Bunds

Mark Kerssenfischer
Deutsche Bundesbank

Konrad Lucke
Goethe University

Willy Scherrieble
Deutsche Bundesbank

March 2026

Abstract

This paper assembles a comprehensive non-anonymous transaction-level dataset for German Bunds. We document key market structure facts, present new evidence on cross-sectional Bund pricing, and evaluate two prominent reform proposals. The data reveal a market that is fragmented across venues and bonds, segmented across investor groups, and hierarchical around a dealer core. Liquidity concentrates in a small set of recently issued, futures-deliverable bonds, while many Bunds trade only sporadically and often not at all on a given venue-day. Trading is dispersed across hundreds of venues; even bilateral OTC, the single largest venue type, accounts for only a fraction of activity. Repo specialness is capitalized into cash prices, but heavily traded Bunds are cheap: on-the-run Bunds are less special due to supply expansions from frequent re-issuances and hence trade at a discount. Studying a mandatory CCP scenario, we find larger potential benefits for Bunds than for U.S. Treasuries and U.K. Gilts. Using a structural model of dealer–investor trading, we show that venue centralization generates welfare gains only when dealer competition is sufficiently strong.

Keywords: sovereign bonds, market microstructure.

JEL classification: G12, G14, G23

Contact address: mark.kerssenfischer@bundesbank.de and willy.scherrieble@bundesbank.de. We thank Jean-Charles Wijnandts (disussant) and participants at the 14th Bundesbank Term Structure Workshop. 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.

1 Introduction

Sovereign bonds sit at the core of the financial system. Yet recent stress episodes have exposed cracks: Covid in 2020, the Gilt crisis in 2022, and Liberation Day in 2025. Each showed how even 'safe' sovereign bond markets can seize up, and how fragile their plumbing can be. These events have raised new questions about market design, regulation, and reform.[1] Answering them requires clean and comprehensive data.

This paper (i) provides the most comprehensive transaction-level dataset on the German Bund market to date, (ii) documents key features of its market structure, (iii) presents new evidence on cross-sectional Bund pricing, and (iv) evaluates two prominent reform proposals.

Using millions of non-anonymous MiFIR transaction reports, we construct and share a cleaned dataset suitable for economic analysis. Because regulatory trade reports are not designed for research use, the raw data contain duplicates, heterogeneous reporting formats, and reporting errors. We therefore develop a transparent sequence of normalization, aggregation, and de-duplication steps that exploit institutional features of the reporting regime and recover a reliable transaction record. We enrich these trades with bond-, investor-, and venue-level information and link them to transaction-level data from the repo and futures market for Bunds.

The data reveal a market that is *fragmented*, *segmented*, and *hierarchical*. Trading concentrates in a small set of bonds—especially recently issued on-the-run bonds that are deliverable into Bund futures—while most outstanding bonds trade only sporadically. At the same time, activity is dispersed across a large number of venues, with no single platform dominating overall volume, and with widespread bond–venue inactivity on a typical day. Market access is uneven: different investor sectors systematically concentrate their activity in different venue types, reflecting segmentation in participation and trading technology. Finally, the trading network features a tight core of primary dealers that intermediate most activity and connect a broad periphery of smaller participants.

Our cross-sectional pricing results reveal a positive relationship between repo rates and cheapness. Bonds with lower repo rates are richer in the cash market, consistent with the view that repo specialness reflects a convenience yield that is capitalized into bond prices. Surprisingly, however, there is a similar positive relationship between trading activity and cheapness: bonds with higher trading turnover are *cheaper* in the cash market. We trace this counter-intuitive pattern primarily to on-the-run (OTR) Bunds. The outstanding volume of OTR bonds is regularly increased through tap auctions. These supply expansions reduce repo specialness and depress cash prices, as investors anticipate the additional collateral supply. This supply effect outweighs liquidity or search-driven demand effects, resulting in an OTR discount in Bunds, in contrast to the OTR premium observed in U.S. Treasuries. A second force shaping cross-sectional pricing arises from futures deliverability. Cheapest-to-deliver (CTD) bonds become highly special in repo markets close to futures delivery, reflecting strong delivery-related borrowing demand.

[1] Proposed and ongoing reforms include mandatory central clearing of sovereign cash and repo trades by 2027 in the United States (SEC); faster settlement, with the EU moving from T+2 to T+1 by October 2027 (ESMA); venue consolidation, including calls for a single European stock exchange (Reuters); a consolidated tape for bonds in the EU (ESMA); and adjusted capital regulation affecting dealer intermediation capacity (Federal Reserve).

Unlike the persistent issuance effects driving the OTR discount, however, this delivery squeeze is short-lived and largely confined to repo markets.

Finally, we use the granular transaction data to evaluate current reform proposals. We conduct a mandatory CCP scenario analysis that quantifies the mechanical reduction in dealers' gross and net settlement obligations under central clearing, following the logic of multilateral netting through a central counterparty. The potential reductions we estimate for Bunds exceed those documented by Fleming and Keane (2021) for U.S. Treasuries and Baranova, Holbrook, MacDonald, Rawstorne, Vause, et al. (2023) for U.K. Gilts. Beyond central clearing, we investigate the welfare effects of centralizing bilateral trading onto electronic platforms. Replicating the structural model of Allen and Wittwer (2023) for the Bund market, we find that the qualitative patterns documented for Canadian government bonds carry over: market centralization decreases welfare when dealers have market power and investors value bilateral relationships. Only under perfect competition does centralization unambiguously lead to positive welfare effects.

2 Data

2.1 Transaction-Level

Our transaction-level dataset for German sovereign bonds ('Bunds') is based on MiFIR transaction reports. MiFIR provides near-complete market coverage but is not designed for economic analysis. The raw data contain duplicated records, heterogeneous reporting formats, and reporting errors, all of which are amplified by the relatively high trading frequency in the Bund market. We therefore apply a sequence of normalization and deduplication steps. This section summarizes the main design choices; full technical details and diagnostics are relegated to Appendix C. To ensure reproducibility and to foster further research in this area, we will share the final dataset with eligible researchers.

Under MiFIR, investment firms must report each bond transaction to their national competent authority. For German sovereign bonds, all transactions are reported to BaFin, irrespective of trading venue. Each report contains the ISIN, execution timestamp, price, quantity, execution venue or off-venue indicator, and identifiers for the buyer, seller, and executing firm. We restrict the sample to plain-vanilla German sovereign bonds and exclude short-term Bubills and any non-standard bonds. The resulting sample contains roughly 10 million transaction reports for 130 bonds over eight years (2018-2025).

Raw MiFIR reports require extensive harmonization. We standardize intermediary trades, aggregate mechanically related reports, conservatively screen for implausible prices and quantities, and identify duplicate reports arising from two-sided MiFIR reporting. Duplicate matching exploits institutional features of the reporting regime and uses one-to-one maximum-weight matching across prices, quantities, timestamps, venues, and counterparties. Matched duplicates are used to correct reporting inconsistencies. The final dataset is a normalized transaction-level record of Bund trading suitable for economic analysis. We enrich the MiFIR transactions as

follows.

2.2 Bond-Level

We obtain bond-level information from the Finanzagentur: issue day, maturity day, first coupon day, dated day (start of accrual), and issuance volume.[2]

Each day, we identify which bonds are deliverable into any of the four Bund future front-end contracts at the time. To be eligible for the 2-year FGBS contract, a bond's remaining maturity must lie between 1.75 and 2.25 years, measured as of the futures' delivery day; for the 5-year FGBM contract, the maturity must lie between 4.5 and 5.5 years; for the 10-year FGBL contract, between 8.5 and 10.5 years; and for the 30-year FGBX contract, between 24 and 35 years. Except for the 30-year contract, the bond must have an original maturity of no more than 11 years in order to be deliverable. Since 1 December 2022, the bond must also meet a minimum day requirement: bonds must be issued prior to the first calendar day of the previous delivery month to be deliverable into the current future contract.[3] To identify the bonds that are cheapest-to-deliver out of all deliverable bonds, we use Bloomberg data. We further use auction history data published by the Finanzagentur to determine on-the-run (OTR) bonds and 'taps', i.e. re-issuances of existing bonds.

2.3 Investor-Level

We construct a comprehensive investor reference dataset for the full universe of over three million Legal Entity Identifiers (LEIs) from GLEIF.[4] We map LEIs into eleven sectors: primary dealers (participants in the primary auctions of Bunds); banks (commercial, investment, savings, and cooperative banks); inter-dealer brokers (IDB; voice and electronic brokers intermediating between dealers); investment funds (UCITS, money market funds, ETFs, and asset managers); hedge funds; insurance companies and pension funds (ICPF); central counterparties (CCPs); non-bank financial intermediaries (NBFI; e.g. 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).

Appendix C.8 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.

Investor Grouping To consolidate ownership hierarchies across entities, we construct a comprehensive investor grouping based on the Global Legal Entity Identifier Foundation (GLEIF)

[2] Finanzagentur bond metadata are more comprehensive and reliable than the Centralised Securities Database (CSDB). Hence, we use the CSDB only to obtain monthly outstanding volumes of each bond.

[3] From September 2022 onwards, we use the official list of deliverable bonds webscraped from the Eurex website. Prior to that, we re-construct the deliverability status by applying the above-mentioned criteria to the universe of outstanding Bunds.

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

Level 2 'Who owns Whom' relationship data.[5] This dataset records bilateral relationships between entities and their direct and ultimate parents, together with their relationship status and validity periods. All valid and active relationships are collected into a unified table of parent–child links.

We first prioritize ultimate parent relationships to ensure a flat hierarchy whenever both direct and ultimate parent information is available. Entities without an ultimate parent but with a valid direct parent are added next. The resulting mapping is filtered to include only general and branch entities, excluding other GLEIF categories such as international branches, funds, or umbrella structures.

Additional group relationships are then incorporated from ECB sources. Specifically, we add parent mappings from the lists of Significant Institutions (SIs) and Monetary Financial Institutions (MFIs), restricted to banks. Duplicates are removed to form the final set of parent–child relationships.

Next, a recursive function traverses the ownership chain to identify the ultimate parent for each entity, thereby assigning a unique group identifier to all members of the same group. Circular ownerships are detected and resolved by terminating recursion at the first repetition. The result is a complete mapping from every LEI to its group head.

Entities without any valid ownership links are treated as singletons, forming their own group. Consistency checks ensure that no entity belongs to multiple groups and that every group head corresponds to an existing LEI.

Finally, the group mapping is merged back into the main investor reference dataset. Each entity thus carries both its individual attributes and the characteristics of its group head—specifically, the group-level name, sector, and country. This step harmonizes the classification of investors operating within the same corporate structure and ensures that analytical groupings reflect consolidated ownership.

2.4 Venue-Level

We combine three external reference datasets on trading venues: the ESMA Financial Instruments Transparency System (FITS) register,[6] the UK Financial Conduct Authority (FCA) register,[7] and the ISO 10383 list of Market Identifier Codes (MICs).[8] All sources provide venue identifiers (MIC), legal entity identifiers (LEI), names, and venue type classifications.

Our approach regarding venue types is hierarchical. We use the ESMA register as the primary source. For MIC codes not classified by ESMA, we rely on the FCA register. For the remaining venues, we turn to ISO data. We harmonize venue types to ESMA taxonomy. This procedure ensures we cover all MICs present in MiFIR.

A small number of venues only listed in the ISO data do not carry explicit type classifications. For these, we apply another hierarchical classification: (i) we match the LEI code of each

[5] https://www.gleif.org/en/lei-data/access-and-use-lei-data/level-2-data-who-owns-whom
[6] https://registers.esma.europa.eu/publication/searchRegister?core=esma_registers_upreg
[7] https://register.fca.org.uk/servlet/servlet.FileDownload?file=0150X000006gbb1
[8] https://www.iso20022.org/sites/default/files/ISO10383_MIC/ISO10383_MIC.xlsx

MIC code to our sector classification (classifying venues run by inter-dealer brokers as OTF and those run by primary dealers as SI), (ii) we parse through the 'description' field for unambiguous keywords ('systematic internaliser', 'approved publication arrangement', 'electronic trading platform', and 'organised trading facility'), and (iii) we classify based on well-known institution names (e.g., Tradeweb and Bloomberg mapped to MTF, and major IDBs mapped to OTF). Only very few insignificant MIC codes remain unclassified after this procedure.

This process yields a harmonized classification of trading venues to venue types and a consistent mapping of MIC to LEI codes.

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 (e.g., Bloomberg, Tradeweb, MarketAxess). These 'sibling venues' largely reflect regulatory and reporting requirements rather than 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.

2.5 Cash Market Pricing

To measure (mis)pricing on the Bund market, we first synchronize prices to account for the fact that trades occur throughout the day. For each trade, we retrieve the bond's last mid-quote on MTS prior to the trade and prior to 11:45 a.m. on that day and add the mid-quote change to the observed transaction price.[9] We then convert each synchronized price into an implied yield, using an internal rate of return solver, correctly accounting for the $T+2$ settlement cycle, TARGET holidays, and irregular first or last coupon periods (stubs).

We compare these transaction yields to three distinct fair yield benchmarks. Our first benchmark uses the official Bundesbank term structure based on the Svensson (1994) model. We use that model's zero-coupon spot rates to discount each bonds' cash flows into a fair price and convert this to a fair yield. We use these fair yields as a global filter for all methods: any transaction where the implied yield deviates from the maturity-matched fair value by 25 basis points or more is excluded from our analysis. These outliers represent likely reporting errors. We further exclude bonds with less than 0.25 years of remaining maturity and when-issued trades. These filters drop less than 1% of total observations.

As a second benchmark, we implement the non-parametric Kernel Ridge approach by Filipović, Pelger, and Ye (2025), which applies the above-mentioned discounting framework with

[9]This approach follows Duffie, Fleming, Keane, Nelson, Shachar, and Van Tassel (2023). Just like them, we require each trade to have a quote for the same bond within at least one hour before the trade to be included in the sample. The actual median time difference is only a few seconds. Note also that this filter effectively restricts our sample to trades between 8:00 a.m. and 7:00 p.m., the trading hours on MTS. We choose a 11:45 a.m. fixing for the Bund yield curve for two reasons. First, this choice aligns with the official Bundesbank yield curve estimation, see Bundesbank website. Second, 11:45 a.m. CET is a relatively low-news period (corresponding to 5:45 a.m. US ET), see Kerssenfischer and Schmeling (2024). US Treasury curves are typically fitted based on CRSP end-of-day quotes, which refer to 4 p.m. US ET from October 16, 1996, till September 3rd 2024, and to 5 p.m. since then, see CRSP documenation.

state-of-the-art machine learning methods, see Appendix B for details. Third and lastly, we apply a cubic spline interpolation directly to the implied yields across the maturity spectrum. We apply both methods day-by-day using the volume-weighted average yield of each bond.

Lastly, we implement a leave-one-out interpolation approach for bonds flagged as on-the-run (OTR), eligible for futures delivery, or cheapest-to-deliver (CTD). For these bonds, we compute an additional 'fair yield' benchmark via linear maturity-based interpolation using the two nearest neighboring bonds that do not share that specific bond's status.

Our final day-bond level 'cheapness' measure is the volume-weighted transaction-implied yield minus the average 'fair yield' across all of the above-mentioned methods. Positive values mean the bond is cheap, having an excessively high yield (low price). Negative values mean the bond is rich.

2.6 Repo and Futures market

We augment our cash market data with transaction-level data from the repo market. The data source is the euro area wide version of the money market statistical reporting (MMSR) dataset. The data covers all overnight repo trades using German Bunds as collateral.

In the MMSR data, the same repo transaction may be reported twice, once as a borrowing and once as a lending trade, when both counterparties are reporting institutions. We remove these duplicates by matching observations that share the same deal rate, collateral ISIN, and transaction nominal amount. For non-centrally cleared trades, we additionally require that the reporting agent and counterparty identifiers are reversed across the matched pair. For centrally cleared trades, where the counterparty is the clearing house on both sides, matching is based on trade characteristics and the clearing house identifier. In each case, matched pairs are treated as duplicate reports of the same transaction, and only one observation is retained.

Lastly, we augment data from Bund futures, which trade on Eurex. In particular, we have daily data on trading volumes and open interest for each of the four Bund futures. We always use the closest-to-maturity contract except on its last trading day, when we roll-over to the next-nearest contract.

3 Stylized Market Structure Facts

3.1 Dealer Dominance

Figure 1 provides a network graph of the Bund market. This figure empirically confirms and quantifies the often asserted importance of dealers in the microstructure of sovereign bond markets; see Bessembinder, Spatt, and Venkataraman (2020) for a recent survey. A handful of dealers sit in the center of the trading network, with thousands of smaller clients in the periphery. Most trading is intermediated by dealers, and only few clients trade directly with each other. Recall that we define dealers as all primary dealers participating in Bund auctions, while all other investors are deemed to be clients.

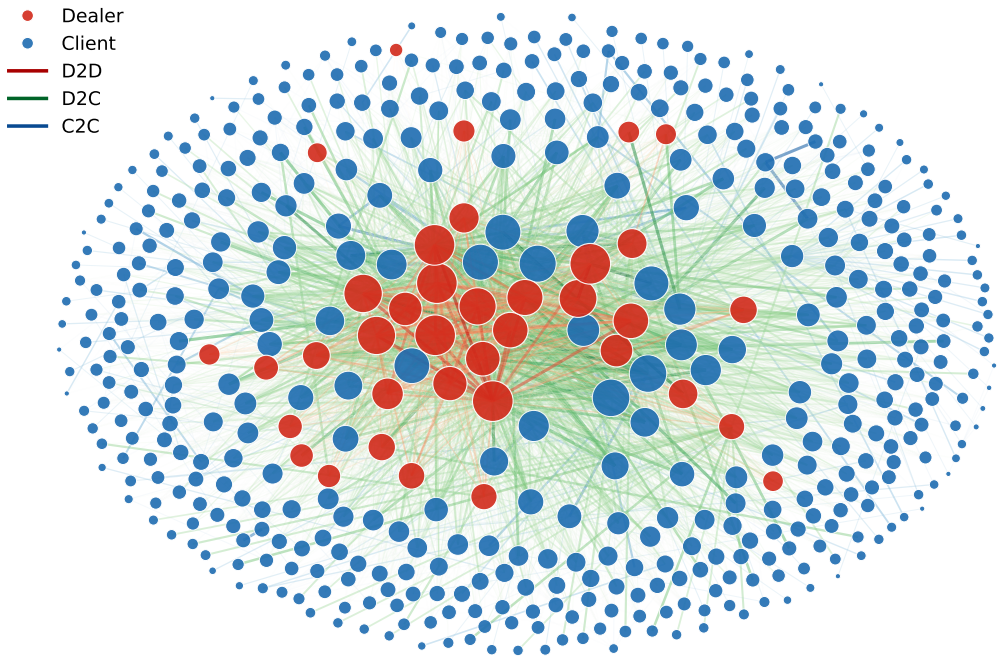


Figure 1: The Bund market network. Undirected graph of top 500 investors (by degree) with at least 50 trades per edge. Node position reflects eigenvector centrality. Node size scales with centrality, with top 5 investors sharing maximum value to ensure confidentiality. Edges scale with the log number of trades between each investor pair. Intragroup trades are omitted.

