identify duplicates using institutional features of each regime. The full pipeline and diagnostics are in Appendices D.1–D.8.

MiFID II/MiFIR is the core market structure sample. Since 3 January 2018, investment firms have reported German sovereign bond transactions to BaFin irrespective of trading venue, with LEI-based buyer, seller, and executing-firm identifiers. Each report names the buyer and the seller, so we observe both counterparties of bilateral trades, and two-sided reporting lets us match their duplicate reports. On anonymous platforms each member reports only its own side; we link the two reports where both are in the data (Section D.8). MiFID I provides a longer historical window but has more limited client-side coverage because reporting was predominantly one-sided and used BaFin institutional codes rather than LEIs. We resolve those codes to LEIs where possible, but the market structure analysis uses the MiFID II subsample unless otherwise noted.

We cross-check the regulatory reports against commercial transaction records from MTS and other trading venues. The comparisons cover prices, quantities, timestamps, and execution counts within overlapping samples. Appendix D.9 describes the sources and method.

3.2 Bond-Level

We obtain bond-level information from the Finanzagentur: issue date, maturity date, first coupon date, dated date (start of accrual), and gross outstanding volume (cumulative auction and syndication volume, including the Finanzagentur's own-account quota).[3]

Each day, we identify whether each outstanding bond is deliverable into the front contract of the 2-year, 5-year, 10-year, or 30-year Eurex Bund futures, applying the contract eligibility rules detailed in Section D. The cheapest-to-deliver (CTD) bond of each contract is the bond most delivered into it, according to Eurex statistics on notified bonds (Bloomberg data before

[3]Finanzagentur bond metadata are more comprehensive and reliable than the Centralised Securities Database (CSDB). Hence, we use the CSDB to validate monthly outstanding volumes of each bond and, together with the Finanzagentur's year-end statistics, to supply outstanding volumes for bonds issued before the auction history begins in 1999.

2009). The on-the-run bond in each maturity segment is the most recent new issue in the Finanzagentur's auction and syndication history, which also identifies tap reissuances.

3.3 Investor-Level

Our investor reference data cover over three million Legal Entity Identifiers (LEIs) from the Global Legal Entity Identifier Foundation (GLEIF).[4] We map LEIs into eleven sectors: primary dealers (participants in the primary auctions of Bunds); banks (commercial, investment, savings, and cooperative banks); interdealer brokers (IDB; voice and electronic brokers intermediating between dealers); investment funds (UCITS, money market funds, ETFs, and asset managers); hedge funds, including principal trading firms (PTFs)[5]; insurance companies and pension funds (ICPF); central counterparties (CCPs); non-bank financial intermediaries (NBFI; e.g. execution intermediaries, trading venue operators, and payment service providers); official sector entities (e.g. central banks and sovereign wealth funds); non-financial corporations (NFC); and households (individual retail investors).[6] Primary issuance defines the dealer core institutionally. The Finanzagentur sells Bunds in multiple-price auctions in which only members of the Bund Issues Auction Group (Bietergruppe) may bid. Membership requires winning at least 0.05% of the duration-weighted amounts allotted in a calendar year; members that fall short are excluded, and the Finanzagentur publishes a semiannual ranking of members.[7] We distinguish dealer-to-dealer, dealer-to-client, and client-to-client trades using a point-in-time dealer flag constructed from Finanzagentur's semiannual auction group rankings and the entry and exit information reported in their footnotes.[8] This dealer definition is intentionally institutional rather than activity-based: some large banks and market makers outside the auction group are classified as clients, not dealers. The secondary market network and concentration facts therefore validate the dealer core in the data rather than defining it by trading volume ex post. Because the auction member is often one entity of a larger banking group, we define dealers for investor groups under "Dealers and Clients".

Appendix D.10 details the mapping, which combines multiple regulatory, commercial, and manually curated sources. To ensure consistent country codes, we use the GLEIF register to map LEIs to countries. The resulting investor reference dataset has near-complete coverage for the Bund market.

Counterparty Coverage by Regime MiFID I usually identifies reporting firms and regulated counterparties, not end clients. MiFID II names buyers and sellers directly. On the client

[4] https://www.gleif.org/en/lei-data/access-and-use-lei-data/level-1-data-who-is-who

[5] Both groups trade in arbitrage, price discovery, and liquidity provision (Aragon and Strahan, 2012; Menkveld, 2013). Hedge funds manage delegated capital; PTFs trade their own balance sheets. The pooled category describes a business type, not a common strategy or funding mechanism.

[6] Natural persons have no retained individual identifier and cannot be linked across trades. They enter sector-level statistics but not group-level counterparty measures. They appear on 32% of MiFID II records (13.1% outside client allocations), but account for 0.16% of notional; the median ticket is €1,441.

[7] See the auction rules.

[8] For example, see Finanzagentur.

side of dealer-to-client records, the share of volume with a sector rises from 17% in 2013–2017 to 62% in 2018–2022 (Table A14). Our sector-level analyses therefore use MiFID II.

Investor Grouping We consolidate LEIs to the group level, using GLEIF ownership relationships augmented with ECB registers, to avoid treating subsidiaries, branches, and parent entities of the same institution as independent investors. We consolidate only operating companies and keep fund vehicles separate, so each fund remains its own group. Consolidating funds would collapse the many separately identified vehicles advised by a single manager into a few mega-groups that mechanically dominate any concentration, dispersion, or network statistic, erasing the cross-fund variation the fund-level reports are designed to capture. The rule keeps investment funds and hedge funds at the entity level (99% are singleton groups), while banks, insurers, interdealer brokers, and primary dealers are consolidated wherever ownership links exist. Asset managers keep the investment fund sector whoever owns them, because they trade as agents for the funds and mandates they manage.

The transaction data are then merged to both entity-level and group-level sector and country identifiers; unless otherwise noted, the analyses use group-level identities. Appendix D.10 reports the data sources, hierarchy rules, and validation checks.

Dealers and Clients Each investor thus carries three labels: the sector of its own LEI, its investor group, and the sector of that group. Apart from asset managers, a group takes the sector of its head entity. A group with an auction member at any point since 2018 is in the primary dealer sector (Appendix D.10), except for its asset management and insurance subsidiaries, which keep their own sectors. The remaining entities of a primary dealer group form its dealer desk: the auction member and its sister trading entities, for example the London entity of a group whose auction member is its EU bank.

We use two dealer definitions. The primary dealer sector is time-invariant and describes what an institution is. The dealer flag is point-in-time and describes whether it acts as an auction group dealer on the day: an entity is a dealer on a trade date if it belongs to the dealer desk of a group with an auction member on that date. Asset management and insurance arms are never flagged; they are clients in their own sectors, also when they trade with their parent. A trade record is dealer-to-dealer, dealer-to-client, or client-to-client by the number of its two legs that carry the flag. A leg without an identified account owner is never flagged, so a dealer's one-sided record of an anonymous platform trade is a dealer-to-client record without an identified client. The client-level analyses—information, execution costs, and the platform and home dealer tests—use the flag and keep records with an identified client outside the primary dealer sector; clients in that sector are dealer groups trading outside their membership windows. The venue type test also reports all dealer-to-client records (Appendix A.2). Statistics on the dealer core as a set of institutions—its concentration, its network, the clearing exercise, and the welfare model—use the primary dealer sector.

Records between two entities of one investor group are intragroup bookings, almost all of them between the entities of one dealer group, such as its EU bank and its U.K. broker-dealer;

they carry 26% of MiFID II record volume. The network and dealer concentration statistics exclude them; the market size and venue statistics include them. Between different groups, dealer-to-dealer records carry 4–7% of record volume a year.[9]

3.4 Venue-Level

We classify trading venues by combining three external reference datasets—the ESMA Financial Instruments Transparency System (FITS) register, the U.K. Financial Conduct Authority (FCA) register, and the ISO 10383 list of Market Identifier Codes (MICs)—each providing venue identifiers (MICs), legal entity identifiers (LEIs), names, and venue types. We harmonize venue types to the ESMA taxonomy of multilateral trading facilities (MTFs), organized trading facilities (OTFs), systematic internalizers (SIs), and regulated markets (RMs), applying the sources hierarchically (ESMA, then FCA, then ISO), which covers all MICs present in MiFIR; Appendix D.11 details the sources, venue group construction, and residual classification of the few venues without an explicit type. Each MIC also receives a protocol family: request-for-quote (RFQ), central limit order book (CLOB), other, or unknown. The family describes the trading model of a venue or segment, not how an individual trade executed; Appendix D.11 states the rules and the coverage of each label.

Venue Grouping Due primarily to Brexit, several major venue operators maintain parallel trading venues in the European Union and the United Kingdom, resulting in distinct MIC codes for what is effectively the same trading infrastructure (for example, Bloomberg, Tradeweb, MarketAxess). These "sibling venues" largely reflect regulatory and reporting requirements, not economically meaningful differences in market structure, and are typically indistinguishable from the perspective of market participants. Hence, we merge MICs of the same venue type that belong to the same operator into a common venue group.

3.5 Price Synchronization and Fair Value

We synchronize transaction prices to the Bundesbank's 11:45 a.m. fixing using changes in MTS mid-quotes, following Duffie, Fleming, Keane, Nelson, Shachar, and Van Tassel (2023). Quote-quality checks and separately flagged fallbacks address unreliable or missing quotes; effective half-spreads always require a pre-trade quote for the same bond. We convert synchronized prices into yields using the settlement cycle and the ICMA coupon convention. Individual-price screening is centered on the bond's own daily yield, with a separate guard against gross errors relative to the Bundesbank curve. This preserves bond-specific yield deviations for analysis. Appendix B details synchronization, sample selection, and curve estimation.

[9] This record share understates interdealer trading. On an anonymous broker or order book venue, each member files a one-sided record against the venue, and from 2021 a U.K. member's record is outside the data. Executions with dealers of two groups identified carry 9% of MiFID II execution volume in 2018–2020 and 7% from 2021. Adding executions in which one dealer group faces an interdealer broker or order book venue and no other party is identified raises these shares to at most 29% and 23%.

We distinguish three pricing measures. The effective half-spread is the deviation of the transaction price from the pre-trade MTS mid, signed from the client's side and reported in price basis points. The synchronized implied yield is the yield-to-maturity implied by the synchronized price, reported in yield basis points. The curve-relative pricing error is the synchronized implied yield minus the kernel ridge fair yield; positive values indicate that the bond trades cheap relative to the kernel ridge curve.

Our baseline fair value benchmark is the kernel ridge Bund yield curve, estimated day by day from each bond's volume-weighted synchronized price following Filipović, Pelger, and Ye (2025); Appendix B gives implementation details. We also estimate the Svensson (1994) curve published by the Bundesbank and a spline curve for diagnostics, including the yield curve fitting error comparison in Figure A12. The execution cost analyses use effective half-spreads, the price dispersion analysis uses synchronized implied yields directly, and the informed flow and venue price discovery tests use changes in curve-relative pricing errors. The platform centralization model also uses the kernel ridge curve as the fair value reference for markups. The clearing exercises are separate and measured in notional euros and settlement obligations.

4 Market Structure

4.1 Dealer Core and Relationships

The central market structure fact is hierarchy. Figure 2 plots the Bund trading network in the MiFID II subsample.[10] A small set of primary dealers sits at the center, intermediating most activity and connecting a broad client periphery. Direct trading between two identified clients is rare: it carries at most 4.0% of MiFID II execution volume in 2018–2020 and 1.6% from 2021. The dealer core is therefore not a label imposed by the primary market roster; it is visible in secondary market transactions.

A concentrated dealer core The network is not only dealer-centered but highly concentrated within the dealer group itself. Across the 38 primary dealers active in our sample, intermediation is far from evenly distributed: the five most active dealers account for 44% of all dealer-intermediated volume, and the top ten for 71%. The auction design sets who belongs to this group, not how concentrated intermediation within it is: the rules impose a minimum allotment share, and the same core emerges from secondary market behavior alone (Table 2). Figure 3 plots the disclosure-safe concentration curve from rank five onward and, as a benchmark, the analogous curve for the client sector. Concentration is far lower among clients: the top twenty clients account for only 30% of client volume, against 94% for the top twenty dealers. This contrast confirms that the concentration is specific to the dealer core, not a feature of the market as a whole. The effective number of intermediating dealers—the inverse of the activity

[10] Throughout the paper, statistics based on protected investor entities are reported only when the underlying cell contains at least five protected entities and the two largest protected entities account for no more than 85% of the reported aggregate. Cells failing these publication restrictions are suppressed or aggregated; exact individual investor statistics are never reported.

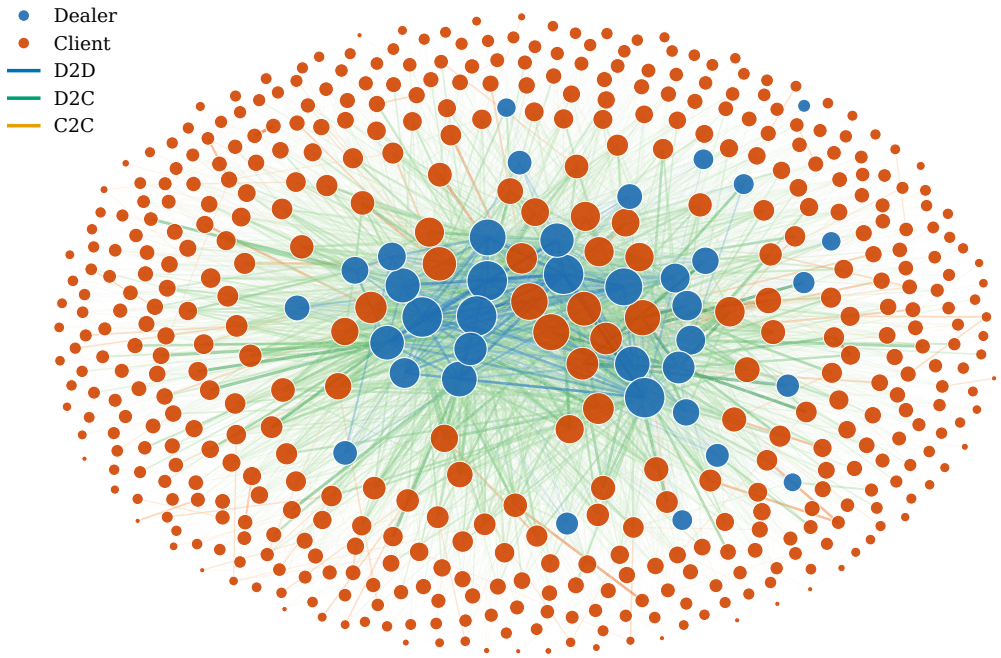


Figure 2: The Bund market network. Undirected graph of investor groups in the MiFID II subsample, restricted to relationships with at least 50 trades per edge. The plotted component is the largest connected component after selecting the top 500 groups by graph degree, with ties broken by weighted degree and group LEI. Node position reflects eigenvector centrality. Node size scales with centrality, with the largest nodes top-coded at a common maximum. Edges scale with the log number of trades between each investor pair. Intragroup trades are omitted. Dealers are investor groups in the primary dealer sector; all other investors are clients.

