Internet Appendix
for
Anatomy of a Safe Asset: Bund Market Structure and Reform

by Mark Kerssenfischer, Konrad Lucke and Willy Scherrieble

Section A provides additional results.
Section B details how we adapt the yield curve fitting method of Filipović et al. (2025) to Bunds.
Section C provides descriptive statistics and cross-checks of our dataset.
Section D details how we clean and process the MiFID I and MiFID II/MiFIR transaction reports, accounting for the complexity of the Bund market.
References are listed in the main paper.

Appendix A Further Evidence

A.1 The Dealer System

This appendix supplements Section 4.2. Figure A1 shows the display behind the club fact: the Bietergruppe members' share of intermediated client volume by half-year, under the baseline behavioral screen and across screening configurations. The paragraphs below document two further properties of the dealer hierarchy: its geographic division of labor and its stability in acute stress windows.

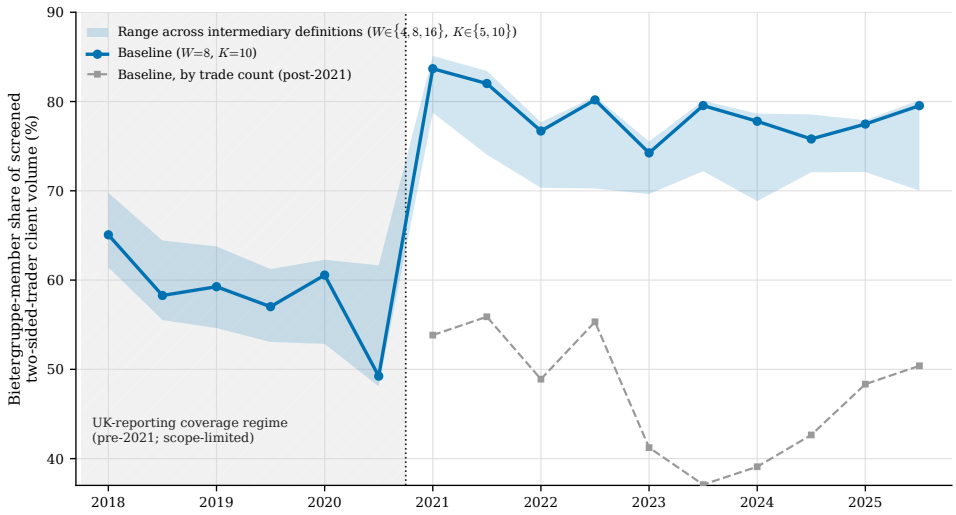


Figure A1: Bietergruppe members' share of intermediated client volume by half-year. The line is the baseline behavioral screen (round-trips in at least $W = 8$ ISIN-weeks and at least $K = 10$ counterparty groups per half-year); the shaded band spans the six screening configurations ($W \in \{4, 8, 16\}$, $K \in \{5, 10\}$) and reflects definitional robustness, not sampling uncertainty. The thin line is the members' share of trade count after 2020. The screened slice carries 25.8% of trading volume outside intragroup bookings (Table 2). Values before 2021 reflect the UK-reporting coverage regime (hatched) and are not comparable to later years.

The hierarchy sorts geography The division of labor across the hierarchy is geographic. German headquartered clients account for 19.4% of the top tercile dealers' institutional client volume but 55.6% at the bottom tercile—a top-minus-bottom gradient of -36.2 percentage points—negative in every half year after 2020 and within every client sector, and surviving comparison within venue, trade size, and maturity cells (-31.6 points; Figure A2). Part of the gradient reflects composition—the bottom tercile is disproportionately populated by small German headquartered dealers—but at most a third: holding dealer nationality fixed alongside venue, size, and maturity leaves a gradient of -25.4 points. The client sector mix differs across tiers as well: the dissimilarity index between top and bottom tercile sector shares is 14.0 percentage points, against a permutation benchmark of 3.1. The bottom tercile executes 1.8% of client volume, so the division of labor is a statement about whom the tail serves, not about how volume is split: the core of the hierarchy intermediates for foreign investors on electronic venues, while the tail is a domestic, voice-leaning niche. Geographic segmentation of bond portfolios is well documented at the investor level—home bias in sovereign holdings (Coeurdacier and Rey, 2013; Maggiori, Neiman, and Schreger, 2020) and gravity in cross-border asset trade (Portes and Rey, 2005)—and euro area capital markets are less integrated than headline holdings suggest (Beck, Coppola, Lewis, Maggiori, Schmitz, and Schreger, 2024). We locate the same segmentation one

layer deeper, in who intermediates for whom: even in the euro area's most homogeneous safe asset, the dealer system sorts clients by geography.

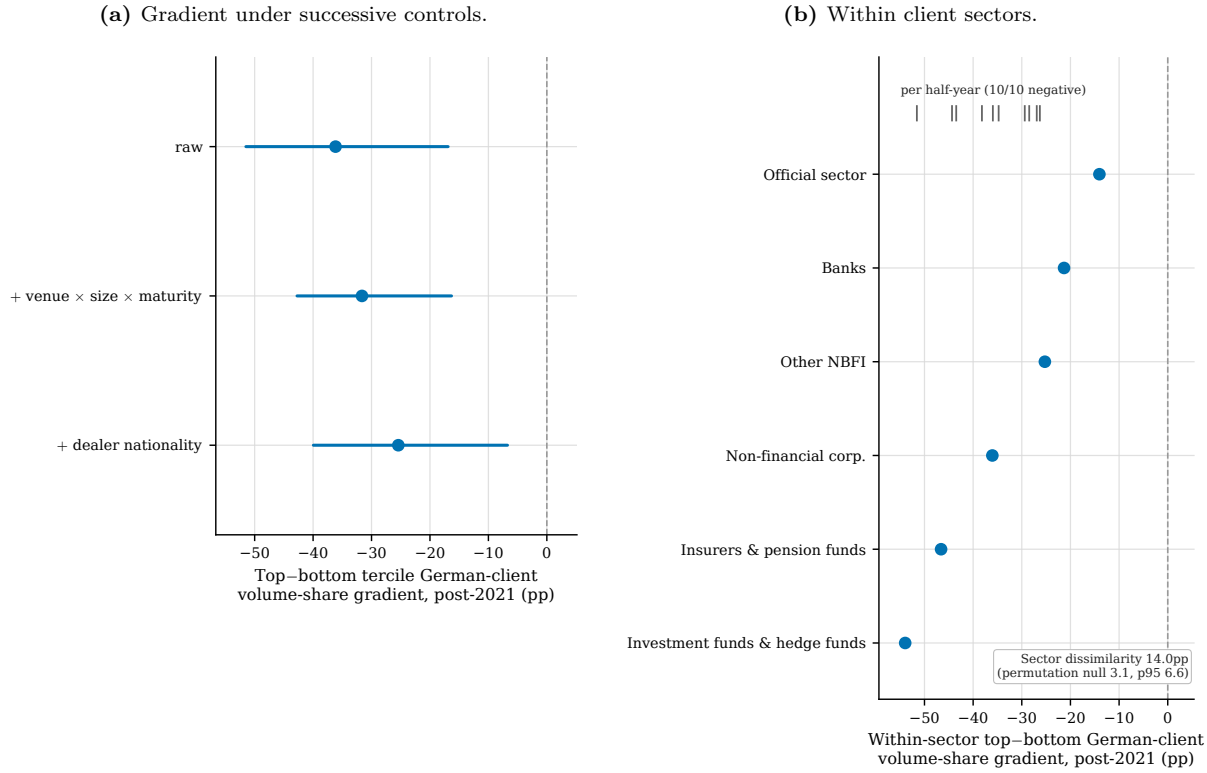


Figure A2: Dealer hierarchy and client geography after 2020. Panel (a): difference between top and bottom centrality terciles in the German headquartered share of institutional client volume—raw, within venue–size–maturity cells, and additionally holding dealer nationality fixed—with 95% dealer-clustered confidence intervals. The control ladder bounds composition rather than decomposing it: a shift-share split attributes between roughly none and a third of the raw gradient to dealer-nationality composition. Client and dealer nationality are the countries of the investor group's ultimate parent, so German arms of foreign parents count as foreign. Panel (b): the same gradient within each client sector, shown descriptively without confidence intervals; inference is in panel (a). The small "other" sector is omitted; its gradient is the smallest.

The core under stress In the three acute stress windows studied in Section 5.1, the dealer hierarchy does not systematically move. The top tercile's share of institutional client volume changes by 2.0 percentage points relative to calm weeks before each episode (95% confidence interval $[-3.2, +7.3]$; Figure A3). The calm weeks are the 40 business days that end one week before each episode, leaving out the days without account owner information (Section D.2); for 2025 they start on 18 November 2024. The level is high in both states: the top tercile intermediates 75.1–81.8% of institutional client volume in calm weeks and 77.3–84.5% in stress weeks. Placebo March windows in calm years move by similar amounts. The stress episodes therefore reshuffle who buys and sells Bunds, not which dealer tier intermediates the flow. This is not a claim that dealer balance sheets are unconstrained; it is a narrower null on reallocation within the dealer hierarchy.

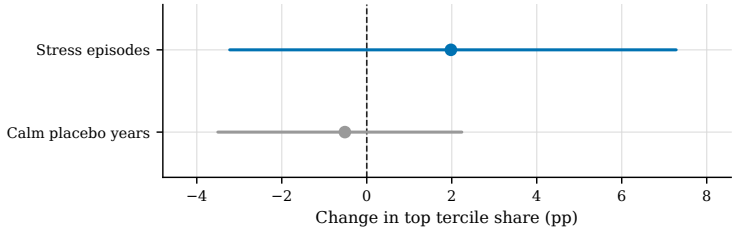


Figure A3: Top tercile dealer share in stress and placebo windows. Dots report pooled changes in the top tercile's share of institutional dealer–client volume relative to the 40 business days that end one week before each episode, with 95% dealer group bootstrap confidence intervals. The stress estimate pools the three acute stress windows; the placebo estimate pools same calendar March windows in calm years.

A.2 Price Discovery by Venue Type and Venue Protocol Family

The informed flow test in Figure 13 can also be cut by where trades execute. We repeat the permanent impact exercise by venue type and, as a secondary descriptive cut, by venue protocol family: for each category, we sign dealer-to-client flow from the client's side, aggregate to bond–day–category net flow, and test whether net buying predicts a subsequent decline in the curve-relative pricing error over one to twenty trading days, absorbing ISIN and maturity bucket×day fixed effects and controlling for the bond's lagged curve-relative pricing error and its two lagged daily changes, as in Figure 13. These cuts are diagnostic, not operator-level claims: venue type and venue protocol family are MIC-level classifications pooled across venues and dealers, and no venue operator is named.

The venue protocol family does not restate the venue types (Table A1). A MIC enters the RFQ family or the order book family only when the venue's own rulebook or its operator documents that model; venue type alone never assigns either family. Systematic internalizers, venues that combine an order book with RFQ trading, and venues without a source are unknown; they account for 20.5% of dealer-to-client volume (Appendix D.11). Among the classified families, RFQ family net flow has a permanent impact of +0.037 bps at the ten-day horizon ($t = 6.8$), and net flow in the residual other family, mostly off-venue reports and OTF trading, has +0.005 bps ($t = 1.0$). Order book family flow, 1.3% of classified volume and almost entirely on records without an identified client, is close to zero at that horizon (+0.007 bps, $t = 0.4$). These estimates compare venue families; they do not show how individual trades executed.

Venue type	RFQ	Other	CLOB	Unknown
MTF	87	4	2	7
OTC	0	100	0	0
SI	0	0	0	100
RM	0	45	40	15

Table A1: Venue protocol family by venue type. Each row reports the composition of a venue type's dealer-to-client volume by venue protocol family (row-normalized percentages over all four labels). The family is a MIC-level label: RFQ family, order book family (CLOB), other (residual), or unknown. Systematic internalizers are unknown by construction. Pooled across all operators of each venue type. MiFID II subsample.

Venue Type The venue-type cut compares multilateral electronic platforms (MTFs), bilateral over-the-counter venues (OTC), single-dealer systematic internalizers (SIs), and regulated markets (RMs), mostly retail exchanges; OTF trades enter only the protocol family cut, in the residual other family. Figure A4 shows that price discovery differs across venue types. Net flow on MTFs, including platforms such as Bloomberg and Tradeweb, predicts a permanent price impact that builds over the first ten days—0.022 bps after one day, 0.039 bps at the ten-day horizon ($t = 7.6$)—and remains at 0.030 bps after twenty. Bilateral OTC and single-dealer SI flow are also informative and permanent (0.020 and 0.033 bps at ten days; $t = 3.5$ and 5.0). Regulated market flow, by contrast, contributes little: the RM price impact is close to zero, −0.022 bps ($t = 1.1$), and is not distinguishable from zero at every horizon.

This is not simply sector composition. Holding the client sector fixed, the same flow is more informative on MTF than OTC: within hedge fund flow alone, the MTF impact is 0.039 bps ($t = 6.6$), versus 0.022 bps ($t = 2.8$) on OTC. The regulated market null may partly reflect lower power—RM is only 0.5% of volume—but the point estimate is near zero through the ten-day horizon, not merely imprecise. The results do not depend on the anonymous interdealer records that the dealer-to-client filter admits: restricted to records with an identified non-dealer client,

MTF, OTC, and SI flow remain significant and permanent at every horizon and regulated market flow is still not distinguishable from zero, although SI flow then carries as much information as MTF flow.

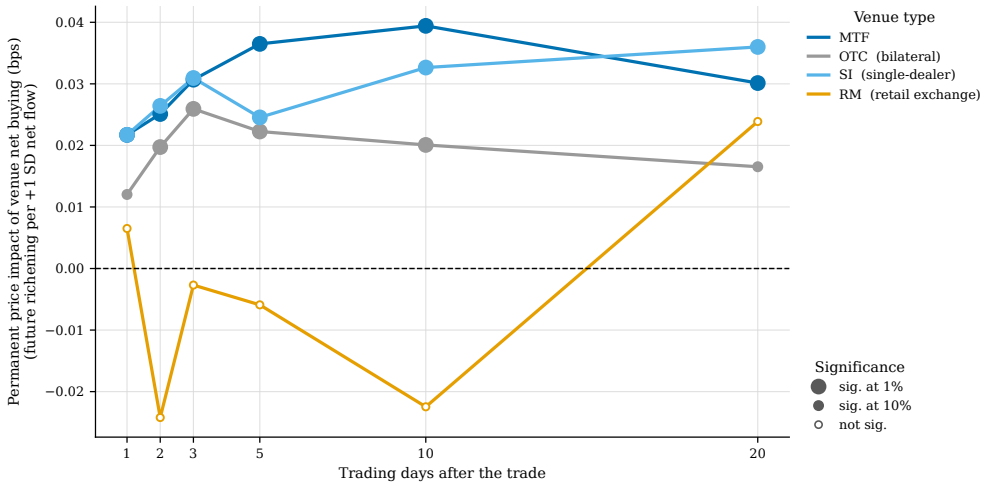


Figure A4: Permanent price impact by venue type. For each trading venue category, the figure reports the cumulative future change in the curve-relative pricing error, defined as synchronized implied yield minus the kernel ridge fair yield, in bps, per net purchase of the bond equal to one standard deviation of the venue type's absolute net purchases, by trading days after the trade, signed so that positive values denote richening. Venue-type-by-horizon regressions absorb ISIN and maturity bucket×day fixed effects and control for the bond's lagged curve-relative pricing error and its two lagged daily changes; marker size and fill encode significance using ISIN-clustered standard errors. Dealer-to-client trades, MiFID II subsample (2018–2025).

A.3 Informed Flow Measurement Robustness

The benchmark informed flow test in Figure 13 uses future changes in the curve-relative pricing error, because the economic question is whether a sector selects bonds whose curve-relative pricing error subsequently declines. For sector k and horizon h, we estimate $e_{i,t+h} - e_{i,t} = \beta_{k,h} f_{i,t,k} + \gamma' z_{i,t} + \alpha_{m(i,t),t} + \alpha_i + \varepsilon_{i,t}$ by unweighted least squares on the bond-days with net flow from sector k, where e is the curve-relative pricing error, $f_{i,t,k}$ is the sector's signed euro net purchases divided by the standard deviation of its absolute net purchases and capped at five, $z_{i,t}$ holds the lagged pricing error and its two lagged daily changes, $\alpha_{m(i,t),t}$ is a fixed effect for each day and bucket $m(i,t)$ of the bond's remaining maturity (under 1, 1–3, 3–7, 7–15, and over 15 years), and α_i is an ISIN fixed effect. The maturity bucket by day effects absorb yield and fitted curve moves common to bonds of similar maturity. Horizons count the bond's subsequent trading days in the sample, so they skip days without trades in the bond and the days without account owner information (Section D.2). Standard errors are clustered by ISIN. Throughout, a one standard deviation net purchase means $f_{i,t,k} = 1$.

Figure A5 repeats the same sector-level permanent impact exercise with the future change in the bond-day mean synchronized implied yield as the outcome, which does not use fitted curves. The outcome averages the trade-level yields unweighted, the same aggregation as the

benchmark pricing error; the fixed effects and the lagged pricing error controls are those of the benchmark.[22]

The hedge fund result carries over: hedge fund net purchases predict a future yield decline of 0.05 bps at the ten-day horizon ($t = 6.8$). So does the official sector result: official sector net purchases predict a decline of 0.06 bps ($t = 3.8$). Investment fund flow, transitory in the benchmark, is not distinguishable from zero here either (0.01 bps, $t = 0.6$). Bank and ICPF flow do not carry over: their ten-day estimates are not distinguishable from zero ($t = 0.7$ and 0.9). Retail household flow has no estimate different from zero at one to five days, 0.11 bps at ten days ($t = 2.1$), and an estimate not distinguishable from zero at twenty days ($t = 1.0$). We keep the curve-relative pricing error as the benchmark because raw yield changes are less directly tied to within-segment security selection.