Figure 2 shows how the extent of dealer intermediation differs across trading venue types. On organised trading facilities (OTF, capturing inter-dealer broker platforms), dealers account for 75% of trading volume. This share drops only slightly for systematic internaliser (single-dealer) platforms and for bilateral OTC trades. On multilateral trading facilities (e.g. Bloomberg and Tradeweb), primary dealers intermediate roughly 50% of volume. Only on regulated markets (regular exchanges) are primary dealers rather unimportant, intermediating only 15% of volume.

3.2 Segmentation

The Bund market is segmented, meaning different investors (have to) trade on different venues. Figure 3 confirms that sector activity is unevenly distributed across venue types. Dealers are active across virtually all venues except exchanges, inter-dealer brokers concentrate on OTFs, and funds and official institutions predominantly trade on multilateral platforms. Retail activity is largely confined to exchanges. Segmentation across venues is consistent with fixed-income evidence that trading behavior differs systematically across investor types and that investor identity interacts strongly with trading costs (Czech, Huang, Lou, and Wang, 2021; Pintér, Wang, and Zou, 2024).

3.3 Concentration and Fragmentation

Trading in Bunds is *concentrated* in a handful of bonds but *fragmented* across trading venues.

Figure 4 shows that the most-traded bond on a given day unites roughly 15% of total volume.

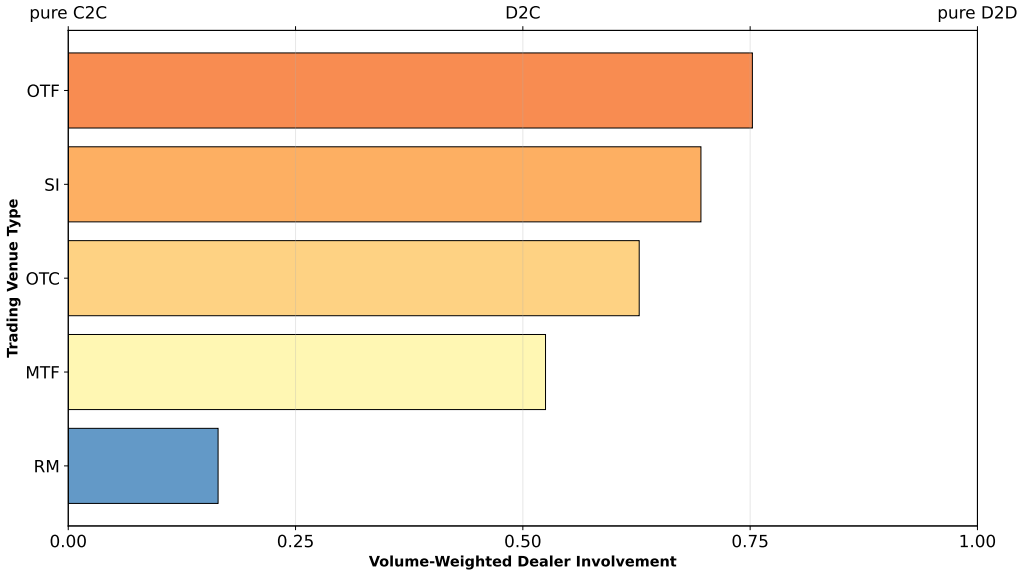


Figure 2: Dealer involvement across venue types. Share of trading volume involving at least one primary dealer.

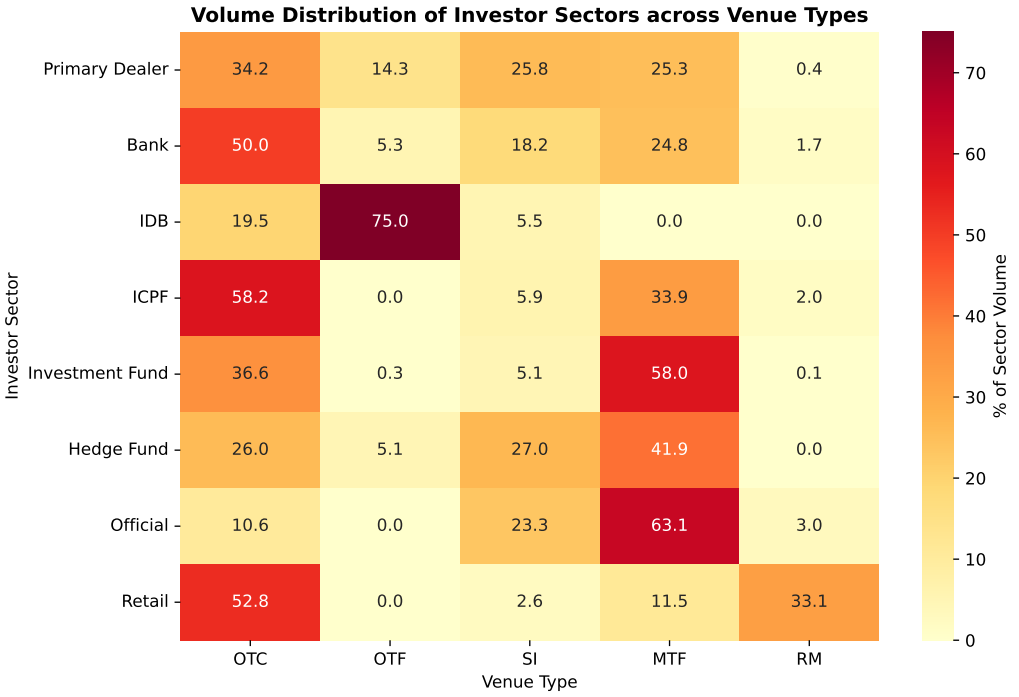


Figure 3: Sector-by-venue trading activity. Trading volume shares by investor sector across venue types.

The top 15 bonds unite 80%. We will show that these are recently-issued on-the-run Bunds with benchmark maturities (2, 5, 10, and 30-year) and those eligible for delivery into Bund futures. The remaining bonds are barely traded.

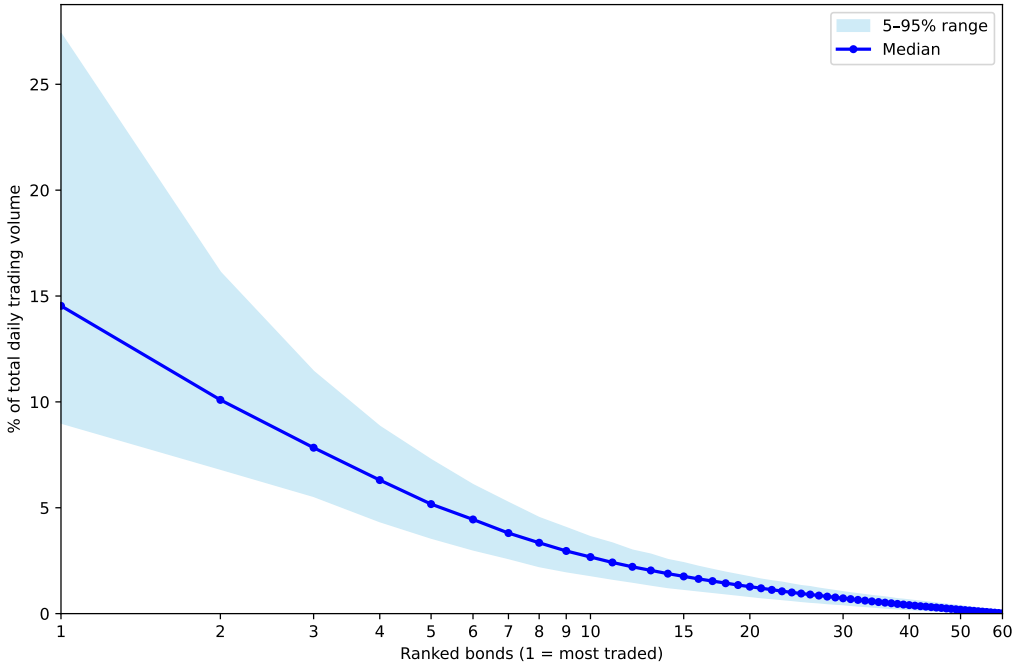


Figure 4: Concentration across bonds. The horizontal axis refers to the daily rank of a bond, from most to least traded. The line refers to the median and the shaded area to the 5% and 95% percentiles across days.

Figure 5 shows how fragmented the Bund market is across trading venues. The figure break down the aggregate trading volume by venue type and by individual venues within each type. Bilateral OTC trading — which by definition is decentral — unites the largest share of volume, followed by MTFs (electronic platforms), OTFs (inter-dealer broker platforms), SIs (single-dealer platforms) and RMs (regular exchanges). Within the MTF and OTF segment, trading is rather concentrated in a handful of major venues. The SI and RM segments, in contrast, are both fragmented across many individual venues.

Figure 6 shows how venue choice depends on trade size. Bilateral OTC trading is widespread across ticket sizes; small trades (up to €100k) are also executed on-exchange; medium trades (up to €50mn) are frequently executed on electronic platforms; and the largest trades concentrate on inter-dealer 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, 2025).

The fragmentation becomes most visible at the bond–venue level, see Figure 7. For each trading day and venue type, we first identify the single venue within that type with the largest trading volume. We then restrict transactions to this daily 'top venue' and count, for each bond ISIN, the number of trades executed on that venue that day. Within each (day, venue type), bonds are ranked by this trade count, and the heatmap plots these ranked bonds over time:

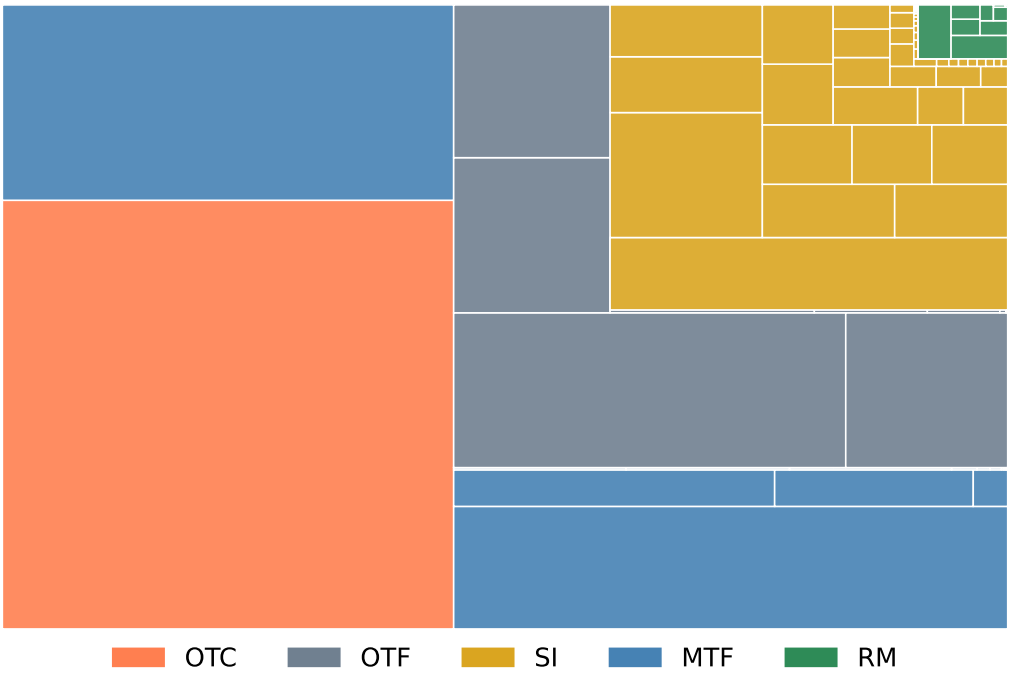


Figure 5: Fragmentation across venues. Trading volumes across venue types and across venues within types.

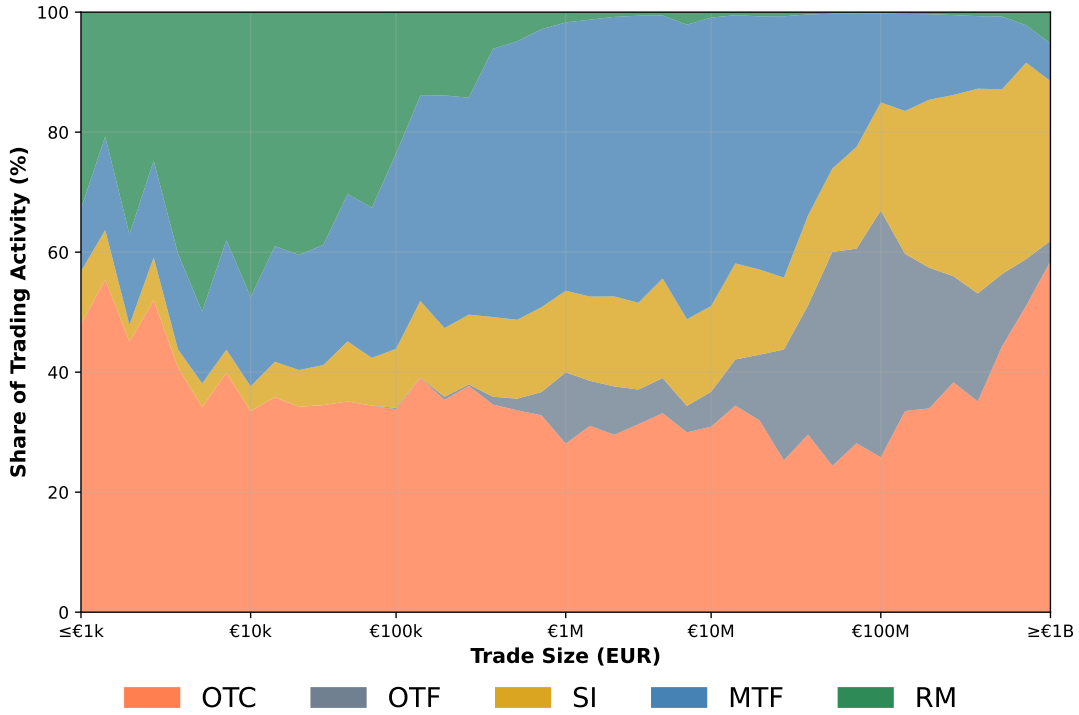


Figure 6: Fragmentation across venues by trade size. Venue composition as a function of ticket size.

each cell reports the number of trades for the bond at a given rank on a given day (with zeros indicating no trades). We see that almost all outstanding Bunds trade at least five times per day via bilateral OTC and on MTFs, while widespread inactivity prevails across all other venue types.

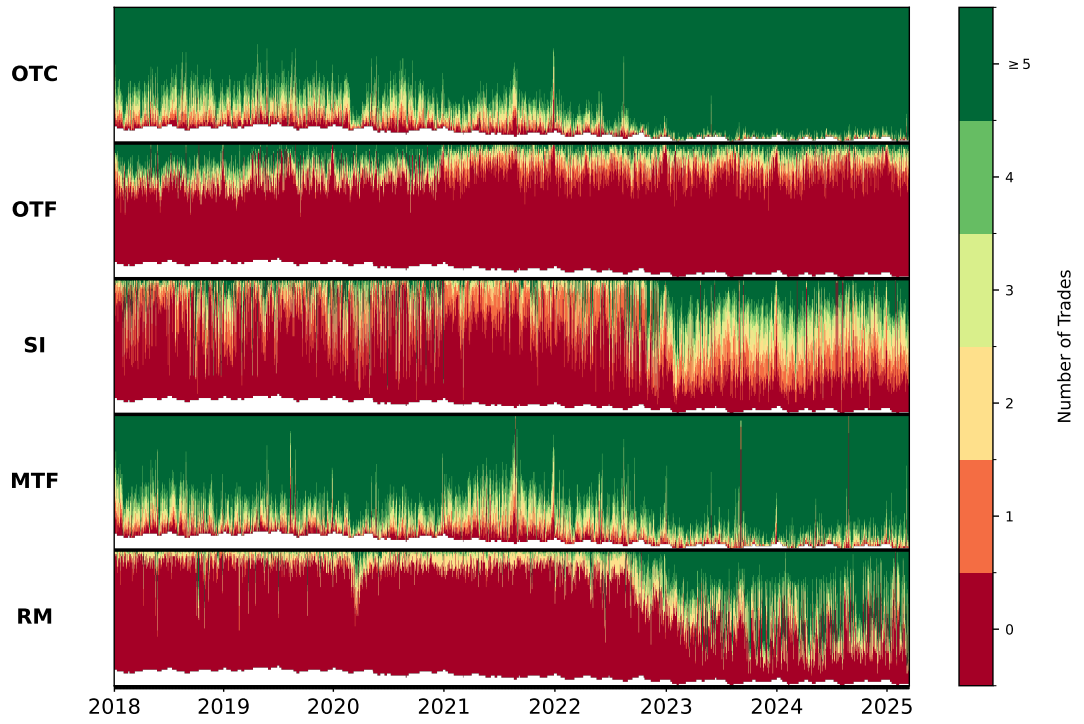


Figure 7: Bond-level trading activity by venue type. For each day and venue type, the figure shows the number of trades on the most-active venue (by volume). Each x-axis spans all outstanding Bunds at that point in time.