Herfindahl, $10{,}000/611 \approx 16$—is small, even though primary dealers as a group stand on at least one side of 84% of all volume. This quantifies the visual hierarchy in Figure 2: the Bund market rests on a thin layer of core intermediaries. The theory interpretation is a competition-versus-inventory tradeoff: a concentrated core can improve inventory offsetting and execution capacity, but it can also concentrate market power and settlement obligations (Li and Schürhoff, 2019; Colliard et al., 2021). The same concentration also sizes the central clearing exercise in Section 6.1: multilateral netting compresses settlement obligations precisely because so few dealers stand on so many trades. Section 4.2 sharpens this picture: the thin layer coincides with the official bidder group itself, and its boundary, internal hierarchy, distribution role, clientele, and stress behavior can be characterized directly.

Sticky dealer–client relationships The periphery connects to the core through narrow, persistent relationships rather than broad search. Figure 4 compares bilateral OTC and MTF trading for active clients.[11] OTC clients usually trade with few dealers and concentrate a large share of volume with a home dealer. Platforms broaden access, but they do not turn the market into all-to-all trading. The persistence diagnostics in Table 1 show that this is not just a small-sample count artifact: home dealer identities are stable across active years, and within-year dealer breadth remains below a trade-count-matched random-matching null. These

[11] We require at least ten trades per client and venue type: a one-trade client necessarily has one dealer. Without this filter, 86% of OTC clients use one dealer. The filter retains the vast majority of traded volume.

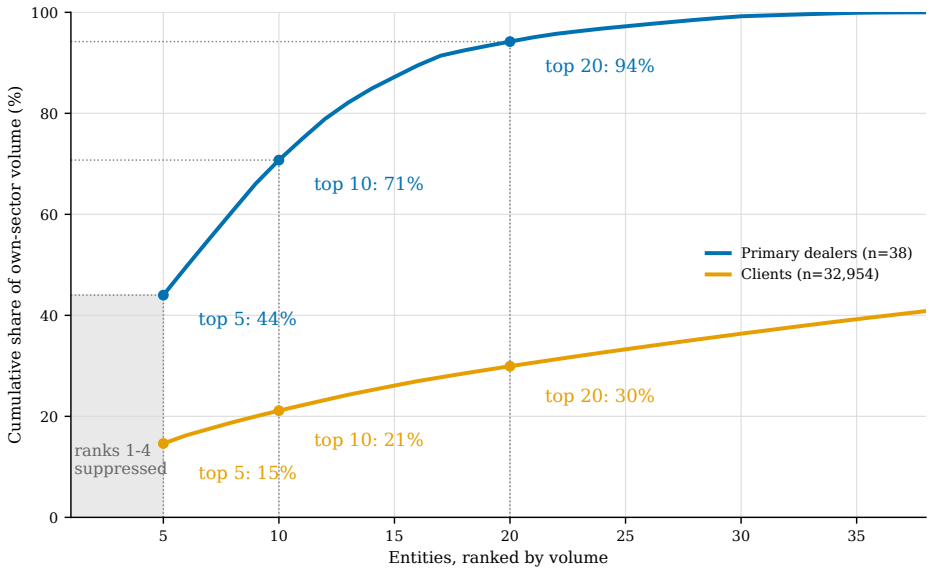


Figure 3: Concentration of dealer intermediation, with the client sector as a benchmark. For each sector, the cumulative share of own-sector volume is plotted against entities ranked from most to least active (blue: primary dealers; orange: clients). A dealer's intermediation volume sums the notional of all trades in which it participates on either side; a client's volume sums its non-dealer legs. Intragroup trades are excluded. The plotted curves start at the top-five aggregate; the shaded left region marks unplotted ranks. The horizontal axis spans the 38 active dealers. Blue markers highlight the top-5, top-10, and top-20 dealer shares; the orange marker the top-20 client share.

persistent relationships support the relevance of the relationship value mechanism in the venue centralization analysis of Section 6.2, but do not by themselves identify the relationships' welfare value.

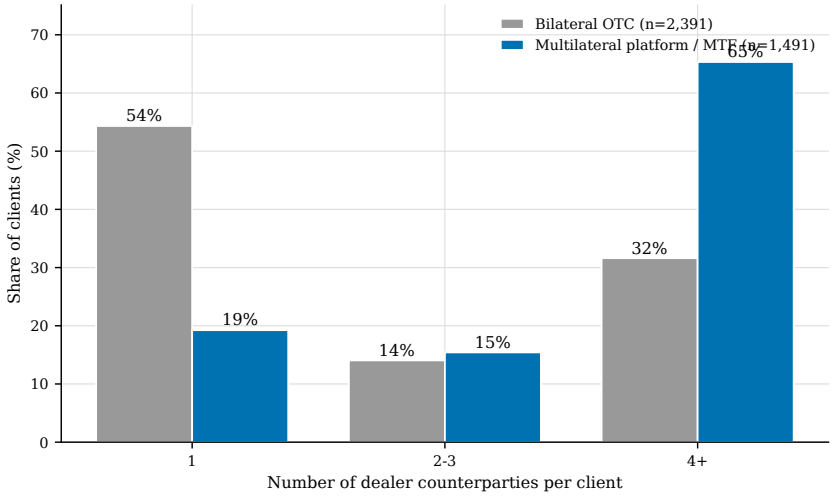


Figure 4: Concentration of dealer–client relationships. Distribution of the number of distinct dealer counterparties per client in dealer-to-client trading, comparing bilateral OTC with MTF execution; bins are 1, 2–3, and 4+ dealers. Computed at the investor-group level on the MiFID II subsample, using all MTF venues and all dealers, and restricted to clients with at least ten trades in the venue.

	Bilateral OTC	Platform (MTF)
Distinct dealers per client-year, observed	2.5	6.2
expected, trade-count-matched null	3.2	7.8
ratio (observed / null)	0.78	0.80
Same home dealer in consecutive years (%)	68	42
Active clients ($\geq$ 10 trades)	2391	1491

Table 1: Relationship persistence over time. Active dealer-to-client clients (investor groups with at least ten trades of that venue type), MiFID II. The trade-count-matched null is the expected number of distinct dealers a client would touch in a year if its lifetime trades were randomly assigned across its active years, holding each year's trade count and the client's lifetime dealer mix fixed; a ratio below one indicates within-year concentration beyond the mechanical shorter-window effect. "Same home dealer" is the share of consecutive active-year pairs in which the client's largest-volume dealer is unchanged—an identity measure free of any count dependence.

Bigger clients use more dealers—but far fewer than chance How many dealers does a client use, and does that number grow with trading footprint? Figure 5a answers yes, but slowly. Most institutional clients use one dealer; the busiest clients use more. This is the same direction as in the Gilt market evidence of Kondor and Pintér (2022): larger and more active clients are better connected.

The relevant benchmark is not zero. A client with more trades mechanically has more chances to meet more dealers. We therefore compare observed breadth with a random-matching null in which each of the client's trades draws a dealer in proportion to dealer activity. Real clients fall far below that null across the activity distribution. They are not spreading trades across the dealer core; they are adding relationships only gradually. The growth that remains is economically meaningful: among clients with the same trade count, larger tickets are associated with more dealer relationships, which a count-only random model cannot generate.

Dealer breadth is descriptive, not causal. Bund clients are sticky but not fixed: larger footprints buy some additional access while remaining far from random matching. That is the relationship stickiness the platform centralization model needs. Nor is breadth a sophistication index: Figure A15 in Appendix C shows that the cross-sectional gradient of flow informativeness in dealer breadth is mostly sector composition—the share of hedge funds and official sector accounts rises with breadth—and that, once those sectors are excluded, only the most connected clients retain a smaller, significant impact. The Appendix also implements the within-client design of Kondor and Pintér (2022)—time variation in a client's connections as an information proxy—and finds no effect in Bunds in that design either, precisely enough estimated to rule out the Gilt market magnitude. Dealer breadth marks who a client is, not how informed its trading currently is.

Appendix C gives complementary market structure evidence. Bund trading is predominantly cross-border—domestic German counterparties account for only 10% of trade-leg volume—and the Appendix maps where the counterparties that trade Bunds are domiciled. It also reports dealer trading capacity and time-series diagnostics on dealer concentration, Brexit-related venue coverage, retail participation, and algorithmic execution.

(a) Dealer breadth by trading footprint.

(b) Dealer breadth by average ticket size.

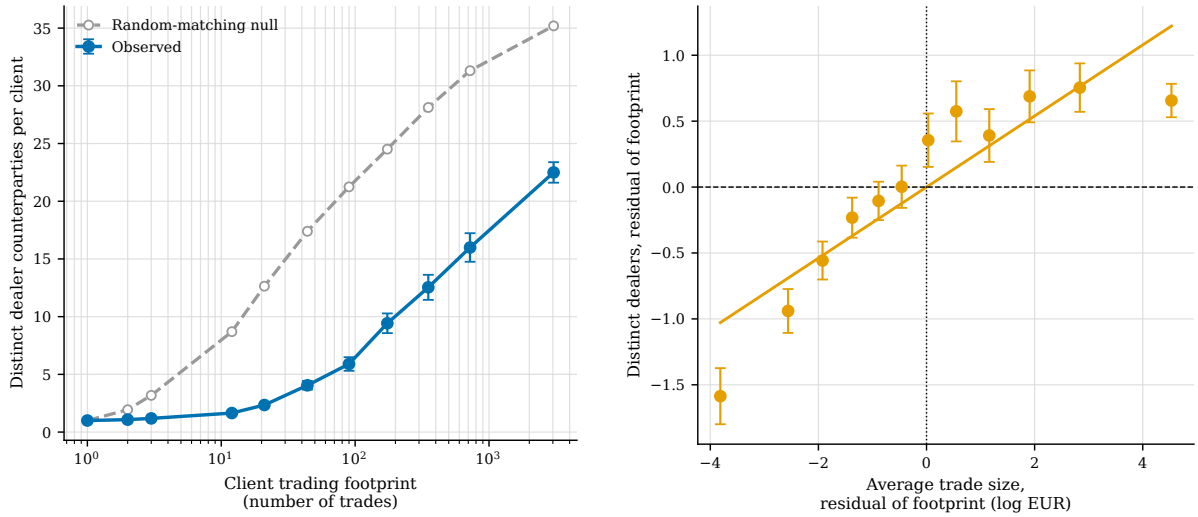


Figure 5: Distinct dealer counterparties per client against trading footprint. The left panel reports the mean number of distinct dealer groups per institutional client by footprint bin (number of dealer-to-client trades, log scale), observed (blue) versus an i.i.d. random-matching null (gray dashed), $\sum_d [1 - (1 - p_d)^N]$, the expected distinct count if each of a client's N trades drew a dealer independently from the dealer-popularity distribution $\{p_d\}$. The right panel reports mean distinct dealers against average trade size, both residualized on footprint-decile fixed effects so that the trade count, and hence the null, is held fixed; the line is the fitted slope. Error bars are 95% confidence intervals. Institutional client groups (a valid legal entity identifier; retail clients and clients in the primary dealer sector excluded), dealer-to-client trades, MiFID II subsample (2018–2025, 13,262 clients).

4.2 The Dealer System

Client identity on every trade, the Finanzagentur's per-auction allotments, and official Bietergruppe rankings let us characterize the dealer system directly: who intermediates, how the hierarchy is organized, and how new supply travels from auctions to investors.

Primary dealers are the intermediaries Defining intermediaries agnostically as groups that buy and sell the same bond within the week, across many bonds and with many distinct counterparties, recovers the Bietergruppe almost exactly (Table 2). Essentially every active member group qualifies; the roughly thirty members handle 78.5% of intermediated client volume after 2020 (71.6–79.6% across screening parameters, stable in every half-year; Figure A1). Primary dealer status is thus a function, not merely a label. By trade count the members' share is 45.6%, so the club's dominance is a large-ticket phenomenon. The residual is not competing intermediation: 84.7% of it is investment funds and hedge funds trading in both directions—the balance sheet absorption that Branger, Muck, and Pütz (2024) document—and the bank-sector two-sided traders outside the club carry 2.0% of intermediated volume. Unlike U.S. Treasuries, where principal trading firms supply a large share of interdealer liquidity (Joint Staff Report, 2015) and proposals for all-to-all trading aim to add intermediation capacity beyond the primary dealers (Duffie, 2020), the Bund market has little intermediation capacity outside the club.

Functional intermediaries	Groups	Share of intermediated volume
Bietergruppe members	32	78.5%
Investment funds & hedge funds	95	18.2%
Banks outside the Bietergruppe	104	2.0%
Other	35	1.3%

Table 2: Functional intermediaries and the Bietergruppe after 2020. Groups qualifying as functional intermediaries under the baseline behavioral screen—same-ISIN round-trips with third parties in at least eight distinct ISIN-weeks and at least ten distinct counterparty groups per half-year—split by Bietergruppe membership and, for non-members, by sector; membership is attached only after the intermediary set is fixed. Group counts are averages per half-year; shares are of intermediated client volume pooled across half-years. On average 33 member groups trade in the screened slice per half-year, of which 32 qualify. Across the six screening configurations the member share ranges from 71.6% to 79.6%. Investment funds and hedge funds are pooled; "Other" pools the remaining client sectors. The screened slice, trades between identified groups with exactly one intermediary side, carries 25.8% of all trading volume outside intragroup bookings, including interdealer trades and anonymous venue, IDB, and CCP legs.

One hierarchy, uncontested Within the club, large dealers are central dealers. Dealers with more secondary market client volume, more trades, more active days, broader client reach, larger auction allotments, and stronger Bietergruppe rank all sit closer to the center of the dealer–client network. A principal component analysis of ranked centrality, secondary market activity, and primary market status measures summarizes this common component: the first component explains 69.7% of their variation.

The more informative question is whether primary market status still matters after holding secondary market size fixed. It does: among dealers with similar secondary market volume, the dealer with the stronger Bietergruppe rank is still more central in the dealer–client network (partial rank correlation 0.28; 0.33 within volume quintiles; Figure 6). Thus the official primary market ranking and the secondary market network ranking are linked, not identical: primary rank contains information about secondary market position that is not explained by secondary market volume alone. The hierarchy is persistent and uncontested in our data: semiannual changes in rank and centrality are uncorrelated, so dealers do not appear to move up the primary market ranking in the same periods in which they gain secondary market centrality. The over-the-counter intermediation literature documents a stable core-periphery dealer network in which central dealers post lower and less dispersed spreads (Li and Schürhoff, 2019; Hollifield, Neklyudov, and Spatt, 2017; Üslü, 2019) and in which heterogeneous risk-bearing capacity among the core dealers shapes prices (Eisfeldt, Herskovic, Rajan, and Siriwardane, 2023), and the Treasury auction literature documents persistent heterogeneity between primary dealers and other bidders (Hortaçsu, Kastl, and Zhang, 2018) alongside durable dealer-client relationships (Hendershott, Li, Livdan, and Schürhoff, 2020). Those literatures study the primary and secondary markets separately; we show they are one object, and that the resulting hierarchy is uncontested at the frequency our data resolve.