The benchmark specification controls for the bond's cheapness before the flow day because some sectors buy bonds that have just cheapened, and cheapness mean-reverts. Hedge fund and official sector purchases follow a cheapening of the bond on the previous trading day (0.013 and 0.011 bps per standard deviation of net flow, $t = 4.7$ and 2.9); investment fund purchases do not ($t = 0.1$). On the same sample without these controls, the ten-day estimates are larger for the official sector (0.07 bps, $t = 6.7$) and investment funds (0.03 bps, $t = 3.9$) and essentially unchanged for hedge funds (0.04 bps). Controlling instead for the flow day's own pricing error and its changes on the flow day and the day before, which the flow itself moves, gives a lower bound: 0.02, 0.03, and 0.01 bps for hedge funds, the official sector, and investment funds ($t = 6.0$, 3.3, and 1.8). Under all three control sets, hedge fund flow has the most precisely estimated impact and official sector flow the largest.

With day and ISIN fixed effects in place of the maturity bucket×day fixed effects, and the same controls, the ten-day estimates are 0.04 bps for hedge funds ($t = 8.2$), 0.05 bps for the official sector ($t = 6.2$), 0.02 bps for banks ($t = 2.0$), and 0.08 bps for retail households ($t = 3.0$). Bank and retail purchases coincide with moves common to bonds of similar maturity; within the maturity segment their estimates are not distinguishable from zero. The retail estimate with day fixed effects is flat for three days and then rises into a hump that is not significant from forty days on.

A.4 Price dispersion

A natural question is whether the market structure taxonomy also organizes transaction prices: do trades in the same bond on the same day price systematically differently across venue types, dealer–client segments, and client sectors? Decomposing the within-bond-day dispersion of synchronized implied yields, we find that it does not. A naive variance decomposition attributes a sizable share of the dispersion to these groupings, but that share is almost entirely a counting artifact: a placebo that randomly reassigns each bond–day's trades to the *same number* of groups reproduces nearly all of it. The genuine, incremental contribution of the real groupings over this same-cardinality null is at most a few percentage points—and is negative for the dealer–client segment, whose observed share falls below the random benchmark. The groupings do not predict dispersion out of sample, and the largest systematic within-bond-day yield gap between any two venue, segment, or sector groups is around 0.1 bp.

[22] A volume-weighted bond-day yield is pulled toward the flow day's large trades and their price concession (Figure 15), so part of the regressor sits in the outcome: a sector's net purchases predict the flow day's gap between the volume-weighted and the unweighted yield for investment funds, hedge funds, and the official sector, and not the gap ten days later. On the volume-weighted outcome, the investment fund estimate is -0.06 bps ($t = -5.1$), a yield rise after their purchases, while hedge funds and the official sector keep 0.02 and 0.01 bps. Measured from the day before the flow, the volume-weighted estimates are 0.03, 0.04, and 0.06 bps for investment funds, hedge funds, and the official sector.

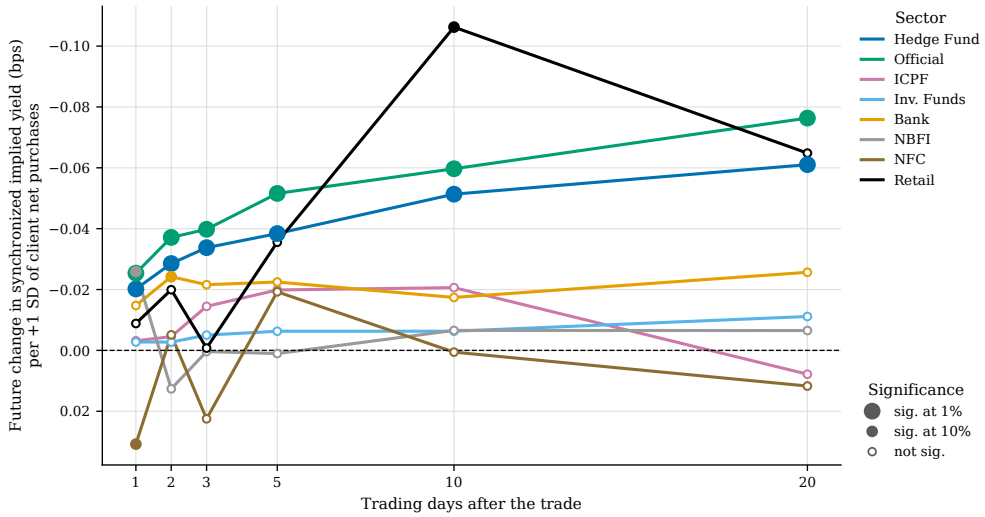


Figure A5: Permanent price impact using future synchronized implied yields. For each identified client sector, the figure reports the cumulative future change in a bond's synchronized implied yield, averaged without weights over the bond-day's trades, in bps, per net purchase of the bond equal to one standard deviation of the sector's absolute net purchases, by trading days after the trade. Sector-by-horizon regressions absorb ISIN and maturity bucket×day fixed effects and control for the bond's lagged curve-relative pricing error and its two lagged daily changes; the y-axis is inverted so that upward denotes richening through a lower future yield. Marker size encodes significance (ISIN-clustered standard errors): large filled = 1%, medium filled = 10%, small hollow = not significant. Dealer-to-client trades, MiFID II subsample (2018–2025).

For trades synchronized using the bond's own pre-trade and 11:45 MTS mids, the synchronized price equals the common fixing mid plus the transaction's deviation from its pre-trade mid. Within such a bond-day, price dispersion therefore reflects variation in transaction concessions. The yield transformation and observations synchronized using fallback quotes qualify this identity for the full sample. The dispersion exercise nevertheless concerns differences among trades within a bond-day; it does not identify whether the bond's common price level is close to fundamental value.

We stress what this is *not*. It is not a claim that trading is frictionless—the spread component is real and falls unevenly across investors (Figure 14). Nor is it a statement about futures anchoring. Bund futures plausibly anchor the *level* of cash prices across each maturity neighborhood through cash–futures basis arbitrage, but anchoring concerns where the reference mid sits relative to fundamental value, whereas within-bond dispersion concerns variation among transactions around their daily center—a different moment, silent on the level. Empirically, dispersion is essentially flat across the maturity curve: it is no tighter at the two-, five-, ten-, and thirty-year deliverable segments than in the gap between the ten- and thirty-year futures, which trades about half as often—so the futures footprint shows up in trading *activity*, not in execution cost uniformity. Testing the anchor itself requires the cash–futures basis, a separate question that this exercise does not take up; nor is it an independent test of the law of one price across venues.

A.5 Venue Centralization: Diagnostics and Sensitivity

This appendix collects diagnostics for the centralization counterfactual of Section 6.2: the venue mapping behind the model's two trading categories, two checks of its relationship-pricing premises, and the sensitivity of the welfare estimate to design choices.

Venue mapping Table A2 reports the composition of the structural estimation sample across the venue groups the model retains: bilateral trading pools OTC and single-dealer SI trades, and platform trading pools the Bloomberg and Tradeweb multilateral platforms, each of which spans several MIC codes of the same parent. Venue types the two-category model does not represent—most importantly OTF trading, which is close to an interdealer broker market, and the retail exchange—are excluded from the estimation sample. As in Allen and Wittwer (2023), quotes, values, markups, and the loyalty benefit are in yield basis points; a dealer's weekly platform quote is the average of its platform transaction yields across bonds after removing ISIN-by-week fixed effects.

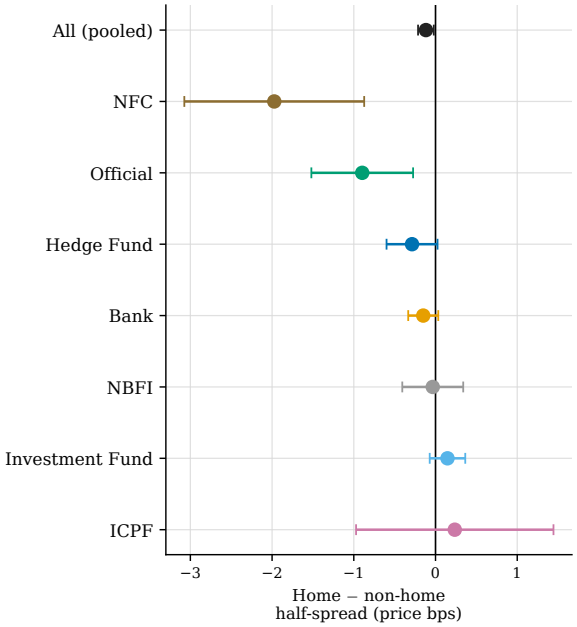
Venue group	Model	Trades	Share	Vol. (EUR bn)	Share	Customers
Panel A: all dealer-to-client trades						
OTC	Bilateral	271,058	26.4%	465.8	11.1%	6,416
SI	Bilateral	88,869	8.7%	928.1	22.2%	1,770
Bloomberg MTFs	Platform	288,186	28.1%	772.0	18.5%	1,941
Tradeweb MTFs	Platform	378,270	36.9%	2,014.6	48.2%	1,336
Bilateral total		359,927	35.1%	1,393.9	33.3%	7,587
Platform total		666,456	64.9%	2,786.6	66.7%	2,827
Panel B: institutional customers						
OTC	Bilateral	218,678	22.8%	387.2	10.2%	923
SI	Bilateral	71,932	7.5%	618.6	16.3%	671
Bloomberg MTFs	Platform	288,186	30.1%	772.0	20.4%	1,941
Tradeweb MTFs	Platform	378,270	39.5%	2,014.6	53.1%	1,336
Bilateral total		290,610	30.4%	1,005.8	26.5%	1,152
Platform total		666,456	69.6%	2,786.6	73.5%	2,827

Table A2: Venue mapping for the platform centralization model. The table reports the composition of the dealer-to-client structural estimation sample before collapsing venue groups into the model's two categories. Bilateral trading combines OTC and SI trades. Platform trading combines MTF trades on Bloomberg and Tradeweb. The Bloomberg and Tradeweb MTF groups each cover multiple MIC codes belonging to the same parent company. Institutional clients are those that trade on a platform at least once during the estimation window.

Matched home dealer half-spread differential The centralization model treats home dealer loyalty as a non-price relationship value rather than a separate bilateral price discount. To check that premise, we compare bilateral trades with a client's home dealer to the same client's other bilateral trades on the Allen and Wittwer (2023) model sample, using a leave-one-out home definition and matching on bond×day, trade size, and side. The model sample differential is small but statistically distinct from zero (-0.12 price bps; 95% confidence interval $[-0.21, -0.02]$), and the canonical investment fund investors show essentially no home discount (+0.20 bps). The full dealer cross-section gives a more precise but still small discount (-0.16 price bps), so the evidence supports treating relationship value as primarily non-price.

Affiliated-fund execution with the parent dealer The most captive clients provide a direct test of relationship pricing: dealers' asset management and insurance arms trading with

(a) Matched home dealer differential, by client sector.



(b) Raw and matched estimates.

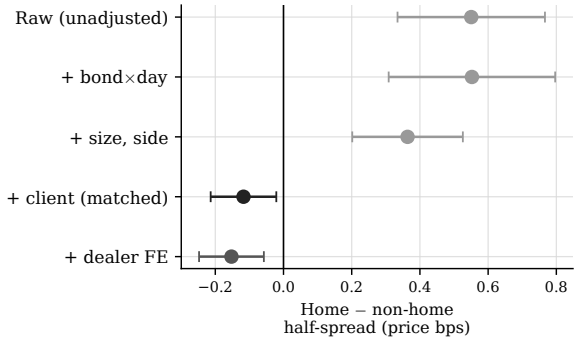


Figure A6: Matched home dealer half-spread differential. The left panel plots the within-client effective half-spread differential between a client's home (largest-volume) dealer and its other dealers, matched on bond×day, trade size, and side, pooled and by client sector (95% ISIN-clustered confidence intervals; negative values mean lower effective half-spreads with the home dealer). The right panel shows the same differential as controls are added, from the unadjusted comparison to the matched within-client specification. The effective half-spread is the client-signed deviation of the transaction price from the pre-trade MTS mid, in price basis points (negative = inside the mid). Bilateral dealer-to-client trades on the Allen and Wittwer (2023) model diagnostic sample (2023–2025, platform-using institutional clients, 545 multi-dealer clients).

their parent. These arms belong to their parent's investor group but not to its dealer desk, so an in-house trade is a dealer-to-client trade in which the client's investor group equals the dealer's. Within a given dealer, bond, and day, these affiliated funds pay their parent 0.52 price bps more on the effective half-spread than the dealer's other fund clients on OTC trades ($t = 1.4$; 95% confidence interval $[-0.23, +1.27]$), and no premium distinguishable from zero on the electronic venues (MTF and SI; Figure A7). The estimate draws on 21 parent groups; dropping any one of them gives estimates between -0.12 and 0.66 price bps. From the client side the gap is similar: within a given fund, bond, and day, the fund pays 0.86 price bps more when it routes to its parent than to an outside dealer ($t = 2.7$). The premium is imprecisely estimated, but the interval excludes any material discount: affiliated funds do not buy more cheaply from their own parent.[23] Bai et al. (2026) document an affiliation premium of 4–5 basis points inside bank holding companies in the Treasury repo market; in the Bund cash market the estimate is far smaller and confined to the opaque bilateral channel. Bilateral relationships do not, on average, give affiliated clients a better price, and the most captive client, if anything, pays more, not less.

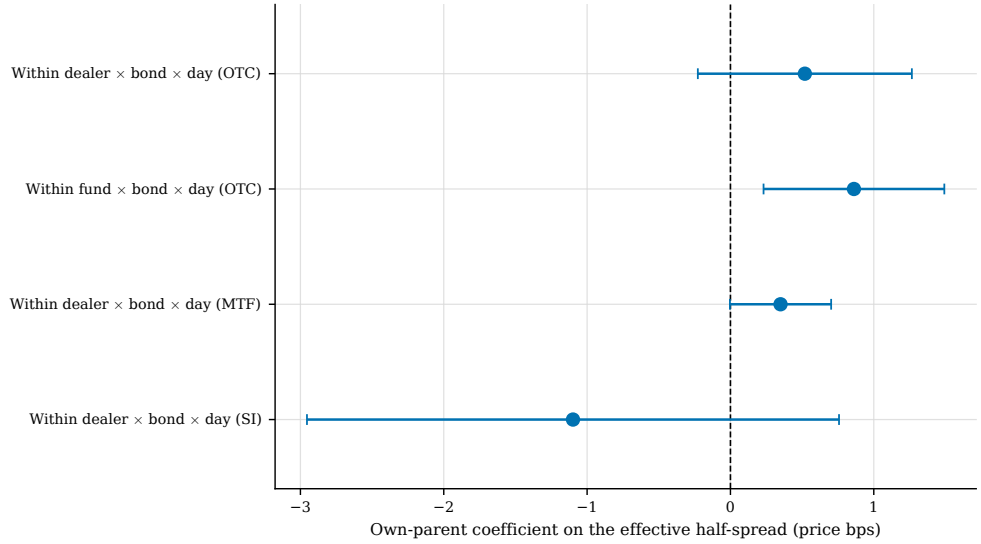


Figure A7: Own-parent coefficient on the effective half-spread for dealer-affiliated funds, by venue and specification. Each marker is the coefficient on an indicator for an affiliated fund trading with its parent dealer, from a regression of the client-signed effective half-spread, in price basis points, on that indicator and log trade size with the fixed effects listed; whiskers are 95% confidence intervals from standard errors clustered by ISIN. The within dealer specification compares an affiliated fund's trades with its parent to the parent's other fund clients on the same bond and day; the within fund specification compares the same fund's trades with its parent to its trades with other dealers on the same bond and day. Dealer-to-client trades of investment funds and insurers affiliated with a primary dealer, MiFID II (2018–2025).

Welfare sensitivity The welfare results in Table 4 hold the demand elasticity parameter at the imported value $1/\hat{\sigma} = 1.765$; Appendix A.6 documents why the Bund data do not identify this parameter. Table A3 reports the Central counterfactual under the "With loyalty" and imperfect competition specification, varying one design choice at a time: the dealer set, the dealer weights, the reference dealer, and the demand elasticity. The elasticity rows use the estimates of Allen and Wittwer (2023) for alternative specifications (0.68 and 1.44), the imported

[23] We read the estimate as descriptive. A positive wedge is consistent with the internal transfer price of Bai, Bostrom, Infante, and Ivashina (2026), but where an affiliated arm manages third-party assets the same wedge is instead a cost borne by outside investors; our data do not separate proprietary from third-party money, so we take no stand on which.

value, the upper end of its 95% confidence interval, and twice the imported value. Bisection places the reversal from loss to gain between $1/\hat{\sigma} = 2.94$ and 3.09, at or just above the upper end of the imported estimate's 95% confidence interval (standard error 0.599). Across the entire interval, centralization therefore lowers welfare relative to a fully bilateral market; its magnitude depends on the imported elasticity. These exercises do not evaluate centralization against the current access regime.

The loss is present in every dealer set, weighting and reference dealer we can estimate, but its size varies. With each of the four largest dealers as reference, the centralization effect at the imported elasticity lies between −13.28% and −1.61%. Trade size does not enter the model: welfare is a per-trade average, and every trade is priced at its dealer's weekly platform quote, so a trade size restriction changes only the estimation sample. Restricting the sample to trades of at most €50 m, the range over which the platform share of trades does not fall with size and every dealer in the sample trades on the platform, gives a Central effect of −14.15%; a cap of €100 m gives −13.08%. The loss therefore does not rest on the largest bilateral trades, for which the platform shows few comparable trades. Dealer sets of 30 or more group-level dealers do not identify the counterfactual: the added dealers appear in the transaction data, but for many weeks we do not observe the electronic platform price at which they compete, and full centralization would require imputing those prices.