4 Cross-sectional Bund pricing

This section studies the relative pricing of Bunds along the yield curve and links it to repo markets, futures deliverability, and trading activity.

Specialness premium and turnover discount Figure 8 shows two striking cross-sectional patterns. First, bonds with lower repo rates are richer in the cash market: specialness reflects a convenience yield that is capitalized into bond prices. Second, and more surprisingly, the most actively traded bonds are also the cheapest. Bonds with higher turnover trade at higher yields relative to a fitted curve. This runs counter to the usual intuition that more trading activity reflects greater liquidity and should therefore be associated with higher prices (Amihud, 2002). Instead, we show that these patterns are driven by supply pressure in on-the-run benchmark bonds.

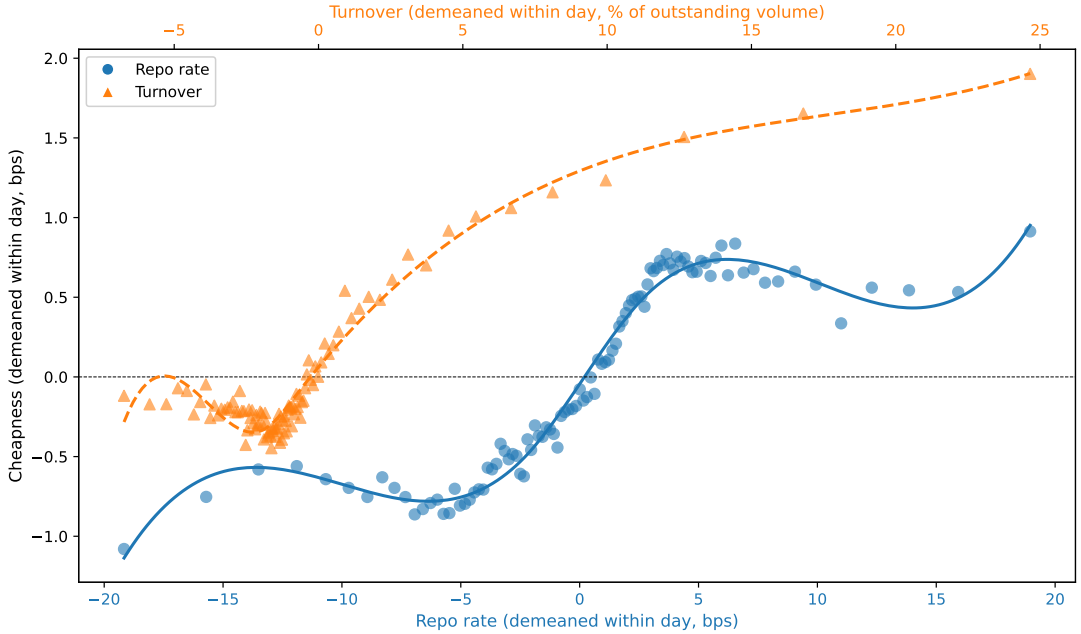


Figure 8: Cross-sectional relationship between Bund yields, repo rates and trading turnover. This figure provides a binned scatter of bond-level yield fitting errors (y-axis) and repo rates (blue circles, bottom x-axis) and turnover (orange triangles, top x-axis). All variables are demeaned within day, stripping out aggregate time-series movements to isolate purely cross-sectional variation across bonds on the same day. The sample is winsorized at the 1% and 99% percentiles to remove extreme outliers. Repo rates and turnover are binned into 100 quantiles. Fitted lines are cubic splines.

On-the-run discount Table 1 confirms that the patterns in Figure 8 are largely driven by on-the-run Bunds.[10] The most recently issued benchmark bonds are both cheap in the cash market and less special in repo markets, while also trading substantially more than other bonds. This OTR discount in Bunds contrasts with the well-known OTR premium in U.S. Treasuries (Duffie, 1996; Krishnamurthy, 2002; D'Amico and Pancost, 2022).

The key driver is the Finanzagentur's tapping policy. Newly issued benchmark bonds are repeatedly reissued, expanding their outstanding supply.[11] This additional supply reduces repo specialness and depresses cash prices. These supply effects dominate the liquidity premium that typically benefits on-the-run securities.[12]

[10]The regression includes daily fixed effects to absorb shifts in overall market conditions, like ECB policy surprises or end-of-quarter window dressing, which are a major driver of repo rates. We omit bond fixed effects because our goal is to explain cross-sectional variation in repo rates between bonds, not just time-series variation within a given bond. Many bonds do not switch deliverability, CTD, or OTR status over time, so including bond fixed effects would effectively exclude these bonds from identifying the coefficients of interest, introducing selection bias. To be precise, in our sample of 140 bonds, the deliverability status is constant for 22% of bonds, CTD status for 41%, and OTR status for 45%. We also avoid using maturity-bucket fixed effects because any reasonably narrow buckets would be highly collinear with Eurex eligibility criteria and thus with deliverability status.

[11]At the 2-, 5-, and 10-year maturity, over 98% of taps occur in OTR bonds, whereas for 30-year bonds, about half of all taps are in off-the-run securities. Appendix A provides further results across time and maturities.

[12]Vayanos and Weill (2008) e.g. argue that on-the-run securities are easier to find than off-the-run securities, increasing demand to source them in repo markets in order to short and, consequently, increasing their specialness. Huh and Infante (2021) argue that higher specialness leads to a more illiquid market. Because specialness is a cost dealers must bear to intermediate client order flow, a higher degree of specialness will translate into a higher transaction cost for the client and, hence, lower liquidity and lower prices. Fleckenstein and Longstaff (2024)

	Pricing		Turnover	
	(1) cash cheapness	(2) repo rate	(3) cash market	(4) repo market
on-the-run	1.39***	2.12**	1.35***	0.76***
	[7.02]	[2.27]	[20.52]	[22.18]
deliverable	0.12	-0.74	0.22***	0.07
	[0.53]	[-1.60]	[6.32]	[1.49]
cheapest-to-deliver	0.42	-1.75**	0.63***	0.47***
	[1.61]	[-2.29]	[11.66]	[7.88]
Day FE	Yes	Yes	Yes	Yes
Orig. Maturity FE	Yes	Yes	Yes	Yes
Observations	118,616	118,616	118,616	118,616
R^2	0.144	0.997	0.480	0.593

Table 1: Pricing and turnover on the cash and repo market. This table reports regressions of repo rates, cash market cheapness, and trading turnover on both markets, on status indicators. The sample includes all German government bonds matched across cash and repo markets at the day-bond level. T-statistics based on standard errors two-way clustered by day and ISIN are reported in brackets.

Figure 9 shows that this discount persists throughout the on-the-run period. Cash cheapness and trading activity remain elevated as the bond ages. Repo rates, however, normalize quickly: newly issued bonds are initially easy to borrow but their repo rates become indistinguishable from other bonds within about two weeks. This divergence is consistent with the supply mechanism: while repo markets reflect current collateral scarcity, cash prices incorporate expectations about future supply created by repeated taps.

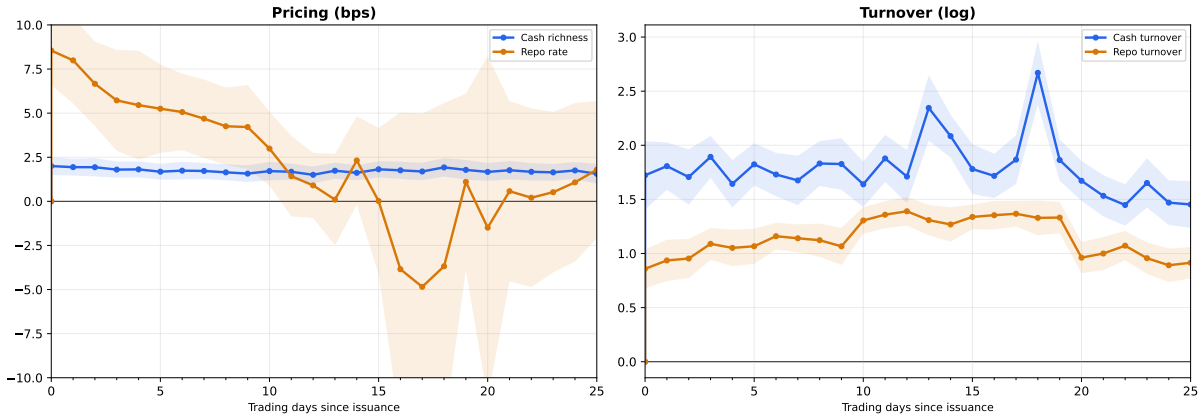


Figure 9: On-the-run effect by days since issuance. This figure plots the effect of on-the-run status as a function of bond age, estimated from a single pooled OLS regression on the full bond-day panel, interacting daily age bins with an on-the-run dummy. Each regression absorbs calendar-day and original-maturity-bucket fixed effects and controls for deliverability status. Coefficients are measured relative to all off-the-run bond-day observations and plotted with 95% confidence intervals based on two-way clustered standard errors (ISIN × day).

To isolate the role of issuance more directly, Figure 10 studies tap auctions. Repo rates drop just before a tap, reflecting temporary collateral scarcity, and jump upward once new supply

show that the OTR premium has disappeared in the US after 2008.

enters the market. Cash prices respond little: bonds remain cheap both before and after the tap. The spike in cash turnover on the tap date, lastly, reflects primary dealers distributing the newly issued bonds.

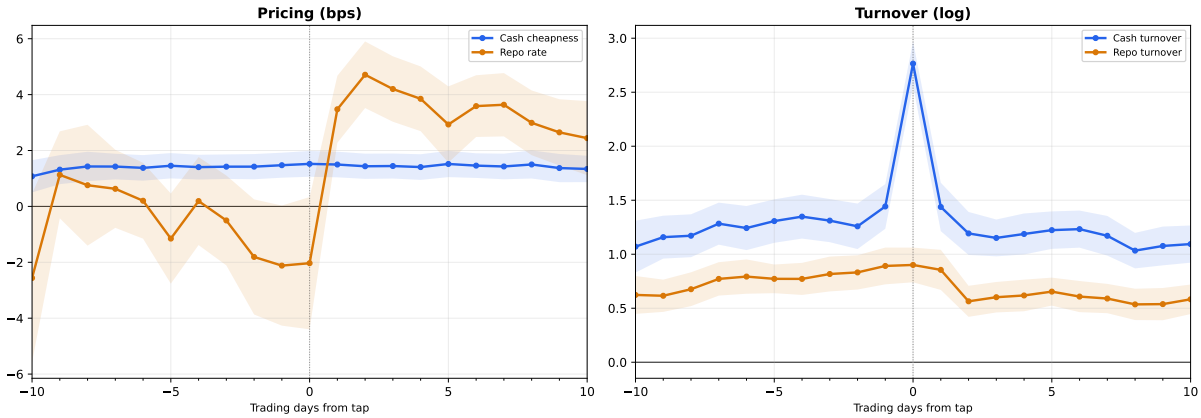


Figure 10: On-the-run effect around re-issuances. This figure plots cash richness, the repo rate, and market turnover as a function of trading days relative to the nearest tap date, estimated from a single pooled OLS regression on the full bond-day panel. Each regression absorbs calendar-day and original-maturity-bucket fixed effects. Coefficients are measured relative to all bond-day observations outside the ±10 trading-day window around any tap. The left panel shows pricing effects in basis points and the right panel shows turnover effects in logs. Shaded bands are 95% confidence intervals based on two-way clustered standard errors (ISIN × day).

To sum up, the asymmetrical pricing of on-the-run Bunds suggests that cash prices incorporate *expected* collateral scarcity, while repo rates respond primarily to *realised* scarcity.

Deliverability premium Besides on-the-run status, a key characteristic of Bunds is their deliverability into Bund futures, i.e. whether a specific bond can be delivered and whether it is the cheapest-to-deliver (CTD) bond. Table 1 shows that CTD bonds have significantly lower repo rates (in line with Baltzer, Schlepper, and Speck, 2025) and higher trading volumes on the cash and repo markets, with no statistically significant effect on cash market pricing.

Figure 11 shows that these effects are a function of the time to delivery. CTD bonds experience a sharp increase in trading activity in both the cash and repo markets as delivery approaches. At the same time, their repo rates fall dramatically, by roughly 25 basis points shortly before delivery. This pattern reflects a 'boxing' behavior, where investors with short futures positions secure CTD bonds on the repo market to ensure delivery, given the high penalties Eurex imposes for delivery fails (Eurex Clearing AG, 2012). This inelastic demand for CTD bonds near delivery amplifies their specialness in repo markets (Huh and Infante, 2021). Other deliverable bonds do not become more expensive to borrow on the repo market, in line with the fact that CTD bonds are delivered more than 99% of the time. Since this repo specialness is short-lived, it does not spill over to the cash market. Moreover, futures arbitrage disciplines bond prices: if a CTD bond became sufficiently rich, another deliverable bond would replace it as cheapest-to-deliver (Hemler, 1990). As a result, the CTD delivery squeeze is largely confined to repo markets.

CTD bonds nevertheless trade much more intensively as delivery approaches, without a

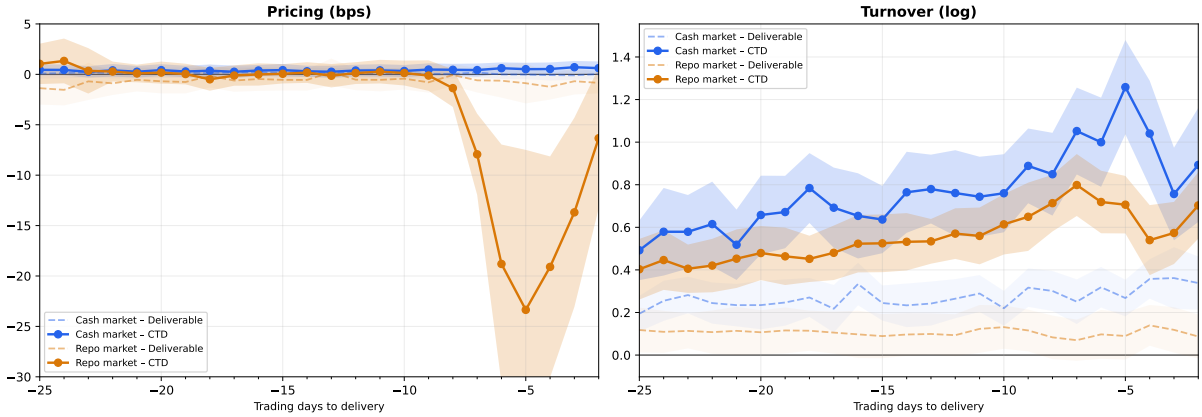


Figure 11: Deliverability and CTD effects by trading days to futures delivery. This figure plots the effect of deliverability and CTD status as a function of trading days to the Bund futures delivery date, estimated from a single pooled OLS regression on the full bond-day panel, interacting each status indicator with trading-day-to-delivery dummies. Each regression absorbs calendar-day and original-maturity-bucket fixed effects, and coefficients are measured relative to all bond-day observations outside the delivery window. Plotted coefficients cover the window from 25 to two days before delivery (last futures trading day), with 95% confidence intervals based on two-way clustered standard errors (ISIN × day).

corresponding price premium. This reinforces the turnover–cheapness pattern documented in Figure 8: heavily traded Bunds need not be rich in the cash market.

Taken together, these results highlight two distinct forces shaping Bund pricing. Tap issuance depresses the cash prices of benchmark bonds, whereas futures delivery primarily affects repo markets and trading activity. We next examine how these mechanisms relate to pricing efficiency in the cash market.

Pricing efficiency Table 2 shows that deliverable and CTD bonds exhibit substantially lower transaction-level price dispersion. Across three dispersion measures—standard deviation, volume-weighted RMSE, and median absolute deviation—deliverable and CTD bonds display significantly tighter pricing, consistent with the quasi-experimental evidence in Kerssenfischer and Helmus (2024): bond futures serve as benchmark prices for the cash market. On-the-run bonds, in contrast, exhibit larger absolute deviations from fitted yield curves despite their tight trade-level dispersion. This pattern reflects the excess supply created by taps, which pushes yields higher relative to the rest of the curve.

5 Reform Proposals

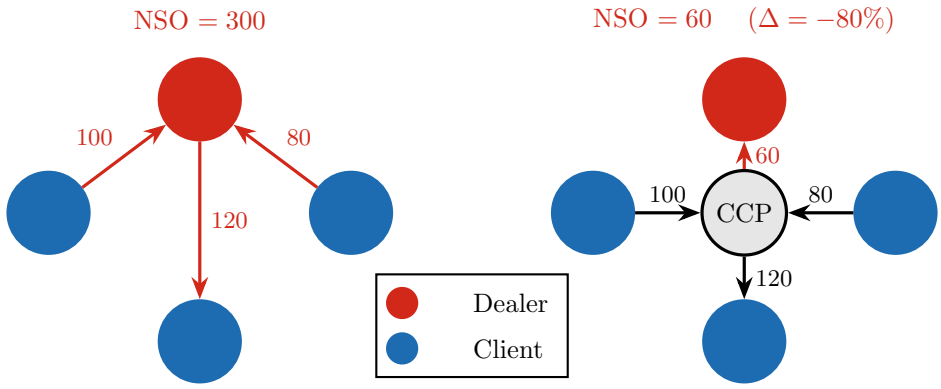
5.1 Mandatory Central Clearing

Central clearing has been widely proposed as a reform to improve the resilience and efficiency of sovereign bond markets. For US Treasuries, central clearing of cash and repo trades will become mandatory by 2027 (SEC press release). A large literature highlights the potential benefits, including a reduction in "daisy-chain" settlement failures, lower counterparty risk, and a market structure more conducive to all-to-all trading (see e.g. Duffie, 2025; Chaboud,

	Trade-level dispersion			Abs. fitting error
	Std. dev.	RMSE	MAD	|PE|
on-the-run	-0.153*	-0.087	-0.211***	0.775***
	[-1.68]	[-1.25]	[-3.09]	[5.64]
deliverable	-0.303***	-0.153***	-0.253***	-0.184*
	[-5.39]	[-3.81]	[-5.88]	[-1.80]
cheapest-to-deliver	-0.276***	-0.151***	-0.269***	0.074
	[-3.63]	[-2.71]	[-4.96]	[0.54]
Day FE	Yes	Yes	Yes	Yes
Orig. Mat. FE	Yes	Yes	Yes	Yes
Observations	118,864	118,976	118,976	118,973
R^2	0.227	0.184	0.224	0.179
R^2 (within)	0.023	0.009	0.033	0.013

Table 2: Yield disagreement and absolute fitting errors by bond status. This table reports regressions of yield disagreement measures and absolute yield curve fitting errors on bond status indicators at the bond-day level. The yield disagreement measures— standard deviation, root mean squared error (RMSE) around the volume-weighted mean, and mean absolute deviation (MAD)—capture the trade-level yield dispersion in a given bond on a given day. The absolute fitting errors refer to our cheapness measure, see Section 2.5. All measures are in basis points. The sample includes bond-days with at least two cash market transactions. T-statistics based on standard errors two-way clustered by day and ISIN are reported in brackets.