The distribution chain Auction wins show up immediately in client flow. Prior work sees auction cycles in prices or aggregate bidder group flows (Lou, Yan, and Zhang, 2013; Beetsma,

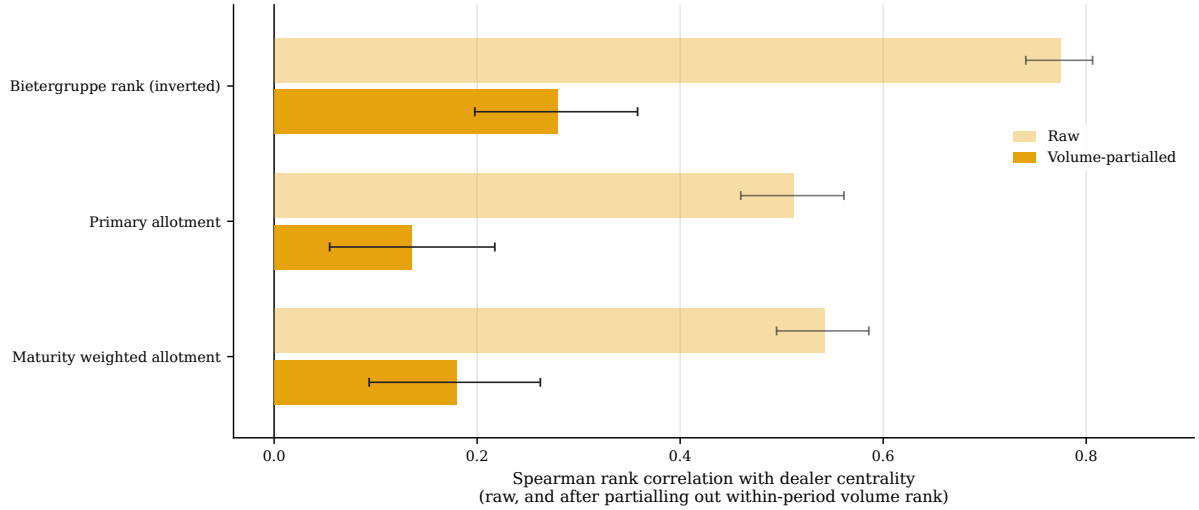
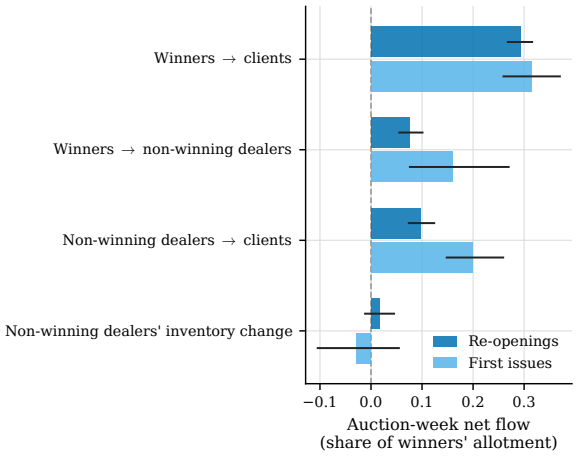


Figure 6: Dealer centrality and primary market status, raw and conditional on scale. Bars are pooled Spearman correlations between half-yearly eigenvector centrality in the dealer–client network and each primary market characteristic, with 95% confidence intervals from resampling half-years. Pale bars are raw correlations; solid bars condition on within-period secondary market volume rank. Bietergruppe rank is inverted, so positive values correspond to stronger official rank.

Giuliodori, Hanson, and de Jong, 2018; Green, Hollifield, and Schürhoff, 2007; Branger et al., 2024). Client identity lets us trace the quantity chain dealer by dealer. In the auction week, winning dealers, the groups with a positive allotment, sell 2–7 percentage points more of the auctioned bond to clients than their share of dealer-to-client sales of other bonds in the same week would predict. This is a re-opening result: first issues show no detectable excess in that week, although with 75 first issues the confidence interval is wide. The contrast reflects demand, not distribution (Figure 7a): winning dealers sell down first issues at the same pace and pass at least as large a fraction of the allotment to other dealers in the interdealer market. Demand for a new line reaches winning and non-winning dealers alike, so non-winners on-sell bonds sourced from winners within the week: net client sales of 19.9% of the winners' allotment, against 9.7% at re-openings, leaving winners' share of client sales at its normal level. At a re-opening, the week's client demand for the bond sits with the dealers that bid for it. Conditional on winning, a larger allotment does not predict more excess selling. What matters is whether the dealer won the auction. The buyer side is concentrated in investment funds and hedge funds, whose share of winners' auction week sales is 10.6 percentage points above winners' usual client mix (Figure 7b). Distribution is fast: winners' net sales to clients equal 30% of the allotment in the auction week and 45% within twelve weeks. New supply therefore moves through dealers as flow, not as inventory they warehouse.

Two further properties of the dealer system—the geographic division of labor across the hierarchy and the stability of the core's share of client volume in acute stress windows—are documented in Section A.1.

(a) Auction week net flows through winning and non-winning dealers.



(b) Shift in the client sector mix of winners' sales.

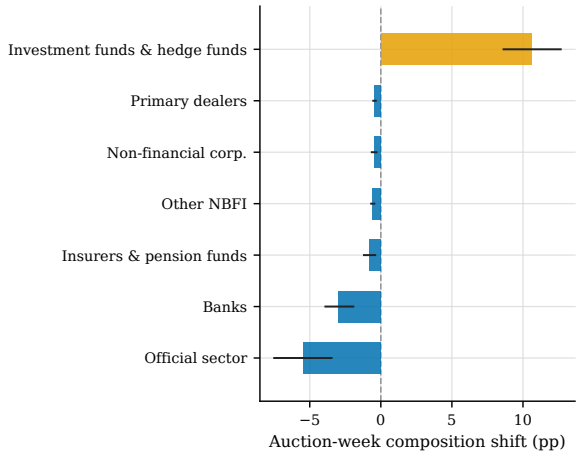


Figure 7: The distribution chain from auction to investors. Panel (a): auction week net flows of the auctioned bond as a share of the winning dealers' allotment, separately for re-openings and first issues: winning dealers' net sales to clients and to non-winning dealers, non-winning dealers' net sales to clients, and non-winning dealers' inventory change, which also includes their net purchases from the Bundesbank or Finanzagentur and from records without an identified counterparty. Winning dealers are the groups with a positive allotment; non-winning dealers are the other auction group members. Panel (b): client sector composition of winners' auction week sales of the auctioned bond relative to their sales of all bonds in the same week. Intervals are 95% auction bootstrap confidence intervals.

4.3 Fragmentation, Concentration, and Segmentation

Fragmentation Venue structure reinforces the hierarchy; it does not replace it. Figure 8a shows that OTC, MTF, OTF, and SI trading all carry different mixtures of dealer access, venue protocol family (Table A1), execution cost environment, and information environment. OTFs are close to interdealer broker markets. SIs and OTC trading remain centered on dealers: even leaving out bookings between entities of one dealer group, a dealer is party to 97% of SI and 81% of OTC execution volume. MTFs broaden access but do not eliminate dealer intermediation, and regulated markets are peripheral in volume but matter for retail access. Fragmentation is thus economic, not just a legal taxonomy.

Concentration Figure 8b uses venue groups rather than raw MIC codes, because Brexit led operators such as Bloomberg, Tradeweb, and MarketAxess to maintain parallel EU and U.K. MICs for the same trading infrastructure; Appendix D.11 details the grouping. On this grouped basis, every venue type is concentrated. The effective number of venue groups is only 2.4 for OTFs and 2.7 for MTFs, and the top three groups account for 97% and 91% of volume, respectively. SIs and regulated markets are less concentrated, with effective group counts of 9.4 and 6.5, but they are not diffuse.

Trading in Bunds is concentrated in a handful of bonds but fragmented across trading venues. Figure 9 combines the bond concentration and bond-venue sparsity facts. About 62 Bunds are outstanding on a typical day, but the liquid core is a handful of recently issued on-the-run Bunds with benchmark maturities and bonds eligible for Bund futures delivery. In Figure 9a,

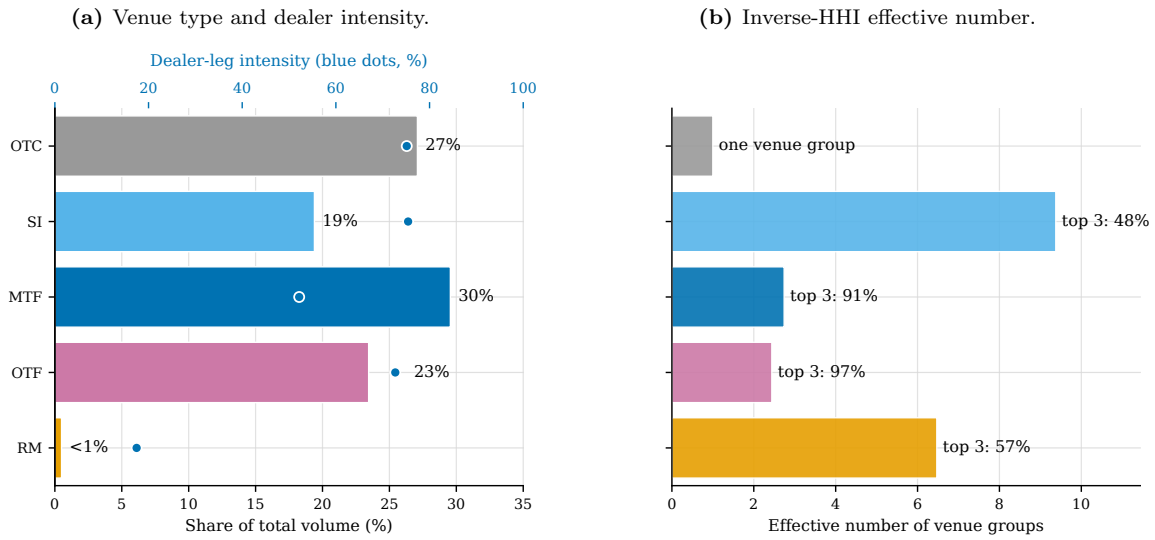


Figure 8: Venue architecture and dealer intermediation. The left panel reports each venue type's share of aggregate Bund trading volume (bars, lower axis) and dealer-leg intensity (blue dots, upper axis): the volume-weighted share of record legs held by investor groups in the primary dealer sector, in percent, so that a client-to-client record scores 0, a dealer-to-client record 50, and a dealer-to-dealer record 100. Volume bars use all positive-notional MiFID II executions with an identified venue type; dealer-intensity dots use the positive-notional trade records with both buyer and seller groups known, including records between two entities of one investor group. The right panel reports the inverse-HHI effective number of venue groups within each venue type, $1/\sum_g s_g^2$, where s_g is venue group g's volume share within the venue type. Venue groups merge MICs of the same type that belong to the same operator; text labels give the combined share of the top three venue groups. Bilateral OTC is treated as one venue group. MiFID II subsample.

each on-the-run bond carries 6.5% of daily volume on average, against 2.9% for each other deliverable bond and 0.9% for each remaining bond, so the four on-the-run benchmarks and other deliverable issues together account for 56% of daily volume. Figure 9b summarizes, for each venue type's daily top venue group, how much of the outstanding bond universe actually trades there. OTC and MTF top venues usually cover almost the full outstanding universe (97% and 98% of bonds with at least one trade), but even there the set of bonds with at least five same-day trades is smaller (63% and 69%). The other venue types are much sparser: the daily top SI, OTF, and regulated market venues cover only 48%, 34%, and 15% of outstanding bonds at the median. Fragmentation therefore does not mean the same liquid order flow is replicated everywhere; trading opportunities are specific to each venue for much of the outstanding Bund universe.

Figure 10 shows how venue choice depends on trade size. Bilateral OTC trading is widespread across ticket sizes; small trades (up to €100,000) are also executed on-exchange; medium trades (up to €50 million) are frequently executed on electronic platforms; and the largest trades concentrate on interdealer platforms and on single-dealer platforms. The trade size patterns across venue types mirror evidence from the interdealer Bund market, where different protocols coexist and brokered OTC trading is particularly relevant for executing large, anonymous trades (de Roure, Moench, Pelizzon, and Schneider, 2026).

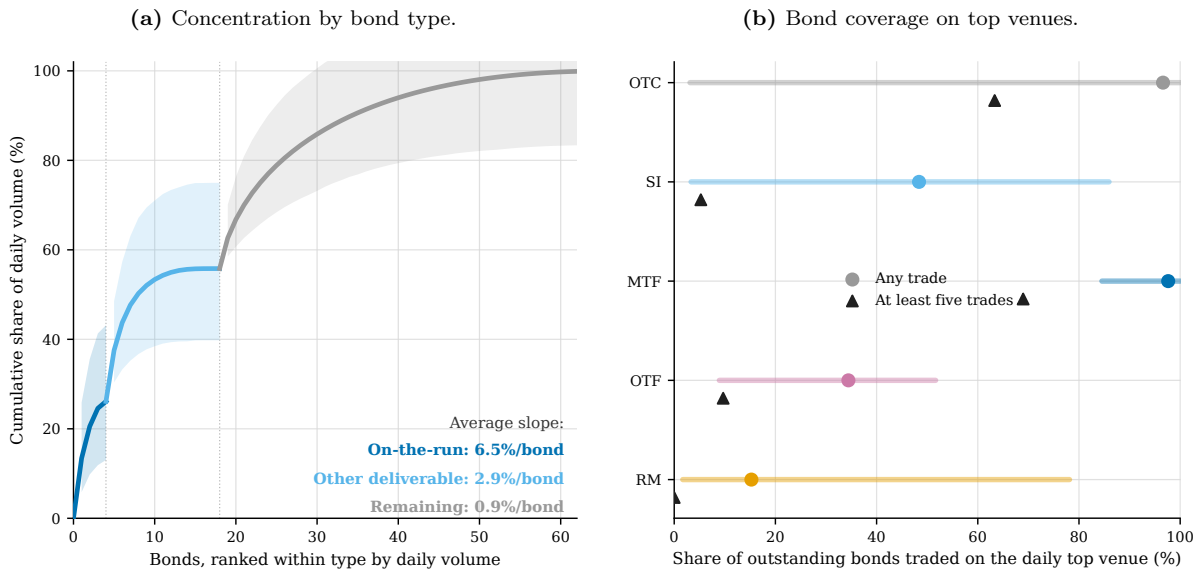


Figure 9: Bond concentration and bond-venue coverage. The left panel splits bonds into three types—on-the-run, other futures-deliverable, and remaining—ranks bonds within each type by daily volume, and concatenates the mean cumulative share of daily trading volume each type contributes; shaded bands are the 5th–95th percentile range across days. The right panel summarizes bond coverage on each venue type's daily top venue group. Circles report the median share of outstanding bonds with at least one trade on that top venue; horizontal bands are the 5th–95th percentile range. Triangles report the median share with at least five trades. Volumes and trade counts are of executions. MiFID II subsample.