Dimension	Specification	Central welfare (%)
Dealer set size	Top 10 dealers, equal weights	−20.80
	Top 15 dealers, equal weights	−19.38
	Top 20 dealers, equal weights (benchmark)	−13.28
	Top 25 dealers, equal weights	−32.01
	Top 30 or more dealers	Not identified
Dealer weighting	Top 20 dealers, volume weights	−12.09
Reference dealer	Largest dealer (benchmark)	−13.28
	Second-largest dealer	−13.16
	Third-largest dealer	−10.26
	Fourth-largest dealer	−1.61
Trade size	Trades up to €50 m	−14.15
	Trades up to €100 m	−13.08
Demand elasticity	$1/\hat{\sigma} = 0.68$ (no-loyalty estimate)	−20.44
	$1/\hat{\sigma} = 1.44$ (trade size estimate)	−16.90
	$1/\hat{\sigma} = 1.765$ (benchmark, imported)	−13.28
	$1/\hat{\sigma} = 2.94$ (upper 95% CI bound)	−0.10
	$1/\hat{\sigma} = 3.53$ (stress, 2× benchmark)	+6.09

Table A3: Sensitivity of platform centralization welfare to one design choice at a time. Entries are percentage welfare changes relative to a fully bilateral market for the Central counterfactual under the "With loyalty" and imperfect competition specification. The benchmark uses the top 20 group-level continuously active Bietergruppe dealers, excludes in-house trades, ranks dealers by dealer-to-client volume, applies equal dealer weights, normalizes the quality of the largest dealer to zero, and imports $1/\hat{\sigma} = 1.765$. Reference dealer rows instead normalize the quality of the second-, third- or fourth-largest dealer. Trade size rows drop the trades above the stated size from the estimation sample and re-estimate the model; the dealer numbering and each investor's platform access stay as in the benchmark.

A.6 Identification of the Demand Elasticity

Allen and Wittwer (2023) identify the demand elasticity from primary auctions: a dealer who wins an unexpectedly large share of the amount she demanded carries more inventory, and the induced variation in her platform quotes is plausibly unrelated to platform demand shocks. Their instrument is the share won of the amount demanded per dealer and week, with a first-stage F-statistic of 16; their appendix also validates a coarser variant that uses the amount won alone.

The Bund data contain the allotment side of this design but not the bids. MiFID II reports the primary allocation legs (Appendix D.6), which we aggregate to amounts won per dealer group and auction. The resulting allotments are well measured: ranked within half-years, they match the official Bietergruppe ranking with a Spearman rank correlation of 0.84, and the fit improves further when allotments are duration weighted, mirroring the weighted allotment criterion of the official ranking. Without the bids, the preferred share-won-of-demand instrument cannot be constructed; the amount-won variant and a surprise variant, which measures wins relative to the expectation implied by the previous half-year ranking, can.

These instruments have no power in Bunds. The win variation is real and survives dealer and week fixed effects, yet it is essentially uncorrelated with dealer platform quotes: first-stage F-statistics stay below 1.5 in the aggregate specification, with duration-weighted allotments, and in maturity-segment designs that match a win to quotes in the auctioned segment. The uninstrumented estimate is attenuated toward zero, as Allen and Wittwer (2023) also document in Canadian data. Measurement explains only part of the gap to the Canadian first stage: platform quotes are weekly averages of transaction yields whose cross-dealer dispersion is below one basis point, with a split-half reliability of roughly 0.25, so even a Canadian-strength pass-through would produce a first-stage F-statistic of only about five here. The observed pass-through is an order of magnitude below that benchmark.

The remaining gap is plausibly market structure. Bund dealers hedge cash inventory in the deep Bund futures market (Kerssenfischer and Helmus, 2025), so an auction allotment need not be worked off through cash quotes; the Canadian government bond market has no comparably deep futures segment. We do not test this mechanism here: the pass-through is absent in every dealer size tier, so no subgroup of dealers identifies σ in the Bund data.

We therefore import the Canadian point estimate and treat the demand elasticity as an identification limit of the Bund setting: Table A3 reports how the centralization welfare estimate depends on it. Estimating a Bund-specific elasticity would require dealer-level bid schedules from the auction system, which the transaction reports do not contain.

Appendix B Price Synchronization and Yield Curve Estimation

Price synchronization At the trade time and the 11:45 a.m. fixing, we select the most recent same-bond MTS quote within the preceding hour, preferring quotes with a bid–ask spread no greater than 10 yield basis points. We choose 11:45 a.m. to align with the Bundesbank yield curve fixing and a relatively quiet news period; see Kerssenfischer and Schmeling (2024). Where comparable bonds are available, we reject shifts that differ by more than 10 yield basis points from the median shift of up to four nearby maturities, requiring at least two. We also reject the bond-day when the median residual across at least three five-minute bins exceeds 10 basis points in absolute value. When a required quote is unavailable, we use a subsequent quote within an hour, subject to the observed trading session for quotes after the trade, or the median yield shift of nearby quoted bonds. These fallback observations are flagged separately and do not undergo the neighbor-shift check. Subsequent quotes and neighboring bonds are not substituted for the pre-trade same-bond benchmark in effective half-spreads.

Yield conversion and sample selection We convert synchronized prices into yields using the settlement cycle, TARGET holidays, irregular first coupons, and unadjusted coupon dates under the ICMA street convention. For older bonds without published first coupon and accrual start dates, we apply the federal convention for long first coupons: accrual starts on the settlement date, or on the preceding maturity anniversary when the bond settles within about two months after it, and the first coupon falls on the first maturity anniversary at least roughly 300 days after accrual starts. The yield sample excludes when-issued trades, bonds with less than three months of remaining maturity, and bonds with at most one year of original maturity. We screen individual prices against the bond's own daily yield center, using a tolerance of 25 yield basis points or the yield equivalent of 0.10 price points, whichever is larger. The center is the median synchronized yield across at least three distinct timestamp–price observations. With fewer observations, it is the yield of the bond's 11:45 MTS mid-quote. A tight 11:45 quote also replaces the median when the two disagree by more than the tolerance. The screen applies where a center is available; we retain a separate flag recording whether the center agrees with the quote. Observations without a synchronized implied yield or a Bundesbank curve yield cannot enter the yield sample. A separate gross-error guard excludes yields at least 100 basis points from the Bundesbank's Svensson (1994) curve. For spline and kernel ridge estimation, we use bond-days whose volume-weighted yield is within 35 basis points of that curve, while evaluating the fitted curves at all retained bond-days.

Yield curve estimation 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 semiannual coupons, and coupon schedules are often irregular, especially for earlier issues and bonds with long first or last coupon periods. We perform the estimation day by day using the volume-weighted average synchronized price of each bond, with pricing errors weighted by the inverse squared product of the bond's duration and dirty price, so that the fit approximately minimizes yield errors.

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 on the last trading day of each month, yielding 216 evaluation dates from January 2008 to December 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: $\alpha \in \{0.01, 0.02, \ldots, 0.10\}$, $\delta \in \{0, 10^{-4}, 5 \cdot 10^{-4}, 10^{-3}, 5 \cdot 10^{-3}, 0.01, 0.05, 0.1, 0.5, 1\}$, and $\lambda \in \{0.001, 0.01, 0.1, 0.5, 1, 5, 10, 50, 100\}$. 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 $(\alpha, \delta, \lambda) = (0.02, 0, 1.0)$, with a mean RMSE of €0.205 per €100 of face value. The criterion is flat near this optimum: the ten best specifications lie within 0.4% of its RMSE. Relative to the baseline parametrization $(\alpha, \delta, \lambda) = (0.05, 0, 1.0)$ that Filipović et al. (2025) use for U.S. Treasuries, the mean RMSE is 1.4% lower, a modest improvement that holds on 59% of the evaluation dates.

Appendix C 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 compares execution volumes, because a MiFID I record is one execution while a MiFID II execution can span several records (Section D.8). MiFID II covers about 6.2 times as much annual trading volume in Bunds as the earlier MiFID I dataset. MiFID II covers at least as much volume as the survey-based dataset of the Finanzagentur, Germany's debt management office, to which each member of the auction group reports its own trades (Branger et al., 2024): over the 80 months both datasets cover in full, MiFID II execution volume is 1.04 times the Finanzagentur total. Because each member reports its own side, a trade between two members may enter the Finanzagentur total once or twice; if twice, the ratio understates the coverage of MiFID II. Figure A8 compares the monthly series. A key difference is that the FA dataset covers only trades involving a primary dealer. Unlike our dataset, it misses trades between European non-dealers but may cover trades between two non-European investors, e.g. a trade between a U.S. branch of a German dealer bank and a U.S. investment fund.

Figure A9 shows an abrupt change in the relative shares of venue types in 2018. Some of this reflects transactions classified as OTC under MiFID I receiving more specific venue labels under MiFID II. The large increase in aggregate volumes, however, suggests that the primary driver is the expanded reporting of previously unobserved MTF, OTF, and SI transactions.

Figure A10 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. In the 2008–2017 sample, there is also a large spike shortly before noon.

Figure A11 shows the wide dispersion of transaction sizes. About 6% of MiFID II executions are smaller than €1,000, while 11% are €10 million or larger (MiFID I: 2% and 10%).

Figure A12 compares daily Bund yield curve fitting errors with the U.S. Treasury noise measure of Hu, Pan, and Wang (2013). The Bund series are similar in magnitude and time-series variation to U.S. Treasuries.

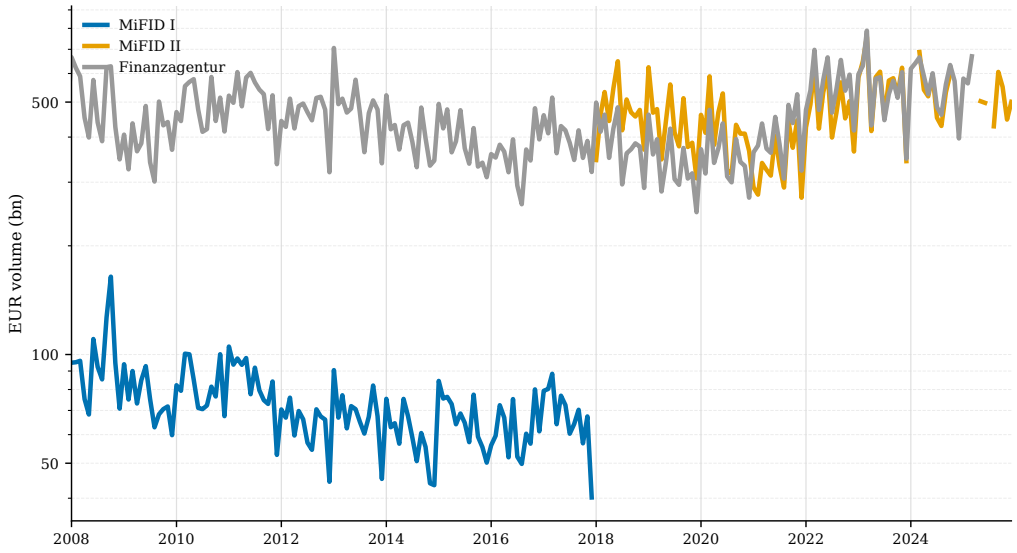


Figure A8: Monthly trading volume in MiFID and Finanzagentur data. MiFID volumes are execution volumes. MiFID II months that lose days without account owner information are left out (Section D.2).

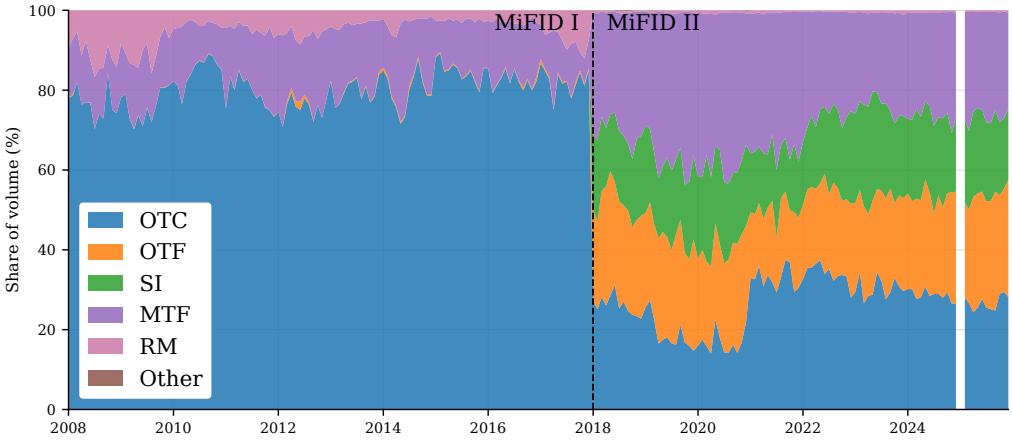


Figure A9: Volume shares across venue types. Shares of execution volume; an execution is assigned the venue type of its most specific venue record.

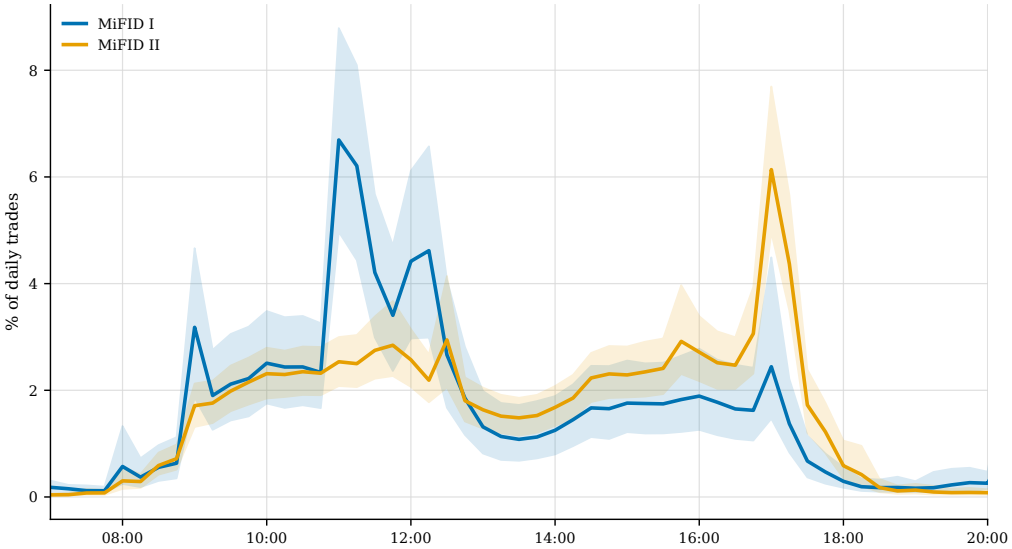


Figure A10: Number of trades across the day. Median and interquartile range across days of each 15-minute interval's share of the day's executions.

Additional market structure descriptives Securities holdings statistics already show who *holds* Bunds; what they cannot show is who *trades* them, with whom, and from where. A non-anonymous dataset linked to the GLEIF register reveals where the counterparties that trade Bunds are domiciled—a first-order question for a sovereign issuer.[24] We map each counterparty to its country of domicile and group countries into Germany, the rest of the euro area, the United Kingdom, the United States, and the rest of the world.[25] Figure A13 shows that Bund trading is predominantly cross-border: domestic German counterparties account for only 10% of trade-leg volume, and only 15% of client-side volume once primary dealers and other intermediaries are

[24]The sample covers reportable EU MiFIR transactions. U.K. and other non-EU venue activity is undercounted, especially after the reporting perimeter narrowed in 2021 (Figure A19). Domicile comes from the investor group's ultimate parent, but country shares still depend on this coverage.

[25]Country is assigned at the investor-group level from GLEIF, i.e. the domicile of the group's ultimate parent; a region is mapped for 83% of trade-leg volume. The unit is the trade leg (buyer and seller side), so a cross-border trade contributes to both regions.

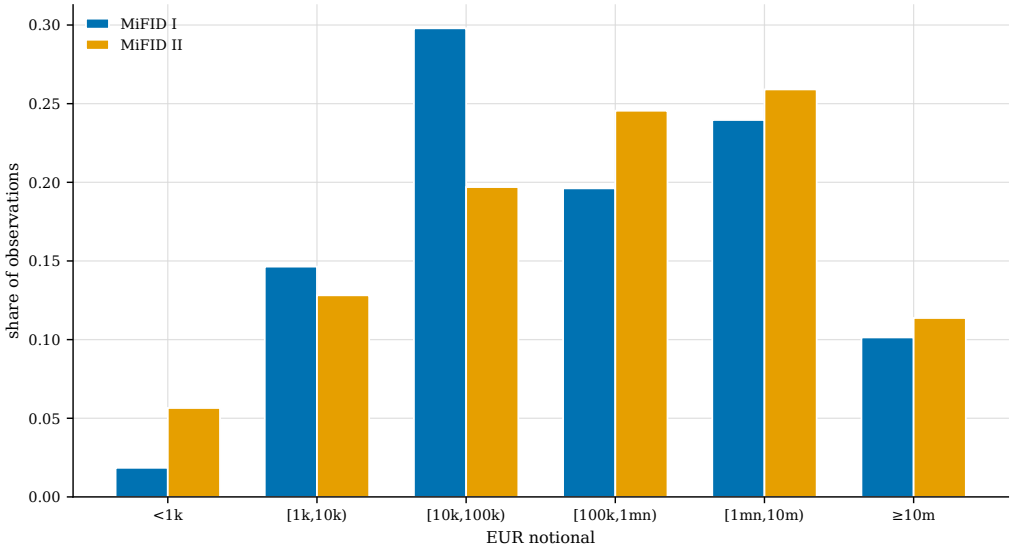


Figure A11: Distribution of transaction sizes. Share of executions by execution size in euros.