Cox, Correia Golay, Fleming, Huh, Keane, Lee, Schwarz, Vega, and Windover, 2025). A direct implication of central clearing is a reduction in dealers' liquidity and balance-sheet needs, as bilateral exposures are replaced by multilateral netting through a central counterparty (CCP). We quantify these potential benefits for Bunds, replicating the analysis by Fleming and Keane (2021) for U.S. Treasuries and Baranova et al. (2023) for U.K. Gilts.



(a) Bilateral settlement (b) Central clearing

Figure 12: Illustrative effect of central clearing on net settlement obligations (NSO). Panel (a) illustrates bilateral settlement with gross obligations across counterparties. Panel (b) shows multilateral netting through a central counterparty (CCP).

The mechanism is in Figure 12. Bilateral settlement forces dealers to fund gross obligations across counterparties. Central clearing novates trades to a CCP and nets them multilaterally, leaving each dealer with a single net obligation. In the example, NSO fall by 80%, close to what

we estimate in the data, see Figure 13. The benefits for Bunds exceed those for US Treasuries, reflecting the near absence of CCP clearing in Bunds. The benefits also exceed those for U.K. Gilts, reflecting the lack of bilateral netting agreements among Bund dealers.

We estimate these effect using our transaction-level Bund data. The sample includes all trades with at least one primary dealer and excludes intra-group trades. We map each trade into dealer settlement positions at $T+2$ (dealer-to-client trades generate one dealer position; dealer-to-dealer trades generate two). Gross settlement obligations are $\text{GSO}_t = \sum |\text{Notional}|$ across dealer positions settling on date t. Net settlement obligations depend on clearing. In the status quo, only CCP-cleared trades net; all other trades settle gross. Under mandatory clearing, all trades net at the CCP by ISIN and settlement date. The difference in NSO across regimes is the implied balance-sheet saving, computed separately for a pre-Covid and Covid-peak sample to match Fleming and Keane (2021) and Baranova et al. (2023). Importantly, our pre-Covid estimate is representative of the entire sample, as we observe rather little time-series variation in potential NSO savings in Bunds.

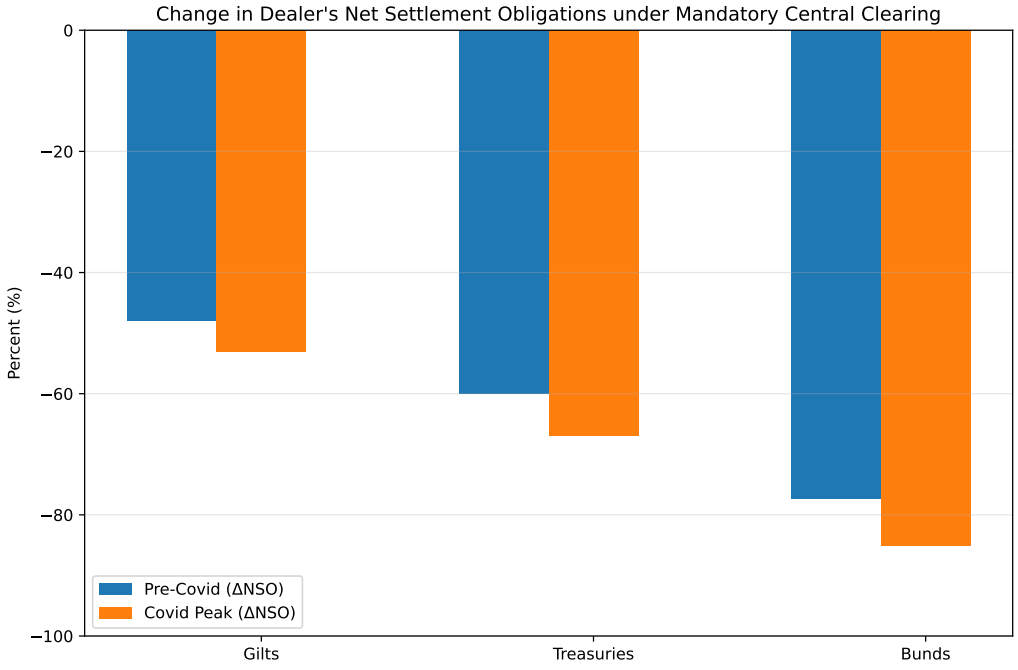


Figure 13: Mandatory central clearing: reduction in dealers' net settlement obligations (NSO). Estimates for US Treasuries are from Table 1 in Fleming and Keane (2021): pre-Covid refers to January 2 till February 25, 2020 and April 1 till 30, 2020; Covid peak to 28 February 2020. Estimates for UK Gilts are from Table 5 in Baranova et al. (2023): pre-Covid refers to January 2 till February 25, 2020; Covid peak to 10 March 2020. Estimates for German Bunds are based on our transaction data, the pre-Covid sample is adopted from US Treasuries; Covid peak refers to 3 March 2020.

The benefits of central clearing come from two margins of netting. First, *within-dealer* netting: each dealer nets its own buys and sells across ISINs. Second, *cross-dealer* netting: offsetting positions across different dealers cancel at the CCP.[13] Figure 14 shows the contribution of each margin of netting for the pre-Covid baseline. Almost 50 percentage points come

[13] We decompose netting benefits by computing net settlement obligations (NSO) under mandatory clearing in two ways. The market-wide CCP measure nets signed positions across dealers, $\text{NSO}_t^{\text{CCP}} =$

from within-dealer netting, the remaining 30 percentage points come from cross-dealer netting. Dealers capture the first component privately, the second is a network effect: it lowers aggregate settlement obligations but does not accrue to any single dealer in isolation. From a policy perspective, this wedge between private and social benefits implies that voluntary CCP adoption is likely to be inefficiently low, providing a rationale for a central clearing mandate.

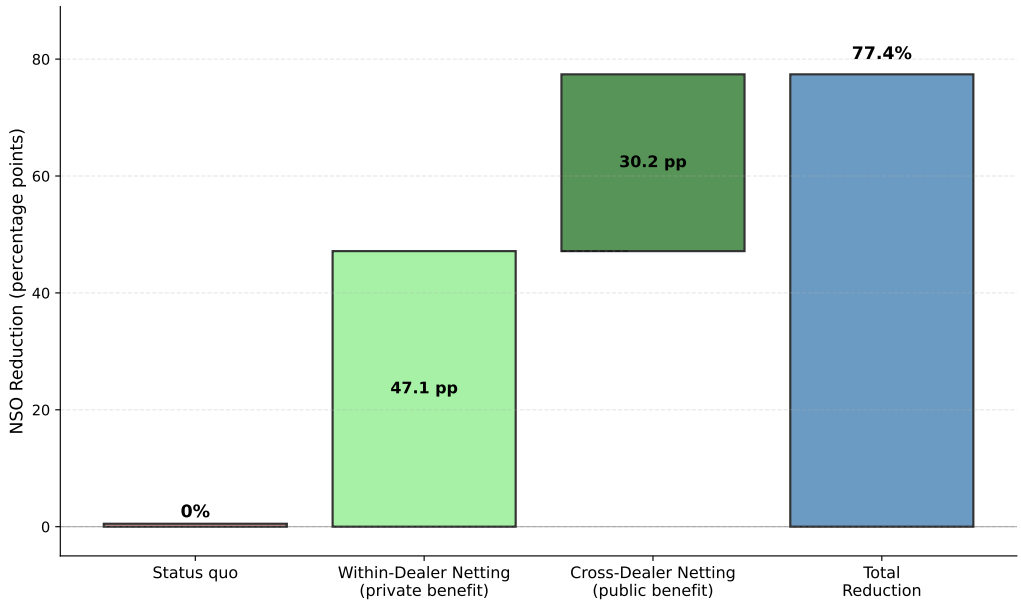


Figure 14: NSO reduction under mandatory clearing, decomposed into within-dealer and cross-dealer netting. The figure is based on the pre-Covid baseline estimate of NSO reductions, see Figure 13.

Figure 15 shows that the netting gains are very unevenly distributed across dealers, ranging from roughly 30% to 90%. Such dispersion may further impede a voluntary migration to central clearing.

A few words of caution are in order. While we document large potential benefits of central clearing for Bunds, it is worth noting some underlying assumptions and institutional constraints. Our system-level (cross-dealer) netting gains are computed as if all cleared positions meet in a single netting set—one CCP, or full cross-CCP interoperability—so fragmentation across CCPs would mechanically attenuate the aggregate benefit. More broadly, a U.S.-style clearing mandate is harder to implement in the euro area: the sovereign bond market is institutionally and infrastructurally fragmented across issuers, venues, clearing arrangements, and supervisory jurisdictions. A recent European Systemic Risk Board (2024) report therefore recommends an incentives-based approach—harmonising risk management between bilateral and CCP clearing (including margins), removing prudential disincentives for NBFIs, broadening CCP access models, and strengthening CCP service offerings—rather than an immediate, uniform mandate. In that sense, our estimates quantify the mechanical gains available in the data; realizing them requires institutional convergence.

$\sum_i \left|\sum_d \text{SignedNotional}_{d,i,t}\right|$, while the dealer-by-dealer measure nets within each dealer and then sums, $\text{NSO}_t^{\text{Dealers}} = \sum_d \sum_i \left|\text{SignedNotional}_{d,i,t}\right|$. The difference $\text{NSO}_t^{\text{Dealers}} - \text{NSO}_t^{\text{CCP}}$ captures cross-dealer netting, a system-level benefit that does not appear in dealer-level calculations.

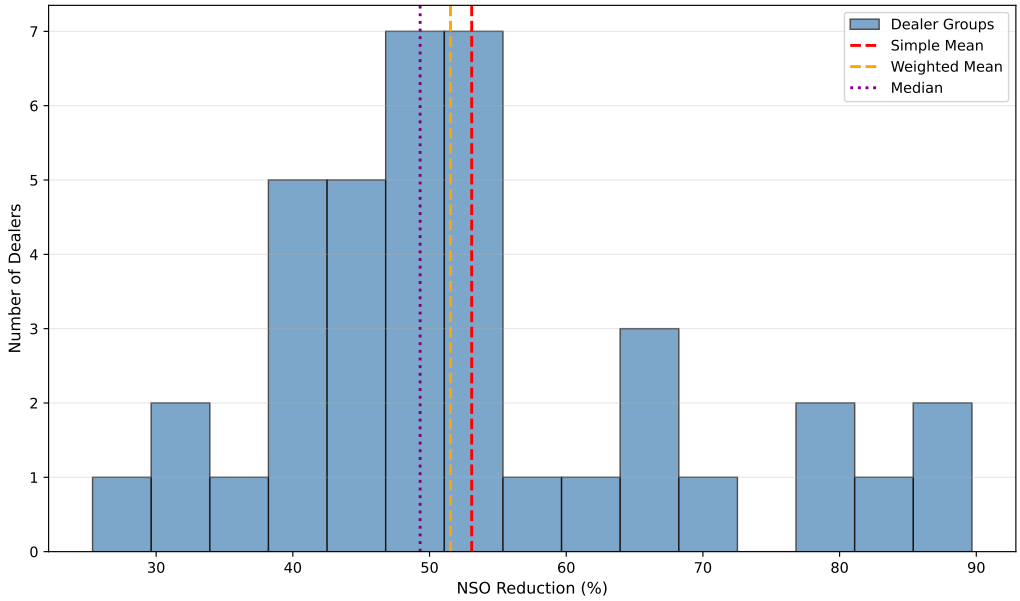


Figure 15: Dealer-level NSO reductions under mandatory clearing. This figure breaks down the private netting benefits across dealers. Estimates only reflect within-dealer netting, cross-dealer netting only appears in the aggregate CCP calculation, see Figure 14.

5.2 Venue Centralization

Central clearing may alter not only settlement mechanics but also the organization of trading itself. By novating trades to a central counterparty, clearing can facilitate more open trading architectures and reduce the need for tightly intermediated dealer networks. In particular, CCP clearing is often viewed as a prerequisite for more *all-to-all* trading structures in which investors can transact more directly with a broader set of counterparties. If clearing reforms move sovereign bond markets toward such structures, an important question is how greater market centralization affects welfare.

Allen and Wittwer (2023) study this question using transaction-level data from the Canadian government bond market to estimate a structural model of dealer–investor trading. We replicate their analysis using our transaction-level dataset to assess whether their findings generalize to the Bund market.

Sample and setup We use dealer-to-client transactions from the last three years of our sample (2023-2025) to reflect current market structures. We classify venues into two categories: bilateral OTC trading, which includes both traditional OTC and SI activity; and platform trading, defined as execution on MTFs such as Bloomberg and Tradeweb, where multiple dealers compete for an investor's order. To mimic the model structure, investors are classified as institutional if they trade at least once on a platform and as retail otherwise. Institutional investors account for 92% of all trades.

Stylized facts Several patterns in Bund trading mirror the Canadian evidence that motivates the structural model.

First, investor-dealer relationships are concentrated. Among investors who trade bilaterally, 79% do so with a single dealer—their "home dealer"—who accounts for 87% of the investor's OTC volume on average. Importantly, home dealers do not offer better yields than non-home dealers on average. This gives rise to the models loyalty benefit mechanism: investors value relationships with their home dealer beyond price discounts, because they offer benefits to their client such as managing the clients financial accounts, providing tax expertise or serving loyal clients in times of stress.

Second, platform trading broadens access to dealers and introduces competition among them. On multilateral platforms, the concentration of trading with a single dealer falls to 57%, and 25% of platform investors access four or more dealers. This broader access is associated with better yields: institutional investors who trade on the platform receive markups that are approximately 0.1 basis points lower than in bilateral OTC trading. This pattern is consistent with platforms reducing search frictions and increasing competition among dealers for investor order flow.

Third, despite the yield advantage of platforms for institutional investors, both OTC and platform trading coexist. This suggests that other factors beyond prices affect trading decisions such as platform costs and loyalty benefits. Platform adoption among institutional investors is notably high in the Bund market: approximately 70% of institutional trades execute on electronic platforms, compared to 30–40% in Canada.

Model outline and estimation The structural model features J dealers and N investors. Each dealer j has a (potentially investor-specific) quality μ_{ij} and quotes a price on the platform that embeds a markup over the dealer's true value v. Each institutional investor i has a "home dealer" and chooses between trading bilaterally or trading on the platform with the dealer offering the best quote. A loyalty benefit r accrues when an investor trades with her home dealer, capturing the non-price value of the bilateral relationship. The key parameters are the demand elasticity σ, which governs how sensitively investors respond to quote differences across dealers, dealer qualities, the loyalty benefit r, a distrbituion of liquidity shocks to investor valuations and a platform cost c.

Estimation proceeds in three stages. First, we assign a home-dealer to each investor. Second, we regress platform quotes on dealer level platform shares to estimate the degree of platform competition, the loyalty benefit with a home dealer dummy and we recover qualities with dealer fixed effects. Third, we estimate the remaining structural parameters via GMM.[14] A key challenge is that we lack primary auction data, which Allen and Wittwer (2023) use to construct instruments for the demand elasticity parameter. We therefore impose the value $1/\hat{\sigma} = 1.765$ from the Canadian estimation and estimate the remaining parameters conditional on it; robustness to this choice is assessed below.

[14]The main estimates are: demand elasticity $\hat{\sigma} = 0.567$, loyalty benefit $\hat{r} = 0.83$ basis points, and platform cost $\hat{c} \approx -0.8$ basis points (a negative cost, implying a net subsidy to platform trading). 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 we observe in Bunds. Dealer quality varies substantially across buy and sell sides and between institutional and retail investors.

Welfare analysis Dealers and investors both have heterogeneous values for bonds and model-implied welfare is measured with gains from trade from counterfactual scenarios. Changes to welfare come entirely from rematching who trades with whom and increases if investors buy from more efficient dealers (better yields, better quality, or loyalty benefit). In particular, welfare gains arise when more investors trade on the platform at lower prices through better matching. However, due to imperfect competition, it is not clear whether welfare increases in accordance to higher platform shares. This is because dealers charge a markup over their true values and the platform itself is imperfect, which leads to inefficient matches.

Table 3 reports the percentage welfare gain from varying degrees of market centralization relative to a fully bilateral (decentralized) baseline. Importantly, dealer values are held constant across all scenarios and the remaining parameters adjust in equilibrium. The columns represent increasing degrees of centralization: *Account I* (institutional investors have platform access but can choose between OTC and platform trades, i.e., the status quo), *Account IR* (all investors including retail have access), *No Fee* (platform fees eliminated), and *Central* (all trades forced onto the platform). The rows vary two model ingredients: the assumption about the loyalty benefit r and the intensity of dealer competition. Within each loyalty assumption, *imperfect competition* allows dealers to set markups optimally, while *perfect competition* forces dealers to quote at true value.