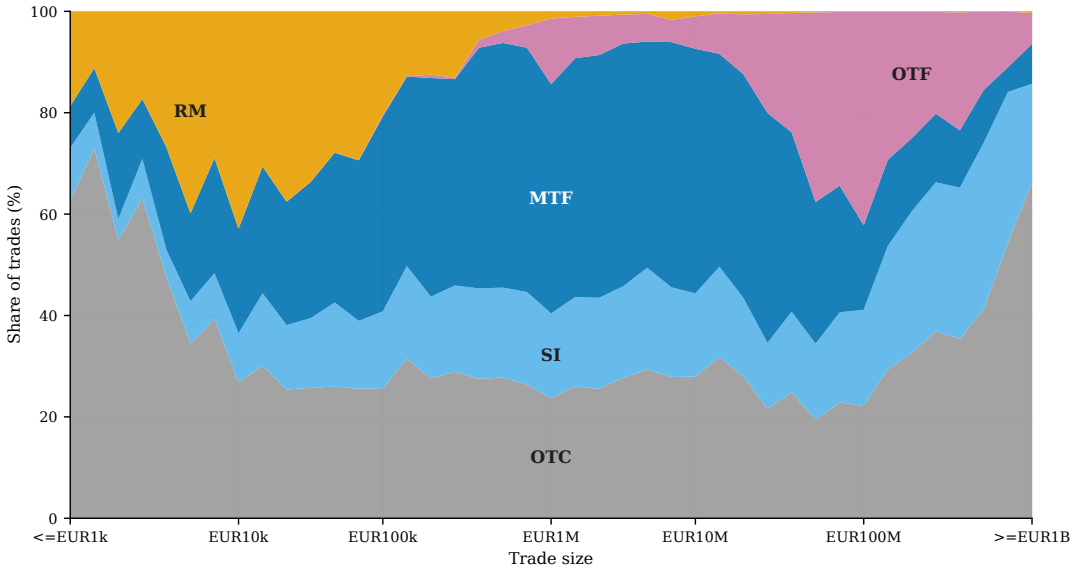


Figure 10: Fragmentation across venues by trade size. Venue composition of executions as a function of execution size.

Segmentation The Bund market is *segmented*: investor sectors reach the market through systematically different venue types. Figure 11a shows that official sector institutions and investment funds trade overwhelmingly on electronic platforms (MTFs and SIs); ICPF trade mostly there. Banks split between bilateral OTC and platforms; retail households are the near mirror image, concentrated on the regulated retail exchange and bilateral execution. This is about *who* trades, not *which* bond: within the same bond and day, knowing a client's sector removes 37.5% of uncertainty about venue type, against 16.8% under a no-segmentation null randomly reassigning sector labels within the bond–day. The 20.7-point gap lies far outside the null and survives conditioning on trade size, restricting to institutions, excluding the official sector, and the 2021 U.K.-venue coverage break.[12]

This segmentation lines up with information. The sectors with the most price-relevant flow in Section 5 trade mostly on MTFs and SIs (Figure 11b), and the venue gradient in permanent price impact holds within hedge fund flow, so it is not only a composition effect (Appendix A.2). The cross-sector gap in effective half-spreads is a within-venue phenomenon that survives venue fixed effects (Figure 14) and is, if anything, larger for the platform-intensive sectors, so segmentation does not deliver it. We read segmentation descriptively, not causally: we do not identify whether investors select venues for their information environment or are sorted onto them by access and intermediation. Its theoretical implication is that venue shares bundle client identity, dealer access, and the information content of flow, not just technology, consistent with evidence that investor identity interacts strongly with information, trading costs, and protocol choice (Czech et al., 2021; Pintér et al., 2024).

These facts identify the market structure: a small dealer core intermediates a segmented periphery through persistent relationships, across fragmented venues.

5 Information and Execution Costs

Market structure matters only if it shapes prices, flows, and trading costs. We therefore turn from who trades with whom and where to who moves prices, who supplies or absorbs stress flow, and who pays execution costs. The information, execution cost, and platform centralization analyses focus on dealer-to-client trades, unlike the market structure and clearing exercises. They require client identity and transaction signs from the client's side; we do not separately model interdealer bargaining or dealer-to-dealer price formation.

5.1 Investor Flows and Intermediation

We first identify the buyers and sellers during sharp Bund price moves to assess whether Bunds were being bought as a hedge, sold for cash, or repriced because expected supply had changed.

[12] The statistic is the conditional mutual information $I(\text{sector}; \text{venue} \mid \text{bond–day})$ on client legs of dealer-to-client trades, normalized by the conditional venue entropy $H(\text{venue} \mid \text{bond–day})$, so it is the fraction of within-bond-day venue uncertainty resolved by knowing the sector. The permutation null preserves each bond–day's sector and venue marginals and breaks only their pairing; because finite-sample mutual information is biased upward in sparse cells, the null is strictly positive, and the observed-minus-null gap is the bias-corrected measure of segmentation. The unconditional sector–venue association is Cramér's $V = 0.34$.

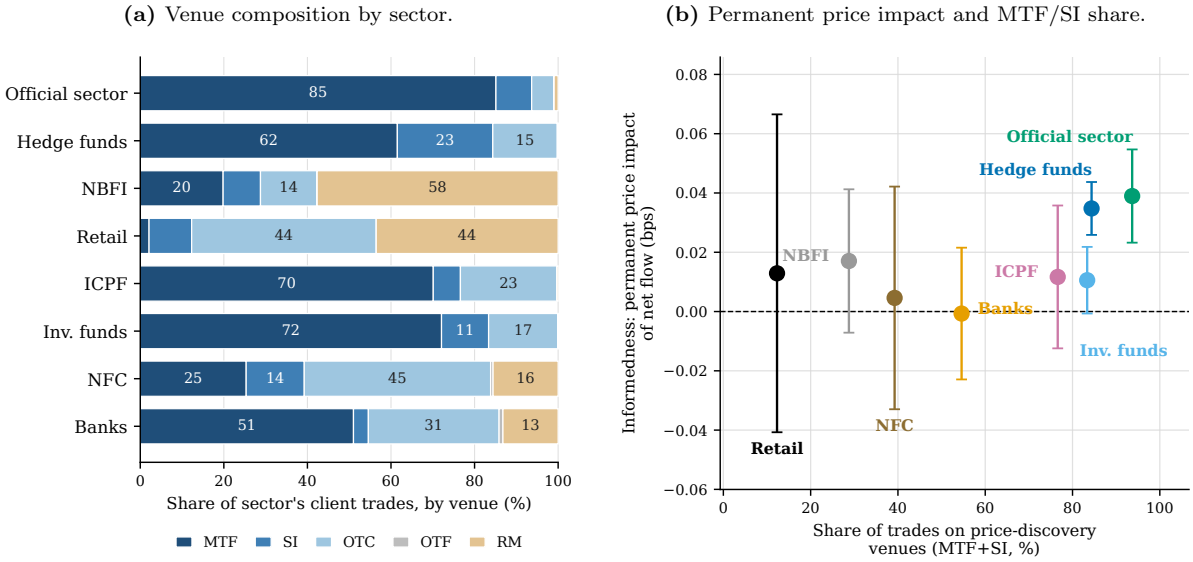


Figure 11: Venue composition and price impact by client sector. The left panel reports the venue composition of client legs in each investor sector's dealer-to-client trades (count shares), with sectors ordered by the permanent price impact of their net flow. Each bar pools a sector's trades across all operators of a venue type. The right panel reports the permanent price impact of each sector's net order flow, estimated as in Figure 13 and signed so that positive values denote richening (with 95% confidence intervals), against the share of its trades on MTFs and SIs. MiFID II subsample, investor-group sector classification, 3,408,392 client legs of dealer-to-client trades.

We therefore study three short event windows that isolate distinct shocks. March 2020 is a global cash shock: Bund prices fell as investors raised liquidity during the Covid dash for cash. March 2023 is a safety shock with the opposite price sign: Bund prices rose during banking-sector turmoil. March 2025 is a domestic fiscal news shock: Bund prices fell after news of substantially higher future German borrowing. The three episodes therefore separate the two safe asset roles emphasized by Duffie (2025) from a non-safe-haven supply repricing.

For each episode, we measure sector-level net Bund purchases over the acute window in which prices moved before the subsequent stabilization or reversal. The Covid window, 9–18 March 2020, ends at the ECB's pandemic emergency purchase program announcement. The banking-stress window, 9–17 March 2023, runs from the Silicon Valley Bank shock to the Friday before Credit Suisse's resolution weekend. For the fiscal news episode, the window is 5–11 March 2025, the first five trading days after the fiscal package was announced on the evening of 4 March. These short windows avoid mixing the dislocation with later policy responses, rebounds, or issuance dynamics. That separation is only feasible with high-frequency transaction flows; monthly or quarterly holdings would net the event against what happened after it.

Figure 12 reports results. Because every transaction has a buyer and a seller, the sector bars sum to zero once retail and unclassified reporting legs are included. We show the unclassified bar for accounting closure only: it captures reporting legs without a mapped investor sector, including missing account owner information, internal account reporting, and identifiers that do not map to an investor sector. In March 2020, the German 10-year yield rose 45 bps. The dominant net seller was the official sector, which sold €11.3 bn net; foreign official sector accounts sold €10.0 bn of it, consistent with reserve managers liquidating Bunds to raise dollars (dash

for cash). Hedge funds were net buyers, and investment funds were small net sellers.[13] Primary dealers were also substantial net buyers in this cash shock.

In March 2023, the price move had the opposite sign: Bund yields fell 46 bps amid the banking turmoil (flight to safety). The flow signature nevertheless repeated. Foreign official sector accounts again sold, and hedge funds bought. The same foreign official seller therefore appears in a cash shock and a safety shock, but the buyers differ: primary dealers are small net sellers in 2023, and hedge funds are the main identified absorbing sector.

The March 2025 fiscal repricing is different. Friedrich Merz, then Chancellor-designate, announced plans to relax Germany's constitutional debt brake for defense and infrastructure spending. Bund prices fell sharply, with the German 10-year yield rising 41 bps. This was not a global dash for cash, and it was not a flight to safety. It was news about future German debt supply. The flows line up with that interpretation. Primary dealers were the main net sellers. Mapped client sectors were net buyers, especially investment funds, banks, and ICPF. Hedge funds, which bought heavily in the March 2020 and March 2023 stress episodes, bought only modestly. Thus dealers were not absorbing customer sales, as in a liquidity shock. They were reducing Bund exposure and passing supply risk to end investors through secondary market trades. Recall that we exclude primary allocation legs, so these flows show the secondary market adjustment to expected supply. As Appendix A.1 shows, the same windows leave the split of client flow across the dealer hierarchy unchanged: the stress reallocation is between investor sectors, not across dealers.

5.2 Who is the informed flow?

A central question in market microstructure is *whose* flow has permanent price content. Because the Bund data identify client sectors, we can rank sectors directly, whereas equity and Treasury data cannot (Brandt and Kavajecz, 2004; Hendershott, Livdan, and Schürhoff, 2015). For each dealer-to-client trade we sign flow from the client's side, aggregate to bond–day–sector net flow, and ask whether a sector's net buying predicts a subsequent decline in the bond's curve-relative pricing error. ISIN fixed effects absorb the bond's average curve-relative pricing error, and maturity bucket×day fixed effects (remaining maturity under 1, 1–3, 3–7, 7–15, and over 15 years) absorb yield and fitted-curve moves common to bonds of similar maturity on a day, so the coefficient is a within-segment security selection effect. The regressions also control for the bond's cheapness before the flow day: its curve-relative pricing error on the previous trading day and the two preceding daily changes in that error. Investors tend to buy bonds that have just cheapened, and curve-relative cheapness mean-reverts, so these controls keep predictable

[13] Quarterly holdings of all debt securities show banks and investment funds trading procyclically and insurers and pension funds countercyclically (Timmer, 2018); in this acute window, insurers and pension funds were net sellers of Bunds, as were investment funds. Quarter-end holdings also include any rebound, while our daily Bund flows isolate the dislocation. A net seller took that side on balance, not necessarily in a forced sale, and primary dealer net purchases measure secondary market flow, not balance sheets.

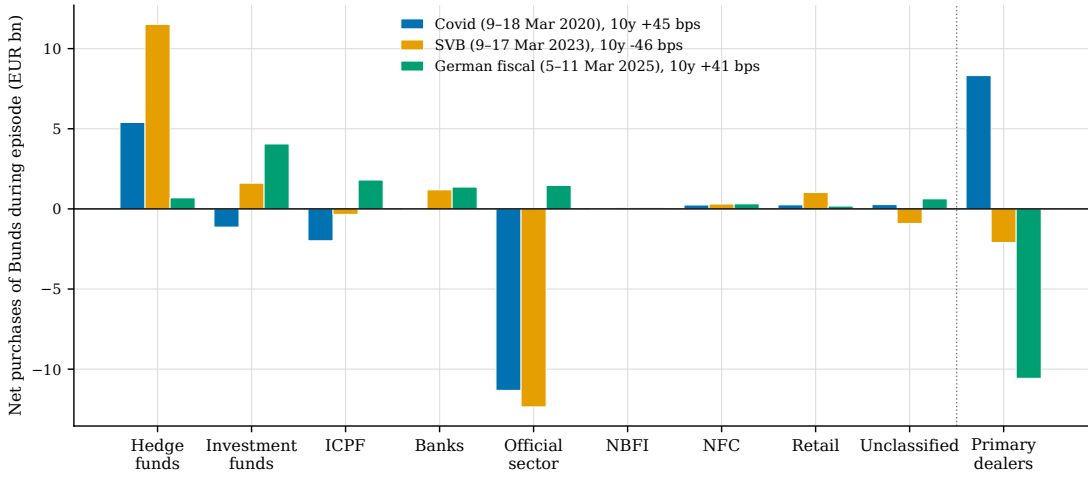


Figure 12: Net Bund purchases by investor sector and reporting category across three stress episodes. Bars report net purchases (trade volume bought minus sold, € bn) over the acute window before intervention in three stress episodes: the Covid dash for cash (9–18 March 2020, Bund yields rising); the March 2023 banking stress around Silicon Valley Bank (9–17 March 2023, Bund yields falling); and the German fiscal reform repricing (5–11 March 2025, Bund yields rising). Negative bars denote net sellers. Legend entries report the change in the German 10-year yield from the last close before each window to the close of its last day. Retail bars collect natural person legs; unclassified bars collect legs without a mapped investor sector. Primary market allocation trades to primary dealers are excluded by construction. The 2020 window predates the United Kingdom's departure from the MiFIR reporting perimeter on 1 January 2021: 60% of its volume is reported on U.K. venues, against 11% and 11% in the 2023 and 2025 windows, so cross-episode magnitudes are not directly comparable.

reversion from being credited to the flow (Appendix A.3). A permanent response indicates price-relevant flow; a response that reverts indicates liquidity demand (Evans and Lyons, 2008).[14]

In Figure 13, hedge funds and the official sector stand out: their net purchases predict the most persistent richening, and their estimates are the most precise. Hedge fund net purchases predict permanent richening of 0.03 bps per standard deviation of net flow at the ten-day horizon ($t = 7.6$). Net purchases by the official sector predict the largest persistent richening, but we read this as policy and reserve management flow rather than private information of the kind associated with hedge funds. Official accounts tend to buy bonds that cheapened on the previous trading day: a one standard deviation net purchase follows a cheapening of 0.011 bps ($t = 2.9$). Part of the official impact estimated without the cheapness controls is therefore reversion of that cheapening (Appendix A.3), and such mechanical value buying fits the policy and reserve management reading. Investment fund flow carries an impact that peaks at three days and fades: 0.01 bps at ten days ($t = 1.9$) and 0.00 bps at twenty days ($t = 0.2$). The impact is transitory, not permanent, the signature of liquidity demand. Bank flow carries a short-lived impact: 0.03 bps at five days ($t = 2.0$), 0.00 bps at ten days ($t = -0.1$), and an estimate not distinguishable from zero at twenty days ($t = 0.0$). Unlike the hedge fund and official sector estimates, it does not carry over to future implied yield changes (Appendix A.3). ICPF, NBFI, NFC, and retail household flow carry no significant ten-day signal; the retail estimate is 0.01 bps

[14] Appendix A.2 repeats this permanent impact exercise at the venue type level rather than the client sector level. MTF, OTC, and SI net flows have positive permanent impact, while regulated market flow is close to zero. A secondary cut by venue protocol family is descriptive and does not restate the venue types (Table A1).