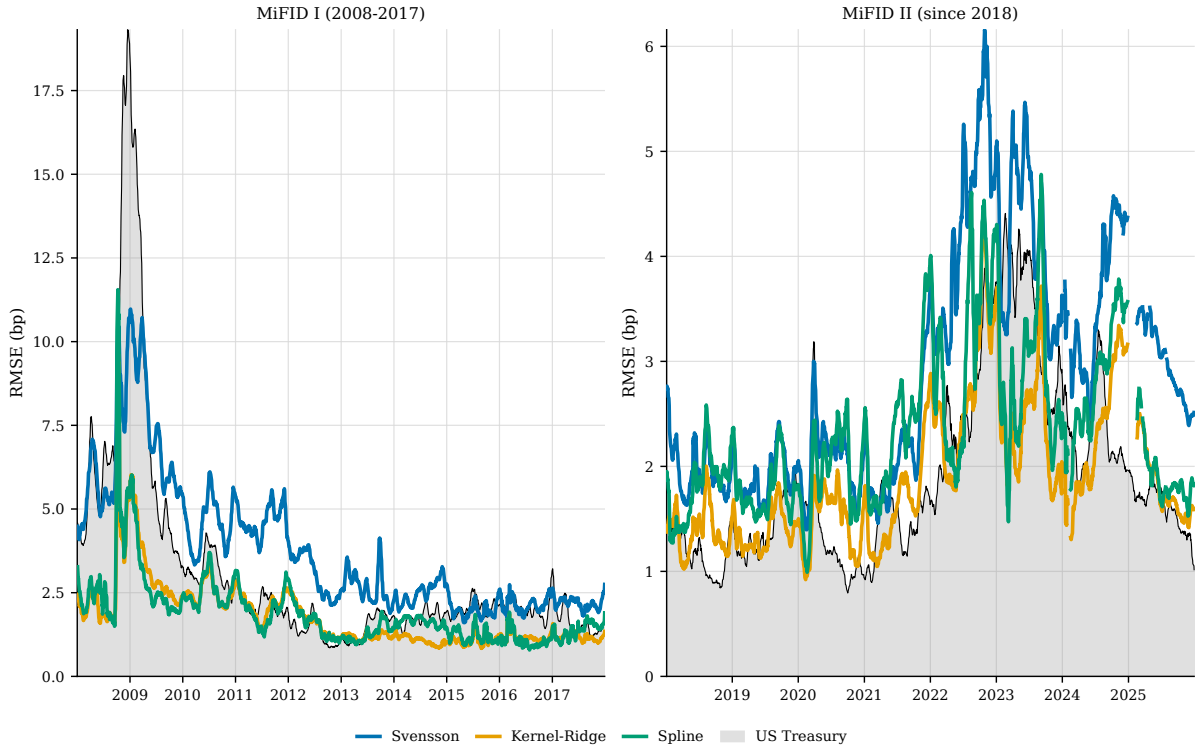


Figure A12: Bund yield curve fitting error over time. Root mean squared error (RMSE) between volume-weighted implied yields and three fitted curves: the kernel ridge curve of Filipović et al. (2025), the Svensson (1994) curve published by the Bundesbank, and a smoothed cubic spline. For each trading day, the RMSE is computed across all Bunds with at least three months to maturity and with observed cash market transactions. Section 3.5 provides details. For reference, the gray shaded area shows the U.S. Treasury noise measure of Hu et al. (2013) based on end-of-day quotes from CRSP, updated by the authors. They also fit yield curves via Svensson (1994), 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.

excluded. The two views differ tellingly: excluding intermediaries, the United States share falls and the rest-of-world share rises sharply, so part of the headline United States footprint is dealer intermediation, not end-investor activity—a distinction between trading and ultimate ownership that transaction data can draw and holdings statistics cannot. Figure A13 and Figure A14 report the full regional split and the region-by-sector composition.

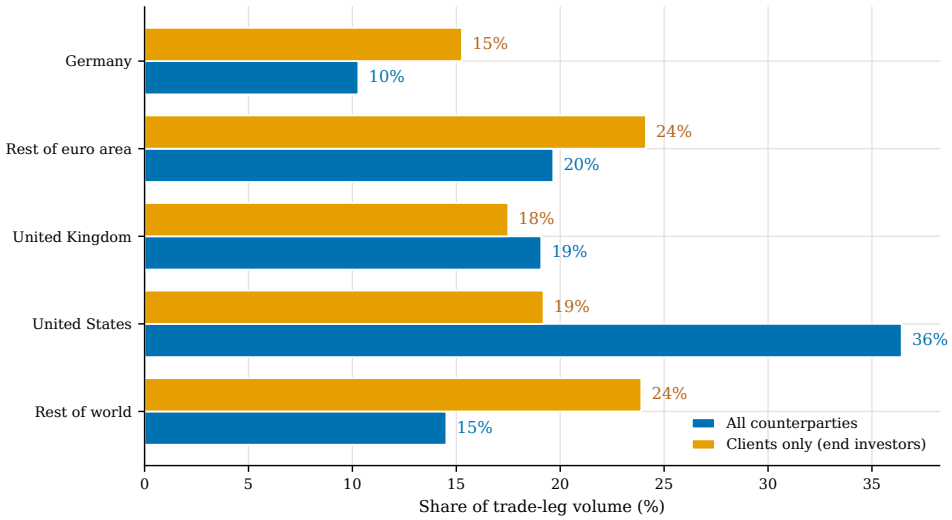


Figure A13: Cross-border participation. Share of trade-leg volume by counterparty region. Dark bars use all counterparties; light bars restrict to client legs, excluding primary dealers and other intermediaries to approximate end-investor geography. MiFID II subsample, investor-group (ultimate-parent) level.

In the all-leg view of Figure A13, the rest of the euro area accounts for 20% of trade-leg volume, the United Kingdom for 19%, and the United States for 36%; in total, 90% of activity involves a non-domestic counterparty. Restricting to client legs, 24% of client-side volume comes from the rest of the euro area, 18% from the United Kingdom, 19% from the United States, and 24% from the rest of the world—85% foreign in total.

Figure A14 reports how the end-investor sector mix varies across counterparty regions. Foreign official sector institutions and banks dominate some non-euro-area participation, while hedge fund activity is more concentrated in the United Kingdom and United States and euro area participation is more heavily weighted toward investment funds.

Dealer breadth and flow informativeness Is a client's dealer breadth a proxy for how informed its flow is? Figure A15 bins institutional clients by their number of distinct dealer counterparties and estimates, for each bin, the permanent price impact of the bin's net flow, exactly as in the sector-level test of Figure 13. In the cross-section the answer appears to be yes: single-dealer clients' flow carries no signal, while flow from the best-connected clients predicts permanent richening (Figure A15a). The gradient is mostly sector composition. The best-connected bins load increasingly on hedge funds and official sector accounts, and once these two sectors are excluded the gradient disappears except in the most connected bin, which keeps a smaller but significant impact (Figure A15b); within the remaining sectors, clients with fewer than ten dealers carry no detectable signal. In the Bund market, cross-sectional connectedness thus proxies for informedness mainly through sector composition: the informed sectors are the broadly connected ones. This complements Kondor and Pintér (2022), who propose a client's dealer connections as a proxy for private information but identify the effect from *time variation* within a given client, absorbing client fixed effects, not comparing connectedness levels across clients. Our cross-sectional result is consistent with that design choice, and it cautions against

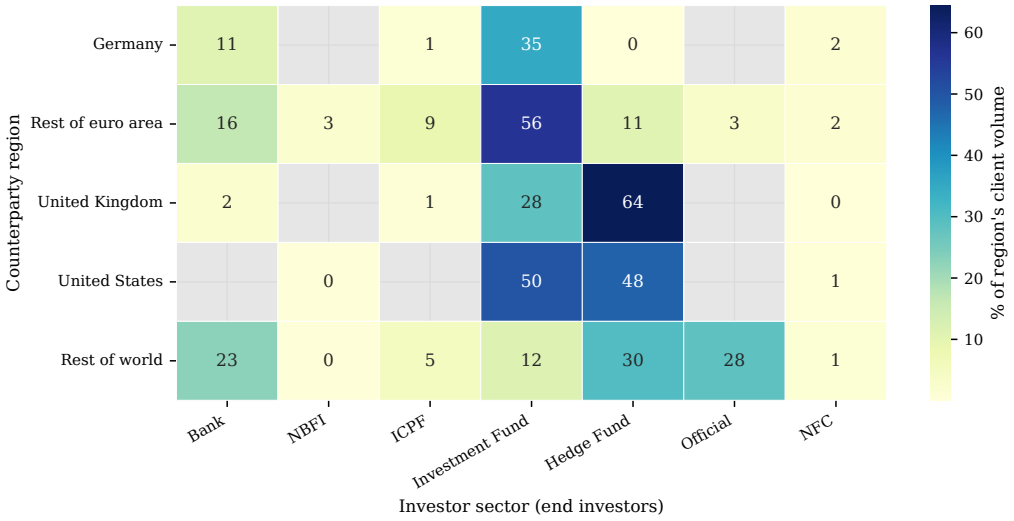


Figure A14: End-investor composition by region. Each cell shows a sector's share of the region's total client-side trade-leg volume (in percent). Intermediary sectors (primary dealers, interdealer brokers) and retail—structurally absent as account owners or without a legal entity identifier—are excluded. Gray cells are omitted; visible cells therefore need not sum to 100% within a row. MiFID II subsample, investor-group level.

reading the level of a client's dealer breadth as a sophistication index without conditioning on sector. Implementing their within-client design directly—regressing a client group's monthly trade performance on the number of dealers it faces that month, with client and month fixed effects, estimated separately before and after the 2021 reporting-perimeter break—we find no within-client effect either: the connections coefficient is small at every horizon and not distinguishable from zero at the five-day horizon (0.004 bps per dealer, $t = 0.9$), precisely enough estimated to rule out an effect of the magnitude documented for Gilts, and equally absent among the 159 largest client groups that account for 80% of institutional volume. In the Bund market, dealer breadth thus indicates who a client is, not how informed its trading currently is.

Whether post-crisis regulation has eroded dealers' willingness to warehouse risk is a central but hard-to-test question, because it requires observing dealers' trading capacity (He and Krishnamurthy, 2013; Duffie et al., 2023). The MiFIR trading capacity field lets us see it. Figure A16 decomposes dealers' volume with identified clients into principal (the dealer trades on own account and warehouses risk), matched principal (a riskless, simultaneous offset), and agency. The field describes the firm that files the record, so we use the records that dealers report themselves (85% of dealer-to-client record volume); on the others it describes the client. Principal trading dominates. Its share averages 97% in 2018–2020 and 93% from 2021 on, with a dip in 2022 and 2023, when the agency share peaks; it is 93% in 2025. The agency share rises from 2% to 5%, and matched principal never exceeds 3%. U.K. venues left the reporting perimeter in 2021, so we do not read the difference between the two periods as a change in dealer behavior. The shares describe how dealers book client trades, not the size of their inventories. They exclude records whose non-dealer side is empty, where the dealer faces an anonymous venue, a CCP, or an IDB (41% of dealer-to-client record volume): nearly all of the matched principal volume on these records belongs to executions that contain no dealer-to-client record with an identified client.

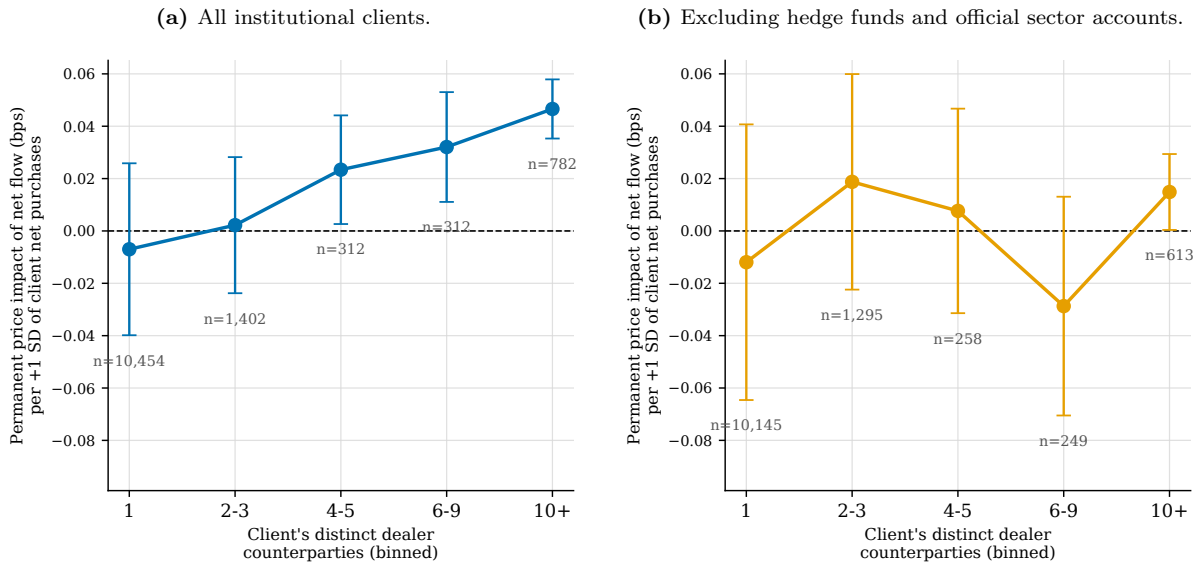


Figure A15: Permanent price impact of client net flow, by dealer breadth. Institutional client groups are binned by their number of distinct dealer counterparties. For each bin, the figure reports the ten-trading-day change in the curve-relative pricing error, defined as synchronized implied yield minus the kernel ridge fair yield, in bps, per net purchase of the bond equal to one standard deviation of the bin's absolute net purchases, signed so that upward denotes richening. Bin-level regressions absorb ISIN and maturity bucket×day fixed effects and control for the bond's lagged curve-relative pricing error and its two lagged daily changes; error bars are 95% ISIN-clustered confidence intervals. The left panel uses all institutional clients; the right panel excludes hedge funds and official sector accounts, on the same y-axis. Text labels report the number of client groups in each bin. Institutional client groups (valid legal entity identifier; retail clients and clients in the primary dealer sector excluded), dealer-to-client trades, MiFID II subsample.

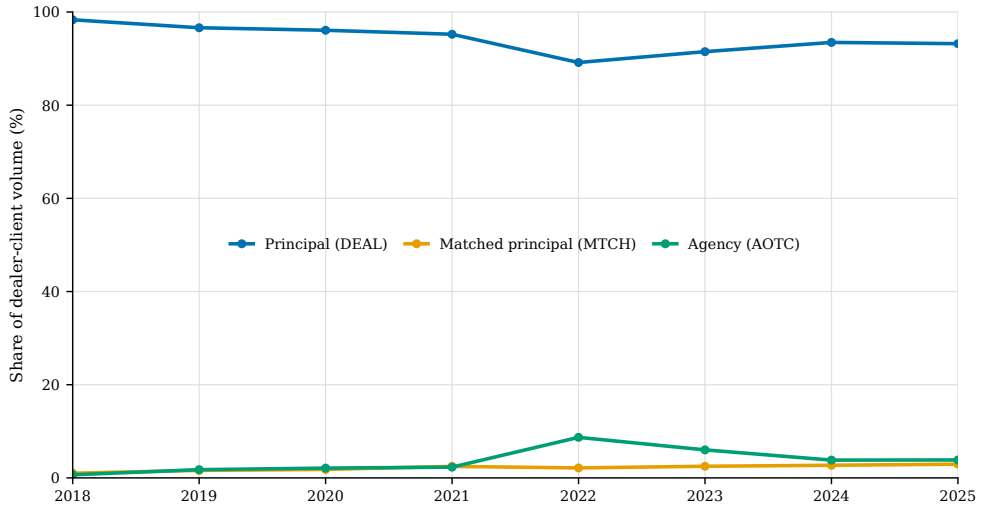


Figure A16: Dealer trading capacity over time. Annual share of dealer-to-client record volume traded in a principal (`DEAL`), matched principal (`MTCH`), or agency (`AOTC`) capacity, from the MiFIR trading capacity field. A record counts as principal only when it is reported as `DEAL` and the dealer is its buyer or seller. Dealer-to-client records reported by the dealer, with an identified client outside the primary dealer sector; MiFID II subsample.

C.1 Cash and Futures Trading Activity

To put the hypothetical scenario of centralizing the OTC market (Section 6.2) into perspective, Figure A17 compares cash market trading activity under the current fragmented structure with a counterfactual centralized market. The left panel shows trading volumes, the right panel trading frequency. For reference, both plots contain the equivalent activity measures for the four Bund futures.[26]

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.

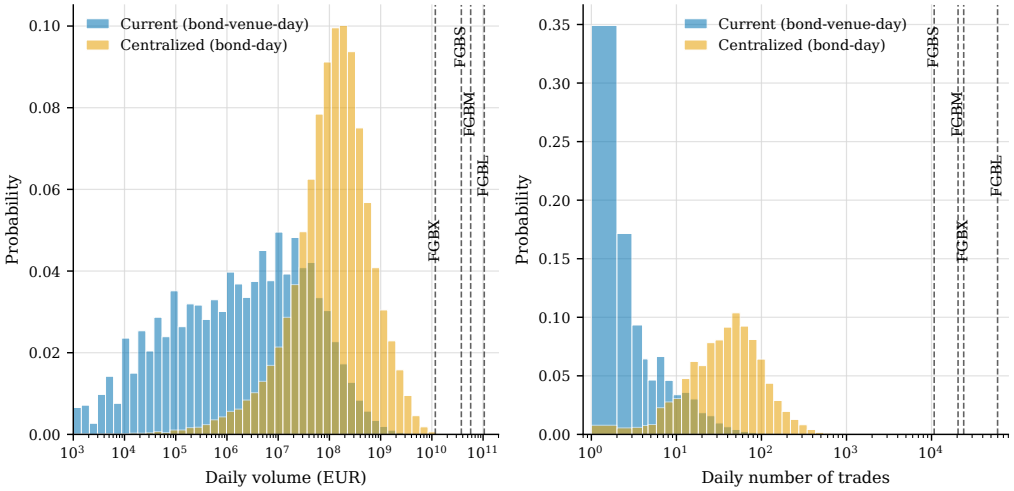


Figure A17: Trading activity in Bunds. The left histogram compares execution volumes in the current market setting (bond-venue-day cells) with a counterfactual fully centralized market (bond-day cells). The right histogram reports the number of executions instead of volume. In both panels, vertical dashed lines show the corresponding median values in the four major Bund futures.