		Account I	Account IR	No Fee	Central
With loyalty	Imperfect	−4.88	−5.72	−6.40	−7.25
	Perfect	+6.82	+7.71	+8.76	+10.15
Switching cost	Imperfect	+9.72	+10.48	+12.79	+16.45
	Perfect	+22.66	+25.24	+28.41	+32.54
No loyalty	Imperfect	+13.23	+14.28	+16.34	+19.12
	Perfect	+34.03	+37.89	+41.17	+45.00

Table 3: Average welfare gain compared with a fully decentralized market. This table replicates Table 6 from Allen and Wittwer (2023) using our Bunds transaction data. Values are percentage welfare gains relative to a fully bilateral market. Imperfect competition: dealers set quotes optimally. Perfect competition: dealers quote at true values ($q = v$, $\sigma = 0$). Account I: only institutional investors have platform access (status quo). Account IR: all investors. No Fee: platform fees set to zero. Central: all trades forced onto the platform. With loyalty: home-dealer benefit r enters welfare. Switching cost: r distorts behavior but does not enter welfare. No loyalty: $r = 0$.

The central finding is that centralization *decreases* welfare under imperfect competition when the loyalty benefit is valued (first row), but *increases* welfare under perfect competition (second row). This reflects two opposing effects operating: centralization enables more efficient investor-dealer matching, but at the same time investors trade less with their home dealer and therefore lose the loyalty benefit. More efficient investor-dealer matching only dominates when platform competition is perfect.

The remaining rows show that dealer-investor relationships play a large role for welfare analysis. If r is reinterpreted as a pure switching cost that distorts behavior but does not enter welfare directly, centralization yields positive welfare gains even under imperfect competition

(rows 3–4). Setting $r = 0$ entirely amplifies the gains further (rows 5–6).[15]

Figure 16 compares the welfare gains between Bunds and Canadian government bonds quantitatively. While the overall pattern of our estimates closely matches Canada, welfare changes appear substantially larger in the Bunds market. Three main factors can explain this. First, greater dispersion of dealer values allows for better matches to form when a high-valued dealer attains a larger market share on a competitive platform. Second, the lower loyalty benefit estimate for Bunds allows more switches to occur even under full centralization, where a large loyalty benefit causes many investors to remain matched with their home dealer. Third, the higher platform shares for Bunds lead to larger gains when platform access is still limited.

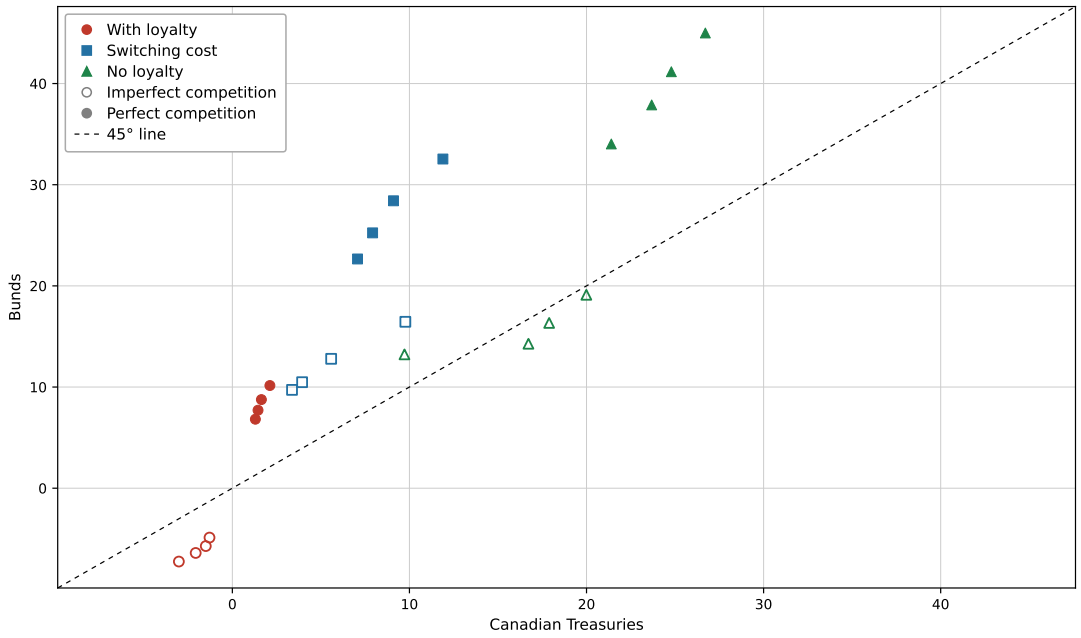


Figure 16: Welfare gains from centralization: Bunds versus Canada. Each bar represents the percentage welfare gain relative to a fully bilateral market under the indicated scenario. Canadian estimates are from Allen and Wittwer (2023), Table 6.

Since the demand elasticity $\hat{\sigma}$ is imported from the Canadian estimation, we conduct a robustness exercise in which we halve $\hat{\sigma}$, corresponding to substantially stronger price competition among dealers. Under this alternative calibration, the qualitative patterns in Table 3 are preserved and the welfare gains become more moderate, moving quantitatively closer to the Canadian estimates.

Discussion A major implication of our analysis is that fostering competition on electronic platforms matters more than centralization itself—a point that is especially relevant for Bunds, where institutional investors already execute the majority of their trades on platforms. Our results, in line with Canada, also suggest that centralization can attenuate frictions such as

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

experienced during Covid when dealers failed to absorb the extraordinary supply of US Treasuries: centralized competitive platforms allow dealers with high values to offload their inventory to new customers. However, several limitations should be borne in mind. The model abstracts from dealer-to-dealer trading, focusing exclusively on dealer-to-client transactions. The binary venue choice—bilateral OTC versus a single centralized platform—is a stylized representation of the fragmented European venue landscape. Furthermore, the model abstracts from the sizeable Bunds future market.

6 Conclusion

This paper studies the German Bund market using a new non-anonymous transaction-level dataset constructed from MiFIR reports. The dataset provides the broadest coverage of the Bund cash market to date and allows us to document key features of market structure and evaluate the effects of proposed reforms. We will make the cleaned dataset available to eligible researchers to facilitate further work on sovereign bond market microstructure and design.

The data reveal a market that is fragmented across venues, segmented across investor groups, and hierarchical around a dealer core. Liquidity concentrates in a small set of bonds—especially newly issued on-the-run bonds eligible for futures delivery—while most outstanding Bunds trade only sporadically and often not at all on a given venue-day. Venue choice is systematic across sectors, and most trading is intermediated by primary dealers.

Turning to prices, we document that repo specialness is capitalized into cash prices: bonds with lower repo rates are richer in cash. At the same time, more heavily traded bonds are cheaper in cash markets. We trace this turnover discount to re-issuances. Because on-the-run Bunds are frequently tapped, their supply expands, making them less special in repo and cheaper in cash. This OTR discount contrasts with the well-known OTR premium in U.S. Treasuries.

Finally, we use the granular transactions data to evaluate policy counterfactuals. In a mandatory CCP scenario, multilateral netting sharply reduces dealers' gross and net settlement obligations, with larger gains for Bunds than those documented for U.S. Treasuries and U.K. Gilts. Using a structural model of dealer–investor trading, we find that centralizing bilateral trading onto electronic platforms can decrease welfare when dealers have market power and investors value bilateral relationships. Only under perfect competition does centralization unambiguously improve welfare.

References

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

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of financial markets 5*(1), 31–56.

Baltzer, M., K. Schlepper, and C. Speck (2025). The eurosystem's asset purchase programmes, securities lending and bund specialness. *Journal of Business Economics*.

Baranova, Y., E. Holbrook, D. MacDonald, W. Rawstorne, N. Vause, et al. (2023). The potential impact of broader central clearing on dealer balance sheet capacity: a case study of uk gilt and gilt repo markets.

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.

Branger, N., M. Muck, and A. Pütz (2024). Unmasking global investors: An examination of the shifting dynamics in sovereign bond trading. *Available at SSRN*.

Chaboud, A. P., C. Cox, E. Correia Golay, 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*. Federal Reserve Bank of New York, 31(2).

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

D'Amico, S. and N. A. Pancost (2022). Special repo rates and the cross-section of bond prices: The role of the special collateral risk premium. *Review of Finance 26*(1), 117–162.

de Roure, C., E. Moench, L. Pelizzon, and M. Schneider (2025). OTC Discount. *Management Science forthcoming*.

Duffie, D. (1996). Special repo rates. *The Journal of Finance 51*(2), 493–526.

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).

Eurex Clearing AG (2012). Clearing Conditions for Eurex Clearing AG.

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.

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

Fleckenstein, M. and F. A. Longstaff (2024). Treasury richness. *The Journal of Finance 79*(4), 2797–2844.

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

Hemler, M. L. (1990). The quality delivery option in treasury bond futures contracts. *Journal of Finance 45*(5), 1565–1586. Pages/issue as in RePEc/EconPapers record.

Hu, G. X., J. Pan, and J. Wang (2013). Noise as Information for Illiquidity. *The 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.

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

Kerssenfischer, M. and C. Helmus (2024). Outages in sovereign bond markets.

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

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

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

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.

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

Vayanos, D. and P.-O. Weill (2008). A search-based theory of the on-the-run phenomenon. *The Journal of Finance 63*(3), 1361–1398.

Internet Appendix
for
Bunds

by Mark Kerssenfischer, Konrad Lucke and Willy Scherrieble

(not for publication)

Appendix A provides additional results.
Appendix B details how we adapt Filipović et al. (2025)'s yield curve fitting method to Bunds.
Appendix C outlines a best-practice methodology for cleaning and processing MiFIR transaction reports, accounting for the complexity of the Bund market. To ensure reproducibility and to foster further research in this area, we will share the final dataset with eligible researchers.
Appendix D provides descriptive statistics and cross-checks of our dataset.

Appendix A Further Evidence

OTR and Deliverability Effects This section provides further evidence on Section 4.

Table A1 confirms that the OTR discount we document is robust to alternative "fair yield" benchmarks. Figure A1 shows the time-series variation in the on-the-run and deliverability discounts. OTR bonds have been excessively cheap on the cash and repo market since around 2022. Deliverable bonds, in contrast, traded at a steep premium in the repo market from 2021 till mid-2023, with slightly *opposite* effects on the cash market. That means deliverable bonds were comparatively cheap in the cash precisely when they were expensive in the repo market.

Figure A2 shows how the repo-cash market link evolved over time. Bonds with higher repo rates have higher cash market yield, i.e. less special bonds are cheap, and this link strenghtened considerably since 2023.

Table A2 estimates the effect of on-the-run and deliverability status separately across the four relevant maturities: 2-, 5- ,10- and 30-years, which are the standard maturities the Finanzgentur issues and for which there are Bund futures on Eurex.

	Average	Yield Curve Error			Interpolation excl.		
		Svensson	Kernel-Ridge	Spline	OTR	Deliv	CTD
on-the-run	1.39***	1.69***	1.12***	0.54***	2.90***		
	(0.20)	(0.31)	(0.21)	(0.10)	(0.28)		
deliverable	0.11	-0.09	0.19	0.07		1.18***	
	(0.22)	(0.42)	(0.21)	(0.08)		(0.21)	
cheapest-to-deliv.	0.42	0.15	0.18	0.39***			1.64***
	(0.26)	(0.49)	(0.24)	(0.09)			(0.30)
Orig. Mat. FE	Yes	Yes	Yes	Yes	No	No	No
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	118973	118973	118973	118973	6065	20826	7754
R^2	0.14	0.14	0.18	0.13	-	-	-
R^2 Within	0.04	0.02	0.02	0.02	-	-	-

Table A1: Robustness of OTR discount to alternative richness measures. This table estimates the effect of on-the-run and deliverability status on cash market richness across alternative fair-yield benchmarks. The first column replicates column (1) of Table 1 using the average across all benchmarks as the dependent variable. The remaining columns use each benchmark individually. All benchmarks are described in Section 2.5: the Bundesbank Svensson curve, Kernel Ridge regression, cubic spline interpolation, and leave-one-out interpolations excluding bonds with a given status. Positive values indicate cheapness (yield above fair value). The sample covers all German government bonds at the day-bond level, with day and original maturity bucket fixed effects. T-statistics based on standard errors two-way clustered by day and ISIN are reported in brackets. The last three columns take a different approach: each column reports the unconditional mean richness of a given bond type using day fixed effects and standard errors clustered by ISIN.

	Pricing		Turnover	
	(1) cash cheapness	(2) repo rate	(3) cash market	(4) repo market
OTR × 2-year	1.25***	5.58**	1.26***	0.79***
	[4.40]	[2.36]	[28.20]	[20.06]
OTR × 5-year	1.42***	4.80***	1.43***	0.84***
	[5.31]	[5.13]	[21.26]	[11.55]
OTR × 10-year	0.90**	0.47	1.32***	0.63***
	[2.42]	[0.35]	[13.37]	[8.40]
OTR × 30-year	1.10**	-0.26	0.93***	0.57***
	[2.06]	[-0.11]	[6.80]	[8.00]
Deliv. × 2-year	-0.22	0.28	-0.08***	-0.03
	[-1.14]	[0.44]	[-3.23]	[-0.81]
Deliv. × 5-year	0.39	-1.07	0.21***	0.18***
	[1.47]	[-1.38]	[4.13]	[3.26]
Deliv. × 10-year	1.82***	0.09	0.89***	0.50***
	[4.33]	[0.08]	[7.97]	[8.63]
Deliv. × 30-year	-0.59	-1.31	0.23**	-0.11
	[-1.05]	[-1.22]	[2.60]	[-0.91]
CTD × 2-year	0.52**	-5.73**	0.80***	0.45***
	[2.16]	[-2.08]	[15.77]	[9.95]
CTD × 5-year	0.13	-1.07	0.61***	0.39***
	[0.36]	[-1.07]	[10.63]	[6.10]
CTD × 10-year	-0.28	-0.96	0.43***	0.18***
	[-0.62]	[-1.05]	[3.69]	[2.81]
CTD × 30-year	0.50	-2.87*	0.39***	0.58***
	[0.73]	[-1.83]	[4.35]	[3.55]
Day FE	Yes	Yes	Yes	Yes
Orig. Maturity FE	Yes	Yes	Yes	Yes
Observations	118,616	118,616	118,616	118,616
R^2	0.165	0.997	0.517	0.610

Table A2: Pricing and turnover on the cash and repo market. This table reports regressions of repo rates, cash market cheapness, and trading turnover on both markets, on status indicators. The sample includes all German government bonds matched across cash and repo markets at the day-bond level. T-statistics based on standard errors two-way clustered by day and ISIN are reported in brackets.

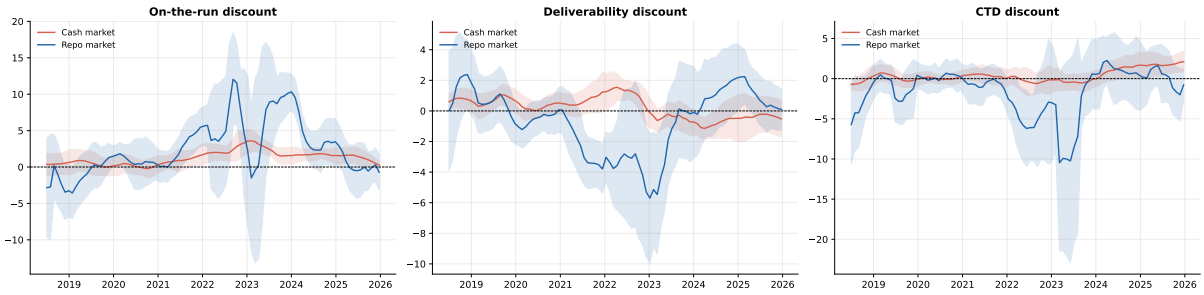


Figure A1: On-the-run and deliverable discounts over time. This figure plots rolling estimates of the cash pricing error and repo rate associated with on-the-run and deliverable status, estimated from a sequence of OLS regressions on 120-calendar-day windows stepped at 20-day intervals. Effects are in basis points. Each window regression absorbs calendar-day and original-maturity-bucket fixed effects. Coefficients are plotted at the end date of each window with 95% confidence intervals based on two-way clustered standard errors (ISIN × day).

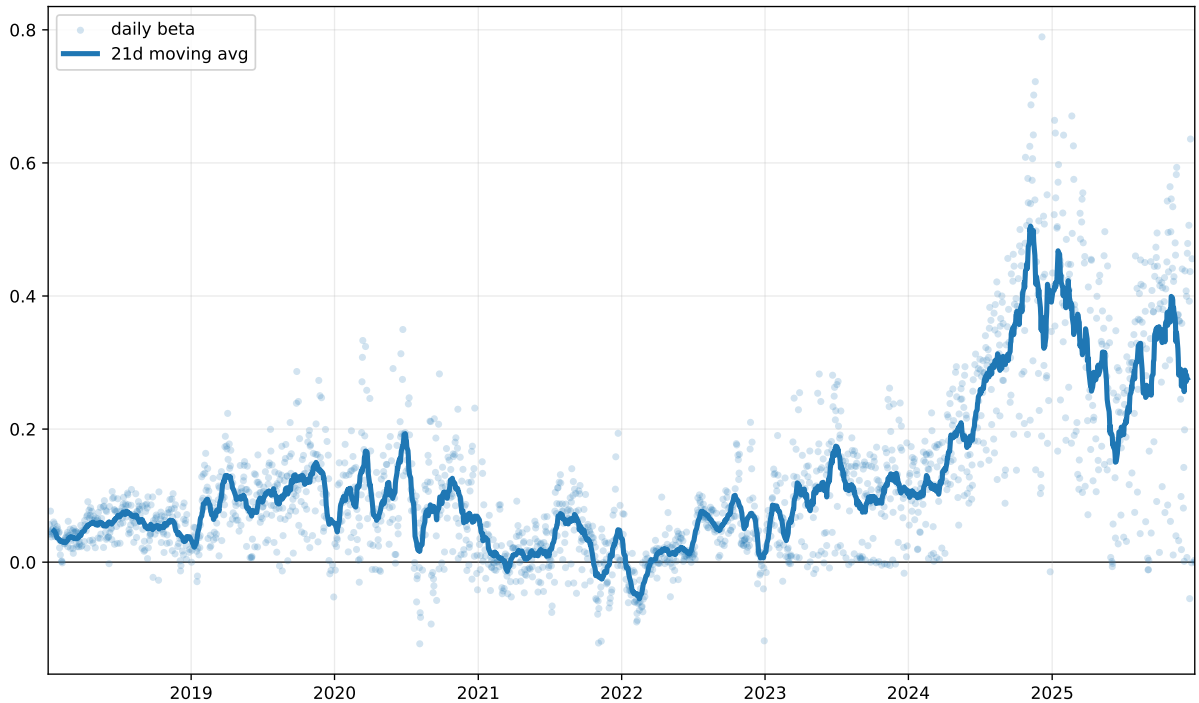


Figure A2: Time-series of cross-sectional repo-cash betas. This figure plots the regression coefficient of repo rates on cash market yield fitting errors, both in basis points.