($t = 0.5$) and is not significant at any horizon from one to twenty days. The ranking is the contribution. The magnitudes are small, measured relative to the kernel ridge curve, and not a tradable strategy net of costs.

These results support the presence of price-relevant sector flow but do not identify whether the information is private research, inventory-revealing flow, policy demand, or dealer learning. The hedge fund result survives several checks: it holds when the March 2020 stress quarter is excluded, has the same sign in every year from 2018 to 2025 and is significant in most of them, survives replacing the curve-relative pricing error with future synchronized implied yield changes (Section A.3), and NFC flow, a placebo sector, carries no significant ten-day impact. The official sector estimate survives the same replacement. The official sector effect is concentrated among foreign reserve managers and is not a Bundesbank/ECB/Finanzagentur artifact.

The venue cut reinforces the interpretation without making operator claims. MTF, OTC, and SI net flow are all informative and permanent, MTF flow the most at the ten-day horizon, while regulated market flow is not distinguishable from zero (Appendix A.2). The gradient is not only sector composition: holding the sector fixed, hedge fund flow carries more permanent price impact on MTFs than over the counter. The ordering is therefore informational, but not a simple electronic-versus-OTC split.

These sector and venue rankings set up the pricing test. A simple adverse selection account would charge the highest effective half-spreads to the sectors whose flow has the most permanent price content. The Bund market does not.

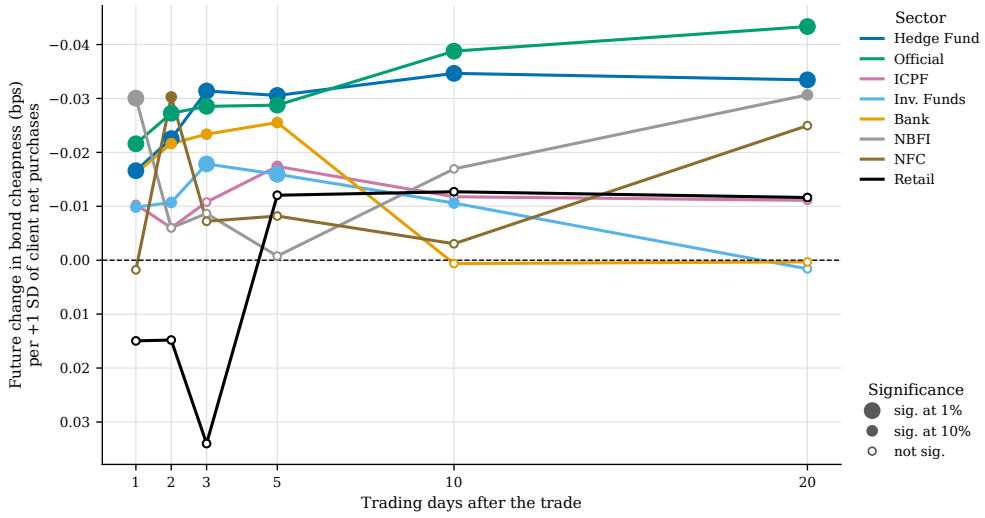


Figure 13: Permanent price impact by client sector. For each identified client sector, the cumulative future change in the curve-relative pricing error, defined as synchronized implied yield minus the kernel ridge fair yield, in bps, per net purchase of the bond equal to one standard deviation of the sector's absolute net purchases, by trading days after the trade. Sector-by-horizon regressions with ISIN and maturity bucket×day fixed effects and controls for the bond's lagged curve-relative pricing error and its two lagged daily changes; the y-axis is inverted so that upward denotes richening. Marker size encodes significance (ISIN-clustered standard errors): large filled = 1%, medium filled = 10%, small hollow = not significant. Dealer-to-client trades, MiFID II subsample (2018–2025).

5.3 Who pays higher effective half-spreads—and is it adverse selection?

The informed flow ranking has a natural counterpart: the price each sector pays for immediate execution.[15] For every dealer-to-client trade we measure the *effective half-spread*—the concession of the execution price relative to the pre-trade MTS top-of-book mid, signed from the client's side and reported in price basis points (Bessembinder et al., 2020; Pintér et al., 2024).[16] Classic adverse selection logic predicts this spread should be highest for the clients whose flow carries the most permanent price content.

Figure 14a shows otherwise: effective half-spreads vary sharply across client sectors. Retail households pay 2.48 bps, roughly 4 times the 0.65 bps paid by banks. The gap *widens* to +3.71 bps after absorbing ISIN×day, venue, and trade size fixed effects. Among institutional sectors the spread is not monotonically increasing in informedness (Figure 14b): hedge funds and official sector accounts—whose net flow has the most permanent price content (Figure 13)—pay only middling spreads (0.70 and 0.65 bps), below ICPF and investment funds, whose permanent price impact is no larger. Retail households, who pay the most, carry no detectable permanent price impact (Figure 14b). Net of venue, size, and bond–day fixed effects, hedge fund flow trades at a spread below banks' (a -0.30 bps difference, $t = 5.2$). In Bunds, the effective half-spread is an execution cost that does not rise monotonically with permanent price content, not a price for adverse selection.

This is a statement about investor cost and the split of surplus, not about welfare: the effective half-spread nets no inventory, hedging, or matching cost, so it is neither a dealer rent nor, absent a model of matching, an efficiency loss. The two firm findings are the retail-versus-institutional cost gap and the absence of any extra charge on the high permanent impact flow; the raw rank correlation between cost and information across the eight client sectors is weak ($\rho = -0.07$) and far from significant. The comparison is across sectors, not a trade-by-trade test. The official sector is read as policy and reserve manager footprint. The sample is the MiFID II period (2018–2025, 3,313,054 dealer-to-client trades).

5.4 The price of size in Bunds

How does the cost of execution vary with trade size? Pintér et al. (2024) separate two offsetting forces in non-anonymous bond data: across clients, larger trades receive a pooled *size discount*, while within a client, larger trades pay a *size penalty*—and they document the pooled discount even in a sovereign market (Gilts). When we measure the client's cost the same way (the signed concession of the execution price to the pre-trade MTS mid) and use the same pooled-versus-within-client contrast, Bunds *confirm the size penalty but, unlike Gilts, exhibit no size discount.* Figure 15 shows effective half-spreads *rising* monotonically with trade size: net of ISIN×day,

[15] We treat the client as the demander of immediate execution. The sign convention provides a check: a client sector that systematically supplied liquidity would show a negative average half-spread, whereas measured half-spreads are positive across sectors. The rules of the regulated retail exchanges make them quote-driven, with designated liquidity providers supplying continuous quotes against which retail orders execute.

[16] The reference mid is the last MTS top-of-book quote within 60 minutes before the trade; trades without a quote in that window are dropped, so the spread is never measured against a stale quote (median quote age is a few seconds; see Section 3.5).

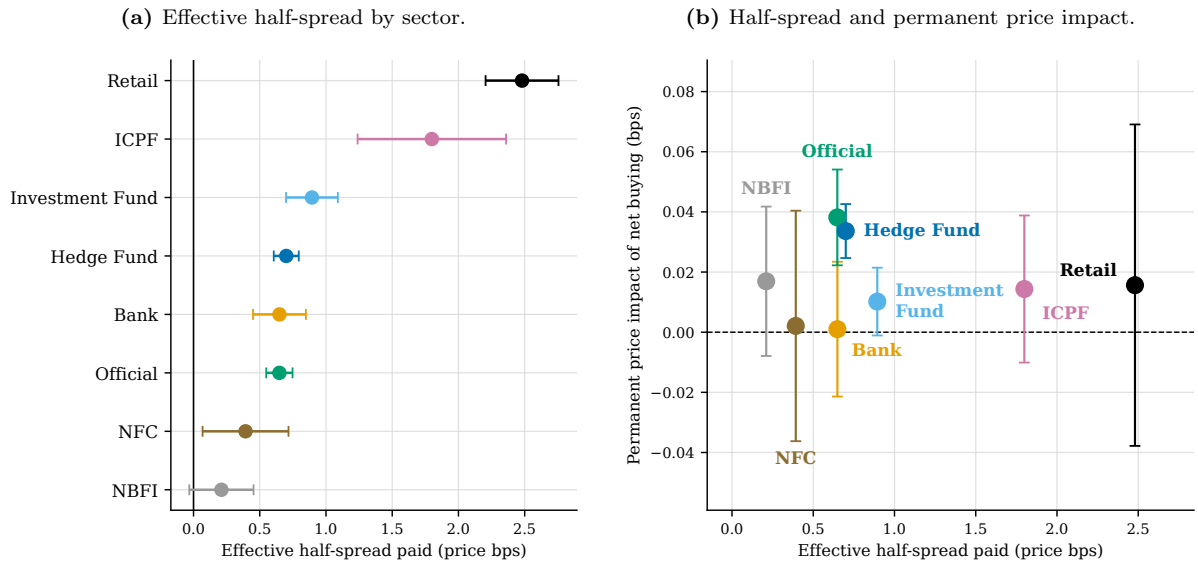


Figure 14: Effective half-spreads and permanent price impact by client sector. The left panel reports the volume-weighted effective half-spread paid by each client sector—the concession of the execution price relative to the pre-trade MTS top-of-book mid, signed from the client's side and reported in price bps; error bars are 95% ISIN-clustered confidence intervals. The right panel plots each sector's half-spread (x) against its ten-day price impact (y), re-estimated on this sample as in Figure 13; positive values denote richening. Vertical bars are 95% ISIN-clustered confidence intervals; half-spread intervals are in the left panel. Dealer-to-client trades, MiFID II subsample (2018–2025).

sector, and venue fixed effects, blocks above €25 million pay an extra 1.36 bps. The coefficient on log notional is 0.138 bps ($t = 13.7$), equivalent to 0.096 bps for a doubling of trade size.

The penalty is genuinely within-investor, and the absence of a discount is not a control artifact. The *raw* cross-client relation—the variation in which Pintér et al. (2024) recover a discount—is essentially flat (−0.002 bps for a doubling of trade size): the pooled-discount force is too weak to turn the net relation into a discount. Absorbing client-*group* (individual investor) fixed effects isolates one investor trading different sizes, the cleanest test of the heterogeneity concern of Pintér et al. (2024). The coefficient stays positive (0.143 bps on log notional, $t = 16.7$), equivalent to 0.099 bps for a doubling of trade size.

That Bunds preserve the within-client penalty but show *no* pooled discount is itself notable. A size discount is one of the more established regularities of over-the-counter trading, usually rationalized by bargaining models in which dealer market power falls with trade size. Its absence here is intuitive for a market that is fragmented yet extremely liquid: dealers are not competed into granting large Bund orders a relationship discount, so the inventory and price impact cost of warehousing a block is passed back to the client. As above, a size penalty is the execution cost of a block, not evidence of mispricing or dealer rents. The sample is the MiFID II period (2018–2025) and the 3,313,054 dealer-to-client trades of Figure 14.[17]

[17] These execution costs are, in turn, essentially all of the within-bond *price* dispersion: Appendix A.4 shows that a trade's venue, dealer–client segment, and client sector explain no more of the within-bond-day synchronized-yield dispersion than a same-cardinality random partition and do not predict it out of sample, so market structure disperses the cost of trading, not the level of prices.

(a) Effective half-spread by trade size.

(b) Size slope by specification and sector.

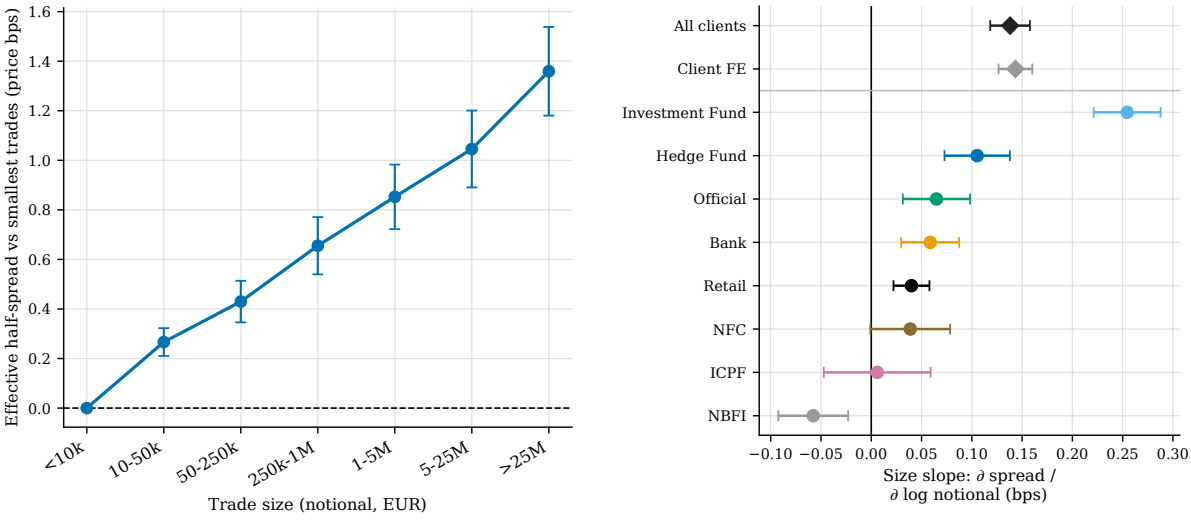


Figure 15: Effective half-spreads by trade size. The left panel reports the effective half-spread paid relative to the smallest trades, by trade size bucket, estimated from size-bucket dummies absorbing ISIN×day, client sector and venue fixed effects (so the profile is composition-free). The right panel reports the all-client size slope with sector and venue fixed effects, the all-client slope with client-group and venue fixed effects, and sector-specific slopes estimated with ISIN×day and venue fixed effects; positive means larger trades pay more. Regressions are unweighted. In both panels the error bars are 95% ISIN-clustered confidence intervals. Dealer-to-client trades, MiFID II subsample (2018–2025).