C.2 Market Structure Dynamics

Because the dataset spans eighteen years and two reporting regimes, it can also document how the market's structure has *changed*. We highlight several dynamics, flagging where a shift reflects evolving data coverage, not economics.

Dealer concentration is cyclical, not secularly rising. Figure A18 plots the inverse-HHI effective number of dealers in dealer-intermediated volume. Within the full-coverage MiFID II window, concentration is far from monotone: the effective dealer count peaks at 16.8 in 2019, falls to 12.7 in 2023 during the inflation and rate-hike regime shift, and recovers to 16.3 by 2025. This is consistent with intermediation concentrating in the largest dealers when conditions tighten, not a secular decline in the number of active intermediaries.

[26] 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 rollover days each quarter.

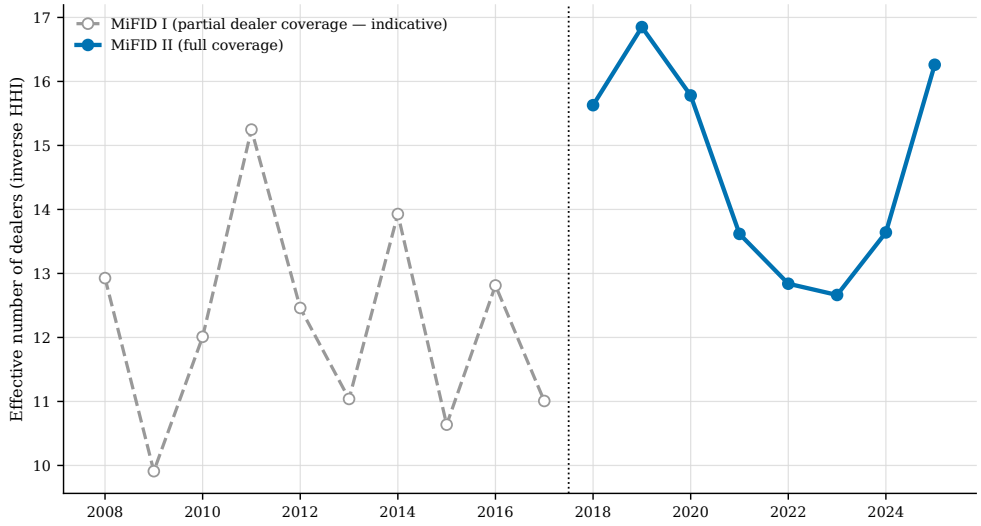


Figure A18: Effective number of dealers over time. Annual inverse-HHI effective number of primary dealers in intermediation volume, $1/\sum_j s_j^2$, where s_j is dealer j's share of primary dealer trading volume, counting each dealer's purchases and sales in all trades in which a primary dealer participates on either side. Lower values mean that intermediation is more concentrated among fewer dealers. The dotted vertical line marks the 2018 MiFID II data break. The MiFID I segment (2008–2017) is shown for completeness but is not directly comparable: that regime is dealer-side weighted and misses roughly eight active primary dealers—including five U.K.-reporting dealers absent throughout—which mechanically lowers the measured effective number.

The end of the Brexit transition reshaped venue coverage. Because the data separate *where* a trade executes from *who* trades, they reveal a sharp break at the end of the Brexit transition period (1 January 2021), when U.K. trading venues left the EU-MiFIR reporting perimeter. Figure A19 shows the share of Bund volume executed on U.K.-registered venues collapsing from about 64% before 2021 to 15% after. Much of this flow migrated to EU-registered venues, whose share rose from about 13% to 52% as operators such as Bloomberg, Tradeweb, and MarketAxess opened EU platforms. But the migration is incomplete: total on-venue volume fell from 79% to 69%, the remainder moving off-venue or out of view. The share of volume with a U.K.-*domiciled* counterparty fell far less (from about 24% to 16%), so the break was largely a relocation of trading *infrastructure* and a contraction of reporting scope, not an exit of U.K. participants. We therefore read the post-2021 venue composition as a change in coverage, not in economics, and treat Bund trading on U.K. and other non-EU venues as undercounted thereafter.

The return of the retail investor. The panel also captures a sharp swing in household participation. Figure A20 plots retail investors' share of Bund trades and of volume over the MiFID II period. Trades here are trade records, and a bank that aggregates household orders reports one record per client allocation. After ECB policy rates turned positive in mid-2022, retail's share of trade records jumped from about 11% to a peak of roughly 49% in 2024—more than a fourfold increase—as households returned to positively yielding Bunds after a decade of near-zero rates. Without the allocations of aggregated orders the share also rises more than fourfold, from 4% to 19% in 2023. Yet retail's share of volume stayed negligible (around 0.2%): the surge is in the number of trades, not in euros, reflecting very small ticket sizes. Retail trading also leans on regulated exchanges—about 36% of retail trades outside aggregated orders execute there (11% of all retail records), far more than for investment funds or insurers—so the episode lifted exchange trade counts without materially affecting market-wide liquidity or pricing.

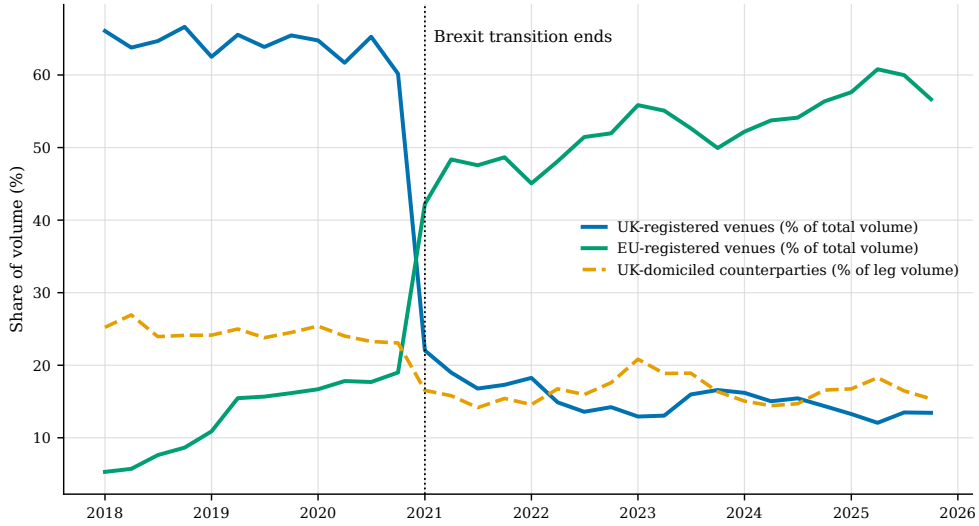


Figure A19: Venue and counterparty shares around the Brexit transition. Quarterly shares of execution volume on U.K.-registered and on EU-registered venues (MTF/OTF/SI/RM), and the share of trade-leg volume with a U.K.-domiciled counterparty (dashed). The dotted line marks the end of the Brexit transition period (1 January 2021). MiFID II subsample.

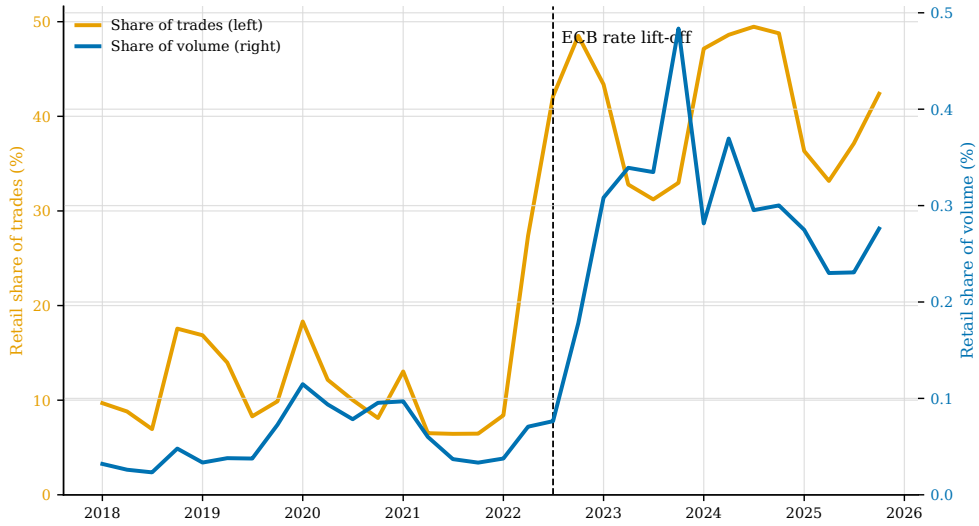


Figure A20: Retail participation over time. Quarterly share of Bund trade records, counting each client allocation (orange, left axis), and of trading volume (blue, right axis) with at least one retail (household) counterparty. The trade-count and volume series use separate axes. The dashed line marks the ECB's first policy-rate increase (July 2022). MiFID II subsample.

Algorithmic execution rose, but humans retook the wheel in stress. The MiFIR "execution within firm" field records whether each trade was executed by an algorithm or by a person, letting us track the machine's advance directly.[27] Figure A21 shows algorithmic execution climbing from roughly 25% of trades in 2018–19 to about 29% by 2025. Measured by volume the share is markedly lower—rising from about 10% to 19%—confirming that algorithms execute disproportionately small tickets while large blocks remain human-handled. The flag describes the firm that reported the record: among records filed by dealers the algorithmic share of trades moves from 42% to 46% over the sample, so the advance is concentrated in the records that clients file. The split by venue protocol family in Table A4 shows that algorithmic execution is concentrated in, but not confined to, order book venues: algorithms account for 55.1% of volume within the order book family, but also for 25.0% within the RFQ family and 10.7% within the residual other family. The advance is also not monotone. Around the March 2020 Covid dislocation the algo share fell from a 29% prior-six-month baseline to 24%, and around the September 2022 Gilt/LDI turmoil from 32% to 21%, as execution reverted to human traders when liquidity was thin; the March 2023 banking stress saw no retreat (27% against 23%). The pattern echoes the man-versus-machine evidence from Gilts (Guimaraes, Jurkatis, Pintér, and Wijnandts, 2025): algorithms have steadily gained share, but at the height of stress humans still step in.

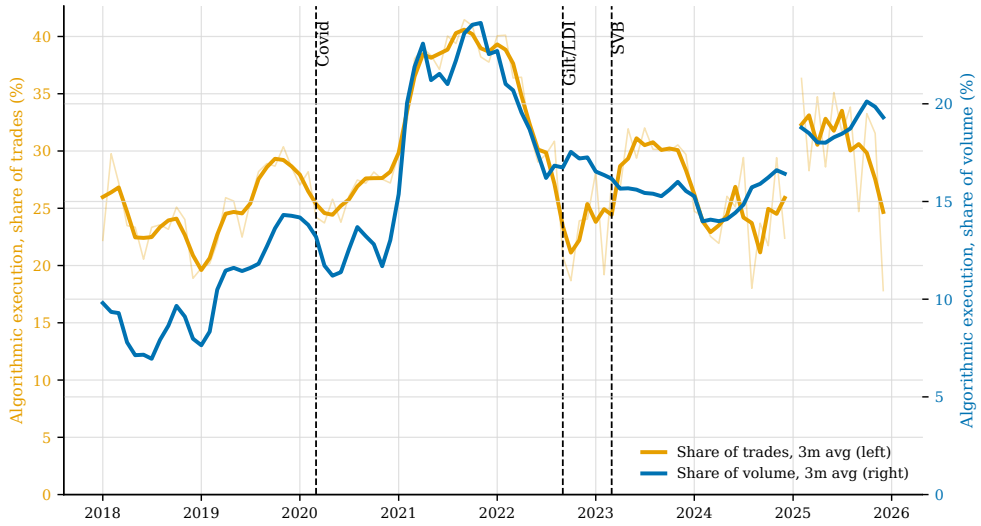


Figure A21: Algorithmic execution over time. Monthly share of Bund trades executed algorithmically (MiFIR "execution within firm" field), as a share of trade count (orange, left axis) and of volume (blue, right axis); both three-month centered averages, with the raw monthly count share in faint orange. Dashed lines mark three stress episodes. MiFID II subsample.

[27] The MiFIR "execution within firm" field (RTS 22) is populated on essentially every MiFID II trade record (100%) and distinguishes algorithmic (`ALGO`) from human (`PRSN`) or client-directed (`CLNT`) execution. Shares are computed on trade records as reported (Section D.8).

Venue protocol family	ALGO	PRSN	CLNT
RFQ family	25.0	72.2	2.7
Other (residual)	10.7	80.0	9.4
Order book (CLOB)	55.1	37.3	7.6
Unknown	11.0	83.7	5.3

Table A4: Execution type by venue protocol family. Each row reports execution-type shares of a venue protocol family's euro volume (row-normalized percentages) among trades with a populated MiFIR "execution within firm" field. MiFID II subsample.

Appendix D Data Details

Article 26 of Regulation (EU) No 600/2014 (MiFIR) requires investment firms to report complete and accurate details of each transaction to their national competent authority by the close of the following business day. For German sovereign bonds, the competent authority is BaFin irrespective of trading venue, so the data provide near-complete coverage of the Bund market. Because MiFIR was not designed for research use, the reports contain duplicated records, heterogeneous formats, and reporting errors. We follow the framework of Jurkatis (2024) for cleaning MiFID transaction data on corporate bonds, adapting it to the Bund market—in particular its far higher trading frequency: a single ISIN can generate up to 65,497 reports on a single day, making duplicate identification substantially more demanding.

Data availability Access to the non-anonymous MiFID I dataset can be requested via the Research Data and Service Centre (RDSC) at the Bundesbank and is subject to strict requirements. MiFID II/MiFIR transaction reports are not available to external researchers through the RDSC. We will maintain the final research dataset and its documentation as a reusable resource for Bundesbank researchers whose projects are authorized to use the underlying data. All other data are available directly from the data vendors.

Pipeline overview This Appendix describes the MiFID II cleaning pipeline in full. The MiFID I pipeline runs in parallel but differs in several respects. First, MiFID I records carry regime-specific administrative flags—regulatory exclusion codes, administrative correction markers, and name-to-follow placeholder legs—that are screened at the outset and have no MiFID II counterpart. Second, trading capacity is derived from proprietary reporting fields rather than an explicit DEAL/MTCH flag, so the report-splitting and aggregation steps described in Section D.3 do not apply. Third, deduplication applies the same cross-reporter record-linkage matching to proprietary (DEAL) pairs—though one-sided MiFID I reporting yields far fewer such duplicates—and adds intra-reporter matched principal pairing and venue-consistent cross-capacity matching. Fourth, investor identity is resolved via a BaFin-ID-to-LEI crosswalk, followed by the LEI enrichment. Both pipelines produce an identical output schema; all downstream processing is applied uniformly to the combined dataset. Figure 1 in the main text summarizes the construction: the raw reports from both regimes pass through the cleaning steps in this Appendix and the cleaned records are then enriched with bond-, investor-, venue-, and price-level reference data.

D.1 Reporting Requirements

MiFID I regime (2008–2017) Prior to MiFIR, transaction reporting in Germany was governed by MiFID I, implemented nationally via the Wertpapierhandelsgesetz (WpHG). Investment firms were required to report transactions in financial instruments to BaFin using BaFin-assigned institutional codes (*BaFin-IDs*) rather than standardized global identifiers such as LEIs. Reporting obligations were primarily one-sided: for a bilateral OTC trade, typically only the selling or executing firm submitted a report. This limits bilateral duplicate detection, which is a central step in the MiFID II cleaning pipeline; we adapt the deduplication logic accordingly. A second convention in MiFID I affects trade direction: the `g_art` buy/sell flag records the reporting party's own direction for proprietary trades but the client's direction for matched principal and agency client business. The pipeline therefore conditions on trading capacity and dealer status when constructing the buyer and seller account owner fields (`bao` and `sao`), swapping the two when the reporting party is the dealer side of a client trade. We resolve BaFin-IDs

to LEIs by combining a BaFin registry lookup, BIC-to-LEI resolution via RIAD, hand-curated assignments for major institutions, and exact and fuzzy name matching against the GLEIF register; no single source achieves full coverage. The resulting map covers more than 99% of MiFID I trade volume.

MiFID II/MiFIR regime (2018–2025) Under MiFIR, which took effect on 3 January 2018, each investment firm must independently report every transaction it executes, identified via LEI. Each report contains the instrument identifier, execution timestamp, price, quantity, venue, trading capacity, and counterparty identities. Because each firm reports its own transactions independently, a bilateral trade between two investment firms generates two reports of the same economic transaction, one from each side. A trade between an investment firm and a client that is not an investment firm, such as a fund, generates a single report. On an anonymous venue, each member reports its own side against the venue or its CCP, so one execution generates two one-sided reports. Trading venues report on behalf of unregulated participants. These asymmetries determine which trades appear once versus twice in the raw data and drive the duplicate-matching strategy in Section D.4.

D.2 Sample Selection

The first step extracts 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 exclude short-term Bubills, green bonds, inflation-linked bonds, and other non-standard bonds. To ensure consistency and remove irrelevant records, we also exclude transactions containing complex trade components, commodity derivatives, securities financing transactions, or changes in derivative notional. The MiFID II/MiFIR reports are as of 21–22 September 2026; the source keeps receiving new reports, corrections and cancellations, so a later extraction can differ.