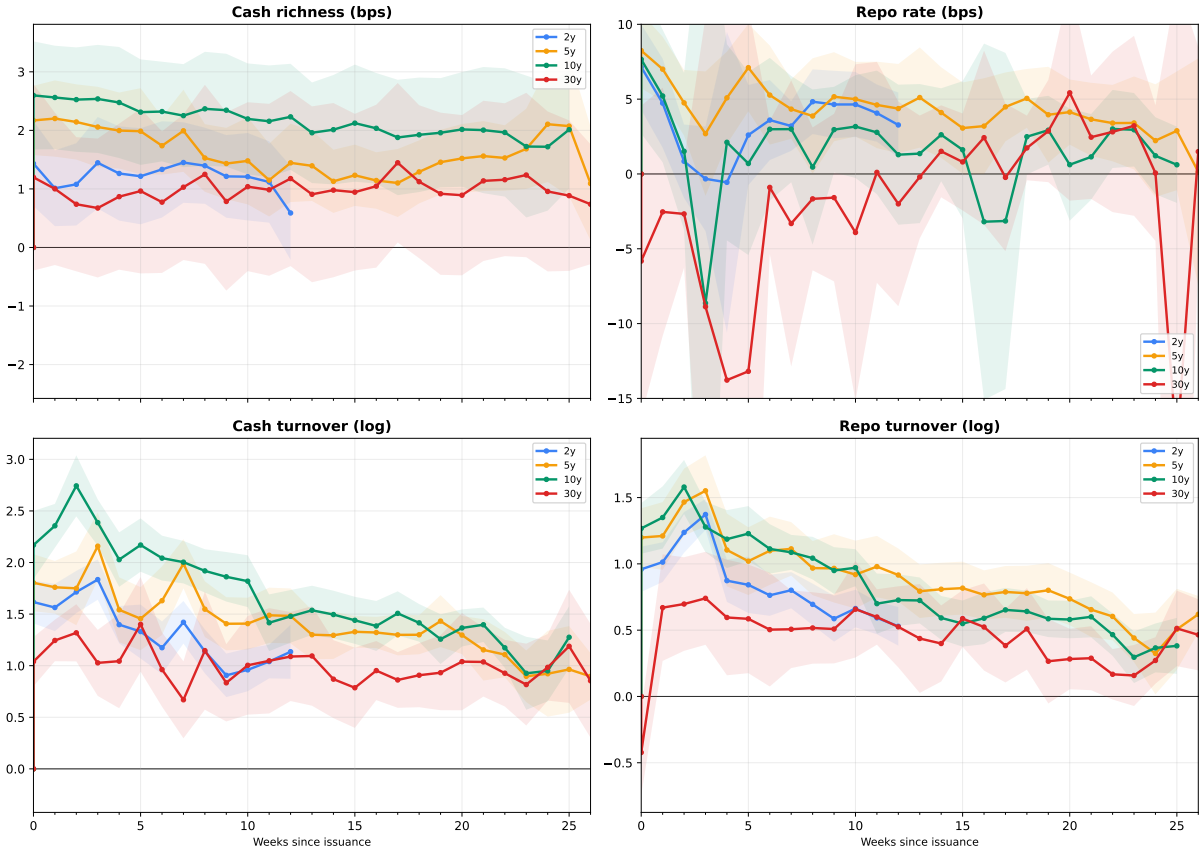


Figure A3: On-the-run effect by week since issuance across maturities. This figure replicates Figure 9, differentiating on-the-run Bunds across the four relevant maturity segments.

Trading Volumes on the Cash and Futures Market To put the hypothetical scenario of centralizing the OTC market (Section 5.2) into perspective, Figure A4 compares cash market trading activity under the current fragmented structure with a counterfactual centralized market. The left plot show trading volumes, the right plot trading frequency. For reference, both plots contain the equivalent activity measures for the four Bund futures.[16]

The figure highlights a stark scale difference: the futures market would dwarf the cash market even under full centralization. This suggests that the welfare and liquidity implications of venue centralization cannot be assessed in isolation from the cash–futures nexus. Bund futures embed institutional features that facilitate very high trading intensity—low initial margins, daily variation margining, and cross-product margining—so that a substantial share of price discovery and risk-sharing may be absorbed by the derivatives market.

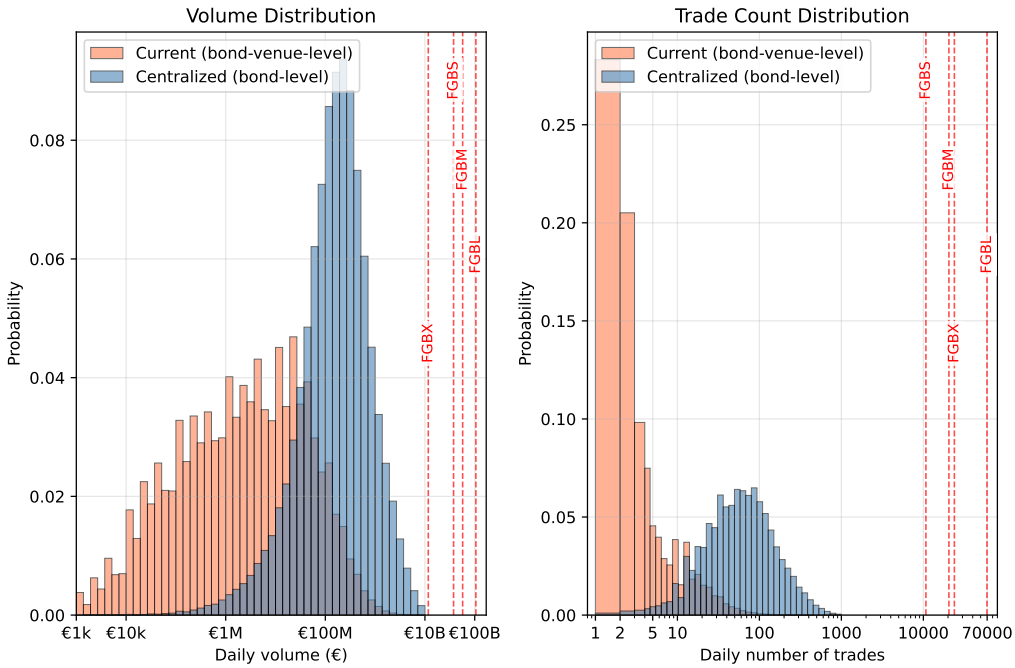


Figure A4: Trading activity in Bunds. The left histogram compares trading volumes in the current market setting (bond-venue-day level) with a counterfactual fully centralized market (day-bond level). The right histogram is in terms of number of trades instead of volume. In both panels, vertical dashed lines show the corresponding median values in the four major Bund futures.

[16] While there are three different contracts available for each Bund future, almost all trading is concentrated in the nearest-to-maturity contract, except for a few roll-over days each quarter.

Appendix B Adapting Filipović et al. (2025) to Bunds

We implement the kernel ridge yield curve estimation of Filipović et al. (2025) using their publicly available Python code. We adapt it to the German Bund market. Several market-specific features require adjustments relative to the U.S. Treasury setting considered in their study. In particular, Bunds pay annual rather than semi-annual coupons, and coupon schedules are often irregular, especially for older issues and bonds with long first or last coupon periods. We perform the estimation day-by-day using the volume-weighted average yield of each bond.

In choosing the three hyperparameters of the kernel ridge regression — namely the kernel decay parameter α, which governs the smoothness of the discount curve, the kernel shift parameter δ, which allows for short-end flexibility, and the regularization parameter λ (ridge), which controls overfitting — we follow the cross-validation approach of Filipović et al. (2025). We apply a 10-fold cross-validation scheme on each evaluation date, where bonds are stratified into maturity buckets (<0.25y, 0.25–1y, 1–2y, 2–3y, 3–4y, 4–5y, 5–7y, 7–10y, 10–20y, >20y) to ensure an approximately equal representation of maturities across folds. All cross-validation is conducted using end-of-month observations, yielding a total of 87 evaluation dates over the period January 2018 to March 2025. For each evaluation date, we randomly partition the cross section of bonds without replacement into ten non-overlapping folds, constructed such that each fold replicates the overall maturity distribution of the sample; nine folds are used for estimation and the remaining fold for evaluation, and this procedure is repeated ten times so that each bond is used exactly once for out-of-sample evaluation.

The hyperparameter grid consists of 100 (α, δ) combinations and 9 ridge values, for a total of 900 candidate specifications. Model performance is evaluated using the root mean squared pricing error (RMSE) in terms of prices (not yields), averaged across dates and folds. The optimal parameter set is found to be $(\alpha, \delta, \lambda) = (0.01, 0.0001, 5.0)$, achieving a mean RMSE of 0.217. Relative to the baseline parametrization $(\alpha, \delta, \lambda) = (0.05, 0, 1.0)$ that Filipović et al. (2025) use for U.S. Treasuries, we achieve a modest improvement of approximately 0.9% in pricing accuracy.

Appendix C Data Details

Article 26 of Regulation (EU) No 600/2014 (MiFIR) mandates that any investment firm trading a financial instrument must report this transaction to the appropriate competent authority. The competent authority is determined by where the reporting entity is regulated as well as by where the traded instrument is regulated. Thus, all transactions in sovereign bonds are reported to the German authority BaFin. The collected data thus provides a complete overview of the Bund market, thereby presenting enormous possibilities for researchers to uncover new insights.

These detailed reporting standards however come with a number of difficulties that have to be addressed for an appropriate use of the data. MiFIR reporting guidelines allow for different reporting formats and duplicate reports are not easily identifiable. This makes even relatively simple tasks like calculating overall trade volumes more difficult. Jurkatis (2024) provides a thorough approach to dealing with these problems for the corporate bond market. We will follow this approach closely, adapting it where necessary to the characteristics of the Bund market.

One major difference of our Bund market dataset with respect to datasets on the corporate bond market is the much higher trading frequency. While the overall number of ISINs is much lower, a single ISIN can be reported more than 50000 times in one day. This makes the task of identifying duplicates much more computationally difficult and requires some extra steps, as will be explained in the following.

C.1 Reporting Requirements

Under MiFIR, which took effect on 3 January 2018, investment firms must report detailed information about each bond trade to regulators by the end of the next business day. These transaction reports include comprehensive data on the trade: the instrument identifier, execution date and exact timestamp, price, volume, and the venue or method of execution. They also identify the parties involved (buyer and seller) via codes such as Legal Entity Identifiers (LEIs) or national IDs, and include other details like the capacity in which the firm acted, any algorithmic trade IDs, and flags for short sales.

In principle, each investment firm is responsible for reporting the transactions it executes. For off-venue trades (over-the-counter deals between two investment firms), both counterparties must submit a report for their side of the trade to avoid gaps in coverage. If only one side of the trade is a MiFIR-regulated firm (for example, a bank trading with a corporate client), then only the investment firm reports the trade. For on-venue trades executed on a regulated market or MTF, the trading venue often provides the transaction details. In fact, MiFIR requires trading venues to report on behalf of any participant that is not itself subject to MiFIR – for example, if a non-EU firm trades on a German exchange, the venue must ensure that transaction is reported to the regulator. Similarly, trades executed through an EU systematic internaliser (SI) or market maker may be reported by that SI (especially when the counterparty is not a reporting entity).

All MiFIR transaction reports go to the respective National Competent Authority (NCA) – in Germany, this is BaFin. German investment firms submit their reports via BaFin's approved systems. ISIN, price, quantity and timestamp have to also be made publicly available.

C.2 Sample Selection

In the first step, we need to extract the relevant universe of Bunds. We collect the ISINs of all plain vanilla German sovereign bonds from Finanzagentur and then filter the MiFIR data for these bonds. We ignore short-term Bubills, green, inflation-linked and other exotic bonds. To manage the large dataset efficiently, the extraction can be performed in smaller time intervals.

To ensure consistency and remove irrelevant records, transactions containing complex trade components, commodity derivatives, securities financing transactions, or changes in derivative notional are excluded.

When an instrument is bought or sold by a joint account, all account owners must be reported (see ESMA Guidelines, p 42). Thus, there will be multiple buyers or sellers for one transaction. Typically, this occurs when an account is owned by natural persons, e.g. a married couple. For 2019, more than 99% of joint accounts were natural persons, while just under 1% were reported as LEIs, and a negligible fraction were reported as internal accounts (INTC) or market identifier codes (MIC).

Each account owner is assigned a "position", starting from one to the number of owners of the account. To avoid double counting of transactions involving joint accounts we only merge the first owner (position = 1) to each transaction. This ensures that we maintain accurate trade volumes but leads to a slight loss of information that could be useful depending on the specific use case of the subsequent analysis.

In 2019, 3% of transactions involved a joint account, and merging all positions to the transactions data would increase the number of overall transactions by around 3.77%. This effect could be substantial, especially for an analysis focusing on the retail market segment. Any effect on aggregate trade volumes would, however, be smaller as the transaction quantities involved are typically much lower.

The same applies to the buyer/seller decision makers, although at a much smaller scale. Only about 20% of transaction reports include a buyer or a seller decision maker. As before we only include the first decision maker for each transaction. We remove any reports where the buyer and seller are identical.

At this stage the data contains 13 mln observations 130 bonds between 2018 and 2025Q1.

C.3 Normalization

The next step involves preparing the data for analysis by handling intermediary trades, aggregating reports, and identifying outliers in prices and quantities. These steps ensure that the dataset accurately reflects executed transactions and is not distorted by inconsistencies in reporting.

Splitting Reports

Some transactions are reported in a format where the executing party is neither the buyer nor the seller but acts as an intermediary. To standardize the dataset, we follow Jurkatis (2024) and split these intermediary reports into two separate ones — one where the executing entity is recorded as the buyer and another where it is recorded as the seller. This approach ensures that each transaction aligns with a uniform reporting structure.

Under MiFIR, firms are required to report the trading capacity of the transaction, i.e., whether the firm was trading on its own account (DEAL), whether it was an agency trade (MTCH), or whether the firm was trading on any other capacity (AOTC). The latter two cases should list the executing party only as an intermediary while buyer and seller are distinct from it. Table A3 below shows how the trading capacity flag aligns with the trades we split according to the criteria above.

As shown in Table A3, for more than 95% of the transactions, the trading capacity flag aligns well with whether the reports have to be split or not. Where it does not match, we assume the trading capacity flag is misreported and adjust it accordingly, all the while keeping track of the original value. Overall, approximately 33.3% of reports are split.

cat_trading_capacity	split	%
DEAL	0	62.9
AOTC	1	28.78
MTCH	1	4.32
AOTC	0	3.66
DEAL	1	0.23
MTCH	0	0.12

Table A3: Accuracy of Trading Capacity Field

Aggregating Reports

As Jurkatis (2024) points out, in some cases, firms report multiple trades with the same counterparty at the exact same time and price. This makes identifying the duplicate report from the counterparty firm difficult, as it might report the same as a single large transaction. To some extent, this problem is self-inflicted from our previous step of splitting the reports. Assume executing party A reports two matched principal transactions with quantities q1 and q2, that it intermediates between executing party B and clients C1 and C2. At the same time, executing party B only reports a single transaction with quantity q1 + q2 between itself and executing party A. Any duplicate matching method will have difficulties matching these two reports as the quantities differ. As our previous step split the reports filed by executing party A into four reports, two for its trade with executing party B and two for the transactions with its executing party A's clients, we can now aggregate the former two without loss of information. The quantity of the aggregate report then matches that reported by executing party B as should all other fields, and we can identify the two transactions as duplicates in our next step.

Around 22% of trades have to be aggregated. Of these 82% are reports that were split in the step before. The aggregation of reports reduces the overall number of reports by around 16%. Table A4 summarizes the characteristics of the trades that are aggregated with respect to their trading capacity, the trade venue where they occur, the buyer and seller account owner type, and their trade volumes. We can see that this largely concerns off-venue trades, commonly involving a natural person or an internal account and a relatively small trade volume.

One issue that arises at this step is that not all fields of the reports that are to be aggregated match. We condition that the executing party, the buyer and seller account owners, the price, and the timestamp are the same and we sum the transacted quantities. However, for differences in any other field, there is no clear way to aggregate them. Table A5 shows the prevalence of disagreement between aggregated reports for each field. Overall, the disagreement is low and mainly concerns fields of minor interest for our analysis. We thus proceed by simply choosing the value from the first report from a group that is to be aggregated.

Price Outliers

We perform a preliminary identification of outliers for prices and quantities to correct these based on the duplicate report in the next step. If two reports are identified as duplicates and one features an unrealistic price or quantity, we can replace these with the value in the second report.