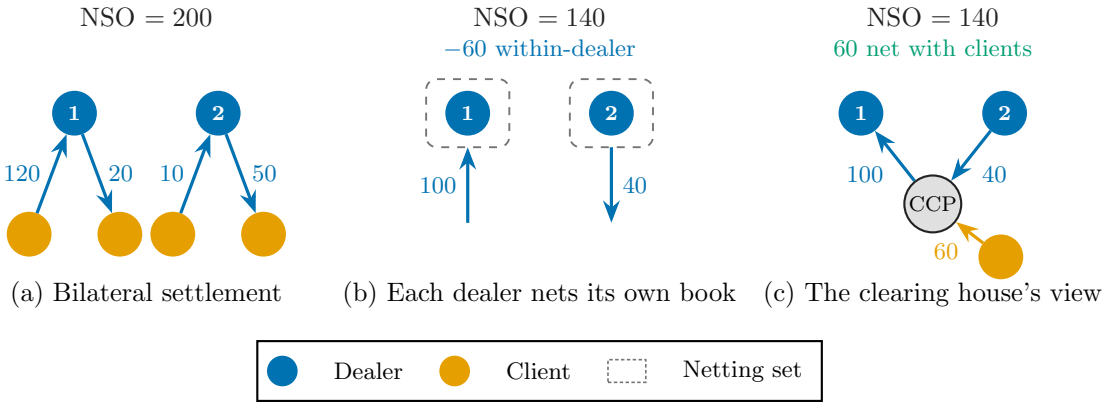
6 Reform Counterfactuals

We use the market structure facts to evaluate two reforms targeting the dealer core at different stages of the trade. Mandatory central clearing is a post-trade settlement counterfactual: dealer concentration and fragmented settlement determine how much netting at a clearing house can reduce dealers' net settlement obligations. Platform centralization is a pre-trade matching counterfactual: platform access, persistent relationships, and dealer competition determine whether broader access raises gains from trade once relationship surplus may be lost. We therefore evaluate clearing as mechanical netting and centralization as equilibrium matching.

6.1 Mandatory Central Clearing

Two facts matter for settlement: a small dealer core stands on most trades, and those trades settle across fragmented venues and counterparties rather than a common netting set. Central clearing has been widely proposed to reduce "daisy-chain" settlement failures and counterparty risk and to enable more all-to-all trading (see, for example, Duffie, 2025; Chaboud et al., 2025); here we isolate its mechanical netting effect. It novates trades to a CCP and nets obligations multilaterally. Holding the observed trade set fixed, we ask how much dealers' net settlement obligations would fall if mandatorily cleared Bund trades entered a CCP. The exercise follows the per-dealer netting of Fleming and Keane (2021) for U.S. Treasuries and Baranova et al. (2023) for U.K. Gilts; the three markets differ in their status quo, which we account for below.

We estimate these effects at the settlement-date level. The sample includes all trade records with at least one leg in the primary dealer sector and excludes intragroup records. Each trade is converted into signed dealer settlement positions at $T+2$: dealer-to-client trades generate one dealer position and dealer-to-dealer trades generate one position for each dealer. Each dealer leg therefore counts as one settlement obligation. Gross settlement obligations are the sum of absolute nominal obligations across dealer legs settling on date s, $\text{GSO}_s = \sum_{\ell \in s} |\text{Notional}_\ell|$. Status quo net settlement obligations net only trades that are already centrally cleared—records that name a CCP and records on three venues we treat as centrally cleared—each dealer against the CCP; all other obligations settle gross. In the mandatory-clearing counterfactual, each dealer nets its own signed positions by ISIN and settlement date against the CCP, as in Fleming and Keane (2021, p. 13) and Baranova et al. (2023, equation 5). We report a pre-Covid baseline, the settlement dates of January 2 to February 25 and April 1 to 30, 2020, which follows the non-crisis window of Fleming and Keane (2021), and the settlement days of the 15 February–20 March 2020 stress window. The pre-Covid estimate is representative of the broader sample: the median share of gross obligations that own-book netting saves is about 34% on settlement days before 2021 and about 33% from 2021.



(a) Bilateral settlement (b) Each dealer nets its own book (c) The clearing house's view

Figure 16: Central clearing and net settlement obligations (NSO) with two dealers. Arrows show deliveries of the security, and numbers are notional amounts. Dealer 1 buys 120 and sells 20; Dealer 2 buys 10 and sells 50. Panel (a) settles every dealer leg bilaterally. Panel (b) nets each dealer's purchases against its own sales. Panel (c) shows the clearing house's view: Dealer 2 delivers 40 to the CCP and Dealer 1 receives 100, so the dealers still owe 140. The 60 that remains after Dealer 2's net sale offsets part of Dealer 1's net purchase is the dealers' net trade with clients, which the CCP passes through on its matched book.

Figure 16 mirrors this accounting. Bilateral settlement leaves dealers funding gross obligations across many counterparties. Central clearing novates trades to the CCP, and each dealer's purchases then offset its own sales. A clearing house settles each member's net separately and keeps a matched book, so one dealer's net purchase does not cancel another dealer's net sale: both dealers still settle with the CCP. Own-book netting at a CCP would lower dealers' pre-Covid net settlement obligations from €17.24 bn to €11.13 bn per settlement day. This cuts 35.4% of the status quo and 35.5% of gross obligations (Figure 17b). The saving of €6.10 bn per settlement day equals 0.57% of the €1.06 tn outstanding in the sample's bonds. Settlement instructions fall from 2,828 to 658 a day, a cut of 77%.

Fleming and Keane (2021) and Baranova et al. (2023) also net each dealer's own purchases against its own sales and add up what each dealer still owes the CCP. Their status quos differ from ours, so a reduction against the status quo is not comparable across markets: the Bund status quo nets 0.14% of gross obligations, against 32% in Treasuries, where interdealer trades are already cleared (Fleming and Keane, 2021, p. 11–12), and about 19% in Gilts, where dealers net bilaterally (Baranova et al., 2023, Table 5). Gross obligations are the common base. Against gross, clearing removes about 74% of dealer obligations in Treasuries from January to April 2020 (Fleming and Keane, 2021, p. 11 and Table 1), about 59% in Gilts before the March 2020 stress (Baranova et al., 2023, Table 5), and 35.5% in Bunds before Covid.[18] Bunds therefore benefit less from central clearing than either market. Against their own status quos, the printed reductions are 60% and 67% for Treasuries (Fleming and Keane, 2021, p. 13–14 and Table 1) and 48% and 53% for Gilts (Baranova et al., 2023, Table 5). Interdealer legs are dealer legs whose other leg is another dealer group or has no account owner, as on anonymous broker and venue trades; their share of dealer gross obligations is similar in Bunds (47%) and Treasuries (51%). But Bund interdealer legs net only 24% within dealer, against 63% (Fleming and Keane, 2021, p. 11 and p. 15), and 25% of Bund dealer gross obligations falls on dealer-bond-days with trades in one direction only, which cannot net. Across the settlement days of the 15 February–20 March 2020 stress window, the median own-book reduction is 35.8% of gross obligations.

(a) Cross-market comparison. **(b)** Bund settlement obligations.

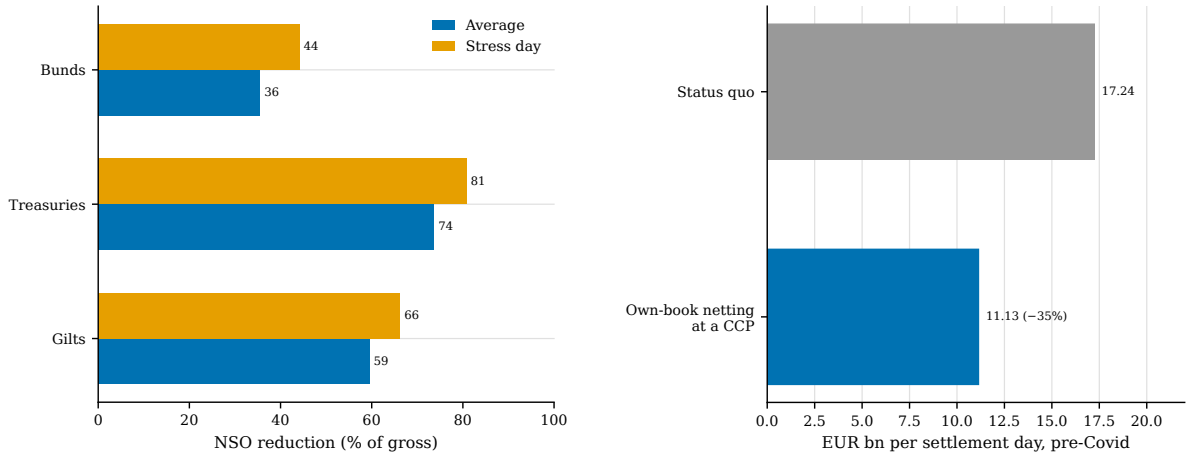


Figure 17: Mandatory central clearing and dealer settlement obligations. The left panel reports reductions in net settlement obligations under mandatory central clearing, with each dealer netting its own trades, as a percentage of gross settlement obligations. Treasury and Gilt bars are computed from the levels on p. 11 and in Table 1 of Fleming and Keane (2021) and in Table 5 of Baranova et al. (2023); Bund bars are based on our transaction data. Average bars: the Treasury window runs from January 2 to April 30, 2020, the only window with a gross figure; the Gilt window runs from January 2 to February 25, 2020; the Bund window runs from January 2 to February 25 and April 1 to 30, 2020. Stress-day bars: 28 February 2020 for Treasuries and 10 March 2020 for Gilts, the papers' peak trading days; 12 March 2020 for Bunds, the settlement date of the trades of 10 March, which is also the largest Bund reduction in the 15 February–20 March 2020 stress window. The right panel reports pre-Covid Bund obligations per settlement day: the status quo and own-book netting at a CCP.

[18] These figures count every reported leg, including pass-through and agency legs. Netting each dealer's positions within one execution before netting across its trades gives a reduction of 32.3% of gross obligations; netting per dealer legal entity instead of per banking group gives 32.4%; both corrections together give 29.2%.

Why would dealers care about net settlement obligations? For cash bonds, the balance sheet channel is weak. Settlement is delivery versus payment over a short window, so the counterparty exposure in an unsettled cash trade is small (Baranova et al., 2023, p. 16). Unsettled cash trades entered the leverage exposure measure only for dealers on trade-date accounting, and the measure now offsets cash receivables against cash payables whenever both trades settle delivery versus payment: under Basel III since January 2023 (Baranova et al., 2023, p. 10, 14, 33) and for EU banks since June 2021 (European Parliament and Council of the European Union, 2019, Article 429g(2)). Clearing such cash trades therefore no longer changes the measure. The material balance sheet gain that Baranova et al. (2023) find is in repo, which our data do not contain. What remains for cash Bunds is operational. Each settlement instruction is a delivery that can fail, and a failed delivery can propagate along a chain of dependent trades; in Treasuries, most fails in specific issues pair off and would disappear under full clearing (Fleming and Keane, 2021, p. 25–26). In our data, settlement dates are computed, not reported, and settlement fails are not observed, so we measure the instruction count, not fails.

Does clearing help most in stress? The share of gross obligations saved is about the same in stress as on other days: on the settlement dates of the trades of 9–18 March 2020, the dash for cash, the median own-book saving is 36.3% of gross obligations, against a median of 34.0% on other settlement days before 2021, and before 2021 the share is uncorrelated with the volatility of the 10-year Bund yield (correlation of −0.031). The euro saving scales with trading volume, so the resilience benefit of clearing is the size of the saving on high-volume days and the simpler settlement, not a countercyclical release of capacity. On those dash-for-cash settlement dates, own-book netting would have removed €8.8 bn of obligations per day, against the cumulative €8.3 bn of Bunds that primary dealers absorbed from clients over the eight trading days (Figure 12). Clearing nets the churn around that absorption, never the absorption itself.

	One CCP	Venue buckets	MIC silos
Status-quo NSO, pre-Covid (EUR bn)	17.24	17.24	17.24
Own-book NSO after clearing, pre-Covid (EUR bn)	11.13	13.49	14.38
Pre-Covid NSO reduction (%)	35.4	21.7	16.6
Share of one-CCP reduction retained (%)	100.0	61.4	46.8
Largest reduction, 15 Feb–20 Mar 2020 (%)	44.1	28.5	22.2
Settlement date of the largest reduction	2020-03-12	2020-02-21	2020-02-24
Largest reduction as share of one-CCP largest (%)	100.0	64.7	50.4

Table 3: Mandatory central clearing with one CCP and with restricted netting sets. In every column, each dealer group nets its own positions by settlement date and ISIN against the CCP. The venue-bucket sensitivity nets only within three buckets: all MTFs, bilateral OTC and SI trades, and one bucket for all other venue types. The MIC-silo sensitivity nets only within observed MIC venue groups, falling back to MIC and venue type when the grouped venue field is missing. The lower panel reports the largest reduction across the settlement dates of 15 February to 20 March 2020.

One CCP for all Bund trades is an upper bound, because actual implementation could leave clearing fragmented across venue buckets or MIC-based venue groups. Table 3 therefore reports restricted netting sensitivities, not observed routing models or political transition paths:

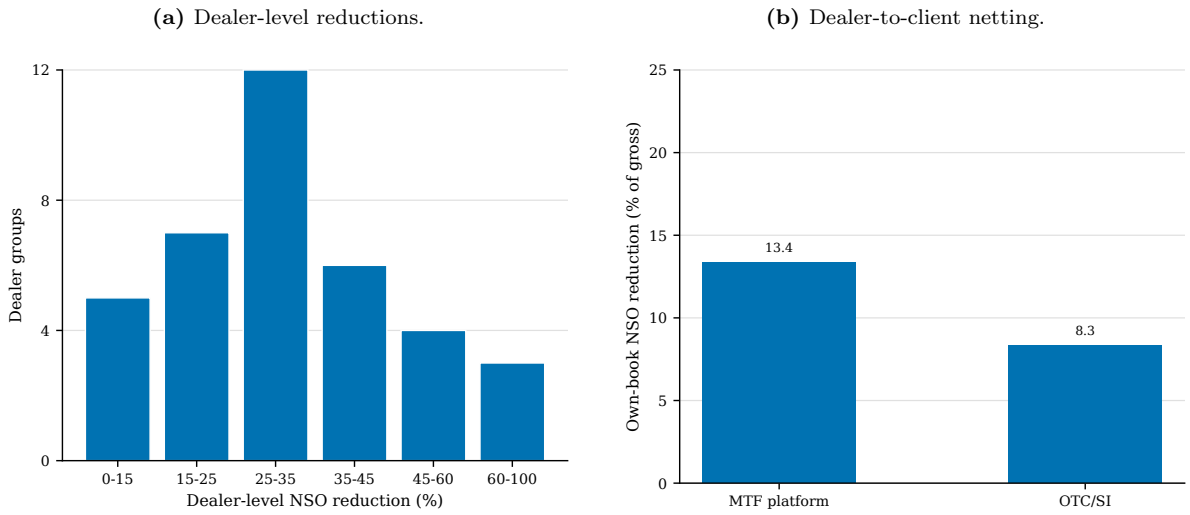


Figure 18: Dealer-level heterogeneity and dealer-to-client netting. The left panel reports the binned distribution of pre-Covid dealer-group NSO reductions under mandatory clearing. Bins are 0–15, 15–25, 25–35, 35–45, 45–60, and 60–100 percentage points. The right panel reports own-book reductions in net settlement obligations, as a percentage of gross obligations, for institutional dealer-to-client trades in the platform-centralization model sample, separately for MTF platform and bilateral OTC/SI trades.

each dealer nets its own positions only within a venue bucket or a venue/MIC silo. The pre-Covid reduction in net settlement obligations declines from 35.4% with one CCP to 21.7% when netting is confined to three broad venue buckets (MTFs, bilateral OTC and SI, and all other venue types), and to 16.6% under the conservative venue/MIC silo proxy. Netting gains are also uneven across dealers. Figure 18a shows five dealer groups with reductions below 15% and three with reductions above 60%; the dealer group median is 34%. Figure 18b gives a related but narrower diagnostic. Within institutional dealer-to-client trades used in the structural model of Section 6.2, MTF platform positions are more internally offset than bilateral OTC/SI positions: own-book netting removes 13.4% of gross obligations on MTF platforms, compared with 8.3% on bilateral OTC/SI trades. We read this as a sample diagnostic, not as causal evidence that rerouting a given bilateral trade to an MTF would lower settlement obligations. The point is institutional: clearing and platform centralization are distinct reforms, but the realized clearing gain depends on how broad the netting set is.