When an instrument is bought or sold by a joint account, all account owners must be reported (see ESMA Guidelines, p. 42). Thus, one transaction can have multiple reported buyers or sellers. In 2019, 86.5% of the owners of joint accounts were reported as LEIs, 0.2% as natural persons, and the remaining 13.3% as internal accounts (INTC) or other identifier types.

Each account owner is assigned a "position", numbered from one to the number of owners of the account. To avoid double counting transactions involving joint accounts, we merge only the first owner (position = 1) into each transaction record. This preserves trade volumes but drops the other owners of a joint account.

In 2019, 0.06% of transactions involved a joint account, and merging all positions to the transaction data would increase the overall number of transaction records by 0.09%. Keeping only the first owner therefore loses almost no information.

The same rule applies to buyer and seller decision makers, although at a much smaller scale. In 2019, only 15% of transaction reports include a buyer or seller decision maker. As before, we include only the first decision maker for each transaction. We remove reports where the executing party is recorded as both the buyer and the seller.

On 64 trading days in 2024 and 2025, at least 20% of the MiFID II reports in our sample carry neither a buyer nor a seller account owner. On these days the account owner information is missing at the source; the data owner has confirmed these data gaps, and their cause is still being investigated. Without account owners, the counterparties of these reports cannot be identified, and each report enters the data as two records without a counterparty (Section D.3). We therefore drop these days from all analyses. They hold 14% of the 2024–2025 records and include every trading day from 2 January to 5 February 2025, so monthly series have no

observation for January 2025. Reports without account owners also occur on other days; they stay in the sample as records without a counterparty.

The raw extraction yields approximately 16.9 million MiFID II reports over 2018–2025 and 2.5 million MiFID I reports over 2008–2017. After the cleaning steps described in the remainder of this Appendix and without these days, these collapse to approximately 14.2 million trade records for 137 bonds and 2.2 million trade records for 146 bonds, respectively. The combined sample covers 16,350,195 trade records in 224 bonds spanning 2008–2025.

D.3 Normalization

The normalization step prepares the data for analysis by handling intermediary trades, aggregating reports, and identifying outliers in prices and quantities. These steps keep the dataset aligned with executed transactions rather than reporting artifacts. The statistics we report on these steps and on duplicate matching refer to the reports of 2019.

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 split these intermediary reports into two separate records—one where the executing entity is recorded as the buyer and another where it is recorded as the seller. This gives each transaction a uniform reporting structure.

Under MiFIR, firms are required to report the trading capacity of the transaction, that is, whether the firm was trading on its own account (`DEAL`), as matched principal (`MTCH`), or in 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 A5 shows how the trading capacity flag aligns with the splitting criteria in this subsection.

cat_trading_capacity	split	%
DEAL	0	64.29
AOTC	1	28.5
AOTC	0	3.34
MTCH	1	3.07
DEAL	1	0.68
MTCH	0	0.12

Table A5: Trading capacity field accuracy.

Table A5 shows that, for 95.9% of transactions, the trading capacity flag aligns with whether the report requires splitting. Where it does not match, we treat the trading capacity flag as misreported and adjust it while retaining the original value. The dealer trading capacity shares in Figure A16 use the reported flag instead: a record reported as `MTCH` or `AOTC` counts as matched principal or agency even when the reporting firm is its buyer or seller, and a split `DEAL` report counts as matched principal. Overall, 32.2% of reports are split.

Aggregating Reports

As Jurkatis (2024) points out, firms sometimes report multiple trades with the same counterparty at the exact same time and price. This makes the duplicate report from the counterparty firm difficult to identify, because the counterparty may report them as a single large transaction.

This problem partly arises from report splitting. Assume executing party A reports two matched principal transactions with quantities q_1 and q_2 that it intermediates between executing party B and clients C1 and C2. At the same time, executing party B reports only a single transaction with quantity $q_1 + q_2$ between itself and executing party A. Any duplicate-matching method will struggle to match these two reports because the quantities differ. Because report splitting turns the reports filed by executing party A into four reports, two for its trade with executing party B and two for transactions with A's clients, we can 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 the two transactions can be identified as duplicates in the next step.

Around 21% of trades require aggregation. Of these, 84% are reports that were split in the previous step. Aggregation reduces the overall number of reports by around 16%. Table A6 summarizes the trading capacity, venue, buyer and seller account owner types, and trade volume of the aggregated trades. The affected trades are mainly off-venue trades with small volume, often involving a natural person or an internal account.

Variable	Split	% in agg trades	% in all data
Split	1	83.96	48.77
	0	16.04	51.23
Variable	Trading Capacity	% in agg trades	% in all data
Trading Capacity	AOTC	81.01	45.63
	DEAL	14.1	49.64
	MTCH	4.89	4.73
Variable	Trade Venue	% in agg trades	% in all data
Trade Venue	XOFF	59.75	28.84
	BMTF	13.54	27.65
	TREU	13.51	16.34
	Other	13.19	27.17
Variable	Counterparty Type	% in agg trades	% in all data
Counterparty Type	INTC	49.52	12.87
	LEI	44.02	77.25
	PRSN	4.99	7.98
	NaN	1.36	1.66
	MIC	0.11	0.24
Variable	Trade Size (mn EUR)	in agg trades	in all data
Trade Size (mn EUR)	Mean	4.9	5.44
	Median	0.11	0.5

Table A6: Characteristics of aggregated trades.

Not all fields of reports selected for aggregation match. We require the executing party, the buyer and seller account owners, the ISIN, the price, and the timestamp to be the same, and we sum the transacted quantities. Some firms file one execution twice, under two report identifiers that agree on the timestamp to the nanosecond, the quantity, every reported field and the venue trade identifier. Summing such copies would double the execution, so we keep one of them; copies without a venue trade identifier stay summed and are flagged. For other fields, there is

no clear aggregation rule. Table A7 reports the prevalence of disagreement between aggregated reports for each field. Overall disagreement is low and mainly concerns fields of minor interest for our analysis, so we keep the value from the first report in the aggregation group.

Column	Groups with differing Values (%)
eur_quantity	69.59
ind_investment_party	5.56
cat_trade_venue	5.37
id_tx_trade_venue	3.44
cat_short_selling	3.08
cat_exec_type	2.97
cat_trading_capacity	2.63
split	2.03
sdc_buyer_decision_maker	1.48
cat_price_type	1.03
cat_quantity_type	0.94
sdc_seller_decision_maker	0.87
sdc_transmitting_buyer	0.81
sdc_transmitting_seller	0.66
ind_transmission	0.14
cat_buyer_decision_maker_type	0.07
cat_inv_dec_type	0.06
cat_seller_decision_maker_type	0.0

Table A7: Disagreement among aggregated trades.

Price Outliers

We preliminarily flag price and quantity outliers so that they can be corrected using duplicate reports in the next step. If two reports are identified as duplicates and one has an unrealistic price or quantity, we replace the flagged value with the value from the paired 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 that, regardless of the reported unit, most prices are given as a percentage. Figure A22 reports the distribution of prices for each unit category. We disregard the reported unit and treat all prices as percentages.

We identify price outliers in two steps. 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 multiplied by 100 lies between 20 and 500 and within 50% of the rolling median, we treat it as a decimal misplacement, for example, a price recorded in units of 100 rather than percentage points, and correct it accordingly. This methodology adapts the corporate bond approach to sovereign bonds. In total, 0.2% 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 and the monetary value. Again, whatever the reported unit, most quantities seem to be reported as nominal values, and we treat all as such (Figure A23).

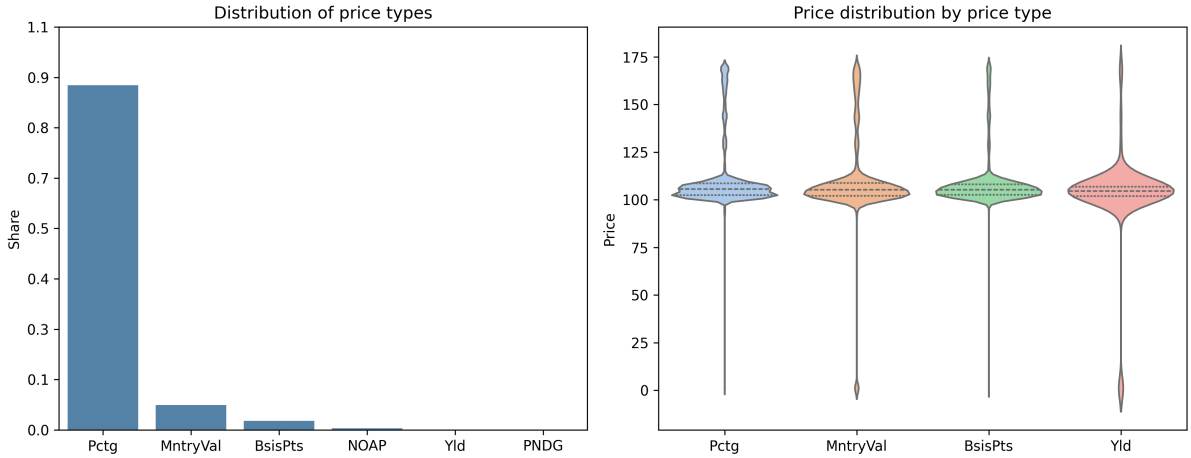


Figure A22: Price distribution by price type.

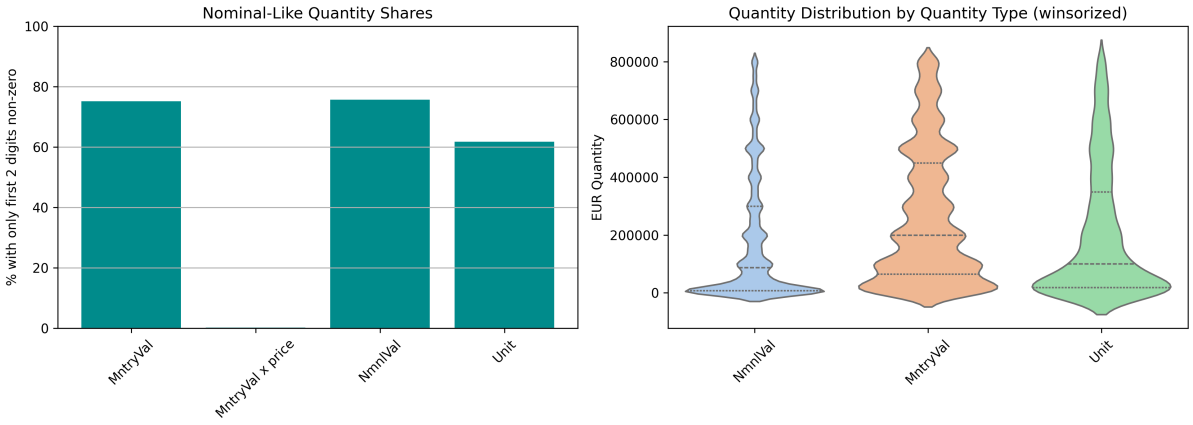


Figure A23: Quantities by quantity type. The left panel reports the share of observations in each category (including monetary value times price trades) that is divisible by a multiple of 10. The right panel reports the distribution of trade sizes for each category for trades below the median trade size.

We identify potential quantity outliers in the same spirit. The first pass compares trade sizes against reasonable thresholds. Extremely small (≤€100) or excessively large (≥€1 billion) trades are initially flagged as potential errors. A more refined pass then considers whether the transaction occurred before or after the bond's issuance date. 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 amount issued and flag as an outlier any transaction above the 99.5th percentile. Post-issuance, we flag any trade whose size exceeds 5% of the gross amount outstanding or whose size-to-outstanding ratio exceeds the 99.5th percentile across that day's trades. Overall, 0.7% of quantities are flagged as outliers. The threshold choices are deliberately conservative and enter the pipeline as correction flags, not as a broad deletion rule: flagged observations are first corrected from duplicate reports where possible, and the remaining global price-synchronization filter removes less than 1% of total observations (Section 3.5). The Appendix reports the affected shares at each stage so that readers can see the scale of the threshold-dependent steps.

Reasonable transaction sizes span a wide range, from small retail transactions to large interdealer trades, making clear outliers hard to identify. This preliminary screen is used only to improve duplicate detection; the analysis-specific outlier screen is applied later. These preparatory steps transform the raw data into a structured dataset suitable for further processing and reduce distortions caused by reporting anomalies. The duplicate-matching step then identifies and removes duplicate reports.

D.4 Duplicate Matching

Once the transaction data have been prepared, the next step is to identify duplicate records. Under MiFID II/MiFIR, both counterparties to a bilateral trade are required to report, so some transactions appear twice in the dataset; removing these duplicates is essential to avoid overestimating trading activity. Under MiFID I, reporting was predominantly one-sided, so fewer trades are reported twice. The duplicate-matching step adds two components specific to MiFID I—intra-reporter matched principal pairs reported by a single intermediary, and venue-consistent cross-capacity pairs—alongside the same cross-reporter record-linkage matching applied to proprietary trades, where it requires a total score of at least five and opposite reported directions. The following description focuses on the MiFID II logic; the MiFID I pipeline follows analogous but simpler steps.

We adapt the pairwise matching approach to sovereign bonds. For all trade reports with the same day and ISIN, we check agreement on price, quantity, timestamp, trade venue, buyer and seller, and whether the reports come from different executing parties. Because Bunds trade much more frequently than corporate bonds, computing all possible pairwise combinations of potential duplicates is infeasible. We therefore reduce the duplicate-search sample using the reporting requirements.

First, the reporting requirements apply only to investment firms, not to natural persons. A transaction side that involves only a natural person will therefore not have a duplicate report. This is the case when both the buyer/seller account owner and decision maker are natural persons, or when only the former is reported. We remove all transactions satisfying these conditions from the duplicate-search sample.

We apply the same logic 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 exclude it from the duplicate search. This is possible because the report-splitting step leaves at most two parties to a trade, one of which is the reporting executing party. We therefore need only check whether the other party also reports.

Removing transactions involving a natural person reduces the sample by approximately 7.4%, while dropping internal account trades removes another 4.2%. These trades can occur in large same-time, same-ISIN batches: an executing party may report up to 7,724 trades at the same time for the same instrument. Removing trades that cannot have duplicates keeps the pairwise comparison problem computationally feasible.

Metric	Batched Trades	All Trades
Number of transactions	56808	1511701
Number of trade batches (>1000)	30	-
Average trade size (EUR)	15,458.55	6,444,532.68
Median trade size (EUR)	2,000.00	792,000.00
Minimum group size	1002	-
Maximum group size	7724	-
Mean group size	1893.60	-
Top counterparty types (%)		
PRSN	98.5%	8.8%
LEI	1.5%	84.8%
INTC	0.1%	4.4%

Table A8: Batched trades.

Table A8 reports the characteristics of batches with more than 1,000 trades. They typically involve natural persons as counterparties and have small trade volumes. There are 30 such batches with 56,808 trades in total. Dropping trades that cannot have duplicates leaves 2,076 trades in batches of this size. About 31% of the dropped trades are part of a large batch. We use a Python record-linkage package to compute the pairwise comparisons. The package allows us to require compared trades to occur on the same date and in the same ISIN, which limits the number of pairs. A second necessary condition is that the executing parties reporting the two trades are distinct. The package does not directly support filtering pairs on the condition that a field differs, so we add this condition in a custom class to reduce both the number of pairs and the memory required for the comparison.

After these reductions, approximately 126 million pairs remain to be compared in 2019.

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 be either 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. Although all investment firms are supposed to report in UTC, time zone discrepancies remain a potential source of error. The score assigned for this field is continuous. The score falls linearly from one at an exact match, or at a shift of exactly one or two hours, to zero at a discrepancy of 20 minutes from either.

Trade quantities and prices are checked for exact matches, as well as potential reporting errors involving common scaling factors, for example, misreporting by a factor of 2, 10, 100, or 1,000. Buyer and seller account owners are checked for exact matches. All four fields receive a binary score; a match after rescaling by one of these factors scores one.

Finally, the trade venue is compared. We usually require an exact match, unless one side reports an off-venue trade ("XOFF"). Because trades are split in the first step, the client side

of a trade may still carry the trade venue of the market side when the two sides were initially reported as one trade. If the client also reports, the client will report only the client side, which occurs outside any trade venue.

The similarity scores from these comparisons are summed, and pairs with a positive timestamp score and a total above four out of six are candidate duplicate matches. This step differs from Jurkatis (2024), who use a probabilistic approach to filter for the most likely duplicates. Their approach accounts for the unconditional probability that a trade has a duplicate and allows for a more precise target rate of identified duplicates. Our threshold approach gives more direct control over match quality. We optimize the threshold by manually checking trades marginally above and below it. The final matched set is high-scoring: among accepted matches, the average score is 5.64 out of six, and Table A9 reports field-level agreement rates, so the accepted threshold can be assessed directly rather than only described qualitatively.

The final step employs graph-based matching to ensure that each transaction is linked to at most one counterparty report, preventing duplicate pairings, while maximizing the overall score across all final matches. Figure A24 illustrates this approach.