The data include a field for the unit in which the price is reported. Prices can be reported as percentages, yields, basis points or monetary values. We confirm the finding in Jurkatis (2024) that regardless of what unit is reported, most prices are given as a percentage. Figure A5 below gives an overview of the distribution of prices for each unit category. As in Jurkatis (2024), we

Variable	Split	% in agg trades	% in all data
Split	1	82.41	50
	0	17.59	50
Variable	Trading Capacity	% in agg trades	% in all data
Trading Capacity	AOTC	78.31	45.92
	DEAL	15.27	47.52
	MTCH	6.42	6.57
Variable	Trade Venue	% in agg trades	% in all data
Trade Venue	XOFF	58.25	29.63
	Other	15.55	29.32
	BMTF	13.28	25.65
	TREU	12.92	15.4
Variable	Counterparty Type	% in agg trades	% in all data
Counterparty Type	INTC	48.77	13.88
	LEI	46.08	76.87
	PRSN	3.99	8.2
	NaN	1.06	0.82
	MIC	0.1	0.24
Variable	Trade Size (mn EUR)	in agg trades	in all data
Trade Size (mn EUR)	Mean	4.82	5.43
	Median	0.13	0.5

Table A4: Characteristics of Trades to be Aggregated

Column	Groups with differing Values (%)
eur_quantity	67.33
cat_trade_venue	6.91
cat_short_selling	5.75
cat_exec_type	4.46
ind_investment_party	4.37
cat_trading_capacity	2.69
sdc_seller_decision_maker	2.47
sdc_buyer_decision_maker	1.59
split	1.52
cat_seller_decision_maker_type	1.3
sdc_transmitting_buyer	0.81
cat_quantity_type	0.63
cat_price_type	0.61
sdc_transmitting_seller	0.59
ind_transmission	0.18
cat_inv_dec_type	0.08
cat_buyer_decision_maker_type	0.08

Table A5: Disagreement Among Aggregated Trades

resolve to disregard the reported unit and treat all prices as percentages.

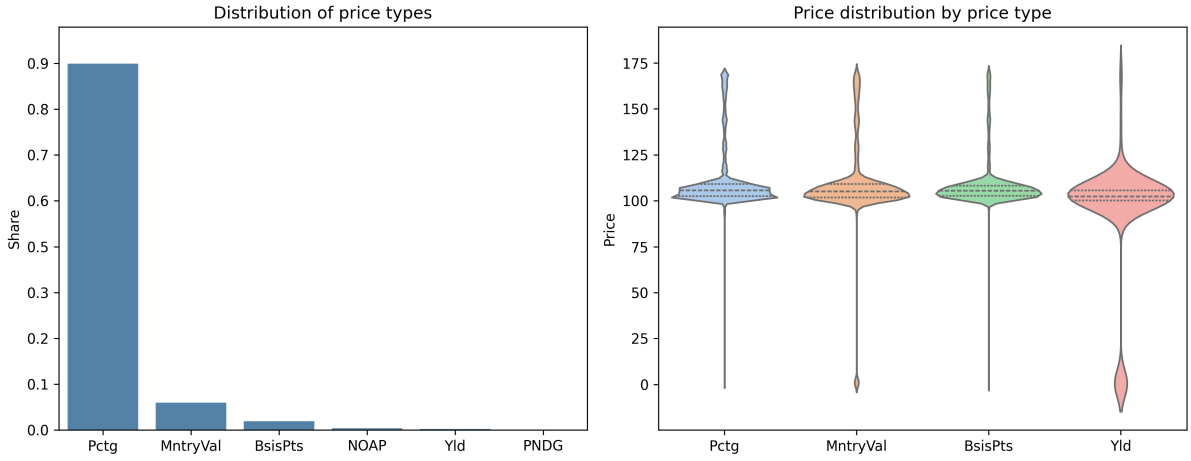


Figure A5: Price Distribution by Price Type

We then proceed to identify price outliers in a two-step process. First, an initial filter flags transactions with prices outside a reasonable range (below 20 or above 500 as a percentage of par). Next, a rolling median price is computed for each bond over a one-week window, and transactions with price deviations exceeding 50% from the rolling median are flagged as potential errors. If the price appears to be a decimal misplacement (e.g., mistakenly recorded in cents rather than percentage points), it is corrected accordingly. This methodology follows the approach outlined in Jurkatis (2024) for corporate bond transactions, adapted here for sovereign bonds. In the end around 0.4% of reported prices are flagged as outliers.

Quantity Outliers

As for prices, transaction quantities can be reported in different units. These are the nominal value, the monetary value. Again, whatever the reported unit, most quantities seem to be reported as nominal values and we treat all as such.

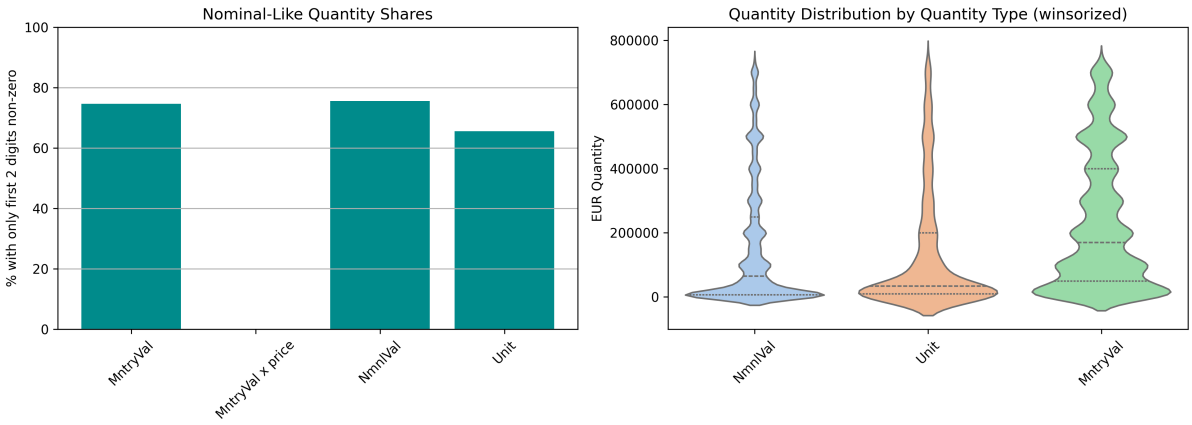


Figure A6: Quantities by Quantity Type: Left panel shows the share of observations in each category (including monetary value times price trades) that is divisible by a multiple of 10. Right panel shows distribution of trade sizes for each category for trades below the median trade size.

We then try to identify potential outliers. Similar to our approach for the price, we start by comparing trade sizes against reasonable thresholds. Extremely small ($\leq$100 EUR) or excessively large ($\geq$1 billion EUR) trades are initially flagged as potential errors. A more refined approach then considers whether the transaction occurred before or after the bond's issuance date. Following Jurkatis (2024), trades reported before issuance are compared against the total issuance volume, and those exceeding this limit are marked as outliers. For transactions near issuance, we divide the reported quantity by the ammount issued and flag as outlier any transaction above the 99.5th quantile. Post-issuance, quantity outliers are identified based on a rolling distribution of trade sizes, ensuring that extreme values are detected and flagged. Overall, less than 1% of quantities are flagged as outliers.

As the dataset at this stage includes anything from small retail transactions to large inter-dealer trades, the range of reasonable transaction sizes is very large making it difficult to identify clear outliers. This step is only meant to improve our algorithm for detecting duplicates. A proper screening for outliers tailored to the analysis will be performed at a later stage. These preparatory steps transform the raw data into a structured dataset suitable for further processing, ensuring consistency and reducing distortions caused by reporting anomalies. The next section will address how duplicate reports are identified and used to improve the data quality.

C.4 Duplicate Matching

Once the transaction data has been prepared, the next step is to identify duplicate records where both counterparties to a trade have reported the same transaction. Since MiFIR requires reporting by all transacting investment firms, some transactions appear twice in the dataset. Removing these duplicates is essential to avoid overestimating trading activity.

Again, we follow the approach developed by Jurkatis (2024) but we adapt it for sovereign bonds. As they do, we compute pairwise comparisons between all trade reports that have the same day and the same ISIN. We then check for each of these pairs to what extent they agree on price, quantity, timestamp, trade venue, buyer and seller, and whether they are reported by different executing parties. However, as these bonds trade much more frequently than corporate bonds we easily run out of memory when trying to compute all possible pair-wise combinations of potential duplicates. We thus try to reduce the sample within which we search for duplicates as much as possible beforehand by taking a closer look at the reporting requirements.

Firstly, the reporting requirements only apply to investment firms, not to natural persons. We thus know that when a transaction on one side only involves a natural person, it will not have a duplicate. This is the case when both the buyer/seller account owner and decision maker are natural persons or only the former while the latter is not reported. All transactions that fulfil these requirements are removed from the dataset in which we match duplicates.

We can proceed in a similar way when the buyer/seller account owner is reported as an internal account ('INTC'). As long as this side of the trade does not list an investment firm as account owner or decision maker we can exclude it from our search for duplicates. Note that this is possible because we split the trades in the first step meaning that there are at most two parties to a trade, one of which is the reporting executing party. We thus only need to check whether the other party also reports.

Removing transactions involving a natural person reduces the sample by approximately 7.2% while dropping internal account trades removes another 5.3%. It is especially useful to remove these trades as they can occur in large batches at the same time in the same ISIN. An executing party may report up to 45 000 trades at the same time for the same instrument. Searching for duplicates among these trades can easily breach the bounds of the available computing power.

Table A6 shows the characteristics of these batches of trades for any batch with more than

Metric	Batched Trades	All Trades
Number of transactions	65647	1768657
Number of trade batches (>1000)	33	-
Average trade size (EUR)	17,048.56	6,470,922.04
Median trade size (EUR)	2,300.00	729,000.00
Minimum group size	1103	-
Maximum group size	7730	-
Mean group size	1989.30	-
Top counterparty types (%)		
PRSN	96.7%	9.2%
LEI	3.3%	84.3%
INTC	0.1%	5.4%

Table A6: Batched Trades

1000 trades. They typically involve natural persons as counterparties and have a small trade volume. There are 33 of these batches with over 65000 trades. Dropping trades that will not have a duplicate reduces this to 2 batches with 3160 trades. Approximately 25% of the trades dropped are part of a large batch of trades. We use a python record linkage package to compute the pairwise comparisons. This package allows us to impose that to be compared two trades must occur on the same date and in the same ISIN. This limits the number of pairs to compare. Another condition we can impose because it must apply for the trades to be duplicates, is that the executing parties reporting the two trades are distinct. The record linkage package does not directly offer the possibility to filter pairs on the condition of a field differing. We thus develop our own class to include this condition and thereby reduce the number of pairs and the memory required for the comparison.

We end up with approximately 149 million pairs to compare per year. Following the steps of reducing the sample size will hopefully prevent any memory errors. Should this still be a problem, one can try reducing the sample period, even potentially running the deduplication section separately for each day in the sample.

To identify duplicates among the selected pairs, we apply a set of comparison functions to assess similarity across key fields, including counterparties, trade venue, trade quantity, price, and timestamp. For each field, we assign a score to measure the similarity of the values reported by the two transaction records. This score can either be binary or continuous depending on the field.

Trade timestamps are matched within a 20-minute window, plus a shift of one or two hours. This accounts for possible reporting delays and time zone discrepancies. While all investment firms are supposed to report in UTC, this could be a potential source of error. The score assigned for this field is continuous. Apart from shifts of exactly one or two hours, we check how close the timestamps are and assign the score based on this. The score is set to zero when the discrepancy in the timestamps exceeds 20 minutes.

Trade quantities and prices are checked for exact matches, as well as potential reporting errors involving common scaling factors (e.g., misreporting by a factor of 10 or 100). Buyer and seller account owner are checked for exact matches. All four fields receive a binary score.

Lastly, the trade venue is compared. Here we usually require an exact match, unless one side reports an off-venue trade ('XOFF'). This is because through splitting trades in the first step, the client side of a trade might still have the trade venue of the market side when the two sides were initially reported as one trade. Assuming the client also reports, he will only report the

client side which occurs outside of any trade venue.

The similarity scores from these comparisons are aggregated, and transactions exceeding a predefined threshold are considered duplicate matches. This step differs from the approach taken by Jurkatis (2024) which uses a probabilistic approach to filter for the most likely duplicates. While this allows for a more precise target rate of identified duplicates, by taking into account the unconditional probability of a trade having a duplicate, our approach allows for more discretionary control over the quality of the matches. We optimize the chosen threshold by manually checking the quality of trades marginally above and below it.

The final step employs graph-based matching to ensure that each transaction is linked to at most one counterpart report, preventing duplicate pairings, while maximizing the overall score across all final matches. Figure A7 illustrates this approach.

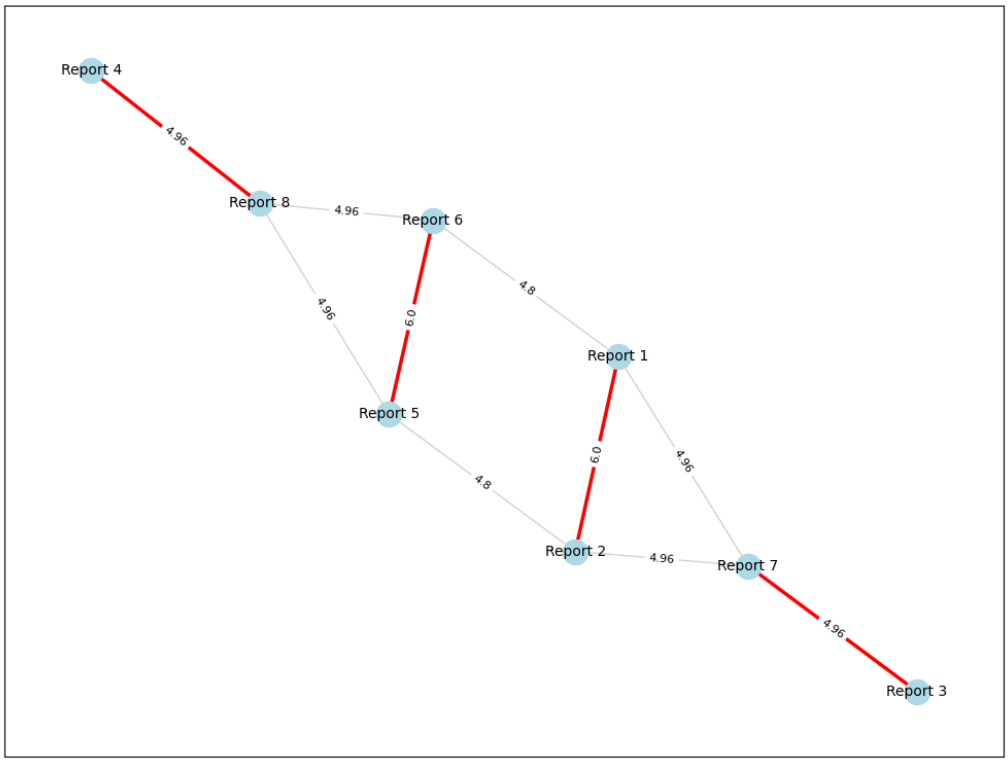


Figure A7: Maximum Weight Matching

We manage to find a duplicate for about 54% of all trades or 61% of the trades that we included in our search for duplicates. The average score of the final set of matches is 5.64 out of a maximum attainable score of six. Table A7 shows the average score across all matches for each field. We can see that matched trades agree on all fields in more than 50% of the cases. For all fields the rate of agreement is at least close to 90%. Matched trades most often disagree on buyer or seller information. This could be caused for example by one side reporting the parent company instead of the subsidiary with whom they traded.

C.5 Correcting Reported Entries

After duplicate trades have been identified, the final step in the data cleaning process is to resolve inconsistencies within matched pairs. Errors in prices, quantities, counterparties, and

Buyer	Seller	Trade Venue	Quantity	Price	Timestamp	Total Score
0.89	0.89	0.92	1.0	0.98	0.97	5.64
0.0	0.0	0.0	0.0	0.0	0.0	4.0
1.0	1.0	1.0	1.0	1.0	1.0	5.0
1.0	1.0	1.0	1.0	1.0	1.0	6.0
1.0	1.0	1.0	1.0	1.0	1.0	6.0
1.0	1.0	1.0	1.0	1.0	1.0	6.0

Table A7: Match Quality

trade venues are corrected using information from both records in a matched pair. This closely resembles the approach taken by Jurkatis (2024).

Price and Quantity Corrections

If only one transaction in a matched pair is flagged as an outlier, its price or quantity is replaced with the corresponding value from the non-outlier transaction. If both transactions in a pair are flagged but report the same price or quantity, the outlier flag is removed. Any remaining differences between prices or quantities after correction are recorded for later evaluation.

Counterparty Adjustments

Discrepancies in reported buyers and sellers are addressed by checking whether one transaction misidentifies the counterparty. If one report lists the executing party in the buyer or seller field while the other does not, the erroneous entry is corrected. If both reports list different counterparties inconsistently, they are flagged as unresolvable mismatches.

Trade Venue Standardization

If a matched transaction reports 'XOFF' as the execution venue, this value is enforced across both records to maintain consistency. This adjustment ensures uniform classification of off-venue trades.

These corrections refine the dataset by leveraging duplicate reports to resolve inconsistencies, improving the accuracy of transaction-level data. Table A8 gives an overview of the corrections made and the remaining differences within a matched pair of transaction reports.

Field	% Disagreement
Price	1.86
Quantity	6.71
Trade Venue	8.48
Buyer	2.2
Seller	2.24
Timestamp	73.54

Table A8: Remaining Disagreement Between Duplicates

C.6 Primary Issuances

The Mifid II data includes primary market issuances. These are reported as primary dealers buying bonds from the federal German government on auction day. Dealers either report Finanzagentur or Bundesbank as sellers (the latter acts as fiscal agent in the issuance process). Since we are interested in the secondary market, we hence drop all observations where Finanzagentur or Bundesbank sell a bond on that bond's auction day to an investor that was primary dealer on that day. We keep all other trades by Finanzagentur and Bundesbank, which mainly capture market maintenance trades and the Eurosystem's asset purchase programs.

C.7 Trades by CCPs, IDBs and venue operators

Due to the reporting standards and our normalization of the data, 'intermediaries' will appear to be buyer or seller even if they do not actually take the bond into inventory at any point in time. This is in particular the case for CCPs and interdealer brokers (IDBs) and to a lesser extent venue operators. We attempt to combine both legs of such intermediated trades in order to properly reflect the role of intermediaries and not to inflate transaction volumes. We proceed as follows.