Interpretation These estimates are mechanical netting gains, not a full cost–benefit analysis of a clearing mandate. They exclude CCP fees, initial margin and liquidity costs, changes in trading flexibility, and private bilateral netting agreements that may already economize on settlement outside observed central clearing. Nor do they show that clearing would relieve a binding constraint on dealers. Implementation is also more complex in the euro area than in the United States because sovereign issuance, venues, clearing arrangements, and supervision are fragmented across jurisdictions. Consistent with this institutional constraint, European Systemic Risk Board (2024) recommends an incentives-based approach that harmonizes risk management between bilateral and CCP clearing, reduces prudential disincentives for NBFIs, broadens CCP access models, and strengthens CCP services. Our estimates therefore quantify the settlement

efficiency gains available in the data before costs, while institutional design determines how much of those gains can be realized.

6.2 Venue Centralization

The centralization exercise uses the relationship and venue facts: platforms broaden dealer access and lower investors' effective half-spreads, but clients retain persistent dealer relationships that may carry non-price value. Platform centralization is therefore different from central clearing. Clearing changes settlement obligations after trades occur. Centralizing trading changes the matching problem before trades occur: investors may reach more dealers, dealers may compete more aggressively, and bilateral relationship value may be lost. Those are welfare questions, not mechanical netting questions.

We evaluate this reform by applying the dealer–investor framework of Allen and Wittwer (2023) to Bund dealer-to-client trades. The earlier facts set the mapping: bilateral OTC and SI trades are relationship-based; MTF trades broaden dealer access; clients have persistent home dealers; and platform use is already high among institutional investors. The counterfactual asks whether forcing more trading onto platforms improves matching enough to offset lost relationship surplus and other costs of restricting bilateral choice.

Sample and setup We use dealer-to-client transactions from 2023–2025 so that the estimation reflects the current venue structure. We collapse venues into two model categories: bilateral trading, which includes traditional OTC and SI activity, and platform trading, which includes MTF execution on venues such as Bloomberg and Tradeweb. This collapse follows the market structure evidence: bilateral trading is relationship-based, while multilateral platforms broaden dealer access and introduce competition among dealers (see Figure 4). It is a binary model mapping, not a claim that SI and traditional OTC execution are economically identical; both enter the bilateral outside option because they preserve dealer-specific access relative to multilateral MTF trading. Following Allen and Wittwer (2023), we exclude in-house trades between dealer groups and their own affiliates, rank group-level continuously active Bietergruppe dealers by dealer-to-client volume, and use the top 20 dealers with equal dealer weights as the preferred specification. Following the model, an investor is classified as institutional if it trades on a platform at least once during the estimation window and as retail otherwise. Table A2 documents this mapping: platforms account for about two thirds of all dealer-to-client trades and 70% among institutional clients, while SI trades are small by trade count but substantial by volume and are grouped with bilateral trading.

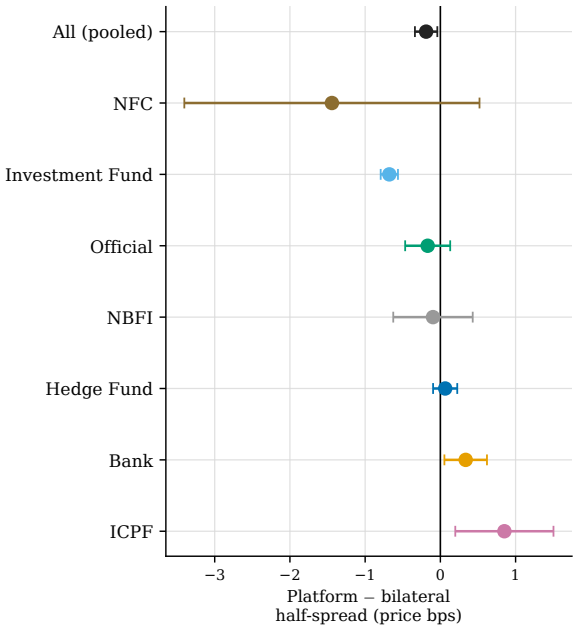
Despite high institutional platform adoption, bilateral and platform trading coexist. The model therefore has to allow platform access to improve dealer competition while preserving non-price relationship values such as account management, balance sheet access, tax and reporting support, or service during stress episodes. This non-price treatment is consistent with the matched home dealer check in Figure A6: for the model's canonical investors, home dealers do

not offer a meaningful effective half-spread discount once we match on client, bond–day, size, and side.

The platform execution advantage Platform access delivers better execution for investors. Within the same client, bond, and day, matched on size and side, institutional clients pay a lower effective half-spread on MTFs than bilaterally. The differential is -0.19 price bps ($t = 2.5$) on the model diagnostic sample and -0.50 bps across the full 34-dealer cross-section, stable in leave-one-out checks (Figure 19). The advantage is concentrated among the genuine buy side—investment funds. It is not mainly an access-to-a-different-dealer effect: holding the dealer fixed leaves most of it intact (-0.18 bps), so the same dealer tends to quote more competitively on the platform.

This diagnostic constrains the model but does not by itself settle welfare. A lower effective half-spread is a transfer from dealer to investor unless it changes matching, participation, or the cost of intermediation; Allen and Wittwer (2023) document the same within-dealer venue wedge in the Canadian market and abstract from it by design (their Fact 2). The structural counterfactual measures welfare as gains from trade, so the platform wedge shows that platforms benefit investors privately, while the welfare result still depends on competition and relationship value.

(a) Matched platform differential, by client sector.



(b) Raw and matched platform differentials.

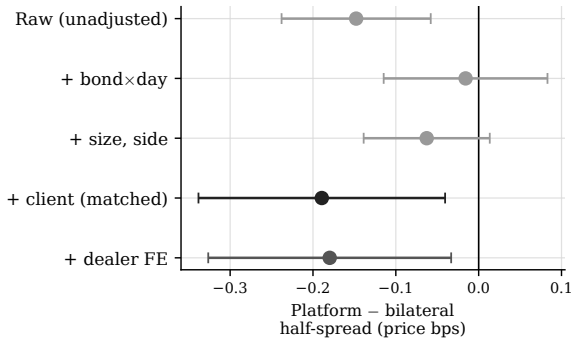


Figure 19: Platform and bilateral execution on the centralization analysis sample. The left panel plots the within-client effective half-spread differential between MTF (platform) and OTC/SI (bilateral) execution, matched on bond×day, trade size, and side, pooled and by client sector (95% ISIN-clustered confidence intervals; negative values mean lower effective half-spreads on the platform). The right panel reports estimates from the unadjusted comparison, the matched within-client specification, and the matched specification with dealer fixed effects. The effective half-spread is the client-signed deviation of the transaction price from the pre-trade MTS mid, in price basis points (negative = inside the mid). Diagnostic sample close to the Allen and Wittwer (2023) model sample (2023–2025: the ten largest dealers by trade count, all MTFs, platform-using clients in institutional sectors, 1,683 clients).

Model outline and estimation The structural model features J dealers and N investors. Dealer j has investor-specific quality μ_{ij} and quotes a platform price equal to value plus a markup. Each institutional investor has a home dealer and chooses between bilateral execution with that dealer and platform execution with the dealer offering the best quote. A loyalty benefit r captures the non-price value of the bilateral relationship. The key parameters are demand sensitivity, dealer qualities, loyalty, liquidity shocks, and platform cost.

We follow Allen and Wittwer (2023) as closely as the Bund data allow: we assign home dealers, use platform prices and dealer shares to discipline competition and quality, and estimate the remaining parameters by the generalized method of moments. The main estimates are: demand sensitivity $\hat{\sigma} = 0.567$, loyalty benefit $\hat{r} = 1.01$ basis points, and platform cost $\hat{c} \approx -0.7$ basis points.[19] Dealer qualities are identified only relative to a reference dealer whose quality is normalized to zero; we use the largest dealer in the estimation sample.[20] For the demand elasticity we import the Canadian calibration of Allen and Wittwer (2023) ($1/\hat{\sigma} = 1.765$) and estimate the rest conditional on it. The Bund data do not identify this parameter: the auction-based instrument that identifies it in Canada has no power in Bunds, plausibly because dealers can lay off auction inventory in futures (Appendix A.6).

Welfare analysis Welfare is measured as model-implied gains from trade: centralization can improve matching by expanding access to efficient dealers, but can lower welfare when platform markups distort matching or investors lose valuable home dealer relationships. Table 4 reports the counterfactuals, which vary platform access (to institutional investors only, or to all investors), platform fees, full centralization, the interpretation of the loyalty benefit, and the intensity of dealer competition; the column and row definitions are given in the table note.

The baseline sign is sharp. In the preferred specification, the home dealer benefit is direct surplus and dealers set platform quotes strategically. Relative to the current platform access regime, forcing all trades onto platforms reduces modeled welfare by 7.73%.

The positive cases in Table 4 are sensitivity exercises, not equally weighted outcomes. Welfare rises under perfect competition because platform access then improves matching without dealer quote distortions. It also rises when the home dealer premium is treated as a switching cost or set to zero, because centralization no longer destroys relationship surplus.[21]

The sign matches the Canadian estimates in Allen and Wittwer (2023); with the largest dealer as reference, the magnitudes are larger. Relative to the current access regime, full centralization

[19] As in Allen and Wittwer (2023), the platform cost enters investors' venue choice but not the welfare measure; in their estimates it is close to the platform fee, a transfer. The loyalty benefit, by contrast, is the value of the relationship service itself. The estimated loyalty benefit is substantially smaller than the Canadian estimate of $\hat{r} = 2.37$ basis points, consistent with the higher platform adoption rate in Bunds. Dealer quality varies substantially across buy and sell sides and between institutional and retail investors.

[20] Allen and Wittwer (2023) report results for a single reference dealer. Because the estimated welfare changes depend on this choice, Table A3 also reports the centralization counterfactual with each of the second- to fourth-largest dealers as reference. The choice moves the welfare difference itself, not only the welfare of the fully bilateral market in the denominator.

[21] The interpretation of the switching cost and no-loyalty rows warrants caution. As noted by Allen and Wittwer (2023) in footnote 22, these rows cannot be interpreted as a policy recommendation because removing loyalty is akin to removing costs from a profit calculation.

		Account I	Account IR	No Fee	Central
With loyalty	Imperfect	−6.01	−10.08	−11.38	−13.28
	Perfect	+9.07	+13.51	+14.52	+15.80
Switching cost	Imperfect	+19.21	+18.76	+22.29	+27.22
	Perfect	+28.21	+36.92	+40.55	+45.17
No loyalty	Imperfect	+26.19	+26.55	+29.54	+32.78
	Perfect	+50.44	+63.03	+66.99	+71.38

Table 4: Welfare counterfactuals relative to a fully decentralized market. Preferred-spec estimate following the Allen and Wittwer (2023) framework using Bund transaction data. Entries are percentage welfare changes relative to a fully bilateral market. Imperfect competition means dealers set quotes optimally; perfect competition means dealers quote at true values ($q = v$). Account I gives platform access to institutional investors only. Account IR gives platform access to all investors. No Fee sets platform fees to zero. Central forces all trades onto the platform. With loyalty means the home dealer benefit r enters welfare directly. Switching cost means r affects choices but is not valued in welfare. No loyalty sets $r = 0$.

lowers welfare by 7.73% in Bunds and by 1.75% in Canadian government bonds. Against the fully bilateral benchmark of Table 4, the current access regime is 6.01% below it in Bunds and full centralization 13.28%, compared with 1.29% and 3.02% in Canadian government bonds. The magnitude is not pinned down by the model alone. Relative to a fully bilateral market, centralization lowers welfare unless the demand elasticity exceeds $1/\hat{\sigma} = 2.94$ to 3.09, at or just above the upper end of the 95% confidence interval of the imported estimate ($1/\hat{\sigma} = 1.765$, standard error 0.599); Table A3 shows that the size of the loss also depends on the reference dealer.

Figure 20 plots every counterfactual in Table 4 against the corresponding estimate in Table 6 of Allen and Wittwer (2023). Every Bund cell has the same sign as its Canadian counterpart and, with the largest dealer as reference, a larger magnitude.

Discussion The model's policy implication is therefore conditional, not a general ranking of market designs. Taking estimated relationship value, imperfect dealer competition, and the benchmark imported demand elasticity at face value, mandatory platform centralization would lower modeled welfare. The positive rows identify the assumptions under which that conclusion reverses, and the sensitivity exercises show that the imported elasticity is especially influential. The model is suggestive about resilience, not a resilience test. In stress, centralized and competitive platforms may help high-value dealers absorb or redistribute inventory across a broader client base; this lies outside the partial equilibrium dealer-to-client model and should be read alongside the separate mechanical netting exercise. Three limitations matter, and their direction of bias is unclear. The institutional settings for Canadian bonds and Bunds differ. The dealer-to-client model captures only a partial equilibrium analysis for Bunds, which are part of a large ecosystem of futures and dealer-to-dealer trading. Lastly, the model assumes a single bond whereas the Bund market consists of many bonds with different maturities, benchmark status, and futures-deliverability status.