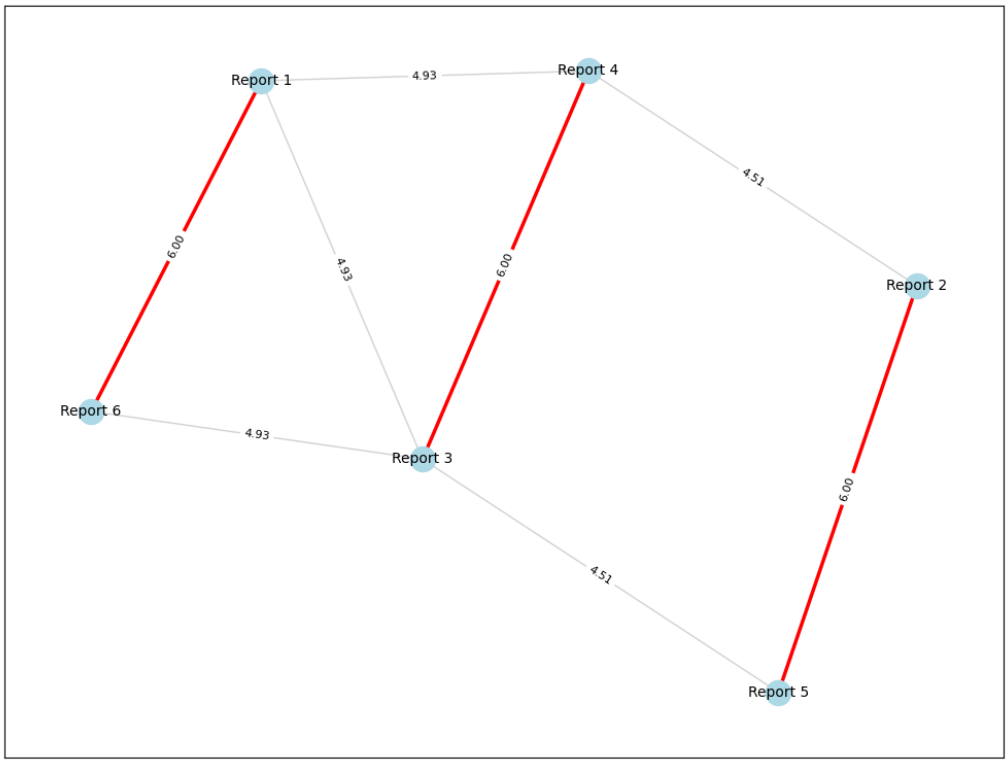


Figure A24: Maximum weight matching.

We find a duplicate for about 53% of all trades, or 60% of the trades included in the duplicate search. The average score of the final match set is 5.64 out of a maximum attainable score of six. Table A9 reports the average score across all matches for each field. Matched trades agree on all fields in 48% of cases, and the field-level agreement rate is at least 88% for every field. Disagreements most often concern buyer or seller information, for example because one side reports the parent company instead of the subsidiary with which it traded. Groups of fifty or more reports filed by the same executing party for the same instrument at an identical timestamp are retained and labeled with their group size. They are the shape MiFIR prescribes

for an aggregated client order—one market execution followed by one allocation report per client at the same price—so we treat them as client allocations, not as spurious reports.

Buyer	Seller	Trade Venue	Quantity	Price	Timestamp	Total Score
0.88	0.88	0.92	1.0	0.99	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 A9: Match quality. Field scores and total score of accepted matches in 2019. Rows: mean, minimum, 25th percentile, median, 75th percentile, and maximum.

D.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 trade venues are corrected using information from both records in a matched pair.

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. The corrected transaction is therefore not an arbitrary kept side. Price and quantity use the non-outlier report when only one side is flagged; counterparty fields are corrected only when one side mechanically reports the executing party; and venue fields follow the off-venue standardization rule. Remaining disagreements are retained as diagnostics rather than resolved by discretion. Table A10 gives an overview of these remaining differences within a matched pair of transaction reports. Of each matched pair we keep the report with the earlier corrected execution time, then the earlier originally reported execution time. Ties favor a dealer report, then the lower report-identifier hash. Reporter-specific fields describe the retained reporting firm.

Field	% Disagreement
Price	1.13
Quantity	6.19
Trade Venue	21.4
Buyer	2.34
Seller	2.4
Timestamp	34.29

Table A10: Remaining disagreement between duplicates.

D.6 Primary Issuances

The MiFID II data include primary market allocations from Finanzagentur auctions and taps. These records appear as primary dealers buying bonds from the German federal government on auction day. Dealers report either Finanzagentur or the Bundesbank as seller; the latter acts as fiscal agent in the issuance process. Because the analysis targets secondary market trading, we drop observations in which Finanzagentur or the Bundesbank sells a bond on that bond's auction day to a primary dealer. We keep all other trades by Finanzagentur and the Bundesbank, which mainly capture market maintenance trades and the Eurosystem's asset purchase programs.

Figure A25 validates this filter in the raw transaction reports. In MiFID II, cumulative BBK/FA sales to primary dealers rise sharply after the auction close around 11:30 CET and reach most of the auction size by noon. In MiFID I, the same diagnostic is much smaller because some dealers do not report their primary market allocations and about eight primary dealers (including five U.K.-reporting dealers) do not report any trades.

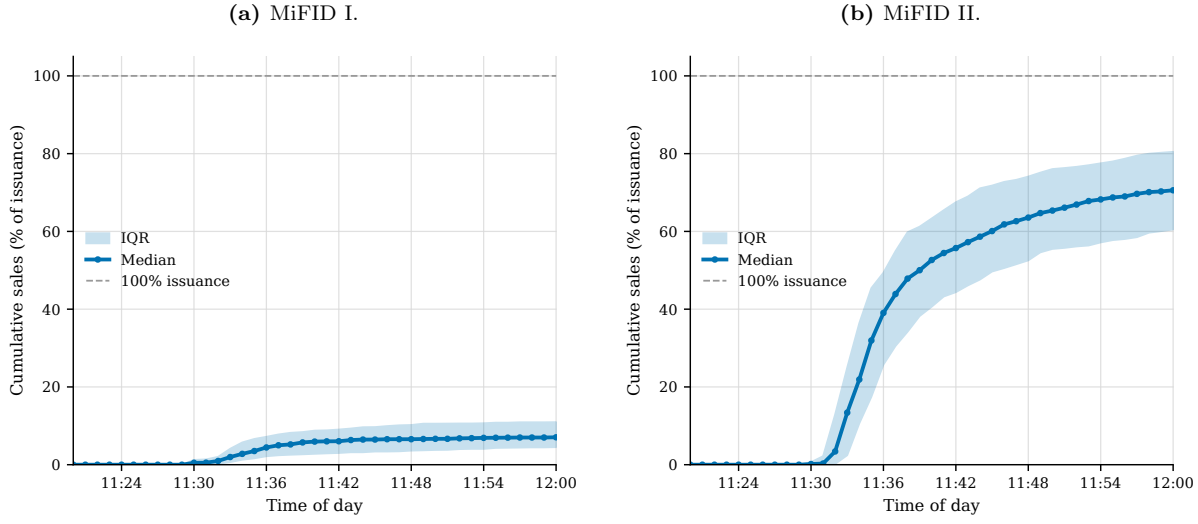


Figure A25: Primary issuance allocation legs in raw transaction reports. The panels plot, separately for MiFID I and MiFID II, the median and interquartile range of cumulative BBK/FA sales to primary dealers on auction days, normalized by the corresponding issuance volume. Each auction event is observed on a balanced one-minute grid from 11:20 to 12:00 CET. The horizontal dashed line marks 100% of issuance volume. The MiFID I panel is shown as a partial-coverage comparison under dealer-side reporting.

Figure A26 provides an independent, dealer-level check on these allocation legs across the full MiFID II sample. We aggregate each primary dealer's MiFID II auction allotments over a half-year, weighting each allocation by the bond's remaining maturity to match the Finanzagentur's own ranking basis, and compare the resulting ordering with the semiannual Bietergruppe ranking

the Finanzagentur publishes. The maturity-weighted allotment share falls monotonically with the official rank (panel a), and the two orderings correlate strongly in every half-year (panel b). The residual gap reflects the ranking's broader scope: it covers all federal securities, including Bubills and inflation-linked and green bonds outside our nominal-Bund universe.

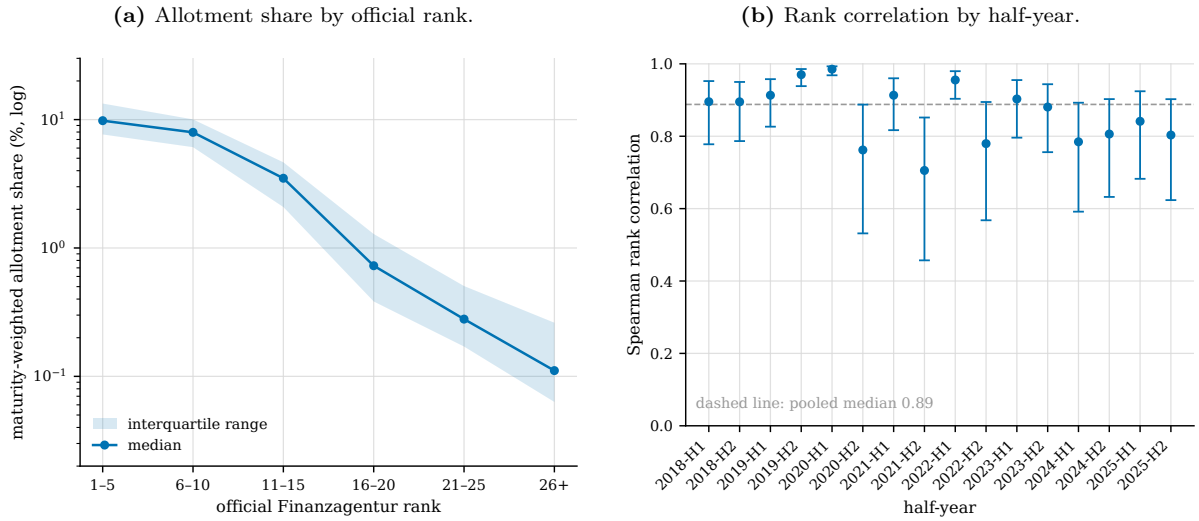


Figure A26: MiFID-derived primary-auction allotments against the official Finanzagentur Bietergruppe ranking over the MiFID II sample (2018–2025), both maturity-weighted (allotment volume times remaining years to maturity, matching the Finanzagentur's ranking basis). Panel (a): pooled median and interquartile range of each dealer's maturity-weighted allotment share of the half-year, by official-rank bin (bins of five). Panel (b): Spearman rank correlation between the MiFID-derived and official orderings per half-year, with 95% confidence intervals; the dashed line marks the pooled median.

D.7 Futures Deliverability and Delivery Legs

Deliverability classification Each day, we flag whether each outstanding bond is deliverable into the front contract of the 2-year (FGBS), 5-year (FGBM), 10-year (FGBL), or 30-year (FGBX) Eurex Bund futures. Eligibility is based on remaining maturity at the futures delivery date: 1.75–2.25 years for FGBS, 4.5–5.5 years for FGBM, 8.5–10.5 years for FGBL, and 24–35 years for FGBX. Except for FGBX, deliverable bonds must have an original maturity of no more than 11 years. Since 1 December 2022, Eurex has also imposed a seasoning requirement: a bond must have been issued before the first calendar day of the previous delivery month. From September 2022 onward, we use the official list of deliverable bonds scraped from the Eurex website; before that date, we reconstruct deliverability by applying these eligibility rules to the universe of outstanding Bunds, counting a bond as deliverable only from its first auction onward.

Delivery legs The raw MiFID II reports also contain physical delivery legs of Eurex Bund futures. MiFIR excludes the delivery of an underlying after the exercise or expiry of a derivative from transaction reporting (European Securities and Markets Authority, 2016, Sections 5.6.2.6–5.6.2.7), so most firms correctly never report their Bund futures deliveries. A minority of clearing members report them nonetheless, typically as one large ticket per bond against Eurex Clearing, dated the futures contract's last trading day or the next business day. These rows are not secondary market cash trades. The pipeline therefore drops MiFID II rows dated either day when Eurex Clearing is the central counterparty, the bond is in the expiring contract's deliverable basket, and the quantity is at least €1 million. The size threshold separates delivery reports from the retail-sized exchange trades that Eurex Clearing also clears on the same days. We do

not apply an analogous delivery-specific filter in MiFID I: in that regime, Eurex-related rows are handled through the general exchange and clearing-member normalization, and the raw data do not show a clean delivery-day signal that can be isolated with the same rule.

Figure A27 documents the filter. Figure A27a shows that raw rows involving Eurex Clearing have a sharp MiFID II spike on the last trading day, two business days before physical delivery, while the corresponding MiFID I series remains within its normal Eurex-related range. Figure A27b compares the volume selected by the MiFID II delivery filter with Eurex's public notified delivery volumes. The coverage distribution is a diagnostic, not an accounting reconciliation: because most firms do not report deliveries, raw MiFIR usually captures only part of the notified face amount, and a small tail exceeds one-to-one coverage when both CCP-side reports or duplicate raw reports are visible.

(a) Delivery-window spike. **(b)** Coverage against Eurex notifications.

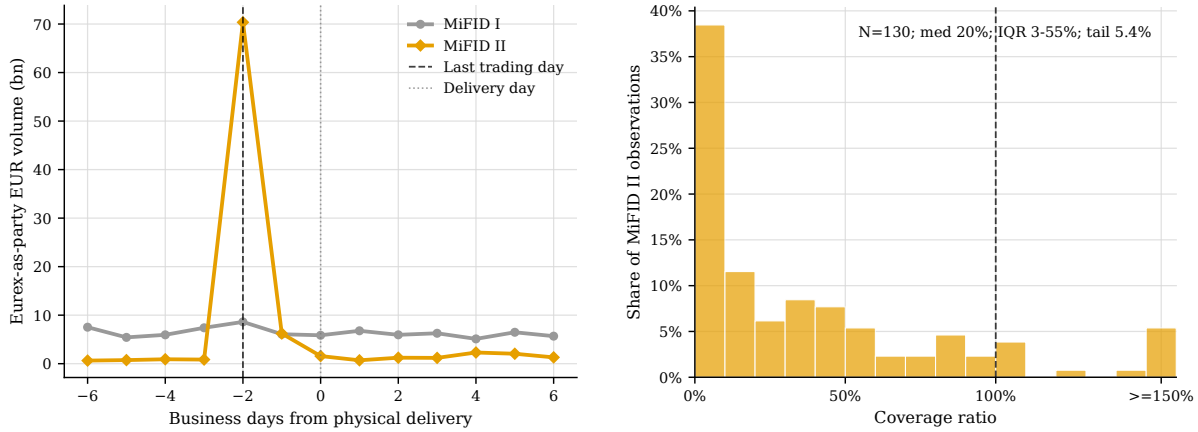


Figure A27: Futures-delivery filter diagnostic. The left panel plots aggregate raw € volume of rows involving Eurex Clearing around physical delivery dates, separately for MiFID I and MiFID II. Day 0 is physical delivery; the dashed vertical line marks the last trading day ($T = -2$). The right panel reports the MiFID II-only distribution of the coverage ratio, defined at the delivery-month–contract-root level as raw MiFIR € quantity satisfying the delivery filter divided by Eurex notified contracts times €100,000 face value. Bars are normalized to sum to 100%; bins are 10 percentage points and the final bin collects ratios of at least 150%. The histogram is a coverage diagnostic, not an exact reconciliation to delivered face amount.

D.8 Trades by CCPs, IDBs and venue operators

Because of reporting standards and our normalization of the data, "intermediaries" can appear as buyer or seller even if they do not take the bond into inventory. This is particularly relevant for CCPs and interdealer brokers (IDBs), and to a lesser extent venue operators. We pair the two sides of an intermediated trade where the reports identify the pairing, keep every other record, and record separately which records belong to one execution, so that reported-record volume can be distinguished from reconstructed execution volume. 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 in two steps:

First, we recombine reports split in the normalization step, provided they have not been aggregated in the meantime, joining the two records of one report on the report identifier and the trade date. A record that lists the executing party as both buyer and seller—a venue operator

reporting against its own market identifier code—stands for the side its report lacks, and only for that side. We drop duplicates associated with either record of a recombined report.

Second, for the remaining reports, we group the intermediary's buy and sell records by bond, timestamp, and price. Within a group, a quantity with exactly one buy record and one sell record is paired into a single transaction between the two members, with the intermediary retained as executing party; Table A11 illustrates this case. Where one record faces several records on the other side whose quantities sum to no more than its own, each of the smaller records is paired with it at its own quantity, and any remainder stays as a residual record facing the intermediary. All other groups—several buy records against several sell records, or records on one side only—are left as they are: each record keeps its member and faces the intermediary. Member reports that the duplicate matching linked to a paired record are deleted, since the pair already carries that trade; a member report linked to a record that remains unpaired stays until the final deduplication, which keeps one report per matched pair.

(a) Original reports.

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) Resulting report.

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

Table A11: Case 1: original and resulting reports.

Case 2: The intermediary is listed as buyer or seller but not as the executing party of the transaction. The preceding steps concern reports where the executing party is an intermediary. Some reports instead list the intermediary as buyer or seller even though another party executed the trade, because the reporting executing party may not know the actual buyer or seller. These reports are kept as reported. In general, they should have a duplicate filed by the intermediary; if the duplicate-linking procedure identifies it, the final deduplication keeps one of the two reports. Otherwise the record stays as a one-sided trade: the intermediary is moved from the buyer or seller field into a separate counterparty field, and that side of the trade is left empty. The records that Case 1 leaves unpaired are treated the same way.

We keep the reported trading venue whenever it identifies a regulated market, MTF, or OTF. Trades intermediated by IDBs, such as BGC Brokers and GFI, usually carry the MIC code of the broker's OTF. When a report against an IDB or another venue operator instead carries an off-venue or systematic internalizer code, we assign the broker's most frequent trading venue. A central counterparty is not a trading venue, so reports against one keep their reported venue: LCH, for instance, clears trades executed both on MTS and on BrokerTec.

These steps do not eliminate pure intermediaries from the data as counterparties. Where the reports do not identify the pairing, the two members' reports of one intermediated trade survive as two one-sided records, each facing the intermediary. The execution layer described next records that they are one execution.