Case 1: The intermediary is listed as buyer or seller and as the executing party of the transaction. This case concerns any transactions where the intermediary reports as an executing party in our data. We proceed hierarchically in two steps:

First, we recombine trades we had split in the first chapter as long as they have not been aggregated in the meantime. We drop any duplicates associated with either leg of the recombined trades. Second, for the remaining trades we take all transactions where the executing party is an intermediary and match buy and sell transactions where they involve the same bond and happen at the same time. One-to-one matches are easily recombined into a single transaction. For one-to-many matches, we split the largest trade to match the several trades of the other leg. In few cases, several buy transactions are matched to several sell transactions or vice versa. For simplicity, we disregard these. Lastly, again we delete all duplicates of either leg of a intermediated transaction that we were able to match. This provides the most consistent format for our data and avoids double counting transactions. Our approach is illustrated in Table A9 below.

Table A9: Case 1: Intermediated Trades — Original and Result

(a) Original

Tx_id	Dup_id	Exe	Buyer	Seller	Time	Price	Qty
1	1	A	CCP	A	10:00:00	103	100,000
2	1	CCP	CCP	A	10:00:00	103	100,000
3	2	CCP	B	CCP	10:00:00	103	100,000
4	2	B	B	CCP	10:00:00	103	100,000

(b) Result

Exe	Buyer	Seller	Time	Price	Qty
CCP	B	A	10:00:00	103	100,000

Case 2: The intermediary is listed as buyer or seller but not as the executing party of the transaction. The steps above concern any reports where the executing party is an

intermediary. However, occasionally it happens that the intermediary is reported as the buyer or seller of the transaction by another executing party. This is because the reporting executing party might not know the actual buyer or seller. In general, these reports should have a duplicate filed by the intermediary. If, through our duplicated linking procedure, we managed to identify this duplicate, we keep it and delete the other report. This allows us to get rid of reports where the intermediary is incorrectly listed as the buyer/seller of the transaction.

If we do not have a duplicate for such a transaction, we create one, as shown in Table A10. This new duplicate report will be identical to the original one except that the executing party becomes the intermediary. We then delete the original report and apply the procedure described in Case 1 to the synthetic duplicate in order to match to the other side of the intermediated trade.

Table A10: Case 2: Intermediated Trades — Original, Synthetic, and Result

(a) Original

Tx_id	Dup_id	Exe	Buyer	Seller	Time	Price	Qty
1	1	A	CCP	A	10:00:00	103	100,000
3	2	CCP	B	CCP	10:00:00	103	100,000
4	2	B	B	CCP	10:00:00	103	100,000

(b) Synthetic Report Creation

Tx_id	Dup_id	Exe	Buyer	Seller	Time	Price	Qty
1	1	A	CCP	A	10:00:00	103	100,000
Synthetic	1	CCP	CCP	A	10:00:00	103	100,000
3	2	CCP	B	CCP	10:00:00	103	100,000
4	2	B	B	CCP	10:00:00	103	100,000

(c) Result

Exe	Buyer	Seller	Time	Price	Qty
CCP	B	A	10:00:00	103	100,000

These steps do not entirely eliminate the pure intermediaries from our data. Where we are unable to match both legs of the intermediated transaction or to link the report filed by the intermediary to its duplicate (filed by the actual buyer or seller), this sector will still appear as a transacting party. However, we believe our approach is helpful to reduce the extent of this problem.

In fact, we confirm that for the trades we link, the trading venue aligns with the intermediary. For instance, trades where LCH Clearnet was the intermediary almost always occur on MTS, whereas for Eurex Clearing it is Frankfurter Wertpapier Boerse. The same is true for trades intermediated by IDBs, like BGC Brokers and GFI, where the venue field is filled with the MIC code of the respective OTF. To ensure consistency, we overwrite the trading venue field in the few cases they do not align.

C.8 Investor-Sector Mapping

We construct an exhaustive mapping of LEIs to the following investor sectors: primary dealers, banks, inter-dealer broker (IDB), investment funds, insurance companies and pension funds

(ICPF), official (incl. central banks), non-financial corporations (NFC), and non-bank financial intermediaries (NBFI).

Section C.8 lists our data sources. We prioritize high-quality regulatory data. The list of primary dealers participating in Bund auctions is available from Finanzagentur. We augment this with the European Central Bank's (ECB) Single Supervisory Mechanism (SSM) lists of Significant Institutions (SIs) and Less Significant Institutions (LSIs) and the ECB's list of Monetary Financial Institutions (MFIs), Investment Funds (IFs), Pension Funds (PFs), and Insurance Corporations (ICs). As the final regulatory data source, we use the UK's Prudential Regulation Authority (PRA) register of regulated banks.

For non-European investors, reliable LEI-based investor classifications are not readily available. Hence, we proceed with two proprietary data sources. First, we use all available LEI-sector mappings from Moody's BvD BankFocus database. This data is especially useful to identify large global banks. Second, we use Morningstar data to identify hedge funds. Next, we apply a manual cross-check for the most important investors in our dataset.

For any remaining LEIs, we first use the GLEIF Level 1 "Who is Who" register[17] and classify investors whose entity category is 'FUND' as investment fund or as hedge fund if the fund is domiciled in the Cayman Islands. Lastly, we use the ECB's Register of Institutions and Affiliates Database (RIAD, see e.g. Perrella and Catz, 2020), which contains a mapping from LEI to sectors (based on the European system of accounts framework).[18]

Even using all these data sources, we noticed suspiciously many large investors were either unclassified or classified as NBFI. Hence, we apply a final correction based on the investor-grouping procedure described below. Specifically, for group heads classified as NBFI or unknown, we infer their sector from the trading activity of their subsidiaries. If a single sector accounts for more than 80% of the group's total trades, and this sector is neither NBFI nor missing, we reassign the group head's classification accordingly. This adjustment corrects holding-company artifacts and ensures that active financial groups are consistently represented in their dominant sector.

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

[18] We avoid using RIAD as much as possible as it is notoriously unreliable. Apart from omitting many important investors, RIAD contains dubious sector classifications. Many Global Systemically Important Banks (G-SIBs) outside the euro area, e.g., are classified as non-banks (ESA codes S.125-S.127), either because they are no depository institution in the euro area or because they operate as a holding company.

Table A11: Investor Reference Data Sources (Hierarchical Order)

Source Name	Description	Link
Finanzagentur	Official auction group list by Germany's debt management office	link
ECB SSM (SIs)	ECB list of systemically important credit institutions under direct supervision	link
ECB SSM (LSIs)	ECB list of smaller credit institutions under national supervision	link
ECB (MFIs)	ECB list of monetary financial institutions, including credit institutions and money market funds	link
ECB (IFs)	ECB quarterly register of investment funds by LEI and investment policy	link
ECB (PFs)	ECB register of pension funds by country of residence and LEI	link
ECB (ICs)	ECB register of insurance corporations by country of residence and LEI	link
UK PRA	UK Prudential Regulation Authority list of banks authorized in the UK	link
Moody's BvD BankFocus	Proprietary database incl. LEI-sector mapping	
Morningstar Hedge Funds	Proprietary database of hedge funds incl. LEI	
Manual Classification	Manually curated LEI-sector mapping	
GLEIF Funds	Global Legal Entity Identifier Foundation (GLEIF) Level 1 entities categorized as investment funds	link
ECB RIAD	ECB Register of Institutions and Affiliates Database (RIAD)	

Appendix D Descriptive Statistics and Cross-Checks

Since our transaction-level dataset is new to the literature, this section provides some descriptive statistics and cross-checks with other established datasets.

Figure A8 shows that our Mifid II dataset covers an order of magnitude more Bund trading volume than the older Mifid I dataset studied e.g. by de Roure et al. (2025). Our data also covers slightly more volume than the survey-based dataset of the Finanzagentur, Germany's debt management office, see Branger, Muck, and Pütz (2024). A key difference is that the FA dataset only covers trades in which a primary dealer is involved. In contrast to our dataset, trades between European non-dealers are not covered while trades between two non-European investors might be covered, e.g. a trade between a US branch of a German dealer bank and a US investment fund.

Figure A9 shows that Bund trading activity follows a distinct intraday pattern, with low volumes before 9 a.m., a midday lull between noon and 2 p.m., and a surge in trading around 5 p.m. followed by a sudden stop at 6 p.m.

Figure A10 shows the wide dispersion of transaction sizes. Roughly 13% of trades have a volume of less than €1,000 while 9% have a volume of more than €10 million.

Figure A11 shows that the 'noise' in Bunds, i.e. the mispricing of individual bonds around a fitted yield curve is of similar magnitude compared to US Treasuries.

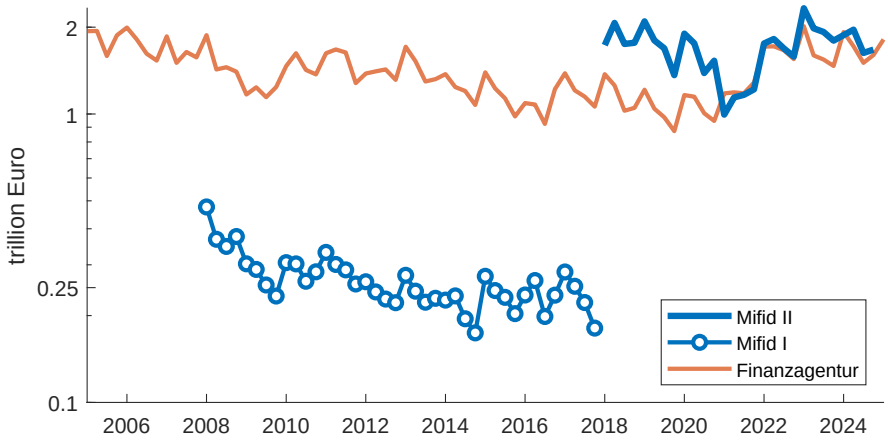


Figure A8: Monthly Trading Volume.

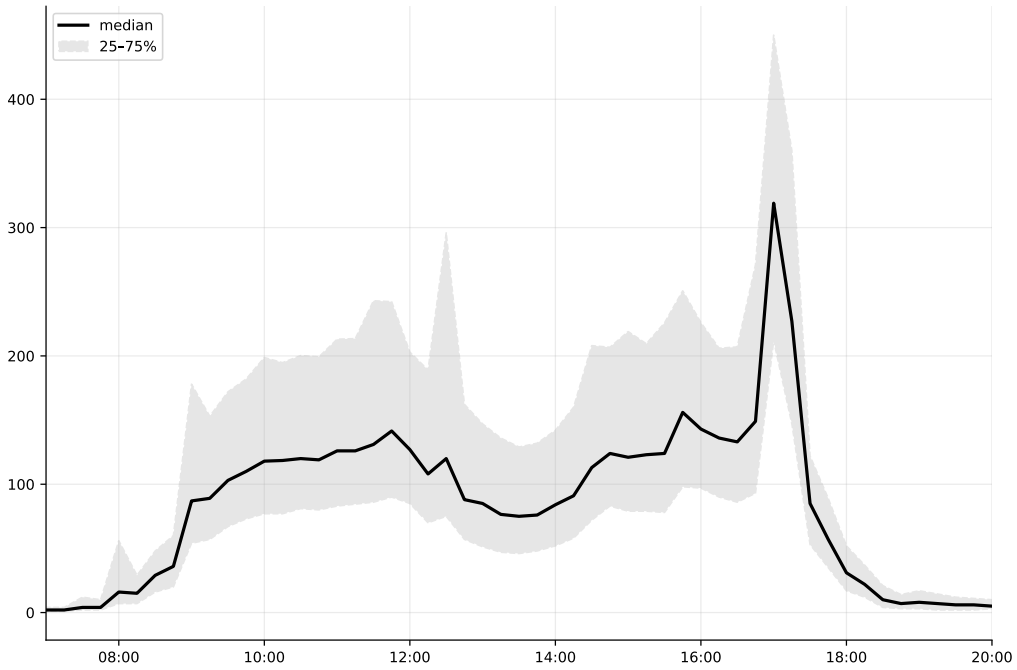


Figure A9: Number of trades across the day. This figure shows the median and interquartile range of the number of trades per 15-minute interval throughout the day.

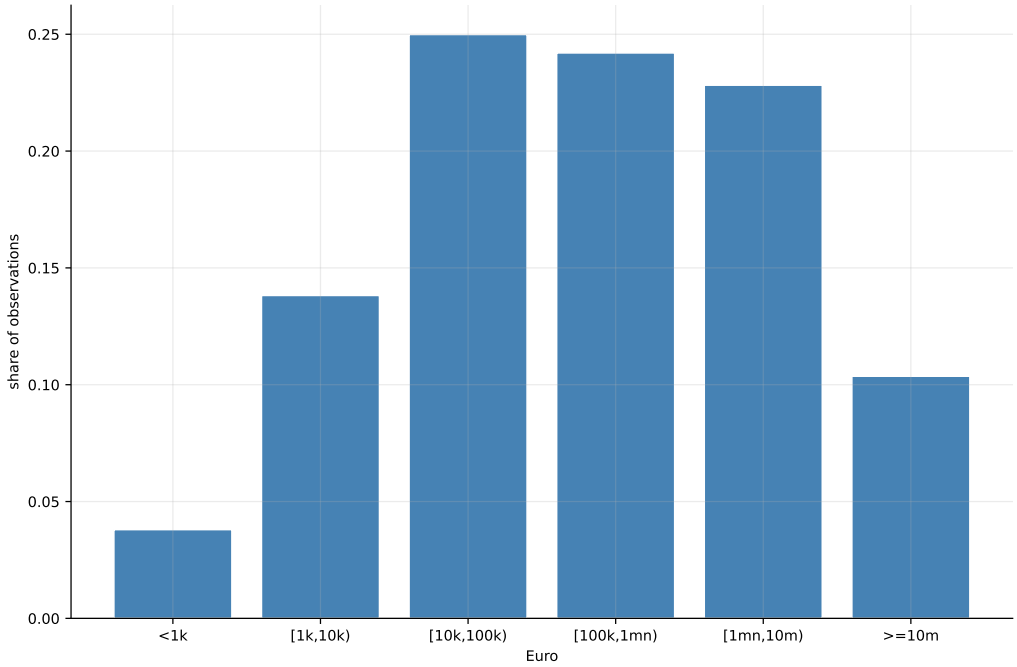


Figure A10: Distribution of transaction sizes. This figure shows the distribution of transaction sizes in Euros.

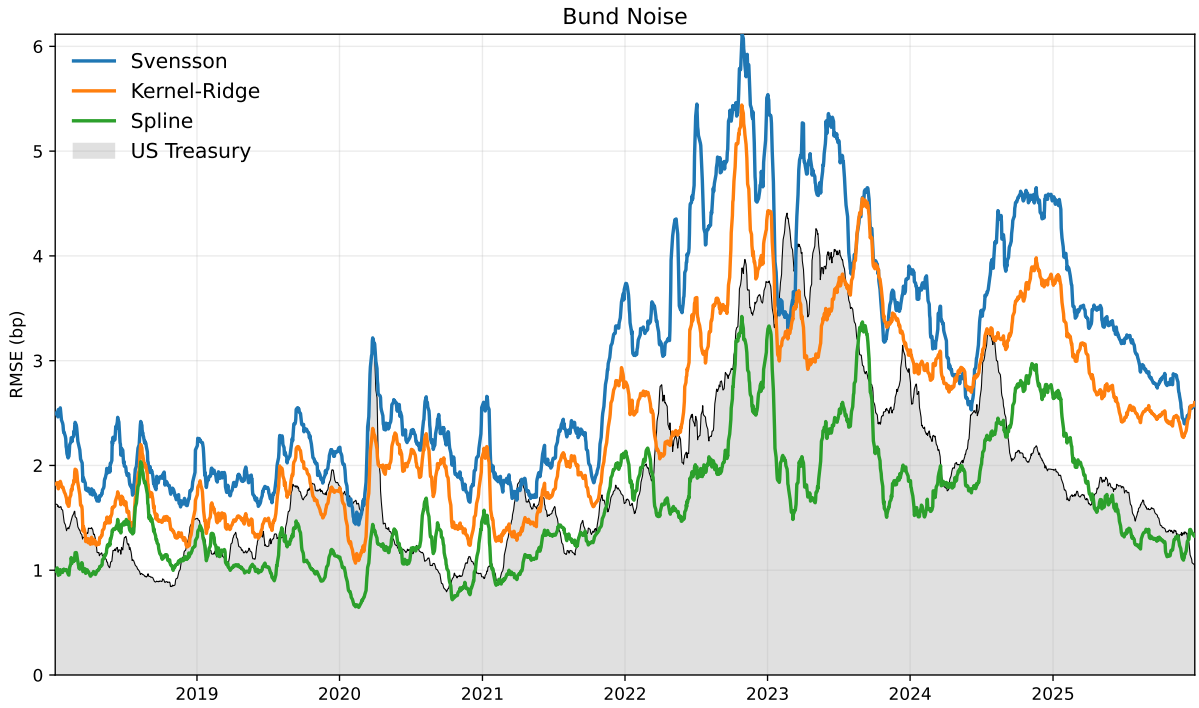


Figure A11: Bund market noise over time. This figure plots the root mean squared error (RMSE) between transaction-weighted implied yields and three fitted yield curves: using the Bundesbank's Svensson (1994) model, Filipović et al. (2025)'s kernel-ridge regression, and a smoothed cubic spline. For each trading day, the RMSE is computed across all Bund ISINs with observed cash market transactions. Section 2.5 provides details. For reference, the grey shaded area shows the US Treasury noise measure of Hu et al. (2013), updated by the authors. They use end-of-day quotes from CRSP and a Svensson (1994) fitted yield curve, but in contrast to our Bund measures, they exclude bonds with less than one or more than ten years to maturity and bonds whose yield is four standard deviations away from the model-implied yield. All series are shown as 14-day moving averages.