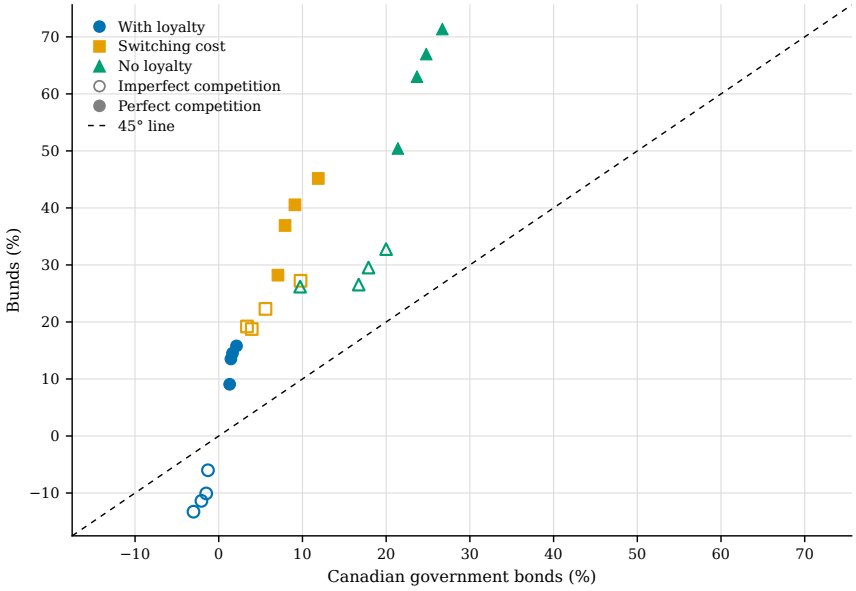


Figure 20: Welfare counterfactuals for Bunds and Canadian government bonds. Each point pairs one welfare change for Bunds from Table 4 with the corresponding Canadian government bond estimate in Table 6 of Allen and Wittwer (2023), in percent relative to a fully bilateral market. The horizontal axis reports Canadian estimates and the vertical axis reports Bund estimates. Hollow markers denote imperfect competition; filled markers denote perfect competition. The dashed line is the 45-degree line.

7 Conclusion

This paper studies the German Bund market using a new non-anonymous transaction-level dataset combining MiFID I and MiFID II transaction reports spanning 2008–2025. The data identify both sides of each trade and reveal the euro area's benchmark safe asset market structure. Trading is hierarchical around a small dealer core, relationship-based between dealers and clients, and segmented across investor sectors, with order flow fragmented across venues. We will maintain the cleaned dataset and its documentation as a reusable resource for authorized research within the Bundesbank.

That structure shapes stress flows, price discovery, and execution costs. Stress flows differ by shock: foreign official accounts sell and hedge funds buy in both the Covid dash for cash and the 2023 flight to quality. Primary dealers absorb material flow only in the cash shock and distribute Bunds after the 2025 fiscal news. In normal times, hedge funds and official sector accounts carry the most permanent price-relevant flow, but they do not pay the highest effective half-spreads. Retail households pay much more, and large trades pay a size penalty. The Bund market therefore rejects a simple adverse selection account in which the most price-relevant flow is charged the highest execution cost.

The reform implications follow from these facts. Mandatory central clearing is arithmetic: given observed trades, netting each dealer's book at a clearing house would cut dealers' net settlement obligations by 35%, less than comparable estimates for U.S. Treasuries and U.K. Gilts. Platform centralization changes dealer access, relationship surplus, and competition in equilibrium. In the benchmark structural counterfactual, forced platform centralization lowers

welfare: broader dealer access does not compensate for the loss of relationship surplus under imperfect dealer competition, in Bunds as in Canadian government bonds. Even in a benchmark safe asset market, market structure determines who moves prices, who pays execution costs, and which reforms create gains.

References

Allen, J., A. Hortaçsu, E. Richert, and M. Wittwer (2026). Entry and Exit in Treasury Auctions. Working Paper 34646, National Bureau of Economic Research.

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

Allen, J. and M. Wittwer (2025). Bundling Trades in Over-the-Counter Markets. Working paper, Columbia Business School.

Aragon, G. O. and P. E. Strahan (2012). Hedge Funds as Liquidity Providers: Evidence from the Lehman Bankruptcy. *Journal of Financial Economics 103*(3), 570–587.

Bai, J., E. Bostrom, S. Infante, and V. Ivashina (2026). Liquidity Flows to Bank-Affiliated Broker Dealers: Insights from Volumes and Prices. *NBER Working Paper 34844*.

Baranova, Y., E. Holbrook, D. MacDonald, W. Rawstorne, N. Vause, and G. Waddington (2023). The potential impact of broader central clearing on dealer balance sheet capacity: a case study of UK gilt and gilt repo markets. Staff Working Paper 1026, Bank of England.

Beck, R., A. Coppola, A. J. Lewis, M. Maggiori, M. Schmitz, and J. Schreger (2024). The Geography of Capital Allocation in the Euro Area. NBER Working Paper 32275, National Bureau of Economic Research.

Beetsma, R., M. Giuliodori, J. Hanson, and F. de Jong (2018). Cross-Border Auction Cycle Effects of Sovereign Bond Issuance in the Euro Area. *Journal of Money, Credit and Banking 50*(7), 1401–1440.

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.

Biais, B. and R. Green (2019). The microstructure of the bond market in the 20th century. *Review of Economic Dynamics 33*, 250–271.

Brandt, M. W. and K. A. Kavajecz (2004). Price discovery in the US Treasury market: The impact of orderflow and liquidity on the yield curve. *Journal of Finance 59*(6), 2623–2654.

Branger, N., M. Muck, and A. Pütz (2024). Unmasking global investors: An examination of the shifting dynamics in sovereign bond trading. Working Paper 5057388, SSRN Working Paper. SSRN preprint, posted 15 December 2024; co-authored at the Deutsche Finanzagentur.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2017). The Safe Assets Shortage Conundrum. *Journal of Economic Perspectives 31*(3), 29–46.

Chaboud, A. P., E. Correia Golay, C. Cox, M. J. Fleming, Y. Huh, F. M. Keane, K. Lee, K. B. Schwarz, C. Vega, and C. Windover (2025). All-to-all trading in the U.S. treasury market. *Economic Policy Review 31*(2), 1–27. Federal Reserve Bank of New York.

Cheung, Y. C., B. Rindi, and F. De Jong (2005). Trading European sovereign bonds: the microstructure of the MTS trading platforms. *Available at SSRN 424936*.

Coeurdacier, N. and H. Rey (2013). Home Bias in Open Economy Financial Macroeconomics. *Journal of Economic Literature 51*(1), 63–115.

Colliard, J.-E., T. Foucault, and P. Hoffmann (2021). Inventory management, dealers' connections, and prices in over-the-counter markets. *Journal of Finance 76*(5), 2199–2247.

Czech, R., S. Huang, D. Lou, and T. Wang (2021). Informed trading in government bond markets. *Journal of Financial Economics 142*(3), 1253–1274.

de Roure, C., E. Moench, L. Pelizzon, and M. Schneider (2026). OTC Discount. *Management Science 72*(6), 4633–4651.

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

Duffie, D. (2020). Still the World's Safe Haven? Redesigning the U.S. Treasury Market After the COVID-19 Crisis. Hutchins Center Working Paper 62, Brookings Institution.

Duffie, D. (2023). Resilience redux in the US Treasury market. *Jackson Hole Symposium, Federal Reserve Bank of Kansas City*.

Duffie, D. (2025). How U.S. Treasuries can remain the world's safe haven. *Journal of Economic Perspectives 39*(2), 195–214.

Duffie, D., M. J. Fleming, F. M. Keane, C. Nelson, O. Shachar, and P. Van Tassel (2023). Dealer capacity and US Treasury market functionality. *FRB of New York Staff Report* (1070).

Duffie, D., S. F. Singh, and C. Wang (2026). An Efficient Liquidity Savings Mechanism. Working paper, Graduate School of Business, Stanford University.

Duffie, D. and C. Wang (2017). Efficient Contracting in Network Financial Markets. Working paper, Graduate School of Business, Stanford University.

Dugast, J., S. Üslü, and P.-O. Weill (2022). A Theory of Participation in OTC and Centralized Markets. *The Review of Economic Studies 89*(6), 3223–3266.

Eisfeldt, A. L., B. Herskovic, S. Rajan, and E. Siriwardane (2023). OTC Intermediaries. *Review of Financial Studies 36*(2), 615–677.

European Parliament and Council of the European Union (2019). Regulation (EU) 2019/876 of 20 May 2019 amending Regulation (EU) No 575/2013 as regards the leverage ratio, the net stable funding ratio, requirements for own funds and eligible liabilities, counterparty credit risk, market risk, exposures to central counterparties, exposures to collective investment undertakings, large exposures, reporting and disclosure requirements. Official Journal of the European Union, L 150, 7 June 2019, pp. 1–225.

European Securities and Markets Authority (2016). Guidelines on Transaction Reporting, Order Record Keeping and Clock Synchronisation under MiFID II. Guidelines ESMA/2016/1452, European Securities and Markets Authority.

European Systemic Risk Board (2024). A System-Wide Approach to Macroprudential Policy: ESRB Response to the European Commission's Consultation Assessing the Adequacy of Macroprudential Policies for Non-Bank Financial Intermediation. Report, European Systemic Risk Board.

Evans, M. D. and R. K. Lyons (2008). How is macro news transmitted to exchange rates? *Journal of Financial Economics 88*(1), 26–50.

Filipović, D., M. Pelger, and Y. Ye (2025). Stripping the discount curve—a robust machine learning approach. *Management Science 72*(9), 8057–8120.

Fleming, M. J. and F. M. Keane (2021). The netting efficiencies of marketwide central clearing. *FRB of New York Staff Report* (964).

Fleming, M. J. and E. M. Remolona (1999). Price formation and liquidity in the US Treasury market: The response to public information. *Journal of Finance 54*(5), 1901–1915.

Gorton, G. (2017). The History and Economics of Safe Assets. *Annual Review of Economics 9*(1), 547–586.

Green, R. C., B. Hollifield, and N. Schürhoff (2007). Dealer Intermediation and Price Behavior in the Aftermarket for New Bond Issues. *Journal of Financial Economics 86*(3), 643–682.

Guimaraes, R., S. Jurkatis, G. Pintér, and J.-C. Wijnandts (2025). Government bond market making: Man vs machine. Work in progress.

Hagströmer, B. (2022). Market Fragmentation in Europe. *Available at SSRN*.

Hagströmer, B. and A. J. Menkveld (2019). Information revelation in decentralized markets. *Journal of Finance 74*(6), 2751–2787.

Hagströmer, B. and A. J. Menkveld (2023). Trades, Quotes, and Information Shares. *Available at SSRN*.

Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *Journal of Finance 50*(4), 1175–1199.

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *American Economic Review 103*(2), 732–70.

He, Z., A. Krishnamurthy, and K. Milbradt (2019). A Model of Safe Asset Determination. *American Economic Review 109*(4), 1230–1262.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020). Relationship Trading in Over-the-Counter Markets. *Journal of Finance 75*(2), 683–734.

Hendershott, T., D. Livdan, and N. Schürhoff (2015). Are institutions informed about news? *Journal of Financial Economics 117*(2), 249–287.

Hollifield, B., A. Neklyudov, and C. Spatt (2017). Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations. *Review of Financial Studies 30*(9), 3048–3085.

Hortaçsu, A., J. Kastl, and A. Zhang (2018). Bid Shading and Bidder Surplus in the US Treasury Auction System. *American Economic Review 108*(1), 147–169.

Hu, G. X., J. Pan, and J. Wang (2013). Noise as Information for Illiquidity. *Journal of Finance 68*(6), 2341–2382.

Huh, Y. and S. Infante (2021). Bond market intermediation and the role of repo. *Journal of Banking & Finance 122*, 105999.

Joint Staff Report (2015). The U.S. Treasury Market on October 15, 2014. Technical report, U.S. Department of the Treasury and Board of Governors of the Federal Reserve System and Federal Reserve Bank of New York and U.S. Securities and Exchange Commission and U.S. Commodity Futures Trading Commission.

Jurkatis, S. (2024). An approach to cleaning MiFID II corporate bond transaction reports. *Bank of England Staff Working Paper No. 1,071.*

Kerssenfischer, M. (2026). The On-the-Run Discount: Tap Issuance and Collateral Supply in German Bunds. *Bundesbank Discussion Paper forthcoming.*

Kerssenfischer, M. and C. Helmus (2025). Outages in sovereign bond markets. Working paper, revised December 2025.

Kerssenfischer, M. and M. Schmeling (2024). What Moves Markets? *Journal of Monetary Economics 145*, 103560.

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

Kutai, A., D. Nathan, and M. Wittwer (2025). Exchanges for Government Bonds? Evidence During COVID-19. *Management Science 71*(10), 8948–8966.

Lee, T., D. Nathan, and C. Wang (2024). China Walls. Working paper, SSRN. SSRN preprint.

Lee, T. and C. Wang (2025). Regulating Over-the-Counter Markets. *Journal of Finance 80*(4), 1929–1962.

Lee, T., C. Wang, and A. Zawadowski (2025). The Modern Bond Market. Working paper, SSRN. SSRN preprint.

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

Lou, D., H. Yan, and J. Zhang (2013). Anticipated and Repeated Shocks in Liquid Markets. *Review of Financial Studies 26*(8), 1891–1912.

Maggiori, M., B. Neiman, and J. Schreger (2020). International Currencies and Capital Allocation. *Journal of Political Economy 128*(6), 2019–2066.

Menkveld, A. J. (2013). High Frequency Trading and the New Market Makers. *Journal of Financial Markets 16*(4), 712–740.

Menkveld, A. J., B. Z. Yueshen, and H. Zhu (2017). Shades of darkness: A pecking order of trading venues. *Journal of Financial Economics 124*(3), 503–534.

Mizrach, B. and C. J. Neely (2008). Information shares in the US Treasury market. *Journal of Banking & Finance 32*(7), 1221–1233.

Pasquariello, P. and C. Vega (2007). Informed and Strategic Order Flow in the Bond Markets. *Review of Financial Studies 20*(6), 1975–2019.

Perrella, A. and J. Catz (2020). Integrating microdata for policy needs: the ESCB experience. *ECB Statistics Paper* (33).

Pintér, G. (2023). An anatomy of the 2022 gilt market crisis. *Bank of England Staff Working Paper No. 1019*.

Pintér, G., E. Siriwardane, and D. Walker (2026). Fire Sales of Safe Assets. *Journal of Financial Economics*. Forthcoming.

Pintér, G., C. Wang, and J. Zou (2022). Information Chasing versus Adverse Selection. Staff Working Paper 971, Bank of England.

Pintér, 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.

Portes, R. and H. Rey (2005). The determinants of cross-border equity flows. *Journal of International Economics 65*(2), 269–296.

Svensson, L. E. (1994). Estimating and interpreting forward interest rates: Sweden 1992-1994. *National Bureau of Economic Research Working Paper No. 4871*.

Timmer, Y. (2018). Cyclical investment behavior across financial institutions. *Journal of Financial Economics 129*(2), 268–286.

Üslü, S. (2019). Pricing and Liquidity in Decentralized Asset Markets. *Econometrica 87*(6), 2079–2140.

Wang, C. (2016). Core-Periphery Trading Networks. Working paper, The Wharton School, University of Pennsylvania.

Wang, C. (2022). Search in Over-the-Counter Markets. Working paper; revise and resubmit, Journal of Finance.

Wang, C. (2023). The Limits of Multi-Dealer Platforms. *Journal of Financial Economics 149*(3), 434–450.

Wittwer, M. and A. Uthemann (2025). Specialization in Financial Markets. Working paper.