Trade records and executions The unit of the cleaned file is the trade record. One execution between two account owners can appear as two, three, or four records: the two records of an agency or matched principal report that the normalization step splits, an internal account booked in and out within the same second, or the two members' one-sided reports of an anonymous venue, CCP, or IDB trade. For example, when a bank buys a bond from a dealer as agent for a fund, its report becomes two records, dealer to bank and bank to fund; summing their quantities would count the bond twice. No record is deleted. Instead, each MiFID II record carries an execution identifier shared by linked records, a flag for whether the record's quantity counts toward execution volume, and the rule that linked it. The rules act on groups of records that share a node—an identified account owner, an executing firm's internal account, or an empty side facing a venue, CCP, or IDB—in the same bond, second, and price. The two records of a split report are always linked, and the record on which the executing party sells stops counting. An identified party whose purchases and sales within the group balance, and that reported at least one of the records in an agency or matched principal capacity, is a pass-through: its selling records stop counting. A balanced pair that the party reports as principal on both sides is flagged but keeps counting on both records. For internal accounts and empty sides, the smaller side of the group by volume stops counting; without a venue trade identifier, this can merge two distinct opposite trades in the same bond, second, and price, so execution counts and execution volume are, if anything, understated. Where the reports carry a venue trade identifier, one-sided records of two reporters that share it in opposite directions are linked as one on-venue trade, and one-sided records that repeat the same identifier, party, direction, and quantity are counted once. When a split agency report and the other venue member's report form a three-record group, these rules can otherwise discount the same outgoing record twice and leave two records counting for one execution. We retain one counting record only when the raw-report lineage establishes the split and all three records agree exactly on price and quantity, with the venue-facing records sharing a known CCP or a specific MIC. This correction preserves the records and their economic fields; it does not collapse principal resale chains, partial fills, or larger ambiguous groups. Each final MiFID II record can be traced to its contributing raw reports through the cleaning transformations. Execution volume is the sum of quantities over counting records. MiFID I records are not linked and count once each.

Statistics that describe the size and venue composition of the market are computed on executions: trading volume and its comparison with MiFID I and the Finanzagentur, trade counts and sizes, venue-type shares, bond and venue concentration, and the venue shares around the Brexit transition. MiFID II has 7.7 million executions against 14.2 million records. The normalization step can record the client record of an agency chain, or one member's report of a venue trade, as off-venue, so an execution takes the venue type of its most specific venue record. Execution weighting lowers the OTC share of volume and raises the SI and OTF shares. Statistics defined by a side of the trade—sectors, counterparties, dealers, relationships, networks, net flows, settlement obligations, the dealer capacity mix with identified clients, and the regressions of execution costs and price impact—are computed on records: an execution has no single buyer or seller, and which record stops counting depends on the direction of the trade. There a "trade" is a record. Retail participation counts client allocations, one record per allocation. Execution type is read from each record as reported.

D.9 External transaction checks

We cross-check the regulatory reports against available MTS Bund transactions over 2008–2025 and records from Tradeweb, MTS BondVision, TP ICAP, and Fenics on six trading days in 2020 obtained for Kerssenfischer and Helmus (2025). After aligning price and quantity units

and timestamps, we compare transaction prices, quantities, and execution counts across the sources. These comparisons check our reconstruction of executions, including the treatment of intermediated trades described in Appendix D.8. We use overlapping transactions and account for differences between commercial execution records and regulatory reports. Within the overlapping MiFID II samples, prices, quantities, and timestamps broadly agree with MTS, Tradeweb, and BondVision records.

D.10 Investor Sector Mapping

We construct an exhaustive mapping of LEIs to the following investor sectors: primary dealers, banks, interdealer brokers (IDB), investment funds, hedge funds, insurance companies and pension funds (ICPF), central counterparties (CCPs), official sector entities, non-financial corporations (NFC), and non-bank financial intermediaries (NBFI).

Table A12 lists our data sources. We prioritize high-quality regulatory data. The list of primary dealers participating in Bund auctions is available from Finanzagentur; we classify an LEI as a primary dealer if its auction-group membership was active during the MiFID II sample, that is, ending after 1 January 2018 or still open-ended, so that institutions that had left the auction group before then are not classified as dealers. 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 Prudential Regulation Authority (PRA) register of U.K. 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-to-sector mappings from Moody's BvD BankFocus database. These data are especially useful for identifying large global banks. Second, we use Morningstar data to identify hedge funds. We used OpenAI's `gpt-6.1-sol` model in Codex with high reasoning effort to assist sector cross-checks for 362 observed identifiers, prioritizing large reported volumes and many reported records and consulting regulator and issuer sources where available; the authors adjudicated the final label decisions.

For any remaining LEIs, we first use the GLEIF Level 1 "Who is Who" register[28] 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, for example, Perrella and Catz, 2020), which contains a mapping from LEIs to sectors (based on the European system of accounts framework).[29]

We use sector concentration within an ownership group as a diagnostic for classification discrepancies. Trading activity alone does not establish an entity's sector, so we retain sourced classifications and apply explicit corrections supported by regulatory or issuer evidence.

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

[29] We avoid using RIAD where possible because it is notoriously unreliable. Apart from omitting many important investors, RIAD contains dubious sector classifications. For example, many Global Systemically Important Banks (G-SIBs) outside the euro area are classified as non-banks (ESA codes S.125-S.127), either because they are not depository institutions in the euro area or because they operate as holding companies.

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
U.K. PRA	U.K. Prudential Regulation Authority list of banks authorized in the U.K.	link
Moody's BvD BankFocus	Proprietary database incl. LEI-to-sector mapping	
Morningstar Hedge Funds	Proprietary database of hedge funds incl. LEI	
Manual Classification	Manually curated LEI-to-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)	

Table A12: Investor reference data sources, hierarchical order.

Table A13 and Table A15 report descriptive diagnostics for this investor sector mapping. The unit is an account-owner side of a cleaned Bund transaction, so each transaction contributes once for the buyer and once for the seller. These diagnostics document the coverage and composition of the mapping used in the paper's downstream analyses.

Panel A: group-level sector coverage by sample

Sample	Sides	Vol. (EUR bn)	Class. sides (%)	Class. vol. (%)	Groups	Class. groups (%)
MiFID I	3,998,530	16,562.0	77.8	64.1	1,475	97.6
MiFID II	25,867,408	101,554.2	92.9	82.3	41,172	77.4
Full sample	29,865,938	118,116.2	90.8	79.8	41,482	77.5

Panel B: entity-source tier mix, MiFID II account-owner sides

Source tier	Sides	Side share (%)	Vol. (EUR bn)	Vol. share (%)	Groups
Regulatory lists	13,168,387	50.9	47,747.5	47.0	7,749
Hand/manual	5,692,808	22.0	32,186.1	31.7	1,337
Commercial data	598,315	2.3	2,461.8	2.4	3,466
GLEIF category	167,979	0.6	375.3	0.4	9,549
Name/MIC heuristic	284,334	1.1	344.5	0.3	769
RIAD fallback	231,566	0.9	563.2	0.6	18,874
Retail flag	3,966,190	15.3	71.5	0.1	0
Unclassified/no source	1,757,829	6.8	17,804.2	17.5	0

Table A13: Investor sector classification coverage and source provenance. Panel A reports group-level sector coverage by reporting regime for account-owner sides of cleaned Bund transactions in the analysis sample. Volume is side-volume, so each transaction contributes once for the buyer and once for the seller. Classified sides have a nonmissing group-level sector. Group coverage excludes retail sides without LEIs. Panel B reports the entity-level source tier recorded in the investor reference data for MiFID II account-owner sides. Source tiers are grouped for readability.

Counterparty coverage by regime The two regimes name different counterparties. A MiFID I report names the reporting firm and, by BaFin code, a counterparty that is itself subject to reporting; an end client usually stays unnamed. A MiFID II report names the buyer and the seller themselves. Table A14 shows the consequence for the client side of dealer-to-client records. In 2013–2017, 59% of its record volume carries no client identifier, 23% names only a CCP, interdealer broker, or venue, and 17% carries a sector. In 2018–2022, these shares are 15%, 22%, and 62%. The identified MiFID I clients are mostly banks and official sector entities: investment funds, hedge funds, insurers and pension funds, other non-bank financial intermediaries, and non-financial corporations together account for 1.6% of client-side volume in 2013–2017, against 39.2% in 2018–2022. Over the same two windows, the number of investment fund groups on the client side rises from 76 to 2,050, and the number of hedge fund groups from 6 to 389. Sector information therefore exists before 2018 (59–68% of buyer and seller volume carries a sector in each MiFID I year), but it describes regulated counterparties, not end clients.

Sector	Full sides (%)	Full vol. (%)	MiFID II sides (%)	MiFID II vol. (%)	Groups
Primary Dealer	34.5	54.9	35.5	56.7	38
IDB	0.0	0.0	0.0	0.0	0
CCP	0.0	0.0	0.0	0.0	0
Bank	21.9	3.9	23.0	2.8	2,118
Investment Fund	9.8	8.0	11.2	9.2	9,324
ICPF	1.3	0.8	1.4	0.9	1,546
Hedge Fund	2.0	6.8	2.3	7.9	732
Official	1.2	4.3	1.2	4.1	412
NBFI	3.8	0.4	2.1	0.2	2,597
NFC	0.7	0.3	0.7	0.3	15,578
Retail	15.6	0.4	15.3	0.1	0
Unclassified	9.2	20.2	7.1	17.7	9,345

Table A15: Investor sector composition of account-owner sides. The table reports the distribution of group-level investor sectors used in the paper's downstream analyses. The IDB and CCP rows are zero by construction: interdealer brokers and central counterparties are part of the entity-level mapping, but the normalization in Appendix D.8 pairs the two sides of the trades they intermediate or moves the intermediary into a separate counterparty field, so they do not appear as account owners in the cleaned file. The full sample spans 2008–2025; the MiFID II sample spans 2018–2025. Volume shares are side-volume shares.

Panel A: identification of account owners by year, share of record volume (%)

	All sides	Client side of dealer-to-client records			
		No client identifier			
Year	Sector	CCP, IDB, venue	None	Identifier, no sector	Sector
2008	62	16	64	1	19
2009	68	14	57	2	27
2010	66	13	61	1	24
2011	65	14	62	2	22
2012	67	14	57	2	27
2013	65	24	50	2	24
2014	63	29	53	1	17
2015	59	22	65	1	13
2016	60	22	65	0	13
2017	66	21	61	1	17
2018	80	26	12	0	61
2019	81	23	14	0	63
2020	83	16	20	0	63
2021	88	24	9	0	66
2022	86	22	18	0	60
2023	81	18	34	0	48
2024	80	20	30	0	50
2025	81	20	29	0	51

Panel B: client side of dealer-to-client records by sector

	MiFID I, 2013–2017		MiFID II, 2018–2022	
Client side	Volume (%)	Groups	Volume (%)	Groups
Banks	10.0	447	7.5	1,109
Investment funds	0.4	76	16.3	2,050
Hedge funds	0.2	6	20.1	389
ICPF	0.1	13	1.4	507
Official sector	5.0	32	15.6	160
NBFI	0.5	25	0.8	574
NFC	0.3	20	0.6	2,524
Households	0.0	–	0.0	–
Identifier without sector	1.1		0.3	
CCP, IDB or venue	23.3		22.3	
No client identifier	59.0		15.0	

Table A14: Counterparty identification in MiFID I and MiFID II. Buyer and seller sides of cleaned records in the analysis sample; shares of record volume, so each record counts once for the buyer and once for the seller. All sides: share of volume whose account owner has a group-level sector. A dealer-to-client record has exactly one side on a primary dealer's trading desk on the trade date; the other side is the client side. CCP, IDB, venue: the client side is a central counterparty, interdealer broker, or trading venue (Appendix D.8). None: no account owner, or an internal account, including reports filed without any account owner, which occur in every year and more often from 2023. Identifier, no sector: an LEI or BaFin code without a sector. Sector includes natural persons; banks include primary dealer sector banks outside their auction group window on the trade date. Panel B uses five-year windows; groups are distinct investor groups with at least one client-side record in the window.

Investor Grouping The group hierarchy uses GLEIF Level 2 direct and ultimate-parent relationships, augmented with ECB lists of Significant Institutions and Monetary Financial Institutions (the latter for banks). We prioritize ultimate-parent links, resolve remaining direct-parent chains recursively, terminate circular chains at the first repeated entity, and treat entities without valid ownership links as singleton groups. Only operating companies—entities GLEIF classifies as *general* or *branch*—are consolidated; fund vehicles (GLEIF category *fund*) stay out of the hierarchy. Consolidating funds would merge the many separately identified vehicles advised by a single manager into one investor: in our reference data BlackRock alone spans roughly three thousand distinct fund LEIs, including some 1,900 iShares ETFs, which consolidation would collapse into a single "BlackRock" investor. The asset-management companies themselves are still consolidated—a manager's discretionary trading arms are grouped—but the funds they advise remain distinct investors; interdealer brokers, by contrast, collapse to roughly one group per two dozen LEIs.

Within a primary dealer group, every entity except the subsidiaries whose own sector is asset management or insurance belongs to the dealer desk. It carries the dealer flag on every date on which an entity of the group is an auction member, whichever entity holds the membership. The excluded subsidiaries remain clients in their own sectors, so their trades with the parent's desk enter as dealer-to-client trades (the margin used in Figure A7). The same principle applies outside primary dealer groups: an asset manager keeps the investment fund sector at the group level even when an insurer, bank, pension fund, or non-financial corporation owns it, and execution subsidiaries that route a manager's fund orders to dealers are classified with the manager. A small number of manual corrections align group assignments with in-sample corporate structure: entities retired during the sample are attached to their in-sample successor (for example, Banca IMI to Intesa Sanpaolo), and Credit Suisse is kept separate from UBS, whose 2023 acquisition the current register snapshot would otherwise back-date onto the entire sample.

D.11 Venue Classification

We combine three external reference datasets on trading venues: the ESMA Financial Instruments Transparency System (FITS) register,[30] the U.K. Financial Conduct Authority (FCA) register,[31] and the ISO 10383 list of Market Identifier Codes (MICs).[32] We use the ESMA register as the primary source; for MIC codes not classified by ESMA, we rely on the FCA register; and for the remaining venues, we rely on ISO data. All venue types are harmonized to the ESMA taxonomy. 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 MIC code to our sector classification (classifying venues run by interdealer brokers as OTF and those run by primary dealers as SI), (ii) we parse 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 (for example, Tradeweb and Bloomberg mapped to MTF, and major IDBs mapped to OTF). Approved publication arrangements and off-venue MICs are collapsed into OTC/off-venue, while residual unclassified MICs are retained as Other. This yields a harmonized classification of trading venues to venue types and a consistent mapping from MICs to LEIs.

We separately assign each MIC a venue protocol family. The family describes the trading model of a venue or segment; it does not record how an individual trade executed. A MIC is

[30] https://registers.esma.europa.eu/publication/searchRegister?core=esma_registers_upreg
[31] https://register.fca.org.uk/servlet/servlet.FileDownload?file=0150X000006gbb1
[32] https://www.iso20022.org/sites/default/files/ISO10383_MIC/ISO10383_MIC.xlsx

in the request-for-quote (RFQ) family or the order book family (CLOB) only when the venue's rulebook or its operator's description documents that model for the venue or segment itself; documents for a sister venue of the same operator in another jurisdiction are not transferred, and venue type alone never assigns either family. Venues whose documented model is neither, such as specialist and market maker segments of exchanges, and all OTC and OTF codes form the residual other family. The remaining MICs are unknown: every systematic internalizer, because SI is a regulatory status, not a trading model; venues documented as combining an order book with RFQ trading; and venues without a source. Each label holds for the whole sample period. The panel contains 276 distinct MICs. The 118 material MICs, which together account for 99.8% of execution volume, received a recorded disposition, and the curated RFQ, order book, and other labels rest on primary sources; the remaining MICs follow the residual rule. We used OpenAI's `gpt-6.1-sol` model in Codex with high reasoning effort to assist the source review of the material MICs; the authors adjudicated the final labels. The family labels are not taken from the ESMA, FCA, or ISO registers. Table A16 reports the share of execution volume and of executions in each family and evidence class, separately for MiFID I, MiFID II before 2021, and MiFID II from 2021; the 2021 split marks the exit of U.K. venues from the reporting perimeter.

		Venue protocol family				Evidence				
Period	Share of	RFQ	CLOB	Other	Unknown	Rulebook	Operator	Type	Mixed	None
MiFID I, 2008–2017	EUR volume	0.6	13.6	83.4	2.4	1.7	15.9	80.0	0.0	2.4
	Executions	0.4	3.2	90.5	5.9	15.7	27.7	50.7	0.1	5.9
MiFID II, 2018–2020	EUR volume	29.2	0.7	48.9	21.3	30.4	2.5	45.8	0.0	21.3
	Executions	47.3	2.7	33.1	17.0	45.5	6.4	31.1	0.7	16.3
MiFID II, 2021–2025	EUR volume	24.5	0.6	53.5	21.3	16.9	8.6	53.2	0.0	21.3
	Executions	28.0	5.2	45.4	21.5	22.3	17.8	38.3	0.1	21.3

Table A16: Coverage of the venue protocol family. Shares of execution volume (EUR volume rows) and of executions (Executions rows) by venue protocol family and by evidence class, in percent of the period total. Evidence classes: rulebook of the venue or segment; operator description; type (OTC or OTF residual); mixed (order book plus RFQ); none. Executions are reconstructed from trade records; MiFID I records are one execution each.

For venue group statistics, we merge MICs of the same venue type when they belong to the same operator. This grouping matters after Brexit because operators such as Bloomberg, Tradeweb, and MarketAxess maintain parallel EU and U.K. venues with distinct MICs but similar trading functionality. Statistics based on raw MIC codes would mechanically split those operator groups and understate concentration, especially for MTFs.